

Optimal Collision Security in Double Block Length Hashing with Single Length Key

Bart Mennink

Dept. Electrical Engineering, ESAT/COSIC, KU Leuven, and IBBT, Belgium
bart.mennink@esat.kuleuven.be

Abstract. The idea of double block length hashing is to construct a compression function on $2n$ bits using a block cipher with an n -bit block size. All optimally secure double length hash functions known in the literature employ a cipher with a key space of double block size, $2n$ -bit. On the other hand, no optimally secure compression functions built from a cipher with an n -bit key space are known. Our work deals with this problem. Firstly, we prove that for a wide class of compression functions with two calls to its underlying n -bit keyed block cipher collisions can be found in about $2^{n/2}$ queries. This attack applies, among others, to functions where the output is derived from the block cipher outputs in a linear way. This observation demonstrates that all security results of designs using a cipher with $2n$ -bit key space crucially rely on the presence of these extra n key bits. The main contribution of this work is a proof that this issue can be resolved by allowing the compression function to make one extra call to the cipher. We propose a family of compression functions making three block cipher calls that asymptotically achieves optimal collision resistance up to $2^{n(1-\varepsilon)}$ queries and preimage resistance up to $2^{3n(1-\varepsilon)/2}$ queries, for any $\varepsilon > 0$. To our knowledge, this is the first optimally collision secure double block length construction using a block cipher with single length key space.

1 Introduction

Double (block) length hashing is a well-established method for constructing a compression function with $2n$ -bit output based only on n -bit block ciphers. The idea of double length hashing dates back to the work of Meyer and Schilling [19], with the introduction of the MDC-2 and MDC-4 compression functions in 1988. In recent years, the design methodology got renewed attention in the works of [2,4,7,9,10,12,16,21,27]. Double length hash functions have an obvious advantage over classical block cipher based functions such as Davies-Meyer and Matyas-Meyer-Oseas [22,26]: the same type of underlying primitive allows for a larger compression function. Yet, for double length compression functions it is harder to achieve optimal n -bit collision and $2n$ -bit preimage security.

We focus on the simplest and most-studied type of compression functions, namely functions that compress $3n$ to $2n$ bits. Those can be classified into two classes: compression functions that internally evaluate a $2n$ -bit keyed block cipher $E : \{0, 1\}^{2n} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ (which we will call the DBL ^{$2n$} class), and

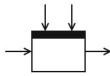
ones that employ an n -bit keyed block cipher $E : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ (the DBL^n class). The DBL^{2n} class is well understood. It includes the classical compression functions Tandem-DM and Abreast-DM [8] and Hirose’s function [6], as well as Stam’s supercharged single call Type-I compression function design [25,26] (reconsidered in [14]) and the generalized designs by Hirose [5] and Özen and Stam [21]. As illustrated in Table 1, all of these functions provide optimal collision security guarantees (up to about 2^n queries), and Tandem-DM, Abreast-DM, and Hirose’s function are additionally proven optimally preimage resistant (up to about 2^{2n} queries). These bounds also hold in the iteration, when a proper domain extender is applied [1]. Lucks [15] introduced a compression function that allows for collisions in about $2^{n/2}$ queries, but achieves optimal collision resistance in the iteration. Members of the DBL^n class are the MDC-2 and MDC-4 compression functions [19], the MJH construction [10], and a construction by Jetchev et al. [7]. For the MDC-2 and MJH compression functions, collisions and preimages can be found in about $2^{n/2}$ and 2^n queries, respectively¹. The MDC-4 compression function achieves a higher level of collision and preimage resistance than MDC-2 [16], but contrary to the other functions it makes four block cipher calls. Jetchev et al.’s construction makes two block cipher calls and achieves $2^{2n/3}$ collision security. Stam also introduced a design based on two calls, and proved it optimally collision secure in a restricted security model where the adversary must fix its queries in advance. Therefore we did not include this design in the table. Further related results include the work of Nandi et al. [20], who presented a $3n$ -to- $2n$ -bit compression function making three calls to a $2n$ -to- n -bit one-way function, achieving collision security up to $2^{2n/3}$ queries. They extended this result to a $4n$ -to- $2n$ -bit function using three $2n$ -bit keyed block ciphers.

Unlike the DBL^{2n} class, for the DBL^n class no optimally secure compression function is known. The situation is the same for the iteration, where none of these designs has been proven to achieve optimal security. Determinative to this gap is the difference in the underlying primitive: in the DBL^{2n} class, the underlying primitive maps $3n$ bits to n bits and thus allows for more compression. In particular, if we consider Tandem-DM, Abreast-DM, and Hirose’s function, the first cipher call already compresses the entire input to the compression function, and the second cipher call is simply used to assure a $2n$ -bit output. In fact, these designs achieve their level of security merely due to this property, for their proofs crucially rely on this (see also Sect. 4).

Thus, from a theoretical point of view it is unreasonable to compare DBL^{2n} and DBL^n . But the gap between the two classes leaves us with an interesting open problem: starting from a single block cipher $E : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$, is it possible to construct a double length compression function that achieves optimal collision and preimage security? This is the central research question of this work. Note that Stam’s bound [25] does not help us here: it claims that collisions can be found in at most $(2^n)^{(2r-1)/(r+1)}$ queries, where r denotes the

¹ In the iteration collision resistance is proven up to $2^{3n/5}$ queries for MDC-2 [27] and $2^{2n/3}$ queries for MJH [10].

Table 1. Asymptotic ideal cipher model security guarantees of known double length compression functions in the classes DBL^{2n} (first) and DBL^n (next). A more detailed comparison of some of these functions can be found in [3, App. A].

compression function	E -calls	collision security	preimage security	underlying cipher
Lucks'	1	$2^{n/2}$	2^n	
Stam's	1	2^n [26]	2^n [26]	
Tandem-DM	2	2^n [12]	2^{2n} [2,13]	
Abreast-DM	2	2^n [4,9]	2^{2n} [2,13]	
Hirose's	2	2^n [6]	2^{2n} [2,13]	
Hirose-class	2	2^n [5]	2^n [5]	
Özen-Stam-class	2	2^n [21]	2^n [21]	
MDC-2	2	$2^{n/2}$	2^n	
MJH	2	$2^{n/2}$	2^n	
Jetchev et al.'s	2	$2^{2n/3}$ [7]	2^n [7]	
MDC-4	4	$2^{5n/8}$ [16]	$2^{5n/4}$ [16]	
Our proposal	3	2^n	$2^{3n/2}$	

number of block cipher calls, which results in the trivial bound for $r \geq 2$. For $r \geq 2$, denote by $F^r : \{0, 1\}^{3n} \rightarrow \{0, 1\}^{2n}$ a compression function that makes r calls to its primitive E .

