The Iterated Random Function Problem[*]

Ritam Bhaumik[1], Nilanjan Datta[2], Avijit Dutta[1], Nicky Mouha[3,4], and Mridul Nandi[1]

[1] Indian Statistical Institute, Kolkata, India.
[2] Indian Institute of Technology, Kharagpur, India.
[3] National Institute of Standards and Technology, Gaithersburg, MD, USA.
[4] Project-team SECRET, Inria, France.
bhaumik.ritam@gmail.com, nilanjan_isi_jrf@yahoo.com,
avirocks.dutta13@gmail.com, nicky@mouha.be, mridul.nandi@gmail.com

Abstract. At CRYPTO 2015, Minaud and Seurin introduced and studied the *iterated random permutation* problem, which is to distinguish the r-th iterate of a random permutation from a random permutation. In this paper, we study the closely related *iterated random function* problem, and prove the first almost-tight bound in the adaptive setting. More specifically, we prove that the advantage to distinguish the r-th iterate of a random function from a random function using q queries is bounded by $O(q^2 r (\log r)^3 / N)$, where N is the size of the domain. In previous work, the best known bound was $O(q^2 r^2 / N)$, obtained as a direct result of interpreting the iterated random function problem as a special case of CBC-MAC based on a random function. For the iterated random function problem, the best known attack has an advantage of $\Omega(q^2 r / N)$, showing that our security bound is tight up to a factor of $(\log r)^3$.

Keywords: Iterated random function, random function, pseudorandom function, password hashing, Patarin, H-coefficient technique, provable security.

1 Introduction

Take any n-bit hash function h. Assuming that this hash function can be modelled as a random function, the probability that the outputs of h collide given $q \ll 2^{n/2}$ distinct inputs is about $q^2/2^n$: the well-known birthday attack.

Now let us consider another hash function g, defined as the r-th iterate of h, i.e. $g(m) = h(h(\dots h(m)))$, where h is applied r times. For the same number of queries $q \ll 2^{n/2}$, the birthday attack has about an r times higher probability to succeed for g than for h (see e.g. Preneel and van Oorschot [18, Lemma 2]).

Iteration is of fundamental importance in many cryptographic constructions. For example, a "possibly weak" function may be iterated to improve its resistance

---

[*] Certain algorithms and commercial products are identified in this paper to foster understanding. Such identification does not imply recommendation or endorsement by NIST, nor does it imply that the algorithms or products identified are necessarily the best available for the purpose.

against various cryptanalysis attacks, or a password hashing function may be iterated to slow down dictionary attacks. But quite surprisingly, the security of iterating a random function is not yet a well-understood problem.

In the aforementioned (non-adaptive) birthday attack, the distinguishing advantage between a random function and an iterated random function increases by about a factor r. But what happens if we consider adaptive collision-finding attacks as well? Or in general, what if we want to consider any adaptive attack, not necessarily a collision-finding attack? Could there be more efficient attacks that have not yet been discovered?

Recently at CRYPTO 2015, Minaud and Seurin [15] put this possibility to rest for the iterated random permutation problem. They proved that the advantage to distinguish an iterated random permutation from a random permutation using q queries is bounded by $O(qr/N)$, where N is the size of the domain, and showed that their bound is almost tight by providing a matching attack.

In this paper, we will do the same for the iterated random function problem. Whereas the best bound in previous work is $O(q^2r^2/N)$, we will prove a bound of $O(q^2r(\log r)^3/N)$, where log is the logarithm to the base e. Our bound is tight up to a factor of about $(\log r)^3$, and thereby rules out the possibility of better attacks.

NOTE. We will focus on asymptotic bounds for large r, as this is parameter range where large improvements over the currently best-known bounds can be achieved. Although our bounds hold for any $r \geq 2$, we will apply generous relaxations to derive an easy-to-see bound that only improves the currently-known bounds for larger, but nevertheless practically-relevant values of r. Also, we will only consider the iteration of a uniformly random function in an information-theoretic setting. A simple hybrid argument can be used to extend this result to the pseudorandom function (prf) advantage in a computational setting, as shown by Minaud and Seurin [15, Theorem 1] for the iterated random permutation problem.

APPLICATIONS. In spite of the frequent use of iterated random functions in practice, this paper is the first to study this problem without relying on the trivial CBC-MAC bound. The most obvious application of iterated random functions is in password hashing, where a hash function is iterated in order to slow down brute force attacks. This idea is used in PKCS #5's PBKDF1 and PBKDF2. In typical password-based key derivation functions, the iteration count is often quite high, ranging from several hundreds of thousands [9], to even ten million [19], as suggested by NIST for critical keys. To analyse the effect of iteration in these constructions, it is common to model the secret low-entropy password as a random-but-known key [11], or even an adversarially-chosen input [20]. But also small values of r, such as $r = 2$, appear in practical applications. In the book "Practical cryptography" [13], Ferguson and Schneier suggest to use SHA-256(SHA-256(m)) to avoid length-extension attacks. They use this construction in their RSA encryption implementation, as well as in their Fortuna random number generator. Interestingly, about 2^{64} evaluations

of SHA-256(SHA-256(m)) are performed *every second* as part of bitcoin mining [21].

RELATED WORK. The security of an iterated random function was first analysed by Yao and Yin [22,23], when they analysed the security of the password-based key derivation functions PBKDF1 and PBKDF2. Their work is parallel to that of Wagner and Goldberg [20], who analysed the security of an iterated random permutation in the context of the Unix password hashing algorithm. Bellare et al. [4] extended these results, and also pointed out some problems in the proofs of Yao and Yin.

As Wagner and Goldberg explain in [20], it is possible to interpret the iterated random permutation problem as a special case of CBC-MAC where the iteration count r equals the number of message blocks, and all message blocks except for the first one are all-zero. The same holds for the iterated random function problem, except that a random function instead of a random permutation is used inside the CBC-MAC construction.

A first proof of the security of CBC-MAC was given by Bellare et al. in [1,2]. For CBC-MAC with a random function, they prove that the advantage of an information-theoretic adversary that makes at most q queries is upper bounded by $1.5r^2q^2/N$. Using the well-known prp-prf switching lemma [5], they derive from this an upper bound of $2r^2q^2/N$ for CBC-MAC with a random permutation. The simplicity of CBC-MAC makes it a good test case for various proof techniques. Of particular interest is the short proof of CBC-MAC by Bernstein [7]. For a more detailed proof using the same technique, we refer to Nandi [16].

In [3], Bellare et al. proved a security bound that is linear in r, instead of quadratic in r as in previous proofs. They point out that their analysis only applies to CBC-MAC with a random permutation, and not with a random function: such a bound is ruled out by an attack by Berke [6]. However, Berke's attack cannot be translated to the iterated random function problem, as the number of message blocks for each of the queries in the attack is not constant.

The iterated random function problem is similar to the nested iterated (NI) construction that Gaži et al. [14] analysed at CRYPTO 2014. However, the analysis of the NI construction critically relies on the use of two *different* random functions, or more precisely on the use of a pseudo-random function (prf) with two different keys. Our analysis applies to the case where only *one* random function is iterated. As we will show, the iterated random function problem will require a more complicated analysis of collision probabilities, in order to avoid ending up with a bound that is quadratic in r.

MAIN RESULTS. The main results of this paper are the proofs of two theorems. Theorem 1 bounds the success probability of a common class of collision adversaries, and Theorem 2 bounds the advantage of distinguishing an iterated random function from a random function. In these theorems, the function $\phi(q, r)$ is defined as

$$\phi(q,r) := 2\left(\frac{q^2\sqrt{r}}{N}\right) + 2\sqrt{\frac{q^2 r \log r}{N}} + 16\left(\frac{q^2 r \log r}{N}\right)^2 + 49(\log r)^2\left(\frac{q^2 r \log r}{N}\right).$$

3

Theorem 1. *Let f be a random function, and let $\mathcal{A}$ be a collision-finding adversary that makes q queries to f^r as follows: every query is either chosen from a set (of size $m \leq q$) of predetermined points, or is the response of a previous query. Under the assumption that $N \log r > 90$, the following bound holds for the success probability $\mathsf{cp}^r[q]$ of $\mathcal{A}$:*

$$\mathsf{cp}^r[q](\mathcal{A}) \leq \phi(q, r).$$

Theorem 2. *Let f be a random function, and let $\mathcal{A}$ be an adversary trying to distinguish f^r from f through q queries. Then, under the assumption that $N \log r > 90$, we have*

$$\mathbf{Adv}_{f,f^r}(q) \leq \frac{q^2 r}{N} + \frac{2q^2}{N} + \phi(q, r).$$

A NOTE ON THE SETTING. We should point out that our results are in an indistinguishability setting. Our goal is to distinguish, in a black-box way, between an iterated random function and a random function. In the indifferentiability setting, the adversary also has access to the underlying random function, or to a simulator that tries to mimic its behaviour. Dodis et al. [12] proved that indifferentiability for an iterated random function holds only with poor concrete security bounds, as they provide a lower bound on the complexity of any successful simulator.

OUTLINE. Notation and preliminaries are introduced in Sect. 2. We study the probabilities to find various types of collisions in a random function in Sect. 3. These results are used in Sect. 4 to bound the probabilities of single-trail attacks and two-trail collision attacks, and eventually to also bound a more general collision attack on an iterated random function. The advantage of distinguishing an iterated random function from a random function is bounded in Sect. 5. For readability, we defer the technical proof of Lemma 7 of Sect. 4 to Sect. 6. We conclude the paper in Sect. 7.

2 Notation and Preliminaries

In this section, we will state some simple lemmas without proof. The proofs of these lemmas can be found in the full version of this paper [8].

FUNCTIONS. Let $f : \mathcal{D} \to \mathcal{D}$ be a function over a domain $\mathcal{D}$ of size N. A collision for a function f is defined as a pair $(x, x') \in \mathcal{D}$ with $x \neq x'$ such that $f(x) = f(x')$. A three-way collision is a triple (x, x', x'') such that $f(x) = f(x') = f(x'')$ for distinct x, x' and x''. For a positive integer r, the r-th iterate f^r of a function f is defined inductively as follows:

$$f^1 = f,$$
$$f^r = f \circ f^{r-1}, r > 1.$$

4

By convention, let f^0 be the identity function. In the remainder of this paper, we will assume that $r \geq 2$. Let a random function denote a function that is drawn uniformly at random from the set of all functions of the same domain and range.

