Multi-property-preserving Domain Extension Using Polynomial-Based Modes of Operation

Jooyoung Lee[1] and John Steinberger[2],[*]

[1] The Attached Institute of Electronics and Telecommunications Research Institute,
Daejeon, Korea
`jlee05@ensec.re.kr`
[2] Institute of Theoretical Computer Science, Tsinghua University, Beijing, China
`jpsteinb@gmail.com`

Abstract. In this paper, we propose a new double-piped mode of operation for multi-property-preserving domain extension of MACs (message authentication codes), PRFs (pseudorandom functions) and PROs (pseudorandom oracles). Our mode of operation performs twice as fast as the original double-piped mode of operation of Lucks [15] while providing comparable security. Our construction, which uses a class of polynomial-based compression functions proposed by Stam [22, 23], makes a single call to a $3n$-bit to n-bit primitive at each iteration and uses a finalization function f_2 at the last iteration, producing an n-bit hash function $H[f_1, f_2]$ satisfying the following properties.

1. $H[f_1, f_2]$ is unforgeable up to $O(2^n/n)$ query complexity as long as f_1 and f_2 are unforgeable.
2. $H[f_1, f_2]$ is pseudorandom up to $O(2^n/n)$ query complexity as long as f_1 is unforgeable and f_2 is pseudorandom.
3. $H[f_1, f_2]$ is indifferentiable from a random oracle up to $O(2^{2n/3})$ query complexity as long as f_1 and f_2 are public random functions.

To our knowledge, our result constitutes the first time $O(2^n/n)$ unforgeability has been achieved using only an unforgeable primitive of n-bit output length. (Yasuda showed unforgeability of $O(2^{5n/6})$ for Lucks' construction assuming an unforgeable primitive, but the analysis is suboptimal; in the appendix, we show how Yasuda's bound can be improved to $O(2^n)$.)

In related work, we strengthen Stam's collision resistance analysis of polynomial-based compression functions (showing that unforgeability of the primitive suffices) and discuss how to implement our mode by replacing f_1 with a $2n$-bit key blockcipher in Davies-Meyer mode or by replacing f_1 with the cascade of two $2n$-bit to n-bit compression functions.

1 Introduction

The Merkle-Damgård transform has been the most popular method to build a cryptographic hash function from a fixed-size compression function. A major

advantage of this construction is that it preserves collision resistance with an appropriate padding algorithm, allowing one to focus on the construction of secure compression functions. However, Joux showed that if computing collisions becomes somehow feasible for the underlying compression function, then the hash function may fail worse than expected: for a hash function based on a compression function of n-bit output, one can find a t-multicollision only with $O(2^{n/2} \log t)$ complexity, which is much smaller than $O(2^{(t-1)n/t})$ required for an ideal random function. This observation led to several generic attacks such as long-message second preimage attacks [13] and herding attacks [12]. Lucks observed that these weaknesses can be mitigated by increasing the size of the internal state and claimed that the internal state size should be seen as a security parameter of its own right [15]. Since a secure compression function of a larger output size might be harder to construct than the hash function itself, Lucks proposed to use a "narrow" compression function in a double-piped mode. In a subsequent paper [24], Yasuda rigorously analyzed the security of the double-piped mode of operation as a multi-property-preserving domain extension. Specifically, he showed that Lucks' double-piped mode of operation preserves unforgeability up to $O(2^{5n/6})$ query complexity, and indistinguishability and indifferentiability both up to $O(2^n)$ query complexity. Moreover it was later noticed by several researchers that Yasuda's unforgeability bound could be increased to $O(2^n)$ with a slightly modified proof. (See appendix B.)

As such Lucks' construction turned out to provide nearly optimal security. However, the fact that Lucks' compression function uses two applications of a (fairly strong) primitive remains a drawback. Stam [22, 23] recently proposed a class of wide-pipe compression functions making a *single* call to an equal primitive (we call these *polynomial-based* compression functions). In this paper we analyze the security properties of double-piped modes using Stam's polynomial-based compression functions, focusing on MAC-preservation, PRF-preservation and PRO-preservation. Except for PRO-preservation (where we only achieve $O(2^{2n/3})$ security), our bounds are comparable to those found by Yasuda for Lucks' original construction (and even better for unforgeability, given the suboptimality of Yasuda's bound in that case, though the "corrected" unforgeability bound exceeds ours by a factor of n) even though our construction has twice the rate.

Besides performance, a second concern that arises for Lucks' double-pipe construction is the rather strong primitive it assumes: a $3n$-bit to n-bit function (note that careful consideration is typically already given for the construction of $2n$-bit to n-bit compression functions from smaller or more available primitives). Here we also tackle this problem and show our double-piped polynomial-based mode can be implemented with a blockcipher of $2n$-bit key in Davies-Meyer mode, in either the ideal-cipher model or the weaker "unpredictable cipher" model (see Section 5) without significant loss of security. We also prove MAC-preservation and PRF-preservation for a compression function obtained by replacing the $3n$-bit to n-bit primitive with the cascade of two $2n$-bit to n-bit primitives. This latter result potentially opens the door to implementing the

(a) The compression function $\phi[f_1]$

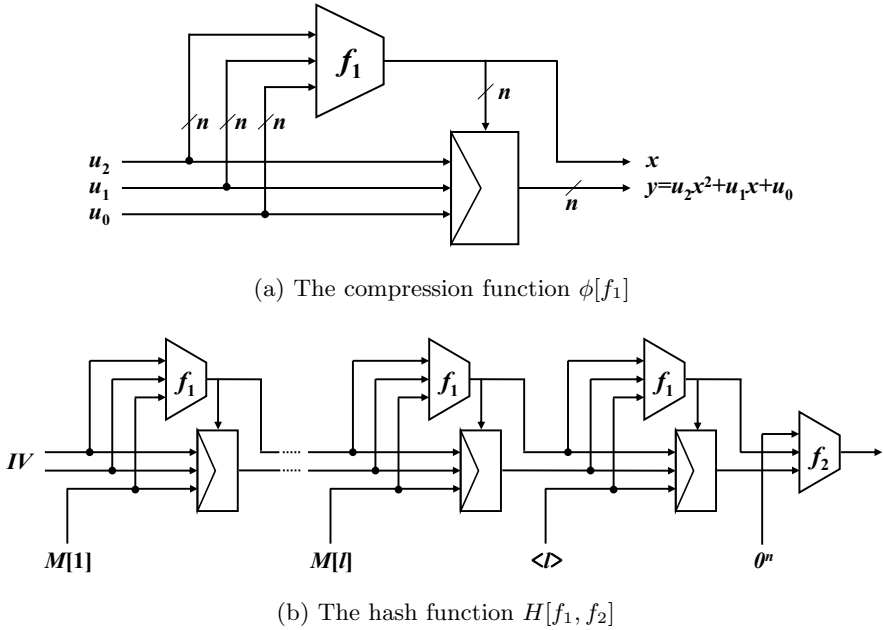(b) The hash function $H[f_1, f_2]$

Fig. 1. The polynomial-based mode of operation for $c = n$

compression function with two calls to an n-bit key blockcipher in Davies-Meyer mode (which would be the first time, to our knowledge, that a $3n$-bit to $2n$-bit compression function using two calls to an n-bit key blockcipher is proved secure nearly up to the birthday bound).

CONSTRUCTION AND RESULTS. To keep our construction comparably general to Lucks' [15] and Yasuda's [24], we discuss a hash function obtained by iterating a $(2n + c)$-bit to $2n$-bit compression function $\phi[f_1]$ where the primitive f_1 used by the compression function is a $2n + c$-bit to n-bit compression function (the "expected" setting of the parameters is $c = n$).

The compression function $\phi[f_1]$ is illustrated in Fig. 1(a) for the case $c = n$ of a $3n$-bit to $2n$-bit compression function. Let $u \in \{0, 1\}^{2n+c}$ and let $u_d || \ldots || u_0$ be the segmentation of u into n-bit blocks $u_0, \ldots, u_{d-1}$ and a block u_d of no more than n bits (so $d = \lceil \frac{2n+c}{n} \rceil - 1$). Then $\phi[f_1](u)$ is defined by

$$\phi[f_1](u) = x||y,$$

where

$$x = f_1(u),$$
$$y = u_d x^d + u_{d-1} x^{d-1} + \cdots + u_1 x + u_0,$$

with all field operations taking place in $\mathbb{F}_{2^n}$ (and u_d being viewed as embedded in $\{0, 1\}^n$). We call $\phi[f_1]$ a *polynomial-based compression function*. This design is due to Stam [22, 23].

Given an independent compression function $f_2 : \{0,1\}^{2n+c} \longrightarrow \{0,1\}^n$, we define a hash function

$$H[f_1, f_2] : \{0,1\}^* \longrightarrow \{0,1\}^n$$

$$M \longmapsto f_2\left(0^c || v\right),$$

where $v = MD[\phi[f_1]](M)$, the Merkle-Damgård iteration of $\phi[f_1]$ on message M (with the usual "strengthened" padding for M that appends the length of the message—see Section 2 for details). The scheme is pictured for $c = n$ in Figure 1(b).

We comment at this point that our mode of operation uses two distinct primitives instead of a single primitive f_1 as do Lucks and Yasuda. As such, our construction explicitly follows the framework of An and Bellare [1] for proving unforgeability whereas Yasuda adopts it implicitly: with some extra work, one can use $f_1 = f_2$ because the f_2-queries are (with very high likelihood) all independent from f_1-queries, due to the presence of the 0^c input segment. (This technique for reducing key material was first used by Maurer and Sjödin [19].) We opt for using two primitives because it simplifies the proofs and allows separation of the security properties required by f_1 and f_2 (the security requirements for f_1 being often much less than those for f_2).

The following points summarize our results on $\phi[f_1]$ and $H[f_1, f_2]$. For this summary we say that f_i is *unforgeable* to mean that a computationally bounded adversary with oracle access to f_i has low probability of predicting the output of f_i on an unqueried value when f_i is sampled from a keyed function family (as low as for a random function of the same range). The *query complexity* of an attack on a variable input length (VIL) function is the number of queries to the underlying primitive necessary to compute the answers to the adversary's queries.