As a first contribution, we consider F^2 , and prove that for a very large class of functions of this form one expects collisions in approximately $2^{n/2}$ queries. Covered by the attack are among others designs with linear finalization function (the function that produces the $2n$ -bit output given the $3n$ -bit input and the block cipher responses). We note that the compression function by Jetchev et al. [7] is not vulnerable to the attack due to its non-linear finalization function. Nevertheless, these results strengthen the claim that no practical optimally collision secure F^2 function exists. Motivated by this, we increase the number of calls to E , and consider F^3 . In this setting, we derive a family of compression functions which we prove asymptotically optimal collision resistant up to $2^{n(1-\varepsilon)}$ queries and preimage resistant up to $2^{3n(1-\varepsilon)/2}$ queries, for any $\varepsilon > 0$. Our compression function family, thus, achieves the same level of collision security as the well-established Tandem-DM, Abreast-DM, and Hirose's function, albeit based on a much weaker assumption. In the DBL^n class, our design clearly compares favorably to MDC-4 that makes four block cipher evaluations, and from a provable security point of view it beats MDC-2 and MJH, still, an extra E evaluation has to be made which results in an efficiency loss. The introduced class of compression functions is simple and easy to understand: they are defined by 4×4 matrices over the field $GF(2^n)$ which are required to comply with easily satisfied conditions. Two example compression functions in this class are given in Fig. 1.

The security proofs of our compression function family rely on basic principles from previous proofs, but in order to accomplish optimal collision security (and




Fig. 1. Two example compression functions from the family of functions introduced and evaluated in this work. For these constructions, all wires carry $n = 128$ bits, and the arithmetic is done over $GF(2^{128})$. We further elaborate on these designs and their derivations in Sect. 4.

as our designs use n -bit keyed block ciphers) our proofs have become significantly more complex. The security proofs of all known DBL^{2n} functions (see Table 1) crucially rely on the property that one block cipher evaluation defines the input to the second one. For F^3 this cannot be achieved as each primitive call fixes at most $2n$ bits of the function input. Although one may expect this to cause an optimal proof to become unlikely, this is not the case. Using a new proof approach—we smartly apply the methodology of “wish lists” (by Armknecht et al. and Lee et al. [2,13]) to collision resistance—we manage to achieve asymptotically the close to 2^n collision security for our family of functions. Nonetheless, the bound on preimage resistance does not reach the optimal level of 2^{2n} queries. One can see this as the price we pay for using single key length rather than double key length block ciphers: a straightforward generalization of the pigeonhole-birthday attack of Rogaway and Steinberger [24] shows that, when the compression function behaves “sufficiently random”, one may expect a preimage in approximately $2^{5n/3}$ queries (cf. Sect. 2). The asymptotic preimage bound of $2^{3n/2}$ found in this work closely approaches this generic bound.

Outline. We present and formalize the security model in Sect. 2. Then, in Sect. 3 we derive our impossibility result on F^2 . We propose and analyze our family of compression functions in Sects. 4 and 5. This work is concluded in Sect. 6.

2 Security Model

For $n \geq 1$, we denote by $\text{Bloc}(n)$ the set of all block ciphers with a key and message space of n bits. Let $E \in \text{Bloc}(n)$. For $r \geq 1$, let $F^r : \{0, 1\}^{3n} \rightarrow \{0, 1\}^{2n}$

be a double length compression function making r calls to its block cipher E . We can represent F^r by mappings $f_i : \{0, 1\}^{(i+2)n} \rightarrow \{0, 1\}^{2n}$ for $i = 1, \dots, r+1$ as follows:

```

 $F^r(u, v, w)$ 
for  $i = 1, \dots, r$ :
     $(k_i, m_i) \leftarrow f_i(u, v, w; c_1, \dots, c_{i-1}),$ 
     $c_i \leftarrow E(k_i, m_i),$ 
return  $(y, z) \leftarrow f_{r+1}(u, v, w; c_1, \dots, c_r).$ 
    
```

For $r = 3$, the F^r compression function design is depicted in Fig. 2. This generic design is a generalization of the permutation based hash function construction described by Rogaway and Steinberger [24]. In fact, it is straightforward to generalize the main findings of [24] to our F^r design and we state them as preliminary results. If the collision- and preimage-degeneracies are sufficiently small (these values intuitively capture the degree of non-randomness of the design with respect to the occurrence of collisions and preimages), one can expect collisions after approximately $2^{n(2-2/r)}$ queries and preimages after approximately $2^{n(2-1/r)}$ queries. We refer to [24] for the details. First of all, these findings confirm that at least two cipher calls are required to get 2^n collision resistance. More importantly, from these results we can conclude that F^r can impossibly achieve optimal 2^{2n} preimage resistance. Yet, it may still be possible to construct a function that achieves optimal collision resistance and almost-optimal preimage resistance.

Throughout, we consider security in the ideal cipher model: we consider an adversary \mathcal{A} that is a probabilistic algorithm with oracle access to a block cipher $E \stackrel{\$}{\leftarrow} \text{Bloc}(n)$ randomly sampled from $\text{Bloc}(n)$. \mathcal{A} is information-theoretic: it has unbounded computational power, and its complexity is measured by the number of queries made to its oracles. The adversary can make forward queries and inverse queries to E , and these are stored in a query history Q as indexed tuples of the form (k_i, m_i, c_i) , where k_i denotes the key input, and (m_i, c_i) the plaintext/ciphertext pair. For $q \geq 0$, by Q_q we define the query history after q queries. We assume that the adversary never makes queries to which it knows the answer in advance.




Fig. 2. $F^3 : \{0, 1\}^{3n} \rightarrow \{0, 1\}^{2n}$ making three block cipher evaluations

A collision-finding adversary \mathcal{A} for F^r aims at finding two distinct inputs to F^r that compress to the same range value. In more detail, we say that \mathcal{A} succeeds if it finds two distinct tuples $(u, v, w), (u', v', w')$ such that $F^r(u, v, w) = F^r(u', v', w')$ and Q contains all queries required for these evaluations of F^r . We define by

$$\text{adv}_{F^r}^{\text{coll}}(\mathcal{A}) = \Pr \left(\begin{array}{l} E \xleftarrow{\$} \text{Bloc}(n), (u, v, w), (u', v', w') \leftarrow \mathcal{A}^{E, E^{-1}} : \\ (u, v, w) \neq (u', v', w') \wedge F^r(u, v, w) = F^r(u', v', w') \end{array} \right)$$

the probability that \mathcal{A} succeeds in this. By $\text{adv}_{F^r}^{\text{coll}}(q)$ we define the maximum collision advantage taken over all adversaries making q queries.