FALLING FACTORIAL POWERS AND THE β FUNCTION. We use the falling factorial powers notation, where for a non-negative integer $i \leq N$, $N^{\underline{i}}$ is defined as

$$N^{\underline{i}} := \frac{N!}{(N-i)!} = N(N-1)\cdots(N-i+1). \tag{1}$$

Note that $N^{\underline{i}}$ denotes the number of permutations of N items taken i at a time, or the number of ways to choose a sample of size i without replacement from a population of size N. When $i > N$, we define $N^{\underline{i}} := 0$. We also define a function $\beta(i)$ that we will frequently encounter:

$$\beta(i) := \frac{N^{\underline{i}}}{N^i}. \tag{2}$$

Again, we define $\beta(i) := 0$ for $i > N$. We derive below a simple bound on $\beta(i)$.

Lemma 1. *Let $\alpha > 0$ be a real number. Then, for $i \geq \sqrt{2\alpha N} + 1$, we have*

$$\beta(i) \leq e^{-\alpha}.$$

PARTIAL SUMS OF THE HARMONIC SERIES. The divergent infinite series

$$\sum_{i=1}^{\infty} \frac{1}{i} = 1 + \frac{1}{2} + \frac{1}{3} + \frac{1}{4} + \cdots$$

is known as the harmonic series. We will be interested in partial sums of the series of the form

$$\sum_{i=a+1}^{b} \frac{1}{i} = \frac{1}{a+1} + \frac{1}{a+2} + \cdots + \frac{1}{b-1} + \frac{1}{b}.$$

We will use the following simple bound for this sum. Throughout this paper, let log denote the natural logarithm, that is the logarithm to the base e.

Lemma 2. *For any two positive integers a and b with $b \geq a$,*

$$\sum_{i=a+1}^{b} \frac{1}{i} \leq \log\left(\frac{b}{a}\right)$$

COUNTING DIVISORS. For a positive integer a and an integer b we use the notation $a|b$ to denote a *divides* b, i.e., $ak = b$ for some integer k. We write $a \nmid b$ when a does not divide b. The number of divisors of b is denoted $\mathsf{d}(b)$. We will use the following simple bound on $\mathsf{d}(b)$.

Lemma 3. *For any positive integer b,*

$$d(b) < 2\sqrt{b}.$$

THE σ FUNCTION. The function $\sigma(b)$ defined as

$$\sigma(b) := \sum_{a|b} a$$

denotes the sum of the divisors of b. We will use the following simple lemma about $\sigma(b)$.

Lemma 4. *For any positive integer b,*

$$\sum_{a|b} \frac{b}{a} = \sigma(b).$$

A simple bound on $\sigma(b)$ can be obtained as follows.

Lemma 5. *For any positive integer $b \geq 2$,*

$$\sigma(b) < 3b \log b.$$

3   Random Function Collisions

In this section, we look at different approaches to find collisions on a random function f. We will bound their success probabilities, and use them in Sect. 4 to get bounds on the success probabilities of collision attacks on an iterated random function f^r.

3.1   Single-Trail Attack

SINGLE-TRAIL ATTACK. Let $[q]$ denote the set $\{1, \ldots, q\}$. The single-trail attack works by starting with an arbitrary initial point x and producing a *trail* of points, hoping to find a collision. A trail is uniquely defined by q queries $f^{i-1}(x)$ for $i \in [q]$, where the i-th query $f^{i-1}(x)$ has response $f^i(x)$. We assume that the attack does not stop when a collision is found, but makes q queries and then checks for collisions. If a collision is found, it will appear as a rho-shaped trail, as illustrated in Fig. 1. Therefore, a collision obtained through a single-trail attack will be called a ρ-collision.

TERMINOLOGY. Suppose the q-query single-trail attack finds a collision. For some t, c, suppose it takes $t + c$ queries to find this collision, so that

$$f^{t+c}(x) = f^t(x),$$

i.e., the output of the $(t+c)$-th query is identical to the output of the t-th query. Then, t is called the tail length of the ρ-collision, and c is called the cycle length.

Fig. 1. Single-trail attack starting from x, resulting in a ρ collision with tail length t and cycle length c. We call the probability of this collision $\mathsf{cp}_\rho(t, c)$.

For fixed t, c, we want to bound the probability that a q-query single-trail attack gives a ρ-collision on f with tail length t and cycle length c. Call this probability $\mathsf{cp}_\rho[q](t, c)$.

BOUNDING $\mathsf{cp}_\rho[q](t, c)$. To get a ρ-collision on f with tail length t and cycle length c, we need to call f at $t + c$ distinct values. Thus, if $q < t + c$, $\mathsf{cp}_\rho[q](t, c) = 0$. So suppose $q \geq t + c$. Out of these $t + c$ calls to f, the first $t + c - 1$ give distinct outputs, and the last coincides with the t-th output. Thus, the number of different ways this can happen is $N^{\underline{t+c-1}}$, out of the total N^{t+c} possible outcomes for the $t + c$ calls to f. Thus,

$$\mathsf{cp}_\rho[q](t, c) = \frac{N^{\underline{t+c-1}}}{N^{t+c}} = \frac{\beta(t + c - 1)}{N}.$$

This is just a function of t and c (since the queries made after the collision is found are of no consequence), so we will use the simpler notation $\mathsf{cp}_\rho(t, c)$, with the implicit assumption that $q \geq t + c$. For a fixed real $\alpha > 0$, when $t + c \geq \sqrt{2\alpha N} + 2$, Lemma 1 gives us the bound

$$\mathsf{cp}_\rho(t, c) \leq \frac{e^{-\alpha}}{N}. \tag{3}$$

When $t + c < \sqrt{2\alpha N} + 2$, we will simply use the bound

$$\mathsf{cp}_\rho(t, c) \leq \frac{1}{N}. \tag{4}$$

3.2 Two-Trail Attack

TWO-TRAIL ATTACK. In the two-trail attack, we start with two different points x_1 and x_2, and produce two trails: the trail $f^{i-1}(x_1)$ for $i \in [q_1]$, and the trail $f^{i-1}(x_2)$ for $i \in [q_2]$, hoping to find a collision. In total $q_1 + q_2$ queries are made, where the i-th query for $i \in [q_1]$ is $f^{i-1}(x_1)$, with response $f^i(x_1)$, and the

7

$(q_1 + i)$-th query for $i \in [q_2]$ is $f^{i-1}(x_2)$, with response $f^i(x_2)$. If a collision is found, the two trails will form a lambda shape, as illustrated in Fig. 2. Therefore, a collision obtained through a two-trail attack will be called a λ-collision.


Fig. 2. Two-trail attack starting from x_1 and x_2, resulting in a λ-collision with foot lengths t_1 and t_2, respectively. We call the probability of this collision $\mathsf{cp}_\lambda(t_1, t_2)$.

TERMINOLOGY. Suppose the (q_1, q_2)-query two-trail attack finds a λ-collision, regardless of whether a ρ-collisions has occurred on either trail. Suppose that a λ-collision is found after making t_1 queries along the first trail and t_2 queries along the second, i.e.,

$$f^{t_1}(x_1) = f^{t_2}(x_2).$$

t_1 and t_2 are called the foot lengths of the λ-collision. For fixed t_1, t_2, we want to bound the probability that a (q_1, q_2)-query two-trail attack finds a λ-collision with foot lengths t_1 and t_2. Denote this probability as $\mathsf{cp}_\lambda[q_1, q_2](t_1, t_2)$.

BOUNDING $\mathsf{cp}_\lambda[q_1, q_2](t_1, t_2)$. To get a λ-collision on f with foot lengths t_1 and t_2, we need to call f at t_1 distinct values on the first trail and t_2 distinct values on the second trail. Thus, if $q_1 < t_1$ or $q_2 < t_2$, $\mathsf{cp}_\lambda[q_1, q_2](t_1, t_2) = 0$. So we assume $q_1 \geq t_1$ and $q_2 \geq t_2$. Out of these $t_1 + t_2$ queries, the first $t_1 - 1$ in one trail and the first $t_2 - 1$ in the other trail give distinct outputs, and the last calls on the two trails coincide on a value distinct from all the earlier ones, i.e., the $t_1 + t_2$ calls lead to $t_1 + t_2 - 1$ distinct outputs, and one collision. Thus, the number of different ways this can happen is $N^{\underline{t_1+t_2-1}}$, out of the total $N^{t_1+t_2}$ possible outcomes for the $t_1 + t_2$ calls to f. Thus,

$$\mathsf{cp}_\lambda[q_1, q_2](t_1, t_2) = \frac{N^{\underline{t_1+t_2-1}}}{N^{t_1+t_2}} = \frac{\beta(t_1 + t_2 - 1)}{N}.$$

Again, this is only a function of t_1 and t_2 (since the queries made after the collision is found are of no consequence), so we will use the simpler notation $\mathsf{cp}_\lambda(t_1, t_2)$, with the implicit assumption that $q_1 \geq t_1$ and $q_2 \geq t_2$. For our purposes it will be enough to use the bound

$$\mathsf{cp}_\lambda(t_1, t_2) \leq \frac{1}{N}. \tag{5}$$

3.3 A $\lambda\rho$-Double-Collision on a Two-Trail Attack

When a two-trail attack leads to two collisions, a double-collision is said to occur. In Sect. 4, in addition to the above bounds, we also need a bound on the probability of two closely related double-collisions. We deal with a $\lambda\rho$-double-collision in this section, and a ρ'-double-collision in the next. A $\lambda\rho$-double-collision takes place when a two-trail attack leads to a λ-collision, and then the combined trail becomes the tail of a ρ-collision, as shown in Fig 3.[5]


Fig. 3. Two-trail attack starting from x_1 and x_2, resulting in a $\lambda\rho$-collision. First, there is a λ-collision with foot lengths t_1 and t_2, respectively. Then, the combined trail continues for Δt queries, and completes a cycle of length c, after which a ρ-collision occurs. We call the probability of this double-collision $\mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c)$.

TERMINOLOGY. We assign four parameters to this collision: the foot lengths t_1 and t_2 of the λ, the intervening length Δt between the two collisions, and the cycle length c of the ρ. Note that Δt can be seen as the tail length of the ρ-collision if we imagine it to have resulted from a single-trail attack beginning at the point of the λ-collision. For fixed $t_1, t_2, \Delta t, c$ we want to find the probability that a (q_1, q_2)-query two-trail attack finds a $\lambda\rho$-double-collision with foot lengths t_1 and t_2, intervening length Δt and cycle length c. Call this probability $\mathsf{cp}_{\lambda\rho}[q_1, q_2](t_1, t_2, \Delta t, c)$.

