

To Hash or Not to Hash Again?

(In)differentiability Results for H^2 and HMAC

Yevgeniy Dodis* Thomas Ristenpart† John Steinberger‡ Stefano Tessaro§

July 9, 2012

Abstract

We show that the second iterate $H^2(M) = H(H(M))$ of a random oracle H cannot achieve strong security in the sense of indistinguishability from a random oracle. We do so by proving that indistinguishability for H^2 holds only with poor concrete security by providing a lower bound (via an attack) and a matching upper bound (via a proof requiring new techniques) on the complexity of any successful simulator. We then investigate HMAC when it is used as a general-purpose hash function with arbitrary keys (and not as a MAC or PRF with uniform, secret keys). We uncover that HMAC's handling of keys gives rise to two types of weak key pairs. The first allows trivial attacks against its indistinguishability; the second gives rise to structural issues similar to that which ruled out strong indistinguishability bounds in the case of H^2 . However, such weak key pairs do not arise, as far as we know, in any deployed applications of HMAC. For example, using keys of any fixed length shorter than $d - 1$, where d is the block length in bits of the underlying hash function, completely avoids weak key pairs. We therefore conclude with a positive result: a proof that HMAC is indistinguishable from a RO (with standard, good bounds) when applications use keys of a fixed length less than $d - 1$.

Keywords: Indistinguishability, Hash functions, HMAC.

* Department of Computer Science, New York University, dodis@cs.nyu.edu

† Department of Computer Sciences, University of Wisconsin–Madison, rist@cs.wisc.edu

‡ Institute of Theoretical Computer Science, Tsinghua University, jpsteinb@gmail.com

§ CSAIL, Massachusetts Institute of Technology, tessaro@csail.mit.edu

Contents

1	Introduction	3
1.1	The Second Iterate Paradox	3
1.2	HMAC with Arbitrary Keys	5
1.3	Discussion	6
1.4	Prior Work	6
2	Preliminaries	6
3	Second Iterates and their Security	8
3.1	Hash chains using Second Iterates	8
3.2	An Example (Vulnerable) Application: Mutual Proofs of Work	11
3.3	An Indifferentiability Distinguisher for any Second Iterate	12
3.4	Indifferentiability Upper Bound for a Second Iterate	14
4	HMAC as a General-purpose Keyed Hash Function	18
4.1	Weak Key Pairs in HMAC	19
4.2	Colliding Key Pairs and the Indifferentiability of HMAC	20
4.3	Ambiguous Key Pairs and the Indifferentiability of HMAC	21
4.4	Indifferentiability Upper Bound for HMAC with Restricted Keys	25
A	Internal Collision Probabilities	34
B	Proof of Theorem 3.2	35
C	Proof of Theorem 4.2	40
D	Proof of Theorem 3.3	43

1 Introduction

Cryptographic hash functions such as those in the MD and SHA families are constructed by extending the domain of a fixed-input-length compression function via the Merkle-Damgård (MD) transform. This applies some padding to a message and then iterates the compression function over the resulting string to compute a digest value. Unfortunately, hash functions built this way are vulnerable to extension attacks that abuse the iterative structure underlying MD [24, 40]: given the hash of a message $H(M)$ an attacker can compute $H(M \parallel X)$ for some arbitrary X , even without knowing M .

In response, suggestions for shoring up the security of MD-based hash functions were made. The simplest is due to Ferguson and Schneier [22], who advocate a hash-of-hash construction: $H^2(M) = H(H(M))$, the second iterate of H . An earlier example is HMAC [5], which similarly applies a hash function H twice, and can be interpreted as giving a hash function with an additional key input. Both constructions enjoy many desirable features: they use H as a black box, do not add large overheads, and appear to prevent the types of extension attacks that plague MD-based hash functions.

Still, the question remains whether they resist other attacks. More generally, we would like that H^2 and HMAC behave like random oracles (ROs). In this paper, we provide the first analysis of these functions as being indifferentiable from ROs in the sense of [14, 33], which (if true) would provably rule out most structure-abusing attacks. Our main results surface a seemingly paradoxical fact, that the hash-of-hash H^2 cannot be indifferentiable from a RO with good bounds, *even* if H is itself modeled as a keyed RO. We then explore the fall out, which also affects HMAC.

INDIFFERENTIABILITY. Coron et al. [14] suggest that hash functions be designed so that they “behave like” a RO. To define this, they use the indifferentiability framework of Maurer et al. [33]. Roughly, this captures that no adversary can distinguish between a pair of oracles consisting of the construction (e.g., H^2) and its underlying ideal primitive (an ideal hash H) and the pair of oracles consisting of a RO and a simulator (which is given access to the RO). A formal definition is given in Section 2. Indifferentiability is an attractive goal because of the MRH composition theorem [33]: if a scheme is secure when using a RO it is also secure when the RO is replaced by a hash construction that is indifferentiable from a RO. The MRH theorem is widely applicable (but not ubiquitously, c.f., [37]), and so showing indifferentiability provides broad security guarantees.

While there exists a large body of work showing various hash constructions to be indifferentiable from a RO (c.f., [1, 7, 12–14, 18, 19, 26]), none have yet analyzed either H^2 or HMAC. Closest is the confusingly named HMAC construction from [14], which hashes a message by computing $H^2(0^d \parallel M)$ where H is MD using a compression function with block size d bits. This is not the same as HMAC proper nor H^2 , but seems close enough to both that one would expect that the proofs of security given in [14] apply to all three.

1.1 The Second Iterate Paradox

Towards refuting the above intuition, consider that $H^2(H(M)) = H(H^2(M))$. This implies that an output of the construction $H^2(M)$ can be used as an intermediate value to compute the hash of the message $H(M)$. This property does not exist in typical indifferentiable hash constructions, which purposefully ensure that construction outputs are unlikely to coincide with intermediate values. However, and unlike where extension attacks apply (they, too, take advantage of outputs being intermediate values), there are no obvious ways to distinguish H^2 from a RO.

Our first technical contribution, then, is detailing how this structural property might give rise to vulnerabilities. Consider computing a hash chain of length ℓ using H^2 as the hash function. That is, compute $Y = H^{2\ell}(M)$. Doing so requires 2ℓ H -applications. But the structural property of H^2 identified above means that, given M and Y one can compute $H^{2\ell}(H(M))$ using only one H -application:

$H(Y) = H(H^{2\ell}(M)) = H^{2\ell}(H(M))$. Moreover, the values computed along the first hash chain, namely the values $Y_i \leftarrow H^{2i}(M)$ and $Y'_i \leftarrow H^{2i}(H(M))$ for $0 \leq i \leq \ell$ are disjoint with overwhelming probability (when ℓ is not unreasonably large). Note that for chains of RO applications, attempting to cheaply compute such a second chain would not lead to disjoint chains. This demonstrates a way in which a RO and H^2 differ.

We exhibit a cryptographic setting, called *mutual proofs of work*, in which the highlighted structure of H^2 can be exploited. In mutual proofs of work, two parties prove to each other that they have computed some asserted amount of computational effort. This task is inspired by, and similar to, client puzzles [20, 21, 27, 28, 39] and puzzle auctions [41]. We give a protocol for mutual proofs of work whose computational task is computing hash chains. This protocol is secure when using a random oracle, but when using instead H^2 an attacker can cheat by abusing the structural properties discussed above.

INDIFFERENTIABILITY LOWER BOUND. The mutual proofs of work example already points to the surprising fact that H^2 does not “behave like” a RO. In fact, it does more, ruling out proofs of indifferenciability for H^2 with good bounds. (The existence of a tight proof of indifferenciability combined with the composition theorem of [33] would imply security for mutual proofs of work, yielding a contradiction.) However, we find that the example does not surface well why simulators must fail, and the subtlety of the issues here prompt further investigation. We therefore provide a direct negative result in the form of an indifferenciability distinguisher. We prove that should the distinguisher make q_1, q_2 queries to its two oracles, then for any simulator the indifferenciability advantage of the distinguisher is lower-bounded by $1 - (q_1 q_2)/q_S - q_S^2/2^n$. (This is slightly simpler than the real bound, see Section 3.3.) What this lower bound states is that the simulator must make very close to $\min\{q_1 q_2, 2^{n/2}\}$ queries to prevent this distinguisher’s success. The result extends to structured underlying hash functions H as well, for example should H be MD-based.

To the best of our knowledge, our results are the first to show lower bounds on the number of queries an indifferenciability simulator must use. That a simulator must make a large number of queries hinders the utility of indifferenciability. When one uses the MRH composition theorem, the security of a scheme when using a monolithic RO must hold up to the number of queries the simulator makes. For example, in settings where one uses a hash function needing to be collision-resistant and attempts to conclude security via some (hypothetical) indifferenciability bound, our results indicate that the resulting security bound for the application can be at most $2^{n/4}$ instead of the expected $2^{n/2}$.

UPPER BOUNDS FOR SECOND ITERATES. We have ruled out good upper bounds on indifferenciability, but the question remains whether weak bounds exist. We provide proofs of indifferenciability for H^2 that hold up to about $2^{n/4}$ distinguisher queries (our lower bounds rule out doing better) when H is a RO. We provide some brief intuition about the proof. Consider an indifferenciability adversary making at most q_1, q_2 queries. The adversarial strategy of import is to compute long chains using the left oracle, and then try to “catch” the simulator in an inconsistency by querying it on a value at the end of the chain and, afterwards, filling in the intermediate values via further left and right queries. But the simulator can avoid being caught if it prepares long chains itself to help it answer queries consistently. Intuitively, as long as the simulator’s chains are a bit longer than q_1 hops, then the adversary cannot build a longer chain itself (being restricted to at most q_1 queries) and will never win. The full proofs of these results are quite involved, and so we defer more discussion until the body. We are unaware of any indifferenciability proofs that requires this kind of nuanced strategy by the simulator.

1.2 HMAC with Arbitrary Keys

HMAC was introduced by Bellare, Canetti, and Krawczyk [5] to be used as a pseudorandom function or message authentication code. It uses an underlying hash function H ; let H have block size d bits and output length n bits. Computing a hash $\text{HMAC}(K, M)$ works as follows [30]. If $|K| > d$ then redefine $K \leftarrow H(K)$. Let K' be K padded with sufficiently many zeros to get a d bit string. Then $\text{HMAC}(K, M) = H(K' \oplus \text{opad} \parallel H(K' \oplus \text{ipad} \parallel M))$ where opad and ipad are distinct d -bit constants. The original (provable security) analyses of HMAC focus on the setting that the key K is honestly generated and secret [3, 5]. But what has happened is that HMAC's speed, ubiquity, and assumed security properties have lead it to be used in a wide variety of settings.

Of particular relevance are settings in which existing (or potential) proofs of security model HMAC as a keyed RO, a function that maps each key, message pair to an independent and uniform point. There are many examples of such settings. The HKDF scheme builds from HMAC a general-purpose key derivation function [31, 32] that uses as key a public, uniformly chosen salt. When used with a source of sufficiently high entropy, Krawczyk proves security using standard model techniques, but when not proves security assuming HMAC is a keyed RO [32]. PKCS#5 standardizes password-based key derivation functions that use HMAC with key being a (low-entropy) password [36]. Recent work provides the first proofs of security for the HMAC-based key derivation function, should HMAC be modeled as a keyed RO [9]. Ristenpart and Yilek [38], in the context of hedged cryptography [4], use HMAC in a setting whose cryptographic security models allow adversarially specified keys. Again, proofs model HMAC as a keyed RO.

As mentioned previously, we would expect a priori that one can show that HMAC is indistinguishable from a keyed RO even when the attacker can query arbitrary keys. Then one could apply the composition theorem of [33] to derive proofs of security for the settings just discussed.

WEAK KEY PAIRS IN HMAC. We are the first to observe that HMAC has weak key pairs. First, there exist $K \neq K'$ for which $\text{HMAC}(K, M) = \text{HMAC}(K', M)$. These pairs of keys arise because of HMAC's ambiguous encoding of differing-length keys. Trivial examples of such "colliding" keys include any K, K' for which either $|K| < d$ and $K' = K \parallel 0^s$ (for any $1 \leq s \leq d - |K|$), or $|K| > d$ and $K' = H(K)$. Colliding keys enable an easy attack that distinguishes $\text{HMAC}(\cdot, \cdot)$ from a random function $\mathcal{R}(\cdot, \cdot)$, which also violates the indistinguishability of HMAC. On the other hand, as long as H is collision-resistant, two keys of the same length can never collide. Still, even if we restrict attention to (non-colliding) keys of a fixed length, there still exist weak key pairs, but of a different form that we term ambiguous. An example of an ambiguous key pair is K, K' of length d bits such that $K \oplus \text{ipad} = K' \oplus \text{opad}$. Because the second least significant bit of ipad and opad differ (see Section 4) and assuming $d > n - 2$, ambiguous key pairs of a fixed length k only exist for $k \in \{d - 1, d\}$. The existence of ambiguous key pairs in HMAC leads to negative results like those given for H^2 . In particular, we straightforwardly extend the H^2 distinguisher to give one that lower bounds the number of queries any indistinguishability simulator must make for HMAC.

UPPER BOUNDS FOR HMAC. Fortunately, it would seem that weak key pairs do not arise in typical applications. Using HMAC with keys of some fixed bit length smaller than $d - 1$ avoids weak key pairs. This holds for several applications, for example the recommendation with HKDF is to use n -bit uniformly chosen salts as HMAC keys. This motivates finding positive results for HMAC when one avoids the corner cases that allow attackers to exploit weak key pairs.

Indeed, as our main positive result, we prove that, should H be a RO or an MD hash with ideal compression functions, *HMAC is indistinguishable from a keyed RO for all distinguishers that do not query weak key pairs*. Our result holds for the case that the keys queried are of length d or less. This upper bound enjoys the best, birthday-bound level of concrete security possible (up to small constants),

and provides the *first positive result about the indistinguishability of the HMAC construction*.

1.3 Discussion

The structural properties within H^2 and HMAC are, in theory, straightforward to avoid. Indeed, as mentioned above, Coron et al. [14] prove indistinguishable from a RO the construction $H^2(0^d \parallel M)$ where H is MD using a compression function with block size d bits and chaining value length $n \leq d$ bits. Analogously, our positive results about HMAC imply as a special case that $\text{HMAC}(K, M)$, for any fixed constant K , is indistinguishable from a RO.

We emphasize that we are unaware of any deployed cryptographic application for which the use of H^2 or HMAC leads to a vulnerability. Still, our results show that future applications should, in particular, be careful when using HMAC with keys which are under partial control of the attacker. More importantly, our results demonstrate the importance of provable security in the design of hash functions (and elsewhere in cryptography), as opposed to the more common “attack-fix” cycle. For example, the hash-of-hash suggestion of Ferguson and Schneier [22] was motivated by preventing the extension attack. Unfortunately, in so doing they accidentally introduced a more subtle (although less dangerous) attack, which was not present on the original design.¹ Indeed, we discovered the subtlety of the problems within H^2 and HMAC, including our explicit attacks, only after attempting to prove indistinguishability of these constructions (with typical, good bounds). In contrast, the existing indistinguishability proofs of (seemingly) small modifications of these hash functions, such as $H^2(0^d \parallel M)$ [14], provably rule out these attacks.

1.4 Prior Work

There exists a large body of work showing hash functions are indistinguishable from a RO (c.f., [1, 7, 12–14, 18, 19, 26]), including analyses of variants of H^2 and HMAC. As mentioned, a construction called HMAC was analyzed in [14] but this construction is not HMAC as standardized. Krawczyk [32] suggests that the analysis of $H^2(0 \parallel M)$ extends to the case of HMAC, but does not offer proof.² HMAC has received much analysis in other contexts. Proofs of its security as a pseudorandom function under reasonable assumptions appear in [3, 5]. These rely on keys being uniform and secret, making the analyses inapplicable for other settings. Analysis of HMAC’s security as a randomness extractor appear in [17, 23]. These results provide strong information theoretic guarantees that HMAC can be used as a key derivation function, but only in settings where the source has a relatively large amount of min-entropy. This requirement makes the analyses insufficient to argue security in many settings of practical importance. See [32] for further discussion.

2 Preliminaries

NOTATION AND GAMES. We denote the empty string by λ . If $|X| < |Y|$ then $X \oplus Y$ signifies that the X is padded with $|Y| - |X|$ zeros first. For set \mathcal{X} and value x , we write $\mathcal{X} \stackrel{\cup}{\leftarrow} x$ to denote $\mathcal{X} \leftarrow \mathcal{X} \cup \{x\}$. For non-empty sets $Keys$, Dom , and Rng with $|Rng|$ finite, a random oracle $f: Keys \times Dom \rightarrow Rng$ is a function taken randomly from the space of all possible functions $Keys \times Dom \rightarrow Rng$. We will sometimes refer to random oracles as keyed when $Keys$ is non-empty, whereas we omit the first parameter when $Keys = \emptyset$.

¹We note the prescience of the proposers of H^2 , who themselves suggested further analysis was needed [22].

²Fortunately, the HKDF application of [32] seems to avoid weak key pairs, and thus our positive results for HMAC appear to validate this claim [32] for this particular application.

We use code-based games [10] to formalize security notions and within our proofs. In the execution of a game G with adversary \mathcal{A} , we denote by $G^{\mathcal{A}}$ the event that the game outputs true and by $\mathcal{A}^G \Rightarrow y$ the event that the adversary outputs y . Fixing some RAM model of computation, our convention is that the running time $\text{Time}(\mathcal{A})$ of an algorithm \mathcal{A} includes its code size. Queries are unit cost, and we will restrict attention to the absolute worst case running time which must hold regardless of queries are answered.

HASH FUNCTIONS. A *hash function* $H[P]: \text{Keys} \times \text{Dom} \rightarrow \text{Rng}$ is a family of functions from Dom to Rng , indexed by a set Keys , that possibly uses (black-box) access to an underlying primitive P (e.g., a compression function). We call the hash function *keyed* if Keys is non-empty, and *key-less* otherwise. (In the latter case, we omit the first parameter.) We assume that the number of applications of P in computing $H[P](K, M)$ is the same for all K, M with the same value of $|K| + |M|$. This allows us to define the *cost* of computing a hash function $H[P]$ on a key and message whose combined length is ℓ , denoted $\text{Cost}(H, \ell)$, as the number of calls to P required to compute $H[P](K, M)$ for K, M with $|K| + |M| = \ell$. For a keyed random oracle $\mathcal{R}: \text{Keys} \times \text{Dom} \rightarrow \text{Rng}$, we fix the convention that $\text{Cost}(\mathcal{R}, \ell) = 1$ for any ℓ for which there exists a key $K \in \text{Keys}$ and message $M \in \text{Dom}$ such that $|K| + |M| = \ell$.

A *compression function* is a hash function for which $\text{Dom} = \{0, 1\}^n \times \{0, 1\}^d$ and $\text{Rng} = \{0, 1\}^n$ for some numbers $n, d > 0$. Our focus will be on keyless compression functions, meaning those of the form $f: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^n$. Our results lift in a straightforward way to the dedicated-key setting [8]. The ℓ -th iterate of $H[P]$ is denoted $H^\ell[P]$, and defined for $\ell > 0$ by $H^\ell[P](X) = H[P](H[P](\dots H[P](X))\dots)$ where the number of applications of H is ℓ . We let $H^0[P](X) = X$. We will often write H instead of $H[P]$ when the underlying primitive P is clear or unimportant.

MERKLE-DAMGÅRD. Let $\text{Pad}: \{0, 1\}^{\leq L} \rightarrow (\{0, 1\}^n)^+$ be an injective padding function. The one used in many of the hash functions within the SHA family outputs $M \parallel 10^r \parallel \langle |M| \rangle_{64}$ where $\langle |x| \rangle_{64}$ is the encoding of the length of M as a 64-bit string and r is the smallest number making the length a multiple of d . This makes $L = 2^{64} - 1$. The function $\text{MD}[f]: (\{0, 1\}^n)^+ \rightarrow \{0, 1\}^n$ is defined as

$$\text{MD}[f](M) = f(f(\dots f(f(\text{IV}, M_1), M_2), \dots), M_k)$$

where $|M| = kd$ and $M_1 \parallel \dots \parallel M_k$. The function $\text{SMD}[f]: \{0, 1\}^{\leq L} \rightarrow \{0, 1\}^n$ is defined as $\text{SMD}[f](M) = \text{MD}[f](\text{Pad}(M))$.

INDIFFERENTIABILITY FROM A RO. Let $\mathcal{R}: \text{Keys} \times \text{Dom} \rightarrow \text{Rng}$ be a random oracle. Consider a hash construction $H[P]: \text{Keys} \times \text{Dom} \rightarrow \text{Rng}$ from an ideal primitive P . Let game $\text{Real}_{H[P]}$ be the game whose main procedure runs an adversary $\mathcal{A}^{\text{Func, Prim}}$ and returns the bit that \mathcal{A} outputs. The procedure **Func** on input $K \in \text{Keys}$ and $M \in \text{Dom}$ returns $H[P](K, M)$. The procedure **Prim** on input X returns $P(X)$. For a simulator \mathcal{S} , let game $\text{Ideal}_{\mathcal{R}, \mathcal{S}}$ be the game whose main procedure runs an adversary $\mathcal{A}^{\text{Func, Prim}}$ and returns the bit that \mathcal{A} outputs. The procedure **Func** on input $K \in \text{Keys}$ and $M \in \text{Dom}$ returns $\mathcal{R}(K, M)$. The procedure **Prim** on input X returns $\mathcal{S}^{\mathcal{R}}(X)$. The indifferenciability advantage of \mathcal{D} is defined as

$$\text{Adv}_{H[P], \mathcal{R}, \mathcal{S}}^{\text{indiff}}(\mathcal{D}) = \Pr \left[\text{Real}_{H[P]}^{\mathcal{D}} \Rightarrow y \right] - \Pr \left[\text{Ideal}_{\mathcal{R}, \mathcal{S}}^{\mathcal{D}} \Rightarrow y \right] .$$

We focus on simulators that must work for any adversary, though our negative results extend as well to the weaker setting in which the simulator can depend on the adversary. The total query cost σ of an adversary \mathcal{D} is the cumulative cost of all its **Func** queries plus q_2 . (This makes σ the total number of P uses in game $\text{Real}_{H[P]}$. In line with our worst-case conventions, this means the same maximums hold in $\text{Ideal}_{\mathcal{R}, \mathcal{S}}$ although here it does not translate to P applications.)

We note that when Keys is non-empty, indifferenciability here follows [8] and allows the distin-

<p>main $\text{PrA}_{H,P,\mathcal{E}}$ $\mathbf{V} \leftarrow \perp ; \alpha \leftarrow \lambda$ $(K, M) \leftarrow_{\mathcal{S}} \mathcal{A}^{\text{Prim,Ext}}$ $z \leftarrow H[P](K, M)$ $\text{Ret } ((K, M) \neq \mathbf{V}[z] \wedge \mathbf{Q}[z] = 1)$</p>	<p>procedure $\text{Prim}(u)$: $v \leftarrow P(u) ; \alpha \leftarrow \alpha \parallel (u, v) ; \text{Ret } v$</p> <p>procedure $\text{Ext}(z)$: $\mathbf{Q}[z] \leftarrow 1 ; \mathbf{V}[z] \leftarrow \mathcal{E}(z, \alpha) ; \text{Ret } \mathbf{V}[z]$</p>
---	---

Figure 1: The game for defining preimage awareness.

guisher to choose keys during an attack. This reflects the desire for a keyed hash function to be indistinguishable from a keyed random oracle for arbitrary uses of the key input.

PREIMAGE AWARENESS [19]. An *extractor* is a deterministic algorithm that takes as input a point z and a string α (called the advice string) and returns a string or \perp . The game $\text{PrA}_{H[P],\mathcal{E}}$ shown in Figure 1 defines the security experiment for preimage awareness. For a function H using an ideal primitive P , an extractor \mathcal{E} , and an adversary \mathcal{A} we define the advantage function

$$\text{Adv}_{H[P],\mathcal{E}}^{\text{pra}}(\mathcal{A}) = \Pr \left[\text{PrA}_{H[P],\mathcal{E}}^{\mathcal{A}} \Rightarrow \text{true} \right] .$$

3 Second Iterates and their Security

Our investigation begins with the second iterate of a hash function, meaning $H^2(M) = H(H(M))$ where $H: \text{Dom} \rightarrow \text{Rng}$ for sets $\text{Dom} \supseteq \text{Rng}$. For simplicity, let $\text{Rng} = \{0, 1\}^n$ and assume that H is itself modeled as a RO. Is H^2 good in the sense of being like a RO? Given that we are modeling H as a RO, we would expect that the answer would be “yes”. The truth is more involved. As we’ll see in Section 4, similar subtleties exist in the case of the related HMAC construction.

We start with the following observations. When computing $H^2(M)$ for some M , we refer to the value $H(M)$ as an intermediate value. Then, we note that the value $Y = H^2(M)$ is in fact the intermediate value used when computing $H^2(X)$ for $X = H(M)$. Given $Y = H^2(M)$, then, one can compute $H^2(H(M))$ directly by computing $H(Y)$. That the hash value Y is also the intermediate value used in computing the hash of another message is cause for concern: other hash function constructions that are indifferentiable from a RO (c.f., [2, 7, 8, 14, 26]) explicitly attempt to ensure that outputs are *not* intermediate values (with overwhelming probability over the randomness of the underlying idealized primitive). Moreover, prior constructions for which hash values are intermediate values have been shown to *not* be indifferentiable from a RO. For example Merkle-Damgård-based iterative hashes fall to extension attacks [14] for this reason. Unlike with Merkle-Damgård, however, it is not immediately clear how an attacker might abuse the structure of H^2 .

3.1 Hash chains using Second Iterates

We turn our attention to hash chains, where potential issues arise. Hash chains, formed by repeatedly applying the hash function to some message, are used in a variety of settings such as password-based cryptography [36] and forward-secure pseudorandom number generators [11]. For a hash function H , we define a hash chain $Y = (Y_0, \dots, Y_\ell)$ to be a sequence of $\ell + 1$ values where Y_0 is a message and $Y_i = H(Y_{i-1})$ for $1 \leq i \leq \ell$. Likewise when using H^2 a hash chain $Y = (Y_0, \dots, Y_\ell)$ is a sequence of $\ell + 1$ values where Y_0 is a message and $Y_i = H^2(Y_{i-1})$ for $1 \leq i \leq \ell$. We refer to Y_0 as the *start* of the hash chain and Y_ℓ as the *end*. Two chains Y, Y' are *non-overlapping* if no value in one chain occurs in the other, meaning $Y_i \neq Y'_j$ for all $0 \leq i \leq \ell \leq j$.

For any hash function and given the start and end of a hash chain $Y = (Y_0, \dots, Y_\ell)$, one can readily




Figure 2: Diagram of two hash chains $Y = (Y_0, \dots, Y_\ell)$ and $Y' = (Y'_0, \dots, Y'_\ell)$ for hash function H^2 .

compute the start and end of a new chain with just two hash calculations. That is, set $Y'_0 \leftarrow H(Y_0)$ and $Y'_\ell \leftarrow H(Y_\ell)$. However, the chain $Y' = (Y'_0, \dots, Y'_\ell)$ and the chain Y overlap. For good hash functions (i.e., ones that behave like a RO) computing the start and end of a non-overlapping chain given the start and end of a chain Y_0, Y_ℓ requires at least ℓ hash computations (assuming $\ell \ll 2^{n/2}$).

Now consider H^2 . Given the start and end of a chain $Y = (Y_0, \dots, Y_\ell)$, one can readily compute a non-overlapping chain $Y' = (Y'_0, \dots, Y'_\ell)$ using just two hash computations instead of the expected 2ℓ computations. Namely, let $Y'_0 \leftarrow H(Y_0)$ and $Y'_\ell \leftarrow H(Y_\ell)$. Then these are the start and end of the chain $Y' = (Y'_0, \dots, Y'_\ell)$ because

$$H^{2\ell}(Y'_0) = H^{2\ell}(H(Y_0)) = H(H^{2\ell}(Y_0))$$

which we call the chain-shift property of H^2 . Moreover, assuming H is itself a RO outputting n -bit strings, the two chains Y, Y' do not overlap with probability at least $1 - (2\ell + 2)^2 / 2^n$. Figure 2 provides a pictorial diagram of the two chains Y and Y' .

The above discussion points out a way in which H^2 and a RO differ. We now formalize this difference via a toy security game. Game $\text{CHAIN}_{H[P],n,\ell}$ is shown in Figure 3. It tasks an attacker with computation of a hash chain of length ℓ for a hash function $H[P]$. The attacker is additionally given the start and end of an honestly generated hash chain, and the attacker can make at most $\ell \cdot \text{Cost}(H, n) - 1$ queries to the primitive underlying the hash construction — one less query than that needed to compute a new hash chain directly. The attacker succeeds if it can nevertheless compute the start and end of a hash chain that does not overlap with the honestly generated one. We define the advantage of an $\text{CHAIN}_{H[P],n,\ell}$ -adversary \mathcal{A} by

$$\text{Adv}_{H[P],n,\ell}^{\text{chain}}(\mathcal{A}) = \Pr \left[\text{CHAIN}_{H[P],n,\ell}^{\mathcal{A}} \Rightarrow \text{true} \right].$$

We consider below two constructions:

- The second iterate $H^2[P]$ defined by computing $P(P(M))$ and where $P: \text{Dom} \rightarrow \{0, 1\}^n$ that is a random oracle. To win, the attacker must somehow compute a length ℓ hash chain for $H^2[P]$ using at most $2\ell - 1$ queries to P .
- The hash function is a random oracle, formally $H[P]$ is defined by $P(M)$ where $P: \text{Dom} \rightarrow \{0, 1\}^n$ is a RO. To win, the attacker must somehow compute a length ℓ hash chain for $H[P]$ using at most $\ell - 1$ queries to P .

CHAIN SECURITY OF A RO. We start with the latter bullet point: We show that no attacker can achieve good advantage in the $\text{CHAIN}_{H,P,n,\ell}$ game for $H[P](M) = P(M)$ when $\ell \ll 2^{n/2}$.

Claim 3.1 Let $H[P]$ be the hash that applies a RO $P: \text{Dom} \rightarrow \{0, 1\}^n$ to each input with $\text{Dom} \supseteq \{0, 1\}^n$ and let $\ell \in [1 .. 2^n/6]$. Then for any \mathcal{A} it holds that $\text{Adv}_{H,P,n,\ell}^{\text{chain}}(\mathcal{A}) \leq \frac{(3\ell)^2}{2^n}$. \square

<p>main CHAIN_{H,P,n,ℓ}:</p> <p>$i \leftarrow 0$ $Y_0 \leftarrow \mathfrak{s} \{0, 1\}^n$; For $i = 1$ to ℓ do $Y_i \leftarrow H[P](Y_{i-1})$ $(Y_0^*, Y_\ell^*) \leftarrow \mathfrak{s} \mathcal{A}^{\text{Prim}}(Y_0, Y_\ell)$ $Y'_0 \leftarrow Y_0^*$; For $i = 1$ to ℓ do $Y'_i \leftarrow H[P](Y'_{i-1})$ $\mathcal{Y} \leftarrow \{Y_0, \dots, Y_\ell\}$ $\mathcal{Y}' \leftarrow \{Y'_0, \dots, Y'_\ell\}$ If $(Y'_\ell = Y_\ell^*) \wedge ((\mathcal{Y} \cap \mathcal{Y}') = \emptyset)$ then Ret true Ret false</p>	<p>procedure Prim(M):</p> <p>$i \leftarrow i + 1$ If $i \geq \ell \cdot \text{Cost}(H, n)$ then Ret \perp Ret $P(M)$</p>
---	--

Figure 3: The chain-making game.