For preimage resistance, we focus on everywhere preimage resistance [23], which captures preimage security for every point of $\{0, 1\}^{2n}$. Before making any queries to its oracle, a preimage-finding adversary \mathcal{A} first decides on a range point $(y, z) \in \{0, 1\}^{2n}$. Then, we say that \mathcal{A} succeeds in finding a preimage if it obtains a tuple (u, v, w) such that $F^r(u, v, w) = (y, z)$ and Q contains all queries required for this evaluation of F^r . We define by

$$\text{adv}_{F^r}^{\text{epre}}(\mathcal{A}) = \max_{(y, z) \in \{0, 1\}^{2n}} \Pr \left(\begin{array}{l} E \xleftarrow{\$} \text{Bloc}(n), (u, v, w) \leftarrow \mathcal{A}^{E, E^{-1}}(y, z) : \\ F^r(u, v, w) = (y, z) \end{array} \right)$$

the probability that \mathcal{A} succeeds, maximized over all possible choices for (y, z) . By $\text{adv}_{F^r}^{\text{epre}}(q)$ we define the maximum (everywhere) preimage advantage taken over all adversaries making q queries.

3 Impossibility Result for 2-Call Double Length Hashing

We present an attack on a wide class of double block length compression functions with two calls to their underlying block cipher $E : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$. Let F^2 be a compression function of this form. We pose a condition on the finalization function f_3 , such that if this condition is satisfied, collisions for F^2 can be found in about $2^{n/2}$ queries. Although we are not considering all possible compression functions, we cover the most interesting and intuitive ones, such as compression functions with linear finalization function f_3 . Compression functions with non-linear f_3 are covered up to some degree (but we note that the attack does not apply to the compression function of [7], for which collision security up to $2^{2n/3}$ queries is proven).

We first state the attack. Then, by ways of examples, we illustrate its generality. For the purpose of the attack, we introduce the function left_n which on input of a bit string of length $2n$ bits outputs the leftmost n bits.

Proposition 1. *Let $F^2 : \{0, 1\}^{3n} \rightarrow \{0, 1\}^{2n}$ be a compression function as described in Sect. 2. Suppose there exists a bijective function L such that for any $u, v, w, c_1, c_2 \in \{0, 1\}^n$ we have*

$$\text{left}_n \circ L \circ f_3(u, v, w; c_1, c_2) = \text{left}_n \circ L \circ f_3(u, v, w; c_1, 0). \tag{1}$$

Then, one can expect collisions for F^2 after $2^{n/2}$ queries.

Proof. Let F^2 be a compression function and let L be a bijection such that (1) holds. First, we consider the case of L being the identity function, and next we show how this attack extends to the case L is an arbitrary bijection.

Suppose (1) holds with L the identity function. This means that the first n bits of $f_3(u, v, w; c_1, c_2)$ do not depend on c_2 and we can write f_3 as a concatenation of two functions $g_1 : \{0, 1\}^{4n} \rightarrow \{0, 1\}^n$ and $g_2 : \{0, 1\}^{5n} \rightarrow \{0, 1\}^n$ as $f_3(u, v, w; c_1, c_2) = g_1(u, v, w; c_1) \| g_2(u, v, w; c_1, c_2)$. Let $\alpha \in \mathbb{N}$. We present an adversary \mathcal{A} for F^2 . The first part of the attack is derived from [24].

- Make α queries $(k_1, m_1) \rightarrow c_1$ that maximize the number of tuples (u, v, w) with $f_1(u, v, w)$ hitting any of these values (k_1, m_1) . By the balls-and-bins principle², the adversary obtains at least $\alpha \cdot 2^{3n} / 2^{2n} = \alpha 2^n$ tuples $(u, v, w; c_1)$ for which it knows the first block cipher evaluation;
- Again by the balls-and-bins principle, there exists a value y such that at least α tuples satisfy $g_1(u, v, w; c_1) = y$;
- Varying over these α tuples, compute $(k_2, m_2) = f_2(u, v, w; c_1)$ and query (k_2, m_2) to the cipher to obtain a c_2 . \mathcal{A} finds a collision for F^2 if it obtains two tuples $(u, v, w; c_1, c_2)$, $(u', v', w'; c'_1, c'_2)$ that satisfy $g_2(u, v, w; c_1, c_2) = g_2(u', v', w'; c'_1, c'_2)$.

In the last round one expects to find a collision if $\alpha^2 / 2^n = 1$, or equivalently if $\alpha = 2^{n/2}$. In total, the attack is done in approximately $2 \cdot 2^{n/2}$ queries.

It remains to consider the case of L being an arbitrary bijection. Define \overline{F}^2 as F^2 with f_3 replaced by $\overline{f}_3 = L \circ f_3$. Using the idea of equivalence classes on compression functions [18] we prove that F^2 and \overline{F}^2 are equally secure with respect to collisions. Let $\overline{\mathcal{A}}$ be a collision finding adversary for \overline{F}^2 . We construct a collision finding adversary \mathcal{A} for F^2 , with oracle access to E , that uses $\overline{\mathcal{A}}$ to output a collision for F^2 . Adversary \mathcal{A} proceeds as follows. It forwards all queries made by $\overline{\mathcal{A}}$ to its own oracle. Eventually, $\overline{\mathcal{A}}$ outputs two tuples (u, v, w) , (u', v', w') such that $\overline{F}^2(u, v, w) = \overline{F}^2(u', v', w')$. Denote by c_1 the block cipher outcome on input of $f_1(u, v, w)$ and by c_2 the outcome on input of $f_2(u, v, w; c_1)$. Define c'_1 and c'_2 similarly. By construction, as (u, v, w) and (u', v', w') form a collision for \overline{F}^2 , we have $L \circ f_3(u, v, w; c_1, c_2) = L \circ f_3(u', v', w'; c'_1, c'_2)$. Now, bijectivity of L implies that $f_3(u, v, w; c_1, c_2) = f_3(u', v', w'; c'_1, c'_2)$, and hence (u, v, w) and (u', v', w') form a collision for F^2 . (Recall that F^2 and \overline{F}^2 only differ in the finalization function f_3 , the functions f_1 and f_2 are the same.) We thus obtain $\mathbf{adv}_{\overline{F}^2}^{\text{coll}}(q) \leq \mathbf{adv}_{F^2}^{\text{coll}}(q)$. The derivation in reverse order is the same by symmetry. But \overline{F}^2 satisfies (1) for L the identity function. Therefore, the attack described in the first part of the proof applies to \overline{F}^2 , and thus to F^2 . □

We demonstrate the impact of the attack by giving several example functions that fall in the categorization. We stress that the requirement of Prop. 1 is in fact solely a requirement on f_3 ; f_1 and f_2 can be any function.

² If k balls are thrown in l bins, the α fullest bins in total contain at least $\alpha k / l$ balls.

