The collision security of Tandem-DM in the ideal cipher model

Jooyoung Lee[1], Martijn Stam[2*], and John Steinberger[3†]

[1] Faculty of Mathematics and Statistics, Sejong University, Seoul, Korea, jlee05@sejong.ac.kr
[2] Royal Holloway, University of London, Egham, United Kingdom, m.stam@alumnus.tue.nl
[3] Institute of Theoretical Computer Science, Tsinghua University, Beijing, China, jpsteinb@gmail.com

Abstract. We prove that Tandem-DM, one of the two "classical" schemes for turning a blockcipher of $2n$-bit key into a double block length hash function, has birthday-type collision resistance in the ideal cipher model. A collision resistance analysis for Tandem-DM achieving a similar birthday-type bound was already proposed by Fleischmann, Gorski and Lucks at FSE 2009 [3]. As we detail, however, the latter analysis is wrong, thus leaving the collision resistance of Tandem-DM as an open problem until now.

1 Introduction

The Tandem-DM compression function is a $3n$-bit to $2n$-bit compression function based on two applications of a blockcipher of $2n$-bit key and n-bit word length (Fig. 1). While Tandem-DM was proposed by Lai and Massey in 1992 [7] the first proof of collision security for Tandem-DM (in the ideal cipher model, as is usual for all such proofs) was only proposed in 2009 by Fleischmann, Gorski and Lucks [3]. Unfortunately, as we detail in Section 3, the "FGL proof" (as we shall refer to it) has a number of serious flaws which make it false and nonobvious to repair. The purpose of this paper is to offer a correct collision resistance analysis of Tandem-DM. We show that, as claimed in [3], Tandem-DM does indeed have birthday-type collision security (necessitating at least $2^{120.8}$ queries to break when the output length is $2n = 256$ bits). A nice feature of our work is that the analysis is relatively simple compared to typical results in this area. In Section 5 we also give a preimage resistance analysis for Tandem-DM, as the preimage analysis of [3] suffers from similar flaws as the collision analysis.

RELATED WORK ON 2-CALL CONSTRUCTIONS. Another classical scheme for turning a $2n$-bit key blockcipher into a $3n$-bit to $2n$-bit compression function is Abreast-DM, pictured in Fig. 2, which was proposed by Lai and Massey in the same paper as Tandem-DM [7]. The collision resistance of Abreast-DM was independently resolved by Fleischmann, Gorski and Lucks [4] and Lee and Kwon [8], who both showed birthday-type collision resistance for Abreast-DM. Before that, Hirose [5] had given a collision resistance analysis for a general class of compression functions that included Abreast-DM as a special case, but under the assumption that the top and bottom blockciphers of the diagram be distinct (this considerably simplifies the analysis). The work by Hirose was further generalized by Özen and Stam [11], who additionally discuss schemes that are only secure in the iteration.

Another $3n$-bit to $2n$-bit compression function making two calls to a blockcipher of $2n$-bit key was proposed by Hirose [6], who proved birthday-type collision resistance for his construction in the ideal cipher model. Hirose's construction (Fig. 3) is simpler than either Abreast-DM or Tandem-DM and in particular uses a single keying schedule for the top and bottom blockciphers. It is noteworthy that while Hirose introduced his construction over 10 years after Abreast-DM and Tandem-DM his collision resistance analysis pre-dates similar collision resistance analyses for Abreast-DM and Tandem-DM.

RELATED WORK ON 1-CALL CONSTRUCTIONS. Stam [16] proposed a class of "polynomial-based" $3n$-bit to $2n$-bit compression functions making a single call to a $2n$-bit key blockcipher, and subsequently proved [17] birthday-type collision resistance for this construction. Lee and Steinberger [9] proved collision resistance for the same compression

* Part of the work performed while at LACAL, École Polytechnique Fédérale de Lausanne, Switzerland.
† Supported by the National Natural Science Foundation of China Grant 60553001 and by the National Basic Research Program of China Grant 2007CB807900, 2007CB807901.

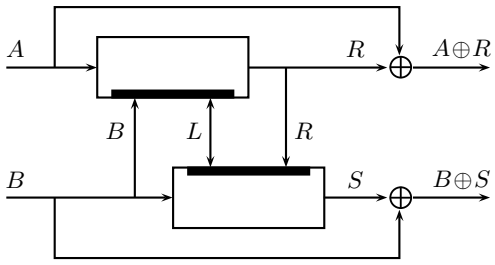

Fig. 1: The Tandem-DM compression function. All wires carry n-bit values. The top and bottom blockciphers are the same. Each has a $2n$-bit key and n-bit input/output. The wire marked L is an input to the compression function (along with A and B).

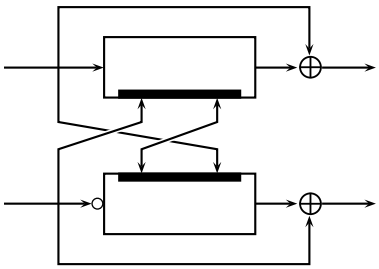

Fig. 2: The Abreast-DM compression function. The empty circle at bottom left denotes bit complementation.

function in the weaker "unpredictable cipher" model. Lucks [10] proposed a double length hash function using a $3n$-bit to $2n$-bit compression function making a single call to a blockcipher of $2n$-bit key, and proved this hash function collision resistant in the ideal cipher model (see [11] for a generalization). However, Lucks' construction is only secure in the iteration, as the compression function itself is collision insecure.

Earlier, Yi and Lam [19] had proposed a $3n$-bit to $2n$-bit compression function making a single call to a $2n$-bit key blockcipher whose design was somewhat similar to Stam's polynomial-based construction but which used a single integer addition operation instead of several field multiplication operations. However, this construction was broken by Wagner [18].

COMPARISON. Of the three well-known $3n$-bit to $2n$-bit compression functions making two calls to a $2n$-bit key blockcipher—those being Tandem-DM, Abreast-DM and Hirose's construction—the two constructions whose collision resistance has been successfully resolved (Hirose and Abreast-DM) share the feature that the inputs to the top and bottom blockcipher are bijectively related. For example, for Abreast-DM, if the top blockcipher call is $E_{B\|L}(A)$ then the bottom blockcipher call (for the same input $A\|B$) is $E_{L\|A}(\overline{B})$, where $\overline{B}$ denotes bit complementation of B; thus the inputs to the top and bottom blockciphers are related by the permutation $\pi : \{0,1\}^{3n} \to \{0,1\}^{3n}$, $\pi(X\|Y\|Z) = \overline{Y}\|Z\|X$. (Here the last $2n$ bits are the key.) In Hirose's construction, the inputs to the top and bottom blockciphers are related by the permutation $\pi' : \{0,1\}^{3n} \to \{0,1\}^{3n}$, $\pi'(X\|Y\|Z) = X \oplus c\|Y\|Z$.

By contrast, Tandem-DM exhibits a more subtle relationship between the inputs of the top and bottom blockciphers, as an *output* of the top blockcipher is used to key the bottom blockcipher. It is the presence of this "feedback" within the construction, it seems, that has complicated efforts to prove a collision resistance bound. On the other hand, Tandem-DM still has the agreeable feature that the top and bottom blockcipher calls uniquely determine each other in the following sense: given the key $B\|L$ and output R of the top cipher one can determine the key $L\|R$ and the input B of the bottom cipher, and vice-versa. This contrasts with constructions such as MDC-2 which use two calls to a blockcipher of n-bit key, and in which the top and bottom blockcipher calls do not uniquely determine each other. Typically, collision resistance analyses are much harder for the latter kind of compression functions. (MDC-2 can

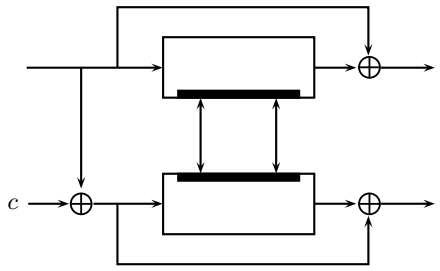

Fig. 3: Hirose's compression function. The bottom left-hand wire is not an input; it carries an arbitrary nonzero constant c.

only be proved nontrivially collision resistant in the iteration, and the current best bound of $O(2^{\frac{3}{5}n})$ queries due to Steinberger [15] is undoubtedly suboptimal.)

We note that the permutations π and π' discussed above share the common feature of having *small cycle lengths*—all cycles of π have length 6 and all cycles of π' have length 2—which constitutes another strong similarity between Abreast-DM and Hirose's scheme. In fact, due to this reason, Hirose's collision resistance proof and the Abreast-DM collision resistance proof can be seen as special cases of the same framework, as noted in [4, 8]. Building on this observation, Fleischmann et al. [4] defined a general class of compression functions called 'Cyclic-DM' that are amenable to collision resistance analyses and that include Hirose's scheme and Abreast-DM as special cases. Similarly, one can define collision-resistant generalizations of Tandem-DM by isolating those properties of Tandem-DM that are used in our proof. While defining the most all-encompassing possible collision resistant generalization of Tandem-DM is not the goal of this paper we do briefly discuss these key properties and the corresponding collision-resistant generalizations of Tandem-DM in Section 6, without proof of security.

A SECOND PROOF. The proof of collision resistance that we provide in this paper is very slick, but slightly mysterious in its efficacy because it relies on a subtle trick that cuts out a large portion of the case analysis that "would have been there" in a more standard proof. As a pedagogical bonus, and to provide some perspective on our proof, we also show how to prove the collision security of Tandem-DM without this trick in Appendix A. We note this second, "brute force" proof yields a slightly weaker bound.

FURTHER POSSIBLE IMPROVEMENTS. We note that our collision resistance has the form $\tilde{O}(q/(2^n - q))$ rather than $\tilde{O}(q^2/(2^n - q)^2)$. Both bounds reach constant values when $q = \Omega(2^n)$, however $q^2/(2^n - q)^2$ grows slower than $q/(2^n - q)$ since our bound is (only) "linear birthday" rather than true "quadratic birthday". We leave it as an open problem to prove "quadratic birthday"-type collision resistance for Tandem-DM (as exists for Abreast-DM and Hirose's scheme). Moreover, it is an open problem to prove preimage resistance for values of q higher than 2^n for either Abreast-DM, Tandem-DM or Hirose.

VERSION HISTORY. After the initial posting of this (our) work, we became aware of another paper by Fleischmann et al. [2], providing a comprehensive generalization of their earlier works [3,4]. (In particular, a new, tighter collision resistance claim for Tandem-DM is made.) Unfortunately, the problems in the (FSE'09) FGL proof are not addressed and actually carry over to this new generalization (in particular, the crucial "Argument B" of [4] is incorrect), rendering the resulting bounds meaningless. We will discuss it in detail at Appendix B.

2 Definitions

A blockcipher is a function $E : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}^n$ such that $E(K, \cdot)$ is a permutation of $\{0,1\}^n$ for each $K \in \{0,1\}^m$. We call m the *key size* and n the *word size* of the blockcipher. It is customary to write $E_K(X)$ instead of $E(K, X)$ for $K \in \{0,1\}^m$, $X \in \{0,1\}^n$. The function $E_K^{-1}(\cdot)$ denotes the inverse of $E_K(\cdot)$ (as $E_K(\cdot)$ is a permutation).

Given a blockcipher $E : \{0,1\}^{2n} \times \{0,1\}^n \to \{0,1\}^n$ we define the Tandem-DM compression function $TDM^E : \{0,1\}^{3n} \to \{0,1\}^{2n}$ by

$$TDM^E(A\|B\|L) = (A \oplus R)\|(B \oplus S)$$

where

$$R = E_{B\|L}(A),$$
$$S = E_{L\|R}(B).$$

In the collision resistance experiment, a computationally unbounded adversary[4] A is given oracle access to a blockcipher E uniformly sampled among all blockciphers of key length $2n$ and word length n. We allow A to query both E and E^{-1}. After q queries to E, the *query history* of A is the set of triples $\mathcal{Q} = \{(X_i, K_i, Y_i)\}_{i=1}^q$ such that $E_{K_i}(X_i) = Y_i$ and A's i-th query is either $E_{K_i}(X_i)$ or $E_{K_i}^{-1}(Y_i)$ for $1 \le i \le q$. We let $\mathcal{Q}_i = \{(X_j, K_j, Y_j)\}_{j=1}^i$ be the first i elements of the query history; thus $\mathcal{Q} = \mathcal{Q}_q$. We say A *succeeds* or *finds a collision* after its first i queries if there exist distinct $3n$-bit values, $A\|B\|L$, $A'\|B'\|L'$ such that $TDM^E(A\|B\|L) = TDM^E(A'\|B'\|L')$ and such that $\mathcal{Q}_i$ contains both the queries necessary to compute $TDM^E(A\|B\|L)$ and $TDM^E(A'\|B'\|L')$. More formally—and see Fig. 4—we define this event by a predicate $\mathsf{Coll}(\mathcal{Q}_i)$, which is true if and only if there exist n-bit values A, B, L, R, S, A', B', L', R', S' such that

$$A\|B\|L \neq A'\|B'\|L' \tag{1}$$
$$A \oplus R = A' \oplus R' \tag{2}$$
$$B \oplus S = B' \oplus S' \tag{3}$$

and such that

$$(A, B\|L, R), (B, L\|R, S), (A', B'\|L', R'), (B', L'\|R', S') \in \mathcal{Q}_i. \tag{4}$$

We denote by

$$\mathbf{Adv}^{\mathrm{coll}}_{TDM}(q)$$

the maximum chance of an adversary making q queries causing $\mathsf{Coll}(\mathcal{Q})$ to become true. The probability occurs over the uniform choice of E and over A's coin tosses, if any. Also note that n is a hidden parameter.

The "XOR-output" of a query (X_i, K_i, Y_i) is the quantity $X_i \oplus Y_i$. Another predicate which plays an important part in both our proof and the FGL proof is the "many queries with the same XOR-output" predicate $\mathsf{Xor}(\mathcal{Q})$, defined on a query history $\mathcal{Q} = \{(X_i, K_i, Y_i)\}_{i=1}^q$ by

$$\mathsf{Xor}(\mathcal{Q}) \iff \max_{Z \in \{0,1\}^n} |\{i : X_i \oplus Y_i = Z\}| > \alpha.$$

Here α is a free parameter of the analysis which appears in the final collision resistance bound. (In [3] this predicate is named $\textsc{Lucky}(\mathcal{Q})$; in [15] a similar predicate is named $\mathsf{Win0}(\mathcal{Q})$.) Without going into details at this point, we mention that the FGL collision resistance proof—and ours, essentially, as well—upper bounds $\Pr[\mathsf{Coll}(\mathcal{Q})]$ by $\Pr[\mathsf{Xor}(\mathcal{Q})] + \Pr[\mathsf{Coll}(\mathcal{Q}) \wedge \neg\mathsf{Xor}(\mathcal{Q})]$. A larger α implies a lower value for $\Pr[\mathsf{Xor}(\mathcal{Q})]$ and a higher value for $\Pr[\mathsf{Coll}(\mathcal{Q}) \wedge \neg\mathsf{Xor}(\mathcal{Q})]$. The best value of α can be found numerically for a given value of n and q. Generally, readers may think of α as some small constant value (e.g. for $n = 128$ and $q = 2^{120.87}$, $\alpha = 16$).

So far, we have described "infrastructure" that is common to both proofs. We shall now introduce some material proper to our proof. Note a query history $\mathcal{Q} = \{(X_i, K_i, Y_i)\}_{i=1}^q$ does not record whether each triple (X_i, K_i, Y_i) was obtained by the adversary through a forward query $E_{K_i}(X_i)$ or a backward query $E_{K_i}^{-1}(Y_i)$. For this, we maintain two arrays $\mathsf{Fwd}[\cdot]$ and $\mathsf{Bwd}[\cdot]$ where $\mathsf{Fwd}[i] = 1$ if and only if the adversary's i-th query is a forward query and $\mathsf{Bwd}[i] = 1$ if and only if the adversary's i-th query is a backward query. We then define an additional predicate

$$\mathsf{FB}(\mathcal{Q}) \iff \max_{Z \in \{0,1\}^n} |\{i : (Y_i = Z \wedge \mathsf{Fwd}[i] = 1) \vee (X_i = Z \wedge \mathsf{Bwd}[i] = 1)\}| > \alpha. \tag{5}$$

[4] Our notation for the adversary and one of the Tandem-DM inputs collide, but without too much danger of confusion.