BOUNDING $\mathsf{cp}_{\lambda\rho}[q_1, q_2](t_1, t_2, \Delta t, c)$. To get a λ-collision on f with foot lengths t_1 and t_2, we need to call f at t_1 distinct values on the first trail, and t_2 distinct values on the second trail; and to get a ρ-collision on f with tail length Δt and cycle length c, we need to call f at Δt common values on each trail, and a further

---

[5] Note that we only call it a double-collision if both trails continue up to the point of second collision.

c points on the first trail; this adds up to $t_1 + t_2 + \Delta t + c$ distinct values in all. Thus, when $q_1 < t_1 + \Delta t + c$ or $q_2 < t_2 + \Delta t$, $\mathsf{cp}_{\lambda\rho}[q_1, q_2](t_1, t_2, \Delta t, c) = 0$. So we assume $q_1 \geq t_1 + \Delta t + c$ and $q_2 \geq t_2 + \Delta t$. These $t_1 + t_2 + \Delta t + c$ calls lead to $t_1 + t_2 + \Delta t + c - 2$ distinct outputs, and two collisions. Thus, the number of different ways this can happen is $N^{\underline{t_1+t_2+\Delta t+c-2}}$, out of the total $N^{t_1+t_2+\Delta t+c}$ possible outcomes for the $t_1 + t_2 + \Delta t + c$ calls to f. Thus,

$$\mathsf{cp}_{\lambda\rho}[q_1, q_2](t_1, t_2, \Delta t, c) = \frac{N^{\underline{t_1+t_2+\Delta t+c-2}}}{N^{t_1+t_2+\Delta t+c}} = \frac{\beta(t_1 + t_2 + \Delta t + c - 2)}{N^2}.$$

As before, this is only a function of $t_1, t_2, \Delta t$ and c (since the queries made after the ρ collision is found are of no consequence), so we use the simpler notation $\mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c)$, with the implicit assumption that $q_1 \geq t_1 + \Delta t + c$ and $q_2 \geq t_2 + \Delta t$. For a fixed real $\alpha > 0$, when $t_1 + t_2 + \Delta t + c \geq \sqrt{2\alpha N} + 3$, Lemma 1 gives us the bound

$$\mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c) \leq \frac{e^{-\alpha}}{N^2}. \tag{6}$$

When $t_1 + t_2 + \Delta t + c < \sqrt{2\alpha N} + 3$, we will simply use the bound

$$\mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c) \leq \frac{1}{N^2}. \tag{7}$$

3.4 A ρ'-Double-Collision on a Two-Trail Attack

A ρ'-double-collision takes place when a two-trail attack leads to a ρ with two tails. This is shown in Figure 4. We will allow $\Delta t = 0$, in which case a three-way collision occurs.

TERMINOLOGY. As before, we assign four parameters to this collision: the tail lengths t_1 and t_2 of the ρ, the intervening length Δt between the two collisions, and the cycle length c of the ρ. For fixed $t_1, t_2, \Delta t, c$ we want to find the probability that a two-trail attack with sufficiently many queries finds a ρ'-double-collision with tail lengths t_1 and t_2, intervening length Δt, and cycle length c. Call this probability $\mathsf{cp}_{\rho'}[q_1, q_2](t_1, t_2, \Delta t, c)$.

BOUNDING $\mathsf{cp}_{\rho'}[q_1, q_2](t_1, t_2, \Delta t, c)$. The bounding of $\mathsf{cp}_{\rho'}[q_1, q_2](t_1, t_2, \Delta t, c)$ is almost identical to that of $\mathsf{cp}_{\lambda\rho}[q_1, q_2](t_1, t_2, \Delta t, c)$. To get a ρ'-double-collision with tail lengths t_1 and t_2, intervening length Δt, and cycle length c, we need to call f at $t_1 + c - \Delta t$ distinct values on the first trail, t_2 distinct values on the second trail, and Δt common values on each trail, resulting in calls at $t_1 + t_2 + c$ distinct values in all. Thus, when $q_1 < t_1 + c$ or $q_2 < t_2 + \Delta t$, $\mathsf{cp}_{\rho'}[q_1, q_2](t_1, t_2, \Delta t, c) = 0$. So we assume $q_1 \geq t_1 + c$ and $q_2 \geq t_2 + \Delta t$. These $t_1 + t_2 + c$ calls lead to $t_1 + t_2 + c - 2$ distinct outputs. Thus, the number of different ways this can happen is $N^{\underline{t_1+t_2+c-2}}$, out of the total $N^{t_1+t_2+c}$ possible outcomes for the $t_1 + t_2 + c$ calls to f. Thus,

$$\mathsf{cp}_{\rho'}[q_1, q_2](t_1, t_2, \Delta t, c) = \frac{N^{\underline{t_1+t_2+c-2}}}{N^{t_1+t_2+c}} = \frac{\beta(t_1 + t_2 + c - 2)}{N^2}.$$

Fig. 4. Two-trail attack starting from x_1 and x_2, resulting in a ρ'-collision with tail lengths t_1 and t_2, intervening length Δt, and cycle length c. We will allow $\Delta t = 0$, in which case a three-way collision occurs. We call the probability of this double-collision $\mathsf{cp}_{\rho'}(t_1, t_2, \Delta t, c)$.

As before, this is only a function of $t_1, t_2, \Delta t$ and c (since the queries made after the ρ collision is found are of no consequence), so we use the simpler notation $\mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c)$, with the implicit assumption that $q_1 \geq t_1 + \Delta t + c$ and $q_2 \geq t_1 + \Delta t$. Recalling that

$$\mathsf{cp}_{\lambda\rho}(t_1, t_2, 0, c) = \frac{\beta(t_1 + t_2 + c - 2)}{N^2},$$

we conclude that

$$\mathsf{cp}_{\rho'}(t_1, t_2, \Delta t, c) = \mathsf{cp}_{\lambda\rho}(t_1, t_2, 0, c). \tag{8}$$

4   Iterated Random Function Collisions

In this section we revisit the two types of collision attacks described in Sect. 3, and analyse their success probabilities when applied to f^r. The main proof in this paper relies heavily on the results obtained in this section.

A CAUTIONARY NOTE. At first glance, this section may appear to be similarly organised as Sect. 3. It is important to keep in mind that we are now interested in something entirely different. In Sect. 3, we looked at the probabilities of specific ρ- and λ-collisions with fixed parameters. In this section, instead, we focus on the probabilities that single-trail attacks and two-trail attacks of some specified number of queries succeed in finding collisions on f^r. By reducing these collisions to collisions on f, we can use the union bound on the bounds obtained in Sect. 3 to get the desired bounds. To distinguish from the collision probabilities on f, which we denoted $\mathsf{cp}[\cdot]$, we now use the notation $\mathsf{cp}^r[\cdot]$ for the collision probabilities on f^r.

4.1 Single-Trail Attack

We want to bound the probability that a q-query single-trail attack finds a collision on f^r. Call this probability $\mathsf{cp}_\rho^r[q]$.

REDUCING TO COLLISION ON f. Suppose the q-query single-trail attack finds a ρ-collision on f^r with tail length t' and cycle length c'. Observe that this collision necessarily arises out of a ρ-collision on f, with tail length t and cycle length c for some t, c. This can happen in two ways:

– DIRECT COLLISION. This happens when r divides c. Then, define k such that rk is the first multiple of r that is not less than t, i.e.,

$$k := \left\lceil \frac{t}{r} \right\rceil,$$

then $rk + c$ is also a multiple of r, and since $f^{t+c}(x) = f^t(x)$, and $rk \geq t$, we also have

$$f^{rk+c}(x) = f^{rk}(x).$$

Writing

$$k' = \frac{c}{r},$$

we have

$$(f^r)^{k+k'}(x) = (f^r)^k(x),$$

our ρ-collision on f^r. Note that according to this notation,

$$t' = k = \left\lceil \frac{t}{r} \right\rceil, c' = k' = \frac{c}{r}.$$

Loosely speaking, in a direct collision, the first collision on f arrives *in phase* with r, i.e.,

$$t = t + c \bmod r,$$

so that this first collision on f leads immediately to a collision on f^r at the next multiple of r.

– DELAYED COLLISION. A *delayed collision* occurs when r does not divide c, i.e., the first collision arrives *out of phase*. Then we need to keep cycling about the ρ of f till the phase is adjusted, and only then we arrive at the next multiple of r and find a collision on f^r. Suppose it cycles around η times. For the phase to be adjusted, $c\eta$ should be a multiple of r. The smallest value of η that satisfies this is

$$\eta = \frac{r}{d}$$

where $d = \mathsf{gcd}(c, r)$ is the greatest common divisor of c and r. Let $k = \left\lceil \frac{t}{r} \right\rceil$ as before, and let

$$k' = \frac{c}{d}.$$

As before, since we have $f^{t+c\eta}(x) = f^t(x)$, and $rk \geq t$, we have

$$f^{rk+c\eta}(x) = f^{rk}(x),$$

which gives us the ρ-collision

$$(f^r)^{k+k'}(x) = (f^r)^k(x),$$

as before. Again, according to this notation,

$$t' = k = \left\lceil \frac{t}{r} \right\rceil, c' = k' = \frac{c}{d}.$$

REQUIRED CONDITIONS. Observing that a direct collision can be seen as a special case of delayed collision, where $d = \gcd(c, r) = r$, we can summarise the above as follows: a ρ-collision on f with tail length t and cycle length c *eventually* leads to a ρ-collision on f^r with tail length t' and cycle length c' where

$$t' = k = \left\lceil \frac{t}{r} \right\rceil, c' = k' = \frac{c}{d},$$

with $d = \gcd(c, r)$ as before. Thus, for a ρ-collision on f to result in a ρ-collision on f^r, the only required condition is that q is sufficiently large, i.e.,

$$t' + c' \leq q.$$

In terms of t and c, this becomes

$$\left\lceil \frac{t}{r} \right\rceil + \frac{c}{d} \leq q.$$

Recall that we are trying to bound the probability $\mathsf{cp}_\rho^r[q]$ of finding a ρ-collision on f^r in q queries. This is equivalent to the probability of finding a ρ-collision on f with the parameters t and c satisfying the above condition. Recall that in Sect. 3, we bounded this probability for a fixed (t, c), which we called $\mathsf{cp}_\rho(t, c)$. We can now use the union bound to get a bound on $\mathsf{cp}_\rho^r[q]$.