Suppose F^2 uses a linear finalization function f_3 . Say, f_3 is defined as follows:

$$\begin{pmatrix} \mathbf{a}_{11} & \mathbf{a}_{12} & \mathbf{a}_{13} & \mathbf{a}_{14} & \mathbf{a}_{15} \\ \mathbf{a}_{21} & \mathbf{a}_{22} & \mathbf{a}_{23} & \mathbf{a}_{24} & \mathbf{a}_{25} \end{pmatrix} (u, v, w, c_1, c_2)^\top = (y, z)^\top,$$


where addition and multiplication is done over the field $GF(2^n)$. Now, if $\mathbf{a}_{25} = 0$ we set $L = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ which corresponds to swapping y and z . If $\mathbf{a}_{25} \neq 0$, we set $L = \begin{pmatrix} 1 & -\mathbf{a}_{15}\mathbf{a}_{25}^{-1} \\ 0 & 1 \end{pmatrix}$, which corresponds to subtracting the second equation $\mathbf{a}_{15}\mathbf{a}_{25}^{-1}$ times from the first one. The attack also covers designs whose finalization function f_3 rotates or shuffles its inputs, such as MDC-2, where one defines L so that the rotation gets undone. We elaborate on this in the full version [17]. In general, if f_3 is a sufficiently simple add-rotate-xor function, it is possible to derive a bijective L that makes (1) satisfied. Up to a degree, the attack also covers general non-linear finalization functions. However, it clearly does not cover all functions and it remains an open problem to either close this gap or to come with a (possibly impractical) F^2 compression function that provably achieves optimal collision resistance. One direction may be to start from the compression function with non-linear finalization f_3 by Jetchev et al. [7], for which collision resistance up to $2^{2n/3}$ queries is proven.

4 Double Length Hashing with 3 E-calls

Motivated by the negative result of Sect. 3, we target the existence of double length hashing with three block cipher calls. We introduce a family of double length compression functions making three cipher calls that achieve asymptotically optimal 2^n collision resistance and preimage resistance significantly beyond the birthday bound (up to $2^{3n/2}$ queries). We note that, although the preimage bound is non-optimal, it closely approaches the generic bound dictated by the pigeonhole-birthday attack (Sect. 2).

Let $GF(2^n)$ be the field of order 2^n . We identify bit strings from $\{0, 1\}^n$ and finite field elements in $GF(2^n)$ to define addition and scalar multiplication over $\{0, 1\}^n$. In the family of double block length functions we propose in this section, the functions f_1, f_2, f_3, f_4 of Fig. 2 will be linear functions over $GF(2^n)$. For two tuples $x = (x_1, \dots, x_l)$ and $y = (y_1, \dots, y_l)$ of elements from $\{0, 1\}^n$, we define by $x \cdot y$ their inner product $\sum_{i=1}^l x_i y_i \in \{0, 1\}^n$.

Before introducing the design, we first explain the fundamental consideration upon which the family is based. The security proofs of all DBL^{2n} functions known in the literature (cf. Table 1) crucially rely on the property that one block cipher evaluation defines the input to the other one. For DBL^{2n} functions this can easily be achieved: any block cipher evaluation can take as input the full $3n$ -bit input state (u, v, w) . Considering the class of functions DBL^n , and F^r of Fig. 2 in particular, this can impossibly be achieved: one block cipher “processes” at most $2n$ out of $3n$ input bits. In our design, we slightly relax this requirement, by requiring that any *two* block cipher evaluations define the input to the third one. Although from a technical point of view one may expect that



$$F_A^3(u, v, w) = (y, z), \text{ where:}$$

$$c_1 \leftarrow E(u, v),$$

$$k_2 \leftarrow a_1 \cdot (u, v, c_1),$$

$$m_2 \leftarrow a_2 \cdot (u, v, c_1, w),$$

$$y \leftarrow E(k_2, m_2) + m_2,$$

$$k_3 \leftarrow a_3 \cdot (u, v, c_1),$$

$$m_3 \leftarrow a_4 \cdot (u, v, c_1, w),$$

$$z \leftarrow E(k_3, m_3) + m_3.$$

Fig. 3. The family of compression functions F_A^3 where A is a 4×4 matrix as specified in the text. Arithmetics is done over $GF(2^n)$.

this change causes optimal collision resistance to be harder or even impossible to be achieved, we will demonstrate that this is not the case due to new proof techniques employed to analyze the collision resistance.

Based on this key observation we propose the compression function design F_A^3 of Fig. 3. Here,

$$A = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} & a_{13} & 0 \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & 0 \\ a_{41} & a_{42} & a_{43} & a_{44} \end{pmatrix} \quad (2)$$

is a 4×4 matrix over $GF(2^n)$. Note that, provided A is invertible and $a_{24}, a_{44} \neq 0$, any two block cipher evaluations of F_A^3 define (the inputs of) the third one. For instance, evaluations of the second and third block cipher fix the vector $A(u, v, c_1, w)^T$, which by invertibility of A fixes (u, v, c_1, w) and thus the first block cipher evaluation. Evaluations of the first and second block cipher fix the inputs of the third block cipher as $a_{24} \neq 0$. For the proofs of collision and preimage resistance, however, we will need to posit additional requirements on A . As we will explain, these requirements are easily satisfied.

In the remainder of this section, we state our results on the collision resistance of F_A^3 in Sect. 4.1 and on the preimage resistance in Sect. 4.2.

4.1 Collision Resistance of F_A^3

We prove that, provided its underlying matrix A satisfies some simple conditions, F_A^3 satisfies optimal collision resistance. In more detail, we pose the following requirements on A :

- A is invertible;
- $a_{12}, a_{13}, a_{24}, a_{32}, a_{33}, a_{44} \neq 0$;
- $a_{12} \neq a_{32}$ and $a_{13} \neq a_{33}$.

We refer to the logical AND of these requirements as `colreq`.

Theorem 1. *Let $n \in \{0, 1\}^n$. Suppose A satisfies `colreq`. Then, for any positive integral values t_1, t_2 ,*

$$\text{adv}_{F_A^3}^{\text{coll}}(q) \leq \frac{2t_2^2q + 3t_2q + 11q + 3t_1t_2^2 + 7t_1t_2}{2^n - q} + \frac{q^2}{t_1(2^n - q)} + 3 \cdot 2^n \left(\frac{eq}{t_2(2^n - q)} \right)^{t_2}. \tag{3}$$

The proof is given in Sect. 5. The basic proof idea is similar to existing proofs in the literature (e.g. [16,27]) and is based on the usage of thresholds t_1, t_2 . For increasing values of t_1, t_2 the first term of the bound increases, while the second two terms decrease. Although the proof derives basic proof principles from literature, for the technical part we deviate from existing proof techniques in order to get a bound that is “as tight as possible”. In particular, we introduce the usage of wish lists in the context of collisions, an approach that allows for significantly better bounds. Wish lists have been introduced by Armknecht et al. [2] and Lee et al. [11,13] for the preimage resistance analysis of DBL^{2n} functions, but they have never been used for collision resistance as there never was a need to do so. Our analysis relies on this proof methodology, but as for collisions more block cipher evaluations are involved (one collision needs six block cipher calls while a preimage requires three) this makes the analysis more technical and delicate.