Proof: We can assume that \mathcal{A} makes exactly $\ell - 1$ queries to Prim. Let $\tilde{P}: \text{Dom} \rightarrow \{0, 1\}^n$ map each input to an output selected uniformly without replacement. Then a standard argument shows that

$$\mathbf{Adv}_{H,P,n,\ell}^{\text{chain}}(\mathcal{A}) \leq \mathbf{Adv}_{H,\tilde{P},n,\ell}^{\text{chain}}(\mathcal{A}) + \frac{(3\ell - 1)^2}{2^n}. \quad (1)$$

To win the game gainst $H[\tilde{P}]$, the adversary \mathcal{A} must output Y_0^*, Y_ℓ^* such that the associated chain Y'_0, \dots, Y'_ℓ is non-overlapping with Y_0, \dots, Y_ℓ and with $Y'_0 = Y_0^*$ and $Y'_\ell = Y_\ell^*$. But \mathcal{A} observes $\ell - 1$ of the ℓ values Y'_0, \dots, Y'_ℓ . Let $i \in [1.. \ell]$ be the index such that Y'_i was not returned to \mathcal{A} by Prim. First consider the case that $i = \ell$. Then, Y'_ℓ is a fresh choice made after \mathcal{A} finishes execution, and so $Y'_\ell = Y_\ell^*$ with probability $1/(2^n - (3\ell - 1))$. (Since we are disallowing collisions, Y'_ℓ is uniformly selected from a set of size $2^n - (3\ell - 1)$.) Now consider the case that $i < \ell$. Then, for the chain to be completed, the value Y'_i , chosen after \mathcal{A} finishes executing, must equal the value M queried by \mathcal{A} that had response Y'_{i+1} . But this occurs, again, with probability at most $1/(2^n - (3\ell - 1))$. Since we restricted attention to $1 \leq \ell \leq 2^n/6$, it holds that $\mathbf{Adv}_{H,\tilde{P},n,\ell}^{\text{chain}}(\mathcal{A}) \leq 2/2^n$. Substituting into (1) above yields the claim. \blacksquare

ATTACK AGAINST H^2 . Now let the hash function be $H^2[P] = P(P(M))$ for $P: \text{Dom} \rightarrow \{0, 1\}^n$ a RO. We give an adversary \mathcal{A} that wins the CHAIN_{H²[P],n,ℓ} game with probability $1 - \ell^2/2^n$. Let \mathcal{A} work as follows. Upon execution with inputs Y_0, Y_ℓ , it queries $Y_0^* \leftarrow \text{Prim}(Y_0)$ and $Y_\ell^* \leftarrow \text{Prim}(Y_\ell)$. It then returns (Y_0^*, Y_ℓ^*) . Then, letting $Y'_0 = Y_0^*$ and Y'_1, \dots, Y'_ℓ be the values computed in the main procedure of the CHAIN_{H²[P],n,ℓ} game, we see that the chain-shift property of H^2 means that $Y'_\ell = Y_\ell^*$. Moreover, no element in Y'_0, \dots, Y'_ℓ collides with an element in Y_0, \dots, Y_ℓ with probability $1 - \ell^2/2^n$, this probability being over the random coins used by P . The adversary therefore achieves the stated advantage.

DISCUSSION. All the above exhibits a way in which H^2 fails to behave like a RO. Moreover, the attack against H^2 generalizes to other hash function constructions that are second iterates, for example when using the second iterate of the Merkle-Damgård construction [16, 34].

But what does it imply about indifferentiability? Recall that the composition theorem of Maurer et al. [33], as discussed further by Ristenpart et al. [37], states —informally speaking— that a cryptosystem secure relative to any single-stage³ game in the ROM will remain secure when replacing the RO with a hash construction that is indifferentiable from a RO. We have above given a game in which no attacker succeeds in the ROM, but there exists an adversary that succeeds when the RO is replaced with H^2 . This might seem to directly rule out H^2 being indifferentiable from a RO. But we must

³The security games we consider fall into this category, which just mandates that the adversary maintains state throughout the entire experiment. See [37] for further discussion.

be careful. The gap between security and insecurity above is conditioned upon the limited number of queries, and such limitations must be carefully handled when indistinguishability is considered. In particular, the results above only would lead to contradicting a positive indistinguishability result using a simulator that does not make so many queries that the bound in Claim 3.1 becomes close to one.

We will tackle these subtleties head-on in Section 3.3, by giving a distinguisher that differentiates H^2 from a RO for any simulator that does not make sufficiently many queries. Before that, we first explore the implications of the surfaced structural property of H^2 : Might it lead to vulnerabilities in applications?

3.2 An Example (Vulnerable) Application: Mutual Proofs of Work

In the last section we saw that the second iterate fails to behave like a RO in the context of hash chains. But the security game detailed in the last section may seem far removed from real protocols. For example, it's not clear where an attacker would be tasked with computing hash chains in a setting where it, too, was given an example hash chain. We suggest that just such a setting could arise in protocols in which parties want to assert to each other, in a verifiable way, that they performed some amount of computation. Such a setting could arise when parties must (provably) compare assertions of computational power, as when using cryptographic puzzles [20, 21, 27, 28, 39, 41]. Or this might work when trying to verifiably calibrate differing computational speeds of the two parties' computers. We refer to this task as a *mutual proof of work*.

MUTUAL PROOFS-OF-WORK. For the sake of brevity, we present an example hash-chain-based protocol and dispense with a more general treatment of mutual proofs of work. Consider the two-party protocol shown in the left diagram of Figure 4. Each party initially chooses a random nonce and sends it to the other. Then, each party computes a hash chain of some length—chosen by the computing party—starting with the nonce chosen by the other party, and sends the chain's output along with the chain's length to the other party. At this point, both parties have given a witness that they performed a certain amount of work. So now, each party checks the other's asserted computation, determining if the received value is the value resulting from chaining together the indicated number of hash applications and checking that the hash chains used by each party are non-overlapping. Note that unlike puzzles, which require fast verification, here the verification step is as costly as puzzle solution.

The goal of the protocol is to ensure that the other party did compute exactly their declared number of iterations. Slight changes to the protocol would lead to easy ways of cheating. For example, if during verification the parties did not check that the chains are non-overlapping, then \mathcal{P}_2 can easily cheat by choosing X_1 so that it can reuse a portion of the chain computed by \mathcal{P}_1 .

Security would be achieved should no cheating party succeed at convincing an honest party using less than ℓ_1 (resp. ℓ_2) work to compute Y_1 (resp. Y_2). The game $\text{POW}_{H[P],n,\ell_1}$ formalizes this security goal for a cheating \mathcal{P}_2 ; see the right portion of Figure 4. We let $\text{Adv}_{H[P],n,\ell_1}^{\text{pow}}(\mathcal{A}) = \Pr \left[\text{POW}_{H[P],n,\ell_1}^{\mathcal{A}} \right]$. Note that the adversary \mathcal{A} only wins should it make $q < \ell_2 \cdot \text{Cost}(H, n)$ queries, where ℓ_2 is the value it declared and $\text{Cost}(H)$ is the cost of computing H . Again we will consider both the hash function $H[P](M) = P(M)$ that just applies a RO P and also $H^2[P](M) = P(P(M))$, the second iterate of a RO. In the former case the can make only $\ell_2 - 1$ queries and in the latter case $2\ell_2 - 1$.

When $H[P](M) = P(M)$, no adversary making $q < \ell_2$ queries to Prim can win the $\text{POW}_{H[P],n,\ell_1}$ game with high advantage. Intuitively, the reason is that, despite being given X_1 and Y_1 where $Y_1 = P^{\ell_1}(X_1)$, a successful attacker must still compute a full ℓ_2 -length chain and this requires ℓ_2 calls to P . A treatment of this follows closely Claim 3.1 and its proof, and so we omit the details. Intuitively, the reason is that, despite being given X_1 and Y_1 where $Y_1 = P^{\ell_1}(X_1)$, a successful attacker




Figure 4: Example protocol (**left**) and adversarial \mathcal{P}_2 security game (**right**) for mutual proofs of work.

must still compute a full ℓ_2 -length chain and this requires ℓ_2 calls to P .

ATTACK AGAINST ANY SECOND ITERATE. Now let us analyze this protocol's security when we use as hash function $H^2[P] = P(P(M))$ for a RO $P: \text{Dom} \rightarrow \text{Rng}$ with $\text{Rng} \subseteq \text{Dom}$. We can abuse the chain-shift property of H^2 in order to win the POW $_{H^2,P,n,\ell_1}$ game for any $n > 0$ and $\ell_1 > 2$. Our adversary \mathcal{A} works as follows. It receives X_2 and then chooses its nonce as $X_1 \leftarrow \text{Prim}(X_2)$. When it later receives $Y_1 = P^{2\ell_1}(X_1)$, the adversary proceeds by setting $\ell_2 = \ell_1 + 1$ and setting $Y_2 \leftarrow \text{Prim}(Y_1)$. Then by the chain-shift property we have that

$$Y_2 = P(Y_1) = P(P^{2\ell_1}(X_1)) = P(P^{2\ell_1}(P(X_2))) = P^{2\ell_1+2}(X_2) = P^{2\ell_2}(X_2).$$

The two chains will be non-overlapping with high probability (over the coins used by P). Finally, \mathcal{A} makes only 2 queries to Prim, so the requirement that $q < 2\ell_2$ is met whenever $\ell_1 > 1$.

DISCUSSION. As far as we are aware, mutual proofs of work have not before been considered — the concept may indeed be of independent interest. A full treatment is beyond the scope of this work. We also note that, of course, it is easy to modify the protocols using H^2 to be secure. Providing secure constructions was not our goal, rather we wanted to show protocols which are insecure using H^2 but secure when H^2 is replaced by a monolithic RO. This illustrates how, hypothetically, the structure of H^2 could give rise to subtle vulnerabilities in an application.

3.3 An Indifferentiability Distinguisher for any Second Iterate

In this section we prove that *any* indifferentiability proof for the double iterate H^2 is subject to inherent quantitative limitations. Recall that indifferentiability asks for a simulator \mathcal{S} such that no adversary can distinguish between the pair of oracles $H^2[P], P$ and \mathcal{R}, \mathcal{S} where P is some underlying ideal primitive and \mathcal{R} is a RO with the same domain and range as H^2 . The simulator can make queries to \mathcal{R} to help it in its simulation of P . Concretely, building on the ideas behind the above attacks in the context of hash chains, we show that in order to withstand a differentiating attack with q queries, any simulator for $H^2[P]$, for *any* hash construction $H[P]$ with output length n , must issue *at least* $\Omega(\min\{q^2, 2^{n/2}\})$ queries to the RO \mathcal{R} . As we explain below, such a lower bound severely limits the concrete security level which can be inferred by using the composition theorem for indifferentiability, effectively neutralizing the benefits of using indifferentiability in the first place.

THE DISTINGUISHER. In the following, we let $H = H[P]$ be an *arbitrary* hash function with n -bit outputs relying on a primitive P , such as a fixed input-length random oracle or an ideal cipher. We are therefore addressing an arbitrary second iterate, and not focusing on some particular ideal primitive P (such as a RO as in previous sections) or construction H . Indeed, H could equally well be Merkle-Damgård and P an ideal compression function, or H could be any number of indifferentiable hash constructions using appropriate ideal primitive P .

Recall that Func and Prim are the oracles associated with construction and primitive queries to $H^2 = H^2[P]$ and P , respectively. Let w, ℓ be parameters (for now, think for convenience of $w = \ell$). The attacker $\mathcal{D}_{w,\ell}$ starts by issuing ℓ queries to Func to compute a *chain* of n -bit values $(x_0, x_1, \dots, x_\ell)$ where $x_i = H^2(x_{i-1})$ and x_0 is a random n -bit string. Then, it also picks a random index $j \in [1..w]$, and creates a list of n -bit strings $\mathbf{u}[1], \dots, \mathbf{u}[w]$ with $\mathbf{u}[j] = x_\ell$, and all remaining $\mathbf{u}[i]$ for $i \neq j$ are chosen uniformly and independently. Then, for all $i \in [1..w]$, the distinguisher $\mathcal{D}_{w,\ell}$ proceeds in asking all Prim queries in order to compute $\mathbf{v}[i] = H(\mathbf{u}[i])$. Subsequently, the attacker compute $y_0 = H(x_0)$ via Prim queries, and also computes the chain $(y_0, y_1, \dots, y_\ell)$ such that $y_i = H^2(y_{i-1})$ by making ℓ Func queries. Finally, it decides to output 1 if and only if $y_\ell = \mathbf{v}[j]$ and x_ℓ as well as $\mathbf{v}[i]$ for $i \neq j$ are not in $\{y_0, y_1, \dots, y_\ell\}$. The attacker $\mathcal{D}_{w,\ell}$ therefore issues a total of 2ℓ Func queries and $(2w + 1) \cdot \text{Cost}(H, n)$ Prim queries.

In the real-world experiment, the distinguisher $\mathcal{D}_{w,\ell}$ outputs 1 with very high probability, as the condition $y_\ell = \mathbf{v}[j]$ *always* holds by the chain-shifting property of H^2 . In fact, the only reason for \mathcal{D} outputting 0 is that one of x_ℓ and $\mathbf{v}[i]$ for $i \neq j$ incidentally happens to be in $\{y_0, y_1, \dots, y_\ell\}$. The (typically small) probability that this occurs obviously depends on the particular construction $H[P]$ at hand; it is thus convenient to define the shorthand

$$p(H, w, \ell) = \Pr [\{x_\ell, H(U_1), \dots, H(U_{w-1})\} \cap \{y_0, y_1, \dots, y_\ell\} \neq \emptyset] ,$$

where $x_0, y_0, x_1, \dots, y_{\ell-1}, x_\ell, y_\ell$ are the intermediate value of a chain of $2\ell + 1$ consecutive evaluations of $H[P]$ starting at a random n -bit string x_0 , and U_1, \dots, U_{w-1} are further independent random n -bit values. In Appendix A we prove that for $H[P] = P = \mathcal{R}$ for a random oracle $\mathcal{R} : \{0, 1\}^* \rightarrow \{0, 1\}^n$ we have $p(H, w, \ell) = \Theta((w\ell + \ell^2)/2^n)$. Similar reasoning can be applied to essentially all relevant constructions.

In contrast, in the ideal-world experiment, we expect the simulator to be completely ignorant about the choice of j as long as it does not learn x_0 , and in particular it does not know j while answering the Prim queries associated with the evaluations of $H(\mathbf{u}[i])$. Consequently, the condition required for $\mathcal{D}_{w,\ell}$ to output 1 appears to force the simulator, for all $i \in [1..w]$, to prepare a distinct chain of ℓ consecutive \mathcal{R} evaluations ending in $\mathbf{v}[i]$, hence requiring $w \cdot \ell$ random oracle queries.

The following theorem quantifies the advantage achieved by the above distinguisher $\mathcal{D}_{w,\ell}$ in differentiating against any simulator for the construction $H[P]$. Its proof is given in Appendix B.

Theorem 3.2 [Attack against H^2] Let $H[P]$ be an arbitrary hash construction with n -bit outputs, calling a primitive P , and let $\mathcal{R} : \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a random oracle. For all integer parameters $w, \ell \geq 1$, there exists an adversary $\mathcal{D}_{w,\ell}$ making 2ℓ Func -queries and $(w + 1) \cdot \text{Cost}(H, n)$ Prim -queries such that for all simulators \mathcal{S} ,

$$\text{Adv}_{H^2[P], \mathcal{R}, \mathcal{S}}^{\text{indiff}}(\mathcal{D}_{w,\ell}) \geq 1 - p(H, w, \ell) - \frac{5\ell^2}{2^{n+1}} - \frac{q_{\mathcal{S}}\ell}{2^n} - \frac{q_{\mathcal{S}}^2}{2^n} - \frac{q_{\mathcal{S}}}{w \cdot \ell} - \frac{1}{w} ,$$

where $q_{\mathcal{S}}$ is the overall number of \mathcal{R} queries by \mathcal{S} when replying to $\mathcal{D}_{w,\ell}$'s Prim queries. ■

DISCUSSION. We now elaborate on Theorem 3.2. If we consider the distinguisher $\mathcal{D}_{w,\ell}$ from Theorem 3.2, we observe that by the advantage lower bound in the theorem statement, if $\ell, w \ll 2^{n/4}$ and consequently $p(H, w, \ell) \approx 0$, the number of queries made by the simulator, denoted $q_{\mathcal{S}} = q_{\mathcal{S}}(2\ell, w + 1)$

must satisfy $q_S = \Omega(w \cdot \ell) = \Omega(q_1 \cdot q_2)$ to ensure a sufficiently small indifferenciability advantage. This in particular means that in the case where both q_1 and q_2 are large, the simulator must make a *quadratic* effort to prevent the attacker from distinguishing. Below, in Theorem 3.3, we show that this simulation effort is essentially optimal.

In many scenarios, this quadratic lower bound happens to be a problem, as we now illustrate. As a concrete example, let $\mathcal{SS} = (\text{key}, \text{sign}, \text{ver})$ be an arbitrary signature scheme signing n bits messages, and let $\widetilde{\mathcal{SS}}[\mathcal{R}] = (\widetilde{\text{key}}^{\mathcal{R}}, \widetilde{\text{sign}}^{\mathcal{R}}, \widetilde{\text{ver}}^{\mathcal{R}})$ for $\mathcal{R} : \{0, 1\}^* \rightarrow \{0, 1\}^n$ be the scheme obtained via the hash-then-sign paradigm such that $\widetilde{\text{sign}}^{\mathcal{R}}(sk, m) = \text{sign}(sk, \mathcal{R}(m))$. It is well known that for an adversary \mathcal{B} making q_{sign} signing and $q_{\mathcal{R}}$ random oracle queries, there exists an adversary \mathcal{C} making q_{sign} signing queries such that

$$\mathbf{Adv}_{\widetilde{\mathcal{SS}}[\mathcal{R}]}^{\text{uf-cma}}(\mathcal{B}^{\mathcal{R}}) \leq \frac{(q_{\text{sign}} + q_{\mathcal{R}})^2}{2^n} + \mathbf{Adv}_{\mathcal{SS}}^{\text{uf-cma}}(\mathcal{C}), \quad (2)$$

where $\mathbf{Adv}_{\widetilde{\mathcal{SS}}[\mathcal{R}]}^{\text{uf-cma}}(\mathcal{B}^{\mathcal{R}})$ and $\mathbf{Adv}_{\mathcal{SS}}^{\text{uf-cma}}(\mathcal{C})$ denote the respective advantages in the standard uf-cma game for security of signature schemes (with and without a random oracle, respectively). This in particular means that $\widetilde{\mathcal{SS}}$ is secure for q_{sign} and $q_{\mathcal{R}}$ as large as $\Theta(2^{n/2})$, provided \mathcal{SS} is secure for q_{sign} signing queries. However, let us now replace \mathcal{R} by $H^2[P]$ for an arbitrary construction $H = H[P]$. Then, for all adversaries \mathcal{A} making q_P queries to P and q_{sign} signing queries, we can combine the concrete version of the MRH composition theorem proven in [37] and (2) to infer that there exists an adversary \mathcal{C} and a distinguisher \mathcal{D} such that

$$\mathbf{Adv}_{\widetilde{\mathcal{SS}}[H^2[P]]}^{\text{uf-cma}}(\mathcal{A}^P) \leq \Theta\left(\frac{(q_{\text{sign}} \cdot q_P)^2}{2^n}\right) + \mathbf{Adv}_{\mathcal{SS}}^{\text{uf-cma}}(\mathcal{C}) + \mathbf{Adv}_{H^2[P], \mathcal{R}, \mathcal{S}}^{\text{indiff}}(\mathcal{D}),$$

where \mathcal{C} makes q_{sign} signing queries. Note that even if the term $\mathbf{Adv}_{H^2[P], \mathcal{R}, \mathcal{S}}^{\text{indiff}}(\mathcal{D})$ is really small, this new bound can only ensure security for the resulting signature scheme as long as $q_{\text{sign}} \cdot q_P = \Theta(2^{n/2})$, i.e., if $q_{\text{sign}} = q_P$, we only get security up to $\Theta(2^{n/4})$ queries, a remarkable loss with respect to the security bound in the random oracle model.

We note that of course this does *not* mean that $H^2[P]$ for a concrete H and P is unsuitable for a certain application, such as hash-then-sign. In fact, $H^2[P]$ may well be optimally collision resistant. However, our result shows that a sufficiently strong security level cannot be inferred from *any* indifferenciability statement via the composition theorem, taking us back to a direct ad-hoc analysis and completely losing the one main advantage of having indifferenciability in the first place.

3.4 Indifferenciability Upper Bound for a Second Iterate

Our negative results do not rule out positive results completely: there could be indifferenciability upper bounds, though for simulators that make around $\mathcal{O}(q^2)$ queries. Ideally, we would like upper bounds that match closely the lower bounds given in prior sections. We do so for the special case of $H^2[g](M) = g(g(M))$ for $g : \{0, 1\}^n \rightarrow \{0, 1\}^n$ being a RO. We have the following theorem.

Theorem 3.3 Let $q_1, q_2 \geq 0$ and $N = 2^n$. Let $g : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and $\mathcal{R} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be uniform random functions. Then there exists a simulator \mathcal{S} such that

$$\mathbf{Adv}_{G[g], \mathcal{R}, \mathcal{S}}^{\text{indiff}}(\mathcal{D}) \leq \frac{2((4q_1 + 3)q_2 + 2q_1)^2}{N} + \frac{2((4q_1 + 3)q_2 + 2q_1)(q_1 + q_2)}{(N - 2q_2 - 2q_1)}$$

for any adversary \mathcal{D} making at most q_1 queries to its left oracle and at most q_2 queries to its right oracle. Moreover, for each query answer that it computes, \mathcal{S} makes at most $3q_1 + 1$ queries to RO and runs in time $O(q_1)$. \square

procedure OnRightQuery(x):	subroutine FillInRungs(x, y)	subroutine MakeLadder(x)
$x_0 \leftarrow x$	$x_0 \leftarrow x, x_1 \leftarrow y$	$s_{-q_1} \leftarrow s \{0, 1\}^n$
For $i = 0$ to q_1	SetTable(\mathbf{g}, x_0, x_1)	For $i = -q_1$ to $q_1 - 1$
If $\mathbf{g}[x_i] \neq \perp$ then	For $i = 1$ to $2q_1 + 1$	SetTable($\mathbf{G}, s_i, \text{RO}(s_i)$)
$y \leftarrow \mathbf{G}^{-i}[\mathbf{g}[x_i]]$	SetTable($\mathbf{G}, x_{i-1}, \text{RO}(x_{i-1})$)	$s_{i+1} \leftarrow \mathbf{G}[s_i]$
If $y = \perp$ then Abort	$x_{i+1} \leftarrow \mathbf{G}[x_{i-1}]$	FillInRungs(x, s_0)
FillInRungs(x, y)	SetTable(\mathbf{g}, x_i, x_{i+1})	
Ret y		subroutine SetTable(\mathbf{T}, x, y)
SetTable($\mathbf{G}, x_i, \text{RO}(x_i)$)		If $\mathbf{T}[x]$ and $\mathbf{T}[x] \neq y$ then Ret
$x_{i+1} \leftarrow \mathbf{G}[x_i]$		$\mathbf{T}[x] = y$
MakeLadder(x)		$\mathbf{T}^{-1}[y] = x$
Ret $\mathbf{g}[x]$		

Figure 5: Simulator \mathcal{S} used in the proof of Theorem 3.3.

We note that the security bound of Theorem 3.3 is approximately $(q_1 q_2)^2 / N$, implying that Theorem 3.3 guarantees security up to $q_1 q_2 \approx 2^{n/2}$. For example, $G[g]$ would be indistinguishable to an attacker making $o(2^{\frac{1}{3}n})$ left oracle queries and $o(2^{\frac{1}{6}n})$ right oracle queries, or $O(1)$ left oracle queries and $o(2^{n/2})$ right oracle queries, or, oppositely, $o(2^{n/2})$ left oracle queries $O(1)$ right oracle queries, and so on. If all that is known is the attacker’s total number of queries, then all one can say is that security is guaranteed up to $\approx 2^{n/4}$ queries.

PROOF SKETCH. The simulator \mathcal{S} referred to in Theorem 3.3 is implemented by the procedure OnRightQuery(\cdot) shown in Figure 5. Here we give some intuition about \mathcal{S} —what it does and why—along with a brief overview of the full proof, which appears in Appendix D.

We use the notation $\text{RO}(x) \rightarrow y$ to indicate that the left oracle (in the ideal world) returns y when queried at x . Consider an attacker \mathcal{D} that starts by making a sequence of two queries $\text{RO}(t_{-2}) \rightarrow t_{-1}$, $\text{RO}(t_{-1}) \rightarrow t_0$ for a randomly chosen $t_{-2} \in \{0, 1\}^n$, and then queries the simulator at t_0 . At this point the simulator has no knowledge of any of the adversary’s queries—it has no “outside data”—and so a naïve simulator might answer a random value, which we denote as s_0 . In this case, if \mathcal{D} queries t_{-2} to the simulator, the simulator is stuck: it can see for itself that $\text{RO}^2(t_{-2}) = t_0$ (it knows both t_{-2} and t_0 , at this point), so to be consistent it knows that it must answer a value s_{-2} such that $\text{RO}^2(s_{-2}) = s_0$ (indeed, the identity $\text{RO} = g^2$ implies $\text{RO}^2(g) = g(\text{RO}^2)$), but finding such a value s_{-2} is computationally infeasible for the simulator, given that s_0 is random.

Considering that the above attack can be generalized by having \mathcal{D} build a long chain $t_{-\ell}, \dots, t_0$ such that $t_{i+1} = \text{RO}(t_i)$ before querying \mathcal{S} at t_0 and then at $t_{-\ell}$, we conclude that the simulator should prepare a chain of RO queries of length⁴ q_1 , the number of RO-queries available to \mathcal{D} . More precisely, when the simulator receives its first query t_0 it chooses a random value s_{-q_1} , computes $s_{i+1} = \text{RO}^i(s_i)$ for $-q_1 \leq i < 0$, and sets $g(t_0) = s_0$. Later, if \mathcal{D} queries the simulator at any point t_{-i} , $0 \leq i \leq q_1$, the simulator can notice that $\text{RO}^i(t_{-i}) = t_0$ (by iterating $\text{RO}(\cdot)$ on t_{-i} until it sees t_0), and then answer s_{-i} . (Note in passing that while the adversary \mathcal{D} only has q_1 RO-queries to make, the simulator does not know for which points the adversary has prepared a long chain. Hence by querying the simulator at $\approx q_2$ random points among which is inserted the endpoint t_0 of some RO-chain of length $\approx q_1$, the adversary can force the simulator to make $\approx q_1 q_2$ queries, to prepare a chain for each point. This is the intuition behind Theorem 3.2.)

⁴One could argue that a chain of length only $q_1/2$ is necessary, given that the adversary must also check the parallel chain to complete its attack; however, such fiddling is not worth the cost of complications in the proof.

Note that once the simulator determines $g(t_0) = s_0$ this also determines, for example, that $g(s_0) = g^2(t_0) = \text{RO}(t_0)$; thus the value $t_1 := g(s_0)$ is fixed; likewise the value $s_1 := g(t_1)$ is fixed because $s_1 = g^2(s_0) = \text{RO}(s_0)$, and so on. Thus if the adversary subsequently computes, for example, $t_{10} := \text{RO}^{10}(t_0)$ on its own and then queries t_{10} to the simulator, the simulator should notice that $t_{10} = \text{RO}^{10}(t_0)$ and answer $\text{RO}^{10}(s_0)$. The simulator can simplify its life in this regard if, right after setting $g(t_0) = s_0$, it precomputes the values t_1, \dots, t_{q_1+1} and s_1, \dots, s_{q_1} where $t_i := \text{RO}^i(t_0)$, $s_i := \text{RO}^i(s_0)$, setting $g(t_i) = s_i$ and $g(s_i) = t_{i+1}$ in the process (see Figure 6). Then, subsequently, it will “automatically” know that $g(t_{10}) = s_{10}$ without having to “notice” that $t_{10} = \text{RO}^{10}(t_0)$.

We call the sequence of values $t_0, s_0, t_1, \dots, s_{q_1}, t_{q_1+1}$ a *g-chain*, since $g(t_i) = s_i$ and $g(s_i) = t_{i+1}$. We note such a chain is just long enough that \mathcal{D} cannot compute the endpoint t_{q_1+1} “on its own” (using only RO queries), having only queried the simulator at s_0 , and at no further points along the chain⁵. Moreover, each time \mathcal{D} queries \mathcal{S} at a point in such a chain, \mathcal{S} extends the chain to always keep the furthest point of the chain at least $(q_1 + 1)$ RO queries away from the last input queried by \mathcal{D} , thus keeping the *g-chain*’s endpoint permanently unknown to \mathcal{D} .

Altogether, three possible scenarios can play out when our simulator \mathcal{S} answers a query x : (a) if \mathcal{S} finds that it has already chosen a value for $g(x)$, it returns this value, and possibly extends the *g-chain* containing the query; (b) otherwise, (for this case picture x as t_{-i} above, for some $1 \leq i \leq \ell$), if \mathcal{S} finds that $g(\text{RO}^i(x))$ is defined for some $1 \leq i \leq q_1$, it answers $y := \text{RO}^{-i}(g(\text{RO}^i(x)))$ if it can compute this value from the previous queries it has made to RO, and aborts otherwise if it cannot; moreover, if it does not abort, then using its newly defined input-output pair $g(x) = y$ it precomputes a *g-chain* of length $2q_1 + 1$ starting at x (this chain will actually “rear-end” a pre-existing *g-chain*, with no damage, since $g(\text{RO}^i(x))$ is already defined); (c) if neither (a) nor (b) occur, the simulator computes a chain $s_{-q_1}, \dots, s_{-1}, s_0, \dots, s_{q_1}$ where s_{-q_1} is chosen randomly and where $s_{-q_1+i} = \text{RO}^i(s_{-q_1})$ for $i \geq 1$, then computes⁶ the chain t_0, \dots, t_{q_1+1} where $t_i = \text{RO}^i(x)$, and finally sets⁷ $g(t_i) = s_i$, $g(s_i) = t_{i+1}$ for $0 \leq i \leq q_1$, before returning $s_0 (= g(t_0))$.

We call the query structure resulting from scenario (c) above a *ladder*; see Figure 6 (A). When the adversary queries a point t_{-i} to the simulator, where $\text{RO}^i(t_{-i}) = t_0$, the simulator “fills in” the values $g(t_{-j}) = s_{-j}$ and $g(s_{-j}) = t_{-j+1}$ for $-i \leq -j < 0$ (Figure 6 (B)), corresponding to scenario (b) above. Note that if the adversary queries a point s_{-j} to the simulator before querying any point t_{-k} with $k \geq j$, the simulator will abort, being unable to invert RO beyond the leftmost point t_{-i} such that the adversary has queried t_{-i} to the simulator (Figure 6 (C)). However, this latter event happens with low probability given that the adversary must guess the value s_{-j} out of thin air: indeed, the adversary’s only way of discovering s_{-j} is to query \mathcal{S} at t_{-i} for some $i \geq j$ (this is formally argued in the proof). Finally, a case not depicted in Figure 6 occurs when the adversary queries a point t_i with $i > 0$ or s_i with $i \geq 0$. This corresponds to scenario (a): the simulator will return the predetermined value of g , after “extending” the ladder such that the ladder’s furthest point is at a distance of at least $(q_1 + 1)$ RO calls from the last point queried by \mathcal{D} . A full specification of our simulator is given in Figure 5.

The above outlines the simulator \mathcal{S} . As for the indistinguishability proof itself, we use a sequence of games, where each game presents the adversary with a two-oracle environment. In the first game, the environment is equivalent to the pair (RO, \mathcal{S}) ; in the last, to the pair $(G[g], g)$ for a random g . For each pair of adjacent games G_i, G_{i+1} , the adversary \mathcal{D} ’s distinguishing advantage $\Pr[\mathcal{D}^{G_i} \Rightarrow 1] - \Pr[\mathcal{D}^{G_{i+1}} \Rightarrow 1]$ is upper bounded, where \mathcal{D}^{G_i} notates \mathcal{D} run in the two-oracle environment of game G_i . Summing these

⁵This does not preclude \mathcal{D} from knowing beforehand the value of some s_i ’s with $i > 0$, since it could make these RO queries before querying the simulator at s_0 , but this costs \mathcal{D} the same number of RO queries as computing these values afterward.