USING THE UNION BOUND ON $\mathsf{cp}_\rho^r[q]$. Let $\mathcal{S}$ be the set of (t, c) values that satisfy the requirement

$$\left\lceil \frac{t}{r} \right\rceil + \frac{c}{\gcd(c, r)} \leq q.$$

For a fixed $\alpha > 0$, we can split $\mathcal{S}$ into two parts:

$$\mathcal{S}^+[\alpha] := \left\{ (t, c) \in \mathcal{S} \mid t + c \geq \sqrt{2\alpha N} + 2 \right\},$$
$$\mathcal{S}^-[\alpha] := \left\{ (t, c) \in \mathcal{S} \mid t + c < \sqrt{2\alpha N} + 2 \right\}.$$

13

Applying the union bound with bounds (3) and (4) obtained for $\mathsf{cp}_\rho(t,c)$ gives

$$
\begin{aligned}
\mathsf{cp}_\rho^r[q] &\leq \sum_{\mathcal{S}} \mathsf{cp}_\rho(t,c) \\
&= \sum_{\mathcal{S}^+[\alpha]} \mathsf{cp}_\rho(t,c) + \sum_{\mathcal{S}^-[\alpha]} \mathsf{cp}_\rho(t,c) \\
&\leq \sum_{\mathcal{S}^+[\alpha]} \frac{e^{-\alpha}}{N} + \sum_{\mathcal{S}^-[\alpha]} \frac{1}{N} \\
&= \#\mathcal{S}^+[\alpha] \cdot \frac{e^{-\alpha}}{N} + \#\mathcal{S}^-[\alpha] \cdot \frac{1}{N}
\end{aligned}
\tag{9}
$$

BOUNDING $\#\mathcal{S}^-[\alpha]$. We observe that whenever $(t,c) \in \mathcal{S}^-[\alpha]$,

$$
t < \sqrt{2\alpha N} + 2,
$$

and

$$
c < q \cdot \gcd(c,r).
$$

If we count the number of (t,c) satisfying these conditions, it will give us an upper bound on $\#\mathcal{S}^-[\alpha]$. There are at most $\sqrt{2\alpha N} + 2$ values of t satisfying $t < \sqrt{2\alpha N} + 2$. For a fixed $d = \gcd(c,r)$, c has to be a multiple of d not exceeding qd. The number of such values of c is q. Since d must be a factor of r, we get the total number of values of c satisfying $c < q \cdot \gcd(c,r)$ to be at most $q \cdot \mathsf{d}(r)$. Putting it all together we get

$$
\#\mathcal{S}^-[\alpha] \leq (\sqrt{2\alpha N} + 2) \cdot q \cdot \mathsf{d}(r).
\tag{10}
$$

BOUNDING $\#\mathcal{S}^+[\alpha]$. For $(t,c) \in \mathcal{S}^+[\alpha]$, it will be enough for our purposes to consider the bounds

$$
t \leq qr,
$$

and

$$
c < q \cdot \gcd(c,r).
$$

Using the same reasoning as before, the number of values of c that satisfy $c < q \cdot \gcd(c,r)$ is at most $q \cdot \mathsf{d}(r)$. For t there are now at most qr values. Thus, we obtain the bound

$$
\#\mathcal{S}^+[\alpha] \leq q^2 r \cdot \mathsf{d}(r).
\tag{11}
$$

FINAL BOUND FOR $\mathsf{cp}_\rho^r[q]$. We can now plug (10) and (11) into (9):

$$
\begin{aligned}
\mathsf{cp}_\rho^r[q] &\leq \#\mathcal{S}^+[\alpha] \cdot \frac{e^{-\alpha}}{N} + \#\mathcal{S}^-[\alpha] \cdot \frac{1}{N} \\
&\leq q^2 r \cdot \mathsf{d}(r) \cdot \frac{e^{-\alpha}}{N} + (\sqrt{2\alpha N} + 2) \cdot q \cdot \mathsf{d}(r) \cdot \frac{1}{N}
\end{aligned}
$$

for any real $\alpha > 0$. We will simplify it by plugging in a suitable value of α.

SIMPLIFYING THE BOUND. We know from Lemma 3 that

$$\mathsf{d}(r) < 2\sqrt{r}.$$

We put $\alpha = \log r$. Then we have

$$\sqrt{2\alpha N} = \sqrt{2N \log r},$$

and

$$e^{-\alpha} = \frac{1}{r}.$$

When $N \log r \geq 16$, we have

$$\begin{aligned}
\sqrt{2\alpha N} + 2 &= \sqrt{2N \log r} + 2 \\
&= 2\sqrt{N \log r} - \left[(2 - \sqrt{2}) \cdot \sqrt{N \log r} - 2 \right] \\
&\leq 2\sqrt{N \log r} - \left[(2 - \sqrt{2}) \cdot 4 - 2 \right] \\
&= 2\sqrt{N \log r} - \left[6 - \sqrt{2} \right] \\
&< 2\sqrt{N \log r}.
\end{aligned}$$

Thus,

$$\mathsf{cp}_\rho^r[q] \leq 2 \cdot \left(\frac{q^2 \sqrt{r}}{N} \right) + 2 \cdot \sqrt{\frac{q^2 r \log r}{N}}.$$

This gives us a bound for the success probability of a q-query single-trail attack on f^r. We state the result as a lemma.

Lemma 6. *Under the assumption that $N \log r \geq 16$, we have*

$$cp_\rho^r[q] \leq 2 \cdot \left(\frac{q^2 \sqrt{r}}{N} \right) + 2 \cdot \sqrt{\frac{q^2 r \log r}{N}}.$$

4.2   Two-Trail Attack

We want to bound the probability that a (q_1, q_2)-query two-trail attack finds a λ-collision on f^r. Call this probability $\mathsf{cp}_\lambda^r[q_1, q_2]$.

REDUCING TO COLLISION ON f. Suppose the (q_1, q_2)-query two-trail attack finds a λ-collision on f^r with foot lengths t_1' and t_2'. As in the case of the ρ-collision on f^r, this can only arise from a λ-collision on f, say with foot lengths t_1 and t_2, which can again happen in two ways:

–   DIRECT COLLISION. A direct collision takes place when the two f-trails collide in phase, i.e.,

$$t_1 = t_2 \bmod r.$$

15

When this happens, the two trails continue till the next multiple of r, where they give a λ-collision on f^r. This collision takes place at

$$t_1' = \left\lceil \frac{t_1}{r} \right\rceil, t_2' = \left\lceil \frac{t_2}{r} \right\rceil.$$

– DELAYED COLLISION. A delayed collision takes place when the two f-trails collide out of phase, i.e.,

$$t_1 \neq t_2 \text{ mod } r.$$

If one of the trails results in a ρ-collision on f^r, this implies that a successful single-trail attack has been carried out on f^r. Here, we will only focus on the scenario where a λ-collision on f^r can still happen. But then one of the two f-trails must have entered into a cycle, otherwise both f-trails will remain out of phase. This can only happen in one of two ways:

- After the λ-collision on f, the combined trail forms the tail of a ρ collision on f, that is, they form a $\lambda\rho$-collision on f as in Fig. 3. One of the trails, say the one from x_1, cycles around the ρ enough number of times to adjust the phase, and then the two f-trails continue to the next multiple of r, giving a λ-collision on f^r;[6]

- After the λ-collision on f, one of the two f-trails, say the one from x_1, continues and collides with the trail from x_2, that is, they form a ρ'-collision on f as in Fig. 4. When $\Delta t = 0$, a three-way collision on f occurs. The trail from x_1 cycles around the ρ enough number of times to adjust the phase, giving a λ-collision on f^r.

In our calculations, we assume that it is the trail from x_1 that cycles multiple times, while the one from x_2 waits for the collision on f^r to happen. We obtain a bound which is symmetric over q_1 and q_2, and thus also holds for the case when the two trails reverse roles. Let τ_1 and τ_2 be the respective lengths of the two trails till *the point of waiting*, i.e., the point of ρ-collision of the trail from x_1. Calling Δt the distance between the two collision points, we simply have

$$\tau_1 = t_1 + \Delta t, \tau_2 = t_2 + \Delta t$$

for the $\lambda\rho$-collision, and

$$\tau_1 = t_1, \tau_2 = t_2 + \Delta t$$

for the ρ'-collision. Let the cycle length of this ρ be c (note that its tail length is τ_1 with respect to this trail). Suppose this trail cycles η times about the ρ in order to adjust the phase difference. Then η is the smallest number that satisfies

$$\tau_1 + c\eta = \tau_2 \text{ mod } r.$$

---

[6] This is indeed a (delayed) λ-collision on f^r: from the point of view of f^r, neither of the two trails could be seen to enter into a cycle.

Suppose k is such that

$$\tau_1 + c\eta = \tau_2 + rk.$$

Also, let

$$k_2 = \left\lceil \frac{\tau_2}{r} \right\rceil.$$

From our definition of τ_1 and τ_2, we have that

$$f^{\tau_1}(x_1) = f^{\tau_2}(x_2),$$

and from the ρ-collision $f^{\tau_1+c}(x_1) = f^{\tau_1}(x_1)$, it follows that

$$f^{\tau_1+c\eta}(x_1) = f^{\tau_1}(x_1).$$

From these two we get

$$f^{\tau_1+c\eta}(x_1) = f^{\tau_2}(x_2).$$

From the definition of k we have

$$f^{\tau_2+rk}(x_1) = f^{\tau_2}(x_2).$$

Continuing on to rk_2, we get a λ-collision on f^r as

$$(f^r)^{k+k_2}(x_1) = (f^r)^{k_2}(x_2).$$

According to this notation we have a λ-collision on f^r with foot lengths t'_1 and t'_2, such that

$$t'_1 = k + k_2 = \left\lceil \frac{\tau_1 + c\eta}{r} \right\rceil, t'_2 = k_2 = \left\lceil \frac{\tau_2}{r} \right\rceil.$$

When this comes from a $\lambda\rho$-collision, we have

$$t'_1 = \left\lceil \frac{t_1 + \Delta t + c\eta}{r} \right\rceil, t'_2 = \left\lceil \frac{t_2 + \Delta t}{r} \right\rceil.$$

When this comes from a ρ'-collision, we have

$$t'_1 = \left\lceil \frac{t_1 + c\eta}{r} \right\rceil, t'_2 = \left\lceil \frac{t_2 + \Delta t}{r} \right\rceil.$$

We will treat these two cases separately, even though they are closely related.