The goal now is to find a good threshold between the first term and the latter two terms of (3). To this end, let $\varepsilon > 0$ be any parameter. We put $t_1 = q$ and $t_2 = 2^{n\varepsilon}$ (we can assume t_2 to be integral). Then, the bound simplifies to

$$\text{adv}_{F_A^3}^{\text{coll}}(q) \leq \frac{5 \cdot 2^{2n\varepsilon}q + 10 \cdot 2^{n\varepsilon}q + 11q}{2^n - q} + \frac{q}{2^n - q} + 3 \cdot 2^n \left(\frac{eq}{2^{n\varepsilon}(2^n - q)} \right)^{2^{n\varepsilon}}.$$

From this, we find that for any $\varepsilon > 0$ we have $\text{adv}_{F_A^3}^{\text{coll}}(2^n/2^{3n\varepsilon}) \rightarrow 0$ for $n \rightarrow \infty$. Hence, the F_A^3 compression function achieves close to optimal 2^n collision security for $n \rightarrow \infty$. For $n = 128$, we evaluate the bound in more detail in [17]. The advantage hits 1/2 for $\log_2 q \approx 118.3$, relatively close to the threshold 127.5 for $q(q + 1)/2^{2n}$. For larger values of n this gap approaches 0.

4.2 Preimage Resistance of F_A^3

In this section we consider the preimage resistance of F_A^3 . Though we do not obtain optimal preimage resistance—which is impossible to achieve after all, due to the generic bounds of the pigeonhole-birthday attack (Sect. 2)—we achieve preimage resistance up to $2^{3n/2}$ queries, much better than the preimage bounds on MDC-2 and MDC-4 [16], relatively close to the generic bound. Yet, for the proof to hold we need to put slightly stronger requirements on A.

- $A - \begin{pmatrix} & 0 & 0 \\ B_1 & 0 & 0 \\ & 0 & 0 \\ B_2 & 0 & 0 \end{pmatrix}$ is invertible for any $B_1, B_2 \in \left\{ \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \right\}$. In the remainder, we write $[B_1/B_2]$ to denote the subtracted matrix;
- $a_{12}, a_{13}, a_{24}, a_{32}, a_{33}, a_{44} \neq 0$;
- $a_{12} \neq a_{32}, a_{13} \neq a_{33},$ and $a_{24} \neq a_{44}$.

We refer to the logical AND of these requirements as **prereq**. We remark that **prereq** \Rightarrow **colreq**, and that matrices satisfying **prereq** are easily found. Simple matrices complying with these conditions over the field $GF(2^{128})$ are

$$\begin{pmatrix} 0 & 1 & 2 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 2 & 1 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix}, \quad \begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 1 & 0 & 1 \\ 0 & 2 & 3 & 0 \\ 1 & 0 & 2 & 2 \end{pmatrix}. \tag{4}$$

These are the matrices corresponding to the compression functions of Fig. 1. Here, we use $x^{128} + x^{127} + x^{126} + x^{121} + 1$ as our irreducible polynomial and we represent bit strings as polynomials in the obvious way ($1 = 1, 2 = x, 3 = 1 + x$). Note that the choice of matrix A influences the efficiency of the construction. The first matrix of (4) has as minimal zeroes as possible, which reduces the amount of computation.

Theorem 2. *Let $n \in \{0, 1\}^n$. Suppose A satisfies **prereq**. Then, for any positive integral value t , provided $t \leq q$,*

$$\mathbf{adv}_{F_A^3}^{\text{epre}}(q) \leq \frac{6t^2 + 18t + 26}{2^n - 2} + 4 \cdot 2^n \left(\frac{4eq}{t2^n} \right)^{t/2} + 8q \left(\frac{8eq}{t2^n} \right)^{\frac{t2^n}{4q}}. \tag{5}$$

The proof is given in the full version of this paper [17]. As for the bound on the collision resistance (Thm. 1), the idea is to make a smart choice of t to minimize this bound. Let $\varepsilon > 0$ be any parameter. Then, for $t = q^{1/3}$, the bound simplifies to

$$\mathbf{adv}_{F_A^3}^{\text{epre}}(q) \leq \frac{6q^{2/3} + 18q^{1/3} + 26}{2^n - 2} + 4 \cdot 2^n \left(\frac{4eq^{2/3}}{2^n} \right)^{q^{1/3}/2} + 8q \left(\frac{8eq^{2/3}}{2^n} \right)^{\frac{2^n}{4q^{2/3}}}.$$

From this, we find that for any $\varepsilon > 0$ we have $\mathbf{adv}_{F_A^3}^{\text{epre}}(2^{3n/2}/2^{n\varepsilon}) \rightarrow 0$ for $n \rightarrow \infty$. Hence, the F_A^3 compression function achieves close to $2^{3n/2}$ preimage security for $n \rightarrow \infty$. For $n = 128$, we evaluate the bound in more detail in [17]. The advantage hits 1/2 for $\log_2 q \approx 180.3$, relatively close to the threshold 191.5 for $q^2/2^{3n}$. For larger values of n this gap approaches 0.

The result shows that F_A^3 with A compliant to **prereq** satisfies preimage resistance up to about $2^{3n/2}$ queries. We note that our proof is the best possible for this design, by demonstrating a preimage-finding adversary that with high probability succeeds in at most $O(2^{3n/2})$ queries. Let $\alpha \in \mathbb{N}$. The adversary proceeds as follows.

- Make $\alpha 2^n$ queries to the block cipher corresponding to the bottom-left position of Fig. 3. One expects to find α tuples (k_2, m_2, c_2) that satisfy $m_2 + c_2 = y$;
- Repeat the first step for the bottom-right position. One expects to find α tuples (k_3, m_3, c_3) satisfying $m_3 + c_3 = z$;
- By invertibility of A , any choice of (k_2, m_2, c_2) and (k_3, m_3, c_3) uniquely defines a tuple (u, v, c_1, w) for the F_A^3 evaluation. Likely, the emerged tuples (u, v, c_1) are all different, and we find about α^2 such tuples;
- Varying over all α^2 tuples (u, v, c_1) , query (u, v) to the block cipher. If it responds c_1 , we have obtained a preimage for F_A^3 .

In the last round one expects to find a preimage if $\alpha^2/2^n = 1$, or equivalently if $\alpha = 2^{n/2}$. The first and second round both require approximately $2^{3n/2}$ queries, and the fourth round takes 2^n queries. In total, the attack is done in approximately $2 \cdot 2^{3n/2} + 2^n$ queries.

5 Proof of Thm. 1

The proof of collision resistance of F_A^3 follows the basic spirit of [16], but crucially differs in the way the probability bounds are computed. A new approach here is the usage of wish lists. While the idea of wish lists is not new—it has been introduced by Armknecht et al. [2] and Lee et al. [11,13] for double block length compression functions, and used by Mennink [16] for the analysis of MDC-4—in these works wish lists are solely used for the analysis of preimage resistance rather than collision resistance. Given that in a collision more block cipher evaluations are involved, the analysis becomes more complex. At a high level, wish lists rely on the idea that in order to find a collision, the adversary must at some point make a query that “completes this collision” together with some other queries already in the query history. Wish lists keep track of such query tuples, and the adversary’s goal is to ever obtain a query tuple that is in such wish list. A more technical treatment can be found in the proof of Lem. 1.