1. We prove that $\phi[f_1]$ is collision resistant up to $O(2^n/n)$ queries to f_1 as long as f_1 is unforgeable. This result also implies the collision resistance of $\phi[f_1]$ against an information-theoretic adversary if f_1 is a random function. It also implies $H[f_1, f_2]$ is unforgeable up to $O(2^n/n)$ query complexity as long as f_1 and f_2 are unforgeable, and that $H[f_1, f_2]$ is weakly collision resistant up to $O(2^n/n)$ query complexity as long as f_1 is unforgeable and f_2 is weakly collision resistant.
2. We prove that $H[f_1, f_2]$ is pseudorandom up to $O(2^n/n)$ query complexity as long as f_1 and f_2 are pseudorandom. In the complexity-theoretic model, we can weaken the assumption so that f_1 is unforgeable.
3. We prove that $\phi[f_1]$ is preimage aware[1] up to $O(2^{2n/3})$ query complexity as long as f_1 is a public random function. By the results of [6], this implies $H[f_1, f_2]$ is indifferentiable from a random oracle up to $O(2^{2n/3})$ query complexity as long as f_1 and f_2 are public random functions.

---

[1] *Preimage awareness* is a security notion introduced by Dodis, Ristenpart and Shrimpton [6]. The Merkle-Damgård iteration of a preimage aware compression function composed with a random function results in a construction that is indifferentiable from a random oracle, up to the preimage awareness security of the compression function and the maximum message length queried to the iterated construction.

(a) The quadratic blockcipher-based function

(b) The quadratic cascade function

Fig. 2. Variants of the quadratic compression function

REFINEMENTS. As mentioned, we also investigate two variants of the $3n$-to-$2n$ bit polynomial-based compression function (a.k.a. the "quadratic" compression function) with a view towards concrete implementations of the mode. These alternate constructions are shown in Figure 2. The first variant replaces f_1 by a blockcipher E of $2n$-bit key in Davies-Meyer mode. We show this compression function $\psi[E]$ is collision resistant up to $O(2^n/n)$ queries as long as E is "unpredictable", a notion we discuss in Section 5. Similar corollaries follow on the security of the hash function obtained by iterating $\psi[E]$.

The second is obtained by replacing f_1 with the cascade of two $2n$-bit to n-bit compression functions h_1 and h_2. We show this compression function, denoted $\tau[h_1, h_2]$, is collision resistant up to $O(2^n/n^3)$ queries as long as h_1 and h_2 are unforgeable. It follows that the hash function $G[h_1, h_2, f_2]$ obtained by iterating $\tau[h_1, h_2]$ (defined like $H[f_1, f_2]$ but substituting $\tau[h_1, h_2]$ for $\phi[f_1]$) has unforgeability security up to $O(2^n/n^3)$ query complexity when h_1, h_2 and f_2 are unforgeable, has collision security up to $O(2^n/n^3)$ query complexity when h_1, h_2 are unforgeable and f_2 is collision resistant, and is indistinguishable from a PRF up to $O(2^n/n^3)$ query complexity when h_1, h_2 are unforgeable and f_2 is pseudorandom.

RELATED WORK. All the compression functions discussed in this paper, including the cascaded and blockcipher variants, were proposed by Stam [22,23]. In [23] Stam proves the collision resistance of polynomial-based compression functions of degrees two and three in the random function model, and also proves the collision security of the quadratic blockcipher mode in the ideal cipher model. Here our contribution is that we weaken the model by showing collision resistance is already assured when f_1 and E are unforgeable/unpredictable rather than random. (It is this weakening of the model that allows us to prove MAC-preservation results for the resulting hash functions.) Regarding the quadratic cascade compression function, Stam proves collision resistance for a special class of non-adaptive adversaries assuming random primitives. Our analysis supports fully adaptive adversaries and once again weakens the model to unforgeable primitives.

Lucks [16] recently proposed a double-pipe hash function iterating a $3n$-bit to $2n$-bit compression function which, like the quadratic blockcipher-based mode, uses a single call to a blockcipher of $2n$-bit key. However, by contrast to the quadratic blockcipher compression function, Lucks' compression function is neither collision resistant nor preimage resistant. As a consequence, collision and

preimage security can only be proved in the iteration (higher security notions like indifferentiability are unaddressed). On the other hand, for $n = 128$ Lucks gives a better explicit collision security bound than we do for the quadratic blockcipher compression function: 2^{122} versus 2^{119} queries, respectively.

This paper can be seen as an extension of Yasuda's work [24] since our main achievement is to double the rate of that construction while maintaining comparable MAC-preservation and PRF-preservation properties. However, from a technical standpoint we owe most to Dodis and Steinberger [7], whose "multicollision-to-forgery" (MTF) balls-in-bins games are used in nearly all of our analyses (the sole exception being the preimage awareness bound for polynomial-based compression functions). Indeed, the main "message" of this paper may well be the versatility and power of MTF games.

2   Preliminaries

$\mathbb{F}_{2^n}$ denotes a finite field of order 2^n. Throughout our work, we will identify $\mathbb{F}_{2^n}$ and $\{0,1\}^n$, assuming a fixed mapping between the two sets. For $u \in \{0,1\}^*$, $|u|$ is the length in bits of u. For two bitstrings u and v, $u||v$ denotes the concatenation of u and v. For a set U, we write $u \xleftarrow{\$} U$ to denote uniform sampling from U and assignment to u.

Let $M \in \{0,1\}^*$ and let $c \geq 1$ be message block length (as c will denote throughout the paper). Then $\mathsf{pad}(M) := M||10^b||\langle l \rangle$ where b is the least integer such that $|M||10^b|$ is a multiple of c and where l is the number of c-bit blocks in $M||10^b$. (This representation is possible as long as $l < 2^c$, but this is not a restriction for most applications.) The main property of $\mathsf{pad}(\cdot)$ is that it gives a suffix-free encoding of messages.

The (strengthened) *Merkle-Damgård transform* produces a VIL-function $MD[F] : \{0,1\}^* \to \{0,1\}^n$ from a FIL-function $F : \{0,1\}^{n+c} \to \{0,1\}^n$. Given a predetermined constant $IV \in \{0,1\}^n$, the function $MD[F]$ is defined as follows.

> **Function** $MD[F](M)$
>
> $z[0] \leftarrow IV$
>   Break $\mathsf{pad}(M)$ into c-bit blocks, $\mathsf{pad}(M) = M[1]||\ldots||M[l+1]$
>   **for** $i \leftarrow 1$ to $l+1$ **do**
> $z[i] \leftarrow F(z[i-1]||M[i])$
>   **return** $z[l+1]$

3   Security Definitions

Unforgeability and Weak Collision Resistance. A *function family* is a map $f : \{0,1\}^\kappa \times \mathsf{Dom}(f) \to \{0,1\}^n$ where $\mathsf{Dom}(f) \subset \{0,1\}^*$. The bitstrings in $\{0,1\}^\kappa$ are the *keys* of f and we write $f_k(M)$ for $f(k,M)$ for $k \in \{0,1\}^\kappa$ and

$M \in \mathsf{Dom}(f)$. The function f_k is called a *member* of f. The unforgeability of f as a secure message authentication code (MAC) is estimated by the experiment $\mathbf{Exp}_{\mathcal{A}}^{\mathsf{mac}}$ described in Figure 3(a). In the experiment, an adversary $\mathcal{A}$ has oracle access to $f_k(\cdot)$ and tries to produce a valid tag z for a new message M. Here we call a message M "new" if it has not been queried to oracle $f_k(\cdot)$. The *forgery advantage* of $\mathcal{A}$ is defined by

$$\mathbf{Adv}_f^{\mathsf{mac}}(\mathcal{A}) = \mathbf{Pr}\left[\mathbf{Exp}_{\mathcal{A}}^{\mathsf{mac}} = 1\right]. \tag{1}$$

The probability is taken over the random choice of k and $\mathcal{A}$'s coins (if any). We define $\mathbf{Adv}_f^{\mathsf{mac}}(t, q, \mu)$ as the maximum of $\mathbf{Adv}_f^{\mathsf{mac}}(\mathcal{A})$ over all adversaries $\mathcal{A}$ making at most q queries whose total combined length is at most μ bits (including the forgery produced by $\mathcal{A}$) and of "running time" at most t. The "running time" is defined to be the total running time of the experiment, including the time required to compute the answers to $\mathcal{A}$'s queries. We write $\mathbf{Adv}_f^{\mathsf{mac}}(t, q)$ for $\mathbf{Adv}_f^{\mathsf{mac}}(t, q, \mu)$ if f is a family of fixed input length functions, as in this case μ is automatically determined by q.

The weak collision resistance (WCR) of f is estimated by the experiment $\mathbf{Exp}_{\mathcal{A}}^{\mathsf{wcr}}$ described in Figure 3(b). In contrast to the definition of collision resistance (in the dedicated-key setting) where $\mathcal{A}$ is provided key k, $\mathcal{A}$ is allowed only oracle access to $f_k(\cdot)$. Let

$$\mathbf{Adv}_f^{\mathsf{wcr}}(\mathcal{A}) = \mathbf{Pr}\left[\mathbf{Exp}_{\mathcal{A}}^{\mathsf{wcr}} = 1\right]. \tag{2}$$

Then the weak collision resistance of f, denoted $\mathbf{Adv}_f^{\mathsf{wcr}}(t, q, \mu)$, is defined to be the maximum of $\mathbf{Adv}_f^{\mathsf{wcr}}(\mathcal{A})$ over all adversaries $\mathcal{A}$ making at most q queries whose total combined length is at most μ bits and of running time at most t. When f is a family of fixed input length functions we likewise write $\mathbf{Adv}_f^{\mathsf{wcr}}(t, q)$ instead of $\mathbf{Adv}_f^{\mathsf{wcr}}(t, q, \mu)$.

Our security proof for unforgeability will follow the approach developed by An and Bellare [1]. One of their results is that $f_2 \circ MD[f_1]$ is a VIL-MAC if f_1 is a FIL-WCR and f_2 is a FIL-MAC. With a slight modification, we summarize Lemma 4.2 and Lemma 4.3 in [1] as the following lemma.

Lemma 1. *Let* $f^1 : \{0,1\}^{\kappa} \times \{0,1\}^{n+c} \to \{0,1\}^n$ *and* $f^2 : \{0,1\}^{\kappa'} \times \{0,1\}^n \to \{0,1\}^m$ *be function families. Then,*

$$\mathbf{Adv}_{f^2 \circ MD[f^1]}^{\mathsf{mac}}(t, \tilde{q}, \mu) \leq \mathbf{Adv}_{f^2}^{\mathsf{mac}}(t, \tilde{q}) + \mathbf{Adv}_{f^1}^{\mathsf{wcr}}\left(t, \left\lfloor \frac{\mu}{c} \right\rfloor + 2\tilde{q}\right).$$

Remark 1. $\left\lfloor \frac{\mu}{c} \right\rfloor + 2\tilde{q}$ is the maximum number of queries to f_1 required to compute $MD[f_1](x_i)$ for $x_1, \ldots, x_{\tilde{q}}$ such that $|x_1| + \ldots + |x_{\tilde{q}}| \leq \mu$.