⁶In fact, the values t_0, \dots, t_{q_1+1} are already computed during step (b), given that step (b) is unsuccessful.

⁷Note one could also set $g(s_{-1}) = t_0$, but this changes little.




Figure 6: Illustration of the simulator for Theorem 3.3. Horizontal arrows show queries to RO, while upward and downward diagonal arrows show g -queries defined by the simulator. Dashed arrows show RO queries known only to the attacker. The last point queried by the attacker and answered by the simulator is circled. Bold black arrows indicate the internal path followed by the simulator to answer the attacker's last query. Top (A): The simulator builds a ladder after being queried at t_0 . Middle (B): The simulator “fills in” values of $g(\cdot)$ after being queried at t_{-i} . Bottom (C): The simulator aborts after being queried at s_{-j} .

upper bounds gives the final indistinguishability bound. Our proof uses 23 games in all.

4 HMAC as a General-purpose Keyed Hash Function

HMAC [5] uses a hash function to build a keyed hash function, i.e. one that takes both a key and message as input. Fix some hash function⁸ $H: \{0, 1\}^* \rightarrow \{0, 1\}^n$. HMAC assumes this function H is built by iterating an underlying compression function with a message block size of $d \geq n$ bits. We define the following functions:

$$\begin{aligned} F_K(M) &= H((\rho(K) \oplus \text{ipad}) \parallel M) \\ G_K(M) &= H((\rho(K) \oplus \text{opad}) \parallel M) \end{aligned} \quad \text{where } \rho(K) = \begin{cases} H(K) & \text{if } |K| > d \\ K & \text{otherwise.} \end{cases}$$

The two constants used are $\text{ipad} = 0 \times 36^{d/8}$ and $\text{opad} = 0 \times 5c^{d/8}$. These constants are given in hexadecimal, translating to binary gives $0 \times 36 = 0011\ 0110_2$ and $0 \times 5c = 0101\ 1100_2$. Recall that we have defined the \oplus operator so that, if $|K| < d$, it first silently pads out the shorter string by sufficiently many zeros before computing the bitwise xor. It will also be convenient to define $\text{xpad} = \text{ipad} \oplus \text{opad}$. The function $\text{HMAC}: \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ is defined by

$$\text{HMAC}(K, M) = G_K(F_K(M)) = (G_K \circ F_K)(M).$$

We sometimes write $\text{HMAC}_d[P]$, HMAC_d , or $\text{HMAC}[P]$ instead of HMAC when we want to make the reliance on the block size and/or an underlying ideal primitive explicit.

SOME APPLICATIONS OF HMAC. The HMAC construction was originally suggested for use as a PRF or MAC — settings in which K is a secret value generated by a trusted party. And while HMAC is still used as such, it has also come to be used (or suggested for use) in settings where K is public [25, 29, 32] or a non-uniform secret [36]. There also exist settings whose theoretical security models allow adversarially generated keys [38].

- *HKDF*. The HKDF scheme builds from HMAC a general key derivation function [32]. It is currently in the process of being standardized [31]. The construction follows an extract-then-expand approach. In the extract phase, compute $Y = \text{HMAC}(sa, msk)$ using a public, application-dependent salt sa and a secret source of entropy msk . It is strongly recommended [31] that sa be a uniformly selected n -bit string. In the expand phase, a key $K = \text{HMAC}(Y, da)$ is derived using application-specific context information da . IPsec key exchange similarly utilizes HMAC with public information as key [25, 29].
- *TLS*. The TLS protocol uses HMAC as a key-derivation function. One derives a key as $\text{HMAC}(pms, da)$ where da consists of public nonces and public session-dependent information and pms is either a 48 byte string consisting of a version number and 46 random bytes (when using RSA key transport) or a Diffie-Hellman value (when using Diffie-Hellman key exchange) of length specified by the group size. The value pms is therefore secret but not necessarily a secret, uniform bit string from an adversary’s perspective.
- *PKCS#5*. The PKCS#5 standard uses HMAC for password-based key derivation. A derived key for some password pw and salt sa is defined by $K = \bigoplus_{i=1}^{\ell} \text{HMAC}^{\ell}(pw, sa)$. In practice, passwords are chosen by users in all kinds of ways, and typically have relatively little entropy.

⁸RFC 2104 defines HMAC over strings of bytes, but we chose to use bits to provide more general positive results — all our negative results lift to a setting in which only byte strings are used. Note also that for simplicity we assumed H with domain $\{0, 1\}^*$. In practice hash functions often do have some maximal length (e.g., 2^{64}), and in this case HMAC must be restricted to smaller lengths.

- *Hedged cryptography.* Ristenpart and Yilek [38] suggest a general way of using HMAC to modify cryptographic routines so as to hedge against randomness failures. For example, let $\mathcal{E}(pk, M; R)$ denote encrypting message M under public-key pk and using randomness R . They suggest to modify encryption to instead proceed via $\mathcal{E}(pk, M; \text{HMAC}(R, pk \parallel M))$. (Here we show only the case that $|R| = n$.) This merges two prior suggestions due to Bellare et al. [4] and Yilek [42]. One of the security goals targeted is chosen-distribution attack (CDA) security [4], which requires no partial information about high min-entropy messages is leaked even in the presence of adversarially-specified randomness R .

Common to the examples above is that security proofs of HMAC when used with uniform secret keys [3, 5] do not (directly) apply. In some cases, standard model proofs have nevertheless been given. Particularly, positive results about HMAC’s security as a randomness extractor [15, 17, 23, 32], which are applicable to its use in HKDF, IPsec, and TLS. But in other cases it seems unlikely that standard model analyses are possible. The security of TLS key transport is one example, as discussed in [35]. Another example is when using lower entropy secrets with HKDF, as discussed in [32]. The same situation faces PKCS#5 because passwords are often short. In these two settings, the results of [15, 17, 23, 32] are inapplicable because they require sources with high (computational) min-entropy. Finally, security proofs for the hedged encryption construction seems to fundamentally rely on the “programmability” of ROs [4, 38].

Prior work has therefore turned to assuming HMAC is a keyed RO to achieve positive results. A variant of TLS key exchange is analyzed under this assumption by Morisse et al. [35]. Krawczyk analyzes HKDF as a randomness extractor under this assumption [32]. Ristenpart and Yilek analyze the security of the hedged encryption construction under this assumption [38]. Validating this assumption as used in these works would require proving indistinguishability of HMAC from a keyed RO for distinguisher’s that query arbitrary keys.

USING INDIFFERENTIABILITY. Our focus is these last analyses in the keyed ROM. We want to understand if the positive results carry over to a setting in which the structure of HMAC is taken into account. Namely, when one uses $\text{HMAC}[P]$ for some underlying ideal primitive P (e.g., an ideal compression function). This is important because security proofs taking into account HMAC’s structure would rule out subtle, exploitable interactions between HMAC’s design and its use in these applications.

Indistinguishability provides a mechanism for this kind of analysis: should $\text{HMAC}[P]$ be indistinguishable from a keyed RO \mathcal{R} with good bounds, then the MRH composition theorem [33] could be applied to theorems from [32, 35, 38] to give corollaries that security holds when using $\text{HMAC}[P]$.⁹ To simultaneously cover all the settings of [32, 35, 38], indistinguishability would have to hold even for adversaries that can query arbitrary keys and messages.

In the following sections, we will therefore analyze the security of HMAC in the sense of being indistinguishable from a keyed RO. As we will see, the story is more involved than one might expect.

4.1 Weak Key Pairs in HMAC

Towards understanding the indistinguishability of HMAC, we start by observing that the way HMAC handles keys gives rise to two worrisome classes of weak key pairs.

- **Colliding key pairs:** We say that keys $K \neq K'$ *collide* if

$$\rho(K) \parallel 0^{d-|\rho(K)|} = \rho(K') \parallel 0^{d-|\rho(K')|} .$$

⁹Technically, this would only apply to a slightly weaker security model for [38] that does not allow hash-dependent randomness, message distributions. See [37] for a discussion.

For any message M and colliding keys K, K' it holds that $\text{HMAC}(K, M) = \text{HMAC}(K', M)$.

Colliding keys exist because of HMAC's ambiguous encoding of different-length keys. Examples of colliding keys include any K, K' for which $|K| < d$ and $K' = K \parallel 0^s$ where $1 \leq s \leq d - |K|$. Or any K, K' such that $|K| > d$ and $K' = H(K)$. As long as H is collision-resistant, two keys of the same length can never collide.

Colliding keys enable a simple attack against indistinguishability: query Func on (K, M) and (K', M) for K, K' colliding and see if the outputs are equal. Colliding key pairs may also have implications for other settings. We discuss this all more in the next section.

The second form of weak key pair we term ambiguous:

- **Ambiguous key pairs:** A pair of keys $K \neq K'$ is *ambiguous* if

$$\rho(K) \oplus \text{ipad} = \rho(K') \oplus \text{opad} .$$

For any X , both $F_K(X) = G_{K'}(X)$ and $G_K(X) = F_{K'}(X)$ when K, K' are ambiguous.

An example such pair is K, K' of length d bits for which $K \oplus K' = \text{xpad}$. For any key K , there exists one key K' that is easily computable and for which K, K' are ambiguous: set $K' = \rho(K) \oplus \text{xpad}$. Finding a third key K'' that is also ambiguous with K is intractable should H be collision resistant. The easily-computable K' will not necessarily have the same length as K . In fact, there exist ambiguous key pairs of the same length k only when $k \in \{d - 1, d\}$. For a fixed length shorter than $d - 1$, no ambiguous key pairs exist due to the fact that the second least significant bit of xpad is 1. For a fixed length longer than d bits, if $n < d - 1$ then no ambiguous key pairs exist and if $n \geq d - 1$ then producing ambiguous key pairs would require finding K, K' such that $H(K) \oplus H(K')$ equals the first n bits of xpad . This is intractable for any reasonable hash function H .

Unlike colliding key pairs, ambiguous key pairs, at first glance, may not seem to be problematic for security. But in fact they give rise to a chain-shift-like property for HMAC that, as with H^2 , can lead to insecurities in some settings. We explore this in Section 4.3.

SUMMARY. We uncover two types of weak key pairs in HMAC. We will discuss further in Section 4.1 how colliding keys trivially rule out indistinguishability. Colliding key pairs are avoided by using fixed-length keys. However, even here, we have ambiguous key pairs. We will show in Section 4.3 how HMAC when ambiguous key pairs are allowed cannot be proven indistinguishability with good concrete security. The underlying structural issue is similar to that of second iterate constructions. We leave as an open question showing a weak upper bound for the indistinguishability of HMAC with ambiguous key pairs, but suspect that the techniques from the proof of Theorem 3.3 might be applicable.

Finally, and fortuitously, most applications of HMAC appear to avoid both kinds of weak key pairs. As we will show in Section 4.4, we can prove indistinguishability holds with standard, good bounds for some cases in which weak key pairs are avoided. For example, the common case of using keys of a fixed length less than $d - 1$ provides security. A summary of all these results is given in Figure 7.

4.2 Colliding Key Pairs and the Indistinguishability of HMAC

Colliding key pairs give rise to a simple attack against the indistinguishability of HMAC. We have the following theorem:

Theorem 4.1 Let $\text{HMAC}[H[P]]$ be the HMAC construction for an arbitrary underlying hash function $H[P]$ and let \mathcal{R} be a keyed RO. Then there exists an adversary \mathcal{A} making two queries and running in a small constant amount of time such that for any simulator \mathcal{S} it holds that

$$\text{Adv}_{\text{HMAC}[H[P]], \mathcal{R}, \mathcal{S}}^{\text{indiff}}(\mathcal{A}) \geq 1 - \frac{1}{2^n} . \quad \square$$

Key space includes	Indifferentiable	Queries	Section
Colliding key pairs	No	2	§4.1
Ambiguous key pairs	At most weakly	$\mathcal{O}(2^{n/4})$	§4.3
Only keys K of fixed length $ K < d - 1$	Yes	$\mathcal{O}(2^{n/2})$	§4.4

Figure 7: Summary of indistinguishability of HMAC from a keyed RO for various restrictions on the key space. The “Queries” column indicates the number of queries used by an attacker to gain good advantage against any simulator.

Proof: Distinguisher \mathcal{A} first picks two keys $K \neq K'$ that collide and picks an arbitrary message M . It then queries its Func oracle on (K, M) and (K', M) to retrieve two values Y, Y' . If $Y = Y'$ then it returns 1 (guessing that it is in game $\text{Real}_{\text{HMAC}[H[P]], \mathcal{R}}$) and returns 0 otherwise (guessing that it is in game $\text{Ideal}_{\mathcal{R}, \mathcal{S}}$). The advantage of \mathcal{A} is equal to $1 - 2^{-n}$ regardless of the simulator \mathcal{S} , which is never invoked. ■

Colliding key pairs endanger security of any application of HMAC that uses variable-length keys. The formal security of several applications is ruled out when colliding key pairs are allowed. For PKCS#5, consider the slight simplification of using $\text{HMAC}(pw, sa)$ to derive a key from a password pw and salt sa . Then if the set of passwords from which pw was drawn includes colliding key pairs (which is allowed by, e.g., the password-based key derivation function security definition given by [9]), then the search space of a dictionary attack will be reduced. For hedged cryptography, if one extends the security notion of [4] to allow the adversary to pick variable lengths of randomness, then there exists a simple adversary that violates the chosen-distribution attack security of the HMAC-based constructions from [38].

That said, we are unaware of any exploitable security vulnerabilities in practice due to colliding key pairs. The PKCS#5 example above would require having passwords in the set that are encoded to binary strings that end in zero bytes. HMAC as actually used in HKDF, TLS, and hedged cryptography all use fixed-length keys.

RELATED-KEY ATTACKS. We digress for a moment to consider the setting of related-key attacks (RKAs) against HMAC as a PRF (i.e., using uniformly selected, secret keys). Recall that RKA-PRF security [6] asks that no attacker can distinguish between two oracles to which it can make adaptive queries. The first oracle allows the attacker to query a related key function ϕ from some allowed set Φ and a message M . It returns $\text{HMAC}(\phi(K), M)$ for a randomly chosen key K . The second oracle has the same interface but returns $\rho(\phi(K), M)$ where ρ is a family of random functions.

Colliding keys give rise to an RKA against HMAC for any Φ that includes both the identity function and a function ϕ for which the keys $K, \phi(K)$ collide. The adversary queries the identity function and a message M and in a second query ϕ and the same message M . If the returned values are the same it guesses that it is interacting with HMAC and otherwise it guesses it is interacting with ρ . The adversary achieves advantage $1 - 2^{-n}$.

4.3 Ambiguous Key Pairs and the Indistinguishability of HMAC

We now turn to ambiguous key pairs and show that these also lead to lower bounds on the indistinguishability of HMAC. Recall that with H^2 , problems arose because outputs of H^2 on some message were valid intermediate values used in computing H^2 on some other message. HMAC is the same due to ambiguous key pairs. Let M be some message and K, K' be an ambiguous key pair. Then, we have

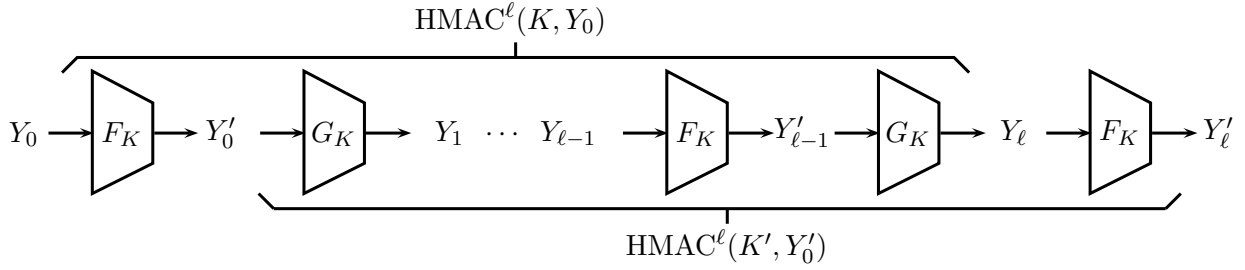


Figure 8: Diagram of two hash chains $(K, Y) = (Y_0, \dots, Y_\ell)$ and $(K', Y') = (Y'_0, \dots, Y'_\ell)$ for HMAC where $\rho(K') = \rho(K) \oplus \text{xpad}$.

that $\rho(K') = \rho(K) \oplus \text{xpad}$ and so $F_K(M) = G_{K'}(M)$. Thus,

$$\text{HMAC}(K', F_K(M)) = G_{K'}(\text{HMAC}(K, M))$$

This property does not appear immediately exploitable in attacks against, for example, the HMAC applications mentioned above. We will thus follow the same path as we did with H^2 to highlight how this structural property affects the ability to show that HMAC is indistinguishable from a RO — even when colliding key pairs do not arise. The result will be our ruling out strong indistinguishability bounds when ambiguous key pairs arise, and thus limiting the scope of applicability of composition-based proofs of security for applications of HMAC. As with H^2 , we will also detail HMAC applications in which ambiguous key pairs can be exploited by attackers. As before, these revolve around hash chains.

KEYED HASH CHAINS. We lift our notions of hash chains from Section 3 to the setting of keyed hash functions. Let $H: \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a keyed hash function. A hash chain $Y = (K, Y_0, \dots, Y_\ell)$ is a key K , a message Y_0 , and a sequence of ℓ values $Y_i = H(K, Y_{i-1})$ for $1 \leq i \leq \ell$. So a keyed hash chain $Y = (K, Y_0, \dots, Y_\ell)$ for HMAC has $Y_i = \text{HMAC}(K, Y_{i-1})$ for $1 \leq i \leq \ell$. We refer to Y_0 as the *start* of the keyed hash chain and to Y_ℓ as the *end*. Two keyed hash chains Y, Y' are *non-overlapping* if $Y_i \neq Y'_j$ for all $0 \leq i \leq j \leq \ell$.

Let $\text{HMAC}[P](K, M)$ be HMAC using as underlying hash function a random oracle $P: \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ and extend the definitions of ρ, F_K, G_K to use P in the natural way. Given the start and end of a chain $Y = (K, Y_0, \dots, Y_\ell)$, it is easy for an adversary to compute the start and end of a new chain $Y' = (K', Y'_0, \dots, Y'_\ell)$. To do so, the adversary chooses K' so that $\rho(K') = \rho(K) \oplus \text{xpad}$ and then computes $Y'_0 \leftarrow F_K(M)$ and $Y'_\ell \leftarrow F_K(Y_0)$. By the choice of K' it holds that both $F_K(X) = G_{K'}(X)$ and $G_K(X) = F_{K'}(X)$ for all X . Thus Y_0, Y'_ℓ are valid start and end points for a chain because

$$\begin{aligned} (G_{K'} \circ F_{K'})^\ell(Y'_0) &= G_{K'}(F_{K'}(\dots F_{K'}(Y'_0) \dots)) \\ &= G_{K'}(F_{K'}(\dots F_{K'}(F_K(M)) \dots)) \\ &= F_K(G_K(\dots G_K(F_K(M)) \dots)) \\ &= F_K((G_K \circ F_K)^\ell(K, M)) \\ &= F_K(Y_\ell) \\ &= Y'_\ell \end{aligned}$$

We refer to the equivalence above as the chain-shift property of HMAC. A diagram of the two HMAC chains involved appears in Figure 8. Finally, we note that with overwhelming probability (over the coins of P) Y and Y' will not overlap.

To capture the gap between HMAC and a RO in a formal way, we extend the $\text{CHAIN}_{H[P], n, \ell}$ game

<p>main CHAIN_{H[P],k,n,ℓ}:</p> <p>$i \leftarrow 0$; $K \leftarrow_{\\$} \{0, 1\}^k$ $Y_0 \leftarrow_{\\$} \{0, 1\}^n$; For $i = 1$ to ℓ do $Y_i \leftarrow H[P](K, Y_{i-1})$ $(K', Y_0^*, Y_\ell^*) \leftarrow_{\\$} \mathcal{A}^{\text{Prim}}(K, Y_0, Y_\ell)$ $Y'_0 \leftarrow Y_0^*$; For $i = 1$ to ℓ do $Y'_i \leftarrow H[P](K', Y'_{i-1})$ $\mathcal{Y} \leftarrow \{Y_0, \dots, Y_\ell\}$ $\mathcal{Y}' \leftarrow \{Y'_0, \dots, Y'_\ell\}$ If $(Y'_\ell = Y_\ell^*) \wedge ((\mathcal{Y} \cap \mathcal{Y}') = \emptyset)$ then Ret true Ret false</p>	<p>procedure Prim(M):</p> <p>$i \leftarrow i + 1$ If $i \geq \ell \cdot \text{Cost}(H, n)$ then Ret \perp Ret $P(M)$</p>
--	--

Figure 9: A keyed version of the chain-making game.

from Section 3 to work for keyed chains; see Figure 9. The only change is that a randomly chosen challenge key of k bits is generated by the game and the adversary now outputs not only an attempted start and end point of a chain, but a chosen key as well. We define advantage as

$$\text{Adv}_{H[P],k,n,\ell}^{\text{chain}}(\mathcal{A}) = \Pr \left[\text{CHAIN}_{H[P],k,n,\ell}^{\mathcal{A}} \Rightarrow \text{true} \right].$$

We compare the CHAIN security of HMAC[P] to the CHAIN security of a keyed RO \mathcal{R} . Technically, we look at the the achieved security of CHAIN_{HMAC[P],k,n,ℓ} to that of CHAIN_{H[ℛ],k,n,ℓ} where $H[\mathcal{R}](K, M) = \mathcal{R}(K, M)$ implements a keyed RO. Security in the latter case follows in the same manner that CHAIN security was established for a (non-keyed) RO, as per Claim 3.1. Namely,

$$\text{Adv}_{H[\mathcal{R}],k,n,\ell}^{\text{chain}}(\mathcal{A}) \leq \frac{(3\ell)^2}{2^n}$$

for all adversaries \mathcal{A} . On the other hand, an adversary \mathcal{B} exploiting the chain-shift property of HMAC achieves

$$\text{Adv}_{\text{HMAC}[P],k,n,\ell}^{\text{chain}}(\mathcal{B}) \geq 1 - \frac{\ell^2}{2^n}.$$

The adversary \mathcal{B} lets $K' = \rho(K) \oplus \text{xpad}$ and picks $Y'_0 = F_K(Y_0)$ and $Y'_\ell = F_K(Y_\ell)$. As with H^2 , we see a gap between the CHAIN security achieved by the ideal object (a keyed RO) and HMAC.

Note that \mathcal{B} may output a K' that is of different length than K . But the attack extends to a setting requiring $|K| = |K'|$ for some, but not all, choices of the parameter k . In particular, the adversary \mathcal{B} can always find a suitable K' with $|K'| = k$ for $k \in \{d-1, d\}$.

MUTUAL PROOFS OF WORK WITH KEYS. We similarly lift mutual proofs of work to the setting of keyed chains. Referring to Figure 10, we modify the protocol to allow both parties to choose a key and a message as challenge for the other party. Likewise, the security game POW_{H[P],n,ℓ₁} is modified as shown in the right hand side of Figure 10 to give game POW_{H[P],k,n,ℓ₁}. An adversary can use the chain-shift-like property of HMAC to mount a successful attack against POW_{HMAC[P],k,n,ℓ₁} for any $\ell_1 > 1$ and for P a RO. Consider the following adversary \mathcal{A} . When it first receives a nonce K_2, X_2 , it chooses K_1 such that $\rho(K_1) = \rho(K_2) \oplus \text{xpad}$ and lets $X_1 \leftarrow F_{K_2}(X_2)$. Later when it receives $Y_1 = \text{HMAC}^{\ell_1}[P](K_1, X_1)$ it computes its response as $Y_2 \leftarrow F_{K_1}(Y_1)$, sets $\ell_2 = \ell_1 + 1$, and returns ℓ_2, Y_2 .




Figure 10: Example protocol (**left**) and adversarial \mathcal{P}_2 security game (**right**) for mutual proofs of work with keyed hash functions.

By the chain-shift property of HMAC, we get that

$$\begin{aligned}
Y_2 &= F_{K_1}((G_{K_1} \circ F_{K_1})^{\ell_1}(X_1)) \\
&= F_{K_1}((G_{K_1} \circ F_{K_1})^{\ell_1}(F_{K_2}(X_2))) \\
&= F_{K_1}(G_{K_1}(\cdots F_{K_1}(F_{K_2}(X_2)) \cdots)) \\
&= G_{K_2}(F_{K_2}(\cdots G_{K_2}(F_{K_2}(X_2)) \cdots)) \\
&= \text{HMAC}^{\ell_1+1}[P](K_2, X_2) .
\end{aligned}$$

The two chains will be non-overlapping with probability close to one and \mathcal{A} makes at most two P -applications, in turn ensuring that $q < 2\ell_2$ when $\ell_1 > 1$.

DIRECT INDIFFERENTIABILITY LOWER BOUNDS. All the above suggests that HMAC is not indifferntiable from a keyed RO because of ambiguous key pairs and regardless of the strength of the primitive underlying HMAC. As before, this requires careful interpretation, in the same sense as discussed for H^2 in Section 3. We therefore provide a direct indifferntiability adversary. It relies on the quantity $p'(H, w, \ell)$ that represents the probability of certain collisions associated to the queries the distinguisher makes. The quantity depends on the particulars of the construction $H[P]$ and is defined as

$$p'(H, w, \ell) = \Pr[\{H(\text{ipad} \parallel U_1), \dots, H(\text{ipad} \parallel U_{w-1})\} \cap \{y_0, y_1, \dots, y_\ell\} \neq \emptyset] ,$$

for U_1, \dots, U_{w-1} being $w - 1$ independent n -bit strings, and y_0, y_1, \dots, y_ℓ are out of a chain of n -bit values $x_0, y_0, \dots, x_\ell, y_\ell$ where x_0 is chosen at random, and $y_i = H(\text{ipad} \parallel x_i)$ for $i = 0, \dots, \ell$ and $x_i = H(\text{opad} \parallel y_{i-1})$ for $i = 1, \dots, \ell$. We show in Appendix A that, for typical H constructions, $p'(H, w, \ell) \leq 2(w - 1)(\ell + 2)/2^n + 2(\ell + 1)^2/2^n$.

Theorem 4.2 [Attack against HMAC] Let $H[P]$ be an arbitrary hash construction with n -bit outputs, calling a random primitive P , and let $\mathcal{R} : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a keyed random oracle. For all integer parameters $w, \ell \geq 1$, there exists an adversary $\tilde{\mathcal{D}}_{w, \ell}$ making 2ℓ Func-queries and

$(w + 1) \cdot \text{Cost}(H, n + d)$ Prim-queries such that for all simulators \mathcal{S} ,

$$\text{Adv}_{\text{HMAC}[H[P]], \mathcal{R}, \mathcal{S}}^{\text{indiff}}(\tilde{\mathcal{D}}_{w, \ell}) \geq 1 - p'(H, w, \ell) - \frac{5\ell^2}{2^{n+1}} - \frac{q_S \ell}{2^n} - \frac{q_S}{2^n} - \frac{q_S}{w \cdot \ell} - \frac{1}{w},$$

where q_S is the overall number of \mathcal{R} queries by \mathcal{S} when replying to $\tilde{\mathcal{D}}_{w, \ell}$'s Prim queries. ■

The proof is given in Appendix C. The interpretation of the theorem is analogous to that of Theorem 3.2.

4.4 Indifferentiability Upper Bound for HMAC with Restricted Keys

We have seen that HMAC's construction gives rise to two kinds of weak key pairs that can be abused to show that HMAC is not indifferentiable from a keyed RO (with good bounds). But weak key pairs are serendipitously avoided in most applications. For example, the recommended usage of HKDF [32] specifies keys of a fixed length less than $d - 1$. Neither kind of weak key pairs exist within this subset of the key space. For the other applications mentioned at the beginning of Section 4, our positive results should be applicable, with some caveats: PKCS#5 for passwords sufficiently restricted, hedged cryptography with certain lengths of randomness, and TLS for particular premaster secret sizes.

Our positive results will focus primarily on the case mentioned above. That is, we restrict attention to keys K for which $|K| = k$ and k is a fixed integer less than $d - 1$. In fact, we prove a positive indifferentiability bound for a slightly more general key space, described next.

RESTRICTED KEY SPACES FOR HMAC. We first provide some definitions regarding restricted key spaces for HMAC. We here focus on the case where all keys are of length d or less. Let $\mathcal{K} \subseteq \{0, 1\}^{\leq d}$ be a set of keys. We say \mathcal{K} is *allowed* if there exists a function $\text{GetKey}: \{0, 1\}^d \rightarrow \{0, 1\}^*$ such that

$$\text{GetKey}(K \oplus \text{ipad}) = K \quad \text{and} \quad \text{GetKey}(K \oplus \text{opad}) = K.$$

The function GetKey implies the existence of a predicate $\text{IsOuter}: \{0, 1\}^d \rightarrow \{0, 1\}$ such that for any $K \in \mathcal{K}$,

$$\text{IsOuter}(K \oplus \text{ipad}) = 0 \quad \text{and} \quad \text{IsOuter}(K \oplus \text{opad}) = 1.$$

The predicate IsOuter can determine which pad was used with a key, while the function GetKey can invert xor'ing by ipad or opad . For any allowed \mathcal{K} it must be that $K \oplus \text{ipad} \neq K' \oplus \text{opad}$ for all $K, K' \in \mathcal{K}$. We let \mathcal{K} -restricted HMAC be the function $\text{HMAC}: \mathcal{K} \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ that is undefined for $K \notin \mathcal{K}$ and for $K \in \mathcal{K}$ is defined equivalently to $\text{HMAC}: \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^n$. In our proofs below, we assume that indifferentiability adversaries do not query keys $K \notin \mathcal{K}$ to the Func oracle.

One example of an allowed \mathcal{K} is all keys K of length equal to d with second least significant bit equal to 0. For this class, $\text{IsOuter}(X) = 1$ iff the second least significant bit of X is 0. (Recall that the second least significant bit of ipad is 1 and of opad is 0.) Another example is exactly the set of widest consequence: the set of all keys of a fixed length that is less than $d - 1$.

HMAC USING A RO. We start with the simpler case, proving that \mathcal{K} -restricted HMAC is indifferentiable from a keyed RO when the underlying hash function is modeled as a RO.

Theorem 4.3 Fix $d, n > 0$. Let $P: \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a RO, and consider \mathcal{K} -restricted $\text{HMAC}_d[P]$ for an allowed key set \mathcal{K} . Let $\mathcal{R}: \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a keyed RO. Then there exists a simulator \mathcal{S} such that for any distinguisher \mathcal{A} whose total query cost is σ it holds that

$$\text{Adv}_{\text{HMAC}_d[P], \mathcal{R}, \mathcal{S}}^{\text{indiff}}(\mathcal{A}) \leq \mathcal{O}\left(\frac{\sigma^2}{2^n}\right)$$

\mathcal{S} makes at most q_2 queries and runs in time $\mathcal{O}(q_2 \log q_2)$ where q_2 is the number of Prim queries made by \mathcal{A} . □

The use of $\mathcal{O}(\cdot)$ just hides small constants. Combining Theorem 4.3 with the indifferntiability composition theorem allows us to conclude security for $\text{HMAC}_d[H]$ for underlying hash function H that is, itself, indifferntiable from a RO. For example, should H be one of the proven-indifferntiable SHA-3 candidates.