We consider any adversary that has query access to its oracle E and makes q queries stored in a query history Q_q . Its goal is to find a collision for F_A^3 , in which it by definition only succeeds if it obtains a query history Q_q that satisfies configuration $\text{coll}(Q_q)$ of Fig. 4. This means,

$$\mathbf{adv}_{F_A^3}^{\text{coll}}(q) = \mathbf{Pr}(\text{coll}(Q_q)). \tag{6}$$

For the sake of readability of the proof, we label the block cipher positions in Fig. 4 as follows. In the left F_A^3 evaluation (on input (u, v, w)), the block ciphers are labeled $1L$ (the one on input (u, v)), $2L$ (the bottom left one), and $3L$ (the bottom right one). The block ciphers for the right F_A^3 evaluation are labeled $1R, 2R, 3R$ in a similar way. When we say “a query $1L$ ”, we refer to a query that in a collision occurs at position $1L$.

For the analysis of $\mathbf{Pr}(\text{coll}(Q_q))$ we introduce an auxiliary event $\text{aux}(Q_q)$. Let $t_1, t_2 > 0$ be any integral values. We define $\text{aux}(Q_q) = \text{aux}_1(Q_q) \vee \dots \vee \text{aux}_4(Q_q)$,




Fig. 4. Configuration $\text{coll}(Q)$. The configuration is satisfied if Q contains six (possibly the same) queries that satisfy this setting. We require $(u, v, w) \neq (u', v', w')$.

where

$$\begin{aligned} \text{aux}_1(Q_q) &: \left| \left\{ (k_i, m_i, c_i), (k_j, m_j, c_j) \in Q_q : i \neq j \wedge m_i + c_i = m_j + c_j \right\} \right| > t_1; \\ \text{aux}_2(Q_q) &: \max_{z \in \{0,1\}^n} \left| \left\{ (k_i, m_i, c_i) \in Q_q : \mathbf{a}_1 \cdot (k_i, m_i, c_i) = z \right\} \right| > t_2; \\ \text{aux}_3(Q_q) &: \max_{z \in \{0,1\}^n} \left| \left\{ (k_i, m_i, c_i) \in Q_q : \mathbf{a}_3 \cdot (k_i, m_i, c_i) = z \right\} \right| > t_2; \\ \text{aux}_4(Q_q) &: \max_{z \in \{0,1\}^n} \left| \left\{ (k_i, m_i, c_i) \in Q_q : m_i + c_i = z \right\} \right| > t_2. \end{aligned}$$

By basic probability theory, we obtain for (6):

$$\Pr(\text{coll}(Q_q)) \leq \Pr(\text{coll}(Q_q) \wedge \neg \text{aux}(Q_q)) + \Pr(\text{aux}(Q_q)). \tag{7}$$

We start with the analysis of $\Pr(\text{coll}(Q_q) \wedge \neg \text{aux}(Q_q))$. For obtaining a query history that fulfills configuration $\text{coll}(Q_q)$, it may be the case that a query appears at multiple positions. For instance, the queries at positions $1L$ and $2R$ are the same. We split the analysis of $\text{coll}(Q_q)$ into essentially all different possible cases, but we do this in two steps. In the first step, we distinct among the cases a query occurs in both words at the same position. We define for binary $\alpha_1, \alpha_2, \alpha_3$ by $\text{coll}_{\alpha_1 \alpha_2 \alpha_3}(Q)$ the configuration $\text{coll}(Q)$ of Fig. 4 restricted to

$$1L = 1R \iff \alpha_1 = 1, \quad 2L = 2R \iff \alpha_2 = 1, \quad 3L = 3R \iff \alpha_3 = 1.$$

By construction, $\text{coll}(Q_q) \Rightarrow \bigvee_{\alpha_1, \alpha_2, \alpha_3 \in \{0,1\}} \text{coll}_{\alpha_1 \alpha_2 \alpha_3}(Q_q)$, and from (6-7) we obtain the following bound on $\text{adv}_{F_A^3}^{\text{coll}}(q)$:

$$\text{adv}_{F_A^3}^{\text{coll}}(q) \leq \sum_{\substack{\alpha_1, \alpha_2, \alpha_3 \\ \in \{0,1\}}} \Pr(\text{coll}_{\alpha_1 \alpha_2 \alpha_3}(Q_q) \wedge \neg \text{aux}(Q_q)) + \Pr(\text{aux}(Q_q)). \tag{8}$$

Note that we did not make a distinction yet whether or not a query occurs at two “different” positions (e.g. at positions $1L$ and $2R$). These cases are analyzed for each of the sub-configurations separately, as becomes clear later. Probabilities $\Pr(\text{coll}_{\alpha_1\alpha_2\alpha_3}(Q_q) \wedge \neg\text{aux}(Q_q))$ for the different choices of $\alpha_1, \alpha_2, \alpha_3$ are bounded in Lems. 1-4. The proofs are rather similar, and we only bound the probability on $\text{coll}_{000}(Q_q)$ in full detail (Lem. 1). A bound on $\Pr(\text{aux}(Q_q))$ is given in Lem. 5. A part of the proof of Lem. 1, and the proofs of Lems. 2-5 are given in [17].

Lemma 1. $\Pr(\text{coll}_{000}(Q_q) \wedge \neg\text{aux}(Q_q)) \leq \frac{t_2q+7q+3t_1t_2^2+3t_1t_2}{2^n-q}$.

Proof. Sub-configuration $\text{coll}_{000}(Q_q)$ is given in Fig. 5. The block cipher queries at positions a and $!a$ are required to be different, and so are the ones at positions $b, !b$ and $c, !c$.




Fig. 5. Configuration $\text{coll}_{000}(Q)$. We require $(u, v, w) \neq (u', v', w')$.

We consider the probability of the adversary finding a solution to configuration $\text{coll}_{000}(Q_q)$ such that Q_q satisfies $\neg\text{aux}(Q_q)$. Consider the i th query, for $i \in \{1, \dots, q\}$. We say this query is a winning query if it makes $\text{coll}_{000}(Q_i) \wedge \neg\text{aux}(Q_i)$ satisfied for any set of other queries in the query history Q_{i-1} . We can assume the i th query does not make $\text{aux}(Q_i)$ satisfied: if it would, by definition it cannot be a winning query.

