

Indifferentiability of Confusion-Diffusion Networks

Yevgeniy Dodis^{1*}, Martijn Stam², John Steinberger^{3**}, and Liu Tianren⁴

dodis@cs.nyu.edu, csxms@bristol.ac.uk, jpsteinb@gmail.com, liutianren@gmail.com

Abstract. We show the first positive results for the indifferentiability security of the confusion-diffusion networks (which are extensively used in the design of block ciphers and hash functions). In particular, our result shows that a constant number of confusion-diffusion rounds is sufficient to extend the domain of a public random permutation.

1 Introduction

In this work we simultaneously address the following two questions:

- **Question 1:** secure domain extension of a *public* random permutation.
- **Question 2:** theoretical soundness of Shannon’s (or Feistel’s) confusion-diffusion paradigm.

DOMAIN EXTENSION OF RPs. The question of domain extension of various cryptographic primitives, such as encryption, signatures, message authentication codes, pseudorandom functions (PRFs), pseudorandom permutations (PRPs), etc., is one of the fundamental questions in cryptography.

In this paper we address a similar question for a *public* random permutation. Namely, given one (or a constant number of) n -bit random permutation(s) $P : \{0, 1\}^n \rightarrow \{0, 1\}^n$, and a number $w \geq 2$, build a wn -bit random permutation $\mathcal{Z} : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$. This question is clearly natural and interesting in its own right, but also seems extremely relevant in practice. Indeed, the random permutation model (RPM) has recently received a lot of attention [2, 12, 26, 29], starting to “compete with” and perhaps even “overtake” the more well known random oracle model (ROM) and the ideal cipher model (ICM). Aside from elegance, one of the reasons for this renewed attention comes from the fact that one can abstract the design of both the block-cipher standard AES and the new SHA-3 standard Keccak as being in the RPM. Namely, AES can be viewed as a 10-round *key-alternating* cipher applied to a concrete (“random-looking”) permutation, while SHA-3 can be viewed as applying a “sponge” mode of operation [2] to a similarly “random-looking” permutation. In fact, in his invited talk at Eurocrypt’13, the designer of both AES and SHA-3 Joan Daemen claimed that the RPM is much closer to the existing practice of designing hash functions and block ciphers than either the ROM or ICM, challenging the cryptographic community to switch to the RPM.

Of course, one must now build those “random looking permutations” \mathcal{Z} on relatively large domains (perhaps from 128 bits, like AES-128, to 1600 bits, like Keccak, or even longer). In practice, we have two well-known methods for accomplishing such a goal. The first method is based on applying several rounds of the Feistel network to some (not necessarily invertible) round functions.

* Partially supported by gifts from VMware Labs and Google, and NSF grants 1319051, 1314568, 1065288, 1017471.

** Supported by National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the National Natural Science Foundation of China Grant 61033001, 61361136003, and by the China Ministry of Education grant number 20121088050.

In our (public) setting, this method was theoretically analyzed only recently by Holenstein et al. [18] (building on an earlier work of [10]), who showed that a 14-round Feistel network is indeed sufficient for building a random permutation (RP), provided the round functions are modeled as (easily made) independent random oracles (ROs). Although very important in theory (i.e., showing the equivalence between ROM and RPM), this method does not seem to be used in practice, as it appears almost as hard,—if not *harder*,—to design “random-looking” non-invertible round functions on large domains as it is to design the desired random-looking permutation \mathcal{Z} .

CONFUSION-DIFFUSION PARADIGM. Instead, practitioners use the second method,—the *confusion-diffusion* (CD) paradigm,¹—which directly connects our motivating Questions 1 and 2. The idea of CD goes back to the seminal paper of Feistel [14] and even² back to Shannon [28]. Abstractly, one splits the input x to \mathcal{Z} into several shorter blocks $x_1 \dots x_w$, and then alternates the following two steps for several rounds: (a) *Confusion*, which consists of applying some fixed short permutations $P_1 \dots P_w$ (called *S-boxes*) to x_1, \dots, x_w ; and (b) *Diffusion*, which consists of applying some “mixing” non-cryptographic permutation $\pi(y_1 \dots y_w)$ (typically, carefully chosen linear function, sometimes also called *D-box*) to the results $y_1 \dots y_w$ of step (a).