Proof of Theorem 4.3: (Sketch) Fix an allowed key set \mathcal{K} . The simulator \mathcal{S} imitates a RO $P: \{0, 1\}^* \rightarrow \{0, 1\}^n$ in a way that is consistent with the keyed RO \mathcal{R} , to which the simulator has oracle access. The simulator works as shown below.

```

algorithm  $\mathcal{S}^{\mathcal{R}}(U)$ :
Parse  $U$  as  $X \parallel Y$  with  $|X| = d$ 
If  $\text{lsOuter}(X) = 0$  then
   $V \leftarrow_{\$} \{0, 1\}^n$ 
   $\mathbf{F}[V] \leftarrow (\text{GetKey}(X), Y)$ 
  Ret  $V$ 
If  $\text{lsOuter}(X) = 1$  and  $\mathbf{F}[Y] \neq \perp$  then
   $Z \leftarrow \mathcal{R}(\mathbf{F}[Y])$ 
  Ret  $Z$ 
Ret  $R \leftarrow_{\$} \{0, 1\}^n$ 

```

In words, the simulator identifies whether a query is associated with an “inner” application or an “outer” application. In the first case it chooses a random response, and records in a table the input associated to the response. If an “outer” application, it looks up in the table whether there is an “inner” input associated with the query. If so, it responds with the keyed RO’s output for the key and message associated with that “inner” query. Otherwise, the simulator outputs a random point.

That keys queried to the construction are in an allowed set means that there is no ambiguity in the lsOuter predicate. Because of this, intuitively, the simulator can only fail in two ways. First, a collision amongst the choices of V across two different queries occur. Second, a query to the simulator with $\text{lsOuter}(X) = 1$ ends up producing a value V such that V was previously queried to the simulator that was an “outer” query. In either cases, the distinguisher can abuse the events to successfully distinguish. Informally, in either the $\text{Real}_{\text{HMAC}_d[P]}$ or game $\text{Ideal}_{\mathcal{R}, \mathcal{S}}$ the probability of a collision occurring is at most $(q_1 + q_2)^2/2^n$ and the probability of the second kind of failure is at most $q_2^2/2^n$. █

HMAC USING MD-BASED HASH FUNCTIONS. The above result does not extend to cover HMAC built from hash functions which are not indifferntiable from a RO. This includes, for example, the SHA family of hash functions and others that use the Merkle-Damgård transform. We therefore treat this special case, investigating $\text{HMAC}_d[\text{SMD}[f]]$ where $f: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^n$ is a random oracle.

Theorem 4.4 Fix $d, n > 0$ with $d \geq n$. Let $f: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^n$ be a RO and consider \mathcal{K} -restricted $\text{HMAC}_d[\text{SMD}[f]]$ for an allowed key set \mathcal{K} . Let $\mathcal{R}: \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a keyed RO. Then there exists a simulator \mathcal{S} such that for any distinguisher \mathcal{A} whose total query cost is $\sigma \leq 2^{n-2}$ it holds that

$$\text{Adv}_{\text{HMAC}_d[\text{SMD}[f]], \mathcal{R}, \mathcal{S}}^{\text{indiff}}(\mathcal{A}) \leq \mathcal{O}\left(\frac{\sigma^2}{2^n}\right)$$

\mathcal{S} makes at most q_2 queries and runs in time $\mathcal{O}(q_2 \log q_2)$ where q_2 is the number of Prim queries by \mathcal{A} . □

We note that the restriction to $\sigma \leq 2^{n-2}$ in the theorem statement is just a technicality to make the bound simpler and likewise the use of $\mathcal{O}(\cdot)$ in the advantage statement hides just a small constant.

Unlike our positive results about H^2 , the bounds provided by Theorems 4.3 and 4.4 match, up to small constants, results for other now-standard indifferentiable constructions (c.f., [14]). First, the advantage bounds both hold up to the birthday bound, namely $\sigma \approx 2^{n/2}$. Second, the simulators are efficient and, specifically, make at most one query per invocation. All this enables use of the indifferentiability composition theorem in a way that yields strong, standard concrete security bounds.

For the proof of Theorem 4.4, we simplify the treatment in two ways. First, let π be the minimal padding length of SMD. For SHA-256, for example, $\pi = 65$. When $d \geq n + \pi$, the outer application of SMD[f] of HMAC $_d$ [SMD[f]] will always consist of exactly two f applications. Otherwise, the outer application of SMD[f] could consist of some larger fixed number of f applications. We will for simplicity focus on the case of two calls; it is easy to extend to the more general case.

Second, we will replace the padding of SMD with adversarially-controlled message bits. For any $K \in \mathcal{K}$, any $X \in (\{0, 1\}^d)^+$, and any $P \in \{0, 1\}^{d-n}$ define the functions

$$\begin{aligned} \mathbf{F}[f](K, X, P) &= (f(\text{IV}, K \oplus \text{opad}), \text{MD}[f]((K \oplus \text{ipad}) \parallel X) \parallel P) \\ \mathbf{fF}[f](K, X, P) &= f(\mathbf{F}[f](K, X, P)) \end{aligned}$$

Here the last bits of X replaces the padding in the first SMD application while P replaces the padding used in the second SMD application. Theorem 4.4 establishes the indifferentiability of $\mathbf{fF}[f]$. In it, the total query cost of an attacker is equal to the sum of the costs (as defined in Section 2) of each query to \mathbf{Func} plus the number of \mathbf{Prim} queries.

Proof of Theorem 4.4: Let $g: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^n$ be a RO and let $\mathbf{gF}[f, g](K, X, P) = g(\mathbf{F}[f](K, X, P))$. Our proof proceeds in two steps. First, Lemma 4.5 below shows that we can restrict attention to showing indifferentiability for the construction $\mathbf{gF}[f, g]$. This step fundamentally relies upon the fact that keys queried by \mathcal{A} are within an allowed set, which in turn enables us to show domain separation between “internal” uses of f within $\mathbf{F}[f]$ and the external application of f in $\mathbf{fF}[f]$. Second, Lemma 4.6 shows that $\mathbf{F}[f](K, X, P)$ is preimage-aware. This step relies upon the fact that we can apply \mathbf{GetKey} to extract K from $K \oplus \text{ipad}$ and from $K \oplus \text{opad}$. To conclude we combine the two lemmas with [19, Th. 4.1], which asserts that the composition of a PrA function and a RO is indifferentiable from a RO.

Lemma 4.5 Let $f, g: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^n$ be random oracles. Let $\mathcal{S}_{\mathcal{B}}$ be a simulator and \mathcal{A} be a distinguisher making at most q_1 \mathbf{Func} queries, q_2 \mathbf{Prim} queries, and whose total query cost is σ . Then there exists a simulator $\mathcal{S}_{\mathcal{A}}$ and adversary \mathcal{B} such that

$$\mathbf{Adv}_{\mathbf{fF}[f], \mathcal{S}_{\mathcal{A}}}^{\text{indiff}}(\mathcal{A}) \leq \mathbf{Adv}_{\mathbf{gF}[f, g], \mathcal{S}_{\mathcal{B}}}^{\text{indiff}}(\mathcal{B}) + \frac{10\sigma^2}{2^n}$$

$\mathcal{S}_{\mathcal{A}}$ makes the same number of queries as $\mathcal{S}_{\mathcal{B}}$ and works in the same time as $\mathcal{S}_{\mathcal{B}}$ plus $\mathcal{O}(q_2)$. Adversary \mathcal{B} runs in time that of \mathcal{A} plus $\mathcal{O}(q_1 + q_2)$, makes the same number of queries as \mathcal{A} , and has the same total query cost. \square

Proof: Let $\mathcal{S}_{\mathcal{B}}$ be an arbitrary simulator for $\mathbf{gF}^{f, g}$. It can be queried on chaining variable, message block pairs (V, M) on either an f interface or g interface. We use the notation $(0, V, M)$ to signal a query to the f interface and $(1, V, M)$ to signify a query to the g interface. We likewise extend the \mathbf{Prim} interface to accept a label for f queries or g queries.

Let \mathcal{A} be an indifferentiability adversary against \mathbf{fF} .

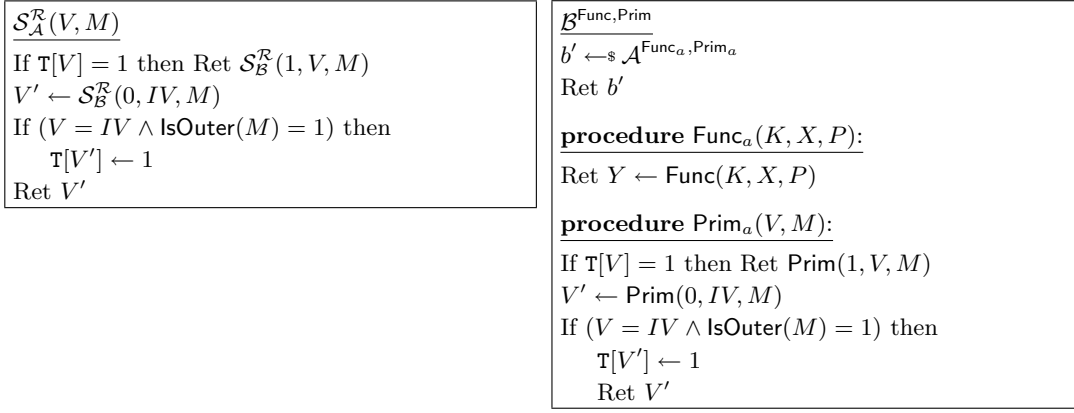


Figure 11: Simulator $\mathcal{S}_{\mathcal{A}}$ and adversary \mathcal{B} used in proof of Lemma 4.5.

We construct a simulator $\mathcal{S}_{\mathcal{A}}$ and an adversary \mathcal{B} as shown in Figure 11. These keep track of the responses V' to f -queries on IV, M with $\mathsf{lsOuter}(M) = 1$ using a table T . A subsequent query made with chaining variable that equals any such response is treated as a g -query. In this way, $\mathcal{S}_{\mathcal{A}}$ and \mathcal{B} “route” queries to either an f or g interface in an appropriate way.

By construction $\mathsf{Ideal}_{\mathcal{R}, \mathcal{S}_{\mathcal{A}}}^{\mathcal{A}}$ and $\mathsf{Ideal}_{\mathcal{R}, \mathcal{S}_{\mathcal{B}}}^{\mathcal{B}}$ are equivalent.

We now use a sequence of games to argue bound the difference between $\Pr[\mathsf{Real}_{\mathsf{fF}[f]}^{\mathcal{A}} \Rightarrow 1]$ and $\Pr[\mathsf{Real}_{\mathsf{gF}[f, g]}^{\mathcal{B}} \Rightarrow 1]$. See Figure 12. Game $\mathsf{G}_0^{\mathcal{A}}$ implements exactly $\mathsf{Real}_{\mathsf{fF}[f]}^{\mathcal{A}}$ — the extra book-keeping code in Prim_a does not affect its behavior. The procedure R_x is used to implement the RO f ; the variable 0 is used as a label. Game G_1 includes the boxed pseudocode, which changes G_0 in that now random choices of chaining variables V are restricted so as to not collide with any chaining variable so-far seen in the game (including IV). Recall that σ is the maximum number of blocks of message bits queried by \mathcal{A} , and this is thus the total number of R_x invocations. Over the course of the game $|\mathcal{V}| \leq 2\sigma$ and so a standard birthday bound argument combined with the discussion above justifies that

$$\Pr \left[\mathsf{Real}_{\mathsf{fF}[f]}^{\mathcal{A}} \Rightarrow 1 \right] = \Pr \left[\mathsf{G}_0^{\mathcal{A}} \Rightarrow 1 \right] \leq \Pr \left[\mathsf{G}_1^{\mathcal{A}} \Rightarrow 1 \right] + \frac{(2\sigma)^2}{2^n} .$$

The next game G_2 is equivalent to G_1 , except that we label random choices in R by either 0 or 1. These labels are added as entries to the table, but lines 31 and 32 ensure an entry already made for a domain point with a different labeling is used by a subsequent call. Thus,

$$\Pr \left[\mathsf{G}_2 \Rightarrow 1 \right] = \Pr \left[\mathsf{G}_1 \Rightarrow 1 \right] .$$

Game G_3 is the same as G_2 but with the boxed statement removed. By the fundamental lemma of game-playing [10] we have that

$$\Pr \left[\mathsf{G}_2 \Rightarrow 1 \right] \leq \Pr \left[\mathsf{G}_3 \Rightarrow 1 \right] + \Pr \left[\mathsf{G}_3 \text{ sets bad} \right] .$$

We now argue that only a call to R on line 21 has non-zero probability of setting bad . We then bound the probability of line 21 causing bad to be set. For the first step, we find it helpful to cast the state of the game’s random choices as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ whose nodes are n -bit strings and whose edges are d -bit strings. When a new entry $V' = \mathsf{R}[x, V, X]$ is chosen, add both V, V' to \mathcal{V} if they are not already present in the graph, and add the edge (V, V') with label X . We additionally color V' by x . Should V' later be returned by $\mathsf{R}(x, V, X)$ and $\mathsf{R}[1-x, V, X] \neq \perp$, then we recolor V' to x . Because of

<p>main G_0 G_1:</p> <p>$b' \leftarrow_s \mathcal{A}^{\text{Func}_a, \text{Prim}_a}$ Ret b'</p> <p>procedure $\text{Func}_a(K, X, P)$:</p> <p>$V_i \leftarrow \text{MD}[\text{R}(0, \cdot, \cdot)](K \oplus \text{ipad} \parallel X)$ $V_o \leftarrow \text{R}(0, IV, K \oplus \text{opad})$ Ret $Y \leftarrow \text{R}(0, V_o, V_i \parallel P)$</p> <p>procedure $\text{Prim}_a(V, M)$:</p> <p>If $\text{T}[V] = 1$ then Ret $\text{R}(0, V, M)$ $V' \leftarrow \text{R}(0, V, M)$ If $(V = IV \wedge \text{lsOuter}(M) = 1)$ then $\text{T}[V'] \leftarrow 1$ Ret V'</p> <p>subroutine $\text{R}(x, V, M)$</p> <p>$\mathcal{V} \stackrel{\leftarrow}{=} V$; $V' \leftarrow_s \{0, 1\}^n \setminus \mathcal{V}$ If $\text{R}[x, V, M] \neq \perp$ then $V' \leftarrow \text{R}[x, V, M]$ $\mathcal{V} \stackrel{\leftarrow}{=} V'$ Ret $\text{R}[x, V, M] \leftarrow V'$</p>	<p>main G_2 G_3:</p> <p>00 $b' \leftarrow_s \mathcal{A}^{\text{Func}_a, \text{Prim}_a}$ 01 Ret b'</p> <p>procedure $\text{Func}_a(K, X, P)$:</p> <p>10 $V_i \leftarrow \text{MD}[\text{R}(0, \cdot, \cdot)](K \oplus \text{ipad} \parallel X)$ 11 $V_o \leftarrow \text{R}(0, IV, K \oplus \text{opad})$ 12 Ret $Y \leftarrow \text{R}(1, V_o, V_i \parallel P)$</p> <p>procedure $\text{Prim}_a(V, M)$:</p> <p>20 If $\text{T}[V] = 1$ then Ret $\text{R}(1, V, M)$ 21 $V' \leftarrow \text{R}(0, V, M)$ 22 If $(V = IV \wedge \text{lsOuter}(M) = 1)$ then 23 $\text{T}[V'] \leftarrow 1$ 24 Ret V'</p> <p>subroutine $\text{R}(x, V, M)$</p> <p>30 $\mathcal{V} \stackrel{\leftarrow}{=} V$; $V' \leftarrow_s \{0, 1\}^n \setminus \mathcal{V}$ 31 If $\text{R}[1-x, V, M] \neq \perp$ then 32 bad $\leftarrow \text{true}$; $x \leftarrow 1-x$ 33 If $\text{R}[x, V, M] \neq \perp$ then 34 $V' \leftarrow \text{R}[x, V, M]$ 35 $\mathcal{V} \stackrel{\leftarrow}{=} V'$ 36 Ret $\text{R}[x, V, M] \leftarrow V'$</p>	<p>main G_4 G_5:</p> <p>$b' \leftarrow_s \mathcal{A}^{\text{Func}_a, \text{Prim}_a}$ Ret b'</p> <p>procedure $\text{Func}_a(K, X, P)$:</p> <p>$V_i \leftarrow \text{MD}[\text{R}(0, \cdot, \cdot)](K \oplus \text{ipad} \parallel X)$ $V_o \leftarrow \text{R}(0, IV, K \oplus \text{opad})$ $\mathcal{X} \leftarrow V_o$ Ret $Y \leftarrow \text{R}(1, V_o, V_i \parallel P)$</p> <p>procedure $\text{Prim}_a(V, M)$:</p> <p>If $\text{T}[V] = 1$ then Ret $\text{R}(1, V, M)$ If $V \in \mathcal{X}$ then bad $\leftarrow \text{true}$ $V' \leftarrow \text{R}(0, V, M)$ If $(V = IV \wedge \text{lsOuter}(M) = 1)$ then $\text{T}[V'] \leftarrow 1$ Ret V'</p> <p>subroutine $\text{R}(x, V, M)$</p> <p>$\mathcal{V} \stackrel{\leftarrow}{=} V$; $V' \leftarrow_s \{0, 1\}^n \setminus \mathcal{V}$ If $\text{R}[x, V, M] \neq \perp$ then $V' \leftarrow \text{R}[x, V, M]$ $\mathcal{V} \stackrel{\leftarrow}{=} V'$ Ret $\text{R}[x, V, M] \leftarrow V'$</p>
--	---	--

Figure 12: Games used in the proof of Lemma 4.5.

the game's restrictions on the choice of V' , we have that \mathcal{G} is throughout the game a forest. The root of one tree is IV and all others are rooted at adversarially-chosen values V (corresponding to queries to Prim_a).

We now argue by case analysis that the probability of recoloring is zero for any call to R but one made on line 21. Recall that we disallow pointless queries, meaning the adversary never queries Func_a or Prim_a twice on the same values.

- Line 10: A recoloring here means that at the time of the query a path IV, V_1, V_2, \dots, V_k for some number k exists in \mathcal{G} with edge (IV, V_1) labeled by $X = K \oplus \text{ipad}$ for the queried K and the color of node V_k was 1. But the only such paths that can exist in \mathcal{G} must have (IV, V_1) labeled by an X with $\text{lsOuter}(X) = 1$. But this contradicts that $K \in \mathcal{K}$.
- Line 11: No paths of length one rooted at IV can exist in \mathcal{G} with the second node colored 1.
- Line 12: A recoloring here means that at the time of the query a path IV, V_1, V_2 exists in \mathcal{G} with edge (IV, V_1) labeled by $X = K \oplus \text{opad}$ for the queried K and V_3 is colored 0. If the path was formed due to a Func_a query, then this must mean that $X = K' \oplus \text{ipad}$ for some other K' and this contradicts that $K, K' \in \mathcal{K}$. If the path was formed due to Prim_a queries, then the check on line 22 means that the $\text{lsOuter}(X) = 0$, but this contradicts that $K \in \mathcal{K}$.
- Line 20: A recoloring here means that at the time of the query a path IV, V_1, V_2 exists in \mathcal{G} with edge (IV, V_1) labeled by X with $\text{lsOuter}(X) = 1$ (by the check on line 22) and V_3 colored by 0.

But the only way such a path can exist is due to execution of line 10 with $X = K \oplus \text{ipad}$ for the value K of this prior query. This contradicts that $K \in \mathcal{K}$.

The flag `bad` in G_3 can only be set due to a query on line 21, and so game G_4 makes this explicit by moving the setting of `bad` to Prim_a . Also, we make the setting of `bad` more liberal by only tracking chaining variable values. These changes have no effect on the values returned to the adversary. We have that

$$\Pr[G_3 \Rightarrow 1] = \Pr[G_4 \Rightarrow 1] \quad \text{and} \quad \Pr[G_3 \text{ sets bad}] \leq \Pr[G_4 \text{ sets bad}] .$$

We now bound the setting of `bad` in G_4 . The flag is set due to a query (V, M) to Prim_a such that: (1) $V \in \mathcal{X}$, meaning $V = V_o$ for the latter chosen in Func_a as the result of $V_o \leftarrow R_0(IV, K \oplus \text{opad})$; and (2) $T[V] \neq 1$, meaning that no previous query to Prim_a was made on IV, M' with $\text{lsOuter}(M') = 1$. But together, (1) and (2) imply that the value $V \in \mathcal{X}$ has not yet been returned to the adversary at the time of the query. This means that the adversary has no knowledge of the coins underlying the choice of V and can narrow it down only by the fact that V is not equal to any other chaining variable returned. Thus, the probability that the adversary can query V is at most $1/(2^n - |\mathcal{V}|)$ where $|\mathcal{V}|$ is the size of \mathcal{V} at the time of the query. Taking a union bound over all queries to Prim_a and using the facts that $|\mathcal{V}| \leq 2\sigma$ and $|\mathcal{X}| \leq q_1$ we have that

$$\Pr[G_4 \text{ sets bad}] \leq \frac{q_1 q_2}{2^n - 2\sigma} \leq \frac{2\sigma^2}{2^n}$$

and where we have additionally used our restriction that $\sigma \leq 2^{n-2}$.

Finally, game G_5 relaxes the restrictions on selection of chaining variables. A birthday-bound argument establishes that

$$\Pr[G_4 \Rightarrow 1] \leq \Pr[G_5 \Rightarrow 1] + \frac{(2\sigma)^2}{2^n} .$$

Combining the above equations gives the bound claimed in the lemma. \blacksquare

The next lemma shows that $F[f]$ is preimage-aware [19]. See Section 2 for the formal definition of preimage-awareness.

Lemma 4.6 Let f be a random oracle $f: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^n$. Then there exists an extractor \mathcal{E} such that for any adversary \mathcal{A} making q_e extraction queries, at most q f -queries, and outputs a message of total length at most $d\ell$. Then it holds that

$$\mathbf{Adv}_{F[f], \mathcal{E}}^{\text{pra}}(\mathcal{A}) \leq \frac{q_e \ell (q + \ell)}{2^n} + \frac{(q + \ell)^2}{2^n} + \frac{q}{2^n} .$$

□

We just sketch the proof, which is straightforward. Recall that

$$F[f](K, X, P) = f(IV, K \oplus \text{ipad}) \parallel \text{MD}[f]((K \oplus \text{opad}) \parallel X) \parallel P$$

We observe that it is trivial to extract P . For K , one can investigate queries of the form $f(IV, M)$ for $\text{lsOuter}(X) = 1$ and determine K via $\text{GetKey}(M)$. Remaining is extracting X , but this follows from the preimage awareness of $\text{MD}[f]$, which is implied by combining [19, Th. B.1] and [19, Th. 3.2]. \blacksquare

Acknowledgments

The authors thank Hugo Krawczyk for providing significant feedback and suggestions, in particular encouraging the authors to include positive results for the indifferentiability of HMAC; Niels Ferguson for in-depth discussions regarding the security of H^2 ; and the anonymous reviewers for their helpful suggestions. Dodis was supported in part by NSF grants CNS-1065134, CNS-1065288, CNS-1017471, CNS-0831299. Ristenpart was supported in part by NSF grant CNS-1065134. Steinberger is supported by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the National Natural Science Foundation of China Grant 61033001, 61061130540, 61073174, and by NSF grant 0994380. Tessaro was supported in part by NSF grants CCF-0915675, CCF-1018064, and DARPA contracts FA8750-11-C-0096, FA8750-11-2-0225.

References

- [1] Elena Andreeva, Bart Mennink, and Bart Preneel. On the indifferentiability of the Grøstl hash function. In Juan A. Garay and Roberto De Prisco, editors, *SCN 10: 7th International Conference on Security in Communication Networks*, volume 6280 of *Lecture Notes in Computer Science*, pages 88–105. Springer, September 2010.
- [2] Elena Andreeva, Gregory Neven, Bart Preneel, and Thomas Shrimpton. Seven-property-preserving iterated hashing: ROX. In Kaoru Kurosawa, editor, *Advances in Cryptology – ASIACRYPT 2007*, volume 4833 of *Lecture Notes in Computer Science*, pages 130–146. Springer, December 2007.
- [3] Mihir Bellare. New proofs for NMAC and HMAC: Security without collision-resistance. In Cynthia Dwork, editor, *Advances in Cryptology – CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 602–619. Springer, August 2006.
- [4] Mihir Bellare, Zvika Brakerski, Moni Naor, Thomas Ristenpart, Gil Segev, Hovav Shacham, and Scott Yilek. Hedged public-key encryption: How to protect against bad randomness. In Mitsuru Matsui, editor, *Advances in Cryptology – ASIACRYPT 2009*, volume 5912 of *Lecture Notes in Computer Science*, pages 232–249. Springer, December 2009.
- [5] Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Keying hash functions for message authentication. In Neal Koblitz, editor, *Advances in Cryptology – CRYPTO’96*, volume 1109 of *Lecture Notes in Computer Science*, pages 1–15. Springer, August 1996.
- [6] Mihir Bellare and Tadayoshi Kohno. A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications. In Eli Biham, editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 491–506. Springer, May 2003.
- [7] Mihir Bellare and Thomas Ristenpart. Multi-property-preserving hash domain extension and the EMD transform. In Xuejia Lai and Kefei Chen, editors, *Advances in Cryptology – ASIACRYPT 2006*, volume 4284 of *Lecture Notes in Computer Science*, pages 299–314. Springer, December 2006.
- [8] Mihir Bellare and Thomas Ristenpart. Hash functions in the dedicated-key setting: Design choices and MPP transforms. In Lars Arge, Christian Cachin, Tomasz Jurdzinski, and Andrzej Tarlecki, editors, *ICALP 2007: 34th International Colloquium on Automata, Languages and Programming*, volume 4596 of *Lecture Notes in Computer Science*, pages 399–410. Springer, July 2007.

- [9] Mihir Bellare, Thomas Ristenpart, and Stefano Tessaro. Multi-instance security and its application to password-based cryptography. In *Advances in Cryptology – CRYPTO ‘12*, Lecture Notes in Computer Science. Springer, 2012.
- [10] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 409–426. Springer, May / June 2006.
- [11] Mihir Bellare and Bennet S. Yee. Forward-security in private-key cryptography. In Marc Joye, editor, *Topics in Cryptology – CT-RSA 2003*, volume 2612 of *Lecture Notes in Computer Science*, pages 1–18. Springer, April 2003.
- [12] Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. On the indistinguishability of the sponge construction. In Nigel P. Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 181–197. Springer, April 2008.
- [13] Donghoon Chang and Mridul Nandi. Improved indistinguishability security analysis of chopMD hash function. In Kaisa Nyberg, editor, *Fast Software Encryption – FSE 2008*, volume 5086 of *Lecture Notes in Computer Science*, pages 429–443. Springer, February 2008.
- [14] Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. Merkle-Damgård revisited: How to construct a hash function. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 430–448. Springer, August 2005.
- [15] Dana Dachman-Soled, Rosario Gennaro, Hugo Krawczyk, and Tal Malkin. Computational extractors and pseudorandomness. In Ronald Cramer, editor, *Theory of Cryptography – TCC ‘12*, volume 7194 of *Lecture Notes in Computer Science*, pages 383–403. Springer, 2012.
- [16] Ivan Damgård. A design principle for hash functions. In Gilles Brassard, editor, *Advances in Cryptology – CRYPTO’89*, volume 435 of *Lecture Notes in Computer Science*, pages 416–427. Springer, August 1990.
- [17] Yevgeniy Dodis, Rosario Gennaro, Johan Håstad, Hugo Krawczyk, and Tal Rabin. Randomness extraction and key derivation using the CBC, cascade and HMAC modes. In Matthew Franklin, editor, *Advances in Cryptology – CRYPTO 2004*, volume 3152 of *Lecture Notes in Computer Science*, pages 494–510. Springer, August 2004.
- [18] Yevgeniy Dodis, Leonid Reyzin, Ronald L. Rivest, and Emily Shen. Indistinguishability of permutation-based compression functions and tree-based modes of operation, with applications to MD6. In Orr Dunkelman, editor, *Fast Software Encryption – FSE 2009*, volume 5665 of *Lecture Notes in Computer Science*, pages 104–121. Springer, February 2009.
- [19] Yevgeniy Dodis, Thomas Ristenpart, and Thomas Shrimpton. Salvaging Merkle-Damgård for practical applications. In Antoine Joux, editor, *Advances in Cryptology – EUROCRYPT 2009*, volume 5479 of *Lecture Notes in Computer Science*, pages 371–388. Springer, April 2009.
- [20] Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In Ernest F. Brickell, editor, *Advances in Cryptology – CRYPTO’92*, volume 740 of *Lecture Notes in Computer Science*, pages 139–147. Springer, August 1993.

- [21] Cynthia Dwork, Moni Naor, and Hoeteck Wee. Pebbling and proofs of work. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 37–54. Springer, August 2005.
- [22] Niels Ferguson and Bruce Schneier. *Practical cryptography*. Wiley, 2003.
- [23] Pierre-Alain Fouque, David Pointcheval, and Sébastien Zimmer. HMAC is a randomness extractor and applications to TLS. In Masayuki Abe and Virgil Gligor, editors, *ASIACCS 08: 3rd Conference on Computer and Communications Security*, pages 21–32. ACM Press, March 2008.
- [24] J. Franks, P. Hallam-Baker, J. Hostetler, P. Leach, A. Luotonen, E. Sink, and L. Stewart. An Extension to HTTP: Digest Access Authentication. RFC 2069 (Proposed Standard), January 1997. Obsoleted by RFC 2617.
- [25] Dan Harkins and Dave Carrel. The Internet Key Exchange (IKE). IETF RFC 2409 (Proposed Standard), 1998.
- [26] Shoichi Hirose, Je Hong Park, and Aaram Yun. A simple variant of the Merkle-Damgård scheme with a permutation. In Kaoru Kurosawa, editor, *Advances in Cryptology – ASIACRYPT 2007*, volume 4833 of *Lecture Notes in Computer Science*, pages 113–129. Springer, December 2007.
- [27] Ari Juels and John G. Brainard. Client puzzles: A cryptographic countermeasure against connection depletion attacks. In *ISOC Network and Distributed System Security Symposium – NDSS’99*. The Internet Society, February 1999.
- [28] Ghassan Karame and Srdjan Capkun. Low-cost client puzzles based on modular exponentiation. In Dimitris Gritzalis, Bart Preneel, and Marianthi Theoharidou, editors, *ESORICS 2010: 15th European Symposium on Research in Computer Security*, volume 6345 of *Lecture Notes in Computer Science*, pages 679–697. Springer, 2010.
- [29] Charlie Kaufman. The Internet Key Exchange (IKEv2) Protocol. IETF RFC 4306 (Proposed Standard), 2005.
- [30] H. Krawczyk, M. Bellare, and R. Canetti. HMAC: Keyed-Hashing for Message Authentication. RFC 2104, February 1997.
- [31] H. Krawczyk and P. Eronen. Hmac-based extract-and-expand key derivation function (hkdf). RFC 5869 (Proposed Standard), January 2010.
- [32] Hugo Krawczyk. Cryptographic extraction and key derivation: The HKDF scheme. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 631–648. Springer, August 2010.
- [33] Ueli M. Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In Moni Naor, editor, *TCC 2004: 1st Theory of Cryptography Conference*, volume 2951 of *Lecture Notes in Computer Science*, pages 21–39. Springer, February 2004.
- [34] Ralph C. Merkle. One way hash functions and DES. In Gilles Brassard, editor, *Advances in Cryptology – CRYPTO’89*, volume 435 of *Lecture Notes in Computer Science*, pages 428–446. Springer, August 1990.