REQUIRED CONDITIONS. Again, we observe that the direct collision is a special case of the delayed collision with $\Delta t = 0$ and $\eta = 0$. *However, there is an important difference.* For the delayed λ-collision, we require *two* collisions on f, unlike all other collisions we have seen so far, which need only one. This case corresponds to the $\lambda\rho$-double-collision and the ρ'-double-collision from Sect. 3,

and requires some special treatment, as we will see in the course of our calculations. The condition needed here is that both trails continue long enough for the collision to happen, i.e.,

$$t_1' \leq q_1, t_2' \leq q_2.$$

In terms of $t_1, t_2, \Delta t, c, \eta$, this translates to

$$\left\lceil \frac{t_1 + \Delta t + c\eta}{r} \right\rceil \leq q_1, \left\lceil \frac{t_2 + \Delta t}{r} \right\rceil \leq q_2$$

for the $\lambda\rho$-double-collision and

$$\left\lceil \frac{t_1 + c\eta}{r} \right\rceil \leq q_1, \left\lceil \frac{t_2 + \Delta t}{r} \right\rceil \leq q_2$$

for the ρ'-double-collision. Recall that we are trying to calculate $\mathsf{cp}_\lambda^r[q_1, q_2]$, the probability of getting a λ-collision on f^r with a (q_1, q_2)-query two-trail attack starting from x_1 and x_2. Based on our observations above, this can happen in two ways:

- A DIRECT λ-COLLISION ON f. This is the direct collision scenario, where the collision is in phase. The foot lengths t_1 and t_2 have the constraints

$$\left\lceil \frac{t_1}{r} \right\rceil \leq q_1, \left\lceil \frac{t_2}{r} \right\rceil \leq q_2, t_1 = t_2 \bmod r.$$

For fixed t_1, t_2, we recall that the probability of this collision is $\mathsf{cp}_\lambda(t_1, t_2)$.

- A $\lambda\rho$-DOUBLE-COLLISION ON f. This is the first case of the delayed collision scenario, where the collision is out of phase. Here, t_1 and t_2 are the foot lengths of the λ, Δt is the distance between the two collision points, c is the cycle length of the ρ, and η is the number of cycles necessary around the ρ. Recall that one of the trails circles around the ρ, while the other waits for the λ-collision on f^r to happen. We continue with our assumption that the one from x_1 does the cycling and the one from x_2 waits, since we will eventually count over all pairs of trails. Now $t_1, t_2, \Delta t, c, \eta$ have the constraints

$$\left\lceil \frac{t_1 + \Delta t + c\eta}{r} \right\rceil \leq q_1, \left\lceil \frac{t_2 + \Delta t}{r} \right\rceil \leq q_2, t_1 + c\eta = t_2 \bmod r.$$

For fixed $t_1, t_2, \Delta t, c, \eta$, we recall that the probability of this $\lambda\rho$-double-collision is $\mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c)$.

- A ρ'-DOUBLE-COLLISION ON f. This is the second case of the delayed collision scenario. Here, t_1 and t_2 are the lengths of the two tails of the ρ, Δt is the distance between the two collision points, c is the cycle length of the ρ, and η is the number of cycles necessary around the ρ. Again, the trail from x_1 circles around the ρ, while the trail from x_2 waits for the λ-collision on f^r to happen. Thus, $t_1, t_2, \Delta t, c, \eta$ have the constraints

$$\left\lceil \frac{t_1 + c\eta}{r} \right\rceil \leq q_1, \left\lceil \frac{t_2 + \Delta t}{r} \right\rceil \leq q_2, t_1 + c\eta = t_2 + \Delta t \bmod r.$$

Our strategy for bounding $\mathsf{cp}_\lambda^r[q_1, q_2]$ will be similar to the one we used for bounding $\mathsf{cp}_\rho^r[q]$: to take the bounds on $\mathsf{cp}_\lambda(t_1, t_2)$ for fixed t_1, t_2, $\mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c)$ for fixed $t_1, t_2, \Delta t, c$ and $\mathsf{cp}_{\rho'}(t_1, t_2, \Delta t, c)$ for fixed $t_1, t_2, \Delta t, c$ obtained in Sect. 3, and then use the union bound over all possible values these parameters can take.

APPLYING THE UNION BOUND TO $\mathsf{cp}_\lambda^r[q_1, q_2]$. Let $\mathcal{S}_1$ be the set of (t_1, t_2) values that satisfy the constraints

$$\left\lceil \frac{t_1}{r} \right\rceil \le q_1, \left\lceil \frac{t_2}{r} \right\rceil \le q_2, t_1 = t_2 \bmod r,$$

and let

$$\mathsf{p}_1 := \sum_{\mathcal{S}_1} \mathsf{cp}_\lambda(t_1, t_2).$$

Let $\mathcal{S}_2$ be the set of $(t_1, t_2, \Delta t, c, \eta)$ values that satisfy the constraints

$$\left\lceil \frac{t_1 + \Delta t + c\eta}{r} \right\rceil \le q_1, \left\lceil \frac{t_2 + \Delta t}{r} \right\rceil \le q_2, t_1 + c\eta = t_2 \bmod r,$$

and let

$$\mathsf{p}_2 := \sum_{\mathcal{S}_2} \mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c).$$

Let $\mathcal{S}_3$ be the set of $(t_1, t_2, \Delta t, c, \eta)$ values that satisfy the constraints

$$\left\lceil \frac{t_1 + c\eta}{r} \right\rceil \le q_1, \left\lceil \frac{t_2 + \Delta t}{r} \right\rceil \le q_2, t_1 + c\eta = t_2 + \Delta t \bmod r,$$

and let

$$\mathsf{p}_3 := \sum_{\mathcal{S}_3} \mathsf{cp}_{\rho'}(t_1, t_2, \Delta t, c).$$

In addition, for the case where the trails reverse roles, we define $\mathcal{S}_4$ as the set of $(t_1, t_2, \Delta t, c, \eta)$ values that satisfy the constraints

$$\left\lceil \frac{t_1 + \Delta t}{r} \right\rceil \le q_1, \left\lceil \frac{t_2 + \Delta t + c\eta}{r} \right\rceil \le q_2, t_1 = t_2 + c\eta \bmod r,$$

and

$$\mathsf{p}_4 := \sum_{\mathcal{S}_4} \mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c).$$

Similarly, we define $\mathcal{S}_5$ as the set of $(t_1, t_2, \Delta t, c, \eta)$ values that satisfy the constraints

$$\left\lceil \frac{t_1 + \Delta t}{r} \right\rceil \le q_1, \left\lceil \frac{t_2 + c\eta}{r} \right\rceil \le q_2, t_1 + \Delta t = t_2 + c\eta \bmod r,$$

and

$$\mathsf{p}_5 := \sum_{\mathcal{S}_5} \mathsf{cp}_{\rho'}(t_1, t_2, \Delta t, c).$$

We state here the following bounds on $\mathsf{p}_1, \mathsf{p}_2, \mathsf{p}_3$, the proof of which we defer to Sect. 6:

Lemma 7. *Under the assumption that $N \log r > 90$,*

$$\mathsf{p}_1 \leq \frac{q_1 q_2 r}{N},$$

$$\mathsf{p}_2 \leq 8 \cdot (\log r)^2 \cdot \left(\frac{q_1 q_2 r}{N}\right)^2 + 24 \cdot (\log r)^3 \cdot \left(\frac{q_1 q_2 r}{N}\right),$$

$$\mathsf{p}_3 \leq 8 \cdot (\log r)^2 \cdot \left(\frac{q_1 q_2 r}{N}\right)^2 + 24 \cdot (\log r)^3 \cdot \left(\frac{q_1 q_2 r}{N}\right).$$

FINAL BOUND FOR $\mathsf{cp}_\lambda^r[q_1, q_2]$. We observe that the bounds for p_2 and p_3 in Lemma 7 are symmetric over q_1 and q_2. Thus, we have

$$\mathsf{p}_4 \leq 8 \cdot (\log r)^2 \cdot \left(\frac{q_1 q_2 r}{N}\right)^2 + 24 \cdot (\log r)^3 \cdot \left(\frac{q_1 q_2 r}{N}\right),$$

$$\mathsf{p}_5 \leq 8 \cdot (\log r)^2 \cdot \left(\frac{q_1 q_2 r}{N}\right)^2 + 24 \cdot (\log r)^3 \cdot \left(\frac{q_1 q_2 r}{N}\right).$$

Using the union bound, we get

$$\mathsf{cp}_\lambda^r[q_1, q_2] \leq \mathsf{p}_1 + \mathsf{p}_2 + \mathsf{p}_3 + \mathsf{p}_4 + \mathsf{p}_5.$$

This gives us the required bound, which we state next in the form of a lemma.

Lemma 8. *When $N \log r > 90$,*

$$\mathsf{cp}_\lambda^r[q_1, q_2] \leq 32 \cdot \left(\frac{q_1 q_2 r \log r}{N}\right)^2 + 97 \cdot (\log r)^2 \cdot \left(\frac{q_1 q_2 r \log r}{N}\right).$$

Proof. As $r \geq 2$, we can relax the bound of p_1 as

$$\mathsf{p}_1 \leq \frac{q_1 q_2 r}{N} \leq \frac{q_1 q_2 r}{N} \cdot (\log r)^3.$$

The rest follows from Lemma 7. □

4.3 A More General Collision Attack

Previously, we looked at two main approaches for a collision attack: the single-trail attack and the two-trail attack, and we bounded their success probabilities. Now, we will bound the success probability of a more general collision attack. More specifically, we consider collision attack subject to the restriction that is given in the statement of Theorem 1 in Sect. 1: every query is either chosen from a set of size m (with $m \leq q$) of predetermined starting points, or is the response of a previous query. First, let us introduce the notion of a transcript.

TRANSCRIPT. Let us consider any adversary $\mathcal{A}$ that interacts with an oracle $\mathcal{O}$. This interaction can be represented as a transcript, that is, as a list of queries made and answers returned. Let the transcript tr be defined as the q-tuple of input-output pairs $\mathsf{tr} = ((x_1, y_1), (x_2, y_2), \ldots, (x_q, y_q))$. Without loss of generality, we do not consider adversaries here that repeat the same query, i.e., all q queries are distinct.

SOURCES AND TRAILS. For $j, j' \in [q], j \neq j'$, we say that $x_{j'}$ is a *predecessor* of x_j if

$$f(x_{j'}) = x_j.$$

We call x_j a *source* if it does not have a predecessor. If there exists a non-empty subset of the queries for which every query has a predecessor that is in the same subset, and no query has a predecessor outside the set, we call this subset a *permutation cycle*. Note that a permutation cycle forms a rho-shape with a tail of length zero. For a permutation cycle, we define the query x_j of the permutation cycle with the smallest index j to be a *source*.

Suppose that there are m sources along the q queries, which we call $z_1, \ldots, z_m$. Then we can see the attack as an m-trail attack, with the m trails starting from $z_1, \ldots, z_m$ and of lengths $q_1, \ldots, q_m$ respectively. Thus, each point that is not a source must be on one of these m trails.

If the collision attack is successful, then for some $i, i' \in [q]$ with $i \neq i'$, we have

$$f(x_i) = f(x_{i'}).$$

In that case, one of the following must hold:

- x_i and $x_{i'}$ are on the same trail, say the one from z_p – in this case, a successful q_p-query single-trail attack starting from z_p has occurred;

- x_i and $x_{i'}$ are on different trails, say the ones from z_p and $z_{p'}$ respectively – in this case, a successful $(q_p, q_{p'})$-query two-trail attack starting from $(z_p, z_{p'})$ has occurred.