Despite its extensive use in practice, the CD paradigm received extremely little attention from the theoretical cryptographic community.³ A notable exception is a beautiful work of Miles and Viola [23], who only looked at the secret-key setting—where the permutations P_1, \dots, P_w are secret—and also primarily considered the “weaker-than-indistinguishability” properties which can be proven about CD (and, more generally, SPN networks). In contrast, we are interested in the public setting, where the permutations P_i are modeled as RPs, and seek to examine the *indifferentiability* properties [8, 21] of the CD paradigm. This leads us to the following more precise reformulation of our motivating Questions 1 and 2:

- **Main question:** *Analyze indifferentiability of the confusion-diffusion paradigm as a way to extend the domain of a (constant number of) random permutation(s).* More precisely, for how many rounds r , and under what conditions on the *D*-boxes $\pi_1 \dots \pi_r$, is the r -round CD paradigm indifferentiable from an nw -bit random permutation \mathcal{Z} ?

Before presenting our results, we make a few remarks. First, we will model the “small permutations” P_i as both random and independent. The independence assumption is crucially used in our current proofs, but does not appear necessary. Unfortunately, the proofs we have are already extremely involved, so we feel this initial simplification is justified. We notice similar abstractions are made by most other papers in the area (including the seminal Luby-Rackoff paper [20] and the indifferentiability results of [10, 18]), though one hopes it might be lifted in future work.

As for modeling P_i as random, it seems inherent if we want to build a random permutation \mathcal{Z} ; e.g., we cannot build it from “nothing” (as this implies $P \neq NP$ and more), and it seems unlikely

¹ This is closely related to the substitution-permutation network (SPN) paradigm. Historically, though, the term SPN usually refers to the design of block ciphers as opposed to a single permutation, where one also XORs some key material in between successive CD rounds. To avoid confusion, we will stick with the term CD and not use the term SPN.

² Shannon [28] introduces “confusion” and “diffusion” into the cryptographic lexicon while Feistel [14] articulates the modern notion of a confusion-diffusion network, crediting Shannon with inspiration. There are some notable gaps between Shannon and the modern viewpoint. In particular Shannon does not seem to view confusion as a local operation, nor does he advocate repeatedly alternating steps of “confusion” and “diffusion”. Instead, Shannon seems to view confusion and diffusion as globally desirable attributes of a cryptographic mixing operation.

³ Of course, there is a lot of cryptanalytic work in the area whose survey is beyond the scope of this work.

rounds (flags)	D -boxes	$\varepsilon (w = 2)$	$q_S (w = 2)$	$\varepsilon (w > 2)$	$q_S (w > 2)$
5 (000)	arbitrary	$q^4/2^n$	q^4	$q^{2w}/2^n$	q^{w^2}
6 (100)	arbitrary	$q^2/2^n$	q^2	$q^2/2^n$	q^w
7 (110)	arbitrary	$q^2/2^n$	q	$q^2/2^n$	q
5 (000)	GF(2^n)-linear	$q^4/2^n$	q^4	1	—
6 (100)	GF(2^n)-linear	$[q^{10/3}/2^n]$	$[q^{10/3}]$	1	—
7 (110)	GF(2^n)-linear	$[q^5/2^n]$	$[q^{25/9}]$	1	—
9 (001)	GF(2^n)-linear	$q^4/2^n$	q^4	$q^{2w}/2^n$	q^{w^2}
10 (101)	GF(2^n)-linear	$[q^{10/3}/2^n]$	$[q^{10/3}]$	$[q^4/2^n]$	$[q^{2w-1/2}]$
11 (111)	GF(2^n)-linear	$[q^5/2^n]$	$[q^{25/9}]$	$[q^6/2^n]$	$[q^4]$

Fig. 1. Summary of security ε and simulator query complexity q_S (as functions of block length n , width w , and the number of distinguisher queries q) across our six main simulators with arbitrary or GF(2^n)-linear permutations. Entries in square brackets are not known, with the value inside the brackets being conjectured based on current best-known bounds. Constants and logarithmic factors are elided for simplicity. The meaning of the bit sequence next to each round number is explained in Section 5.

that any weaker assumption on the P_i will work. However, it does come with an important caveat: the best security bound ε we can naturally get with this approach will certainly be $\varepsilon \gg 2^{-n}$, where n is the domain of the S -boxes P_i . In practice, however, the S -boxes use a very small value of n (e.g., $n = 8$ for the AES), partly so that S -boxes can be easily and efficiently implemented as lookup tables. With such a small value of n , however, our bounds appear “practically meaningless”, irrespective of the number of queries q made by the attacker. This means that none of our results would be directly applicable to any of the “practical” permutations \mathcal{Z} used in the existing hash functions and block ciphers. Still, we believe establishing “structural soundness” of the CD paradigm is an important conceptual contribution—and an overdue sanity check—even with this serious (and inherent) limitation.