('FB' stands for "Forward Backward".) Here α is the same free parameter as above. Note that $\neg\mathsf{FB}(\mathcal{Q})$ implies that

$$\max_{Z\in\{0,1\}^n} |\{i : Y_i = Z \wedge \mathsf{Fwd}[i] = 1\}| \leq \alpha, \tag{6}$$
$$\max_{Z\in\{0,1\}^n} |\{i : X_i = Z \wedge \mathsf{Bwd}[i] = 1\}| \leq \alpha. \tag{7}$$

It is really consequences (6) and (7) of $\neg\mathsf{FB}(\mathcal{Q})$ that interest us, though we define $\mathsf{FB}(\mathcal{Q})$ via (5) because this makes it slightly easier to bound $\Pr[\mathsf{FB}(\mathcal{Q})]$. We will use the bound

$$\begin{aligned}\Pr[\mathsf{Coll}(\mathcal{Q})] &\leq \Pr[\mathsf{Xor}(\mathcal{Q})] + \Pr[\mathsf{Coll}(\mathcal{Q}) \wedge \neg\mathsf{Xor}(\mathcal{Q})] \\ &\leq \Pr[\mathsf{Xor}(\mathcal{Q})] + \Pr[\mathsf{FB}(\mathcal{Q})] + \Pr[\mathsf{Coll}(\mathcal{Q}) \wedge \neg\mathsf{Xor}(\mathcal{Q}) \wedge \neg\mathsf{FB}(\mathcal{Q})].\end{aligned} \tag{8}$$

One should thus think of $\mathsf{FB}(\mathcal{Q})$ and $\mathsf{Xor}(\mathcal{Q})$ as bad events whose nonoccurrence helps bound the probability of $\mathsf{Coll}(\mathcal{Q})$ occurring. We warn that (8) constitutes a slightly oversimplified encapsulation of our proof's high-level structure. We refer to Section 4 for more details.

3 The FGL collision resistance proof

Since the interest of our paper would be substantially diminished (though not nullified, since our proof is much shorter) if the FGL collision resistance proof were correct, we detail here some of our objections to [3]. This material also serves as a good introduction to our own proof, and will give the reader more intuition about Tandem-DM.

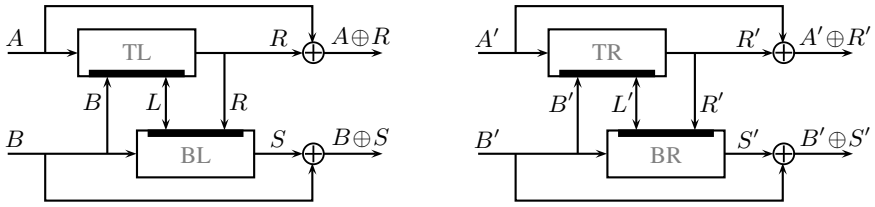

Fig. 4: The collision diagram for Tandem-DM. The adversary must find blockcipher queries to fit both sides of the diagram such that $A \oplus R = A' \oplus R'$, $B \oplus S = B' \oplus S'$ and $A\|B\|L \neq A'\|B'\|L'$. More precisely, the adversary must find four queries of the form $E_{B\|L}(A) = R$, $E_{L\|R}(B) = S$, $E_{B'\|L'}(A') = R'$, $E_{L'\|R'}(B') = S'$ such that the above conditions hold. Each query could either be learned through a forward query (to E) or through a backward query (to E^{-1}). The four queries in the diagram are labeled 'TL', 'BL', 'TR', 'BR' for 'Top Left', 'Bottom Left', etc.

Starting with a q-query collision-finding adversary A, FGL first make the standard assumption that A never makes a query to which it already knows the answer (this could occur two ways: A could make the exact same query twice, or A could query (say) $E_K^{-1}(Y)$ after having received Y as an answer beforehand to a query $E_K(X)$). This ensures each answer A receives comes uniformly at random from a set of size at least $2^n - q$ (since $E_K(\cdot)$ is a random permutation for each K). Moreover, after A makes i queries its query history will contain exactly i distinct elements.

Say A *succeeds* at the i-th query if $\mathsf{Coll}(\mathcal{Q}_i)$ holds but $\mathsf{Coll}(\mathcal{Q}_{i-1})$ and $\mathsf{Xor}(\mathcal{Q}_{i-1})$ do not hold. By upper bounding the probability that A ever succeeds we upper bound $\Pr[\mathsf{Coll}(\mathcal{Q}) \wedge \neg\mathsf{Xor}(\mathcal{Q})]$. (Upper bounding $\Pr[\mathsf{Xor}(\mathcal{Q})]$ is an easy probability exercise that we overlook for the purposes of this proof sketch.) A good analogy is to view A as trying to complete a puzzle where each element of its query history is a puzzle piece it can use to complete the collision diagram of Fig. 4. We use the expressions "A succeeds", "A finds a [puzzle] solution" or "A completes a collision" interchangeably (and we will rarely remind that the condition $\neg\mathsf{Xor}(\mathcal{Q}_{i-1})$ must hold for A to succeed). We refer to the four queries (in any hypothetical puzzle solution (a.k.a. collision)) as 'TL', 'BL', 'TR' and 'BR'; see Fig. 4.

Note the constraint $A\|B\|L \neq A'\|B'\|L'$ does not imply that the queries TL, BL, TR, BR are all distinct. For example, one could have TL = BR (in which case $(A, B\|L, R) = (B', L'\|R', S')$, so $A = B'$, $B = L'$, $L = R'$ and

$R = S'$) or TL = BL (in which case we have the dramatic consequence that $A = B = L = R = S$, as is easy to check). This gives rise to several combinatorially distinct cases to consider; A's chance of obtaining a solution of each kind is upper bounded separately, and these probabilities are added together to form a final upper bound on A's chance of success. (Oddly, FGL include the cases TL = TR and BL = BR in their analysis, while these are impossible since they imply $A\|B\|L = A'\|B'\|L'$. This oversight, however, does not imply an incorrect proof in itself.)

We shall restrict our critique to FGL's analysis of the "generic" case when the queries TL, BL, TR, BR are all distinct. We note that, in these types of analyses, the generic case is usually the hardest to handle as A's job typically grows harder when additional constraints are placed on its solution. (The possibility of reusing the same query in two different positions of the collision diagram does however occasionally prove useful to A, depending on the construction, so all cases must always be considered.) We call a puzzle solution in which TL, BL, TR, BR are distinct a "generic solution."

If A succeeds in finding a generic solution there is a smallest i such that a generic solution can be assembled from the queries in $\mathcal{Q}_i$. The i-th query is then called the "last query" of A's solution. To upper bound A's chance of obtaining a generic solution, FGL consider two cases. The first case is the event that A's last query can be used in position TL of the puzzle solution and the second case is the event that A's last query can be used in position BL (one of these two cases must occur). We shall focus on the first of these two cases, which is also the first analyzed in the order of the FGL proof. We call it the "TL generic" case.

One would typically consider two subcases for the TL generic case (or any other) depending on whether A's last query is a forward query to E or an inverse query to E^{-1}, but FGL lump their analysis into a single argument claiming that the two types of queries can be handled the same (in fact, they make this claim for every case in their proof, and never distinguish between forward and backward queries to E). For clarity, however, we shall restrict ourselves to considering the case of a forward query to E, and discuss how their argument specializes to that case. We also choose to specifically consider the forward query case because this is where FGL's analysis seems to be the most problematic.

The task at hand is thus to upper bound A's chance of completing a generic solution by making a forward query to E that can be used as query TL of such a solution. The usual approach for this, and the one used by FGL, is to consider any given forward query $E_{K_i}(X_i)$ made by A and to upper bound the probability that the answer Y_i to this query is such that the query history element (X_i, K_i, Y_i) can be used in the desired manner; one then multiplies this probability by q since A can make q queries total. With foresight on how we wish to use the query $E_{K_i}(X_i)$ it is convenient to rename K_i as $B\|L$ and X_i as A; thus the query is $E_{B\|L}(A)$. To proceed, one would typically upper bound the number of values $R \in \{0,1\}^n$ such that, if we had $E_{B\|L}(A) = R$, the query $(A, B\|L, R)$ could be used in position TL of a generic solution together with previous elements of the query history, and divide this number by $2^n - q$, since the answer to the query $E_{B\|L}(A)$ will come uniformly at random from a set of size at least $2^n - q$. In turn, the standard, formal way of bounding the number of such R's would be to upper bound the possible number of query triples (BL, BR, TR) in the query history that could potentially be used with the query $E_{B\|L}(A)$ to form a generic solution, as the number of such triples is an upper bound for the number of R's. Note such a triple must have the form $\mathrm{BL} = (B, L\|R, S)$, $\mathrm{BR} = (B', L'\|R', S')$, $\mathrm{TR} = (A', B'\|L', R')$ where $B \oplus S = B' \oplus S'$ (and note that A, B and L are fixed here by the last query).

FGL do not adopt[5] this approach for bounding the number of good R's. Rather, they make the following argument: take the value of R, whatever it is, that is returned by the query $E_{B\|L}(A)$; because $\neg\mathsf{Xor}(\mathcal{Q}_{i-1})$ there will be at most α queries $\mathrm{TR} = (A', B'\|L', R')$ in the query history such that $A \oplus R = A' \oplus R'$; as the TR query uniquely determines the BR query, there are at most α possibilities for the BR query; now "give the query $\mathrm{BL} = (B, L\|R, S)$ for free to the adversary"; then since there are at most α possibilities for the query $\mathrm{BR} = (B', L'\|R', S')$ there is chance at most $\alpha/(2^n - q)$ that $B \oplus S = B' \oplus S'$ for one of the queries BR, so total chance at most $q\alpha/(2^n - q)$ that the adversary ever obtains a TL-generic solution with a forward query, there being at most q queries total.

The fallacy in the above argument can be succinctly summarized by pointing out that *the query* $\mathrm{BL} = (B, L\|R, S)$ *may already be in the query history, in which case there is no randomness left in the value* $B \oplus S$. However, let us review in detail the argument in two different cases: when the query $\mathrm{BL} = (B, L\|R, S)$ is already in the query history prior to the last query, and when it isn't. (Note that query BL only depends on R (besides B and L which are fixed by

[5] Neither do we, in fact. Using a careful trick, we manage to upper bound the number of good R's by only considering the possibilities for the query BL rather than by considering the possible triples (BL, TR, BR). In Appendix A, however, we give for comparison the "brute force" proof which uses the method of upper bounding the number of triples (BL, TR, BR).

the last query), and not on which queries are "chosen" for TR and BR.) In the latter case, when BL = $(B, L\|R, S)$ is not yet in the query history at the i-th query, then A's last query can in any case not succeed in completing a generic TL collision since the query BL is missing; thus there is no need to bound anything (and no need even to "give the query BL for free"). In the case when query BL is already in the query history, on the other hand, all randomness is lost once R is revealed. FGL successfully argue that, for a given value of R, there will be at most α possibilities for the pair (TR , BR), but this does not in any way imply the *non-existence* of such queries TR, BR.

Other issues are raised by FGL's casual comment that the query BL $= (B, L\|R, S)$ is simply "given for free" to the adversary. Indeed, if this query is not yet present, is it added to the query history before or after the i-th query itself? Is this query only made after the value of R is revealed, or is it somehow inserted into the query history before the value of R is revealed? The former might be all right; the latter not, since it would (drastically) alter R's distribution conditioned on the query history, i.e. R would no longer come uniformly at random from a set of size $\geq 2^n - q$. Most importantly, since this free query becomes part of the query history, one should account for the possibility that *this query* (not the i-th query) causes the adversary to succeed (and not necessarily by being used in position BL of a generic solution). Indeed, we are forced to give such credit to the adversary, since we have required the adversary never to make a query to which it already knows the answer, and since the adversary may have wished to subsequently make this query itself; this means the case analysis should be applied recursively to the free query, but if the case analysis requires other queries to be "given for free", then we bite our tail and end up giving an astronomical number of free queries to the adversary (e.g., nearly all possible queries).

Note also that nothing in the FGL argument precludes the possibility that, when the adversary makes its i-th query $E_{B\|L}(A)$, there is not some very large number of distinct values of R—say $2^{0.5n}$—for which there exists a triplet of queries (BL, TR, BR) of the form BL $= (B, L\|R, S)$, BR $= (B', L'\|R', S')$, TR $= (A', B'\|L', R')$ where $B \oplus S = B' \oplus S'$, and such that R does not yet appear as the third coordinate of any query in the query history with key $B\|L$. Certainly, there being such a large number of values of R does not contradict $\neg\mathsf{Xor}(\mathcal{Q}_{i-1})$. Also certainly, the i-th query would have chance $2^{0.5n}/(2^n - q)$ of making the adversary succeed if such a large number of values of R existed, and not chance $\alpha/(2^n - q)$. In other words, one can infer something is wrong with the FGL argument because it does not address the main difficulty of the case at hand.

While we singled out the TL generic case for examination, the same kinds of problems recur throughout the FGL case analysis, essentially invalidating the entire proof. Moreover, since the FGL proof sidesteps the most crucial challenges posed by an analysis of Tandem-DM (see the previous paragraph), it leaves little for any subsequent analysis to build on. We note that the FGL preimage resistance proof suffers from very similar flaws as the collision resistance proof, as briefly discussed in Section 5.

4 Main result: collision resistance of Tandem-DM

It will be easier to explain the form of the probability bound in our main theorem if we explain a few high-level ideas from the proof beforehand. The proof starts by considering an arbitrary q-query collision-finding adversary A for Tandem-DM. We then construct an adversary A' as follows: A' simulates A, but after each forward query $E_{V\|W}(U)$ made by A, A' makes the backward query $E^{-1}_{U\|V}(W)$ *if* it does not already know[6] the answer to this query, and after each backward query $E^{-1}_{U\|V}(W)$ made by A, A' makes the forward query $E_{V\|W}(U)$ *if* it does not already know[7] the answer to this query. (To better understand the relation of these instructions to Tandem-DM, view U, V, W as B, L, R.) Moreover if A ever makes a query to which A' already knows the answer from its query history, A' ignores this query. Thus A' never makes a query to which it knows the answer.

Let $\mathcal{Q}'$ be the query history of A' and $\mathcal{Q}$ be the query history of A. Then $\mathcal{Q} \subseteq \mathcal{Q}'$ and $|\mathcal{Q}'| \leq 2q$. Since $\mathcal{Q} \subseteq \mathcal{Q}'$ we have

$$\Pr[\mathsf{Coll}(\mathcal{Q})] \leq \Pr[\mathsf{Coll}(\mathcal{Q}')] \leq \Pr[\mathsf{Xor}(\mathcal{Q}')] + \Pr[\mathsf{FB}(\mathcal{Q}')] + \Pr[\mathsf{Coll}(\mathcal{Q}') \wedge \neg\mathsf{Xor}(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')]. \quad (9)$$

Our proof uses the inequality above to bound $\Pr[\mathsf{Coll}(\mathcal{Q})]$. Incidentally, we point out that if we construct an adversary A'' from A' the same way A' is constructed from A, then A'' and A' will have the same query history, as is not difficult

[6] More formally, if its query history does not contain any triple of the form $(\cdot, U\|V, W)$.
[7] More formally, if its query history does not contain any triple of the form $(U, V\|W, \cdot)$.

to see. In other words, *every* forward query $E_{V\|W}(U)$ made by A' (including its "own" queries) is followed by the query $E^{-1}_{U\|V}(W)$ unless A' already knows this query, and likewise *every* backward query $E^{-1}_{U\|V}(W)$ made by A' is followed by the forward query $E_{V\|W}(U)$ unless A' already knows the answer to this query. The use of the augmented adversary A' may seem superficially similar to Fleischmann et al.'s idea of "giving away a query for free". However, it will become clear from our case analysis that we exploit the added structure of $\mathcal{Q}'$ entirely differently from the way Fleischmann et al. exploit their free queries. We also point out that the added structure of $\mathcal{Q}'$ enables the main interesting trick of our analysis, found in case 'TL Forward' of Proposition 3 below.