Indifferentiability and Indistinguishability. In the indifferentiability framework, a distinguisher is given two systems $(F[\mathcal{P}], \mathcal{P})$ and $(\mathcal{H}, \mathcal{S}[\mathcal{H}])$. Here $\mathcal{P}$ is an ideal primitive used as a building block for the construction of $F[\mathcal{P}]$. An ideal primitive $\mathcal{H}$ and a probabilistic Turing machine $\mathcal{S}[\mathcal{H}]$ with oracle access

Experiment Exp$_{\mathcal{A}}^{\mathsf{mac}}$

$k \xleftarrow{\$} \{0,1\}^\kappa$
$(M, z) \leftarrow \mathcal{A}^{f_k(\cdot)}$
if M is new and $f_k(M) = z$ **then**
    output 1
else
    output 0

Experiment Exp$_{\mathcal{A}}^{\mathsf{wcr}}$

$k \xleftarrow{\$} \{0,1\}^\kappa$
$(M, M') \leftarrow \mathcal{A}^{f_k(\cdot)}$
if $f_k(M) = f_k(M')$ **then**
    output 1
else
    output 0

(a) Quantification of unforgeability

(b) Quantification of weak collision resistance

Fig. 3. Experiments for quantification of unforgeability and weak collision resistance

to $\mathcal{H}$ have the same interfaces as $F[\mathcal{P}]$ and $\mathcal{P}$, respectively. The *simulator* $\mathcal{S}[\mathcal{H}]$ tries to emulate the ideal primitive $\mathcal{P}$ so that no distinguisher can tell apart the two systems $(\mathcal{H}, \mathcal{S}[\mathcal{H}])$ and $(F[\mathcal{P}], \mathcal{P})$ with non-negligible probability, based on their responses to queries that the distinguisher may send. We say that the construction $F[\mathcal{P}]$ is indifferentiable from $\mathcal{H}$ if the existence of such a simulator is proved. The indifferentiability implies the absence of a generic attack against $F[\mathcal{P}]$ that regards $\mathcal{P}$ merely as a black-box. Now we give a formal definition of indifferentiability in the information-theoretic model. For a more comprehensive introduction of the indifferentiability framework, we refer to [3, 18].

Definition 1. *A Turing machine F with oracle access to an ideal primitive $\mathcal{P}$ is said to be (q, ϵ, t)-indifferentiable from an ideal primitive $\mathcal{H}$ if there exists a simulator $\mathcal{S}$ of running time at most t with oracle access to $\mathcal{H}$ such that for any distinguisher $\mathcal{A}$ making at most q queries, it holds that*

$$\left| \mathbf{Pr}\left[\mathcal{A}^{F[\mathcal{P}], \mathcal{P}} = 1 \right] - \mathbf{Pr}\left[\mathcal{A}^{\mathcal{H}, \mathcal{S}[\mathcal{H}]} = 1 \right] \right| < \epsilon.$$

If $\mathcal{H}$ is a public random function, then $F[\mathcal{P}]$ is called a (q, ϵ, t)-pseudorandom oracle (PRO).

If $\mathcal{A}$ is not allowed to make queries for the underlying primitive, we obtain the definition of indistinguishability.

Definition 2. *A Turing machine F with oracle access to an ideal primitive $\mathcal{P}$ is said to be (q, ϵ)-indistinguishable from an ideal primitive $\mathcal{H}$ if for any distinguisher $\mathcal{A}$ making at most q queries, it holds that*

$$\left| \mathbf{Pr}\left[\mathcal{A}^{F[\mathcal{P}]} = 1 \right] - \mathbf{Pr}\left[\mathcal{A}^{\mathcal{H}} = 1 \right] \right| < \epsilon.$$

If $\mathcal{H}$ is a public random function, then $F[\mathcal{P}]$ is called a (q, ϵ)-pseudorandom function (PRF).

Collision Resistance and Adaptive Preimage Resistance. First, we review the definition of collision resistance *in the information-theoretic model.*

<div style="border:1px solid;">

Experiment $\mathbf{Exp}_{\mathcal{A}}^{\mathsf{coll}}$

  $\mathcal{A}^{\mathcal{P}}$ updates $\mathcal{Q}$

  if $\exists\ u\ \neq\ u',\ z$ s.t. $z\ =\ F_{\mathcal{Q}}(u)\ =$ $F_{\mathcal{Q}}(u')$ **then**

    output 1

  else

    output 0

</div>

(a) Quantification of collision resistance

<div style="border:1px solid;">

Experiment $\mathbf{Exp}_{\mathcal{A}}^{\mathsf{adpr}}$

  $\mathcal{A}^{\mathcal{P}}$ updates $\mathcal{Q}$ and $\mathcal{L}$

  if $\exists\ u$ s.t. $z = F_{\mathcal{Q}}(u)$ for some $z \in \mathcal{L}$ **then**

    output 1

  else

    output 0

</div>

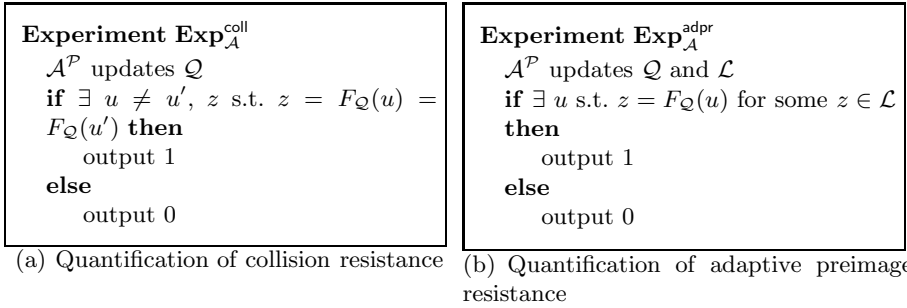(b) Quantification of adaptive preimage resistance

Fig. 4. Experiments for quantification of collision resistance and adaptive preimage resistance

Given a function $F = F[\mathcal{P}]$ and an information-theoretic adversary $\mathcal{A}$ both with oracle access to an ideal primitive $\mathcal{P}$, the collision resistance of F against $\mathcal{A}$ is estimated by the experiment $\mathbf{Exp}_{\mathcal{A}}^{\mathsf{coll}}$ described in Figure 4(a). The experiment records every query-response pair that $\mathcal{A}$ obtains by oracle queries into a *query history* $\mathcal{Q}$. We write $z = F_{\mathcal{Q}}(u)$ if $\mathcal{Q}$ contains all the query-response pairs required to compute $z = F(u)$. At the end of the experiment, $\mathcal{A}$ would like to find two distinct evaluations yielding a collision. The *collision-finding advantage* of $\mathcal{A}$ is defined to be

$$\mathbf{Adv}_F^{\mathsf{coll}}(\mathcal{A}) = \mathbf{Pr}\left[\mathbf{Exp}_{\mathcal{A}}^{\mathsf{coll}} = 1\right]. \tag{3}$$

The probability is taken over the random choice of $\mathcal{P}$ and $\mathcal{A}$'s coins (if any). For $q > 0$, we define $\mathbf{Adv}_F^{\mathsf{coll}}(q)$ as the maximum of $\mathbf{Adv}_F^{\mathsf{coll}}(\mathcal{A})$ over all adversaries $\mathcal{A}$ making at most q queries.

In this section, we also present a new notion of security, called *adaptive preimage resistance*. This notion will be useful for the proof of preimage awareness. Given a function $F = F[\mathcal{P}]$ and an information-theoretic adversary $\mathcal{A}$ both with oracle access to an ideal primitive $\mathcal{P}$, the adaptive preimage resistance of F against $\mathcal{A}$ is estimated by the experiment $\mathbf{Exp}_{\mathcal{A}}^{\mathsf{adpr}}$ described in Figure 4(b). At any point during the experiment, the adversary $\mathcal{A}$ is allowed to choose a commitment element z in the range of F such that the query history $\mathcal{Q}$ has not determined any preimage of z. The experiment $\mathbf{Exp}_{\mathcal{A}}^{\mathsf{adpr}}$ records the element z into a *commitment list* $\mathcal{L}$. Queries and commitments are made in an arbitrarily interleaved order. At the end of the experiment, $\mathcal{A}$ would like to succeed in finding a preimage of some element in the commitment list. The *adaptive preimage-finding advantage* of $\mathcal{A}$ is defined to be

$$\mathbf{Adv}_F^{\mathsf{adpr}}(\mathcal{A}) = \mathbf{Pr}\left[\mathbf{Exp}_{\mathcal{A}}^{\mathsf{adpr}} = 1\right]. \tag{4}$$

For $q_p, q_e > 0$, we define $\mathbf{Adv}_F^{\mathsf{adpr}}(q_p, q_e)$ as the maximum of $\mathbf{Adv}_F^{\mathsf{adpr}}(\mathcal{A})$ over all adversaries $\mathcal{A}$ that make at most q_p queries and at most q_e commitments.

Preimage Awareness. The notion of preimage awareness was first introduced by Dodis, Ristenpart and Shrimpton [6]. This notion is useful for the proof of

indifferentiability of "NMAC" type constructions. Let $F = F[\mathcal{P}]$ be a function with oracle access to an ideal primitive $\mathcal{P}$. In order to estimate the preimage awareness of F, we use the experiment described in Figure 5. Here an adversary $\mathcal{A}$ is provided two oracles P and Ex. The oracle P provides access to the ideal primitive $\mathcal{P}$ and records a query history $\mathcal{Q}$. Note that this oracle is implicitly used in the experiments for collision resistance and adaptive preimage resistance. The *extraction oracle* Ex provides an interface to an *extractor* $\mathcal{E}$, which is a deterministic algorithm that takes as input an element z in the range of F and the query history $\mathcal{Q}$, and returns either $\perp$ or an element in the domain of F. With respect to the extractor $\mathcal{E}$, we define the advantage of $\mathcal{A}$ to be

$$\mathbf{Adv}_F^{\mathsf{pra}}(\mathcal{A}, \mathcal{E}) = \mathbf{Pr}\left[\mathbf{Exp}_{\mathcal{A},\mathcal{E}}^{\mathsf{pra}} = 1\right]. \tag{5}$$

Let $\mathcal{E}^*$ be an algorithm that on input $(z, \mathcal{Q})$ returns an element u if there exists u such that $F_{\mathcal{Q}}(u) = z$ and $\perp$ otherwise. Let

$$\mathbf{Adv}_F^{\mathsf{pra}}(\mathcal{A}, \mathcal{E}^*) = \mathsf{P}_1 + \mathsf{P}_2,$$

where P_1 is the probability that $u \neq \mathsf{V}[z] \neq \perp$ and $\mathsf{L}[z] = 1$ at the end of the experiment and P_2 is the probability that $u \neq \mathsf{V}[z] = \perp$ and $\mathsf{L}[z] = 1$ at the end of the experiment. Then $\mathcal{A}$ can be regarded as a collision-finding adversary such that $\mathbf{Adv}_F^{\mathsf{coll}}(\mathcal{A}) = \mathsf{P}_1$. Furthermore, $\mathcal{A}$ can be transformed into an adaptive preimage-finding adversary $\mathcal{B}$ such that $\mathbf{Adv}_F^{\mathsf{adpr}}(\mathcal{B}) = \mathsf{P}_2$: $\mathcal{B}$ runs $\mathcal{A}$ as a subroutine, asks the same primitive queries as $\mathcal{A}$, and makes commitments z if $\mathcal{A}$ makes a query for $\mathsf{Ex}(z)$ and $\mathcal{Q}$ has not determined any preimage of z. If $\mathcal{A}$ makes at most q_p primitive queries and q_e extraction queries, then it follows that $\mathsf{P}_1 \leq \mathbf{Adv}_F^{\mathsf{coll}}(q_p)$ and $\mathsf{P}_2 \leq \mathbf{Adv}_F^{\mathsf{adpr}}(q_p, q_e)$. We record this observation as the following lemma.