1.1 Overview of Our Results

We give a sequence of results establishing the soundness of the CD paradigm as a method for domain-extension of random permutations. Our indistinguishability results include CD networks of 5, 6, 7, 9, 10 and 11 rounds. These networks achieve different security levels, different query complexities, and place different combinatorial requirements on the D -boxes as well, even within the same network. Fig. 1 summarizes the main bounds achieved for each network length, up to lower-order (e.g., logarithmic) factors, and subject to various caveats to be shortly explained.

To read Fig. 1, recall that n is the block length of the S -boxes and that w is the width of the network. Moreover, q is the number of distinguisher queries, ε is the simulator’s security (i.e., the indistinguishability of the real and simulated worlds), and lastly q_S is the simulator’s query complexity (see Section 2 for definitions). The meaning of the 3-bit sequence next to each round number in the left column will be explained in Section 5.

The first three rows of Fig. 1 present our best bounds (i.e., assuming a “smart” choice of the D -boxes) when the D -boxes are not restricted to be GF(2^n)-linear. One can observe that in this part of the table the bounds for $w = 2$ are simply obtained by plugging in $w = 2$ to the general bounds. The best simulator here, at 7 rounds, achieves “birthday” security of $q^2/2^n$ and an essentially optimal query complexity of q .

Concerning the first three rows of Fig. 1 there is only one caveat: for some networks (specifically those of 6 and 7 rounds, and for both $w = 2$ and $w > 2$) the bounds presume that certain D -boxes have low “conductance”—a new critical property that we introduce and elaborate on below. Such permutations are known to exist on probabilistic grounds, but so far we do not know any explicit constructions, and building explicit permutations with low conductance is indeed one of the more interesting open problems raised by our work. (We do strongly believe, however, that permutations with short descriptions and low conductance exist. The problem is to find them.)

The last six rows of Fig. 1 concern the case when all D -boxes in the network are required to be $\text{GF}(2^n)$ -linear, which is a case of interest because it aligns with most practical constructions. (Even if the use of linear D -boxes in practical constructions seems to be, to a certain extent, a matter of tradition.) In this case, the simulators at width $w > 2$ and at 5, 6 and 7 rounds are not secure at all: $\varepsilon = 1$. In this regime, indeed, our simulator places certain combinatorial requirements on some of the D -boxes that are not satisfied by any $\text{GF}(2^n)$ -linear permutation. Fortunately, these requirements can be relaxed by using four more (i.e., 9, 10 or 11) rounds, so that $\text{GF}(2^n)$ -linear D -boxes become possible again for those round numbers at $w > 2$.

Unfortunately, except for the bounds at 5 rounds for $w = 2$ and the (rather poor) bounds at 9 rounds for $w > 2$, remaining bounds in the $\text{GF}(2^n)$ -linear section of the table are speculative, this being again related to “conductance”—specifically, the issue is that the lowest possible conductance of a $\text{GF}(2^n)$ -linear permutations is not currently known. We show some nontrivial lower bounds on the conductance of *generic* $\text{GF}(2^n)$ -linear permutations in Appendix C, but we have no similar nontrivial upper bounds! The conjectured bounds contained in Fig. 1 are obtained by using our lower bounds as a guess for the actual lowest possible conductance, and by rounding up some inconvenient exponents.⁴ The “true” values in the lower part of the table may well turn out to be *lower* than our conjectured values (if *non-generic* $\text{GF}(2^n)$ -linear permutations with low conductance turn out to exist) or *higher* (if even better lower bounds on the conductance of all $\text{GF}(2^n)$ -linear permutations are proved).

Still in the same part of the table, one can also note that going from 6 to 7 rounds or from 10 to 11 rounds entails a *decrease* in security. In our simulator, indeed, the actual purpose of adding the extra round (from 6 to 7 or from 10 to 11) is to improve the query complexity! (In fact a similar security deterioration also occurs between 6 and 7 rounds for the case of arbitrary D -boxes, but in that case the deterioration is only a logarithmic factor, and so cannot be seen in the table.)