- [35] Paul Morrissey, Nigel P. Smart, and Bogdan Warinschi. A modular security analysis of the TLS handshake protocol. In Josef Pieprzyk, editor, *Advances in Cryptology – ASIACRYPT 2008*, volume 5350 of *Lecture Notes in Computer Science*, pages 55–73. Springer, December 2008.
- [36] PKCS #5: Password-based cryptography standard (rfc 2898). RSA Data Security, Inc., September 2000. Version 2.0.
- [37] Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. Careful with composition: Limitations of the indistinguishability framework. In Kenneth G. Paterson, editor, *Advances in Cryptology – EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 487–506. Springer, May 2011.
- [38] Thomas Ristenpart and Scott Yilek. When good randomness goes bad: Virtual machine reset vulnerabilities and hedging deployed cryptography. In *Network and Distributed Systems Security – NDSS ’10*. ISOC, 2010.
- [39] Douglas Stebila, Lakshmi Kuppusamy, Jothi Rangan, Colin Boyd, and Juan Manuel González Nieto. Stronger difficulty notions for client puzzles and denial-of-service-resistant protocols. In Aggelos Kiayias, editor, *Topics in Cryptology – CT-RSA 2011*, volume 6558 of *Lecture Notes in Computer Science*, pages 284–301. Springer, February 2011.
- [40] Gene Tsudik. Message authentication with one-way hash functions. In *Proceedings IEEE INFOCOM’92*, volume 3, pages 2055–2059. IEEE, 1992.
- [41] XiaoFeng Wang and Michael K. Reiter. Defending against denial-of-service attacks with puzzle auction. In *IEEE Symposium on Security and Privacy*, pages 78–92, 2003.
- [42] Scott Yilek. Resettable public-key encryption: How to encrypt on a virtual machine. In Josef Pieprzyk, editor, *Topics in Cryptology – CT-RSA 2010*, volume 5985 of *Lecture Notes in Computer Science*, pages 41–56. Springer, March 2010.

A Internal Collision Probabilities

We briefly discuss computing the probabilities $p(H, w, \ell)$ and $p'(H, w, \ell)$ for the case where $H[P] = P = \mathcal{R}$.

To compute $p(H, w, \ell)$ for $H^2 = H^2[P]$, let z_0 be a randomly chosen n -bit string, and define z_i such that $z_i = \mathcal{R}(z_{i-1})$ for all $i = 1, \dots, 2\ell + 1$. Note that in particular $z_0, z_1, \dots, z_{2\ell}, z_{2\ell+1}$ correspond to the values $x_0, y_0, \dots, x_\ell, y_\ell$ in the definition of $p(H, w, \ell)$. We first upper bound the probability of the event bad_1 that there is a collision among the z values, which is

$$\begin{aligned} \Pr[\text{bad}] &\leq \sum_{i=0}^{2\ell+1} \Pr[z_i \in \{z_0, z_1, \dots, z_{i-1}\} \mid |\{z_0, z_1, \dots, z_{i-1}\}| = i] \\ &= \sum_{i=0}^{2\ell+1} \frac{i}{2^n} \leq \frac{(2\ell+1)(2\ell+2)}{2 \cdot 2^n} \leq \frac{2 \cdot (\ell+1)^2}{2^n}. \end{aligned}$$

Conditioned on $\overline{\text{bad}}$, we have $x_\ell \notin \{y_0, y_1, \dots, y_\ell\}$. Moreover, let bad_2 be the event that for some

<p>adversary $\mathcal{D}_{w,\ell}^{\text{Func,Prim}}$:</p> <p>$\mathbf{u}[1], \dots, \mathbf{u}[w] \leftarrow_{\\$} \{0, 1\}^n$ $x_0 \leftarrow_{\\$} \{0, 1\}^n ; j \leftarrow_{\\$} [1 .. w]$ For $i = 1$ to ℓ do $x_i \leftarrow \text{Func}(x_{i-1})$ $\mathbf{u}[j] \leftarrow x_\ell$ For $i = 1$ to w do $\mathbf{v}[i] \leftarrow_{\\$} H[\text{Prim}](\mathbf{u}[i])$ $y_0 \leftarrow_{\\$} H[\text{Prim}](x_0)$ For $i = 1$ to ℓ do $y_i \leftarrow_{\\$} \text{Func}(y_{i-1})$ Ret $(y_\ell = \mathbf{v}[j]) \wedge (x_\ell \notin \{y_0, y_1, \dots, y_\ell\})$ $\wedge (\forall i \neq j : \mathbf{v}[i] \notin \{y_0, y_1, \dots, y_\ell\})$.</p>	<p>adversary $\tilde{\mathcal{D}}_{w,\ell}^{\text{Func,Prim}}$:</p> <p>$\mathbf{u}[1], \dots, \mathbf{u}[w] \leftarrow_{\\$} \{0, 1\}^n$ $x_0 \leftarrow_{\\$} \{0, 1\}^n ; j \leftarrow_{\\$} [1 .. w]$ For $i = 1$ to ℓ do $x_i \leftarrow \text{Func}(0^d, x_{i-1})$ $\mathbf{u}[j] \leftarrow x_\ell$ For $i = 1$ to w do $\mathbf{v}[i] \leftarrow_{\\$} H[\text{Prim}](\text{ipad} \parallel \mathbf{u}[j])$ $y_0 \leftarrow_{\\$} H[\text{Prim}](\text{ipad} \parallel x_0)$ For $i = 1$ to ℓ do $y_i \leftarrow_{\\$} \text{Func}(\text{opad} \oplus \text{ipad}, y_{i-1})$ Ret $(y_\ell = \mathbf{v}[j]) \wedge (\forall i \neq j : \mathbf{v}[i] \notin \{y_0, y_1, \dots, y_\ell\})$.</p>
--	--

Figure 13: Left: Adversary $\mathcal{D}_{w,\ell}$ used in the proof of Theorem 3.2. Right: Adversary $\tilde{\mathcal{D}}_{w,\ell}$ used in the proof of Theorem 4.2. The notation $H[\text{Prim}](x)$ indicates the evaluation of the hash construction H on input x in which P queries are replied by the corresponding Prim queries.

$i \in [1 .. w - 1]$ we get $\mathcal{R}(U_i) \in \{y_0, y_1, \dots, y_\ell\}$. Then,

$$\begin{aligned}
\Pr [\text{bad}_2 \mid \overline{\text{bad}}_1] &\leq \sum_{i \neq j} \Pr [U_i \in \{x_0, x_1, \dots, x_\ell\} \mid \overline{\text{bad}}_1] \\
&\quad + \sum_{i \neq j} \Pr [\mathcal{R}(U_i) \in \{y_0, y_1, \dots, y_\ell\} \mid U_i \notin \{x_0, x_1, \dots, x_\ell\} \wedge \overline{\text{bad}}_1] \\
&= \frac{2(w-1) \cdot (\ell+2)}{2^n}.
\end{aligned}$$

Therefore, $p(\mathcal{R}, w, \ell) \leq \Pr [\text{bad}_1] + \Pr [\text{bad}_2 \mid \overline{\text{bad}}_1] \leq \frac{2(w-1) \cdot (\ell+2)}{2^n} + \frac{2 \cdot (\ell+1)^2}{2^n}$.

It is not hard to see that the same upper bound can be computed for p' , where the z -values are the intermediate values with respect to \mathcal{R} and $\text{HMAC}[\mathcal{R}]$.

B Proof of Theorem 3.2

A formal description of the adversary $\mathcal{D}_{w,\ell}$ sketched above is provided in Figure 13.

As the first step, note that by the chain-shift property of $H^2[P]$, the condition $(y_\ell = \mathbf{v}[j])$ is always satisfied in the real-world experiment. Consequently,

$$\Pr \left[\text{Real}_{H^2[P]}^{\mathcal{D}_{w,\ell}} \Rightarrow 1 \right] \geq 1 - p(H, w, \ell).$$

The remainder of this proof consists of upper bounding the probability $\Pr \left[\text{Ideal}_{\text{RO}, S[\text{RO}]}^{\mathcal{D}_{w,\ell}} \Rightarrow 1 \right]$ under the constraint that the simulator makes q_S queries. To this end, it is convenient to introduce a security game – which we refer to as G and is described in Figure 14– involving a (stateful) adversary \mathcal{B} and a random oracle RO . The adversary \mathcal{B} is given a vector \mathbf{u} of w n -bit strings. $n - 1$ of these are chosen uniformly at random and independently. The remaining one, placed in a randomly chosen component $j \in [1 .. w]$, is the output x_ℓ of a chain of ℓ random-oracle invocations starting at a randomly chosen initial n -bit string x_0 . The adversary is asked to output a vector \mathbf{v} of w *distinct* n -bit values, to which

<pre> procedure main: // G phase ← 1 <u>u</u>[1], ..., <u>u</u>[w] ←_s {0, 1}ⁿ x₀ ←_s {0, 1}ⁿ; j ←_s [1 .. w] For i = 1 to ℓ do x_i ←_s RO(x_{i-1}) <u>u</u>[j] ← x_ℓ <u>v</u> ←_s \mathcal{B}^{RO}(<u>u</u>); phase ← 2; y₀ ← \mathcal{B}^{RO}(x₀) For i = 1 to ℓ do y_i ←_s RO(y_{i-1}) Ret (y_ℓ = <u>v</u>[j]) ∧ (x_ℓ ∉ {y₀, y₁, ..., y_ℓ}) ∧ (∀i ≠ j : <u>v</u>[i] ∉ {y₀, y₁, ..., y_ℓ}) procedure RO(x): // G, G₀ - G₂, G₅ If $\mathbb{R}[x] = \perp$ then $\mathbb{R}[x] \leftarrow_s \{0, 1\}^n$ Ret $\mathbb{R}[x]$ </pre>	<pre> procedure main: // G₀ phase ← 1 <u>u</u>[1], ..., <u>u</u>[w] ←_s {0, 1}ⁿ x₀ ←_s {0, 1}ⁿ; j ←_s [1 .. w] For i = 1 to ℓ do x_i ←_s RO(x_{i-1}) <u>u</u>[j] ← x_ℓ <u>v</u> ←_s \mathcal{B}^{RO}(<u>u</u>); phase ← 2 y₀ ← \mathcal{B}^{RO}(x₀) Ret (FC(j) ≠ ∅) subroutine FC(i) // G₀ - G₇ P ← ∅ For all (y'₀, y'₁, ..., y'_ℓ) do If (phase = 2) ∧ (y'_ℓ = <u>v</u>[i]) ∧ (∀k ∈ [1 .. ℓ] : $\mathbb{R}[y'_{k-1}] = y'_k$) ∧ (∀i' ≠ i : <u>v</u>[i'] ∉ {y'₀, y'₁, ..., y'_ℓ}) ∧ (x_ℓ ∉ {y'₀, y'₁, ..., y'_ℓ}) then P ←_s {(y₀, y₁, ..., y_ℓ)} Ret P </pre>
--	--

Figure 14: Games G and G₀. Note that the adversary \mathcal{B} keeps a state across its first and second invocations.

it commits. At this point, x_0 is revealed to \mathcal{B} (and hence, indirectly, also j), which now needs to output a value y_0 such that the output y_ℓ of an chain of RO invocations starting at y_0 equals $\mathbf{v}[j]$, and such that x_ℓ , as well as $\mathbf{v}[i]$ for $i \neq j$, is not part of this chain.

The following lemma establishes the rather direct relation between the task of building a simulator for $H^2[P]$ and providing a good adversary \mathcal{B} for the game G.

Lemma B.1 For all simulators \mathcal{S} making $q_{\mathcal{S}}$ queries, there exists an adversary \mathcal{B} making $q' = q_{\mathcal{S}} + \ell$ queries such that

$$\Pr \left[\text{Ideal}_{\mathcal{R}, \mathcal{S}}^{\mathcal{D}_{w, \ell}} \Rightarrow 1 \right] = \Pr \left[G^{\mathcal{B}} \Rightarrow 1 \right] .$$

Moreover, whenever \mathcal{B} outputs y_0 , then it has issued all RO queries to compute the chain y_0, y_1, \dots, y_ℓ starting in y_0 .

Proof of of Lemma B.1: The adversary \mathcal{B} , upon receiving the vector \mathbf{u} , runs an execution of the simulator \mathcal{S} , feeding it with Prim queries $\mathbf{u}[1], \dots, \mathbf{u}[m]$, and obtaining replies $\mathbf{v}[1], \dots, \mathbf{v}[m]$. It outputs \mathbf{v} to conclude the first phase of G. In particular, RO queries by the simulator are replied directly by the oracle in the game G. Then, upon receiving x_0 , \mathcal{B} continues the execution of the simulator (recall that \mathcal{B} is safe), asking Prim queries to \mathcal{S} in order to compute $y_0 = H[P](x_0)$. It continues by making all RO queries to evaluate the chain y_0, y_1, \dots, y_ℓ such that $y_i = \text{RO}(y_{i-1})$. Finally, \mathcal{B} outputs y_0 . It is not hard to verify that the probability that \mathcal{B} wins the game is exactly $\Pr \left[\text{Ideal}_{\mathcal{R}, \mathcal{S}}^{\mathcal{D}_{w, \ell}} \Rightarrow 1 \right]$. ■




Figure 15: Pseudocode descriptions of Games $G_1 - G_4$.

In the following, we focus on showing an upper bound on $\Pr [G^{\mathcal{B}} \Rightarrow 1]$ for a q' -query adversary \mathcal{B} which, without loss of generality, outputs a value y_0 such that it has asked all RO queries defining a chain $(y_0, y_1, \dots, y_\ell)$ starting in y_0 . As a first step, we consider the game G_0 , depicted in Figure 14, which is similar to G , but possibly slightly easier to win. In G_0 , we introduce a sub-routine, called FC, which on input i returns the sets of tuples $(y'_0, y'_1, \dots, y'_\ell)$ such that, with respect to RO queries asked so far, define a chain such that $y'_\ell = \mathbf{v}[i]$ and $y'_i = \text{RO}(y'_{i-1})$ for $i = 1, \dots, \ell$, and x_ℓ and $\mathbf{v}[i]$ for $i \neq j$ is not part of this chain. In particular, if \mathbf{v} has not been defined yet (i.e., phase = 1), then FC always returns the empty set. We also modify the winning condition so that it returns true as long as there is a chain $(y'_0, y'_1, \dots, y'_\ell) \in \text{FC}(j)$ (i.e., it does not need to be the one starting at the value y_0 output by \mathcal{B}). Then, clearly

$$\Pr [G^{\mathcal{B}} \Rightarrow \text{true}] \leq \Pr [G_0^{\mathcal{B}} \Rightarrow \text{true}] ,$$

as \mathcal{B} winning in G implies \mathcal{B} winning G_0 because of the fact that \mathcal{B} has asked all queries corresponding to the chain starting at the output string y_0 .

We continue with Game G_1 (cf. Figure 15) which is equivalent to G_0 . The (merely syntactical) difference is that it starts by first choosing all components of \mathbf{u} uniformly at random. It then compute the chain x_0, x_1, \dots only up to $x_{\ell-1}$, and then sets $\mathbf{R}[x_{\ell-1}]$ to equal $\mathbf{u}[j]$ if it is undefined, whereas otherwise it overwrites $\mathbf{u}[j]$ as $\mathbf{u}[j] \leftarrow \mathbf{R}[x_{\ell-1}]$, and sets the flag bad. In both cases, x_ℓ equals $\mathbf{u}[j]$. The next Game G_2 (also in Figure 15) is then obtained from G_1 by modifying the latter case so that the value $\mathbf{u}[j]$ is *not* overwritten. Clearly, G_1 and G_2 are equivalent until bad. With $\text{Adv}(G_1^{\mathcal{B}}, G_2^{\mathcal{B}}) =$

<pre> procedure main: // G₅, G₆ phase ← 1 <u>u</u>[1], . . . , u[w] ←_s {0, 1}ⁿ v ←_s $\mathcal{B}^{\text{RO}}(\mathbf{u})$; phase ← 2 x₀ ←_s {0, 1}ⁿ; j ←_s [1..w]; x_ℓ ← u[j] For i = 1 to ℓ − 1 do x_i ←_s RO(x_{i−1}) If R[x_{ℓ−1}] = ⊥ then R[x_{ℓ−1}] ← u[j] If (FC(j) ≠ ∅) ∧ (FC'(j) = ∅) then bad ← true x_ℓ ← u[j] y₀ ← $\mathcal{B}^{\text{RO}}(x_0)$ Ret (FC(j) ≠ ∅) </pre>	<pre> procedure main: // G₇ phase ← 1 <u>u</u>[1], . . . , u[w] ←_s {0, 1}ⁿ v ←_s $\mathcal{B}^{\text{RO}}(\mathbf{u})$; phase ← 2 x₀ ←_s {0, 1}ⁿ; j ←_s [1..w]; x_ℓ ← u[j] For i = 1 to ℓ − 1 do x_i ←_s RO(x_{i−1}) If R[x_{ℓ−1}] = ⊥ then R[x_{ℓ−1}] ← u[j] If (FC(j) ≠ ∅) ∧ (FC'(j) = ∅) then bad ← true y₀ ← $\mathcal{B}^{\text{RO}}(x_0)$ Ret (FC'(j) ≠ ∅) </pre>
<pre> subroutine FC'(i): // G₆ − G₇ P ← ∅ For all (y'₀, y'₁, . . . , y'_ℓ) do If (phase = 2) ∧ (y'_ℓ = v[i]) ∧ (∀k ∈ [1..ℓ]: R'[y'_{k−1}] = y'_k) ∧ (∀i' ≠ i: v[i'] ∉ {y'₀, y'₁, . . . , y'_ℓ}) ∧ (x_ℓ ∉ {y'₀, y'₁, . . . , y'_ℓ}) then P ←_⊔ {(y₀, y₁, . . . , y_ℓ)} Ret P </pre>	<pre> procedure RO(x): // G₆ − G₇ If R[x] = ⊥ then R[x] ←_s {0, 1}ⁿ If phase = 1 then R'[x] ← R[x] If (phase = 2) ∧ (FC(j) ≠ ∅) ∧ (FC'(j) = ∅) then bad ← true Ret R[x] </pre>

Figure 16: Pseudocode descriptions of Games G₅ - G₇.

$\Pr [G_1^{\mathcal{B}} \Rightarrow \text{true}] - \Pr [G_2^{\mathcal{B}} \Rightarrow \text{true}]$, by the fundamental lemma of game playing we obtain

$$\mathbf{Adv}(G_1^{\mathcal{B}}, G_2^{\mathcal{B}}) \leq \Pr [G_1^{\mathcal{B}} \text{ sets bad }] \leq \sum_{i=0}^{\ell-2} (i+1) \cdot 2^{-n} \leq \frac{\ell^2}{2^{n+1}},$$

as conditioned on x_0, \dots, x_i being distinct, each output $\text{RO}(x_i)$ for $i \in [0.. \ell - 2]$ is in $\{x_0, \dots, x_i\}$ with probability $(i+1)/2^n$.

Note that in G_2 , the input \mathbf{u} to \mathcal{B} is random and independent of everything else, as long as \mathcal{B} does not query $x_{\ell-1}$ to RO . We first transition from G_2 into new games G_3 and G_4 (also on Figure 15). In G_3 , the game sets **bad** if the query $\text{RO}(x_{\ell-1})$ is made before \mathcal{B} commits to \mathbf{v} . Clearly, $\mathbf{Adv}(G_2^{\mathcal{B}}, G_3^{\mathcal{B}}) = 0$. Additionally, G_4 is modified so that the value $\mathbf{R}[x_{\ell-1}]$ is set to equal $\mathbf{u}[j]$ only *after* \mathcal{B} has output \mathbf{v} . Hence, G_3 and G_4 are equivalent until **bad**, and

$$\mathbf{Adv}(G_3^{\mathcal{B}}, G_4^{\mathcal{B}}) \leq \Pr [G_4^{\mathcal{B}} \text{ sets bad }] .$$

We postpone an analysis of $\Pr [G_4^{\mathcal{B}} \text{ sets bad }]$ to a later point in the proof, and now continue with the main sequence of games.

The next game, Game G_5 (cf. Figure 16), simply rearranges the contents of Game G_4 for better readability, but is otherwise fully equivalent. In particular, we have now postponed the computation of the values x_1, \dots, x_{ℓ} to after \mathcal{B} outputs \mathbf{v} , which clearly does not affect the game. Then, we transition

to a game G_6 which takes into account (via a second procedure FC') which chains have been created before the adversary outputs \mathbf{v} , and after this, sets the condition \mathbf{bad} as soon as some new $\mathbf{R}[\cdot]$ entry is defined such that $FC'(j) = \emptyset$ but $FC(j) \neq \emptyset$ for the chosen j . Finally, G_7 is the same as G_6 , but the winning condition checks for $(y'_0, y'_1, \dots, y'_\ell) \in FC'(j)$. Clearly, G_6 and G_7 are equivalent until \mathbf{bad} , and thus

$$\mathbf{Adv}(G_6^{\mathcal{B}}, G_7^{\mathcal{B}}) \leq \Pr [G_7^{\mathcal{B}} \text{ sets } \mathbf{bad}] .$$

Finally, combining all transitions, we obtain

$$\Pr [G^{\mathcal{B}} \Rightarrow \text{true}] \leq \frac{\ell^2}{2^{n+1}} + \Pr [G_4^{\mathcal{B}} \text{ sets } \mathbf{bad}] + \Pr [G_7^{\mathcal{B}} \text{ sets } \mathbf{bad}] + \Pr [G_7^{\mathcal{B}} \Rightarrow \text{true}] .$$

To conclude the proof, we now turn to upper bounding the three probabilities on the RHS.

UPPER BOUNDING $\Pr [G_4^{\mathcal{B}} \text{ sets } \mathbf{bad}]$. By inspection, it is not hard to verify that the probability that \mathbf{bad} is set in G_4 equals the probability that \mathcal{B} wins the following game G'_4 :

<p>procedure main: // G'_4</p> <p>$Q \leftarrow \emptyset$; $\mathbf{u}[1], \dots, \mathbf{u}[w] \leftarrow_{\\$} \{0, 1\}^n$</p> <p>$\mathbf{v} \leftarrow_{\\$} \mathcal{B}^{\text{RO}}(\mathbf{u})$</p> <p>$x_0 \leftarrow_{\\$} \{0, 1\}^n$; $j \leftarrow_{\\$} [1 .. w]$</p> <p>For $i = 1$ to $\ell - 1$ do</p> <p style="padding-left: 20px;">If $\mathbf{R}[x_{i-1}] = \perp$ then $\mathbf{R}[x_{i-1}] \leftarrow_{\\$} \{0, 1\}^n$</p> <p style="padding-left: 20px;">$x_i \leftarrow \mathbf{R}[x_{i-1}]$</p> <p>Ret $x_{\ell-1} \in Q$</p>	<p>procedure RO(x): // G'_4</p> <p>$Q \stackrel{\cup}{\leftarrow} \{x\}$</p> <p>If $\mathbf{R}[x] = \perp$ then</p> <p style="padding-left: 20px;">$\mathbf{R}[x] \leftarrow_{\\$} \{0, 1\}^n$</p> <p>Ret $\mathbf{R}[x]$</p>
--	--

This is because first, we can focus on phase 1, and second, we can postpone the generation of the values $x_0, x_1, \dots, x_{\ell-1}$ to after the execution of the adversary \mathcal{B} . To upper bound the probability, let \mathbf{bad}_i be the event that of $\{x_0, x_1, \dots, x_i\} \cap Q \neq \emptyset$. Then,

$$\begin{aligned} \Pr [G_4^{\mathcal{B}} \text{ sets } \mathbf{bad}] &= \Pr [G'_4^{\mathcal{B}} \Rightarrow \text{true}] \\ &\leq \Pr [\mathbf{bad}_{\ell-1}] \\ &\leq \Pr [\mathbf{bad}_0] + \sum_{i=1}^{\ell-1} \Pr [\mathbf{bad}_i \mid \overline{\mathbf{bad}_{i-1}}] \leq \frac{\ell \cdot q'}{2^n} , \end{aligned}$$

since $\Pr [\mathbf{bad}_0] = q'/2^n$, and conditioned on \mathbf{bad}_{i-1} not having occurred, we have two cases when x_{i-1} is defined: (i) $\mathbf{R}[x_{i-1}] = \perp$. Here, $x_i = \text{RO}(x_{i-1}) = \mathbf{R}[x_{i-1}]$ is chosen uniformly, and hence it collides with one of the values in Q with probability at most $q'/2^n$. (ii) $\mathbf{R}[x_{i-1}] \neq \perp$. But then, this means that $x_{i-1} = x_j$ for $j < i - 1$, as $x_j \notin Q$ by $\overline{\mathbf{bad}_{i-1}}$. In turn, it must also be that $\mathbf{R}[x_{i-1}] = \mathbf{R}[x_j] = x_{j+1} \notin Q$ as $\overline{\mathbf{bad}_{i-1}}$ holds, and thus the probability of provoking \mathbf{bad}_i is 0.

UPPER BOUNDING $\Pr [G_7^{\mathcal{B}} \text{ sets } \mathbf{bad}]$. We seek for an upper bound on the probability that a RO query x in the second phase of the game G_7 provokes $FC(j) \neq \emptyset$ when $FC'(j) = \emptyset$. First, note that it can never be that setting $\mathbf{R}[x_{\ell-1}]$ to $\mathbf{u}[j]$ provokes $FC(j)$ to become non-empty: Every new chain $(y'_0, y'_1, \dots, y'_\ell)$ being defined thanks to $\mathbf{R}[x_{\ell-1}]$ being set to $\mathbf{u}[j] = x_\ell$ contains x_ℓ , and hence is not returned by $FC(j)$. Therefore, a new chain $(y'_0, y'_1, \dots, y'_\ell)$ must be defined when $\mathbf{R}[x]$ is set to a fresh random value, for some string x .

Let Z be the set of strings $z_0 \in \{0, 1\}^n$ such that there exists $z_1, \dots, z_{\ell'}$ with $\ell' < \ell$, $z_{\ell'} = \mathbf{v}[j]$, and $\mathbf{R}'[z_{i-1}] = z_i$ for all $i = 1, \dots, \ell'$. Then, note that as long as no random choice of $\mathbf{R}[x]$ in phase 2 hits

one element of Z , the set $\text{FC}(j)$ remains empty. The probability that one such values hits Z is clearly at most $(q' \cdot |Z|)/2^n$. Since $|Z| \leq q'$, then this means

$$\Pr [G_7^{\mathcal{B}} \text{ sets bad}] \leq q' \cdot 2^{-n} .$$

UPPER BOUNDING $\Pr [G_7^{\mathcal{B}} \Rightarrow \text{true}]$. Assume that \mathcal{B} indeed outputs a vector \mathbf{v} with w distinct components after making at most q' queries. Note that any two chains $(y'_0, y'_1, \dots, y'_\ell) \in \text{FC}'(i')$ and $(y''_0, y''_1, \dots, y''_\ell) \in \text{FC}'(i'')$ are disjoint. As each chain is built by means of asking ℓ RO-queries, this means that there are at most $\lfloor \frac{q_S}{\ell} \rfloor$ indices j for which $\text{FC}'(j)$ is not empty. Hence,

$$\Pr [G_7^{\mathcal{B}} \Rightarrow \text{true}] \leq \frac{1}{w} \frac{q'}{\ell} ,$$

as the choice of j is independent of the behavior of the adversary \mathcal{B} so far. ■

C Proof of Theorem 4.2

Let $H = H[P]$ be the hash function with n -bit output, and let $\text{HMAC}[H]$ be the corresponding instantiation of HMAC using H . A description of the adversary $\tilde{\mathcal{D}}_{w,\ell}$ is given in Figure 13, on the right. The attacker first chooses a random n -bit string x_0 , and then builds a chain of values of $(x_0, x_1, \dots, x_\ell)$ so that $x_i = \text{Func}(0^d, x_{i-1})$ for all $i \in [1.. \ell]$. It then chooses a random index $j \in [1..w]$, and sets $\mathbf{u}[j] \leftarrow x_\ell$, whereas $\mathbf{u}[i]$ is set to a random value for all $i \neq j$. The next step has $\tilde{\mathcal{D}}_{w,\ell}$ asking all necessary Prim queries to evaluate H on input $\text{ipad} \parallel \mathbf{u}[i]$ for all $i \in [1..w]$, obtaining outputs $\mathbf{v}[1], \dots, \mathbf{v}[w]$. Subsequently, $\tilde{\mathcal{D}}_{w,\ell}$ also makes all Prim queries needed to evaluate H on input $\text{ipad} \parallel x_0$, and we refer to resulting output n -bit string as y_0 . Finally, $\tilde{\mathcal{D}}_{w,\ell}$ computes y_1, \dots, y_ℓ , where $y_i = \text{Func}(\text{opad} \oplus \text{ipad}, y_{i-1})$ for all $i \in [1.. \ell]$, and outputs 1 if and only if $y_\ell = \mathbf{v}[j]$ (it outputs 0 otherwise).

Let us start by observing that in the real-world experiment $\text{Real}_{\text{HMAC}[H]}^{\tilde{\mathcal{D}}_{w,\ell}}$, by the chain-shift property of HMAC, we will always have $y_\ell = \mathbf{v}[j]$. Therefore, the probability that the attacker $\tilde{\mathcal{D}}_{w,\ell}$ outputs 0 is bounded by the probability that there exists $i \neq j$ such that $\mathbf{v}[i] \in \{y_0, y_1, \dots, y_\ell\}$. The probability that this happens is, by definition, exactly $p'(H, w, \ell)$, and thus

$$\Pr \left[\text{Real}_{\text{HMAC}[H]}^{\tilde{\mathcal{D}}_{w,\ell}} \Rightarrow 1 \right] = 1 - p'(H, w, \ell) .$$

The remainder of this proof analyzes the ideal world experiment $\text{Ideal}_{\mathcal{R},\mathcal{S}}^{\tilde{\mathcal{D}}_{w,\ell}}$. As a first step, we proceed by showing a reduction to a new game, denote $\tilde{\mathcal{G}}$, which we define in Figure 17. This is essentially a version of Game G considered earlier in Figure 14, but using two independent random oracles RO_1 and RO_2 instead.

Lemma C.1 For all simulators \mathcal{S} making q_S queries, there exists an adversary \mathcal{B} making $q' = q_S + \ell$ queries such that

$$\Pr \left[\text{Ideal}_{\mathcal{R},\mathcal{S}}^{\tilde{\mathcal{D}}_{w,\ell}} \Rightarrow 1 \right] = \Pr \left[\tilde{\mathcal{G}}^{\mathcal{B}} \Rightarrow 1 \right] .$$

Moreover, whenever \mathcal{B} outputs y_0 , then it has already issued all RO_2 queries to compute the chain y_0, y_1, \dots, y_ℓ .