Lemma 2. *Let $F = F[\mathcal{P}]$ be a function with oracle access to an ideal primitive $\mathcal{P}$. Then there exists an extractor $\mathcal{E}^*$ such that for any adversary $\mathcal{A}$ it holds that*

$$\mathbf{Adv}_F^{\mathsf{pra}}(\mathcal{A}, \mathcal{E}^*) \leq \mathbf{Adv}_F^{\mathsf{coll}}(q_p) + \mathbf{Adv}_F^{\mathsf{adpr}}(q_p, q_e).$$

The main application of preimage awareness lies in the construction of pseudorandom oracles. In the following lemma which is a combination of Theorem 4.1 and Theorem 4.2 in [6], unpad is an algorithm such that $\mathsf{unpad}(y) = x$ if $y = \mathsf{pad}(x)$ is a valid output of pad and $\mathsf{unpad}(y) = \perp$ otherwise. For any algorithm $\mathcal{F}$, we write $\mathsf{Time}(\mathcal{F}, l)$ for the maximum time required to compute $\mathcal{F}(x)$ for any input x such that $|x| \leq l$. If $\mathcal{F}$ is an algorithm with oracle access to an ideal primitive $\mathcal{P}$, then $\mathsf{NQ}(F, l)$ is the maximum number of queries to $\mathcal{P}$ required to compute $F(x)$ for any input x such that $|x| \leq l$. Without any constraint on the input length, we just write $\mathsf{Time}(\mathcal{F})$ and $\mathsf{NQ}(F)$.

Lemma 3. *Let $F : \{0,1\}^{n+c} \to \{0,1\}^n$ be a function with oracle access to an ideal primitive $\mathcal{P}$ and let $g : \{0,1\}^n \to \{0,1\}^m$ and $\mathcal{H} : \{0,1\}^* \to \{0,1\}^m$ be public random functions for $m \leq n$. For an arbitrary extractor $\mathcal{E}$ with respect to F,*

Experiment $\mathbf{Exp}^{\mathsf{pra}}_{\mathcal{A},\mathcal{E}}$	Oracle $\mathsf{P}(x)$	Oracle $\mathsf{Ex}(z)$
$u \leftarrow \mathcal{A}^{\mathsf{P},\mathsf{Ex}}$	$y \leftarrow \mathcal{P}(x)$	$\mathsf{L}[z] \leftarrow 1$
$z \leftarrow F[\mathcal{P}](u)$	$\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(x,y)\}$	$\mathsf{V}[z] \leftarrow \mathcal{E}(z,\mathcal{Q})$
output 1 if $u \neq \mathsf{V}[z]$ and $\mathsf{L}[z]=1$	**return** y	**return** $\mathsf{V}[z]$

Fig. 5. Experiments for quantification of preimage awareness. Arrays L and V are global, and respectively initialized empty and $\perp$ everywhere.

there exists a simulator $\mathcal{S}$ such that for any distinguisher $\mathcal{A}$ making at most (q_0, q_1, q_2) queries to the three oracle interfaces associated with $(\mathcal{H}, \mathcal{P}, g)$, there exists an adversary $\mathcal{B}$ such that

$$\left| \mathbf{Pr}\left[\mathcal{A}^{g \circ MD[F[\mathcal{P}]], (\mathcal{P},g)} = 1 \right] - \mathbf{Pr}\left[\mathcal{A}^{\mathcal{H}, \mathcal{S}[\mathcal{H}]} = 1 \right] \right| \leq \mathbf{Adv}^{\mathsf{pra}}_F(\mathcal{B}, \mathcal{E}).$$

Let l_{max} be the length in bits of the longest query made by $\mathcal{A}$ to its first oracle, and let $L = \left\lceil \frac{l_{max}+1}{c} + 1 \right\rceil$. Then, simulator $\mathcal{S}$ runs in time $O\left(q_1 + Lq_2\mathsf{Time}(\mathcal{E}) + Lq_2\mathsf{Time}(\mathsf{unpad})\right)$. Adversary $\mathcal{B}$ runs in time $O\left(\mathsf{Time}(\mathcal{A}) + q_0\mathsf{Time}(MD[F], l_{max}) + q_1 + (L+1)q_2\right)$, makes at most $L\mathsf{NQ}(F)(q_0+1)+q_1$ primitive queries, and makes at most Lq_2 extraction queries.

4 Security of the Polynomial-Based Mode of Operation

For this section and the rest of the paper, $\phi[f_1]$ and $H[f_1, f_2]$ refer to the compression function and hash function defined in Section 1. We use "log" to denote the logarithm base 2. For simplicity of notation, we assume that $\log q$ is an integer for the number of queries q.

4.1 Weak Collision Resistance and Unforgeability

We begin with the proof of weak collision resistance for $\phi[f_1]$ such that f_1 is randomly chosen from a function family f.

Theorem 1. Let ϕ be a function family defined by $f : \{0,1\}^\kappa \times \{0,1\}^{2n+c} \to \{0,1\}^n$. Then,

$$\mathbf{Adv}^{\mathsf{wcr}}_{\phi}(t, q) \leq 2q \max(d + \log q, d2^{d+2})\mathbf{Adv}^{\mathsf{mac}}_f\left(t + O(d^2 n^2 q^{d+2}) + \mathsf{Time}_d, q\right),$$

where $d = \left\lceil \frac{2n+c}{n} \right\rceil - 1$ and Time_d is the time required to solve a univariate polynomial equation of degree d over $\mathbb{F}_{2^n}$.

Remark 2. For a univariate polynomial of degree d over $\mathbb{F}_{2^n}$, there is a deterministic algorithm to find zeros using $O(d^{3/2}n)$ field operations (ignoring log factors). See [8, 9].

In order to prove Theorem 1, we use a generalization of the multicollision-to-forgery (MTF) balls-in-bins game first introduced in [7].

MTF game. This game is played by two players $\mathcal{A}$ and $\mathcal{B}$ according to the following rules. Parameters are integers $q > 0$ and $m_2 > m_1 \geq 0$.

1. The game consists of q rounds.
2. At each round, $\mathcal{A}$ publicly places a set of balls into a set of bins such that
   (a) balls placed at the same round go into distinct bins,
   (b) the number of balls that are placed into bins already containing m_1 balls at the moment of placement is finite,
   (c) some bin eventually contains more than m_2 balls.
3. Before each round, $\mathcal{B}$ can secretly "pass" or "guess" a bin that will receive a ball in the next round. $\mathcal{B}$ makes exactly one guess throughout the game.
4. If $\mathcal{B}$ makes a correct guess, then $\mathcal{B}$ wins. Otherwise, $\mathcal{B}$ loses.

Note there is no constraint on the total number of balls or bins.

Lemma 4. *Irrespective of $\mathcal{A}$'s strategy, there exists a strategy for $\mathcal{B}$ to win the above game with probability at least $1/q \cdot 1/(c_{m_1})^{1/(m_2-m_1)}$, where c_{m_1} is the number of balls that are placed into bins already containing m_1 balls at the moment of placement.*

Proof. $\mathcal{B}$'s strategy is simple, as follows:

1. Choose a round $i \in \{1, \ldots, q\}$ and a level $j \in \{m_1 + 1, \ldots, m_2\}$ uniformly at random.
2. Before the i-th round of the game, guess a bin uniformly at random from all bins containing at least j balls already.

Let c_j be the total number of balls that are placed into bins that already have at least j balls in them right before the round when the ball is placed. For a given value of j, each ball placed into a bin with at least j balls in it already gives $\mathcal{B}$ chance at least $1/(qc_{j-1})$ of winning since c_{j-1} is an upper bound for the number of bins that have j balls in them at the end of the game. Therefore $\mathcal{B}$'s chance of winning is at least $c_j/(qc_{j-1})$ for a given value of j, and hence $\mathcal{B}$'s overall chance of winning is at least

$$\frac{1}{m_2 - m_1} \sum_{j=m_1+1}^{m_2} \frac{c_j}{qc_{j-1}} = \frac{1}{q}\mathsf{ArithmeticMean}\left(\frac{c_{m_1+1}}{c_{m_1}}, \ldots, \frac{c_{m_2}}{c_{m_2-1}}\right)$$

$$\geq \frac{1}{q}\mathsf{GeometricMean}\left(\frac{c_{m_1+1}}{c_{m_1}}, \ldots, \frac{c_{m_2}}{c_{m_2-1}}\right)$$

$$\geq \frac{1}{q}\left(\frac{1}{c_{m_1}}\right)^{1/(m_2-m_1)}. \qquad \square$$

Note that Lemma 4 asserts nothing about $\mathcal{B}$'s efficiency—in fact, if a huge number of balls are thrown, it may be difficult to keep track of all bins that have received at least j balls already (which is necessary for sampling uniformly among the bins).

In our case, bins will often be curves defined by polynomials of degree $\leq d$ over $\mathbb{F}_{2^n}$ and balls points in $\mathbb{F}_{2^n} \times \mathbb{F}_{2^n}$, where a ball (x, y) goes into bin $\mathcal{C}$ if $(x, y) \in \mathcal{C}$ (a ball is thus "cloned" into many different bins). In this setting, it becomes easier to keep track of which bins have at least j balls in them when $j \geq d+1$, as $d+1$ points uniquely determine a polynomial of degree $\leq d$. Thus for such a game it may be helpful to set $m_1 = d$, in order to keep the complexity of sampling under control. (Dodis and Steinberger do not have games in which the number of bins containing balls is so large that sampling for small values of j is an issue, and always use MTF games with $m_1 = 0$.) The value of m_2 is then set large enough to make the term $(1/c_{m_1})^{1/(m_2-m_1)}$ small.