COMBINATORIAL PROPERTIES OF D -BOXES. We note that our general theorem statement (found in Section 4) makes no distinction between linear and nonlinear cases; it simply expresses the security, query complexity and runtime of the simulator as a function of various combinatorial metrics (such as conductance mentioned above) of the diffusion permutations that are present in the network. Different metrics matter for different D -boxes, depending on their position in the networks, leading to a subtle (but also modular and fine-grained) result. In our opinion, the identification of precise (and distinct!) combinatorial requirements for each D -box of the network is one of the interesting contributions of this work, as it potentially allows future constructions to optimize the design of D -boxes layer by layer, with respect to the specific metrics targeted by each layer.

Moreover, some of our metrics entirely disappear when more rounds are added (specifically, by going from 5, 6, 7 rounds to 9, 10, 11 rounds respectively), leading to a relaxation of the conditions

⁴ If trivial upper bounds on conductance are applied instead, the rows for the 6- and 7-round networks with $w = 2$ in the second part of the table become the same as the row for the 5-round network in that half, while the rows for the 10- and 11-round networks become the same as the row for the 9-round network.

on the D -boxes at a larger number of rounds. For example, as already mentioned, security cannot be achieved at 5, 6 or 7 rounds with $w > 2$ and with $\text{GF}(2^n)$ -linear D -boxes, but can (unconditionally) be achieved by adding 4 more rounds to each of these networks, because a certain metric that no $\text{GF}(2^n)$ -linear permutation satisfies at $w > 2$ is no longer needed after the addition of the extra four rounds. In fact, although our table for general D -boxes (first 3 rows) in Fig. 1 does not include networks of 9, 10 and 11 rounds (since these networks would have the same security and simulator efficiency as the included networks for 5, 6 and 7 rounds, respectively) such networks are nonetheless considered by our main result, and might indeed be interesting from the point of view of increased efficiency. Namely, the weakened requirements on the D -boxes could potentially make these networks more cheap/fast/space efficient than their shorter, more combinatorially demanding counterparts! (Further such options, at even more rounds than considered by our main result, are explored in Appendix A.)

In all we identify four combinatorial metrics on D -boxes that are useful for our purposes, these being (and for lack of better terminology) “entry-wise randomized preimage resistance” (RPR), “entry-wise randomized collision resistance” (RCR), “conductance” and its cousin “all-but-one conductance”. The full definitions for these metrics can be found in Section 4. The first two metrics—RPR and RCR—are relatively unsurprising for our type of proof. (Briefly, they concern experiments in which all but one of the w input wires are fixed, and the final D -box input wire is drawn at random; the probability of a certain event occurring on the output wires should be low.) Moreover, there is not much “mystery” in RPR and RCR, since it happens that one can construct explicit permutations that achieve essentially optimal bounds for these metrics.⁵

Indeed, we consider conductance to be a more novel and interesting metric. (All-but-one conductance is, conceptually at least, very closely related to conductance.) We expand on this key metric now.

CONDUCTANCE. Briefly, conductance is a function of the number of queries q ; the conductance of a permutation $\pi : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$ at q queries is the maximum over all possible pairs of cartesian products $(U_1 \times \cdots \times U_w, V_1 \times \cdots \times V_w)$, where $U_i, V_i \subseteq \{0, 1\}^n$ and $|U_i| = |V_i| = q$ for each $1 \leq i \leq w$, of the numbers of pairs $(\mathbf{x}, \mathbf{y}) \in \{0, 1\}^{wn} \times \{0, 1\}^{wn}$ such that

$$\pi(\mathbf{x}) = \mathbf{y} \quad \text{and} \quad (\mathbf{x}, \mathbf{y}) \in (U_1 \times \cdots \times U_w, V_1 \times \cdots \times V_w)$$

In other words, one can choose q different values on each input and on each output wire, and one counts the number of input-output pairs (\mathbf{x}, \mathbf{y}) that “entirely fit” inside the induced cartesian products.

Now if one imagines the D -box π to be sandwiched between two rounds of S -boxes—as it will be in the network—the relevance of conductance to our setting can easily be guessed: U_i corresponds to the set of values that are queried outputs of the i -th S -box in the round before π , V_i corresponds to the set of values that are queried inputs of the i -th S -box in the round after π , and the conductance is an upper bound on the number of “all consistent input-output pairs that can be assembled” from these S -box queries under π ’s mapping. Intuitively, if the number of such pairs is low, the job of the indistinguishability simulator is easier, as it has to worry about fewer “consistent chains” that the distinguisher is trying to assemble.