Proof of of Lemma C.1: The adversary \mathcal{B} internally simulates the simulator \mathcal{S} . Simulator random oracle queries of the form (ipad, x) are answered as $\text{RO}_1(x)$, queries of the form (opad, x) are answered


```

procedure main: //  $\tilde{G}$ 
phase  $\leftarrow$  1
 $\mathbf{u}[1], \dots, \mathbf{u}[w] \leftarrow_{\mathcal{B}} \{0, 1\}^n$ 
 $x_0 \leftarrow_{\mathcal{B}} \{0, 1\}^n ; j \leftarrow_{\mathcal{B}} [1 .. w]$ 
For  $i = 1$  to  $\ell$  do
   $x_i \leftarrow_{\mathcal{B}} \text{RO}_1(x_{i-1})$ 
 $\mathbf{u}[j] \leftarrow x_\ell$ 
 $\mathbf{v} \leftarrow_{\mathcal{B}} \mathcal{B}^{\text{RO}_1, \text{RO}_2}(\mathbf{u})$ 
phase  $\leftarrow$  2
 $y_0 \leftarrow \mathcal{B}^{\text{RO}_1, \text{RO}_2}(x_0)$ 
For  $i = 1$  to  $\ell$ 
   $y_i \leftarrow_{\mathcal{B}} \text{RO}_2(y_{i-1})$ 
Ret  $(y_\ell = \mathbf{v}[j]) \wedge (\forall i \neq j : \mathbf{v}[i] \notin \{y_0, y_1, \dots, y_\ell\})$ .

procedure main: //  $\tilde{G}_0$ 
phase  $\leftarrow$  1
 $\mathbf{u}[1], \dots, \mathbf{u}[w] \leftarrow_{\mathcal{B}} \{0, 1\}^n$ 
 $x_0 \leftarrow_{\mathcal{B}} \{0, 1\}^n ; j \leftarrow_{\mathcal{B}} [1 .. w]$ 
For  $i = 1$  to  $\ell$  do
   $x_i \leftarrow_{\mathcal{B}} \text{RO}_1(x_{i-1})$ 
 $\mathbf{u}[j] \leftarrow x_\ell$ 
 $\mathbf{v} \leftarrow_{\mathcal{B}} \mathcal{B}^{\text{RO}_1, \text{RO}_2}(\mathbf{u})$ 
phase  $\leftarrow$  2
 $y_0 \leftarrow \mathcal{B}^{\text{RO}_1, \text{RO}_2}(x_0)$ 
Ret  $(\text{FC}(j) \neq \emptyset)$ 

procedure  $\text{RO}_i(x)$ : //  $\tilde{G}, \tilde{G}_0 - \tilde{G}_2, \tilde{G}_5$ 
If  $\mathbf{R}_i[x] = \perp$  then
   $\mathbf{R}_i[x] \leftarrow_{\mathcal{B}} \{0, 1\}^n$ 
Ret  $\mathbf{R}_i[x]$ 

subroutine  $\text{FC}(i)$  //  $\tilde{G}_0 - \tilde{G}_7$ 
 $P \leftarrow \emptyset$ 
For all  $(y'_0, y'_1, \dots, y'_\ell)$  do
  If  $(\text{phase} = 2) \wedge (y'_\ell = \mathbf{v}[i])$ 
     $\wedge (\forall i \neq j : \mathbf{v}[i] \notin \{y_0, y_1, \dots, y_\ell\})$ 
     $\wedge (\forall k \in [1 .. \ell] : \mathbf{R}_2[y'_{k-1}] = y'_k)$  then
       $P \stackrel{\perp}{\leftarrow} \{(y'_0, y'_1, \dots, y'_\ell)\}$ 
Ret  $P$ 

```

```

procedure main: //  $\tilde{G}_1, \tilde{G}_2$ 
phase  $\leftarrow$  1
 $\mathbf{u}[1], \dots, \mathbf{u}[w] \leftarrow_{\mathcal{B}} \{0, 1\}^n$ 
 $x_0 \leftarrow_{\mathcal{B}} \{0, 1\}^n ; j \leftarrow_{\mathcal{B}} [1 .. w]$ 
For  $i = 1$  to  $\ell - 1$  do
   $x_i \leftarrow_{\mathcal{B}} \text{RO}_1(x_{i-1})$ 
If  $\mathbf{R}_1[x_{\ell-1}] = \perp$  then
   $\mathbf{R}_1[x_{\ell-1}] \leftarrow \mathbf{u}[j]$ 
Else
  bad  $\leftarrow$  true ;  $\mathbf{u}[j] \leftarrow \mathbf{R}_1[x_{\ell-1}]$ 
 $x_\ell \leftarrow \mathbf{u}[j]$ 
 $\mathbf{v} \leftarrow_{\mathcal{B}} \mathcal{B}^{\text{RO}_1, \text{RO}_2}(\mathbf{u})$ 
phase  $\leftarrow$  2
 $y_0 \leftarrow \mathcal{B}^{\text{RO}_1, \text{RO}_2}(x_0)$ 
Ret  $(\text{FC}(j) \neq \emptyset)$ 

procedure main: //  $\tilde{G}_3, \tilde{G}_4$ 
phase  $\leftarrow$  1
 $\mathbf{u}[1], \dots, \mathbf{u}[w] \leftarrow_{\mathcal{B}} \{0, 1\}^n$ 
 $x_0 \leftarrow_{\mathcal{B}} \{0, 1\}^n ; j \leftarrow_{\mathcal{B}} [1 .. w]$ 
For  $i = 1$  to  $\ell - 1$  do
   $x_i \leftarrow_{\mathcal{B}} \text{RO}_1(x_{i-1})$ 
  If  $\mathbf{R}_1[x_{\ell-1}] = \perp$  then  $\mathbf{R}_1[x_{\ell-1}] \leftarrow \mathbf{u}[j]$ 
 $x_\ell \leftarrow \mathbf{u}[j]$ 
 $\mathbf{v} \leftarrow_{\mathcal{B}} \mathcal{B}^{\text{RO}_1, \text{RO}_2}(\mathbf{u})$ 
phase  $\leftarrow$  2
If  $\mathbf{R}_1[x_{\ell-1}] = \perp$  then  $\mathbf{R}_1[x_{\ell-1}] \leftarrow \mathbf{u}[j]$ 
 $y_0 \leftarrow \mathcal{B}^{\text{RO}_1, \text{RO}_2}(x_0)$ 
Ret  $(\text{FC}(j) \neq \emptyset)$ 

subroutine  $\text{RO}_i(x)$ : //  $\tilde{G}_3, \tilde{G}_4$ 
If  $\mathbf{R}_i[x] = \perp$  then
  If  $(i = 1) \wedge (x = x_{\ell-1}) \wedge (\text{phase} = 1)$ 
  then
    bad  $\leftarrow$  true
     $\mathbf{R}_i[x] \leftarrow_{\mathcal{B}} \{0, 1\}^n$ 
Ret  $\mathbf{R}_i[x]$ 

```

Figure 17: Games \tilde{G} and $\tilde{G}_1 - \tilde{G}_4$. Note that the adversary \mathcal{B} keeps a state across its first and second invocations.

as $\text{RO}_2(x)$, whereas \mathcal{B} internally simulates answers to all other types of RO queries. Then, when receiving the vector \mathbf{u} , the adversary \mathcal{B} makes Prim queries to \mathcal{S} to evaluate $H(\text{ipad} \parallel \mathbf{u}[i])$ for all i , and outputting the results $\mathbf{v}[1], \dots, \mathbf{v}[w]$ of these evaluations as a vector. Then, upon receiving x_0 , \mathcal{B} asks Prim queries to \mathcal{S} to evaluate H on input $\text{ipad} \parallel x_0$, and once the result y_0 is computed, first makes all RO_2 queries to evaluate the chain of length ℓ starting at y_0 of length ℓ , and then outputs y_0 . Clearly, the probability that \mathcal{B} outputs 1 is the same as the probability that $\tilde{\mathcal{D}}_{w,\ell}$ outputs 1 in $\text{Ideal}_{\mathcal{R},\mathcal{S}}^{\tilde{\mathcal{D}}_{w,\ell}}$, since Func queries needed to compute the two chains give independent outputs due to $0^d \neq \text{ipad} \oplus \text{opad}$. \blacksquare

The rest of the proof is similar to the one of Theorem 3.2. Specifically we now focus on upper bounding $\Pr \left[\tilde{\mathcal{G}}^{\mathcal{B}} \Rightarrow 1 \right]$ for the adversary \mathcal{B} built from \mathcal{S} as in the lemma statement making q' queries.

As a first step, we consider the game $\tilde{\mathcal{G}}_0$, depicted in Figure 17, which is similar to $\tilde{\mathcal{G}}$, but is potentially easier to win. Concretely, in $\tilde{\mathcal{G}}_0$, we introduce a sub-routine, called FC , which on input i returns the sets of tuples $(y'_0, y'_1, \dots, y'_\ell)$ such that, with respect to RO_2 queries asked so far, define a chain such that $y'_\ell = \mathbf{v}[i]$ and $y'_i = \text{RO}_2(y'_{i-1})$ for $i = 1, \dots, \ell$, and $\mathbf{v}[i]$ for $i \neq j$ is not part of this chain. In particular, if \mathbf{v} has not been defined yet (i.e., phase = 1), then FC always returns the empty set. We also modify the winning condition so that it returns true as long as there is a chain $(y'_0, y'_1, \dots, y'_\ell) \in \text{FC}(j)$ (i.e., it does not need to be the one starting at the value y_0 output by \mathcal{B}). Then, clearly

$$\Pr \left[\tilde{\mathcal{G}}^{\mathcal{B}} \Rightarrow \text{true} \right] \leq \Pr \left[\tilde{\mathcal{G}}_0^{\mathcal{B}} \Rightarrow \text{true} \right] ,$$

as \mathcal{B} winning in $\tilde{\mathcal{G}}$ implies \mathcal{B} winning $\tilde{\mathcal{G}}_0$ because of the fact that \mathcal{B} has asked all queries corresponding to the chain starting at y_0 . We continue with Game $\tilde{\mathcal{G}}_1$, which is equivalent to $\tilde{\mathcal{G}}_0$. The only modification is purely syntactical: The game starts by first choosing all components of \mathbf{u} uniformly at random. It then compute the chain x_0, x_1, \dots only up to $x_{\ell-1}$, and then sets $\text{R}_1[x_{\ell-1}]$ to equal $\mathbf{u}[j]$ if it is undefined, whereas otherwise it overwrites $\mathbf{u}[j]$ as $\mathbf{u}[j] \leftarrow \text{R}_1[x_{\ell-1}]$, and sets the flag bad . In both cases, x_ℓ equals $\mathbf{u}[j]$. The next game, Game $\tilde{\mathcal{G}}_2$, is derived from $\tilde{\mathcal{G}}_1$ by modifying the latter case so that the value $\mathbf{u}[j]$ is *not* overwritten. Clearly, $\tilde{\mathcal{G}}_1$ and $\tilde{\mathcal{G}}_2$ are equivalent until bad , and therefore $\mathbf{Adv}(\tilde{\mathcal{G}}_1^{\mathcal{B}}, \tilde{\mathcal{G}}_2^{\mathcal{B}}) \leq \Pr \left[\tilde{\mathcal{G}}_1^{\mathcal{B}} \text{ sets bad} \right] \leq \sum_{i=0}^{\ell-2} (i+1) \cdot 2^{-n} \leq \frac{\ell^2}{2^{n+1}}$, as conditioned on x_0, \dots, x_i being distinct, each output $\text{RO}_1(x_i)$ for $i \in [0.. \ell-2]$ is in $\{x_0, \dots, x_i\}$ with probability $(i+1)/2^n$.

Note that in $\tilde{\mathcal{G}}_2$, the input \mathbf{u} to \mathcal{B} is random and independent of everything else, as long as \mathcal{B} does not query $x_{\ell-1}$ to RO_1 . We transition from $\tilde{\mathcal{G}}_2$ into new games $\tilde{\mathcal{G}}_3$ and $\tilde{\mathcal{G}}_4$. In $\tilde{\mathcal{G}}_3$, the game sets bad if the query $\text{RO}_1(x_{\ell-1})$ is made before \mathcal{B} commits to \mathbf{v} . Clearly, $\mathbf{Adv}(\tilde{\mathcal{G}}_2^{\mathcal{B}}, \tilde{\mathcal{G}}_3^{\mathcal{B}}) = 0$. Additionally, $\tilde{\mathcal{G}}_4$ is modified so that the value $\text{R}_1[x_{\ell-1}]$ is set to equal $\mathbf{u}[j]$ only *after* \mathcal{B} has output \mathbf{v} . Hence, $\tilde{\mathcal{G}}_3$ and $\tilde{\mathcal{G}}_4$ are equivalent until bad , and $\mathbf{Adv}(\tilde{\mathcal{G}}_3^{\mathcal{B}}, \tilde{\mathcal{G}}_4^{\mathcal{B}}) \leq \Pr \left[\tilde{\mathcal{G}}_4^{\mathcal{B}} \text{ sets bad} \right]$. We will prove an upper bound on this probability below, but for now continue with the main sequence of games.

The next game, Game $\tilde{\mathcal{G}}_5$, simply rearranges the contents of Game $\tilde{\mathcal{G}}_4$ for better readability, but is otherwise fully equivalent. In particular, we postpone the computation of the values x_1, \dots, x_ℓ to after \mathcal{B} outputs \mathbf{v} , which clearly does not affect the game. Then, we transition to a game $\tilde{\mathcal{G}}_6$ which takes into account (via a second procedure FC') those chains that have been created before the adversary outputs \mathbf{v} , and after this, sets the condition bad as soon as some new $\text{R}_2[\cdot]$ entry is defined such that $\text{FC}'(j) = \emptyset$ but $\text{FC}(j) \neq \emptyset$ for the chosen j . Finally, $\tilde{\mathcal{G}}_7$ is the same as $\tilde{\mathcal{G}}_6$, but the winning condition checks for $(y'_0, y'_1, \dots, y'_\ell) \in \text{FC}'(j)$. Clearly, $\tilde{\mathcal{G}}_6$ and $\tilde{\mathcal{G}}_7$ are equivalent until bad , and thus $\mathbf{Adv}(\tilde{\mathcal{G}}_6^{\mathcal{B}}, \tilde{\mathcal{G}}_7^{\mathcal{B}}) \leq \Pr \left[\tilde{\mathcal{G}}_7^{\mathcal{B}} \text{ sets bad} \right]$. Finally, combining all transitions, we obtain

$$\Pr \left[\tilde{\mathcal{G}}^{\mathcal{B}} \Rightarrow \text{true} \right] \leq \frac{\ell^2}{2^{n+1}} + \Pr \left[\tilde{\mathcal{G}}_4^{\mathcal{B}} \text{ sets bad} \right] + \Pr \left[\tilde{\mathcal{G}}_7^{\mathcal{B}} \text{ sets bad} \right] + \Pr \left[\tilde{\mathcal{G}}_7^{\mathcal{B}} \Rightarrow \text{true} \right] .$$

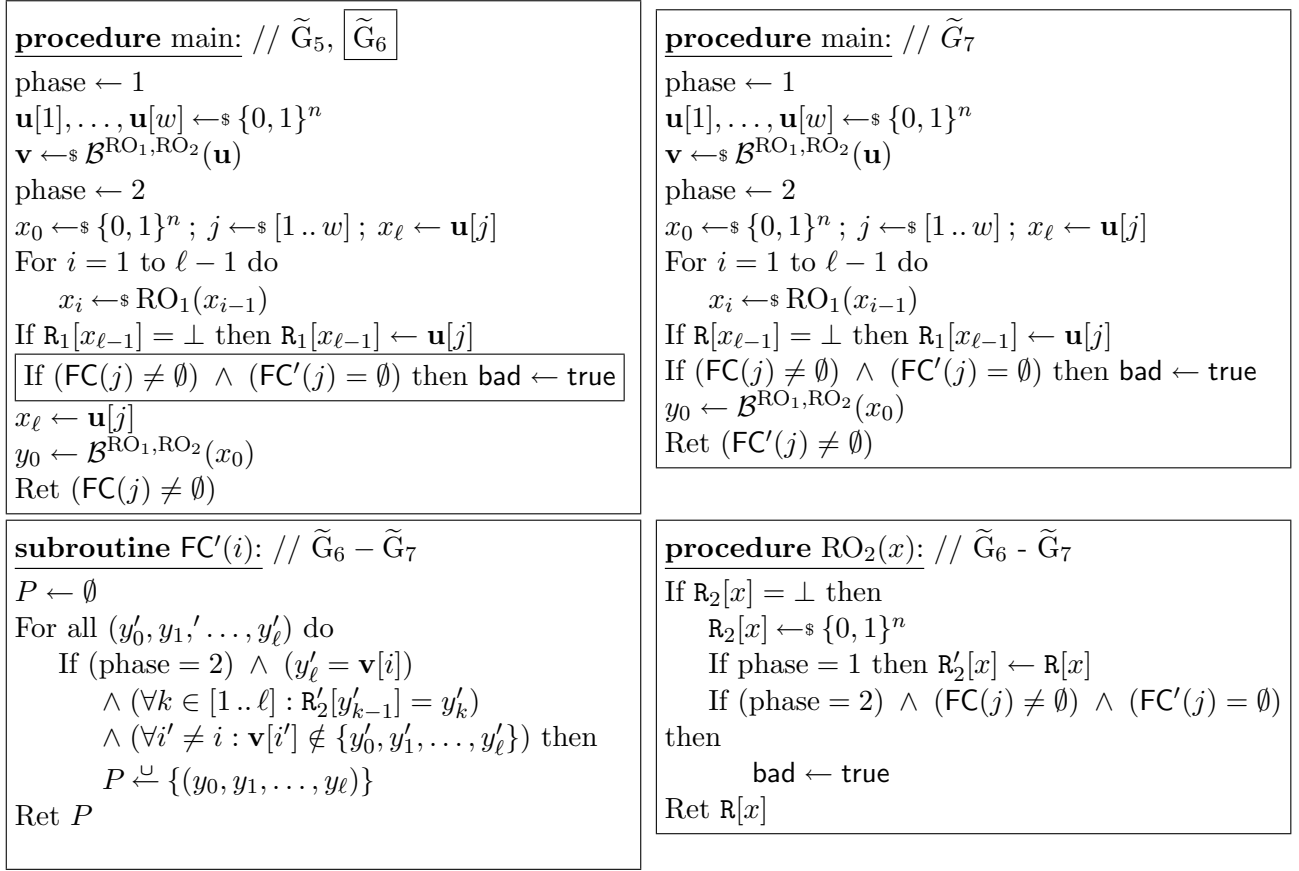


Figure 18: Descriptions of Games $\tilde{G}_5 - \tilde{G}_7$.

To conclude the proof, we need to upper bound the three probabilities on the RHS. This is very similar to the proof of Theorem 3.2, and is omitted.

D Proof of Theorem 3.3

The simulator \mathcal{S} that we use is given by the procedure `OnRightQuery(\cdot)` in game G_0 of Figure 19. The simulator maintains four tables: \mathbf{g} , \mathbf{g}^{-1} , \mathbf{G} and \mathbf{G}^{-1} . These tables, as well as other arrays in subsequent games, are assumed to have their entries initialized to \perp at the start of the game. We note these tables can be interpreted as four separate directed graphs of vertex set $\{0, 1\}^n$, or as a single directed graph of vertex set $\{0, 1\}^n$ with edges of four different colors.

We use sans-serif variable names such as `bad`, `fresh` and `KnownToSim` for boolean values. All boolean values are implicitly initialized to `false`. Notation such as `gProxy[x].fresh` (see e.g. game G_{11}) indicates that each entry in the table `gProxy` holds a bit-value `fresh` besides the n -bit value `gProxy[x]` itself. We emphasize that `fresh` is not an attribute of the *value* `gProxy[x]`, but just an additional bit of data stored at the x -th entry of the table `gProxy[]`. Such bits are also assumed to be initialized to `false` at the start of the game.

It will be convenient to view the internal randomness of each game (to be distinguished from \mathcal{A} 's randomness) as a tape divided into n -bits *blocks*. Thus, the instruction $x \leftarrow_s \{0, 1\}^n$ should be viewed as setting x equal to the next unread block of the random tape.

The boolean value `bad` is called a *flag*. Once set to `true`, `bad` is never reset to `false`. We recall the standard fact that if two games G_i, G_{i+1} differ only for instructions that occur after the flag `bad` is set to `true` (see for example games G_2, G_3) then

$$|\Pr[\mathcal{A}^{G_i} \Rightarrow 1] - \Pr[\mathcal{A}^{G_{i+1}} \Rightarrow 1]| \leq \Pr[\mathcal{A}^{G_j} \text{ sets } \text{bad} \leftarrow \text{true}].$$

for either $j = i$ or $j = i + 1$, where all the probabilities are computed both over \mathcal{A} 's coins and over the random coins used in each game.

The *query history* of an adversary \mathcal{A} after its i -th query has been answered is the sequence $\mathbf{Q}_i = (Q_j)_{j=1}^i$, where each Q_j is a tuple in $\{0, 1\}^n \times \{0, 1\}^n \times \{1, \mathbf{r}\}$. For example, if $Q_j = (x_j, y_j, 1)$, this indicates \mathcal{A} 's j -th query was $x_j \{0, 1\}^n$, asked to its left oracle, and that it received the value $y_j \in \{0, 1\}^n$ in response. Also, \mathcal{A} 's *query sequence* \mathbf{Q}_i^* of first i queries is $((x_j, *_{j}))_{j=1}^i$ if $\mathbf{Q}_i = ((x_j, y_j, *_j))_{j=1}^i$. Namely, the query sequence records which queries were asked, without the answers.

We note that the game G_0 (Figure 19) is equivalent to the oracle pair (RO, \mathcal{S}) and that game G_{23} is equivalent to the oracle pair $(G[g], g)$ for a random g . As written, the simulator given by `OnRightQuery(\cdot)` in game G_0 makes more than $3q_1 + 1$ calls to `RO`, however many of these are redundant, and could be eliminated by having the simulator check its table \mathbf{G} before calling `RO`. It is not difficult to see that once redundant calls are eliminated the simulator never makes more than $3q_1 + 1$ calls to `RO` per query it answers.

We proceed to upper bound \mathcal{A} 's distinguishing advantage $\Delta(G_i, G_{i+1}) := |\Pr[\mathcal{A}^{G_i} \Rightarrow 1] - \Pr[\mathcal{A}^{G_{i+1}} \Rightarrow 1]|$ between games G_i, G_{i+1} for $0 \leq i < 21$.

($G_0 \rightarrow G_1$; Figure 19.) Game G_1 implements the `RO` via lazy sampling, taking place in the subroutine `ROsub`. Moreover, all random sampling is done via the subroutine `Random()`, and the “simulator” (which is still, roughly speaking, the procedure `OnRightQuery()`, though such a distinction will progressively become harder to make) make its queries to `ROsub()` via the subroutine `SimROsub()`. In particular, this means that \mathbf{G}^{-1} is set only for queries made by the simulator, in accordance with game G_0 . All in all the changes from game G_0 to G_1 are syntactical, and so $\Delta(G_0, G_1) = 0$.

Note: Since after game G_1 the simulator no longer “queries” any external oracle, from now on we take the term *query* to exclusively mean “adversarial query”: a query to either `OnLeftQuery(\cdot)` or `OnRightQuery(\cdot)`, made by \mathcal{A} .

($G_1 \rightarrow G_2$; Figures 19 & 20.) Game G_2 adds a number of “bells and whistles” that have no effect except to result in many unused calls to `Random()`. Moreover, G_2 maintains a set X containing all values queried by the adversary, as well as all values sampled in `Random()`. The number of calls to `Random()` per query is normalized via the variable `NumCallsToRandom`, which forces `Random()` to be called exactly twice for every left oracle query and $4q_1 + 3$ times for every right oracle query. (Indeed, it is possible to check for G_2 —as well as for all other games—that `Random()` is called at most twice per left oracle query and at most $4q_1 + 3$ times per right oracle query. The latter maximum is achieved in game G_{11} .) Besides wasting a portion of the random tape these changes have no effect, and so $\Delta(G_1, G_2) = 0$.

Note: For the remaining games, it will be convenient to view the randomness used in `BuildXprime()` as coming from a second, independent random tape; that is, we dedicate a “primary” random tape for the subroutine `Random()`, and use a “secondary” random tape for sampling done by `BuildXprime()`. For every right oracle query there is an associated sequence of $4q_1 + 3$ blocks used for that call on the primary random tape; we call such a sequence a *chunk* (of the primary random tape). We note that chunks are not necessarily contiguous on the primary random tape, given the presence of left oracle

queries and that, moreover, their placement only becomes known at runtime, as the adversary makes queries.

(G₂ → G₃; Figure 20.) Since games G₂ and G₃ are identical up to `bad ← true`, it suffices to upper bound the probability of the latter event. Since $X \cup X'$ has size $(4q_1 + 3)q_2 + 2q_1$, and since `Random()` is called $4q_1 + 3$ times per query to `OnRightQuery()` which is itself called (at most) q_2 times by the adversary, the chance `bad ← true` is triggered by one of the calls to `OnRightQuery()` is at most $((4q_1 + 3)q_2 + 2q_1)(4q_1 + 3)q_2/N$. Likewise, the chance of triggering `bad ← true` by a call to `OnLeftQuery()` is at most $((4q_1 + 3)q_2 + 2q_1)2q_1/N$, since at most q_1 such calls are made, each giving rise to 2 calls to `Random()`. Adding these two bounds, we the chance of `bad ← true` is at most $((4q_1 + 3)q_2 + 2q_1)^2/N$. Thus $\Delta(G_2, G_3) \leq ((4q_1 + 3)q_2 + 2q_1)^2/N$.

(G₃ → G₄; Figures 20 & 21.) Game G₄ introduces a set Y that keeps track of values queried by the adversary and returned to the adversary. (These are, essentially, all the values “known to the adversary”.) This and other minor changes (such as the insertion of a new `bad` flag) do not affect the execution, and thus $\Delta(G_3, G_4) = 0$.

(G₄ → G₅; Figure 21.) We argue that the probability of `bad ← true` in game G₅ is upper bounded by $((4q_1 + 3)q_2 + 2q_1)(q_1 + q_2)/(N - 2q_2 - 2q_1)$. For this, we will show that when \mathcal{A} makes its $(i + 1)$ -th query x_{i+1} to either of its oracles in game either G₄ or G₅ (but assuming the game has not aborted yet), it has no knowledge of the values in the set $X \setminus Y$, in the sense that, conditioned on its query history $\mathbf{Q}_i = ((x_j, y_j, *j))_{j=1}^i$ so far, the set $X \setminus Y$ is equidistributed over all sets of size $(4q_1 + 3)q_{q_2} + 2q_{q_1} - |Y|$ in the complement of Y , where q_{q_1} is the number of left oracle queries and q_{q_2} is the number of right oracle queries among the adversary’s first i queries, with randomness computed with respect to all possible random tapes compatible with \mathbf{Q}_i . To argue this, let $X_{r,s}$ denote the value of the set X when game G₅ is run with primary and secondary random tapes r, s on the query sequence \mathbf{Q}_i^* , assuming that r and s are compatible with \mathbf{Q}_i . Fix such tapes r, s (thus fixing $X_{r,s}$), and let T be an arbitrary subset of $\{0, 1\}^n \setminus Y$ of size $(4q_1 + 3)q_{q_2} + 2q_{q_1} - |Y|$ (where $Y = \{x_1, \dots, y_i\}$). Let $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a permutation taking $X_{r,s} \setminus Y$ to T and fixing points in Y . Letting $\pi(r), \pi(s)$ denote the application blockwise of π to the random tapes r, s , it is easy to check that $X_{\pi(r), \pi(s)} = T \cup Y$, i.e. $X_{\pi(r), \pi(s)} \setminus Y = T$. It follows that $X \setminus Y$ is equidistributed among all sets of size $(4q_1 + 3)q_{q_2} + 2q_{q_1} - |Y|$ in $\{0, 1\}^n \setminus Y$, conditioned on \mathbf{Q}_i . Thus, the probability that `bad ← true` at the i -th query is at most $((4q_1 + 3)q_{q_2} + 2q_{q_1} - |Y|)/(N - |Y|) \leq ((4q_1 + 3)q_2 + 2q_1)/(N - 2q_2 - 2q_1)$. Union bounding over all $q_1 + q_2$ queries made by the adversary gives $\Delta(G_4, G_5) \leq ((4q_1 + 3)q_2 + 2q_1)(q_1 + q_2)/(N - 2q_2 - 2q_1)$.

Note: If `Abort` does not occur in game G₅, then the graphs \mathbf{G} , \mathbf{G}^{-1} and \mathbf{g} can be shown to be fairly structured. For example, it is easy to check that after each query is answered (subject to non-abortion) the following three invariants are maintained in G₅:

- (1) No cycles or collisions (every vertex has indegree¹⁰ at most 1) occur in the directed graph defined by \mathbf{G} , or in the directed graphs defined by \mathbf{G}^{-1} and \mathbf{g} ;
- (2) As the graph \mathbf{G} grows, no edge is ever added whose endpoint is already adjacent to another edge of $\mathbf{G} \cup \mathbf{G}^{-1} \cup \mathbf{g}$;
- (3) \mathbf{G}^{-1} is a reversed subgraph of \mathbf{G} (every edge of \mathbf{G}^{-1} , reversed, becomes an edge of \mathbf{G}).

To discuss further features of the graphs, we make some additional definitions. Let $\bar{\mathbf{G}}$ be the restriction of \mathbf{G} to edges whose reversal is also in \mathbf{G} (these are the edges “known to the simulator”). The edges in

¹⁰The outdegree is automatically at most 1, since these graphs represent functions.

$\mathbb{G} \setminus \bar{\mathbb{G}}$ are called *hidden*. A *ladder* is configuration consisting of an ordered pair of two maximal paths $(s_{-q_1}, \dots, s_0, \dots, s_{m'})$ and $(t_{-k}, \dots, t_0, \dots, t_m)$ in $\bar{\mathbb{G}}$, such that: (i) the paths are vertex disjoint; (ii) $0 \leq k \leq q_1$, $m \geq q_1 + 1$ and $m' \in \{m - 1, m\}$; (iii) (t_i, s_i) , $-k \leq i \leq m'$, and (s_i, t_{i+1}) , $-k \leq i \leq m - 1$, are edges of \mathbb{g} , and these are the only edges of \mathbb{g} adjacent to a vertex in $\{s_{-q_1}, \dots, s_{m'}, t_{-k}, \dots, t_m\}$.

The path $(s_{-q_1}, \dots, s_{m'})$ is called the ladder's *upper side* and the path (t_{-k}, \dots, t_m) is called the ladder's *lower side*¹¹. The vertex t_{-k} is called the *anchor* of a ladder, and t_0 is called the *center*. The edges (t_i, s_i) , (s_i, t_{i+1}) , $-k \leq i < q_1$, are called the *rungs* of the ladder. A rung of the form (t_i, s_i) is *upward*; a rung of the form (s_i, t_{i+1}) is *downward*. The *tip* of a ladder is the last vertex on the unique \mathbb{g} -path starting at the anchor. (The tip is either t_m or $s_{m'}$ depending on whether $m > m'$ or $m = m'$.) The *tip precursor* is the next-to-last vertex on this path.

A maximal path in $\mathbb{G} \setminus \bar{\mathbb{G}}$ ending at a ladder's anchor is called an *anchored hidden chain*. A maximal path in $\mathbb{G} \setminus \bar{\mathbb{G}}$ not adjacent to any edges in $\bar{\mathbb{G}} \cup \mathbb{g}$ is called an *isolated hidden chain*. We make these definitions only for games G_5 and G_6 (afterwards, we will update the definitions as well as the stated properties). For games G_5 , G_6 one can check that the following facts hold after each query is answered, presuming the game is in a non-aborted state:

- (4) Any two ladders are vertex- (and hence also edge-) disjoint;
- (5) Every edge of $\mathbb{G} \cup \mathbb{g}$ is either in a ladder, in a hidden chain anchored to a ladder, or in an isolated hidden chain; moreover (trivially) the endpoints of edges in $\mathbb{G} \setminus \bar{\mathbb{G}}$ are in the set Y ;
- (6) The set of vertices in a ladder that are in the set Y forms a connected component in $\mathbb{G} \cup \mathbb{g}$ that contains the anchor of the ladder; moreover, this set does not contain any upper side vertices non-adjacent to a rung, and does not contain the ladder's tip.

Giving detailed proofs of (4), (5) and (6) for G_5 and G_6 is not difficult, but would take us afield. We note that the (high-level) reason why Y never contains a ladder's tip is that the simulator keeps elongating each ladder such that every ladder's tip is at least distance $q_1 + 1$ in $\mathbb{G} \cup \mathbb{g}$ from any point in the ladder that \mathcal{A} has queried to the simulator, coupled with the fact that \mathcal{A} has only q_1 queries to the RO.

We make a few more definitions. The length ℓ of the longest path in \mathbb{G} ending at the center t_0 of a ladder is called the *extension length* of the ladder; we note the path of length ℓ ending at t_0 is unique, and that its head is either the ladder's anchor or else the head of the (unique) hidden chain ending at the ladder's anchor. It is easy to check from (5), (6) and from the fact that the adversary makes at most q_1 queries that the extension length is at most q_1 . Finally, if $(s_{-q_1}, \dots, s_{q_1})$ is the upper side of a ladder of extension length ℓ , then we call the path $(s_{-q_1}, \dots, s_{-\ell})$ the *excess upper path*.