We typically upper bound c_{m_1} by qM where M is an upper bound on *the total number of bins with at least $m_1 + 1$ balls at the end of the game*. Indeed, because balls are thrown into distinct bins at each round, this definition of M implies that at each round at most M balls are thrown into bins with m_1 balls in them already. We thus have the following corollary:

Corollary 1. *If the number of bins that contain at least $m_1 + 1$ balls at the end of the MTF game is at most M, then $\mathcal{B}$ can win the MTF game with probability at least $1/q \cdot 1/(qM)^{1/(m_2-m_1)}$.*

We note that Corollary 1 is a bit wasteful, in the sense that it is possible to give a better bound for $\mathcal{B}$'s chance of success as a function of m_1, m_2 and M from the relationship $M = c_{m_1} - c_{m_1+1}$ and the fact that $\mathcal{B}$'s chance of success is also lower bounded by $\frac{\alpha}{q(m_2-m_1)}$ where $\alpha = \max\left(\frac{c_{m_1+1}}{c_{m_1}}, \ldots, \frac{c_{m_2}}{c_{m_2-1}}\right)$. However this gain leads to a more complicated statement and is minor enough for us to ignore [2].

Proof (Theorem 1). Let $\mathcal{A}$ be a weak collision-finding adversary such that

$$\mathbf{Adv}_\phi^{\mathrm{wcr}}(\mathcal{A}) = \mathbf{Adv}_\phi^{\mathrm{wcr}}(t, q) = \epsilon.$$

We write $u[i] = u_d[i]|| \cdots ||u_0[i]$ for the i-query that $\mathcal{A}$ makes to $f_1(\cdot)$ and $\phi(u[i]) = (x[i], y[i])$, where $x[i] = f_1(u[i])$ and $y[i] = u_d[i]x[i]^d + \cdots + u_1[i]x[i] + u_0[i]$. The i-th query is associated with a curve

$$\mathcal{C}_i = \{(x, y) \in \mathbb{F}_{2^n}^2 : y = u_d[i]x^d + \cdots + u_1[i]x + u_0[i]\}.$$

Let $\Gamma_i = \{1 \leq j \leq i : (x[j], y[j]) \in \mathcal{C}_i\}$ and let $\gamma = \max_i |\Gamma_i|$. By assumption that $\mathcal{A}$ succeeds to find a collision with probability ϵ, one of the following two events occurs with probability at least $\epsilon/2$.

Case 1: $\mathcal{A}$ finds a collision and $\gamma \leq d + \log q$. For this case, we can construct a forger $\mathcal{B}_1$ for f_1 as follows.

---

[2] More precisely, we have $M = c_{m_1} - c_{m_1+1} \geq c_{m_1}(1 - \alpha)$ or $c_{m_1} \leq M/(1 - \alpha)$. Then $\mathcal{B}$'s chance of success is at least $\frac{1}{q} \max\left(\frac{\alpha}{m_2-m_1}, \left(\frac{1-\alpha}{M}\right)^{1/(m_2-m_1)}\right)$ where we know $0 < \alpha \leq 1$.

1. $\mathcal{B}_1$ chooses $i \in \{1, \ldots, q\}$ and $s \in \{1, \ldots, d + \log q\}$ uniformly at random.
2. $\mathcal{B}_1$ runs $\mathcal{A}$ as a subroutine and faithfully answers the queries made by $\mathcal{A}$ until the $(i-1)$-th query.
3. On the i-query $u[i]$, $\mathcal{B}_1$ presents a forgery $(u[i], x[j_s])$ without making a query to $f_1(\cdot)$, where j_s is the s-th element of Γ_i. (If $|\Gamma_i| < s$, then $\mathcal{B}_1$ presents a random value.)

Note that if there exists a collision $(x[j], y[j]) = (x[i], y[i])$ for $j < i$, then $(x[j], y[j]) \in \mathcal{C}_i$ or equivalently $j \in \Gamma_i$. With this observation, we obtain

$$\mathbf{Adv}_f^{\mathsf{mac}}(\mathcal{B}_1) \geq \frac{\epsilon}{2q(d + \log q)}. \tag{6}$$

Case 2: $\mathcal{A}$ produces $\gamma > d + \log q$. This is the case where we construct a forger $\mathcal{B}_2$ for f_1 using the MTF game: The bins are $(d+1)$-tuples $(u_d, \ldots, u_0) \in \mathbb{F}_{2^n}^{d+1}$ (regarded as a curve in the plane $\mathbb{F}_{2^n}^2$) and the balls are points $(x, y) \in \mathbb{F}_{2^n}^2$. A query $f_1(u_d||\cdots||u_0)$ results in a new ball $(x, y) = (f_1(u_d||\cdots||u_0), u_d x^d + \cdots + u_1 x + u_0)$ being placed into all bins $(v_d, \ldots, v_0)$ such that $v_d x^d + \cdots + v_1 x + v_0 = y$, namely all bins giving the coefficients of a polynomial curve fitting the point (x, y), *except the bin $(u_d, \ldots, u_0)$ itself*. Thus one ball is replicated in $2^{dn} - 1$ different bins. We assume that the i-th query $u[i]$ is known to $\mathcal{B}_2$ before the i-th round of the game. Then, when $\mathcal{B}_2$ correctly guesses a bin $(v_d, \ldots, v_0)$ that will receive the new ball $(x[i], y[i])$, $\mathcal{B}_2$ has a chance to forge f_1 with probability $1/d$ since the intersection of the curves associated with $(u_d, \ldots, u_0)$ and $(v_d, \ldots, v_0)$ contains at most d elements. Here we assume the existence of an algorithm of running time Time_d to find zeros of a univariate polynomial of degree d. Let $m_1 = d$, $m_2 = d + \log q$ and $M = \binom{q}{d+1}$. Since $d + 1$ points determine a unique polynomial of degree d fitting the points, we can apply Corollary 1 to obtain a forger $\mathcal{B}_2$ of success probability

$$\mathbf{Adv}_f^{\mathsf{mac}}(\mathcal{B}_2) \geq \frac{\epsilon}{2} \cdot \frac{1}{d} \cdot \frac{1}{q} \left(\frac{1}{qM}\right)^{1/(m_2 - m_1)} \geq \frac{\epsilon}{2} \cdot \frac{1}{d} \cdot \frac{1}{q} \left(\frac{1}{q^{d+2}}\right)^{1/\log q} = \frac{\epsilon}{dq2^{d+3}}, \tag{7}$$

and of running time $O(d^2 n^2 q^{d+2}) + \mathsf{Time}_d$. From (6) and (7), it follows that

$$\mathbf{Adv}_\phi^{\mathsf{wcr}}(t, q) \leq 2q \max(d + \log q, d2^{d+2}) \mathbf{Adv}_f^{\mathsf{mac}}\left(t + O(d^2 n^2 q^{d+2}) + \mathsf{Time}_d, q\right).$$

$\square$

The following theorem is immediate from Lemma 1 and Theorem 1.

Theorem 2. *Let $H = H[f_1, f_2]$ be a function family where f_1 and f_2 are independently chosen from two function families f^1 and f^2, respectively. Then for $q = \lfloor \mu/c \rfloor + 2\tilde{q}$,*

$$\mathbf{Adv}_H^{\mathsf{mac}}(t, \tilde{q}, \mu) \leq \epsilon,$$

where

$$\epsilon = \mathbf{Adv}_{f^2}^{\mathsf{mac}}(t, \tilde{q}) + 2q \max\left(d + \log q, d2^{d+2}\right) \mathbf{Adv}_{f^1}^{\mathsf{mac}}\left(t + O(d^2 n^2 q^{d+2}) + \mathsf{Time}_d, q\right),$$

and Time_d *is the time required to solve a univariate polynomial equation of degree* d *over* $\mathbb{F}_{2^n}$. *If* $f_1 = f_2$ *are chosen from the same function family* f, *then*

$$\mathbf{Adv}_{H^*}^{\mathsf{mac}}(t, \tilde{q}, \mu) \le \epsilon + q2^{\max\{0, n-c\}} \mathbf{Adv}_f^{\mathsf{mac}}(t, q),$$

where ϵ *is as above with* f *replacing* f^1 *and* f^2.

In the single-key setting, we assume that $IV_1 \ne 0^n$ for two n-bit blocks IV_1 and IV_2 such that $IV = IV_1 || IV_2$. Then we can use the techniques employed in the CS construction [19]. The term $q2^{\max\{0, n-c\}} \mathbf{Adv}_f^{\mathsf{mac}}(t, q)$ comes from the case where $f_1(M[i]) = 0^{\min\{n, c\}} || *$ for some message block $M[i]$ during the Merkle-Damgård iteration.

4.2   Collision Resistance and Indistinguishability

Let Φ_{2n+c}^n be the set of all functions from $\{0,1\}^{2n+c}$ to $\{0,1\}^n$. Then Φ_{2n+c}^n can be regarded as a function family $f^* : \{0,1\}^\kappa \times \{0,1\}^{2n+c} \to \{0,1\}^n$ by identifying Φ_{2n+c}^n and $\{0,1\}^\kappa$ for $\kappa = n2^{2n+c}$. The weak collision resistance of ϕ defined by f^* against a computationally unbounded adversary implies its collision resistance in the information-theoretic model (due to the equivalence of oracle access to either ϕ or f^*). Since

$$\mathbf{Adv}_{f^*}^{\mathsf{mac}}(t, q) = \frac{1}{2^n}$$

for any q and t, the following theorem is immediate from Theorem 1.

Theorem 3. *If* $f_1 : \{0,1\}^{2n+c} \to \{0,1\}^n$ *is a random function, then*

$$\mathbf{Adv}_{\phi[f_1]}^{\mathsf{coll}}(q) \le \frac{\max\left(d + \log q, d2^{d+2}\right) q}{2^{n-1}}.$$

When we construct NMAC type pseudorandom oracles based on preimage aware functions, adaptive preimage resistance is only needed for the case where a distinguisher makes a query to the finalization function. If there is no interface to access the inner primitive, we do not need to worry about adaptive preimage finding. The following lemma shows that any collision resistant function can be combined with a random function producing a pseudorandom function. We give the proof in Appendix A.

Lemma 5. *Let* $F : \{0,1\}^* \to \{0,1\}^n$ *be a function with oracle access to an ideal primitive* $\mathcal{P}$ *and let* $g : \{0,1\}^n \to \{0,1\}^m$ *and* $H : \{0,1\}^* \to \{0,1\}^m$ *be random functions. Then the composite function* $g \circ F$ *is* $(\tilde{q}, \epsilon)$-*indistinguishable from* H, *where* $\epsilon = \mathbf{Adv}_F^{\mathsf{coll}}(q)$, $q = \mathsf{NQ}(F, l_{max})\tilde{q}$ *and* l_{max} *is the length in bits of the longest query made by a distinguisher.*

Since the strengthened Merkle-Damgård transform preserves collision resistance, we obtain the following theorem.