We can now more easily discuss our main result:

Theorem 1. *Let $N = 2^n$, $q < N/2$, $N' = N - 2q$ and let α be an integer, $1 \le \alpha \le 2q$. Then*

$$\mathbf{Adv}^{\mathrm{coll}}_{TDM}(q) \le 2N\left(\frac{2eq}{\alpha N'}\right)^{\alpha} + \frac{4q\alpha}{N'} + \frac{4q}{N'}.$$

The term $2N\left(\frac{2eq}{\alpha N'}\right)^{\alpha}$ in Theorem 1 is an upper bound for $\Pr[\mathsf{Xor}(\mathcal{Q}')] + \Pr[\mathsf{FB}(\mathcal{Q}')]$. In fact $\Pr[\mathsf{Xor}(\mathcal{Q}')] \le N\left(\frac{2eq}{\alpha N'}\right)^{\alpha}$ and $\Pr[\mathsf{FB}(\mathcal{Q}')] \le N\left(\frac{2eq}{\alpha N'}\right)^{\alpha}$. The two remaining terms $4q\alpha/N' + 4q/N'$ are an upper bound for $\Pr[\mathsf{Coll}(\mathcal{Q}') \wedge \neg\mathsf{Xor}(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')]$. To bound $\mathbf{Adv}^{\mathrm{coll}}_{TDM}(q)$ for a given value of n and q one should optimize α numerically. For example, for $n = 128$, Theorem 1 yields that $\mathbf{Adv}^{\mathrm{coll}}_{TDM}(2^{120.87}) < \frac{1}{2}$ using $\alpha = 16$. Asymptotically, Theorem 1 yields the following result:

Corollary 1. $\lim_{n\to\infty} \mathbf{Adv}^{\mathrm{coll}}_{TDM}(N/n) = 0$.

Proof. Let $q = N/n$ and $\alpha = n/\log n$, where the logarithm takes base 2. Since $N' > N/2$ for $n > 4$, we have

$$\begin{aligned}\mathbf{Adv}^{\mathrm{coll}}_{TDM}(q) &\le 2N\left(\frac{2eq}{\alpha N'}\right)^{\alpha} + \frac{4q\alpha}{N'} + \frac{4q}{N'} \le 2N\left(\frac{4eq}{\alpha N}\right)^{\alpha} + \frac{8q\alpha}{N} + \frac{8q}{N} \\ &\le 2N\left(\frac{4e\log n}{n^2}\right)^{\frac{n}{\log n}} + \frac{8}{\log n} + \frac{8}{n} = 2\left(\frac{4e\log n}{n}\right)^{\frac{n}{\log n}} + \frac{8}{\log n} + \frac{8}{n}.\end{aligned}$$

The last expression obviously goes to zero as $n \to \infty$. □

In particular, $\lim_{n\to\infty} \mathbf{Adv}^{\mathrm{coll}}_{TDM}\left(2^{(1-\varepsilon)n}\right) = 0$ for any fixed $\varepsilon > 0$.

The proof of Theorem 1 uses refinements $\mathsf{Coll}_1(\mathcal{Q})$, $\mathsf{Coll}_2(\mathcal{Q})$, $\mathsf{Coll}_3(\mathcal{Q})$ of the collision predicate $\mathsf{Coll}(\mathcal{Q})$, defined as follows:

$\mathsf{Coll}_1(\mathcal{Q})$ occurs if $\mathcal{Q}$ contains a collision with TL, BL, TR, BR distinct.
$\mathsf{Coll}_2(\mathcal{Q})$ occurs if $\mathcal{Q}$ contains a collision with either TL = BL or TR = BR.
$\mathsf{Coll}_3(\mathcal{Q})$ occurs if $\mathcal{Q}$ contains a collision with either TL = BR or BL = TR.

For example, $\mathsf{Coll}_2(\mathcal{Q})$ occurs if there exist values $A, B, L, R, S, A', B', L', R', S'$ such that (1)–(4) hold and such that $(A, B\|L, R) = (B, L\|R, S)$. Since BL $\neq$ BR and TL $\neq$ TR in any collision, we have the following proposition.

Proposition 1. $\mathsf{Coll}(\mathcal{Q}) \implies \mathsf{Coll}_1(\mathcal{Q}) \vee \mathsf{Coll}_2(\mathcal{Q}) \vee \mathsf{Coll}_3(\mathcal{Q})$ *for any query history $\mathcal{Q}$.*

In view of proving Theorem 1, let A be an arbitrary q-query adversary for Tandem-DM, and let A' be obtained from A as outlined above; let $\mathcal{Q}$ be the query history of A and $\mathcal{Q}'$ be the query history of A'. Then by (9) it suffices to show that

$$\begin{aligned}\Pr[\mathsf{Xor}(\mathcal{Q}')] &\le N\left(\frac{2eq}{\alpha N'}\right)^{\alpha} \\ \Pr[\mathsf{FB}(\mathcal{Q}')] &\le N\left(\frac{2eq}{\alpha N'}\right)^{\alpha} \\ \Pr[\mathsf{Coll}(\mathcal{Q}') \wedge \neg\mathsf{Xor}(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')] &\le \frac{4q\alpha}{N'} + \frac{4q}{N'}\end{aligned}$$

since the sum of the above probabilities is an upper bound for $\Pr[\mathsf{Coll}(\mathcal{Q})]$. Moreover, by Proposition 1, $\Pr[\mathsf{Coll}(\mathcal{Q}') \wedge \neg\mathsf{Xor}(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')]$ can be upper bounded by finding upper bounds for $\Pr[\mathsf{Coll}_i(\mathcal{Q}') \wedge \neg\mathsf{Xor}(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')]$ for $i = 1, 2, 3$ and taking the sum of these. We now upper bound these various probabilities in a series of propositions. For these propositions q, N and α are as in Theorem 1, and $\mathcal{Q}'$ is the query history of any adversary A' as just specified. We emphasize that $|\mathcal{Q}'| \leq 2q$ and that probabilities are taken over the random cipher E and over the coins of A', if any (it inherits these from A).

Proposition 2. $\Pr[\mathsf{Xor}(\mathcal{Q}')] \leq N\left(\frac{2eq}{\alpha N'}\right)^{\alpha}$ *and* $\Pr[\mathsf{FB}(\mathcal{Q}')] \leq N\left(\frac{2eq}{\alpha N'}\right)^{\alpha}$.

Proof. Let $\mathcal{Q}' = \{(X'_i, K'_i, Y'_i)\}_{i=1}^{2q}$ denote the query history of A'. Since

$$\Pr[|\{i : X'_i \oplus Y'_i = Z\}| > \alpha] \leq \binom{2q}{\alpha}\left(\frac{1}{N'}\right)^{\alpha},$$

for each $Z \in \{0,1\}^n$, we have

$$\Pr[\mathsf{Xor}(\mathcal{Q}')] \leq N\binom{2q}{\alpha}\left(\frac{1}{N'}\right)^{\alpha} \leq N\left(\frac{2eq}{\alpha N'}\right)^{\alpha}.$$

$\Pr[\mathsf{FB}(\mathcal{Q}')]$ can be bounded similarly. □

Proposition 3. $\Pr[\mathsf{Coll}_1(\mathcal{Q}') \wedge \neg\mathsf{Xor}(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')] \leq 4q\alpha/N'$.

Proof. Let

$$\mathsf{Success}_1(\mathcal{Q}'_i) = \mathsf{Coll}_1(\mathcal{Q}'_i) \wedge \neg\mathsf{Coll}_1(\mathcal{Q}'_{i-1}) \wedge \neg\mathsf{Xor}(\mathcal{Q}'_{i-1}) \wedge \neg\mathsf{FB}(\mathcal{Q}'_{i-1})$$

for $i = 1 \ldots 2q$. Then $\Pr[\mathsf{Coll}_1(\mathcal{Q}') \wedge \neg\mathsf{Xor}(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')] \leq \sum_{i=1}^{2q} \Pr[\mathsf{Success}_1(\mathcal{Q}'_i)]$ and $\Pr[\mathsf{Success}_1(\mathcal{Q}'_i)] \leq \Pr[\mathsf{Coll}_1(\mathcal{Q}'_i) | \neg\mathsf{Coll}_1(\mathcal{Q}'_{i-1}) \wedge \neg\mathsf{Xor}(\mathcal{Q}'_{i-1}) \wedge \neg\mathsf{FB}(\mathcal{Q}'_{i-1})]$.

Fix a value of i, $1 \leq i \leq 2q$. We call the i-th query made by A' the *last query*. If $\mathsf{Success}_1(\mathcal{Q}'_i)$ occurs then either the adversary (i.e. A') can use its last query as query TL or as query BL of a collision in which TL, BL, TR and BR are distinct, by symmetry. Moreover the last query could either be a forward query or a backward query. This gives rise to four possible cases, and we bound $\Pr[\mathsf{Success}_1(\mathcal{Q}'_i)]$ for each separately. (We note the very first case, 'TL forward', is the case we discussed in Section 3.) For each case, we call the last query *successful* if this query completes a collision with TL, BL, TR, BR distinct and where the last query is used in the position stipulated by that case (e.g., for the case 'TL forward', the last query must be used in position TL).

TL forward: Let the last query be $E_{B\|L}(A)$. Call a value R *good* if there exists a query of the form $(B, L\|R, \cdot)$ in $\mathcal{Q}'$ that was obtained by A' as a backward query. We note that because of (7), $\neg\mathsf{FB}(\mathcal{Q}'_{i-1})$ implies there are at most α good R's.

We claim that for the last query to be successful the value R returned as an answer to the query must be good. Indeed, let R be the value returned; then a prerequisite for the query to be successful is that there be a query of the form $(B, L\|R, \cdot)$ in $\mathcal{Q}'_{i-1}$. We claim that this query must have been obtained as a backward query. Indeed, assume that the query $(B, L\|R, \cdot)$ was obtained as a forward query $E_{L\|R}(B)$ by A'. Then, by construction, A' would have immediately followed this query by the query $E^{-1}_{B\|L}(R)$ unless A' already knew the answer to $E^{-1}_{B\|L}(R)$. Either way A' would have the query $(A, B\|L, R)$ in its query history *prior* to the i-th (forward) query $E_{B\|L}(A)$, a contradiction since A' never makes a query to which it knows the answer. Thus the value R returned as an answer to the query $E_{B\|L}(A)$ must be good for the query to be successful.

Since there are at most α good values of R and since A' makes at most $2q$ queries, the probability that the last query is successful is therefore at most $\alpha/(2^n - 2q) = \alpha/N'$.

TL backward: Let the last query be $E^{-1}_{B\|L}(R)$. For the last query to be successful, there must be a (necessarily unique) query BL $= (B, L\|R, S) \in \mathcal{Q}'_{i-1}$, for some value $S \in \{0,1\}^n$. From the condition $B \oplus S = B' \oplus S'$ and from $\neg\mathsf{Xor}(\mathcal{Q}'_{i-1})$ there are at most α possibilities for the query BR. As each query BR uniquely determines the query TR,

there are at most α possibilities for the query TR as well, and thus at most α possibilities for the value $A' \oplus R'$. Thus the value A returned by the last query has chance at most α/N' that $A \oplus R$ will be equal to $A' \oplus R'$ for one of these values $A' \oplus R'$, and so the last query has chance at most α/N' of being successful.

BL forward: A 180° rotation of the collision diagram shows this case is symmetric to the case TL backward. The chance of success in this case is therefore at most α/N'.

BL backward: A 180° rotation of the collision diagram shows this case is symmetric to the case TL forward. The chance of success in this case is therefore at most α/N'.

The chance a forward last query is successful is therefore at most $2\alpha/N'$ (adding the TL and BL forward cases) and likewise the chance that a backward last query is successful is at most $2\alpha/N'$. Thus $\Pr[\mathsf{Success}_1(\mathcal{Q}'_i)] \le 2\alpha/N'$ for all i and $\sum_{i=1}^{2q} \Pr[\mathsf{Success}_1(\mathcal{Q}'_i)] \le 4q\alpha/N'$. □

Proposition 4. $\Pr[\mathsf{Coll}_2(\mathcal{Q}') \wedge \neg\mathsf{Xor}(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')] \le 2q/N'$.

Proof. Note that when TL = BL, $B\|L = L\|R$, so $B = L = R$; moreover $R = S$ and $A = B$, so $A = B = L = R = S$. For the adversary to obtain a collision with TL = BL, therefore, it must obtain a query of the form $(U, U\|U, U)$. The same argument applies to the case TR = BR. The chance of a query $E_{U\|U}(U)$ or of a query $E^{-1}_{U\|U}(U)$ being answered by U is at most[8] $1/N'$. Thus, since $2q$ queries are made total, $\Pr[\mathsf{Coll}_2(\mathcal{Q}')] \le 2q/N'$. □

Proposition 5. $\Pr[\mathsf{Coll}_3(\mathcal{Q}') \wedge \neg\mathsf{Xor}(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')] \le 2q\alpha/N' + 2q/N'$.

Proof. Note that in a collision with TL = BR we must have TL $\neq$ BL and $A \oplus R = B \oplus S$ (since $B \oplus S = B' \oplus S' = A \oplus R$, using TL = BR). Say the event $\mathsf{Coll}'_3(\mathcal{Q}'_i)$ occurs if there exist distinct queries $(A, B\|L, R)$, $(B, L\|R, S)$ in $\mathcal{Q}'_i$ such that $A \oplus R = B \oplus S$. With the same argument applied to the case BL = TR, we have $\mathsf{Coll}_3(\mathcal{Q}'_i) \implies \mathsf{Coll}'_3(\mathcal{Q}'_i)$. Therefore it suffices to show $\Pr[\mathsf{Coll}'_3(\mathcal{Q}') \wedge \neg\mathsf{Xor}(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')] \le 2q\alpha/N' + 2q/N'$.

The analysis now proceeds rather similarly to Proposition 3. Let

$$\mathsf{Success}'_3(\mathcal{Q}'_i) = \mathsf{Coll}'_3(\mathcal{Q}'_i) \wedge \neg\mathsf{Coll}'_3(\mathcal{Q}'_{i-1}) \wedge \neg\mathsf{Xor}(\mathcal{Q}'_{i-1}) \wedge \neg\mathsf{FB}(\mathcal{Q}'_{i-1}).$$

We have $\Pr[\mathsf{Coll}'_3(\mathcal{Q}') \wedge \neg\mathsf{Xor}(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')] \le \sum_{i=1}^{2q} \Pr[\mathsf{Success}'_3(\mathcal{Q}'_i)]$.

Fix a value of i, $1 \le i \le 2q$, and call the i-th query made by A' the *last query*. If $\mathsf{Success}'_3(\mathcal{Q}'_i)$ occurs then either the adversary (i.e. A') can use its last query as query TL or as query BL of its Coll'_3-solution. This gives rise to four possible cases given that the last query could be either forward or backward. In each case, we call the last query *successful* if $\mathsf{Success}'_3(\mathcal{Q}'_i)$ occurs and if the last query can be used in the position prescribed by that case (either TL or BL) in the Coll'_3-solution.

TL forward: We can use exactly the same analysis as in the case 'Forward TL' of Proposition 3. The probability that the last query is successful is therefore at most α/N'.