A WORD ON THE CHOICE OF $q_1, \ldots, q_m$. We note here that since we are allowing the trails to collide and merge with each other, the trail lengths $q_1, \ldots, q_m$ are not necessarily unique, since the queries on the merged trail can be counted on either trail, or both. We can get around this by choosing to count each merged trail as part of any one of the pre-merging trails, while the other is thought to stop at the point of collision. This way, we ensure that $\sum_{j=1}^{m} q_j = q$.

To bound the success probability of this more general collision attack, we can use the previously obtained bounds on the success probabilities of single-trail attacks and two-trail attacks along with the union bound. With notation as above we recall the following bounds:

- SINGLE-TRAIL ATTACK. For a q-query single-trail attack, Lemma 6 gives us the bound

$$\mathsf{cp}_\rho^r[q] \leq 2 \cdot \left(\frac{q^2 \sqrt{r}}{N} \right) + 2 \cdot \sqrt{\frac{q^2 r \log r}{N}}.$$

- TWO-TRAIL ATTACK. For a (q_1, q_2)-query two-trail attack, Lemma 8 gives us the bound

$$\mathsf{cp}_\lambda^r[q_1, q_2] \leq 32 \cdot \left(\frac{q_1 q_2 r \log r}{N} \right)^2 + 97 \cdot (\log r)^2 \cdot \left(\frac{q_1 q_2 r \log r}{N} \right).$$

Let $\mathsf{cp}^r[q](\mathcal{A})$ denote the probability that the collision adversary $\mathcal{A}$ making q queries finds a collision on f^r. For $q_1, \ldots, q_m$, with

$$\sum_{i=1}^{m} q_i = q,$$

and let $\mathsf{cp}^r[q](q_1, \ldots, q_m)$ denote the probability that a collision attack with m trails of lengths $q_1, \ldots, q_m$ finds a collision on f^r. Thus,

$$\mathsf{cp}^r[q](\mathcal{A}) \leq \max_{\sum q_i = q} \mathsf{cp}^r[q](q_1, \ldots, q_m).$$

By the union bound, we have

$$\mathsf{cp}^r[q](q_1, \ldots, q_m) \leq \sum_{i=1}^{m} \mathsf{cp}_\rho^r[q_i] + \sum_{i=1}^{m-1} \sum_{j=i+1}^{m} \mathsf{cp}_\lambda^r[q_i, q_j].$$

We bound the two terms separately.

$$\sum_{i=1}^{m} \mathsf{cp}_\rho^r[q_i] = \sum_{i=1}^{m} \left[2 \cdot \left(\frac{q_i^2 \sqrt{r}}{N} \right) + 2 \cdot \sqrt{\frac{q_i^2 r \log r}{N}} \right]$$

$$= 2 \cdot \left(\frac{\sqrt{r}}{N} \right) \cdot \sum_{i=1}^{m} q_i^2 + 2 \cdot \sqrt{\frac{r \log r}{N}} \cdot \sum_{i=1}^{m} q_i$$

$$\leq 2 \cdot \left(\frac{\sqrt{r}}{N} \right) \cdot q^2 + 2 \cdot \sqrt{\frac{r \log r}{N}} \cdot q$$

$$= 2 \cdot \left(\frac{q^2 \sqrt{r}}{N} \right) + 2 \cdot \sqrt{\frac{q^2 r \log r}{N}};$$

$$\sum_{i=1}^{m-1} \sum_{j=i+1}^{m} \mathsf{cp}_\lambda^r[q_i, q_j] = \sum_{i=1}^{m-1} \sum_{j=i+1}^{m} \left[32 \cdot \left(\frac{q_i q_j r \log r}{N} \right)^2 \right.$$

$$\left. + 97 \cdot (\log r)^2 \cdot \left(\frac{q_i q_j r \log r}{N} \right) \right]$$

$$= 32 \cdot \left(\frac{r \log r}{N} \right)^2 \cdot \sum_{i=1}^{m-1} \sum_{j=i+1}^{m} q_i^2 q_j^2$$

$$+ 97 \cdot (\log r)^2 \cdot \left(\frac{r \log r}{N} \right) \cdot \sum_{i=1}^{m-1} \sum_{j=i+1}^{m} q_i q_j$$

$$\leq 16 \cdot \left(\frac{r \log r}{N} \right)^2 \cdot q^4 + 49 \cdot (\log r)^2 \cdot \left(\frac{r \log r}{N} \right) \cdot q^2$$

$$= 16 \cdot \left(\frac{q^2 r \log r}{N} \right)^2 + 49 \cdot (\log r)^2 \cdot \left(\frac{q^2 r \log r}{N} \right).$$

Since these bounds are free of $q_1, \ldots, q_m$, this proves Theorem 1 of the paper.

5 Bounding the Advantage of Distinguishing f and f^r

5.1 Security Game

THE SETUP. An oracle $\mathcal{O}$ imitating a function g takes q queries $\{x_i \mid i \in [q]\}$ and returns

$$\{y_i = g(x_i) \mid i \in [q]\}.$$

The q-tuple of input-output pairs of the oracle is called the transcript, denoted as

$$\mathsf{tr} = ((x_1, y_1), (x_2, y_2), \dots, (x_q, y_q)).$$

Both the real oracle $\mathcal{O}_{\mathrm{REAL}}$ and the ideal oracle $\mathcal{O}_{\mathrm{IDEAL}}$ will initially select a uniformly random function f. Then, $\mathcal{O}_{\mathrm{REAL}}$ goes on to imitate f^r, while $\mathcal{O}_{\mathrm{IDEAL}}$ imitates f itself. For any adversary $\mathcal{A}$, we want to bound its advantage, defined as

$$\mathbf{Adv}_{f,f^r}(q) = \left| \Pr\left[\mathcal{A}^{\mathcal{O}_{\mathrm{IDEAL}}}(q) \to 1\right] - \Pr\left[\mathcal{A}^{\mathcal{O}_{\mathrm{REAL}}}(q) \to 1\right] \right|.$$

As in the collision attack of Sect. 4.3, we can view the transcript tr as m trails of lengths $q_1, \dots, q_m$ with sources $z_1, \dots, z_m$, possibly with collisions, such that no query is counted in more than one trail, and hence

$$\sum_{j=1}^{m} q_j = q.$$

For $i \in [m]$, we shall use the notation

$$z_{i,1} := \mathcal{O}(z_i),$$
$$z_{i,j} := \mathcal{O}(z_{i,j-1}), 2 \leq j \leq q_i.$$

GOOD AND BAD TRANSCRIPTS. We partition the set of attainable transcripts into a set $\mathcal{T}_{\mathsf{good}}$ of good transcripts, and a set $\mathcal{T}_{\mathsf{bad}}$ of bad transcripts. We say $\mathsf{tr} \in \mathcal{T}_{\mathsf{bad}}$ if either of the following holds:

- For some $i \in [m]$,
$$z_{i,q_i} = z_i,$$

  that is, the i-th trail forms a permutation cycle. Note that, by our construction of the trails, $z_{i_1,j}$ cannot equal z_{i_2} unless $i_1 = i_2$.

- For some $i_1, i_2 \in [m], j_1 \in [q_{i_1}], j_2 \in [q_{i_2}]$ with $(i_1, j_1) \neq (i_2, j_2)$, we have

$$z_{i_1,j_1} = z_{i_2,j_2},$$

  that is, there is a ρ-collision on one of the trails ($i_1 = i_2$), or there is a λ-collision on two of the trails ($i_1 \neq i_2$).

5.2 Applying the H-Coefficient Technique

Let us denote the probability distribution of the transcripts in the real world by $\mathrm{Pr}_{\mathcal{O}_{\mathrm{REAL}}}$, and in the ideal world by $\mathrm{Pr}_{\mathcal{O}_{\mathrm{IDEAL}}}$. Our proof will use Patarin's H-coefficient technique [17].

Lemma 9 (H-Coefficient Technique). *Let $\mathcal{A}$ be an adversary, and let $\mathcal{T} = \mathcal{T}_{good} \cup \mathcal{T}_{bad}$ be a partition of the set of attainable transcripts. Let ε_1 be such that for all $tr \in \mathcal{T}_{good}$:*

$$\frac{\mathrm{Pr}_{\mathcal{O}_{\mathrm{REAL}}}[tr]}{\mathrm{Pr}_{\mathcal{O}_{\mathrm{IDEAL}}}[tr]} \geq 1 - \varepsilon_1.$$

Furthermore, let $\varepsilon_2 = \mathrm{Pr}_{\mathcal{O}_{\mathrm{IDEAL}}}[tr \in \mathcal{T}_{bad}]$. Then $\mathbf{Adv}_{f,f^r}(q) \leq \varepsilon_1 + \varepsilon_2$.

Proof. For a proof and a detailed explanation of this technique, see Chen and Steinberger [10]. □

PROBABILITY OF BAD TRANSCRIPTS IN IDEAL MODEL. We can easily bound the probability that a transcript tr from the ideal oracle $\mathcal{O}_{\mathrm{IDEAL}}$ is in $\mathcal{T}_{\mathsf{bad}}$. Suppose all of the q responses lie outside $\{z_i \mid i \in [m]\}$, and there is no collision between any of the responses. When this happens, tr cannot be in $\mathcal{T}_{\mathsf{bad}}$. The probability of this is at least $1 - \dfrac{2q^2}{N}$: two responses collide with probability at most $\dfrac{q^2}{N}$; and a response collides with a z_i with probability at most $\dfrac{q^2}{N}$, since there are m different values of z_i, and $m \leq q$. Thus,

$$\varepsilon_2 := \mathrm{Pr}_{\mathcal{O}_{\mathrm{IDEAL}}}[\mathsf{tr} \in \mathcal{T}_{\mathsf{bad}}] \leq \frac{2q^2}{N}.$$

PROBABILITY OF GOOD TRANSCRIPTS. We now focus only on transcripts in $\mathcal{T}_{\mathsf{good}}$. Let us consider a good and attainable transcript $\mathsf{tr} \in \mathcal{T}_{\mathsf{good}}$. For the ideal oracle, as the number of distinct inputs is q, we have

$$\mathrm{Pr}_{\mathcal{O}_{\mathrm{IDEAL}}}[\mathsf{tr}] = \frac{1}{N^q}.$$

Now we bound $\mathrm{Pr}_{\mathcal{O}_{\mathrm{REAL}}}[\mathsf{tr}]$ for $\mathsf{tr} \in \mathcal{T}_{\mathsf{good}}$. Consider a $(q_1, \ldots, q_m)$-query m-trail collision attack on f^r, with sources $z_1, \ldots, z_m$ respectively. Theorem 1 tells us that this attack fails with probability at least $1 - \phi(q, r)$, where

$$\phi(q,r) := 2\left(\frac{q^2\sqrt{r}}{N}\right) + 2\sqrt{\frac{q^2 r \log r}{N}} + 16\left(\frac{q^2 r \log r}{N}\right)^2 + 49(\log r)^2\left(\frac{q^2 r \log r}{N}\right).$$