Theorem 4. *If $f_1, f_2 : \{0,1\}^{2n+c} \to \{0,1\}^n$ are random functions, then $H[f_1, f_2]$ is $(\tilde{q}, \epsilon)$-indis-tinguishable from a random function $\mathcal{H} : \{0,1\}^* \to \{0,1\}^n$, where*

$$\epsilon = \frac{\max\left(d + \log q, d2^{d+2}\right) q}{2^{n-1}},$$

and

$$q = \mathsf{NQ}(MD[\phi[f_1]], l_{max})\tilde{q} = \left\lceil \frac{l_{max} + 1}{c} + 1 \right\rceil \tilde{q},$$

for the length in bits l_{max} of the longest query made by a distinguisher. In the single-key setting, $H[f_1, f_1]$ is $(\tilde{q}, \epsilon + \frac{q}{2^c})$-indistinguishable from $\mathcal{H}$.

Lemma 5 holds with $\epsilon = \mathbf{Adv}_F^{\mathsf{wcr}}(t, \tilde{q}, l_{max}\tilde{q})$ when F is a keyed function family (in the complexity-theoretic model). This implies that $H[f_1, f_2]$ is pseudorandom up to $O(2^n/n)$ query complexity as long as f_1 is unforgeable and f_2 is pseudorandom.

4.3   Preimage Awareness and Indifferentiability

We begin with the proof of adaptive preimage resistance for $\phi[f_1]$ where f_1 is a public random function. Let $\mathcal{A}$ be an "optimal" adaptive preimage-finding adversary that makes at most q_p queries and at most q_e commitments. That is, $\mathbf{Adv}_{\phi[f_1]}^{\mathsf{adpr}}(\mathcal{A}) = \mathbf{Adv}_{\phi[f_1]}^{\mathsf{adpr}}(q_p, q_e)$. During the experiment $\mathbf{Exp}_{\mathcal{A}}^{\mathsf{adpr}}$, $\mathcal{A}$ makes queries and commitments in an arbitrarily interleaved order based on a deterministic strategy. Here we can assume that the strategy does not depend on the responses of oracle $f_1(\cdot)$ to queries that $\mathcal{A}$ sends since the probability distribution of the response to a certain query is independent of the previous query-response pairs (as long as $\mathcal{A}$ does not make a redundant query). Therefore, in order to estimate $\mathbf{Adv}_{\phi[f_1]}^{\mathsf{adpr}}(\mathcal{A})$, we can use the following game.

1. $\mathcal{A}$ makes q_p queries $\{\mathcal{C}_1, \ldots, \mathcal{C}_{q_p}\}$ and q_e commitments $\mathcal{L} = \{(x_1, y_1), \ldots, (x_{q_e}, y_{q_e})\}$ based on the optimal strategy. (Here each query is represented by a curve in $\mathbb{F}_{2^n}^2$ as in the analysis of unforgeability.)
2. One point (x_i^*, y_i^*) is chosen from each curve $\mathcal{C}_i$ uniformly at random.
3. If $(x_i^*, y_i^*) \in \mathcal{L}$ for some $i = 1, \ldots, q_e$, then $\mathcal{A}$ wins. Otherwise, $\mathcal{A}$ loses.

The winning probability of $\mathcal{A}$ for the above game is equal to the adaptive preimage-finding advantage of $\mathcal{A}$. Let $\Gamma_i = \mathcal{C}_i \cap \mathcal{L}$ for $i = 1, \ldots, q_p$, and let

$$\Delta_\theta = \{1 \leq i \leq q_p : |\Gamma_i| \geq \theta\},$$

for a parameter θ. Then for $\Delta \subset \Delta_\theta$ and $\delta = |\Delta|$, we have

$$q_e \geq \left| \bigcup_{i \in \Delta} \Gamma_i \right| \geq \sum_{i \in \Delta} |\Gamma_i| - \sum_{i \neq j \in \Delta} |\Gamma_i \cap \Gamma_j| \geq \delta\theta - \binom{\delta}{2} \cdot d.$$

Therefore we conclude that $|\Delta_{\theta^*}| < \delta^*$ for any (θ^*, δ^*) such that

$$\delta^* \theta^* - \binom{\delta^*}{2} \cdot d > q_e. \tag{8}$$

This implies that the number of curves that intersect with $\mathcal{L}$ at $\geq \theta^*$ points is less than δ^*. Thus the winning probability of $\mathcal{A}$ is upper-bounded by

$$\mathbf{Adv}^{\mathsf{adpr}}_{\phi[f_1]}(\mathcal{A}) \leq \delta^* \frac{q_e}{2^n} + (q_p - \delta^*) \frac{\theta^*}{2^n}.$$

By Lemma 2 and Theorem 3, we obtain the following theorem.

Theorem 5. *Let (θ^*, δ^*) satisfy inequality (8). Then for a random function f_1, there exists an extractor $\mathcal{E}^*$ such that for any adversary $\mathcal{A}$ it holds that*

$$\mathbf{Adv}^{\mathsf{pra}}_{\phi[f_1]}(\mathcal{A}, \mathcal{E}^*) \leq \frac{\max\left(d + \log q_p, d2^{d+2}\right) q_p}{2^{n-1}} + \delta^* \frac{q_e}{2^n} + (q_p - \delta^*) \frac{\theta^*}{2^n}.$$

We can use Lemma 3 and Theorem 5 with $(\theta^*, \delta^*) = (dq_e^{1/2}, q_e^{1/2})$ to obtain the following theorem.

Theorem 6. *Let $f_1, f_2 : \{0,1\}^{2n+c} \to \{0,1\}^n$ and $\mathcal{H} : \{0,1\}^* \to \{0,1\}^n$ be public random functions. Then there exists a simulator $\mathcal{S}$ such that for any distinguisher $\mathcal{A}$ making at most (q_0, q_1, q_2) queries to the three oracle interfaces associated with $(\mathcal{H}, f_1, f_2)$,*

$$\left| \mathbf{Pr}\left[\mathcal{A}^{H[f_1, f_2],(f_1, f_2)} = 1\right] - \mathbf{Pr}\left[\mathcal{A}^{\mathcal{H}, \mathcal{S}[\mathcal{H}]} = 1\right] \right| \leq \epsilon,$$

where

$$\epsilon = \frac{\max\left(d + \log\left(Lq_0 + q_1 + L\right), d2^{d+2}\right)}{2^{n-1}} \left(Lq_0 + q_1 + L\right) + \frac{L^{1/2} q_2^{1/2}}{2^n} \left(dLq_0 + dq_1 + Lq_2 + dL\right),$$

l_{max} is the length in bits of the longest query made by $\mathcal{A}$ to its first oracle and $L = \left\lceil \frac{l_{max}+1}{c} + 1 \right\rceil$. Simulator $\mathcal{S}$ runs in time $O\left(q_1 + Lq_2 \mathsf{Time}(\mathcal{E}^*) + Lq_2 \mathsf{Time}(\mathsf{unpad})\right)$, where $\mathcal{E}^*$ is the obvious extractor used in Lemma 2. In the single-key setting, we have

$$\left| \mathbf{Pr}\left[\mathcal{A}^{H[f_1, f_1], f_1} = 1\right] - \mathbf{Pr}\left[\mathcal{A}^{\mathcal{H}, \mathcal{S}'[\mathcal{H}]} = 1\right] \right| \leq \epsilon + \frac{Lq_0}{2^c},$$

where simulator $\mathcal{S}'$ is obtained by a slight modification of $\mathcal{S}$: On input x, $\mathcal{S}'$ returns $f_2(x)$ if $x = 0^c || *$ for some $* \in \{0,1\}^{2n}$ and returns $f_1(x)$ otherwise, by using the interfaces f_1 and f_2 of $\mathcal{S}$.

Tightness of Indifferentiability. The preservation of indifferentiability is guaranteed only up to $O(2^{2n/3})$ query complexity which is beyond the birthday bound but still far from optimal. This bound is dominated by the adaptive preimage resistance, depending on a configuration that consists of q curves in $\mathbb{F}_{2^n}^2$ and q points on the curves (assuming $q = q_p = q_e$). If there exists a subfield

$\mathbb{F}'$ of $\mathbb{F}_{2^n}$ such that $|\mathbb{F}'| = \sqrt{q}$, then we have a configuration that provides tight adaptive preimage resistance: The set of points is $\mathbb{F}' \times \mathbb{F}' \subset \mathbb{F}_{2^n}^2$ and the set of curves consists of q polynomials of degree d with coefficients in $\mathbb{F}'$. However, for the case where $\mathbb{F}_{2^n}$ admits no proper subfield (e.g. with prime n), there remains a question whether a similar construction exists or the bound can be qualitatively improved. We pose this as an open problem.

5   The Quadratic Blockcipher-Based Compression Function

In this section and the next, we discuss how to instantiate $\phi[f_1]$ for $c = n$ (the "quadratic" polynomial mode) by replacing the $3n$-bit to n-bit compression function f_1 with a smaller primitive. First, we discuss a concrete instantiation of the quadratic compression function using a blockcipher with $2n$-bit keys. Given $f_1 : \{0,1\}^{3n} \to \{0,1\}^n$, the compression function $\phi[f_1] : \{0,1\}^{3n} \to \{0,1\}^{2n}$ is defined by $\phi[f_1](u_2||u_1||u_0) = x||y$, where $x = f_1(u_2||u_1||u_0)$ and $y = u_2 x^2 + u_1 x + u_0 \in \mathbb{F}_{2^n}$ for $u_2, u_1, u_0 \in \{0,1\}^n$. In the *quadratic blockcipher-based compression function*, f_1 is implemented using a blockcipher $E : \{0,1\}^{2n} \times \{0,1\}^n \to \{0,1\}^n$, $E(k,x) = E_k(x)$ by letting $f_1(u_2||u_1||u_0) = E_{u_2||u_1}(u_0) + u_0$ as described in Figure 2(a). We write $\psi[E]$ for the resulting compression function. Thus $\psi[E] : \{0,1\}^{3n} \to \{0,1\}^n$ and $\psi[E](u_2||u_1||u_0) = x||y$ where

$$x = E_{u_2||u_1}(u_0) + u_0,$$
$$y = u_2 x^2 + u_1 x + u_0.$$

We can prove that $\psi[E]$ provide similar security as the quadratic mode $\phi[f_1]$ when instantiated with an ideal cipher E, in terms of unforgeability, collision resistance and pseudorandomness. In fact, our results do not actually necessitate an ideal cipher E (which is a set of independent random permutations with one permutation per key) but only an "unpredictable" blockcipher E. For the latter, all that is assumed is that it is difficult for an adversary to fully predict the output of an unqueried value. We call this the *unpredictability* of the blockcipher (which is similar to the unforgeability of a keyed function family, except no keys are involved) and we quantify it by the advantage $\mathbf{Adv}_E^{\mathsf{unp}}(t,q)$ which is the maximum over all adversaries A running in time t and making at most q queries to E of the probability that A can output a tuple $(u_2||u_1, v, w)$ such that $E_{u_2||u_1}(v) = w$ without making queries for $E_{u_2||u_1}(v)$ or $E_{u_2||u_1}^{-1}(w)$.