TL backward: Let $E^{-1}_{B\|L}(R)$ be the last query. For the last query to be successful, there must be a (necessarily unique) query of the form $(B, L\|R, S) \in \mathcal{Q}'_{i-1}$, for some $S \in \{0,1\}^n$. Since the answer A to the last query must be such that $A \oplus R = B \oplus S$ (as per the definition of Coll'_3) and $B \oplus S$ is uniquely determined, the last query has chance at most $1/N'$ of success.

BL forward: A 180° rotation of the collision diagram shows this case is symmetric to the case TL backward. The chance of success in this case is therefore at most $1/N'$.

BL backward: A 180° rotation of the collision diagram shows this case is symmetric to the case TL forward. The chance of success in this case is therefore at most α/N'.

[8] Since for each key there is only one relevant query, the tighter $1/N$ could be used as well.

The chance a forward last query is successful is therefore at most $(\alpha+1)/N'$ (adding the TL and BL forward cases) and likewise the chance that a backward last query is successful is at most $(\alpha+1)/N'$. Thus $\Pr[\mathsf{Success}'_3(\mathcal{Q}'_i)] \leq (\alpha+1)/N'$ for all i and $\sum_{i=1}^{2q} \Pr[\mathsf{Success}_1(\mathcal{Q}'_i)] \leq 2q\alpha/N' + 2q/N'$. (In fact, we even have $\Pr[\mathsf{Coll}_3(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')] \leq 2q\alpha/N' + 2q/N'$ since $\neg\mathsf{Xor}(\mathcal{Q}')$ was never used in the above.) □

Taking the sum of the bounds of Propositions 3, 4 and 5 one obtains that

$$\Pr[\mathsf{Coll}(\mathcal{Q}') \wedge \neg\mathsf{Xor}(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')] \leq \frac{6q\alpha}{N'} + \frac{4q}{N'}.$$

However, cases TL forward, BL backward and cases TL forward, BL backward of Propositions 3 and 5 reference the same events (the adversary is successful in case TL forward of Proposition 3 if and only if it is successful in case TL forward of Proposition 5, and likewise for the BL backward cases), which results in an "overcounting" of the adversary's probability of success by $2q\alpha/N'$. A more careful accounting of the adversary's probability of success thus shows

$$\Pr[\mathsf{Coll}(\mathcal{Q}') \wedge \neg\mathsf{Xor}(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')] \leq \frac{4q\alpha}{N'} + \frac{4q}{N'}. \tag{10}$$

Here we have not established (10) entirely formally, though this is the bound we use for $\Pr[\mathsf{Coll}(\mathcal{Q}') \wedge \neg\mathsf{Xor}(\mathcal{Q}') \wedge \neg\mathsf{FB}(\mathcal{Q}')]$ in Theorem 1. Establishing (10) formally would require dividing the event $\mathsf{Coll}(\mathcal{Q})$ into a different, less intuitive set of events than $\mathsf{Coll}_1(\mathcal{Q})$, $\mathsf{Coll}_2(\mathcal{Q})$, $\mathsf{Coll}_3(\mathcal{Q})$, events that are directly based on those that occur in the case analyses of Propositions 3–5. (For example, one of these events would be the event that the adversary ever obtains a "good R" through a forward or backward query, as defined for forward queries in case TL forward of Proposition 3 and implicitly defined (by symmetry) for backward queries in case BL backward of Proposition 3; another event would cover the cases TL backward and BL forward of Proposition 5; and so on.) The current form of the proof is our best compromise between readability and formality. In any case, the difference between $4q\alpha/N'$ and $6q\alpha/N'$ is relatively minor.

Summing (10) with the bounds of Proposition 2 and using (9), we obtain

$$\Pr[\mathsf{Coll}(\mathcal{Q})] \leq 2N \left(\frac{2eq}{\alpha N'}\right)^{\alpha} + \frac{4q\alpha}{N'} + \frac{4q}{N'}. \tag{11}$$

Since (11) holds for an arbitrary q-query adversary A, this establishes Theorem 1.

5 Preimage resistance

Ideally we would like to prove a strong bound on the everywhere preimage resistance [12] of Tandem-DM. In this notion, the adversary first gets to pick a challenge digest and subsequently (using oracle access to E) needs to find a preimage.

Unfortunately, Tandem-DM has the particularity that the point 0^{2n} is weaker than other range points with respect to preimage resistance. Indeed, to find a preimage of 0^{2n} (given a random blockcipher) an adversary can make queries of the form $E_{U\|U}(U)$ for different values of U until it finds a U such that $E_{U\|U}(U) = U$; then it is easy to see that $TDM^E(U\|U\|U) = 0^{2n}$. The probability (over the choice of E) of this attack succeeding in q queries is $1 - (1-1/N)^q \approx q/N = q/2^n$, since a different key is used for each query. On the other hand, we shall see that all nonzero points in $\{0,1\}^{2n}$ have much better preimage resistance than q/N, at least for q in the range of interest (i.e. $q = o(N), \omega(1)$). We also note this preimage attack on 0^{2n} is nearly matched by an easily-proved preimage resistance bound of $q/N' = q/(2^n - q)$ for 0^{2n} (or any other point in $\{0,1\}^{2n}$); the bound follows from the fact that a necessary condition for inverting 0^{2n} is to find a query with XOR output 0^n.

One solution for avoiding issues associated to 0^{2n} is to have the point-to-invert be chosen at random from $\{0,1\}^{2n}$; in this case there is chance at most $1/2^{2n}$ anyway that 0^{2n} is the image to invert. However, we find it slightly more interesting to emphasize that 0^{2n} is the only "bad" point in the range by letting the adversary choose which point to invert, under the stipulation that the adversary is not allowed to choose 0^{2n} (for which we anyway have the above

q/N' preimage resistance bound which, though worse than the preimage resistance bound we shall prove for nonzero points, is acceptable from a practical standpoint). A more detailed description of the preimage resistance experiment can be found below.

We note that Fleischmann et al. [3] claim to prove preimage resistance of the type $O(q/(2^n - q)^2)$ for Tandem-DM. Unfortunately, their analysis has similar flaws to their collision resistance proof. For example, while examining the case that the adversary's last query may be used in the top row of a solution for the preimage, FGL "give for free" the bottom row query if it is not already in the query history, and claim that the two queries (the last query and the free query) have combined chance of success $1/(2^n - q)^2$, since $A \oplus R$ must equal C_1 and $B \oplus S$ must equal C_2; the issue, once again, is that if the "free" bottom row query is already in the query history, there is no randomness left in the value $B \oplus S$ (whereas if the bottom row query was not in the query history, the adversary could not succeed anyway, so in this case there is in fact no need to give it the bottom query for free). Moreover, Fleischmann et al. ignore the possibility that the adversary may use the same query for the top and bottom row in its attack, which is associated to the issues regarding 0^{2n} discussed above; however, since they work in a model where the range point to invert is chosen at random, this particular omission would be easy to repair.

Our preimage resistance experiment will be as follows: an adversary A with oracle access to a randomly sampled blockcipher $E : \{0,1\}^{2n} \times \{0,1\}^n \to \{0,1\}^n$ selects and announces a point $C_1 \| C_2 \in \{0,1\}^{2n}$, $C_1 \| C_2 \neq 0^{2n}$, before making queries to E. The adversary wins after q queries if its query history $\mathcal{Q} = \{(X_i, K_i, Y_i)\}_{i=1}^q$ contains the means of computing a preimage of $C_1 \| C_2$, in the sense that there exist values A, B, L, R, $S \in \{0,1\}^n$ such that $A \oplus R = C_1$, $B \oplus S = C_2$ and such that the queries $(A, B \| L, R)$, $(B, L \| R, S)$ are in $\mathcal{Q}$. (In this case, we say $\mathcal{Q}$ *contains a preimage* of $C_1 \| C_2$.) We denote by

$$\mathbf{Adv}_{TDM}^{\mathrm{pre}\neq 0}(q)$$

the maximum advantage of any (probabilistic, computationally unbounded) adversary at this game. We note that here, too, n is a hidden parameter of the advantage. Moreover, we let

$$\mathsf{Preim}(\mathcal{Q})$$

be the predicate that is true if and only if $\mathcal{Q}$ contains a preimage of $C_1 \| C_2$, where $C_1 \| C_2$ is an elided-but-understood parameter of the predicate. Thus, $\mathbf{Adv}_{TDM}^{\mathrm{pre}\neq 0}(q)$ is the maximum of $\Pr[\mathsf{Preim}(\mathcal{Q})]$ taken over all q-query adversaries A, the probability being taken over E and the coins of A. We always assume that A is honest in the sense of choosing a nonzero value $C_1 \| C_2$. Now our preimage resistance theorem is the following (note that the definition of N' is different than in Theorem 1):

Theorem 2. *Let $N = 2^n$, $q < N$, $N' = N - q$ and let α be an integer, $1 \leq \alpha \leq q$. Then*

$$\mathbf{Adv}_{TDM}^{\mathrm{pre}\neq 0}(q) \leq 2\left(\frac{eq}{\alpha N'}\right)^{\alpha} + \frac{2\alpha}{N'}.$$

Proof. The "preimage diagram" for Tandem-DM is the left-hand portion of Fig. 4. While there are no "right-hand side" queries for the preimage diagram, we keep the labelling 'TL', 'BL' for the queries on the left-hand side. That is, in the preimage resistance game, the adversary's goal is to solve a "puzzle" by finding queries TL, BL of the form TL $= (A, B \| L, R)$, BL $= (B, L \| R, S)$ such that $A \oplus R = C_1$, $B \oplus S = C_2$. We emphasize at the outset that the case $C_1 = C_2$ does *not* require a separate analysis, and is handled in the same way as the case $C_1 \neq C_2$.

We start by noting that, in any solution of the preimage diagram, the queries TL, BL are necessarily distinct. Indeed, as discussed in Proposition 4, when the queries TL, BL are equal they have the form $(U, U \| U, U)$ and the output of Tandem-DM is 0^{2n}. We also note that if TL $= (A, B \| L, R)$ then BL $= (B, L \| R, B \oplus C_2)$, and, conversely, if BL $= (B, L \| R, S)$ then TL $= (R \oplus C_1, B \| L, R)$. Thus the queries TL, BL uniquely determine each other in the strong sense that *all three* coordinates of BL are determined by the query TL, and vice-versa.

As the preimage adversary A makes queries we maintain two sequences $\mathcal{W}_{\mathrm{TL}}$ and $\mathcal{W}_{\mathrm{BL}}$ called *wish lists*, which are initially empty, as well as two flags flag_1 and flag_2, which are initially zero. For each new query $(X, K_1 \| K_2, Y)$ learned by A we update the wish lists and the flags as follows:

1. If $(X, K_1 || K_2, Y) \in \mathcal{W}_{\mathrm{TL}}$ or $(X, K_1 || K_2, Y) \in \mathcal{W}_{\mathrm{BL}}$ then $\mathsf{flag}_1 \leftarrow 1$.
2. If $X \oplus Y = C_1$ then $(K_1, K_2 || Y, K_1 \oplus C_2)$ is added to $\mathcal{W}_{\mathrm{BL}}$.

3. If $X \oplus Y = C_2$ then $(K_2 \oplus C_1, X||K_1, K_2)$ is added to $\mathcal{W}_{\rm TL}$.
4. If $|\mathcal{W}_{\rm TL}| > \alpha$ or $|\mathcal{W}_{\rm BL}| > \alpha$, then $\mathsf{flag}_2 \leftarrow 1$.

We point out that, as long as A does not make redundant queries (which we assume it does not), the elements of $\mathcal{W}_{\rm TL}$ are all distinct from one another, as are the elements of $\mathcal{W}_{\rm BL}$. Indeed, it is easy to see that each element of $\mathcal{W}_{\rm TL}$ uniquely determines the query $(X, K_1||K_2, Y)$ which caused it to be added to $\mathcal{W}_{\rm TL}$, and likewise for $\mathcal{W}_{\rm BL}$.

We claim $\mathsf{Preim}(\mathcal{Q}) \implies \mathsf{flag}_1$. Indeed, if there are two queries $(A, B\|L, R)$, $(B, L\|R, S)$ in $\mathcal{Q}$ such that $A \oplus R = C_1$, $B \oplus S = C_2$, then one of these two queries was made after the other. Reasoning on both cases, we find that this query was an element of one of the wish lists at the point when it was learned, thus setting flag_1. (The reverse implication $\mathsf{flag}_1 \implies \mathsf{Preim}(\mathcal{Q})$ is also true and is trivial.) We thus have

$$\Pr[\mathsf{Preim}(\mathcal{Q})] = \Pr[\mathsf{flag}_1 = 1] \leq \Pr[\mathsf{flag}_2 = 1] + \Pr[\mathsf{flag}_1 = 1 \wedge \mathsf{flag}_2 = 0]. \tag{12}$$

We can bound

$$\Pr[\mathsf{flag}_2 = 1] \leq 2\left(\frac{eq}{\alpha N'}\right)^{\alpha}. \tag{13}$$

The proof of (13) is similar to Proposition 2, except that one omits the final union bound which results in the multiplication by N.

Let $\mathsf{WishGranted}_{{\rm TL},i}$ be the event that, at any point during the attack, A learns a query $(X, K_1\|K_2, Y)$, such that, at that moment and prior to the updating of the lists for that query, the i-th element of $\mathcal{W}_{\rm TL}$ is equal to $(X, K_1\|K_2, Y)$. Define $\mathsf{WishGranted}_{{\rm BL},i}$ in the same way. We then have

$$\Pr[\mathsf{flag}_1 = 1 \wedge \mathsf{flag}_2 = 0] \leq \sum_{i=1}^{\alpha} \Pr[\mathsf{WishGranted}_{{\rm TL},i}] + \sum_{i=1}^{\alpha} \Pr[\mathsf{WishGranted}_{{\rm BL},i}].$$

However, each wish list element can only be "wished for" once by A, due to the fact that $E_{K_1\|K_2}(\cdot)$ is a permutation. Thus $\Pr[\mathsf{WishGranted}_{{\rm TL},i}], \Pr[\mathsf{WishGranted}_{{\rm BL},i}] \leq 1/N'$ and so

$$\Pr[\mathsf{flag}_1 = 1 \wedge \mathsf{flag}_2 = 0] \leq \frac{2\alpha}{N'}. \tag{14}$$

By (13), (14) and (12) we obtain

$$\Pr[\mathsf{Preim}(\mathcal{Q})] \leq 2\left(\frac{eq}{\alpha N'}\right)^{\alpha} + \frac{2\alpha}{N'}$$

thus establishing the Theorem. □

Here also, α must be optimized numerically for given values of n and q. For $n = 128$, for example, Theorem 2 yields $\mathbf{Adv}^{\rm pre}_{TDM}(2^{127.0}) \leq 10^{-36}$ with $\alpha = 35$, $\mathbf{Adv}^{\rm pre}_{TDM}(2^{127.9}) \leq 10^{-35}$ with $\alpha = 95$ and $\mathbf{Adv}^{\rm pre}_{TDM}(2^{127.99}) \leq 10^{-33}$ with $\alpha = 468$. In fact, for $n = 128$ Theorem 2 gives a non-void upper bound for $\mathbf{Adv}^{\rm pre}_{TDM}(q)$ for values of q up to $\approx 2^{128-2^{-60}}$.

Theorem 2 should be compared with the trivial preimage resistance bound q/N' valid for any range point, that follows from the above-mentioned observation that inverting a point $C_1\|C_2$ in particular implies finding a query $(A, B\|L, R)$ such that $A \oplus R = C_1$ (there is chance at most $1/N'$ of this occurring for any query). Firstly, $q/N' = 1$ when $q = N/2$, whereas the bound of Theorem 2 implies that for δ constant, $\delta < 1$, and $q = \delta N$, $\mathbf{Adv}^{\rm pre}_{TDM}(q) \to 0$ as $n \to \infty$ with any $\alpha(n)$ such that $\lim_{n\to\infty} \alpha(n) = \infty$ and $\lim_{n\to\infty} \alpha(n)/N = 0$. Secondly, q/N' exhibits a linear growth in q for fixed n, whereas the bound of Theorem 2 pinpoints a much more "sudden threshold" of success, located near $q \approx N$; this is illustrated by the two graphs for the case $n = 128$, shown in Fig. 5.