Recall that, although we narrowed down the number of possible positions for a winning query to occur (in $\text{coll}_{000}(Q_q)$ it cannot occur at both $1L$ and $1R$, at both $2L$ and $2R$, or at both $3L$ and $3R$), it may still be the case that such a query contributes to multiple “different” positions, e.g. $1L$ and $2R$. Note that by construction, a winning query can contribute to at most three block cipher positions of Fig. 5. In total, there are 26 sets of positions at which the winning

query can contribute at the same time. Discarding symmetric cases caused by swapping (u, v, w) and (u', v', w') , one identifies the following 13 sets of positions:

$$\begin{aligned} \mathcal{S}_1 &= \{1L\}, & \mathcal{S}_4 &= \{1L, 2L\}, & \mathcal{S}_7 &= \{1L, 2R\}, & \mathcal{S}_{10} &= \{1L, 2L, 3L\}, \\ \mathcal{S}_2 &= \{2L\}, & \mathcal{S}_5 &= \{1L, 3L\}, & \mathcal{S}_8 &= \{1L, 3R\}, & \mathcal{S}_{11} &= \{1L, 2L, 3R\}, \\ \mathcal{S}_3 &= \{3L\}, & \mathcal{S}_6 &= \{2L, 3L\}, & \mathcal{S}_9 &= \{2L, 3R\}, & \mathcal{S}_{12} &= \{1L, 2R, 3L\}, \\ & & & & & & \mathcal{S}_{13} &= \{1L, 2R, 3R\}. \end{aligned}$$

Note that there are many more symmetric cases among these, but we are not allowed to discard those as these may result in effectively different collisions. For $j = 1, \dots, 13$ we denote by $\text{coll}_{000:S_j}(Q)$ configuration $\text{coll}_{000}(Q)$ with the restriction that the winning query *must* appear at the positions in \mathcal{S}_j . By basic probability theory,

$$\Pr(\text{coll}_{000}(Q_q) \wedge \neg \text{aux}(Q_q)) \leq \sum_{j=1}^{13} \Pr(\text{coll}_{000:S_j}(Q_q) \wedge \neg \text{aux}(Q_q)). \quad (9)$$

$\text{coll}_{000:S_1}(Q_q)$. Rather than considering the success probability of the i th query, and then sum over $i = 1, \dots, q$ (as is done in the analysis of [4,5,6,7,9,12,16,21,26], hence all collision security proofs of Table 1), the approach in this proof is to focus on “wish lists”. Intuitively, a wish list is a continuously updated sequence of query tuples that would make configuration $\text{coll}_{000:S_j}(Q_q)$ satisfied. During the attack of the adversary, we maintain an initially empty wish list \mathcal{W}_{S_1} . Consider configuration $\text{coll}_{000}(Q)$ with the query at position $\mathcal{S}_1 = \{1L\}$ left out (see [17] for a graphical intuition). If a new query is made, suppose it fits this configuration for some other queries in the query history (the new query appearing at least once), jointly representing queries at positions $\{2L, 3L, 1R, 2R, 3R\}$. Then the corresponding tuple (u, v, c_1) is added to \mathcal{W}_{S_1} . Note that this tuple is uniquely determined by the queries at $2L$ and $3L$ by invertibility of A , but different combinations of queries may define the same wish. The latter does, however, not invalidate the analysis: this is covered by the upper bound on \mathcal{W}_{S_1} that will be computed later in the proof, and will simply render a slightly worse bound.

As we have restricted to the case the winning query only occurring at the position of \mathcal{S}_1 , we can assume a query never adds itself to a wish list³. Clearly, in order to find a collision for F_A^3 in this sub-configuration, the adversary needs to wish for a query at least once. Suppose the adversary makes a query $E(k, m)$ where $(k, m, c) \in \mathcal{W}_{S_1}$ for some c . We say that (k, m, c) is wished for, and the wish is granted if the query response equals c . As the adversary makes at most q queries, such wish is granted with probability at most $1/(2^n - q)$, and the same for inverse queries. By construction, each element from \mathcal{W}_{S_1} can be wished for only once, and we find that the adversary finds a collision with probability at most $\frac{|\mathcal{W}_{S_1}|}{2^n - q}$.

³ A winning query that would appear at multiple positions is counted in $\text{coll}_{000:S_j}(Q_q)$ for some other set \mathcal{S}_j .

Now, it suffices to upper bound the size of the wish list \mathcal{W}_{S_1} after q queries, and to this end we bound the number of solutions to configuration $\text{coll}_{000:S_j}(Q_q)$. By $\neg\text{aux}_1(Q_q)$, the configuration has at most t_1 choices for $2L, 2R$. For any such choice, by $\neg\text{aux}_2(Q_q)$ we have at most t_2 choices for $1R$. Any such choice fixes w' (as $a_{24} \neq 0$), and thus the query at position $3R$, and consequently z . By $\neg\text{aux}_4(Q_q)$, we have at most t_2 choices for $3L$. The queries at positions $2L$ and $3L$ uniquely fix (u, v, c_1) by invertibility of A . We find $|\mathcal{W}_{S_1}| \leq t_1 t_2^2$, and hence in this setting a collision is found with probability at most $t_1 t_2^2 / (2^n - q)$.

coll_{000:S_j}(Q_q) for $j = 2, \dots, 13$. In [17], $\Pr(\text{coll}_{000:S_j}(Q_q) \wedge \neg\text{aux}(Q_q))$ is bounded by $t_1 t_2^2 / (2^n - q)$ for $j = 2, 3$, $q / (2^n - q)$ for $j = 4, 5, 6, 10, 11, 12, 13$, $t_1 t_2 / (2^n - q)$ for $j = 7, 8$, and $(t_1 t_2 + t_2 q) / (2^n - q)$ for $j = 9$.

The proof is now completed by adding all bounds in accordance with (9). □

Lemma 2. $\Pr(\text{coll}_{100}(Q_q) \wedge \neg\text{aux}(Q_q)) \leq \frac{2q+2t_1 t_2}{2^n - q}$.

Lemma 3. $\Pr(\text{coll}_{\alpha_1 \alpha_2 \alpha_3}(Q_q) \wedge \neg\text{aux}(Q_q)) \leq \frac{t_2^2 q + t_2 q + q + t_1 t_2}{2^n - q}$ for $\alpha_1 \alpha_2 \alpha_3 \in \{010, 001\}$.

Lemma 4. $\Pr(\text{coll}_{\alpha_1 \alpha_2 \alpha_3}(Q_q) \wedge \neg\text{aux}(Q_q)) = 0$ when $\alpha_1 + \alpha_2 + \alpha_3 \geq 2$.

Lemma 5. $\Pr(\text{aux}(Q_q)) \leq \frac{q^2}{t_1(2^n - q)} + 3 \cdot 2^n \left(\frac{eq}{t_2(2^n - q)} \right)^{t_2}$.

From (8) and the results of Lems. 1-5 we conclude the bound of (3). This completes the proof of Thm. 1.

6 Conclusions

In the area of double block length hashing, where a $3n$ -to- $2n$ -bit compression function is constructed from n -bit block ciphers, all optimally secure constructions known in the literature employ a block cipher with $2n$ -bit key space. We have reconsidered the principle of double length hashing, focusing on double length hashing from a block cipher with n -bit message and key space. Unlike in the DBL^{2n} class, we demonstrate that there does not exist any optimally secure design with reasonably simple finalization function that makes two cipher calls. By allowing one extra call, optimal collision resistance can nevertheless be achieved, as we have proven by introducing our family of designs F_A^3 .