Implicitly $\mathbf{Adv}_E^{\mathsf{unp}}(t,q)$ depends on a sampling procedure for E. In the ideal cipher model, E is sampled uniformly at random among all n-bit blockciphers with $2n$-bit keys. Here we allow any sampling procedure for E. Note that $\mathbf{Adv}_E^{\mathsf{unp}}(t,q) \leq 1/(2^n - q)$ if E is an ideal cipher, so we can always revert to that bound by assuming an ideal cipher. Our use of unpredictable blockciphers is somewhat similar to that of [7], with the significant difference that the blockciphers of [7] use fixed keys, and that they are sampled by sampling the fixed keys. The unpredictability then corresponds to the unforgeability of a keyed function family (which happens to be a family of permutations).

We show that the collision resistance of $\psi[E]$ can be effectively bounded in terms of $\mathbf{Adv}_E^{\mathsf{unp}}(t, q)$. Let $\mathbf{Adv}_{\psi[E]}^{\mathsf{coll}}(t, q)$ be the maximum probability that an adversary A of running time t with oracle access to E and E^{-1} outputs a collision (M, M') for $\psi[E]$ which it has verified (i.e. has made the queries necessary to compute $\psi[E](M)$ and $\psi[E](M')$). The following theorems are proved in the full version.

Theorem 7. *Let $\psi[E]$ be the quadratic blockcipher-based compression function, where E is sampled from the set of all n-bit blockciphers with $2n$-bit keys according to an arbitrary fixed distribution. Then,*

$$\mathbf{Adv}_{\psi[E]}^{\mathsf{coll}}(t, q) \leq 2q(\log q + 3)\mathbf{Adv}_E^{\mathsf{unp}}(t + O(n^2 q^4) + \mathsf{Time}_2, q).$$

Furthermore, let $G = G[E, f_2] = f_2 \circ MD[\psi[E]]$ be a function family where f_2 is chosen from a function family $f : \{0,1\}^\kappa \times \{0,1\}^{2n} \to \{0,1\}^n$. Then for $q = \lfloor \mu/c \rfloor + 2\tilde{q}$,

$$\mathbf{Adv}_G^{\mathsf{mac}}(t, \tilde{q}, \mu) \leq \mathbf{Adv}_f^{\mathsf{mac}}(t, \tilde{q}) + 2q(\log q + 3)\mathbf{Adv}_E^{\mathsf{unp}}(t + O(n^2 q^4) + \mathsf{Time}_2, q).$$

Remark 3. The term "Time_2", which represents the time necessary to select a root of a quadratic polynomial over $\mathbb{F}_{2^n}$, is mainly kept to facilitate comparison with Theorem 1.

Theorem 8. *Let $E : \{0,1\}^{2n} \times \{0,1\}^n \to \{0,1\}^n$ be an ideal blockcipher and let $f_2 : \{0,1\}^{2n} \to \{0,1\}^n$ be a random functions. Then, $\mathbf{Adv}_{\psi[E]}^{\mathsf{coll}}(q) \leq \epsilon(q)$ for*

$$\epsilon(q) = \frac{2q(\log q + 3)}{2^n - q},$$

and $G[E, f_2] = f_2 \circ MD[\psi[E]]$ is $(\tilde{q}, \epsilon(\bar{q}))$-indistinguishable from a random function $\mathcal{H} : \{0,1\}^* \to \{0,1\}^n$, where

$$\bar{q} = \mathsf{NQ}(MD[\psi[E]], l_{max})\tilde{q} = \left\lceil \frac{l_{max} + 1}{n} + 1 \right\rceil \tilde{q},$$

for the length in bits l_{max} of the longest query made by a distinguisher.

6   The Quadratic Cascade Compression Function

In this section, we discuss a concrete instantiation of the quadratic compression function (polynomial-based compression function of degree $d = 2$) using the cascade of two $2n \to n$ functions. In the *quadratic cascade compression function*, the compression function f_1 is implemented by the cascade of two compression functions $h_1, h_2 : \{0,1\}^{2n} \times \{0,1\}^n \to \{0,1\}^n$ by letting $f_1(u_2||u_1||u_0) = h_2(h_1(u_2||u_1)||u_0)$ as described in Figure 2(b). We write $\tau[h_1, h_2]$ for the resulting compression function. Thus $\tau[h_1, h_2] : \{0,1\}^{3n} \to \{0,1\}^n$ and $\tau[h_1, h_2](u_2||u_1|| u_0) = x||y$, where

$$x = h_2(h_1(u_2||u_1)||u_0),$$
$$y = u_2 x^2 + u_1 x + u_0.$$

Now we have the following theorems. Due to lack of space, the proof will be given in the full version of this paper. (The proof has similar—but slightly worse—complexity as the security proof for SS-NMAC in [7]. It involves 13 different cases for the forger, 8 of which use MTF games.)

Theorem 9. *Let $\tau = \tau[h_1, h_2]$ be a function family where h_1 and h_2 are independently chosen from a function family h, respectively. Then,*

$$\mathbf{Adv}_{\tau}^{\mathsf{wcr}}(t, q) \leq 26q \log q \left(\log q + 1\right)^2 \mathbf{Adv}_h^{\mathsf{mac}}\left(t + O(n^2 q^4 \log^4 q) + \mathsf{Time}_2, q\right).$$

Furthermore, let $G = G[h_1, h_2, f_2] = f_2 \circ MD\left[\tau[h_1, h_2]\right]$ be a function family where f_2 is chosen from a function family f. Then for $q = \lfloor \mu/c \rfloor + 2\tilde{q}$,

$$\mathbf{Adv}_G^{\mathsf{mac}}(t, \tilde{q}, \mu) \leq \mathbf{Adv}_f^{\mathsf{mac}}(t, \tilde{q}) + 26q \log q \left(\log q + 1\right)^2 \mathbf{Adv}_h^{\mathsf{mac}}\left(t + O(n^2 q^4 \log^4 q) + \mathsf{Time}_2, q\right).$$

Theorem 10. *If $h_1, h_2, f_2 : \{0,1\}^{2n} \rightarrow \{0,1\}^n$ are random functions, then* $\mathbf{Adv}_{\tau[h_1, h_2]}^{\mathsf{coll}}(q) \leq \epsilon(q)$ *for*

$$\epsilon(q) = \frac{26q \log q \left(\log q + 1\right)^2}{2^n},$$

and $G[h_1, h_2, f_2] = f_2 \circ MD\left[\tau[h_1, h_2]\right]$ is $(\tilde{q}, \epsilon(\bar{q}))$-indistinguishable from a random function $\mathcal{H} : \{0,1\}^* \rightarrow \{0,1\}^n$, where

$$\bar{q} = \mathsf{NQ}(MD[\tau[h_1, h_2]], l_{max})\tilde{q} = 2 \left\lceil \frac{l_{max} + 1}{n} + 1 \right\rceil \tilde{q},$$

for the length in bits l_{max} of the longest query made by a distinguisher.

References

1. An, J.H., Bellare, M.: Constructing VIL-MACs from FIL-MACs: Message authentication under weakened assumptions. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 252–269. Springer, Heidelberg (1999)
2. Bellare, M., Ristenpart, T.: Multi-property-preserving Hash Domain Extension and the EMD Transform. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 299–314. Springer, Heidelberg (2006)
3. Coron, J., Dodis, Y., Malinaud, C., Puniya, P.: Merkle-Damgård revisited: How to construct a hash function. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 430–448. Springer, Heidelberg (2005)
4. Damgård, I.: A design principle for hash functions. In: Brassard, G. (ed.) CRYPTO 1989. LNCS, vol. 435, pp. 416–427. Springer, Heidelberg (1990)
5. Dodis, Y., Pietrzak, K., Puniya, P.: A new mode of operation for block ciphers and length-preserving MACs. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 198–219. Springer, Heidelberg (2008)
6. Dodis, Y., Ristenpart, T., Shrimpton, T.: Salvaging Merkle-Damgård for practical applications. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 371–388. Springer, Heidelberg (2009)
7. Dodis, Y., Steinberger, J.: Message authentication codes from unpredictable block ciphers. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 267–285. Springer, Heidelberg (2009)

8. von zur Gathen, J., Panario, D.: Factoring polynomials over finite fields: A survey. J. Symbolic computation 31, 3–17 (2001)
9. von zur Gathen, J., Shoup, V.: Computing Frobenius maps and factoring polynomials. Computational complexity 2, 187–224 (1992)
10. Hirose, S.: Some plausible constructions of double length hash functions. In: Robshaw, M.J.B. (ed.) FSE 2006. LNCS, vol. 4047, pp. 210–225. Springer, Heidelberg (2006)
11. Joux, A.: Multicollisions in iterated hash functions. Application to cascaded constructions. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 306–316. Springer, Heidelberg (2004)
12. Kelsey, J., Kohno, T.: Herding hash functions and the Nostradmus attack. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 183–200. Springer, Heidelberg (2006)
13. Kelsey, J., Schneier, B.: Second preimages on n-bit hash functions for much less than 2^n work. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 474–490. Springer, Heidelberg (2005)
14. Lai, X., Massey, J.: Hash function based on block ciphers. In: Rueppel, R.A. (ed.) EUROCRYPT 1992. LNCS, vol. 658, pp. 55–70. Springer, Heidelberg (1993)
15. Lucks, S.: A failure-freindly design principle for hash functions. In: Roy, B. (ed.) ASIACRYPT 2005. LNCS, vol. 3788, pp. 474–494. Springer, Heidelberg (2005)
16. Lucks, S.: A collision-resistant rate-1 double-block-length hash function. In: Symmetric Cryptography, Dagstuhl Seminar Proceedings 07021 (2007)
17. Merkle, R.: One way hash functions and DES. In: Brassard, G. (ed.) CRYPTO 1989. LNCS, vol. 435, pp. 428–446. Springer, Heidelberg (1990)
18. Maurer, U., Renner, R., Holenstein, R.: Indifferentiability, impossibility results on reductions, and apllications to the random oracle methodology. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 21–39. Springer, Heidelberg (2004)
19. Maurer, U., Sjödin, J.: Single-key AIL-MACs from any FIL-MAC. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 472–484. Springer, Heidelberg (2005)
20. Özen, O., Stam, M.: Another glance at double length hashing. In: Parker, M.G. (ed.) Cryptography and Coding 2009. LNCS, vol. 5921, pp. 176–201. Springer, Heidelberg (2009)
21. Rogaway, P., Steinberger, J.: Constructing cryptographic hash functions from fixed-key blockciphers. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 433–450. Springer, Heidelberg (2008)
22. Stam, M.: Beyond uniformity: Security/efficiency tradeoffs for compression functions. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 397–412. Springer, Heidelberg (2008)
23. Stam, M.: Blockcipher based hashing revisited. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 67–83. Springer, Heidelberg (2009)
24. Yasuda, K.: A double-piped mode of operation for MACs, PRFs and PROs: Security beyond the birthday barrier. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 242–259. Springer, Heidelberg (2009)