Using Theorem 2 we can also derive a preimage resistance bound for the more standard definition of preimage resistance in which the adversary is given a random point in the range to invert. (A third definition, which we do not consider, samples the point to invert by sampling and evaluating a random point in the domain. For further discussion of these definitions and reductions among them, see [12].) Let $\mathbf{Adv}^{\rm pre\$}_{TDM}$ denote the maximum advantage of a q-query adversary at inverting a random point in $\{0,1\}^{2n}$, where the probability of inversion is also taken over the random choice of the point, and where "inverting the point" means, like above, constructing a query history that contains a preimage of the point. As an easy consequence of Theorem 2, we have:

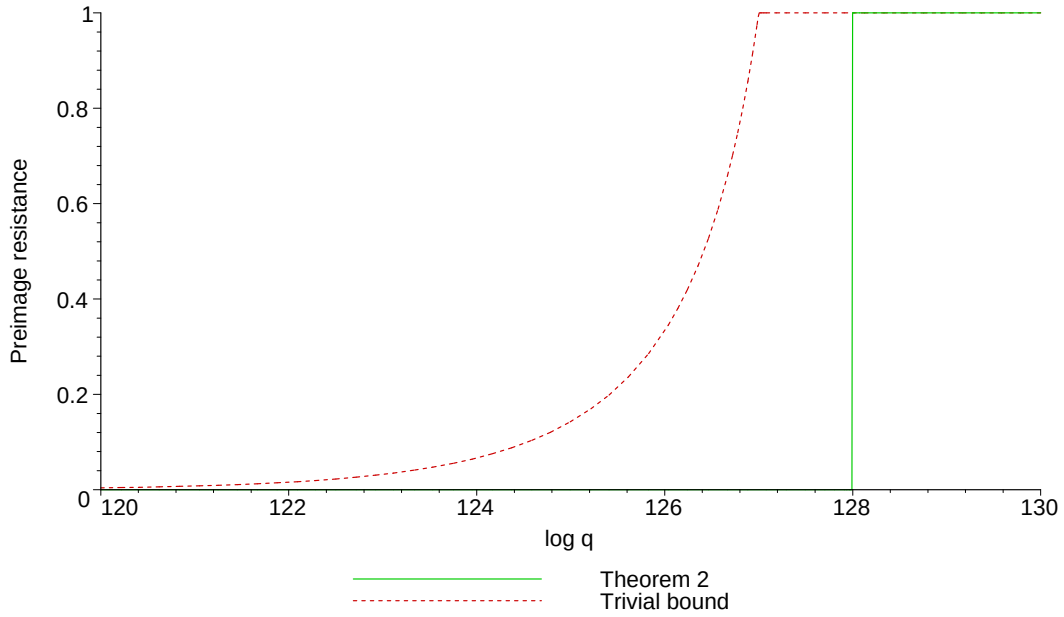

Fig. 5: Comparison between Theorem 2 and the trivial bound for $n = 128$. The Theorem 2 bound has a very sharp inflection point located near $q = 2^{128-2^{-60}}$.

Theorem 3. *Let $N = 2^n$, $q < N$, $N' = N - q$ and let α be an integer, $1 \leq \alpha \leq q$. Then*

$$\mathbf{Adv}^{\mathrm{pre\$}}_{TDM}(q) \leq 2\left(\frac{eq}{\alpha N'}\right)^{\alpha} + \frac{2\alpha}{N'} + \frac{q}{N^2 N'}.$$

Here the additional term q/N^2N' accounts for the event that the point to invert is 0^{2n}. This event happens with probability $1/N^2$, in which case the adversary has chance at most q/N' of success.

6 A generalization

In this section we give (without proof) a generalization of Tandem-DM that has the same level of collision resistance as Tandem-DM and that is subject to the same type of collision resistance analysis as the one we do in this paper. The main purpose of this section is not to propose a new scheme for potential implementation but rather to shed some additional light on Tandem-DM and on our proof by showing which key features enable our analysis.

Let $F_1 : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$, $F_2 : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$ be functions such that $F_i(U,\cdot)$ and $F_i(\cdot,U)$ are permutations of $\{0,1\}^n$ for any constant $U \in \{0,1\}^n$, $i = 1,2$. Let $G : \{0,1\}^n \times \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n \times \{0,1\}^n \times \{0,1\}^n$ be a permutation such that the first coordinate of G's output is determined by the first two coordinates of its input and such that the last coordinate of G's input is determined by the last two coordinates of its output (thus if $G(X,Y,Z) = (U,V,W)$ we can always compute U from X and Y only and always compute Z from V and W only). Moreover let $H : \{0,1\}^{3n} \to \{0,1\}^{3n}$ be an arbitrary permutation. Our generalization is the function $TDM^E_{F_1,F_2,G,H} : \{0,1\}^{3n} \to \{0,1\}^{2n}$ defined by

$$TDM^E_{F_1,F_2,G}(A\|B\|L) = F_1(X_1,Y_1)\|F_2(X_2,Y_2)$$

where

$$\begin{aligned} X_1\|K_1\|K_1' &= H(A\|B\|L) \\ Y_1 &= E_{K_1\|K_1'}(X_1) \\ X_2\|K_2\|K_2' &= G(K_1, K_1', Y_1) \\ Y_2 &= E_{K_2\|K_2'}(X_2) \end{aligned}$$

where the penultimate assignment identifies $(\{0,1\}^n)^3$ with $\{0,1\}^{3n}$. One may think of the value Y_1 as R and of the value Y_2 as S. We note that $TDM^E = TDM^E_{F_1,F_2,G,H}$ when G, H are identity functions and $F_1(X,Y) = F_2(X,Y) = X \oplus Y$.

This generalization is also preimage resistant up to the "trivial" bound of $q/(2^n - q)$. However we do not claim $TDM^E_{F_1,F_2,G,H}$ enjoys the same kind of preimage resistance as offered by Tandem-DM (under, say, the random-point-in-the-range model). Indeed, preimage resistance seems more subtle to bound than collision resistance, mainly because of attacks in which TL = BL and because we are not happy to give up a term $q/(2^n - q)$ for preimage resistance. We leave the worst-case preimage resistance of $TDM^E_{F_1,F_2,G,H}$ as an interesting open problem.

7 Conclusion

In this work, we have shown that an earlier work concerning the security of Tandem-DM was incorrect. However, with a new proof (exploiting new ideas) we have shown that, in the ideal-cipher model, Tandem-DM is collision resistant almost up to the birthday bound and (provably) preimage resistant essentially up to the birthday bound (leaving considerable room for improvement for the latter).

On a high level, our proof of collision resistance adheres to a (by now) standard framework. We first modify the collision-finding adversary by giving it several "free" queries and subsequently we bound the modified adversary's chance of success using a case analysis. This approach allows to easily bound both the number of free queries and the probability of a query (free or not) causing a collision.

In contrast, the FGL proof directly uses a case analysis and subsequently uses free queries within the case analysis. This ad hoc addition of free queries (and its binding to a particular case) is problematic, as it does not allow proper accounting of the free queries. In particular, if a free query is fresh it might cause a collision (or other bad event) elsewhere whereas if the free query has actually been asked before, no new randomness can be extracted from it.

Thus, apart from having established the security of Tandem-DM, we hope that our work also serves as a useful reminder to some of the subtleties involved in ICM proofs and as a guideline on how to avoid certain pitfalls.

References

1. Y. Dodis and J. Steinberger.: Message Authentication Codes from Unpredictable Block Ciphers. Crypto 2009, LNCS 5677, pp. 267–285. Springer, Heidelberg (2010). Full version available at `http://people.csail.mit.edu/dodis/ps/tight-mac.ps`
2. E. Fleischmann, C. Forler, M. Gorski and S. Lucks.: Collision Resistant Double-Length Hashing. ProvSec 2010, LNCS 6401, pp. 102–118. Springer, Heidelberg (2010)
3. E. Fleischmann, M. Gorski and S. Lucks.: On the security of Tandem-DM. FSE 2009, LNCS 5665, pp. 84–103. Springer, Heidelberg (2009)
4. E. Fleischmann, M. Gorski and S. Lucks.: Security of cyclic double block length hash functions, Cryptography and Coding, 12th IMA International Conference, Cirencester, UK, LNCS 5921 pp. 153–175. Springer, Heidelberg (2009)
5. S. Hirose.: Provably secure double-block-length hash functions in a black-box model. ICISC 2004, LNCS 3506, pp. 330–342. Springer, Heidelberg (2005)
6. S. Hirose.: Some plausible constructions of double-block-length hash functions. FSE 2006, LNCS 4047, pp. 210–225. Springer, Heidelberg (2006)
7. X. Lai and J. Massey.: Hash function based on block ciphers. Eurocrypt 1992, LNCS 658, pp. 55–70. Springer, Heidelberg (1993)
8. J. Lee and D. Kwon.: The security of Abreast-DM in the ideal cipher model. `http://eprint.iacr.org/2009/225.pdf`

9. J. Lee and J. Steinberger.: Multi-property preservation using polynomial-based modes of operation, Eurocrypt 2010, LNCS 6110, pp. 573–596. Springer, Heidelberg (2010)
10. S. Lucks.: A collision-resistant rate-1 double-block-length hash function. Symmetric Cryptography, Dagstuhl Seminar Proceedings 07021 (2007)
11. O. Özen and M. Stam.: Another Glance at Double-Length Hashing, Cryptography and Coding, 12th IMA International Conference, Cirencester, UK, LNCS 5921, pp. 94–115. Springer, Heidelberg (2009)
12. P. Rogaway and T. Shrimpton.: Cryptographic hash-function basics: definitions, implications, and separations for preimage resistance, second-preimage resistance, and collision-resistance. FSE 2004, LNCS 3017, pp. 371–388. Springer, Heidelberg (2004)
13. P. Rogaway and J. Steinberger.: Constructing cryptographic hash functions from fixed-key blockciphers. Crypto 2008, LNCS 5157, pp. 433–450. Springer, Heidelberg (2008)
14. T. Shrimpton and M. Stam.: Building a collision-resistant compression function from non-compressing primitives. ICALP 2008, Part II. LNCS 5126, pp. 643–654, Springer, 2008.
15. J. Steinberger.: The collision intractability of MDC-2 in the ideal-cipher model, Eurocrypt 2007, LNCS 4515, pp. 34–51. Springer, Heidelberg (2007)
16. M. Stam.: Beyond uniformity: better security/efficiency tradeoffs for compression functions, Crypto 2008, LNCS 5157, pp. 397–412. Springer, Heidelberg (2008)
17. M. Stam.: Blockcipher-based hashing revisited, FSE 2009, LNCS 5665, pp. 67–83. Springer, Heidelberg (2009)
18. D. Wagner.: Cryptanalysis of the Yi-Lam hash, Asiacrypt 2000, LNCS 1976, pp. 483–488. Springer, Heidelberg (2000)
19. X. Yi and K.-Y. Lam.: A new hash function based on block cipher. ACISP 1997, Second Australasian Conference on Information Security and Privacy, LNCS 1270, pp. 139–146. Springer, Heidelberg (1997)

A A second collision resistance proof: doing without the trick

The collision resistance analysis of Tandem-DM given in Section 4 depends on a rather subtle trick—namely, a modification of the adversary that allows us, later on in the proof, to infer that a particular type of query, if present in the query history, *must have been made in a particular direction* (forward or backward); knowing the query direction then allows us to conclude that very few queries of the given type can exist in the query history. This observation dispatches the most crucial (i.e. difficult) cases of the analysis. We refer the reader back to Section 4, in particular cases Forward TL of Proposition 3 and Forward TL of Proposition 5 for more details.

In this section we give a more standard collision analysis of Tandem-DM that does not use this trick. This analysis does not modify the adversary in any way (in particular, the adversary will make q queries, not $2q$) and resorts to sub-analyses for dealing with difficult cases (i.e. cases that were previously handled via our trick). Some of the sub-analyses require sub-sub-analyses of their own; the work is tedious, but straightforward, following the path laid out by previous recursive analyses of this type, in particular the analysis of MDC-2 [15] (such recursive analyses may also be found in [1, 13]). The purpose of presenting a second proof is purely for general interest: it shows "what the proof looks like" (and in particular its length) when our trick isn't used, and serves as a tutorial and reminder on the use of recursive analyses. In particular, the collision resistance bound derived from this second proof is of birthday-type, but worse than the bound of Theorem 1: while the adversary only makes q queries instead of $2q$, the greater number of cases considered in the proof yields a weaker bound overall. For $n = 128$, our second theorem gives that an adversary making $q = 2^{119.18}$ queries achieves chance < 0.5 of obtaining a collision (Theorem 1 gives $q = 2^{120.87}$).

For the proof, we reuse the predicate $\mathsf{Xor}(\mathcal{Q})$ of Section 2. Writing the query history as $\mathcal{Q} = (X_i, U_i \| V_i, Y_i)_1^q$ (i.e. decomposing K_i as $U_i \| V_i$) we also define two predicates XorKL, XorKR (where 'KL' and 'KR' stand for 'Key Left' and 'Key Right', respectively) as follows:

$$\mathsf{XorKL}(\mathcal{Q}) \iff \max_{C \in \{0,1\}^n} |\{i : X_i \oplus Y_i \oplus U_i = C\}| > \alpha,$$
$$\mathsf{XorKR}(\mathcal{Q}) \iff \max_{C \in \{0,1\}^n} |\{i : X_i \oplus Y_i \oplus V_i = C\}| > \alpha.$$

We emphasize that the parameter α which appears in these definitions is the "same α" as for $\mathsf{Xor}(\mathcal{Q})$. We additionally define five predicates $\mathsf{Triples}^+_{A\|B\|L}$, $\mathsf{Triples}^-_{S\|R\|L}$, $\mathsf{Doubles1}_{A\|B\|L\|S'\|R'\|L'}$, $\mathsf{Doubles2}^+_{A\|B\|L\|A'\|B'\|L'}$ and $\mathsf{Doubles2}^-_{S\|R\|L\|S'\|R'\|L'}$ using two new parameters $\beta > 0$ and $\gamma > 0$. The definitions for these more complicated

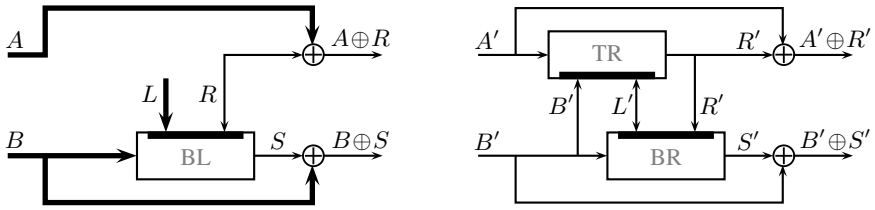

Fig. 6: Definition of $\mathsf{Triples}^{+}_{A\|B\|L}(\mathcal{Q})$. For every $A, B, L \in \{0,1\}^n$, the event $\mathsf{Triples}^{+}_{A\|B\|L}(\mathcal{Q})$ occurs iff there are more than γ values $R \in \{0,1\}^n$ for which there exists an ordered triple of distinct queries $(B, L\|R, S), (A', B'\|L', R'), (B', L'\|R', S') \in \mathcal{Q}$ such that $A \oplus R = A' \oplus R'$ and $B \oplus S = B' \oplus S'$, namely more than γ values R for which the adversary can complete the above partial collision diagram using distinct queries. The wires A, B, L are drawn in bold to indicate that their values are "externally fixed".