In our quest for optimal collision secure compression function designs, we had to resort to designs with three block cipher calls rather than two, which moreover are not parallelizable. This entails an efficiency loss compared to MDC-2, MJH, and Jetchev et al.'s construction. On the other hand, our family of functions is based on simple arithmetic in the finite field: unlike constructions by Stam [25,26], Lee and Steinberger [14], and Jetchev et al. [7], our design does not make use of full field multiplications. The example matrices A given in (4) are designed

to use a minimal amount of non-zero elements. We note that specific choices of A may be more suited for this construction to be used in an iterated design.

This work provides new insights in double length hashing, but also results in interesting research questions. Most importantly, is it possible to construct other collision secure F^3 constructions (beyond our family of functions F_A^3), that achieve optimal $2^{5n/3}$ preimage resistance? Given the negative collision resistance result for a wide class of compression functions F^2 , is it possible to achieve optimal collision security *in the iteration* anyhow? This question is beyond the scope of this work. On the other hand, in line with ideas of [18], is it possible to achieve an impossibility result for F^3 restricted to the xor-only design (where f_1, \dots, f_4 only xor their parameters)?

Acknowledgments. This work has been funded in part by the IAP Program P6/26 BCRYPT of the Belgian State (Belgian Science Policy), in part by the European Commission through the ICT program under contract ICT-2007-216676 ECRYPT II, and in part by the Research Council K.U.Leuven: GOA TENSE. The author is supported by a Ph.D. Fellowship from the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen). The author would like to thank Elena Andreeva and the anonymous ASIACRYPT 2012 reviewers for their valuable help and feedback.

References

1. Andreeva, E., Neven, G., Preneel, B., Shrimpton, T.: Seven-Property-Preserving Iterated Hashing: ROX. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 130–146. Springer, Heidelberg (2007)
2. Armknecht, F., Fleischmann, E., Krause, M., Lee, J., Stam, M., Steinberger, J.: The Preimage Security of Double-Block-Length Compression Functions. In: Lee, D.H. (ed.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 233–251. Springer, Heidelberg (2011)
3. Bos, J.W., Özen, O., Stam, M.: Efficient Hashing Using the AES Instruction Set. In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 507–522. Springer, Heidelberg (2011)
4. Fleischmann, E., Gorski, M., Lucks, S.: Security of Cyclic Double Block Length Hash Functions. In: Parker, M.G. (ed.) Cryptography and Coding 2009. LNCS, vol. 5921, pp. 153–175. Springer, Heidelberg (2009)
5. Hirose, S.: Provably Secure Double-Block-Length Hash Functions in a Black-Box Model. In: Park, C., Chee, S. (eds.) ICISC 2004. LNCS, vol. 3506, pp. 330–342. Springer, Heidelberg (2005)
6. Hirose, S.: Some Plausible Constructions of Double-Block-Length Hash Functions. In: Robshaw, M. (ed.) FSE 2006. LNCS, vol. 4047, pp. 210–225. Springer, Heidelberg (2006)
7. Jetchev, D., Özen, O., Stam, M.: Collisions Are Not Incidental: A Compression Function Exploiting Discrete Geometry. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 303–320. Springer, Heidelberg (2012)
8. Lai, X., Massey, J.L.: Hash Functions Based on Block Ciphers. In: Rueppel, R.A. (ed.) EUROCRYPT 1992. LNCS, vol. 658, pp. 55–70. Springer, Heidelberg (1993)

9. Lee, J., Kwon, D.: The security of Abreast-DM in the ideal cipher model. Cryptology ePrint Archive, Report 2009/225 (2009)
10. Lee, J., Stam, M.: MJH: A Faster Alternative to MDC-2. In: Kiayias, A. (ed.) CT-RSA 2011. LNCS, vol. 6558, pp. 213–236. Springer, Heidelberg (2011)
11. Lee, J., Stam, M., Steinberger, J.: The collision security of Tandem-DM in the ideal cipher model. Cryptology ePrint Archive, Report 2010/409 (2010); full version of [12]
12. Lee, J., Stam, M., Steinberger, J.: The Collision Security of Tandem-DM in the Ideal Cipher Model. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 561–577. Springer, Heidelberg (2011)
13. Lee, J., Stam, M., Steinberger, J.: The preimage security of double-block-length compression functions. Cryptology ePrint Archive, Report 2011/210 (2011)
14. Lee, J., Steinberger, J.: Multi-property-preserving Domain Extension Using Polynomial-Based Modes of Operation. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 573–596. Springer, Heidelberg (2010)
15. Lucks, S.: A collision-resistant rate-1 double-block-length hash function. In: Symmetric Cryptography. Dagstuhl Seminar Proceedings, vol. (07021) (2007)
16. Mennink, B.: On the collision and preimage security of MDC-4 in the ideal cipher model. Cryptology ePrint Archive, Report 2012/113 (2012)
17. Mennink, B.: Optimal collision security in double block length hashing with single length key (2012); full version of this paper
18. Mennink, B., Preneel, B.: Hash Functions Based on Three Permutations: A Generic Security Analysis. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 330–347. Springer, Heidelberg (2012)
19. Meyer, C., Schilling, M.: Secure program load with manipulation detection code. In: Proc. Securicom, pp. 111–130 (1988)
20. Nandi, M., Lee, W., Sakurai, K., Lee, S.: Security Analysis of a 2/3-Rate Double Length Compression Function in the Black-Box Model. In: Gilbert, H., Handschuh, H. (eds.) FSE 2005. LNCS, vol. 3557, pp. 243–254. Springer, Heidelberg (2005)
21. 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)
22. Preneel, B., Govaerts, R., Vandewalle, J.: Hash Functions Based on Block Ciphers: A Synthetic Approach. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 368–378. Springer, Heidelberg (1994)
23. Rogaway, P., Shrimpton, T.: Cryptographic Hash-Function Basics: Definitions, Implications, and Separations for Preimage Resistance, Second-Preimage Resistance, and Collision Resistance. In: Roy, B., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 371–388. Springer, Heidelberg (2004)
24. Rogaway, P., Steinberger, J.: Security/Efficiency Tradeoffs for Permutation-Based Hashing. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 220–236. Springer, Heidelberg (2008)
25. Stam, M.: Beyond Uniformity: Better Security/Efficiency Tradeoffs for Compression Functions. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 397–412. Springer, Heidelberg (2008)
26. Stam, M.: Blockcipher-Based Hashing Revisited. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 67–83. Springer, Heidelberg (2009)
27. Steinberger, J.P.: The Collision Intractability of MDC-2 in the Ideal-Cipher Model. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 34–51. Springer, Heidelberg (2007)