We now observe that when this attack fails, the attack transcript is either isomorphic as a graph to tr, or contains a permutation cycle.[7] A permutation cycle

---

[7] Note that the graph isomorphism follows from a simple relabeling of inputs and outputs, starting with the sources of every trail. This is possible because excluding collisions and permutation cycles means that no two inputs will have the same output, and outputs never correspond to a source.

occurs when queries of f^r collide with a source z_i, which has probability at most $\dfrac{q^2 r}{N}$, since there are m different values of z_i and $m \le q$. Thus, the attack transcript is isomorphic to tr with probability at least

$$1 - \phi(q, r) - \frac{q^2 r}{N}.$$

Now the graph of this attack transcript has $q + m$ nodes, all distinct. Of these, the m sources are already fixed. The rest can take values in $N^{\underline{q}}$ ways. Now all of these $N^{\underline{q}}$ graphs are equally likely to occur in the scenario described above, i.e., when the m-trail attack fails and does not contain a permutation cycle. One of the equally likely $N^{\underline{q}}$ graphs is the graph of tr. Thus,

$$\mathrm{Pr}_{\mathcal{O}_{\mathrm{REAL}}}[\mathsf{tr}] \ge \left(1 - \phi(q, r) - \frac{q^2 r}{N}\right) \cdot \frac{1}{N^{\underline{q}}}.$$

APPLYING THE H-COEFFICIENT TECHNIQUE. Let $R(\mathsf{tr})$ be the ratio of the probabilities of $\mathsf{tr} \in \mathcal{T}_{\mathsf{good}}$ under $\mathcal{O}_{\mathrm{REAL}}$ and $\mathcal{O}_{\mathrm{IDEAL}}$ respectively. Then we have shown above that

$$R(\mathsf{tr}) \ge \left(1 - \phi(q, r) - \frac{q^2 r}{N}\right) \cdot \frac{1}{\beta(q)}.$$

From Lemma 1, we have

$$\beta(q) \le 1.$$

Thus,

$$R(\mathsf{tr}) \ge 1 - \varepsilon_1$$

where

$$\varepsilon_1 := \phi(q, r) + \frac{q^2 r}{N}.$$

Hence, by the H-coefficient technique of Lemma 9, we have

$$\mathbf{Adv}_{f, f^r}(q) \le \varepsilon_1 + \varepsilon_2.$$

This proves Theorem 2 of the paper.

6   Proof of Lemma 7

RECALLING THE SETUP. In Sect. 4 we defined three sets $\mathcal{S}_1$, $\mathcal{S}_2$, and $\mathcal{S}_3$. $\mathcal{S}_1$ is the set of (t_1, t_2) values that satisfy the constraints

$$\left\lceil \frac{t_1}{r} \right\rceil \le q_1, \left\lceil \frac{t_2}{r} \right\rceil \le q_2, t_1 = t_2 \bmod r;$$

$\mathcal{S}_2$ is the set of $(t_1, t_2, \Delta t, c, \eta)$ values that satisfy the constraints

$$\left\lceil \frac{t_1 + \Delta t + c\eta}{r} \right\rceil \le q_1, \left\lceil \frac{t_2 + \Delta t}{r} \right\rceil \le q_2, t_1 + c\eta = t_2 \bmod r,$$

$\mathcal{S}_3$ is the set of $(t_1, t_2, \Delta t, c, \eta)$ values that satisfy the constraints

$$\left\lceil \frac{t_1 + c\eta}{r} \right\rceil \le q_1, \quad \left\lceil \frac{t_2 + \Delta t}{r} \right\rceil \le q_2, t_1 + c\eta = t_2 + \Delta t \text{ mod } r.$$

We further defined the following:

$$\mathsf{p}_1 = \sum_{\mathcal{S}_1} \mathsf{cp}_\lambda(t_1, t_2);$$

$$\mathsf{p}_2 = \sum_{\mathcal{S}_2} \mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c);$$

$$\mathsf{p}_3 = \sum_{\mathcal{S}_3} \mathsf{cp}_{\rho'}(t_1, t_2, \Delta t, c).$$

Lemma 7 claimed the following bounds for $\mathsf{p}_1, \mathsf{p}_2$ and p_3 (as long as $N \log r > 90$):

$$\mathsf{p}_1 \le \frac{q_1 q_2 r}{N},$$

$$\mathsf{p}_2 \le 6 \cdot (\log r)^2 \cdot \left(\frac{q_1 q_2 r}{N} \right)^2 + 18 \cdot (\log r)^3 \cdot \left(\frac{q_1 q_2 r}{N} \right),$$

$$\mathsf{p}_3 \le 6 \cdot (\log r)^2 \cdot \left(\frac{q_1 q_2 r}{N} \right)^2 + 18 \cdot (\log r)^3 \cdot \left(\frac{q_1 q_2 r}{N} \right).$$

In this section, we establish these bounds.

BOUNDING p_1. For this we need to bound $\#\mathcal{S}_1$. This case is very simple. We observe the $t_1 \le q_1 r$, so there are at most $q_1 r$ choices for t_1. Once t_1 is fixed, given the constraints $t_1 = t_2 \text{ mod } r$ and $t_2 \le q_2 r$, there are at most q_2 choices for t_2. Thus, we have

$$\#\mathcal{S}_1 \le q_1 q_2 r,$$

which, using (5), gives the bound

$$\mathsf{p}_1 = \sum_{\mathcal{S}_1} \mathsf{cp}_\lambda(t_1, t_2) \le \#\mathcal{S}_1 \cdot \frac{1}{N} \le \frac{q_1 q_2 r}{N}.$$

TOWARDS BOUNDING p_2: COUNTING OVER t_1, t_2 AND Δt. This is the most involved part of the calculations. For simplicity of notation we define the function

$$\zeta(\alpha) := (\sqrt{2\alpha N} + 3)^2 = 2\alpha N + 6\sqrt{2\alpha N} + 9.$$

Recall that $\mathcal{S}_2$ is the set of all $(t_1, t_2, \Delta t, c, \eta)$ satisfying

$$\left\lceil \frac{t_1 + \Delta t + c\eta}{r} \right\rceil \le q_1, \quad \left\lceil \frac{t_2 + \Delta t}{r} \right\rceil \le q_2, t_1 + c\eta = t_2 \text{ mod } r.$$

We begin by fixing a choice of c and η. We want to bound the number of choices for $(t_1, t_2, \Delta t)$. For this we relax the constraints a little. Let $\mathcal{S}'_2 = \mathcal{S}'_2(c, \eta)$ be the set of values for $(t_1, t_2, \Delta t)$ satisfying

$$t_1 \le q_1 r, \Delta t \le q_2 r, t_2 \le q_2 r, t_1 + c\eta = t_2 \text{ mod } r.$$

Now we fix a real number $\alpha > 0$, and split $\mathcal{S}_2'$ into two disjoint sets:

$$\mathcal{S}_2'^+[\alpha] := \left\{ (t_1, t_2, \Delta t) \in \mathcal{S}_2' \mid \max(t_1, \Delta t) \geq \sqrt{2\alpha N} + 3 \right\},$$

$$\mathcal{S}_2'^-[\alpha] := \left\{ (t_1, t_2, \Delta t) \in \mathcal{S}_2' \mid \max(t_1, \Delta t) < \sqrt{2\alpha N} + 3 \right\}.$$

For $\mathcal{S}_2'^+[\alpha]$, there are at most $q_1 r$ choices for t_1 and at most $q_2 r$ choices for Δt, and for each of these choices, we have at most q_2 choices for t_2. Thus,

$$\# \mathcal{S}_2'^+[\alpha] \leq q_1 q_2^2 r^2.$$

For $\mathcal{S}_2'^-[\alpha]$, there are at most $\sqrt{2\alpha N} + 3$ choices for t_1 and at most $\sqrt{2\alpha N} + 3$ choices for Δt, and for each of these choices, since choosing t_1 also fixes $t_2 \bmod r$, we have at most q_2 choices for t_2. Thus,

$$\# \mathcal{S}_2'^-[\alpha] \leq (\sqrt{2\alpha N} + 3)^2 \cdot q_2 = \zeta(\alpha) \cdot q_2.$$

When $(t_1, t_2, \Delta t) \in \mathcal{S}_2'^+[\alpha]$,

$$t_1 + t_2 + \Delta t + c\eta \geq \sqrt{2\alpha N} + 3,$$

so that according to (6):

$$\mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c) \leq \frac{e^{-\alpha}}{N^2}.$$

When $(t_1, t_2, \Delta t) \in \mathcal{S}_2'^-[\alpha]$, (7) gives us

$$\mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c) \leq \frac{1}{N^2}.$$

Let

$$\mathsf{p}_2(c, \eta) := \sum_{\mathcal{S}_2'} \mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c)$$

$$= \sum_{\mathcal{S}_2'^+[\alpha]} \mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c) + \sum_{\mathcal{S}_2'^-[\alpha]} \mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c)$$

$$\leq q_1 q_2^2 r^2 \cdot \frac{e^{-\alpha}}{N^2} + \zeta(\alpha) \cdot q_2 \cdot \frac{1}{N^2}$$

$$= \frac{q_2}{N^2} \cdot \left[q_1 q_2 r^2 \cdot e^{-\alpha} + \zeta(\alpha) \right].$$

TOWARDS BOUNDING p_2: COUNTING OVER c AND η. We next bound the number of choices for (c, η) that satisfy the constraints. Again, we relax the constraints a little. Let $\mathcal{T}$ be the set of (c, η) values such that

$$c\eta \leq q_1 r.$$

Next we fix $d = \mathsf{gcd}(c, r)$. Let $\mathcal{T}[d]$ denote the set

$$\{(c, \eta) \in \mathcal{T} \mid \mathsf{gcd}(c, r) = d\}.$$

c now takes values over multiples of d. We split the counting into two parts:

- When $c \le q_1 d$, we recall that η is defined as the smallest solution to $t_1 + c\eta = t_2 \bmod r$. From elementary number theory, we have $\eta \le \dfrac{r}{d}$. Thus, there are q_1 choices of c and for each there are $\dfrac{r}{d}$ choices for η, so in all there are $\dfrac{q_1 r}{d}$ such choices for η and c.