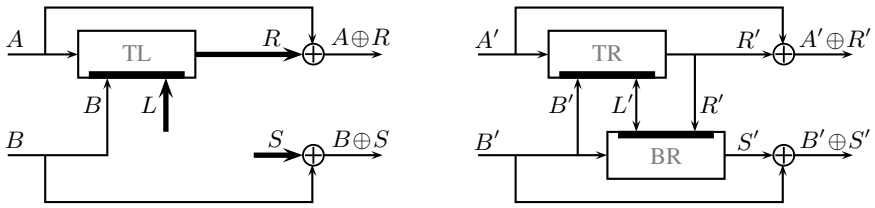

Fig. 7: Definition of $\mathsf{Triples}^{-}_{S\|R\|L}(\mathcal{Q})$. For every $S, R, L \in \{0,1\}^n$, the event $\mathsf{Triples}^{-}_{S\|R\|L}(\mathcal{Q})$ occurs iff there are more than γ values $B \in \{0,1\}^n$ for which there exists an ordered triple of distinct queries $(A, B\|L, R), (A', B'\|L', R'), (B', L'\|R', S') \in \mathcal{Q}$ such that $A \oplus R = A' \oplus R'$ and $B \oplus S = B' \oplus S'$, namely more than γ values B for which the adversary can complete the above partial collision diagram using distinct queries.

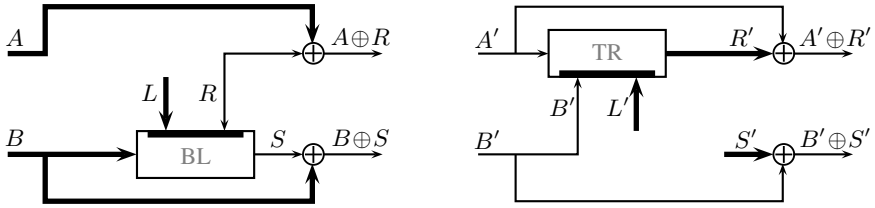

Fig. 8: Definition of the event $\mathsf{Doubles1}_{A\|B\|L\|S'\|R'\|L'}(\mathcal{Q})$. For every A, B, L, S', R', $L' \in \{0,1\}^n$, the event $\mathsf{Doubles1}_{A\|B\|L\|S'\|R'\|L'}(\mathcal{Q})$ occurs iff the adversary obtains more than β solutions to the above diagram, where a "solution" consists of an ordered pair of (distinct) queries $(B, L\|R, S), (A', B'\|L', R') \in \mathcal{Q}$ such that $A \oplus R = A' \oplus R'$ and $B \oplus S = B' \oplus S'$.

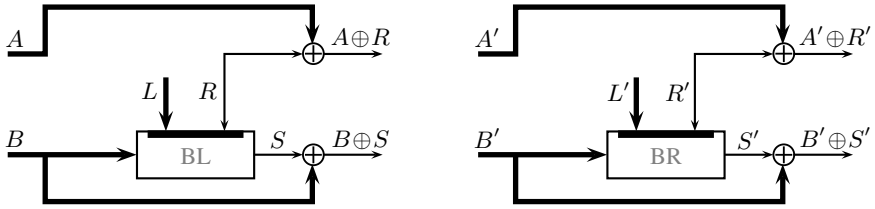

Fig. 9: Definition of the event $\mathsf{Doubles2}^{+}_{A\|B\|L\|A'\|B'\|L'}(\mathcal{Q})$. For every A, B, L, A', B', $L' \in \{0,1\}^n$, the event $\mathsf{Doubles2}^{+}_{A\|B\|L\|A'\|B'\|L'}(\mathcal{Q})$ occurs iff the adversary obtains more than β solutions to the above diagram, where a "solution" consists of an ordered pair of (distinct) queries $(B, L\|R, S), (B', L'\|R', S') \in \mathcal{Q}$ such that $A \oplus R = A' \oplus R'$ and $B \oplus S = B' \oplus S'$.

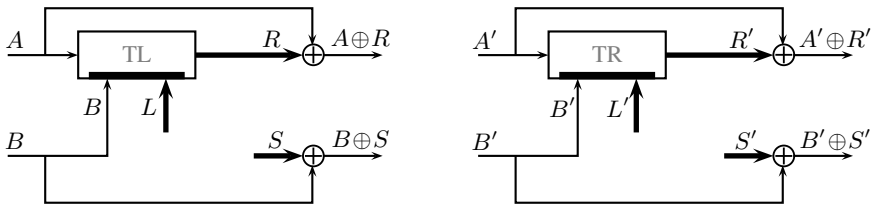

Fig. 10: Definition of the event $\mathsf{Doubles2}^-_{S\|R\|L\|S'\|R'\|L'}(\mathcal{Q})$. For every S, R, L, S', R', $L' \in \{0,1\}^n$, the event $\mathsf{Doubles2}^-_{S\|R\|L\|S'\|R'\|L'}(\mathcal{Q})$ occurs iff the adversary obtains more than β solutions to the above diagram, where a "solution" consists of an ordered pair of (distinct) queries $(A, B\|L, R), (A', B'\|L', R') \in \mathcal{Q}$ such that $A \oplus R = A' \oplus R'$ and $B \oplus S = B' \oplus S'$.

predicates are given in Figures 6–10. We note that the '$\mathsf{Doubles}$' events are defined with respect to the parameter β whereas the '$\mathsf{Triples}$' events are defined with respect to γ.

The reader may wonder as to the "logic" behind which wires are held constant in which diagram. Note, say, for $\mathsf{Triples}^+$, that the wires A, L and B are all those which "would be held constant" if we had fixed a certain forward query $E_{B\|L}(A)$ for position TL whose output R was not yet known; similarly, for $\mathsf{Triples}^-$, the wires S, R and L are those which would be held constant if we had fixed a backward query $E^{-1}_{L\|R}(S)$ for position BL whose output B was not yet known. The wires held constant in the event $\mathsf{Doubles1}$ are similarly obtained by fixing a forward query $E_{B\|L}(A)$ and a backward query $E^{-1}_{L'\|R'}(S')$ of unknown outputs, and so on for the events $\mathsf{Doubles2}^+$, $\mathsf{Doubles2}^-$.

We further define the existentially quantified versions of these predicates:

$$
\begin{aligned}
\mathsf{Triples}^+(\mathcal{Q}) &\iff \text{there exist } A, B, L \in \{0,1\}^n \text{ such that } \mathsf{Triples}^+_{A\|B\|C}(\mathcal{Q}) \\
\mathsf{Triples}^-(\mathcal{Q}) &\iff \text{there exist } S, R, L \in \{0,1\}^n \text{ such that } \mathsf{Triples}^-_{S\|R\|L}(\mathcal{Q}) \\
\mathsf{Doubles1}(\mathcal{Q}) &\iff \text{there exist } A, B, L, S', R', L' \in \{0,1\}^n \text{ such that } \mathsf{Doubles1}_{A\|B\|L\|S'\|R'\|L'}(\mathcal{Q}) \\
\mathsf{Doubles2}^+(\mathcal{Q}) &\iff \text{there exist } A, B, L, A', B', L' \in \{0,1\}^n \text{ such that } \mathsf{Doubles2}^+_{A\|B\|L\|A'\|B'\|L'}(\mathcal{Q}) \\
\mathsf{Doubles2}^-(\mathcal{Q}) &\iff \text{there exist } S, R, L, S', R', L' \in \{0,1\}^n \text{ such that } \mathsf{Doubles2}^+_{S\|R\|L\|S'\|R'\|L'}(\mathcal{Q}).
\end{aligned}
$$

We finally define the following shorthands:

$$
\begin{aligned}
\mathsf{X}(\mathcal{Q}) &= \mathsf{Xor}(\mathcal{Q}) \vee \mathsf{XorKL}(\mathcal{Q}) \vee \mathsf{XorKR}(\mathcal{Q}) \\
\mathsf{Triples}(\mathcal{Q}) &= \mathsf{Triples}^+(\mathcal{Q}) \vee \mathsf{Triples}^-(\mathcal{Q}) \\
\mathsf{Doubles2}(\mathcal{Q}) &= \mathsf{Doubles2}^+(\mathcal{Q}) \vee \mathsf{Doubles2}^-(\mathcal{Q}) \\
\mathsf{Doubles}(\mathcal{Q}) &= \mathsf{Doubles1}(\mathcal{Q}) \vee \mathsf{Doubles2}(\mathcal{Q}).
\end{aligned}
$$

Keeping the predicates $\mathsf{Coll}_1(\mathcal{Q})$, $\mathsf{Coll}_2(\mathcal{Q})$ and $\mathsf{Coll}_3(\mathcal{Q})$ as defined in Section 4, we have the following elementary implications:

$$
\begin{aligned}
\mathsf{Coll}(\mathcal{Q}) &\implies \mathsf{X}(\mathcal{Q}) \vee (\mathsf{Coll}_1(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})) \vee \mathsf{Coll}_2(\mathcal{Q}) \vee (\mathsf{Coll}_3(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})) \\
\mathsf{Coll}_1(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q}) &\implies (\mathsf{Triples}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})) \vee (\mathsf{Coll}_1(\mathcal{Q}) \wedge \neg\mathsf{Triples}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})) \\
\mathsf{Triples}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q}) &\implies (\mathsf{Doubles}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})) \vee (\mathsf{Triples}(\mathcal{Q}) \wedge \neg\mathsf{Doubles}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})) \\
\mathsf{Doubles}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q}) &\implies (\mathsf{Doubles1}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})) \vee (\mathsf{Doubles2}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})).
\end{aligned}
$$

(The first implication follows from Proposition 1.) Thus, we have

$$
\begin{aligned}
\Pr[\mathsf{Coll}(\mathcal{Q})] \le{}& \Pr[\mathsf{X}(\mathcal{Q})] + \Pr[\mathsf{Coll}_2(\mathcal{Q})] + \Pr[\mathsf{Coll}_3(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \\
&+ \Pr[\mathsf{Coll}_1(\mathcal{Q}) \wedge \neg\mathsf{Triples}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] + \Pr[\mathsf{Triples}(\mathcal{Q}) \wedge \neg\mathsf{Doubles}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \\
&+ \Pr[\mathsf{Doubles1}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] + \Pr[\mathsf{Doubles2}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})]. \qquad (15)
\end{aligned}
$$

We now proceed to individually upper bound each of the probabilities in (15). In each of the following propositions, $\mathcal{Q}$ is the query history of a q-query adversary, $N' = N - q = 2^n - q$ and α, β, γ are integers such that $1 \leq \alpha, \beta, \gamma \leq q$ and $\alpha \leq \beta$ and such that $\gamma \equiv 0 \bmod 3$ and $\beta \equiv 0 \bmod 2$. Moreover we let $\mathcal{Q}_i$ denote the adversary's query history after the first i queries (including the answer of the i-th query), as usual.

Proposition 6. $\Pr[\mathsf{X}(\mathcal{Q})] \leq 3N \left(\frac{eq}{\alpha N'}\right)^{\alpha}$.

Proof. We individually have $\Pr[\mathsf{Xor}(\mathcal{Q})] \leq N \left(\frac{eq}{\alpha N'}\right)^{\alpha}$, $\Pr[\mathsf{XorKL}(\mathcal{Q})] \leq N \left(\frac{eq}{\alpha N'}\right)^{\alpha}$ and $\Pr[\mathsf{XorKR}(\mathcal{Q})] \leq N \left(\frac{eq}{\alpha N'}\right)^{\alpha}$. Each of these inequalities can be proved as in Proposition 2. □

Proposition 7. $\Pr[\mathsf{Coll}_2(\mathcal{Q})] \leq q/N'$.

Proof. Same as for Proposition 4 (which does not use the assumptions $\neg\mathsf{Xor}(\mathcal{Q})$, $\neg\mathsf{FB}(\mathcal{Q})$). □

Proposition 8. $\Pr[\mathsf{Coll}_3(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \leq q(1+\alpha)/N'$.

Proof. Recall that, as observed in Proposition 5, $\mathsf{Coll}_3(\mathcal{Q})$ implies that the adversary obtains two distinct queries $(A, B\|L, R)$, $(B, L\|R, S)$ such that $A \oplus R = B \oplus S$. We let $\mathsf{Coll}'_3(\mathcal{Q})$ denote the latter event, and upper bound $\Pr[\mathsf{Coll}'_3(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})]$ instead.

We say the i-th query is *successful* if it can be used either as query TL = $(A, B\|L, R)$ or BL = $(B, L\|R, S)$ of a Coll'_3-solution, where the other query of the solution is in $\mathcal{Q}_{i-1}$. Fixing a value of i, we upper bound separately the probability that the i-th query can be used in position TL and that it can be used in position BL. We further divide each case into forward and backward queries, giving four cases to consider:

TL forward: Let $E_{B\|L}(A)$ be the i-th query of the adversary. For this query to be successful (in this case) there must be a query $(B, L\|R, S)$ in the query history such that $B \oplus S = R \oplus A$. Because $\neg\mathsf{X}(\mathcal{Q}_{i-1}) \implies \neg\mathsf{XorKR}(\mathcal{Q}_{i-1})$, there at most α such queries $(B, L\|R, S)$ in the query history, each determining a unique value of R. Thus the adversary's i-th query has chance of succeeding at most α/N'.

TL backward: Same analysis as case TL backward of Proposition 5, with chance of success at most $1/N'$.

BL forward: Symmetrical to case TL backward, with chance of success at most $1/N'$.

BL backward: Symmetrical to case TL forward, with chance of success at most α/N'.

Since each query must be either forward or backward (but not both), the chance of success of any given query is at most $(1+\alpha)/N'$, and the overall chance of success in q queries is at most $q(1+\alpha)/N'$. □

Proposition 9. $\Pr[\mathsf{Coll}_1(\mathcal{Q}) \wedge \neg\mathsf{Triples}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \leq q(\alpha+\gamma)/N'$.

Proof. By symmetry between the left- and right-hand sides of the collision diagram, we can divide the event $\mathsf{Coll}_1(\mathcal{Q})$ according to whether the last query made by the adversary to complete a Coll_1-type collision is used in position TL or BL. We further divide each of these two cases into forward and backward queries. We say the last (or "i-th") query of the adversary is *successful* if it completes a Coll_1-type collision.

TL forward: Let $E_{B\|L}(A)$ be the i-th query of the adversary. For this query to succeed in completing a Coll_1-type collision at position TL, the answer R be this query must be such that there exists a triple of (distinct) queries $(B, L\|R, S)$, $(A', B'\|L', R')$, $(B', L'\|R', S') \in \mathcal{Q}$ such that $A \oplus R = A' \oplus R'$ and $B \oplus S = B' \oplus S'$. However, $\neg\mathsf{Triples}(\mathcal{Q}_{i-1}) \implies \neg\mathsf{Triples}^{+}(\mathcal{Q}_{i-1})$ implies that there are most γ such values R. The chance of success of the i-th query is therefore at most γ/N'.

TL backward: Same analyis as case TL backward of Proposition 3, with chance of success at most α/N'.

BL forward: Symmetrical to case TL backward, with chance of success at most α/N'.

BL backward: Symmetrical to case TL forward, with chance of success at most γ/N'.

Since each query is either forward or backward, the chance of success of any given query is at most $(\alpha+\gamma)/N'$, and the overall chance of success in q is at most $q(\alpha+\gamma)/N'$. □

Proposition 10. $\Pr[\mathsf{Triples}(\mathcal{Q}) \wedge \neg\mathsf{Doubles}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \leq 6N^3 \left(\frac{3eq\beta}{\gamma N'}\right)^{\gamma/3}$.