A   The Proof of Lemma 5

Let $\mathcal{G}_0$ and $\mathcal{G}_1$ be games with a single interface, as defined in Figure 6. Assume that a distinguisher $\mathcal{A}$ makes no redundant query. Then whenever $\mathcal{A}$ makes
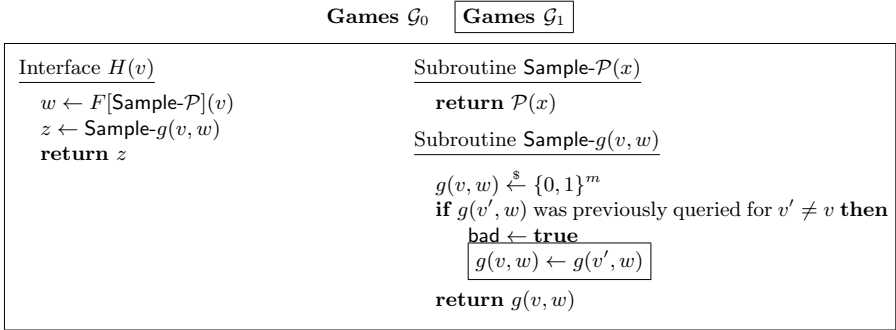
Games $\mathcal{G}_0$ $\boxed{\textbf{Games } \mathcal{G}_1}$

Interface $H(v)$	Subroutine $\mathsf{Sample}\text{-}\mathcal{P}(x)$

Interface $H(v)$

   $w \leftarrow F[\mathsf{Sample}\text{-}\mathcal{P}](v)$
   $z \leftarrow \mathsf{Sample}\text{-}g(v, w)$
   return z

Subroutine $\mathsf{Sample}\text{-}\mathcal{P}(x)$

   return $\mathcal{P}(x)$

Subroutine $\mathsf{Sample}\text{-}g(v, w)$

   $g(v, w) \xleftarrow{\$} \{0, 1\}^m$
   if $g(v', w)$ was previously queried for $v' \neq v$ **then**
      bad $\leftarrow$ **true**
      $\boxed{g(v, w) \leftarrow g(v', w)}$
   return $g(v, w)$

Fig. 6. Games $\mathcal{G}_0$ and $\mathcal{G}_1$. $\mathcal{G}_1$ includes the boxed statement

a query to $H(\cdot)$ in game $\mathcal{G}_0$, it will receive an independent random value in $\{0, 1\}^m$. It means that the interface $H(\cdot)$ faithfully implements a random function $H : \{0, 1\}^* \rightarrow \{0, 1\}^m$ in game $\mathcal{G}_0$. On the other hand, the interface $H(\cdot)$ in game $\mathcal{G}_1$ faithfully implements $g \circ F[\mathcal{P}] : \{0, 1\}^* \rightarrow \{0, 1\}^m$ for a random function $g : \{0, 1\}^n \rightarrow \{0, 1\}^m$ since $\mathsf{Sample}\text{-}g(v, w)$ only depends on the value v. Flag bad sets to true only when $\mathcal{A}$ makes a collision in $F[\mathcal{P}]$. Therefore, for any distinguisher $\mathcal{A}$ that makes at most $\tilde{q}$ queries to $H(\cdot)$, we have

$$\left| \mathbf{Pr}\left[\mathcal{A}^{g \circ F[\mathcal{P}]} = 1\right] - \mathbf{Pr}\left[\mathcal{A}^H = 1\right] \right| \leq \mathbf{Pr}\left(\mathcal{A}^{\mathcal{G}_1} \text{ sets bad}\right) \leq \mathbf{Adv}_F^{\mathsf{coll}}(q),$$

where $q = \mathsf{NQ}(F, l_{max})\tilde{q}$ and l_{max} is the length in bits of the longest query made by a distinguisher.

B Improved Analysis of Lucks' Double-Piped Mode of Operation

We begin with the definition of Lucks' double-piped mode of operation (with a slight modification). First, two $2n + c \rightarrow n$ bit functions f_1 and f_2 are concatenated, yielding the following compression function.

$$F : \{0, 1\}^{2n+c} \longrightarrow \{0, 1\}^{2n}$$

$$u \longmapsto f_1(u) || f_2(u).$$

The pictorial description of F for $c = n$ is shown in Figure 7. Given $F = F[f_1, f_2]$ and an independent compression function $f_3 : \{0, 1\}^{2n+c} \rightarrow \{0, 1\}^n$, the Lucks' mode of operation defines a hash function

$$G[f_1, f_2, f_3] : \{0, 1\}^* \longrightarrow \{0, 1\}^n$$

$$M \longmapsto f_3(0^c || v),$$

where $v = MD[F](M)$.

Fig. 7. Lucks' double-piped compression function

Now we prove the weak collision resistance of $F = F[f_1, f_2]$ where f_1 and f_2 are independently chosen from a function family f.

Theorem 11. *Let F be a function family defined by $f : \{0,1\}^{\kappa} \times \{0,1\}^{2n+c} \to \{0,1\}^n$ as described above. Then,*

$$\mathbf{Adv}_F^{\mathsf{wcr}}(t,q) \leq \sqrt{2}q\mathbf{Adv}_f^{\mathsf{mac}}(t,q).$$

Proof. Let $\mathcal{A}$ be a weak collision-finding adversary such that

$$\mathbf{Adv}_F^{\mathsf{wcr}}(\mathcal{A}) = \mathbf{Adv}_F^{\mathsf{wcr}}(t,q) = \epsilon.$$

We write $u[i]$ for the i-query that $\mathcal{A}$ makes to F and $F(u[i]) = (x[i], y[i])$, where $x[i] = f_1(u[i])$ and $y[i] = f_2(u[i])$. Let

$$\Gamma = \left\{ (j,i) \in \{1,\dots,q\}^2 : j < i \text{ and } x[j] = x[i] \right\},$$

and let $\gamma = |\Gamma|$ be the number of "upper half-collisions". Here we fix an ordering in Γ. By assumption that $\mathcal{A}$ succeeds to find a collision with probability ϵ, at least one of the following two events happens with probability $\geq \epsilon/2$, where θ is a parameter to be optimized later.

Case 1: $\mathcal{A}$ finds a collision and $\gamma > \theta$. For this case, we can construct a forger $\mathcal{B}_1$ for f_1 as follows.

1. $\mathcal{B}_1$ chooses $(j,i) \in \{1,\dots,q\}^2$ such that $j < i$ uniformly at random.
2. $\mathcal{B}_1$ runs $\mathcal{A}$ as a subroutine and faithfully answers the queries made by $\mathcal{A}$ until the $(i-1)$-th query. Here $\mathcal{B}_1$ simulates f_2 by choosing a key for f_2 uniformly at random.
3. On the i-query $u[i]$, $\mathcal{B}_1$ presents a forgery $(u[i], y[j])$ without making a query to $f_1(\cdot)$.

Since the probability that $(j,i) \in \Gamma$ is $\theta/\binom{q}{2}$, we have

$$\mathbf{Adv}_f^{\mathsf{mac}}(\mathcal{B}_1) \geq \frac{\epsilon\theta}{2\binom{q}{2}} \geq \frac{\epsilon\theta}{q^2}. \tag{9}$$

Case 2: $\mathcal{A}$ finds a collision and $\gamma \leq \theta$. For this case, we can construct a forger $\mathcal{B}_2$ for f_2 as follows.

1. $\mathcal{B}_2$ chooses $s \in \{1, \ldots, \theta\}$ uniformly at random.
2. $\mathcal{B}_2$ runs $\mathcal{A}$ as a subroutine: To each query made by $\mathcal{A}$, $\mathcal{B}_2$ faithfully respond by simulating f_1 and making queries to $f_2(\cdot)$.
3. $\mathcal{B}_2$ counts the number of collisions in f_1. At the s-th collision $(j, i) \in \Gamma$, $\mathcal{B}_2$ stops without making a query to $f_2(\cdot)$ and presents a forgery $(u[i], y[j])$.

If there exists a collision $(x[j], y[j]) = (x[i], y[i])$ for $j < i$, then obviously $(j, i) \in \Gamma$. Therefore we have

$$\mathbf{Adv}_f^{\mathsf{mac}}(\mathcal{B}_2) \geq \frac{\epsilon}{2\theta}. \tag{10}$$

From (9) and (10), it follows that

$$\mathbf{Adv}_F^{\mathsf{wcr}}(t, q) \leq \max\left\{\frac{q^2}{\theta}, 2\theta\right\} \mathbf{Adv}_f^{\mathsf{mac}}(t, q).$$

By setting $q^2/\theta = 2\theta$ or $\theta = q/\sqrt{2}$, we obtain

$$\mathbf{Adv}_F^{\mathsf{wcr}}(t, q) \leq \sqrt{2}q\mathbf{Adv}_f^{\mathsf{mac}}(t, q). \qquad \square$$

By Lemma 1 and Theorem 11, we obtain the following theorem.

Theorem 12. *Let $G = G[f_1, f_2, f_3]$ be a function family such that f_1, f_2 and f_3 are independently chosen from a function family f. Then for $q = \lfloor \mu/c \rfloor + 2\tilde{q}$,*

$$\mathbf{Adv}_G^{\mathsf{mac}}(t, \tilde{q}, \mu) \leq \mathbf{Adv}_f^{\mathsf{mac}}(t, \tilde{q}) + \sqrt{2}q\mathbf{Adv}_f^{\mathsf{mac}}(t, q).$$

With a slight modification of the above argument, we can prove that the Lucks' mode of operation *using a single key* also preserves unforgeability up to $O(2^n)$ query complexity, improving the bound $O(2^{5n/6})$ proved by Yasuda [24].