As hinted above, since successive games keep tweaking the data structures, an unavoidable annoyance is that some of the above definitions will become outdated (see, for example, the comment after the transition $G_6 \rightarrow G_7$). We have found no other way, unfortunately, but to update each time those definitions that are still required for the discussion of subsequent games.

We also point out that games G_5 – G_{20} have the following well-behavedness property: *except for the abort event*, the execution path taken during one query answer cycle does not depend on the internal randomness; that is, knowing the state $\mathbb{G}^{\pm 1}$ and \mathbb{g} at the point when a query is asked is enough to know which lines of code will be executed by the game in what order until it returns, assuming Abort does not occur. This property will be useful for “randomness cut-and-paste” arguments used below.

¹¹We note that since ladders are defined as an ordered pair of paths, there is no ambiguity as to which is the upper or lower side; nonetheless, even if a ladder is only given as the subgraph in $\mathbb{G} \cup \mathbb{g}$ induced by the path vertices, one can tell the upper from the lower side from the fact that the lower side contains the head of the \mathbb{g} -path formed by the ladder's rungs.

(G₅ → G₆; Figures 21 & 22.) The unique change in game G₆ is that the line “If $y = \perp$ then Abort” is removed from the procedure OnRightQuery(\cdot). We argue that $\mathbf{G}^{-i}[\mathbf{g}[x_i]]$ at the previous line never returns \perp , so that this change has no effect. Indeed, given property (5) above, the edge $(x_i, \mathbf{g}[x_i])$ must be either an upward or a downward rung of some ladder. If it is a downward rung and $\mathbf{G}^{-i}[\mathbf{g}[x_i]] = \perp$ then $i > 0$ and $x = x_0$ must be an upper side vertex nonadjacent to any downward rung, a contradiction to the fact that Y contains no such vertices and to the abortion condition added in game G₅. If $(x_i, \mathbf{g}[x_i])$ is an upward rung, then $\mathbf{G}^{-i}[\mathbf{g}[x_i]] \neq \perp$ follows from the fact that any ladder’s extension length is at most q_1 . Thus $\Delta(\mathbf{G}_5, \mathbf{G}_6) = 0$.

(G₆ → G₇; Figures 22 & 23.) In game G₇, the call SetTable($\mathbf{G}, x, \text{ROsub}$) is moved from SimROsub() to ROsub(), and SimROsub() is bypassed altogether in favor of ROsub(). The effect of this change is that \mathbf{G}^{-1} becomes the exact reverse graph of \mathbf{G} , instead of being a subgraph of the reversal of \mathbf{G} . However, this enlargement of \mathbf{G}^{-1} has no effect, since no lookups in \mathbf{G}^{-1} (or its iteration) in game G₆ ever give value \perp (since the only use of \mathbf{G}^{-1} occurs in OnRightQuery(), this was in fact argued above). Thus $\Delta(\mathbf{G}_6, \mathbf{G}_7) = 0$.

Definition update. Since \mathbf{G}^{-1} is now the same as \mathbf{G} , $\bar{\mathbf{G}} = \mathbf{G}$ and some definitions must be adjusted. In particular, the lower side of a ladder (t_{-k}, \dots, t_m) is not required to be a maximal path in \mathbf{G} , but only to be “non-extendible at the end” in \mathbf{G} (i.e., that t_m have outdegree 0 in \mathbf{G}). The definitions of upper side, anchor, center, tip and tip precursor, upward and downward rungs are unaffected. We redefine a anchored hidden chain to be a path in \mathbf{G} ending at a ladder’s anchor, that is “non-extendible at the start” (i.e., the path’s head has indegree 0), and that is nonadjacent to any edges of \mathbf{g} except at its endpoint (the ladder’s anchor). Isolated hidden chains are redefined as (maximal) connected components of \mathbf{G} not containing any vertex of a ladder or any vertex adjacent to an edge of \mathbf{g} . Subject to these changes, the above properties (1)–(6) still hold (though the second half of (5) is now void).

(G₇ → G₈; Figures 23 & 24.) In game G₈ MakeLadder() “manually” erases the excess upper path of the ladder being created. This has no effect because these entries of \mathbf{G} and \mathbf{G}^{-1} are never subsequently read. Thus $\Delta(\mathbf{G}_7, \mathbf{G}_8) = 0$.

Definition update. We allow the upper side of a ladder to be shortened. Specifically, the upper and lower sides of a ladder are now paths $(s_{-\ell}, \dots, s_{m'})$ and (t_{-k}, \dots, t_m) in \mathbf{G} such that $0 \leq k \leq \ell$, such that the upper side is a maximal path in \mathbf{G} and t_m has outdegree 0 in \mathbf{G} , such that properties (i), (ii), (iii) above for ladders still hold. Again, one can verify that items (1)–(6) still hold (from G₈ all the way to G₁₆, in fact). We also deprecate the definitions of “center” and “extension length” without updating them. (These will no longer be needed.)

(G₈ → G₉; Figures 24 & 25.) In game G₉ the excess upper path of the ladder is “never created in the first place”. Note this results in fewer calls to Random() from within MakeLadder(). In particular, for example, the call to MakeLadder() has fewer chances of causing the game to abort. Nonetheless, we argue that games G₈ and G₉ are indistinguishable. Note that if one cuts the first $q_1 - \ell$ blocks of (primary tape) randomness used by MakeLadder during a given query and pastes these at the end the same chunk (for the i -th right oracle query, this is the i -th chunk), then, for this one query, and assuming given states of $\mathbf{G}, \mathbf{G}^{-1}$ and \mathbf{g} at the start of the query, running G₈ with the original chunk or else running G₉ with the new chunk produces the exact same result, in the sense that the probability of abortion (computed only over the secondary random tape, and fixing the first) is equal in either world, and that if the game doesn’t abort then the final values of the tables $\mathbf{G}, \mathbf{G}^{-1}, \mathbf{g}$ as well as of the set X are equal in either world at the end of that query (in particular, the blocks that were cut-and-pasted become part of X in either case). From this it is easy to conclude, by induction over

the number of queries, that the two worlds are equidistributed. Thus $\Delta(G_8, G_9) = 0$.

($G_9 \rightarrow G_{10}$; Figures 25 & 26.) The game G_{10} introduces a functionality whereby the value sampled in $\text{ROsub}()$ can be set externally. This functionality is only used by $\text{MakeLadder}()$, which is honestly sampling the second argument to ROsub . Thus the change is syntactical, and $\Delta(G_9, G_{10}) = 0$.

($G_{10} \rightarrow G_{11}$; Figures 26 & 27.) Game G_{11} introduces a new array gProxy that is sampled in $\text{Random}()$, and that is never used anywhere else. The sampling of gProxy in ROsub produces a “translation effect” on the random tape, but one can argue this effect leaves the games equidistributed by using a cut-and-paste argument as in the transition $G_8 \rightarrow G_9$. (We note in passing that the cut-and-paste argument is facilitated by the fact that we know at the outset of a right oracle query exactly which primary random blocks of the current chunk will be assigned to gProxy ; see the last comment before the $G_5 \rightarrow G_6$ transition.) Thus we have $\Delta(G_{10}, G_{11}) = 0$.

($G_{11} \rightarrow G_{12}$; Figures 27 & 28.) In game G_{12} the randomness “stored” in gProxy is accessed by $\text{MakeLadder}()$ instead of using calls to $\text{Random}()$. This requires some justifications. Note that, since every value stored in gProxy is sampled by $\text{Random}()$, and hence is in X , then if one conditions only on knowledge of G , G^{-1} , g and on the fact that the currently queried value x did not cause Abort (via collision with $X \setminus Y$), and one reads a “fresh” (yet-unread) value from gProxy , then this value is distributed randomly at uniform among all values that do not appear in $G^{\pm 1}$, $g^{\pm 1}$, and that are not equal to x —moreover, note the same statement holds if one replaces “reading a fresh value in gProxy ” by “calling $\text{Random}()$, assuming Abort does not occur”. Moreover, the fresh value read from gProxy is also stored in X , just like a value returned by $\text{Random}()$. The two methods of obtaining a random value therefore only differ in that $\text{Random}()$ can cause Abort , whereas reading a fresh value from gProxy cannot. So *conditioned on non-abortion* and on a given state of $G^{\pm 1}$ and g at the point when a query is issued, the answer to the query will be equidistributed in games G_{11} and G_{12} ; moreover, the state of $G^{\pm 1}$ and g will also be equidistributed (allowing to pursue induction on the number of queries). Thus, since Abort anyway occurs with equal probability at each query of each game, $\Delta(G_{11}, G_{12}) = 0$.

Note: Game G_{12} already samples and uses all randomness “in the same order” as the final construction $G[g]$. One consequence of this is that games G_{12} through G_{20} are indistinguishable in the following strong sense: running these games with equal random tapes (primary and secondary) on the same query sequence gives the same exact query answers (or Abort answer). Nonetheless, this must be argued game by game.

($G_{12} \rightarrow G_{13}$; Figures 28 & 29.) Game G_{13} removes the “freshness” and “definedness” checks on gProxy from G_{12} . On the one hand, because a ladder is only created once and ladders are vertex-disjoint, all values $G^{-\ell}[x]$ or $G^{i+1}[x]$ used as indices in $\text{MakeLadder}()$ are distinct, the same entry of gProxy cannot be read by two distinct calls to MakeLadder (or twice during the same call). On the other hand, note a fresh value $\text{gProxy}[u]$ exists as long as a call of the form $\text{ROsub}(u, \perp)$ has been made before $g[u]$ is defined, and before any call of the form $\text{ROsub}(u, y)$, $y \neq \perp$ has been made. For the case of $u = G^{-i}[x]$, $i \geq 1$, such calls have been while answering the adversary’s own queries to the left oracle. For the case of $u = G^i[x]$, $i \geq 0$, such calls in the (necessarily completed) For loop of $\text{OnRightQuery}()$ (and possibly before while answering \mathcal{A} ’s left oracle queries). Hence the gProxy values read by $\text{MakeLadder}()$ are always defined and fresh. Thus, $\Delta(G_{12}, G_{13}) = 0$.

($G_{13} \rightarrow G_{14}$; Figures 29 & 30.) In G_{14} the line $\text{gProxy}[\text{gProxy}[x]] \leftarrow z$ is added to the If -block of $\text{ROsub}()$. Because we have argued above that $\text{MakeLadder}()$ never reads a \perp value from gProxy , it is sufficient to argue that $\text{gProxy}[\text{gProxy}[x]]$ is always \perp right before $\text{gProxy}[\text{gProxy}[x]]$ is set in

$\text{ROsub}()$ in order to argue that the change in $\text{ROsub}()$ has no effect. (In particular, note that it is not our concern whether $\mathbf{gProxy}[\mathbf{gProxy}[x]]$ is overwritten by the line $\mathbf{gProxy}[x] \leftarrow \mathbf{g}[x]$ during a later call to $\text{ROsub}()$ with argument $\mathbf{gProxy}[x]$; we are concerned with $\mathbf{gProxy}[\mathbf{gProxy}[x]]$ being *read*, and influencing the execution, not with it being overwritten.) But since $\mathbf{gProxy}[x]$ has just been returned by $\text{Random}()$, and since $\text{ROsub}()$ is always called with a first argument that is in X (and since $\text{Random}()$ samples outside of X), it is clear that $\mathbf{gProxy}[\mathbf{gProxy}[x]]$ is \perp right before being assigned z . Hence the change to $\text{ROsub}()$ has no effect and $\Delta(\mathbf{G}_{13}, \mathbf{G}_{14}) = 0$.

($\mathbf{G}_{14} \rightarrow \mathbf{G}_{15}$; Figures 30 & 31.) In game \mathbf{G}_{15} the function $\text{ROsub}()$ is rewritten, and takes (again) a single argument. We argue the changes have no effect. More precisely, we argue by induction on the number of calls to $\text{ROsub}()$ that for the same query sequence $\mathbf{Q}_{q_2}^*$ and the same random coins (primary and secondary), the contents of the arrays \mathbf{G} , \mathbf{g} and \mathbf{gProxy} are the exact same after each call to $\text{ROsub}()$ in \mathbf{G}_{14} and \mathbf{G}_{15} . (Obviously, thus, the sequence of query answers will also be the same.) Moreover we also argue by induction on the number of calls to $\text{ROsub}()$ that, for both games \mathbf{G}_{14} and \mathbf{G}_{15} , \mathbf{gProxy} always contains \mathbf{g} , in the sense that $\mathbf{g}[u] \neq \perp \implies \mathbf{gProxy}[u] = \mathbf{g}[u]$ for all $u \in \{0, 1\}^n$, with this invariant being true, in fact, after each line of execution (internal or external to $\text{ROsub}()$, in either game). We note \mathbf{gProxy} obviously contains \mathbf{g} at the start of each game, when both arrays are empty.

Assuming the latter claim by induction, the If-block containing the instruction $\mathbf{gProxy}[x] \leftarrow \mathbf{g}[x]$ obviously has no effect in \mathbf{G}_{14} , so we may ignore it¹². We note that at the point when $\text{MakeLadder}()$ is called on a point x_0 , and assuming ℓ as defined in $\text{MakeLadder}()$, there is a maximal path $(x_{-\ell}, \dots, x_0, \dots, x_{q_1+1})$ existing in \mathbf{G} such that \mathbf{g} is defined at none of the x_i 's, such that $s_i := \mathbf{gProxy}[x_i]$ is defined for $-\ell \leq i \leq q_1$, and such that $\mathbf{gProxy}[s_i] = x_{i+1}$ also for $-\ell \leq i \leq q_1$; the ladder's upper side becomes $(s_{-\ell}, \dots, s_{q_1})$ and \mathbf{gProxy} is unchanged when $\text{MakeLadder}()$ returns, so that \mathbf{gProxy} , in particular, contains all the values of \mathbf{g} set by $\text{FillInRungs}()$ at the end of the call to $\text{MakeLadder}()$. Moreover, it is easy to check that future calls to $\text{FillInRungs}()$ for the same ladder maintain the invariant that \mathbf{g} is a subtable of \mathbf{gProxy} , given that any call to $\text{ROsub}(x)$ or $\text{ROsub}(x, \perp)$ at a point x such that $\mathbf{g}[x]$ is already defined results in setting $\mathbf{gProxy}[\mathbf{g}[x]] = \mathbf{gProxy}[\mathbf{gProxy}[x]]$ to the call's value. Thus, \mathbf{g} remains throughout a subtable of \mathbf{gProxy} .

Note that if $\mathbf{G}[x]$ is defined then calls to $\text{ROsub}()$ with first argument x are obviously equivalent in \mathbf{G}_{14} and \mathbf{G}_{15} . It therefore suffices to consider the cases when $\mathbf{G}[x]$ is undefined.

For calls $\text{ROsub}(x, z)$ with $z = \perp$ (and $\mathbf{G}[x] = \perp$) in game \mathbf{G}_{14} , these are obviously equivalent to calling $\text{ROsub}(x)$ in \mathbf{G}_{15} if $\mathbf{gProxy}[x]$ is undefined; if $y := \mathbf{gProxy}[x]$ is defined then the calls are still equivalent as long as $\mathbf{gProxy}[y]$ is undefined. However it is easy to check that the only calls for which $\mathbf{G}[x] = \perp$ and $\mathbf{gProxy}[x]$, $\mathbf{gProxy}[y]$ are both defined are the calls made via $\text{MakeLadder}()$, for which $z \neq \perp$. This establishes that $\text{ROsub}()$ has the same effect (on \mathbf{G} , \mathbf{gProxy} and \mathbf{g}) in \mathbf{G}_{14} , \mathbf{G}_{15} for all calls to $\text{ROsub}()$ made outside $\text{MakeLadder}()$. For the calls to $\text{ROsub}(x)/\text{ROsub}(x, z)$ made from within $\text{MakeLadder}()$, moreover, it is easy to see that $\mathbf{gProxy}[x]$, $\mathbf{gProxy}[y] = \mathbf{gProxy}[\mathbf{gProxy}[x]]$ are defined beforehand, and that the result of such a call is just to set $\mathbf{G}[x] = z$ where $z = \mathbf{gProxy}[y]$. Hence, all calls to $\text{ROsub}()$ are equivalent in either world, and $\Delta(\mathbf{G}_{14}, \mathbf{G}_{15}) = 0$.

($\mathbf{G}_{15} \rightarrow \mathbf{G}_{16}$; Figures 31 & 32.) \mathbf{G}_{16} replaces the array \mathbf{g} by \mathbf{gProxy} ; the only difference between the two arrays in game \mathbf{G}_{15} is that \mathbf{gProxy} can be defined on more points than \mathbf{g} (as established),

¹²An astute reader may note that this If-block is never used in the analysis of any transition, from the point in game \mathbf{G}_{11} where it appears. In fact, this If-block is superfluous, strictly speaking, but is kept for esthetical reasons. More precisely, it is kept so that $\mathbf{gProxy}[u]$ and $\mathbf{g}[u]$ never contain different non-null values in games \mathbf{G}_{11} – \mathbf{G}_{13} ; without the If-block such inconsistencies could (and would) occur when a ladder is “extended” (via calls to $\text{FillInRungs}()$ made from within $\text{OnRightQuery}()$), but without effect, because not affecting entries of \mathbf{gProxy} subsequently read by $\text{MakeLadder}()$. In \mathbf{G}_{14} , the If-block has no effect at all.

and this matters when $\mathbf{g}[x_i]$ is tested for a value in `OnRightQuery()`. Game G_{16} circumvents this by introducing boolean field `KnownToSim`, which is set to `true` if and only if the corresponding entry of the “old” table \mathbf{g} (in game G_{15}) is set to a non- \perp value. Moreover G_{16} introduces two new calls to `SetTable()` in `ROsub()`, but since the array \mathbf{g}^{-1} is not used these have no effect. Altogether, therefore, $\Delta(G_{15}, G_{16}) = 0$.

($G_{16} \rightarrow G_{17}$; Figures 32 & 33.) In game G_{17} , \mathbf{G} is replaced everywhere by \mathbf{g}^2 and \mathbf{G}^{-1} is replaced everywhere by \mathbf{g}^{-2} . Since it is straightforward to check that, in game G_{16} , $\mathbf{G}[x] = \mathbf{g}^2[x]$ for all x where $\mathbf{G}[x] \neq \perp$, and also $\mathbf{G}^{-1}[x] = \mathbf{g}^{-2}[x]$ for all x where $\mathbf{G}^{-1}[x] \neq \perp$, this change has no effect. Thus $\Delta(G_{16}, G_{17}) = 0$.

($G_{17} \rightarrow G_{18}$; Figures 33 & 34.) Game G_{18} replaces the call to `MakeLadder(x)` with simply `FillInRungs(x)`. Indeed it is easy to check that in game G_{17} the only part of `MakeLadder()` with any effect is the call to `FillInRungs()`. Moreover calls to `SetTable($\mathbf{g}, *, *$)` are dropped in G_{18} since \mathbf{g}^{-1} is no longer used there. Thus the changes have no effect and $\Delta(G_{17}, G_{18}) = 0$.

($G_{18} \rightarrow G_{19}$; Figures 34 & 35.) In game G_{19} calls to `FillInRungs()` have been “folded back” into the `For` loop of `OnRightQuery()`, to equivalent effect. Indeed it is just a matter of case checking to see that that G_{19} produces the same result as G_{18} , whether `FillInRungs()` is called inside or outside of the `For` loop in G_{18} . Thus $\Delta(G_{18}, G_{19}) = 0$.

($G_{19} \rightarrow G_{20}$; Figure 35.) In game G_{20} all `Abort` conditions are dropped. Arguing as in transitions $G_2 \rightarrow G_3$ and $G_4 \rightarrow G_5$ (and using a union bound), one can show that probability of `bad` \leftarrow `true` in G_{19} is at most $\Delta(G_2, G_3) + \Delta(G_4, G_5)$. Thus $\Delta(G_{19}, G_{20}) \leq ((4q_1 + 3)q_2 + 2q_1)^2/N + ((4q_1 + 3)q_2 + 2q_1)(q_1 + q_2)/(N - 2q_2 - 2q_1)$.

($G_{20} \rightarrow G_{21}$; Figures 35 & 36.) Game G_{21} clears up clutter left over by the defunct `Abort` conditions, and removes all references to `KnownToSim`, which was unused (because never tested for) in games G_{19} and G_{20} , and therefore superfluous. The changes have no import and $\Delta(G_{20}, G_{21}) = 0$.

($G_{21} \rightarrow G_{22}$; Figure 36.) Note that G_{21} always sets each (non-defined) entry of \mathbf{g} to a lazy-sampled value, never overwrites entries of \mathbf{g} , and always returns $\mathbf{g}[x]$ to a query `OnRightQuery(x)` and $\mathbf{g}[\mathbf{g}[x]]$ to a query `OnLeftQuery(x)`, like game G_{22} . Thus, the equivalence of games G_{21} can be seen by using *early sampling* for \mathbf{g} , instead of lazy sampling, in which case both games are obviously equivalent. Thus $\Delta(G_{21}, G_{22}) = 0$.

Finally, summing the Δ -values of transitions $G_0 \rightarrow G_1$ through $G_{21} \rightarrow G_{22}$, we find $\Delta(G_2, G_3) + \Delta(G_4, G_5) + \Delta(G_{19}, G_{20}) = 2\Delta(G_2, G_3) + 2\Delta(G_4, G_5)$, which is the bound advertised by the theorem.

<p>procedure <u>OnLeftQuery(x):</u> G_0 Ret RO(x)</p> <p>procedure <u>OnRightQuery(x):</u> $x_0 \leftarrow x$ For $i = 0$ to q_1 If $g[x_i] \neq \perp$ then $y \leftarrow G^{-i}[g[x_i]]$ If $y = \perp$ then Abort FillInRungs(x, y) Ret y SetTable($G, x_i, RO(x_i)$) $x_{i+1} \leftarrow G[x_i]$ MakeLadder(x) Ret $g[x]$</p> <p>subroutine <u>FillInRungs(x, y)</u> $x_0 \leftarrow x, x_1 \leftarrow y$ SetTable(g, x_0, x_1) For $i = 1$ to $2q_1 + 1$ SetTable($G, x_{i-1}, RO(x_{i-1})$) $x_{i+1} \leftarrow G[x_{i-1}]$ SetTable(g, x_i, x_{i+1})</p> <p>subroutine <u>MakeLadder(x)</u> $s_{-q_1} \leftarrow_s \{0, 1\}^n$ For $i = -q_1$ to $q_1 - 1$ SetTable($G, s_i, RO(s_i)$) $s_{i+1} \leftarrow G[s_i]$ FillInRungs(x, s_0)</p> <p>subroutine <u>SetTable(T, x, y)</u> If $T[x]$ and $T[x] \neq y$ then Ret $T[x] = y$ $T^{-1}[y] = x$</p>	<p>procedure <u>OnLeftQuery(x):</u> G_1 Ret ROsub(x)</p> <p>procedure <u>OnRightQuery(x):</u> $x_0 \leftarrow x$ For $i = 0$ to q_1 If $g[x_i] \neq \perp$ then $y \leftarrow G^{-i}[g[x_i]]$ If $y = \perp$ then Abort FillInRungs(x, y) Ret y $x_{i+1} \leftarrow \text{SimROsub}(x_i)$ MakeLadder(x) Ret $g[x]$</p> <p>subroutine <u>FillInRungs(x, y)</u> $x_0 \leftarrow x, x_1 \leftarrow y$ SetTable(g, x_0, x_1) For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow \text{SimROsub}(x_{i-1})$ SetTable(g, x_i, x_{i+1})</p> <p>subroutine <u>MakeLadder(x)</u> $s_{-q_1} \leftarrow \text{Random}()$ For $i = -q_1$ to $q_1 - 1$ $s_{i+1} \leftarrow \text{SimROsub}(s_i)$ FillInRungs(x, s_0)</p> <p>subroutine <u>ROsub(x)</u> If $G[x]$ Ret $G[x]$ Ret $G[x] \leftarrow \text{Random}()$</p> <p>subroutine <u>SimROsub(x)</u> SetTable($G, x, ROsub(x)$) Ret $G[x]$</p> <p>subroutine <u>Random()</u> Ret $\leftarrow_s \{0, 1\}^n$</p> <p>subroutine <u>SetTable(T, x, y)</u> If $T[x]$ and $T[x] \neq y$ then Ret $T[x] = y$ $T^{-1}[y] = x$</p>
--	--

Figure 19: The first two games used in the proof of Theorem 3.3. The adversary’s left and right oracles are implemented by the procedures OnLeftQuery(\cdot) and OnRightQuery(\cdot). Game G_0 is the “Ideal World”, where the left oracle is implemented by a random oracle $RO : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and the right oracle is the simulator \mathcal{S} , whose goal is to mimic a function $g : \{0, 1\}^n \rightarrow \{0, 1\}^n$ such that $RO = g^2$.

<pre> procedure OnLeftQuery(x): G₂, G₃ NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow 2 $X \leftarrow X \cup \{x\}$ Finalization() Ret ROsub(x) procedure OnRightQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow $4q_1 + 3$ $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $g[x_i] \neq \perp$ then $y \leftarrow G^{-i}[g[x_i]]$ If $y = \perp$ then Abort FillInRungs(x, y) Break // (For loop) $x_{i+1} \leftarrow$ SimROsub(x_i) If $i > q_1$ MakeLadder(x) Finalization() Ret $g[x]$ subroutine FillInRungs(x, y) $x_0 \leftarrow x, x_1 \leftarrow y$ SetTable(g, x_0, x_1) For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow$ SimROsub(x_{i-1}) SetTable(g, x_i, x_{i+1}) subroutine MakeLadder(x) $s_{-q_1} \leftarrow$ Random() For $i = -q_1$ to $q_1 - 1$ $s_{i+1} \leftarrow$ SimROsub(s_i) FillInRungs(x, s_0) </pre>	<pre> subroutine ROsub(x) G₂, G₃ (cont.) If $G[x]$ Ret $G[x]$ Ret $G[x] \leftarrow$ Random() subroutine SimROsub(x) SetTable($G, x, ROsub(x)$) Ret $G[x]$ subroutine Random() $y \leftarrow_s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ bad \leftarrow true Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom ++ Ret y subroutine SetTable(T, x, y) If $T[x]$ and $T[x] \neq y$ then Ret $T[x] = y$ $T^{-1}[y] = x$ subroutine BuildXprime() $X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow_s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$ subroutine Finalization() While NumCallsToRandom $<$ MAX_CALLS do Random() </pre>
---	---

Figure 20: Games G_2 and G_3 for the proof of Theorem 3.3. Game G_3 includes the boxed statement, game G_2 does not.

<pre> procedure OnLeftQuery(x): $G_4, \boxed{G_5}$ NumCallsToRandom $\leftarrow 0$; MAX_CALLS $\leftarrow 2$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then bad \leftarrow true Abort $X \leftarrow X \cup \{x\}$ ROsub(x) $Y \leftarrow Y \cup \{x, G[x]\}$ Finalization() Ret $G[x]$ procedure OnRightQuery(x): NumCallsToRandom $\leftarrow 0$; MAX_CALLS $\leftarrow 4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then bad \leftarrow true Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $g[x_i] \neq \perp$ then $y \leftarrow G^{-i}[g[x_i]]$ If $y = \perp$ then Abort FillInRungs(x, y) Break // (For loop) $x_{i+1} \leftarrow$ SimROsub(x_i) If $i > q_1$ MakeLadder(x) $Y \leftarrow Y \cup \{x, g[x]\}$ Finalization() Ret $g[x]$ subroutine FillInRungs(x, y) $x_0 \leftarrow x, x_1 \leftarrow y$ SetTable(g, x_0, x_1) For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow$ SimROsub(x_{i-1}) SetTable(g, x_i, x_{i+1}) </pre>	<pre> subroutine MakeLadder(x) $G_4, \boxed{G_5}$ $s_{-q_1} \leftarrow$ Random() For $i = -q_1$ to $q_1 - 1$ $s_{i+1} \leftarrow$ SimROsub(s_i) FillInRungs(x, s_0) subroutine ROsub(x) If $G[x]$ Ret $G[x]$ Ret $G[x] \leftarrow$ Random() subroutine SimROsub(x) SetTable($G, x, ROsub(x)$) Ret $G[x]$ subroutine Random() $y \leftarrow_s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ then Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom ++ Ret y subroutine SetTable(T, x, y) If $T[x]$ and $T[x] \neq y$ then Ret $T[x] = y$ $T^{-1}[y] = x$ subroutine BuildXprime() $X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow_s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$ subroutine Finalization() While NumCallsToRandom $<$ MAX_CALLS do Random() </pre>
--	---

Figure 21: Games G_4 and G_5 for the proof of Theorem 3.3.

<pre> procedure OnLeftQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow 2 BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ ROsub(x) $Y \leftarrow Y \cup \{x, \mathbf{g}[x]\}$ Finalization() Ret $\mathbf{g}[x]$ procedure OnRightQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow $4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $\mathbf{g}[x_i] \neq \perp$ then $y \leftarrow \mathbf{G}^{-i}[\mathbf{g}[x_i]]$ FillInRungs(x, y) Break // (For loop) $x_{i+1} \leftarrow$ ROsub(x_i) If $i > q_1$ MakeLadder(x) $Y \leftarrow Y \cup \{x, \mathbf{g}[x]\}$ Finalization() Ret $\mathbf{g}[x]$ subroutine FillInRungs(x, y) $x_0 \leftarrow x, x_1 \leftarrow y$ SetTable(\mathbf{g}, x_0, x_1) For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow$ SimROsub(x_{i-1}) SetTable(\mathbf{g}, x_i, x_{i+1}) </pre>	\mathbf{G}_6	<pre> subroutine MakeLadder(x) $s_{-q_1} \leftarrow$ Random() For $i = -q_1$ to $q_1 - 1$ $s_{i+1} \leftarrow$ ROsub(s_i) FillInRungs(x, s_0) subroutine ROsub(x) If $\mathbf{G}[x]$ Ret $\mathbf{G}[x]$ Ret $\mathbf{G}[x] \leftarrow$ Random() subroutine SimROsub(x) SetTable($\mathbf{G}, x, \text{ROsub}(x)$) Ret $\mathbf{G}[x]$ subroutine Random() $y \leftarrow_s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ then Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom ++ Ret y subroutine SetTable(\mathbf{T}, x, y) If $\mathbf{T}[x]$ and $\mathbf{T}[x] \neq y$ then Ret $\mathbf{T}[x] = y$ $\mathbf{T}^{-1}[y] = x$ subroutine BuildXprime() $X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow_s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$ subroutine Finalization() While NumCallsToRandom < MAX_CALLS do Random() </pre>	\mathbf{G}_6 (cont.)
--	----------------	--	------------------------

Figure 22: Game \mathbf{G}_6 for the proof of Theorem 3.3.

<pre> procedure OnLeftQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow 2 BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ ROsub(x) $Y \leftarrow Y \cup \{x, \mathbf{G}[x]\}$ Finalization() Ret $\mathbf{G}[x]$ procedure OnRightQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow $4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $\mathbf{g}[x_i] \neq \perp$ then $y \leftarrow \mathbf{G}^{-i}[\mathbf{g}[x_i]]$ FillInRungs(x, y) Break // (For loop) $x_{i+1} \leftarrow$ ROsub(x_i) If $i > q_1$ MakeLadder(x) $Y \leftarrow Y \cup \{x, \mathbf{g}[x]\}$ Finalization() Ret $\mathbf{g}[x]$ subroutine FillInRungs(x, y) $x_0 \leftarrow x, x_1 \leftarrow y$ SetTable(\mathbf{g}, x_0, x_1) For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow$ ROsub(x_{i-1}) SetTable(\mathbf{g}, x_i, x_{i+1}) </pre>	<p>G_7</p>	<pre> subroutine MakeLadder(x) $s_{-q_1} \leftarrow$ Random() For $i = -q_1$ to $q_1 - 1$ $s_{i+1} \leftarrow$ ROsub(s_i) FillInRungs(x, s_0) subroutine ROsub(x) If $\mathbf{G}[x]$ Ret $\mathbf{G}[x]$ $\mathbf{G}[x] \leftarrow$ Random() SetTable($\mathbf{G}, x, \mathbf{G}[x]$) Ret $\mathbf{G}[x]$ subroutine Random() $y \leftarrow^s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ then Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom ++ Ret y subroutine SetTable(\mathbf{T}, x, y) If $\mathbf{T}[x]$ and $\mathbf{T}[x] \neq y$ then Ret $\mathbf{T}[x] = y$ $\mathbf{T}^{-1}[y] = x$ subroutine BuildXprime() $X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow^s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$ subroutine Finalization() While NumCallsToRandom $<$ MAX_CALLS do Random() </pre>	<p>G_7 (cont.)</p>
---	-------------------------	---	---------------------------------

Figure 23: Game G_7 for the proof of Theorem 3.3.