Proof. We show

$$\Pr[\mathsf{Triples}^{+}(\mathcal{Q}) \wedge \neg\mathsf{Doubles}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \leq 3N^3 \left(\frac{3eq\beta}{\gamma N'}\right)^{\gamma/3}.$$

A similar analysis gives the same upper bound for $\Pr[\mathsf{Triples}^{-}(\mathcal{Q}) \wedge \neg\mathsf{Doubles}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})]$, thus yielding the bound.

We fix arbitrary values $A, B, L \in \{0,1\}^n$. It suffices to show

$$\Pr[\mathsf{Triples}^{+}_{A\|B\|L}(\mathcal{Q}) \wedge \neg\mathsf{Doubles}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \leq 3 \left(\frac{3eq\beta}{\gamma N'}\right)^{\gamma/3} \tag{16}$$

as the desired bound will then follow by a union bound over A, B, L.

We let $\#R(\mathcal{Q}_i)$ be the number of values $R \in \{0,1\}^n$ such that there exists an ordered triple of distinct queries $(B, L\|R, S)$, $(A', B'\|L', R')$, $(B', L'\|R', S') \in \mathcal{Q}_i$ such that $A \oplus R = A' \oplus R'$ and $B \oplus S = B' \oplus S'$. Moreover we let $\#R_{\mathrm{BL}}(\mathcal{Q}_i)$ be the number of values $R \in \{0,1\}^n$ such that a triple of this type exists *where the last query made completing the triple is used in position* BL, *namely where the first element of the triple is added to the query history after the last two elements*. We similarly define $\#R_{\mathrm{TR}}(\mathcal{Q}_i)$ and $\#R_{\mathrm{BR}}(\mathcal{Q}_i)$. Because

$$\#R(\mathcal{Q}) > \gamma \implies (\#R_{\mathrm{BL}}(\mathcal{Q}) > \gamma/3) \vee (\#R_{\mathrm{TR}}(\mathcal{Q}) > \gamma/3) \vee (\#R_{\mathrm{BR}}(\mathcal{Q}) > \gamma/3)$$

it suffices to show:

$$\Pr[\#R_{\mathrm{BL}}(\mathcal{Q}) > \gamma/3 \wedge \neg\mathsf{Doubles}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \leq \left(\frac{3eq\alpha}{\gamma N'}\right)^{\gamma/3} \tag{17}$$

$$\Pr[\#R_{\mathrm{TR}}(\mathcal{Q}) > \gamma/3 \wedge \neg\mathsf{Doubles}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \leq \left(\frac{3eq\beta}{\gamma N'}\right)^{\gamma/3} \tag{18}$$

$$\Pr[\#R_{\mathrm{BR}}(\mathcal{Q}) > \gamma/3 \wedge \neg\mathsf{Doubles}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \leq \left(\frac{3eq\beta}{\gamma N'}\right)^{\gamma/3} \tag{19}$$

to show

$$\Pr[\#R(\mathcal{Q}) > \gamma \wedge \neg\mathsf{Doubles}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \leq 3 \left(\frac{3eq\beta}{\gamma N'}\right)^{\gamma/3} \tag{20}$$

since $\alpha \leq \beta$.

We start by proving (17). Note, firstly, that $\#R_{\mathrm{BL}}(\mathcal{Q}_i) - \#R_{\mathrm{BL}}(\mathcal{Q}_{i-1}) \leq 1$ for all $i \geq 1$, because a query in position BL uniquely determines the value R. We now bound $\Pr[\#R_{\mathrm{BL}}(\mathcal{Q}_i) - \#R_{\mathrm{BL}}(\mathcal{Q}_{i-1}) = 1]$, considering two cases according to whether the i-th query is forward or backward. We write 'BL forward' and 'BL backward' to emphasize the query is to be used in position BL. We say the i-th (or last) query is *successful* if $\#R_{\mathrm{BL}}(\mathcal{Q}_i) - \#R_{\mathrm{BL}}(\mathcal{Q}_{i-1}) = 1$.

BL forward: Let $E_{L\|R}(B)$ be the i-th query. Then because $\neg\mathsf{X}(\mathcal{Q}_{i-1})$ there are at most α queries $(A', B'\|L', R') \in \mathcal{Q}_{i-1}$ such that $A' \oplus R' = A \oplus R$ (recall A is fixed) and each of these queries for position TR uniquely determines a query for position BR. Thus, there are at most α possibilities for the value $B' \oplus S'$ and thus at most α values S that would make the last query successful. The chance of success is therefore at most α/N'.

BL backward: Since the value B is fixed, the chance of success in this case is trivially at most $1/N'$.

Therefore, $\Pr[\#R_{\mathrm{BL}}(\mathcal{Q}_i) - \#R_{\mathrm{BL}}(\mathcal{Q}_{i-1}) = 1] \leq \alpha/N'$. Using a similar bound as in Proposition 2 (with q instead of $2q$, α/N' instead of $1/N'$ and $\gamma/3$ instead of α) we thus get (17).

We now prove (18). Here too we have $\#R_{\mathrm{TR}}(\mathcal{Q}_i) - \#R_{\mathrm{TR}}(\mathcal{Q}_{i-1}) \leq 1$ for all i. Indeed, a given value of $A' \oplus R'$ uniquely determines R, since A is fixed. We bound $\Pr[\#R_{\mathrm{TR}}(\mathcal{Q}_i) - \#R_{\mathrm{TR}}(\mathcal{Q}_{i-1}) = 1]$ considering two cases, according to whether the i-th query is forward or backward. We again say the i-th query is *successful* if $\#R_{\mathrm{TR}}(\mathcal{Q}_i) - \#R_{\mathrm{TR}}(\mathcal{Q}_{i-1}) = 1$.

TR forward: Let $E_{B'\|L'}(A')$ be the i-th query. Then because $\neg\mathsf{Doubles2}^+(\mathcal{Q}_{i-1})$ there are at most β pairs of queries $(B, L\|R, S)$, $(B', L'\|R', S') \in \mathcal{Q}_{i-1}$ such that $A \oplus R = A' \oplus R'$ and $B \oplus S = B' \oplus S'$. Each such pair determines a unique value R', and the output of the i-th query must be one of these values R' for the i-th query to be successful. Thus the i-th query is successful with chance at most β/N'.

TR backward: Let $E^{-1}_{B'\|L'}(R')$ be the i-th query. Then this query uniquely determines the BR query, so uniquely determines the values $B' \oplus S' = B \oplus S$ making at most α possibilities for the query BL (using $\neg\mathsf{X}(\mathcal{Q}_{i-1})$). But each query BL uniquely determines the values $A \oplus R$, so the last query has chance at most α/N' of being successful.

Since $\alpha \leq \beta$ we therefore have $\Pr[\#R_{\mathrm{TR}}(\mathcal{Q}_i) - \#R_{\mathrm{TR}}(\mathcal{Q}_{i-1}) = 1] \leq \beta/N'$, and (18) follows by a similar computation as in Proposition 2 (with q instead of $2q$, β/N' instead of $1/N'$ and $\gamma/3$ instead of α).

We finally prove (19). Once again we have $\#R_{\mathrm{BR}}(\mathcal{Q}_i) - \#R_{\mathrm{BR}}(\mathcal{Q}_{i-1}) \leq 1$ because a given query BR uniquely determines the query TR, which uniquely determines the value $A' \oplus R' = A \oplus R$ and hence the value R. We bound $\Pr[\#R_{\mathrm{BR}}(\mathcal{Q}_i) - \#R_{\mathrm{BR}}(\mathcal{Q}_{i-1}) = 1]$ using the same method and conventions as above:

BR forward: Let $E_{L'\|R'}(B')$ be the i-th query. This query uniquely determines the query TR, so uniquely determines the values of R (as A is fixed and $A' \oplus R' = A \oplus R$) and hence the query BL (as B, L are fixed), and so the value $B \oplus S$ is uniquely determined. The chance of success in this case is thus at most $1/N'$.

BR backward: Let $E^{-1}_{L'\|R'}(S')$ be the i-th query. Because $\neg\mathsf{Doubles1}(\mathcal{Q}_{i-1})$ there are at most β pairs of queries $(B, L\|R, S)$, $(A', B'\|L', R')$ such that $A \oplus R = A' \oplus R'$ and $B \oplus S = B' \oplus S'$. Each such pair of queries uniquely determines a value B', and an output of the i-th query cannot be successful unless it is the B' of such a pair. Thus the i-th query has chance of success at most β/N'.

Thus we have $\Pr[\#R_{\mathrm{BR}}(\mathcal{Q}_i) - \#R_{\mathrm{BR}}(\mathcal{Q}_{i-1}) = 1] \leq \beta/N'$, leading to (19) by the same computation as for (18). This concludes the proof of (20) which is the same as (16), and thus completes the proof of the proposition. □

Proposition 11. $\Pr[\mathsf{Doubles1}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \leq 2N^6 \left(\frac{2eq\alpha}{\beta N'}\right)^{\beta/2}$.

Proof. We fix values $A, B, L, S', R', L' \in \{0,1\}^n$. By a union bound, it suffices to show that

$$\Pr[\mathsf{Doubles1}_{A\|B\|L\|S'\|R'\|L'}(\mathcal{Q})] \leq 2\left(\frac{2eq\alpha}{\beta N'}\right)^{\beta/2}. \tag{21}$$

Let $\#D(\mathcal{Q}_i)$ be the number of pairs of queries $(B, L\|R, S)$, $(A', B'\|L', R') \in \mathcal{Q}_i$ such that $A \oplus R = A' \oplus R'$ and $B \oplus S = B' \oplus S'$. Moreover we let $\#D_{\mathrm{BL}}(\mathcal{Q}_i)$ be the number of such pairs where the query $(B, L\|R, S)$ was made after the query $(A', B'\|L', R')$, and let $\#D_{\mathrm{TR}}(\mathcal{Q}_i)$ be the number of such pairs where the query $(A', B'\|L', R')$ was made after the query $(B, L\|R, S)$. Since

$$\#D(Q) > \beta \implies (\#D_{\mathrm{BL}}(\mathcal{Q}_i) > \beta/2) \vee (\#D_{\mathrm{TR}}(\mathcal{Q}_i) > \beta/2)$$

it suffices to show

$$\Pr[\#D_{\mathrm{BL}}(\mathcal{Q}) > \beta/2 \wedge \neg\mathsf{X}(\mathcal{Q})] \leq \left(\frac{2eq\alpha}{\beta N'}\right)^{\beta/2} \tag{22}$$

$$\Pr[\#D_{\mathrm{TR}}(\mathcal{Q}) > \beta/2 \wedge \neg\mathsf{X}(\mathcal{Q})] \leq \left(\frac{2eq\alpha}{\beta N'}\right)^{\beta/2} \tag{23}$$

in order to show

$$\Pr[\#D(\mathcal{Q}) > \beta \wedge \neg\mathsf{X}(\mathcal{Q})] \leq 2\left(\frac{2eq\alpha}{\beta N'}\right)^{\beta/2}. \tag{24}$$

We show only (22) because the proof of (23) is entirely similar. First note that $\#D_{\mathrm{BL}}(\mathcal{Q}_i) - \#D_{\mathrm{BL}}(\mathcal{Q}_{i-1}) \leq 1$ for all i because a query BL uniquely fixes $A \oplus R$ which uniquely fixes $A' \oplus R'$ and hence uniquely fixes A', whereas the query BL also fixes $B \oplus S = B' \oplus S'$ and hence uniquely fixes B'; since the value L' is already fixed anyway, a query BL uniquely determines a query TR (and vice-versa).

We now bound $\Pr[\#D_{\mathrm{BL}}(\mathcal{Q}_i) - \#D_{\mathrm{BL}}(\mathcal{Q}_{i-1}) = 1]$, considering separately the cases when the i-th query is forward and backward. We label these cases as 'Forward BL' and 'Backward BL'. We say that the i-th query is *successful* if $\#D_{\mathrm{BL}}(\mathcal{Q}_i) - \#D_{\mathrm{BL}}(\mathcal{Q}_{i-1}) = 1$.

Forward BL: Let $E_{L\|R}(B)$ be the i-th query. Then $A \oplus R = A' \oplus R'$ is uniquely determined by the i-th query. In particular, since $\neg\mathsf{X}(\mathcal{Q}_{i-1})$ implies there are at most α queries of XOR output $A' \oplus R'$, there are at most α possibilities for query TR and hence α possibilities for the value $B' \oplus S'$. Hence the i-th query has chance of success at most α/N'.

Backward BL: Trivially, since B is fixed, the i-th query has chance of success at most $1/N'$.

In any case, thus, the chance of success of the i-th query is at most α/N'. The bound (22) then follows from a similar computation as in Proposition 2 (with q instead of $2q$, α/N' instead of $1/N'$ and $\beta/2$ instead of α).

Together with the (identical) proof of (23), this implies (24) which is equivalent to (21), and thus completes the proof. □

Proposition 12. $\Pr[\mathsf{Doubles2}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \leq 4N^5\left(\frac{2eq}{\beta N'}\right)^{\beta/2}$.

Proof. We show that

$$\Pr[\mathsf{Doubles2}^{+}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \leq 2N^5\left(\frac{2eq}{\beta N'}\right)^{\beta/2}$$

since the same bound can be proved for $\Pr[\mathsf{Doubles2}^{-}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})]$ by the same method.

We fix values A, B, L, A', B', $L' \in \{0,1\}^n$. By a union bound, it suffices to show that

$$\Pr[\mathsf{Doubles2}^{+}_{A\|B\|L\|A'\|B'\|B'}(\mathcal{Q}) \wedge \neg\mathsf{X}(\mathcal{Q})] \leq 2\left(\frac{2eq}{\beta N'}\right)^{\beta/2}. \tag{25}$$

(Indeed, since the constraint $A \oplus R = A' \oplus R'$ is equivalent to $R \oplus R' = A \oplus A'$, the two values A, A' are only as relevant as their xor $A \oplus A'$, thus removing one factor of N from the union bound; formally, we should say that we "fix a value of the xor $A \oplus A'$", but this is not notationally convenient.)

Overwriting the notation of Proposition 11, we now define $\#D(\mathcal{Q}_i)$ to be the number of pairs of queries $(B, L\|R, S)$, $(B', L'\|R', S') \in \mathcal{Q}_i$ such that $A \oplus R = A' \oplus R'$ and $B \oplus S = B' \oplus S'$. We further define $\#D_{\mathrm{BL}}(\mathcal{Q}_i)$ and $\#D_{\mathrm{BR}}(\mathcal{Q}_i)$ analogously to previous such definitions. Since

$$\#D(Q) > \beta \implies (\#D_{\mathrm{BL}}(\mathcal{Q}_i) > \beta/2) \vee (\#D_{\mathrm{BR}}(\mathcal{Q}_i) > \beta/2)$$

it suffices to show

$$\Pr[\#D_{\mathrm{BL}}(\mathcal{Q}) > \beta/2 \wedge \neg \mathsf{X}(\mathcal{Q})] \leq \left(\frac{2eq}{\beta N'}\right)^{\beta/2} \tag{26}$$

$$\Pr[\#D_{\mathrm{BR}}(\mathcal{Q}) > \beta/2 \wedge \neg \mathsf{X}(\mathcal{Q})] \leq \left(\frac{2eq}{\beta N'}\right)^{\beta/2} \tag{27}$$

in order to show

$$\Pr[\#D(\mathcal{Q}) > \beta \wedge \neg \mathsf{X}(\mathcal{Q})] \leq 2\left(\frac{2eq}{\beta N'}\right)^{\beta/2}. \tag{28}$$

We prove only (26) since (27) is analogous. We have $\#D_{\mathrm{BL}}(\mathcal{Q}_i) - \#D_{\mathrm{BL}}(\mathcal{Q}_{i-1}) \leq 1$ for all i because a query BL uniquely fixes $A \oplus R$ which uniquely fixes $A' \oplus R'$ and hence uniquely fixes R'. We now bound $\Pr[\#D_{\mathrm{BL}}(\mathcal{Q}_i) - \#D_{\mathrm{BL}}(\mathcal{Q}_{i-1}) = 1]$, considering separately the cases when the i-th query is forward and backward. As usual, we say that the i-th query is *successful* if $\#D_{\mathrm{BL}}(\mathcal{Q}_i) - \#D_{\mathrm{BL}}(\mathcal{Q}_{i-1}) = 1$.