- When $c > q_1 d$, we use the bounds $c \le q_1 r$ and $\eta \le \dfrac{q_1 r}{c}$. Let $z = \dfrac{c}{d}$. Thus, as c runs over all multiples of d from $(q_1+1) \cdot d$ to $q_1 r$, z takes all integer values from $q_1 + 1$ to $\dfrac{q_1 r}{d}$. Thus, the number of choices for η and c with $c > q_1 d$ is

$$
\sum_{z=q_1+1}^{\frac{q_1 r}{d}} \frac{q_1 r}{zd} = \frac{q_1 r}{d} \cdot \sum_{z=q_1+1}^{\frac{q_1 r}{d}} \frac{1}{z} \le \frac{q_1 r}{d} \cdot \log\left(\frac{r}{d}\right),
$$

the last step following from Lemma 2.

Putting these two together, we get

$$
\#\mathcal{T}[d] \le \frac{q_1 r}{d} \cdot \left(1 + \log\left(\frac{r}{d}\right)\right).
$$

Now, d can take values over all factors of r, so we have

$$
\#\mathcal{T} = \sum_{d|r} \#\mathcal{T}[d] \le \sum_{d|r} \frac{q_1 r}{d} \cdot \left(1 + \log\left(\frac{r}{d}\right)\right)
$$
$$
\le \sum_{d|r} \frac{q_1 r}{d} \cdot (1 + \log r) \le q_1 \cdot (1 + \log r) \sum_{d|r} \frac{r}{d}
$$
$$
\le q_1 \cdot (1 + \log r) \cdot \sigma(r),
$$

the last step coming from Lemma 4.

Finally, we observe that whenever $(t_1, t_2, \Delta t, c, \eta) \in \mathcal{S}_2$, we have $(t_1, t_2, \Delta t) \in \mathcal{S}_2'(c, \eta)$, and $(c, \eta) \in \mathcal{T}$. Hence,

$$
\mathsf{p}_2 = \sum_{\mathcal{S}_2} \mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c) \le \sum_{\mathcal{T}} \sum_{\mathcal{S}_2'} \mathsf{cp}_{\lambda\rho}(t_1, t_2, \Delta t, c) = \sum_{\mathcal{T}} \mathsf{p}_2(c, \eta).
$$

This gives us the bound

$$
\mathsf{p}_2 \le \frac{q_1 q_2}{N^2} \cdot (1 + \log r) \cdot \sigma(r) \cdot \left[q_1 q_2 r^2 \cdot e^{-\alpha} + \zeta(\alpha)\right]. \tag{12}
$$

BOUNDING p_3. Recall that $\mathcal{S}_3$ is the set of all $(t_1, t_2, \Delta t, c, \eta)$ satisfying

$$
\left\lceil \frac{t_1 + c\eta}{r} \right\rceil \le q_1, \left\lceil \frac{t_2 + \Delta t}{r} \right\rceil \le q_2, t_1 + c\eta = t_2 + \Delta t \bmod r.
$$

The set $\mathcal{S}_3$ is almost identical to the set $\mathcal{S}_2$. However, the counting arguments are identical to those for p_2, as the relaxation of the constraints is valid for p_2 as well as p_3. Combined with (8), we have

$$
\mathsf{p}_3 = \sum_{\mathcal{S}_3} \mathsf{cp}_{\rho'}(t_1, t_2, \Delta t, c) = \sum_{\mathcal{S}_3} \mathsf{cp}_{\lambda\rho}(t_1, t_2, 0, c) \le \sum_{\mathcal{T}} \sum_{\mathcal{S}_2'} \mathsf{cp}_{\lambda\rho}(t_1, t_2, 0, c).
$$

Thus, we have

$$\mathsf{p}_3 \leq \frac{q_1 q_2}{N^2} \cdot (1 + \log r) \cdot \sigma(r) \cdot \left[q_1 q_2 r^2 \cdot e^{-\alpha} + \zeta(\alpha) \right].$$

SIMPLIFYING THE BOUNDS. Now we make a series of generous relaxations to get a simple easy-to-see bound for p_2 and p_3. Under the assumption that $\sqrt{2\alpha N} + 3 \leq \sqrt{3\alpha N}$, we have $\zeta(\alpha) \leq 3\alpha N$. The assumption can be written as

$$(\sqrt{3} - \sqrt{2}).\sqrt{\alpha N} \geq 3.$$

In other words,

$$\alpha N \geq 9(\sqrt{3} + \sqrt{2})^2 = 9(5 + 2\sqrt{6}).$$

Now, $2\sqrt{6} < 5$, so a sufficient condition to ensure this is $\alpha N \geq 90$. We now put $\alpha = \log r$, and observe in passing that the ensuing assumption that $N \log r \geq 90$ is quite reasonable. For this choice of α, we have

$$\zeta(\alpha) \leq 3N \log r, \tag{13}$$

and

$$e^{-\alpha} = \frac{1}{r}. \tag{14}$$

Since $(5/3) \cdot \log r \geq 1$ for $r \geq 2$, we have

$$1 + \log r < \frac{5}{3} \log r + \log r = \frac{8}{3} \log r. \tag{15}$$

Finally, to bound $\sigma(r)$, we use Lemma 5, which gives us

$$\sigma(r) < 3r \log r. \tag{16}$$

Plugging (13)–(16) into (12), we have

$$\mathsf{p}_2 \leq \frac{q_1 q_2}{N^2} \cdot 3r \log r \cdot \frac{8}{3} \log r \cdot (q_1 q_2 r^2 \cdot \frac{1}{r} + 3N \log r)$$
$$= 8 \cdot (\log r)^2 \cdot \left(\frac{q_1 q_2 r}{N} \right)^2 + 24 \cdot (\log r)^3 \cdot \left(\frac{q_1 q_2 r}{N} \right).$$

Similarly,

$$\mathsf{p}_3 \leq 8 \cdot (\log r)^2 \cdot \left(\frac{q_1 q_2 r}{N} \right)^2 + 24 \cdot (\log r)^3 \cdot \left(\frac{q_1 q_2 r}{N} \right).$$

This completes the proof of Lemma 7.

7 Conclusion and Future Work

We studied the iterated random function problem, and proved the first bound in this setting that is tight up to a factor of $(\log r)^3$. In previous work, the iterated random function problem was seen as a special case of CBC-MAC based on a random function f. We obtained our bound by analysing the probability of a common class of collision attacks, and applying Patarin's H-coefficient technique to bound the advantage of distinguishing f^r from f. Trying to improve the $(\log r)^3$ factor in the security bound is an interesting topic for future work.

References

1. Bellare, M., Kilian, J., Rogaway, P.: The Security of Cipher Block Chaining. In: CRYPTO 1994. LNCS, vol. 839, pp. 341–358 (1994)
2. Bellare, M., Kilian, J., Rogaway, P.: The Security of the Cipher Block Chaining Message Authentication Code. J. Comp. Syst. Sci. 61(3), 362–399 (2000)
3. Bellare, M., Pietrzak, K., Rogaway, P.: Improved Security Analyses for CBC MACs. In: CRYPTO 2005. LNCS, vol. 3621, pp. 527–545 (2005)
4. Bellare, M., Ristenpart, T., Tessaro, S.: Multi-instance Security and Its Application to Password-Based Cryptography. In: CRYPTO 2012. vol. 7417, pp. 312–329. Springer (2012)
5. Bellare, M., Rogaway, P.: The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs. In: EUROCRYPT 2006. LNCS, vol. 4004, pp. 409–426 (2006)
6. Berke, R.: On the security of iterated MACs. Ph.D. thesis, ETH Zürich (2003)
7. Bernstein, D.J.: A short proof of the unpredictability of cipher block chaining (January 2005), http://cr.yp.to/antiforgery/easycbc-20050109.pdf
8. Bhaumik, R., Datta, N., Dutta, A., Mouha, N., Nandi, M.: The Iterated Random Function Problem. ePrint Report 2017/xxx (2017), full version of this paper
9. Bossi, S., Visconti, A.: What Users Should Know About Full Disk Encryption Based on LUKS. In: CANS 2015. LNCS, vol. 9476, pp. 225–237 (2015)
10. Chen, S., Steinberger, J.P.: Tight Security Bounds for Key-Alternating Ciphers. In: EUROCRYPT 2014. LNCS, vol. 8441, pp. 327–350 (2014)
11. Dodis, Y., Gennaro, R., Håstad, J., Krawczyk, H., Rabin, T.: Randomness Extraction and Key Derivation Using the CBC, Cascade and HMAC Modes. In: CRYPTO 2004. LNCS, vol. 3152, pp. 494–510 (2004)
12. Dodis, Y., Ristenpart, T., Steinberger, J.P., Tessaro, S.: To Hash or Not to Hash Again? (In)Differentiability Results for H^2 and HMAC. In: CRYPTO 2012. LNCS, vol. 7417, pp. 348–366 (2012)
13. Ferguson, N., Schneier, B.: Practical Cryptography. Wiley (2003)
14. Gaži, P., Pietrzak, K., Rybár, M.: The Exact PRF-Security of NMAC and HMAC. In: CRYPTO 2014. LNCS, vol. 8616, pp. 113–130 (2014)
15. Minaud, B., Seurin, Y.: The Iterated Random Permutation Problem with Applications to Cascade Encryption. In: CRYPTO 2015. LNCS, vol. 9215, pp. 351–367 (2015)
16. Nandi, M.: A Simple and Unified Method of Proving Indistinguishability. In: INDOCRYPT 2006. LNCS, vol. 4329, pp. 317–334 (2006)
17. Patarin, J.: The "Coefficients H" Technique. In: SAC 2008. LNCS, vol. 5381, pp. 328–345 (2008)
18. Preneel, B., van Oorschot, P.C.: MDx-MAC and Building Fast MACs from Hash Functions. In: CRYPTO 1995. LNCS, vol. 963, pp. 1–14 (1995)
19. Sönmez Turan, M., Barker, E., Burr, W., Chen, L.: Recommendation for Key Derivation Using Pseudorandom Functions (Revised). NIST Special Publication 800-132, National Institute of Standards and Technology (NIST) (December 2010)
20. Wagner, D., Goldberg, I.: Proofs of Security for the Unix Password Hashing Algorithm. In: ASIACRYPT 2000. LNCS, vol. 1976, pp. 560–572 (2000)
21. Wuille, P.: Bitcoin Network Graphs (2017), http://bitcoin.sipa.be/
22. Yao, F.F., Yin, Y.L.: Design and Analysis of Password-Based Key Derivation Functions. In: CT-RSA 2005. LNCS, vol. 3376, pp. 245–261 (2005)
23. Yao, F.F., Yin, Y.L.: Design and Analysis of Password-Based Key Derivation Functions. IEEE Transactions on Information Theory 51(9), 3292–3297 (2005)