<p>procedure <u>OnLeftQuery</u>(x): G_8</p> <p>NumCallsToRandom $\leftarrow 0$; MAX_CALLS $\leftarrow 2$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ ROsub(x) $Y \leftarrow Y \cup \{x, \mathbf{G}[x]\}$ Finalization() Ret $\mathbf{G}[x]$</p> <p>procedure <u>OnRightQuery</u>(x):</p> <p>NumCallsToRandom $\leftarrow 0$; MAX_CALLS $\leftarrow 4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $\mathbf{g}[x_i] \neq \perp$ then $y \leftarrow \mathbf{G}^{-i}[\mathbf{g}[x_i]]$ FillInRungs(x, y) Break // (For loop) $x_{i+1} \leftarrow \text{ROsub}(x_i)$ If $i > q_1$ MakeLadder(x) $Y \leftarrow Y \cup \{x, \mathbf{g}[x]\}$ Finalization() Ret $\mathbf{g}[x]$</p> <p>subroutine <u>FillInRungs</u>(x, y)</p> <p>$x_0 \leftarrow x, x_1 \leftarrow y$ SetTable(\mathbf{g}, x_0, x_1) For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow \text{ROsub}(x_{i-1})$ SetTable(\mathbf{g}, x_i, x_{i+1})</p>	<p>subroutine <u>MakeLadder</u>(x) G_8 (cont.)</p> <p>$s_{-q_1} \leftarrow \text{Random}()$ For $i = -q_1$ to $q_1 - 1$ $s_{i+1} \leftarrow \text{ROsub}(s_i)$ $\ell \leftarrow 0$ While $\mathbf{G}^{(-\ell)}[x] \neq \perp$ $\ell \leftarrow \ell + 1$ For $i = -q_1$ to $-\ell$ $\mathbf{G}^{-1}[s_i] \leftarrow \perp$ If $i > \ell$ then $\mathbf{G}[s_i] \leftarrow \perp$ FillInRungs(x, s_0)</p> <p>subroutine <u>ROsub</u>(x)</p> <p>If $\mathbf{G}[x]$ Ret $\mathbf{G}[x]$ $\mathbf{G}[x] \leftarrow \text{Random}()$ SetTable($\mathbf{G}, x, \mathbf{G}[x]$) Ret $\mathbf{G}[x]$</p> <p>subroutine <u>Random</u>()</p> <p>$y \leftarrow_s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ then Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom $++$ Ret y</p> <p>subroutine <u>SetTable</u>(\mathbf{T}, x, y)</p> <p>If $\mathbf{T}[x]$ and $\mathbf{T}[x] \neq y$ then Ret $\mathbf{T}[x] = y$ $\mathbf{T}^{-1}[y] = x$</p> <p>subroutine <u>BuildXprime</u>()</p> <p>$X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow_s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$</p> <p>subroutine <u>Finalization</u>()</p> <p>While NumCallsToRandom $< \text{MAX_CALLS}$ do Random()</p>
--	--

Figure 24: Game G_8 for the proof of Theorem 3.3.

<pre> procedure OnLeftQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow 2 BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ ROsub(x) $Y \leftarrow Y \cup \{x, \mathbf{G}[x]\}$ Finalization() Ret $\mathbf{G}[x]$ procedure OnRightQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow $4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $\mathbf{g}[x_i] \neq \perp$ then $y \leftarrow \mathbf{G}^{-i}[\mathbf{g}[x_i]]$ FillInRungs(x, y) Break // (For loop) $x_{i+1} \leftarrow$ ROsub(x_i) If $i > q_1$ MakeLadder(x) $Y \leftarrow Y \cup \{x, \mathbf{g}[x]\}$ Finalization() Ret $\mathbf{g}[x]$ subroutine FillInRungs(x, y) $x_0 \leftarrow x, x_1 \leftarrow y$ SetTable(\mathbf{g}, x_0, x_1) For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow$ ROsub(x_{i-1}) SetTable(\mathbf{g}, x_i, x_{i+1}) </pre>	<p>G_9</p>	<pre> subroutine MakeLadder(x) $\ell \leftarrow 0$ While $\mathbf{G}^{(-\ell)}[x] \neq \perp$ $\ell \leftarrow \ell + 1$ $s_{-\ell} \leftarrow$ Random() For $i = -\ell$ to $q_1 - 1$ $s_{i+1} \leftarrow$ ROsub(s_i) FillInRungs(x, s_0) subroutine ROsub(x) If $\mathbf{G}[x]$ Ret $\mathbf{G}[x]$ $\mathbf{G}[x] \leftarrow$ Random() SetTable($\mathbf{G}, x, \mathbf{G}[x]$) Ret $\mathbf{G}[x]$ subroutine Random() $y \leftarrow_s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ then Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom ++ Ret y subroutine SetTable(\mathbf{T}, x, y) If $\mathbf{T}[x]$ and $\mathbf{T}[x] \neq y$ then Ret $\mathbf{T}[x] = y$ $\mathbf{T}^{-1}[y] = x$ subroutine BuildXprime() $X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow_s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$ subroutine Finalization() While NumCallsToRandom $<$ MAX_CALLS do Random() </pre>	<p>G_9 (cont.)</p>
---	-------------------------	--	---------------------------------

Figure 25: Game G_9 for the proof of Theorem 3.3.

<pre> procedure OnLeftQuery(x): G₁₀ NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow 2 BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ ROsub(x, \perp) $Y \leftarrow Y \cup \{x, \mathbf{G}[x]\}$ Finalization() Ret $\mathbf{G}[x]$ procedure OnRightQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow $4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $\mathbf{g}[x_i] \neq \perp$ then $y \leftarrow \mathbf{G}^{-i}[\mathbf{g}[x_i]]$ FillInRungs(x, y) Break // (For loop) $x_{i+1} \leftarrow$ ROsub(x_i, \perp) If $i > q_1$ MakeLadder(x) $Y \leftarrow Y \cup \{x, \mathbf{g}[x]\}$ Finalization() Ret $\mathbf{g}[x]$ subroutine FillInRungs(x, y) $x_0 \leftarrow x, x_1 \leftarrow y$ SetTable(\mathbf{g}, x_0, x_1) For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow$ ROsub(x_{i-1}) SetTable(\mathbf{g}, x_i, x_{i+1}) </pre>	<pre> subroutine MakeLadder(x) G₁₀ (cont.) $\ell \leftarrow 0$ While $\mathbf{G}^{(-\ell)}[x] \neq \perp$ $\ell \leftarrow \ell + 1$ $s_{-\ell} \leftarrow$ Random() For $i = -\ell$ to $q_1 - 1$ $s_{i+1} \leftarrow$ Random() ROsub(s_i, s_{i+1}) FillInRungs(x, s_0) subroutine ROsub(x, z) If $\mathbf{G}[x]$ Ret $\mathbf{G}[x]$ If $z = \perp$ then $z \leftarrow$ Random() SetTable(\mathbf{G}, x, z) Ret $\mathbf{G}[x]$ subroutine Random() $y \leftarrow_s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ then Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom $++$ Ret y subroutine SetTable(\mathbf{T}, x, y) If $\mathbf{T}[x]$ and $\mathbf{T}[x] \neq y$ then Ret $\mathbf{T}[x] = y$ $\mathbf{T}^{-1}[y] = x$ subroutine BuildXprime() $X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow_s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$ subroutine Finalization() While NumCallsToRandom $<$ MAX_CALLS do Random() </pre>
---	---

Figure 26: Game G_{10} for the proof of Theorem 3.3.

<pre> procedure OnLeftQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow 2 BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ ROsub(x, \perp) $Y \leftarrow Y \cup \{x, \mathbf{G}[x]\}$ Finalization() Ret $\mathbf{G}[x]$ procedure OnRightQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow $4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $\mathbf{g}[x_i] \neq \perp$ then $y \leftarrow \mathbf{G}^{-i}[\mathbf{g}[x_i]]$ FillInRungs(x, y) Break // (For loop) $x_{i+1} \leftarrow$ ROsub(x_i, \perp) If $i > q_1$ MakeLadder(x) $Y \leftarrow Y \cup \{x, \mathbf{g}[x]\}$ Finalization() Ret $\mathbf{g}[x]$ subroutine FillInRungs(x, y) $x_0 \leftarrow x, x_1 \leftarrow y$ SetTable(\mathbf{g}, x_0, x_1) For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow$ ROsub(x_{i-1}) SetTable(\mathbf{g}, x_i, x_{i+1}) </pre>	<p style="text-align: right;">G_{11}</p>	<pre> subroutine MakeLadder(x) $\ell \leftarrow 0$ While $\mathbf{G}^{(-\ell)}[x] \neq \perp$ $\ell \leftarrow \ell + 1$ $s_{-\ell} \leftarrow$ Random() For $i = -\ell$ to $q_1 - 1$ $s_{i+1} \leftarrow$ Random() ROsub(s_i, s_{i+1}) FillInRungs(x, s_0) subroutine ROsub(x, z) If $\mathbf{G}[x]$ Ret $\mathbf{G}[x]$ If $z = \perp$ If $\mathbf{g}[x]$ $\mathbf{gProxy}[x] \leftarrow \mathbf{g}[x]$ If $\mathbf{gProxy}[x] = \perp$ $\mathbf{gProxy}[x] \leftarrow$ Random() $\mathbf{gProxy}[x].\text{fresh} \leftarrow \text{true}$ $z \leftarrow$ Random() SetTable(\mathbf{G}, x, z) Ret $\mathbf{G}[x]$ subroutine Random() $y \leftarrow_s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ then Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom ++ Ret y subroutine SetTable(\mathbf{T}, x, y) If $\mathbf{T}[x]$ and $\mathbf{T}[x] \neq y$ then Ret $\mathbf{T}[x] = y$ $\mathbf{T}^{-1}[y] = x$ subroutine BuildXprime() $X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow_s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$ subroutine Finalization() While NumCallsToRandom < MAX_CALLS do Random() </pre>	<p style="text-align: right;">G_{11} (cont.)</p>
---	---	--	---

Figure 27: Game G_{11} for the proof of Theorem 3.3.

<pre> procedure OnLeftQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow 2 BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ ROsub(x, \perp) $Y \leftarrow Y \cup \{x, G[x]\}$ Finalization() Ret $G[x]$ procedure OnRightQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow $4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $g[x_i] \neq \perp$ then $y \leftarrow G^{-i}[g[x_i]]$ FillInRungs(x, y) Break // (For loop) $x_{i+1} \leftarrow$ ROsub(x_i, \perp) If $i > q_1$ MakeLadder(x) $Y \leftarrow Y \cup \{x, g[x]\}$ Finalization() Ret $g[x]$ subroutine FillInRungs(x, y) $x_0 \leftarrow x, x_1 \leftarrow y$ SetTable(g, x_0, x_1) For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow$ ROsub(x_{i-1}) SetTable(g, x_i, x_{i+1}) subroutine Finalization() While NumCallsToRandom $<$ MAX_CALLS do Random() </pre>	<p style="text-align: right;">G₁₂</p> <pre> subroutine MakeLadder(x) $\ell \leftarrow 0$ While $G^{(-\ell)}[x] \neq \perp$ $\ell \leftarrow \ell + 1$ $s_{-\ell} \leftarrow$ GetFresh($gProxy, G^{-\ell}[x]$) For $i = -\ell$ to $q_1 - 1$ $s_{i+1} \leftarrow$ GetFresh($gProxy, G^{i+1}[x]$) ROsub(s_i, s_{i+1}) FillInRungs(x, s_0) subroutine ROsub(x, z) If $G[x]$ Ret $G[x]$ If $z = \perp$ If $g[x]$ $gProxy[x] \leftarrow g[x]$ If $gProxy[x] = \perp$ $gProxy[x] \leftarrow$ Random() $gProxy[x].fresh \leftarrow$ true $z \leftarrow$ Random() SetTable(G, x, z) Ret $G[x]$ subroutine Random() $y \leftarrow_s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ then Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom $++$ Ret y subroutine SetTable(T, x, y) If $T[x]$ and $T[x] \neq y$ then Ret $T[x] = y$ $T^{-1}[y] = x$ subroutine BuildXprime() $X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow_s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$ subroutine GetFresh(T, x) If $T[x]$ and $T[x].fresh = true$ $T[x].fresh = false$ Ret $T[x]$ Ret Random() </pre> <p style="text-align: right;">G₁₂ (cont.)</p>
---	---

Figure 28: Game G₁₂ for the proof of Theorem 3.3.

<p>procedure <u>OnLeftQuery(x):</u> G_{13} NumCallsToRandom $\leftarrow 0$; MAX_CALLS $\leftarrow 2$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ ROsub(x, \perp) $Y \leftarrow Y \cup \{x, G[x]\}$ Finalization() Ret $G[x]$</p> <p>procedure <u>OnRightQuery(x):</u> NumCallsToRandom $\leftarrow 0$; MAX_CALLS $\leftarrow 4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $g[x_i] \neq \perp$ then $y \leftarrow G^{-i}[g[x_i]]$ FillInRungs(x, y) Break // (For loop) $x_{i+1} \leftarrow \text{ROsub}(x_i, \perp)$ If $i > q_1$ MakeLadder(x) $Y \leftarrow Y \cup \{x, g[x]\}$ Finalization() Ret $g[x]$</p> <p>subroutine <u>FillInRungs(x, y)</u> $x_0 \leftarrow x, x_1 \leftarrow y$ SetTable(g, x_0, x_1) For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow \text{ROsub}(x_{i-1})$ SetTable(g, x_i, x_{i+1})</p> <p>subroutine <u>Finalization()</u> While NumCallsToRandom $<$ MAX_CALLS do Random()</p>	<p>subroutine <u>MakeLadder(x)</u> G_{13} (cont.) $\ell \leftarrow 0$ While $G^{(-\ell)}[x] \neq \perp$ $\ell \leftarrow \ell + 1$ $s_{-\ell} \leftarrow \text{gProxy}[G^{-\ell}[x]]$ For $i = -\ell$ to $q_1 - 1$ $s_{i+1} \leftarrow \text{gProxy}[G^{i+1}[x]]$ ROsub(s_i, s_{i+1}) FillInRungs(x, s_0)</p> <p>subroutine <u>ROsub(x, z)</u> If $G[x]$ Ret $G[x]$ If $z = \perp$ If $g[x]$ $\text{gProxy}[x] \leftarrow g[x]$ If $\text{gProxy}[x] = \perp$ $\text{gProxy}[x] \leftarrow \text{Random}()$ $z \leftarrow \text{Random}()$ SetTable(G, x, z) Ret $G[x]$</p> <p>subroutine <u>Random()</u> $y \leftarrow_s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ then Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom $++$ Ret y</p> <p>subroutine <u>SetTable(T, x, y)</u> If $T[x]$ and $T[x] \neq y$ then Ret $T[x] = y$ $T^{-1}[y] = x$</p> <p>subroutine <u>BuildXprime()</u> $X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow_s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$</p>
--	---

Figure 29: Game G_{13} for the proof of Theorem 3.3.

<p>procedure <u>OnLeftQuery</u>(x): G_{14}</p> <p>NumCallsToRandom $\leftarrow 0$; MAX_CALLS $\leftarrow 2$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ ROsub(x, \perp) $Y \leftarrow Y \cup \{x, G[x]\}$ Finalization() Ret $G[x]$</p> <p>procedure <u>OnRightQuery</u>(x):</p> <p>NumCallsToRandom $\leftarrow 0$; MAX_CALLS $\leftarrow 4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $g[x_i] \neq \perp$ then $y \leftarrow G^{-i}[g[x_i]]$ FillInRungs(x, y) Break // (For loop) $x_{i+1} \leftarrow \text{ROsub}(x_i, \perp)$ If $i > q_1$ MakeLadder(x) $Y \leftarrow Y \cup \{x, g[x]\}$ Finalization() Ret $g[x]$</p> <p>subroutine <u>FillInRungs</u>(x, y)</p> <p>$x_0 \leftarrow x, x_1 \leftarrow y$ SetTable(g, x_0, x_1) For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow \text{ROsub}(x_{i-1})$ SetTable(g, x_i, x_{i+1})</p> <p>subroutine <u>Finalization</u>()</p> <p>While NumCallsToRandom $< \text{MAX_CALLS}$ do Random()</p>	<p>subroutine <u>MakeLadder</u>(x) G_{14} (cont.)</p> <p>$\ell \leftarrow 0$ While $G^{(-\ell)}[x] \neq \perp$ $\ell \leftarrow \ell + 1$ $s_{-\ell} \leftarrow \text{gProxy}[G^{-\ell}[x]]$ For $i = -\ell$ to $q_1 - 1$ $s_{i+1} \leftarrow \text{gProxy}[G^{i+1}[x]]$ ROsub(s_i, s_{i+1}) FillInRungs(x, s_0)</p> <p>subroutine <u>ROsub</u>(x, z)</p> <p>If $G[x]$ Ret $G[x]$ If $z = \perp$ If $g[x]$ $\text{gProxy}[x] \leftarrow g[x]$ If $\text{gProxy}[x] = \perp$ $\text{gProxy}[x] \leftarrow \text{Random}()$ $z \leftarrow \text{Random}()$ $\text{gProxy}[g\text{Proxy}[x]] \leftarrow z$ SetTable(G, x, z) Ret $G[x]$</p> <p>subroutine <u>Random</u>()</p> <p>$y \leftarrow_s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ then Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom $++$ Ret y</p> <p>subroutine <u>SetTable</u>(T, x, y)</p> <p>If $T[x]$ and $T[x] \neq y$ then Ret $T[x] = y$ $T^{-1}[y] = x$</p> <p>subroutine <u>BuildXprime</u>()</p> <p>$X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow_s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$</p>
--	--

Figure 30: Game G_{14} for the proof of Theorem 3.3.

<pre> procedure OnLeftQuery(x): G₁₅ NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow 2 BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ ROsub(x) $Y \leftarrow Y \cup \{x, \mathbf{G}[x]\}$ Finalization() Ret $\mathbf{G}[x]$ procedure OnRightQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow $4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $\mathbf{g}[x_i] \neq \perp$ then $y \leftarrow \mathbf{G}^{-i}[\mathbf{g}[x_i]]$ FillInRungs(x, y) Break // (For loop) $x_{i+1} \leftarrow$ ROsub(x_i) If $i > q_1$ MakeLadder(x) $Y \leftarrow Y \cup \{x, \mathbf{g}[x]\}$ Finalization() Ret $\mathbf{g}[x]$ subroutine FillInRungs(x, y) $x_0 \leftarrow x, x_1 \leftarrow y$ SetTable(\mathbf{g}, x_0, x_1) For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow$ ROsub(x_{i-1}) SetTable(\mathbf{g}, x_i, x_{i+1}) subroutine Finalization() While NumCallsToRandom $<$ MAX_CALLS do Random() </pre>	<pre> subroutine MakeLadder(x) G₁₅ (cont.) $\ell \leftarrow 0$ While $\mathbf{G}^{(-\ell)}[x] \neq \perp$ $\ell \leftarrow \ell + 1$ $s_{-\ell} \leftarrow$ gProxy[$\mathbf{G}^{-\ell}[x]$] For $i = -\ell$ to $q_1 - 1$ $s_{i+1} \leftarrow$ ROsub(s_i) FillInRungs(x, s_0) subroutine ROsub(x) If $\mathbf{G}[x]$ Ret $\mathbf{G}[x]$ If gProxy[x] $y \leftarrow$ gProxy[x] else $y \leftarrow$ gProxy[x] \leftarrow Random() If gProxy[y] $z \leftarrow$ gProxy[y] else $z \leftarrow$ gProxy[y] \leftarrow Random() SetTable(\mathbf{G}, x, z) Ret $\mathbf{G}[x]$ subroutine Random() $y \leftarrow_s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ then Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom ++ Ret y subroutine SetTable(\mathbf{T}, x, y) If $\mathbf{T}[x]$ and $\mathbf{T}[x] \neq y$ then Ret $\mathbf{T}[x] = y$ $\mathbf{T}^{-1}[y] = x$ subroutine BuildXprime() $X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow_s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$ </pre>
--	--

Figure 31: Game G_{15} for the proof of Theorem 3.3.

<p>procedure <u>OnLeftQuery</u>(x): G_{16}</p> <p>NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow 2 BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ ROsub(x) $Y \leftarrow Y \cup \{x, \mathbf{G}[x]\}$ Finalization() Ret $\mathbf{G}[x]$</p> <p>procedure <u>OnRightQuery</u>(x):</p> <p>NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow $4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $\mathbf{g}[x_i] \neq \perp$ and $\mathbf{g}[x_i].\text{KnownToSim}$ then FillInRungs(x) Break // (For loop) $x_{i+1} \leftarrow \text{ROsub}(x_i)$ If $i > q_1$ MakeLadder(x) $Y \leftarrow Y \cup \{x, \mathbf{g}[x]\}$ Finalization() Ret $\mathbf{g}[x]$</p> <p>subroutine <u>FillInRungs</u>(x)</p> <p>$x_0 \leftarrow x$ $x_1 \leftarrow \mathbf{g}[x_0]$ $\mathbf{g}[x_0].\text{KnownToSim} \leftarrow \text{true}$ For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow \text{ROsub}(x_{i-1})$ $\mathbf{g}[x_i].\text{KnownToSim} \leftarrow \text{true}$</p> <p>subroutine <u>Finalization</u>()</p> <p>While NumCallsToRandom $<$ MAX_CALLS do Random()</p>	<p>subroutine <u>MakeLadder</u>(x) G_{16} (cont.)</p> <p>$\ell \leftarrow 0$ While $\mathbf{G}^{(-\ell)}[x] \neq \perp$ $\ell \leftarrow \ell + 1$ $s_{-\ell} \leftarrow \mathbf{g}[\mathbf{G}^{-\ell}[x]]$ For $i = -\ell$ to $q_1 - 1$ $s_{i+1} \leftarrow \text{ROsub}(s_i)$ FillInRungs(x)</p> <p>subroutine <u>ROsub</u>(x)</p> <p>If $\mathbf{G}[x]$ Ret $\mathbf{G}[x]$ If $\mathbf{g}[x]$ $y \leftarrow \mathbf{g}[x]$ else $y \leftarrow \text{Random}()$ SetTable(\mathbf{g}, x, y) If $\mathbf{g}[y]$ $z \leftarrow \mathbf{g}[y]$ else $z \leftarrow \text{Random}()$ SetTable(\mathbf{g}, y, z) SetTable(\mathbf{G}, x, z) Ret $\mathbf{G}[x]$</p> <p>subroutine <u>Random</u>()</p> <p>$y \leftarrow_{\\$} \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ then Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom $++$ Ret y</p> <p>subroutine <u>SetTable</u>(T, x, y)</p> <p>If $T[x]$ and $T[x] \neq y$ then Ret $T[x] = y$ $T^{-1}[y] = x$</p> <p>subroutine <u>BuildXprime</u>()</p> <p>$X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow_{\\$} \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$</p>
--	---

Figure 32: Game G_{16} for the proof of Theorem 3.3.

<pre> procedure OnLeftQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow 2 BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ ROsub(x) $Y \leftarrow Y \cup \{x, \mathbf{g}[x]\}$ Finalization() Ret $\mathbf{g}[x]$ procedure OnRightQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow $4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $\mathbf{g}[x_i] \neq \perp$ and $\mathbf{g}[x_i].\text{KnownToSim}$ then FillInRungs(x) Break // (For loop) $x_{i+1} \leftarrow \text{ROsub}(x_i)$ If $i > q_1$ MakeLadder(x) $Y \leftarrow Y \cup \{x, \mathbf{g}[x]\}$ Finalization() Ret $\mathbf{g}[x]$ subroutine FillInRungs(x) $x_0 \leftarrow x$ $x_1 \leftarrow \mathbf{g}[x_0]$ $\mathbf{g}[x_0].\text{KnownToSim} \leftarrow \text{true}$ For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow \text{ROsub}(x_{i-1})$ $\mathbf{g}[x_i].\text{KnownToSim} \leftarrow \text{true}$ subroutine Finalization() While NumCallsToRandom $<$ MAX_CALLS do Random() </pre>	<p style="text-align: right;">G₁₇</p>	<pre> subroutine MakeLadder(x) $\ell \leftarrow 0$ While $\mathbf{g}^{(-2\ell)}[x] \neq \perp$ $\ell \leftarrow \ell + 1$ $s_{-\ell} \leftarrow \mathbf{g}[\mathbf{g}^{-2\ell}[x]]$ For $i = -\ell$ to $q_1 - 1$ $s_{i+1} \leftarrow \text{ROsub}(s_i)$ FillInRungs(x) subroutine ROsub(x) If $\mathbf{g}[x]$ $y \leftarrow \mathbf{g}[x]$ else $y \leftarrow \text{Random}()$ SetTable(\mathbf{g}, x, y) If $\mathbf{g}[y]$ $z \leftarrow \mathbf{g}[y]$ else $z \leftarrow \text{Random}()$ SetTable(\mathbf{g}, y, z) Ret z subroutine Random() $y \leftarrow_s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ then Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom ++ Ret y subroutine SetTable(\mathbf{T}, x, y) If $\mathbf{T}[x]$ and $\mathbf{T}[x] \neq y$ then Ret $\mathbf{T}[x] = y$ $\mathbf{T}^{-1}[y] = x$ subroutine BuildXprime() $X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow_s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$ </pre>	<p style="text-align: right;">G₁₇ (cont.)</p>
--	---	---	---

Figure 33: Game G₁₇ for the proof of Theorem 3.3.

<pre> procedure OnLeftQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow 2 BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ ROsub(x) $Y \leftarrow Y \cup \{x, g[g[x]]\}$ Finalization() Ret $g[g[x]]$ procedure OnRightQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow $4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ then Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 0$ to q_1 If $g[x_i] \neq \perp$ and $g[x_i].\text{KnownToSim}$ then FillInRungs(x) Break // (For loop) $x_{i+1} \leftarrow \text{ROsub}(x_i)$ If $i > q_1$ FillInRungs(x) $Y \leftarrow Y \cup \{x, g[x]\}$ Finalization() Ret $g[x]$ subroutine FillInRungs(x) $x_0 \leftarrow x$ $x_1 \leftarrow g[x_0]$ $g[x_0].\text{KnownToSim} \leftarrow \text{true}$ For $i = 1$ to $2q_1 + 1$ $x_{i+1} \leftarrow \text{ROsub}(x_{i-1})$ $g[x_i].\text{KnownToSim} \leftarrow \text{true}$ subroutine Finalization() While NumCallsToRandom $<$ MAX_CALLS do Random() </pre>	<p style="text-align: right;">G₁₈</p> <pre> subroutine ROsub(x) If $g[x]$ $y \leftarrow g[x]$ else $y \leftarrow g[x] \leftarrow \text{Random}()$ If $g[y]$ $z \leftarrow g[y]$ else $z \leftarrow g[y] \leftarrow \text{Random}()$ Ret z subroutine Random() $y \leftarrow s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ then Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom ++ Ret y subroutine BuildXprime() $X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$ </pre>
--	---

Figure 34: Game G₁₈ for the proof of Theorem 3.3.

<pre> procedure OnLeftQuery(x): G₁₉ G₂₀ NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow 2 BuildXprime() If $x \in (X \cup X') \setminus Y$ bad \leftarrow true Abort $X \leftarrow X \cup \{x\}$ ROsub(x) $Y \leftarrow Y \cup \{x, g[g[x]]\}$ Finalization() Ret $g[g[x]]$ procedure OnRightQuery(x): NumCallsToRandom \leftarrow 0; MAX_CALLS \leftarrow $4q_1 + 3$ BuildXprime() If $x \in (X \cup X') \setminus Y$ bad \leftarrow true Abort $X \leftarrow X \cup \{x\}$ $x_0 \leftarrow x$ For $i = 1$ to $2q_1 + 1$ ROsub(x_{i-1}) $x_i \leftarrow g[x_{i-1}]$ $g[x_{i-1}].\text{KnownToSim} \leftarrow \text{true}$ $g[x_i].\text{KnownToSim} \leftarrow \text{true}$ $Y \leftarrow Y \cup \{x, g[x]\}$ Finalization() Ret $g[x]$ subroutine Finalization() While NumCallsToRandom $<$ MAX_CALLS do Random() </pre>	<pre> subroutine ROsub(x) G₁₉ G₂₀ If $g[x]$ $y \leftarrow g[x]$ else $y \leftarrow g[x] \leftarrow \text{Random}()$ If $g[y]$ $z \leftarrow g[y]$ else $z \leftarrow g[y] \leftarrow \text{Random}()$ Ret z subroutine Random() $y \leftarrow_s \{0, 1\}^n$ BuildXprime() If $y \in X \cup X'$ bad \leftarrow true Abort $X \leftarrow X \cup \{y\}$ NumCallsToRandom ++ Ret y subroutine BuildXprime() $X' \leftarrow \emptyset$ While $X + X' < (4q_1 + 3)q_2 + 2q_1$ do $z \leftarrow_s \{0, 1\}^n \setminus (X \cup X')$ $X' \leftarrow X' \cup \{z\}$ </pre>
---	---

Figure 35: Games G_{19} and G_{20} for the proof of Theorem 3.3.

<pre> procedure <u>OnLeftQuery</u>(x): ROsub(x) Ret $\mathbf{g}[\mathbf{g}[x]]$ procedure <u>OnRightQuery</u>(x): $x_0 \leftarrow x$ For $i = 1$ to $2q_1 + 1$ ROsub(x_{i-1}) $x_i \leftarrow \mathbf{g}[x_{i-1}]$ Ret $\mathbf{g}[x]$ subroutine <u>ROsub</u>(x) If $\mathbf{g}[x] = \perp$ $\mathbf{g}[x] \leftarrow \text{Random}()$ $y \leftarrow \mathbf{g}[x]$ If $\mathbf{g}[y] = \perp$ $\mathbf{g}[y] \leftarrow \text{Random}()$ $z \leftarrow \mathbf{g}[y]$ Ret z subroutine <u>Random</u>() $y \leftarrow_s \{0, 1\}^n$ Ret y </pre>	G ₂₁	<pre> procedure <u>OnLeftQuery</u>(x): ROsub(x) Ret $\mathbf{g}[\mathbf{g}[x]]$ procedure <u>OnRightQuery</u>(x): If $\mathbf{g}[x] = \perp$ $\mathbf{g}[x] \leftarrow_s \{0, 1\}^n$ Ret $\mathbf{g}[x]$ subroutine <u>ROsub</u>(x) If $\mathbf{g}[x] = \perp$ $\mathbf{g}[x] \leftarrow_s \{0, 1\}^n$ $y \leftarrow \mathbf{g}[x]$ If $\mathbf{g}[y] = \perp$ $\mathbf{g}[y] \leftarrow_s \{0, 1\}^n$ Ret $\mathbf{g}[y]$ </pre>	G ₂₂
--	-----------------	--	-----------------

Figure 36: Games G₂₁ and G₂₂ for the proof of Theorem 3.3.