Forward BL: Let $E_{L\|R}(B)$ be the i-th query. Then $A \oplus R = A \oplus R'$ is uniquely determined, so R' is uniquely determined and hence the query BR is uniquely determined. Thus $B' \oplus S'$ is uniquely determined and the i-th query has chance $1/N'$ of success.

Backward BL: Since B is fixed, the i-th query has chance of success at most $1/N'$.

In any case, thus, the chance of success of the i-th query is at most $1/N'$. The bound (22) then follows from a similar computation as in Proposition 2 (with q instead of $2q$ and $\beta/2$ instead of α).

Together with the (identical) proof of (27), this implies (28) which is equivalent to (25), and thus completes the proof. □

Adding together the bounds of Propositions 6–12, we thus obtain the following Theorem:

Theorem 4. *Let $1 \leq q < N$, $N' = N - q$. Let α, β, γ be integers between 1 and q such that $\alpha \leq \beta$, $\beta \equiv 0$ mod 2 and $\gamma \equiv 0$ mod 3. Then*

$$\begin{aligned}\Pr[\mathsf{Coll}(\mathcal{Q})] \leq\ & 3N\left(\frac{eq}{\alpha N'}\right)^{\alpha} + \frac{2q}{N'} + \frac{2q\alpha}{N'} + \frac{q\gamma}{N'} \\ & + 6N^3\left(\frac{3eq\beta}{\gamma N'}\right)^{\gamma/3} + 2N^6\left(\frac{2eq\alpha}{\beta N'}\right)^{\beta/2} + 4N^5\left(\frac{2eq}{\beta N'}\right)^{\beta/2}.\end{aligned}$$

For $n = 128$ and with $\alpha = 13$, $\beta = 156$ and $\gamma = 195$ Theorem 4 shows that an adversary making $q = 2^{119.18}$ achieves chance less than 0.5 of obtaining a collision. Moreover, Theorem 4 has the same "qualitative" corollary as Theorem 1:

Corollary 2. $\lim_{n\to\infty} \boldsymbol{Adv}^{\mathrm{TDM}}_{\mathrm{coll}}(N/n) = 0$.

Proof. Let $q = N/n$, $\alpha = 7n/\log n$ (more precisely, $\alpha = \lceil 7n/\log n \rceil$), $\beta = 2\alpha$ and $\gamma = 3\alpha$, where the logarithm takes base 2. Since $N' > N/2$ for $n > 2$, we have

$$\begin{aligned}\mathbf{Adv}^{\mathrm{coll}}_{TDM}(q) &\leq 3N\left(\frac{eq}{\alpha N'}\right)^{\alpha} + \frac{2q}{N'} + \frac{2q\alpha}{N'} + \frac{q\gamma}{N'} \\ &\quad + 6N^3\left(\frac{3eq\beta}{\gamma N'}\right)^{\gamma/3} + 2N^6\left(\frac{2eq\alpha}{\beta N'}\right)^{\beta/2} + 4N^5\left(\frac{2eq}{\beta N'}\right)^{\beta/2} \\ &\leq 3N\left(\frac{2eq}{\alpha N}\right)^{\alpha} + \frac{4q}{N} + \frac{10q\alpha}{N} + 6N^3\left(\frac{4eq}{N}\right)^{\alpha} + 2N^6\left(\frac{2eq}{N}\right)^{\alpha} + 4N^5\left(\frac{2eq}{\alpha N}\right)^{\alpha} \\ &\leq \left(3N + 4N^5\right)\left(\frac{2e\log n}{7n^2}\right)^{\frac{7n}{\log n}} + \frac{4}{n} + \frac{70}{\log n} + 6N^3\left(\frac{4e}{n}\right)^{\frac{7n}{\log n}} + 2N^6\left(\frac{2e}{n}\right)^{\frac{7n}{\log n}}.\end{aligned}$$

Using the equality $N = n^{n/\log n}$, we can show that the last expression goes to zero as $n \to \infty$. For example, the last term

$$2N^6 \left(\frac{2e}{n}\right)^{\frac{7n}{\log n}} = 2 \cdot n^{\frac{6n}{\log n}} \left(\frac{2e}{n}\right)^{\frac{7n}{\log n}} = 2\left(\frac{2e}{n^{\frac{1}{7}}}\right)^{\frac{7n}{\log n}}$$

goes to zero as $n \to \infty$. □

B A critique of the ProvSec'10 paper [2]

Fleischmann et al. [2] revisited the collision resistance of Tandem-DM in the context of a comprehensive generalization of various earlier works [3, 4, 8, 11]. Their main result of this new paper is Theorem 5 below.

To understand the theorem statement, some notations should be explained: 'Parallel-DL' and 'Serial-DL' are two classes of double block length blockcipher-based compression functions; Tandem-DM is a special case of Serial-DL. Moreover, ζ is a separate, scheme-dependent probability. For Tandem-DM one has $\zeta = 1/N^3$. (In fact Fleischmann et al. omit to explicitly compute ζ for Tandem-DM; see also the comments below.)

Theorem 5 (Fleischmann et al. [2, Theorem 3]). *Let H be a Parallel-DL or Serial-DL compression function as given in Section 2.3 [of [2]]. Let α, β, κ, n be constants such that $\alpha, k > e$ and let $\tau = N'\alpha/q, N = 2^n, N' = N - q$. Then*

$$\mathbf{Adv}_H^{\mathrm{coll}}(q) \le q \cdot \left(2\alpha/N' + 2\beta/(N')^2 + q\zeta/N' + 1/N'\right) + q\kappa/N' + L,$$

where

$$L = 2q2^n e^{\tau q(1-\ln\tau)/N'} + 2q^2/(\beta N') + 2qe^{\kappa(1+\ln\frac{q}{N'}-\ln\kappa)}.$$

(Compared to [2], we have slightly modified the statement of Theorem 5: the term $q\kappa/N'$, erroneously omitted in [2], has been added to the main bound; note that without this term, there would be no disadvantage to setting κ arbitrarily large; see also the comment 'A Missing Term', below.)

By setting $\alpha = 24, \beta = 2^{120}$ and $\kappa = 14$, Fleischmann et al. conclude that for $n = 128$ any adversary asking at most $2^{120.66}$ queries has chance less than 1/2 of obtaining a collision in a Serial-DL compression for which $\zeta = 0$. In Table 1 of their paper, Fleischmann et al. also list $2^{120.66}$ as the collision resistance of Tandem-DM, seemingly as a corollary of this computation, even though Tandem-DM has $\zeta = \frac{1}{N^3}$. (This seems to be an oversight, but since ζ is so small anyway for Tandem-DM, it is unlikely that taking into account the term $q\zeta/N'$ would actually change their bound for Tandem-DM; hence this oversight does not seem significant.)

Unfortunately, [2, Theorem 3] is incorrect; its derivation builds on the earlier mistakes by Fleischmann et al. [3] and adds a few of its own (tellingly, the errors in the original [3] are never mentioned). Below we will give a list of problems, where we will initially concentrate on a fatal one, concerning subcase "OL≠ QR" of "argument B" (we will explain what this means momentarily). Afterwards we will discuss some other (mostly non-fatal, but yet nontrivial to correct) errors in the paper.

THE MAIN FLAW. The main problem in [2, Theorem 3] relates to the terms $2q\beta/(N')^2$ and $2q^2/(\beta N')$. The parameter β itself is introduced to bound the occurrence of a certain configuration. For simplicity we look at the concrete case of Tandem-DM, in which case β is simply a bound on the number of collisions in the (single-length) Davies–Meyer construction. Concretely, define

$$\texttt{NumColl}(\mathcal{Q}) = |\{(i,j) \in \{1,\ldots,q\}^2, i \ne j : X_i \oplus Y_i = X_j \oplus Y_j\}|$$

then the claim is that if $\texttt{NumColl}(\mathcal{Q}) \le \beta$ then a specific type of collision is unlikely to occur. It is shown (in their full version) that $\Pr[\texttt{NumColl}(\mathcal{Q}) > \beta] \le 2q^2/(\beta N')$. So what is the probability of a collision occurring if we know that $\texttt{NumColl}(\mathcal{Q}) \le \beta$?

Argument B relates (among others) to the case that a $E_{B||L}(A)$ (forward) query is used in position TL and subcase OL≠QR in that case translates[9] to BL≠TR. The event $\texttt{NumColl}(\mathcal{Q}) \le \beta$ implies that there are at most β pairs BR, BL

[9] In [2] 'QR', 'OR', 'QL' and 'OL' are acronyms respectively for 'Query row Right', 'Other row Right', 'Query row Left' and 'Other row Left'. Roughly speaking QR ↔ TR, OR ↔ BR, QL ↔ TL, OL ↔ BL under our notation.

that cause a collision in the lower part of the diagram in Fig. 4). Since any query at position BR uniquely determines a corresponding TR query, we can instead consider β triples BR, BL, TR. For the current TL query to cause a collision, it needs to fit with one of these triples. This poses two restrictions on the output R (of TL): it needs to collide with the TR query ($R = A' \oplus R' \oplus A$) *and* it needs to fit with the BL query (R being part of the key input of BL). Since R is essentially random over a set of size (at least) N', the probability of satisfying either condition (for a fixed triple BR, BL, TR) is at most $1/N'$. From this, Fleischmann et al. conclude that the probability of satisfying *both* conditions is at most $(1/N')^2$ (which leads to the claimed $\beta/(N')^2$ after using a union bound). However, this multiplication of probabilities requires independence of events, yet both probabilities exploit exactly the same random choice of R. In fact, if there exists a possible R satisfying both constraints, it will be chosen with probability at least $1/N \gg (1/N')^2$, so clearly capping the probability at $(1/N')^2$ cannot be correct.

An easy fix to this problem is to use the correct bound β/N' instead. Unfortunately, the overall collision upper bound will then feature a term $q\beta/N' + 2q^2/(\beta N')$. It is not hard to see that $\beta \approx N^{1/3}$ is optimal, leading to a vacuous bound for $q > N^{2/3}$. In other words, one cannot prove collision resistance with this method for q beyond $2^{2n/3}$. This is still quite far removed from optimal collision resistance.

A possible solution would be to give a more accurate bound on the number of possible R that could complete a collision. Most of the β triples (BL, BR, TR) with BL, BR colliding are useless because the constraints imposed separately by TR and BL on the output R of the query TL are incompatible (i.e. no such R exists). However, as should be evident from Appendix A, a proper analysis bounding the number of relevant triples (BL, BR, TR) is far from trivial, and is in fact rather laborious.

PROBLEMS WITH THE CASE TL=BL. To upper bound the probability of the adversary obtaining a collision in which TL=BL, Fleischmann et al. [2, Definition 1] introduce the parameter ζ. Roughly speaking, ζ is the probability that once a query is made intended for the top row, that very same query would also fit in the bottom row. Here the probability is taken over the randomness of the ideal cipher, *as well as* over a random choice of the input (the latter being a key/plaintext pair, or a key/ciphertext pair in case of a backward query) to the top-row query.

Later, ζ is used in [2, Proposition 5] to bound the probability that an adversary poses a query that can be used concurrently in the top and bottom row. The mismatch should be obvious: in the definition of ζ the inputs (to the blockcipher) are randomly chosen, whereas for this proposition, they are adversarially chosen. The probability experiment used to define ζ, simply put, has no relation whatsoever to a setting in which queries are made by an adaptive adversary.

In general, one can fix this problem by redefining ζ: instead of randomizing over the blockcipher input, condition the probability on an arbitrary q-query history (subject, potentially, to some "lucky" events not having been triggered) and on an arbitrary input for the last query; only by doing the latter does one accurately capture the success chance of an adaptive adversary. In the specific case of Tandem-DM, we compute exactly such a bound in propositions 4 and 7, where we show the adversary's query has chance at most $1/N'$ of succeeding (one can slightly sharpen this result[10]). But we note that even with such a corrected (and much more stringent[11]) definition of ζ, the proof of Proposition 5 is actually still flawed. Fleischmann et al. multiply the probability ζ that the last query can be used both in positions TL and BL by the probability q/N' that the last query also collides with some query in position TR, whereas the two events (being functions of the last query's output) are of course not independent, rendering multiplication of these probabilities meaningless. We do note, however, that under the more stringent definition of ζ just discussed, $q\zeta$ is a correct and often satisfactory upper bound for the probability of an adversary mounting an attack where a query is used simultaneously in positions TL and BL.

INCORRECT PROOF OF LEMMA 3. In Lemma 3 of [2], Fleischmann et al. upper bound the probability that the adversary obtains queries Q_i and Q_j, $j \neq i$, such that query Q_i can be used both in positions TR and BL and query Q_j used in position BL. Firstly, Fleischmann et al. omit to consider that Q_j may be obtained after the query Q_i (the rest of this discussion presumes the only case they consider, namely that $j < i$). Secondly, letting Q_i be the last query, and

[10] Namely, an adversary whose *sole* purpose was to trigger the event $\mathsf{Coll}_2(\mathcal{Q})$ would only make queries of the form $E_{U\|U}(U)$ or $E_U^{-1}(U\|U)$ for (necessarily distinct) values of U; such an adversary (who is optimal for the purpose of triggering $\mathsf{Coll}_2(\mathcal{Q})$) has advantage exactly q/N. One can use this observation to sharpen propositions 4 and 7. We chose the current form of propositions 4 and 7 to keep our proof more uniform.

[11] For Tandem-DM, the value of ζ is $1/N'$ under the new definition, up from $1/N^3$ under the old.

from the fact that the XOR output of BL (to be filled by query Q_i) equals the XOR output of BR (to be filled by query Q_j for some $j < i$), Fleischmann et al. draw the erroneous conclusion that the last query Q_i has chance at most $1/N'$ of succeeding because, in their words, for a *given* XOR output of query Q_i there exists *at most one* matching XOR output (of a potential Q_j). This is similar to arguing, say, that a ball thrown at random among n slots has chance at most $1/n$ of falling in one of the first $n/2$ slots because for a *given* output of the experiment the ball has fallen into *at most one* of the first $n/2$ slots. Moreover, it is in fact easy to see there exists a partial query history Q_{i-1} with respect to which the i-th query has chance $(i-1)/N'$ of success, thus directly contradicting this claim.

A trivial correct upper bound on the probability of success of Q_i is q/N', as follows by union bounding over j. However this is quite suboptimal. A less trivial (and closer to optimal) upper bound, in the case of Tandem-DM, is α/N', where, here, α is the parameter used to define the events $\mathsf{XorKR}(\mathcal{Q})$, $\mathsf{XorKL}(\mathcal{Q})$ of (our) Appendix A. This presumes that $\neg\mathsf{XorKL}(\mathcal{Q}) \wedge \neg\mathsf{XorKR}(\mathcal{Q})$ holds, so one must separately account for $\Pr[\mathsf{XorKL}(\mathcal{Q}) \vee \mathsf{XorKR}(\mathcal{Q})]$, as in e.g. Proposition 6 of Appendix A.

THE USE OF FREE QUERIES. Argument A deals with the scenario that the input to the last query uniquely determines the input to the other row. Assume that the last query itself is used in position TL, then the corresponding BL query is uniquely determined. Fleischmann et al. "we give the query [BL] to the adversary for free." Under the normal and usual understanding of giving queries for free—namely posing that query without charging the adversary—this runs into the same problems as in their FSE paper (cf. Section 3).

A MISSING TERM. In Argument B, Subcase OL=QR the chance of success for any single query is upper bounded by κ/N'. Using a union bound, this should result in a term $q\kappa/N'$ in the final bound. This term is missing in [2, Theorem 3], as already noted.