It is easy to see from the definition that conductance for any permutation $\pi : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$ lies between q and q^w . In Appendix B we moreover show that a random permutation has

⁵ On the other hand, an interesting research direction might be to find explicit constructions of RPR and RCR permutations that achieve higher speeds than our own naïve constructions!

conductance close to $wqn \approx q$. On the other hand, in Appendix C we show that the conductance of a generic $\text{GF}(2^n)$ -linear permutation is at least $q^{2^{-1}/(2^w-1)}$, which is always super-linear in q . An already-mentioned corollary is that, and at least with respect to our current simulator, having linear D -boxes seems to cause strictly worse security than what is achievable by arbitrary (albeit currently non-constructive!) D -boxes. Since it may well turn out that low conductance is instantiable by D -boxes that are no slower than the current $\text{GF}(2^n)$ -linear D -boxes, our work raises, among others, the question of whether $\text{GF}(2^n)$ -linearity is really the right choice for D -boxes in the CD paradigm.

SUMMARY. Overall, we show the first positive results for the indistinguishability security of the confusion-diffusion paradigm which is extensively used in the design of block ciphers and hash functions. Our result shows that a constant number of confusion-diffusion rounds is sufficient to extend the domain of a public random permutation. In the process, we reduced the indistinguishability properties of the CD network (for a variety of rounds between 5 and 11) to natural and novel combinatorial properties of the D -boxes (such as conductance), which we hope will lead to a better understanding of the confusion-diffusion networks, and will be useful in the future design and analysis of block ciphers and hash functions.

1.2 Other Related Work

The question of domain extension ideal primitives was considered by [8, 22] for the setting of public random functions (ROM), and by [9] for the setting of block ciphers (ICM). While none of these domain extensions directly apply to the RPM (e.g., the result of [9] crucially relies on the existence of a key for the “small ideal cipher”), they can be composed with the series of results showing the equivalence between RPM, ICM and ROM [1, 8, 10–13, 18] to achieve various theoretical domain extension methods in the RPM. For example, one can get a “small RO” from “small RP” [11, 12], extend domain of RO [8], and apply the 14-round Feistel construction to get a large-domain RP [18] (many other combinations of prior results will suffice as well). However, all such combinations of prior work will be *much* less efficient (and elegant) than our natural construction, and, more importantly, such results will not correspond to the way random permutations are built in real life.

The domain extension of *secret-key* random permutations is well studied: examples include PEP [4], XCB [16], HCTR [32], HCH [5] and TET [17] (and even the original Feistel constructions [20, 24] could be viewed as domain doubling techniques in this setting). However, it is easy to see that none of those constructions provide the indistinguishability property in the public permutation setting.

Finally, the design of public permutations is related in spirit to the area of *white-box cryptography* [6, 7], with the idea to “obfuscate” key-dependent parts of the cipher and publish them as lookup tables, making the entire construction public. We refer to [3] for an excellent discussion of this big area of research, as well as a survey of many cryptanalytic efforts attacking various SPN designs with linear diffusion layers.

2 Definitions

BASIC NOTATIONS. We write $[w]$ for the set of integers $\{1, \dots, w\}$. Vectors in \mathbb{F}^w , where \mathbb{F} denotes a finite field, are written in bold letters such as \mathbf{x} , \mathbf{y} . The i -th entry of \mathbf{x} , $i \in [w]$, is written $\mathbf{x}[i]$. We put $N = 2^n$.

Except in Appendices B and C, where for the sake of conceptual correctness we work with an arbitrary field \mathbb{F} , we will set $\mathbb{F} = \text{GF}(2^n)$. Then elements in \mathbb{F} can (under some representation) be

identified with n -bit strings, and moreover a string in $\{0, 1\}^{wn}$ can be identified with an element $\mathbf{x} \in \mathbb{F}^w$.

RANDOM PERMUTATIONS AND IDEAL PRIMITIVES. In this paper, the only ideal primitives under consideration are random permutations. A random permutation from $\{0, 1\}^k$ to $\{0, 1\}^k$ is a permutation drawn uniformly at random from the set of all permutations from $\{0, 1\}^k$ to $\{0, 1\}^k$. Such a permutation is implemented as an oracle accessible to a specified subset of the parties performing a security experiment. In our setting all parties also have access to the inverse oracle (i.e., implementing the inverse of the random permutation).

CONFUSION-DIFFUSION NETWORKS. Fix integers $w, n, r \in \mathbb{N}$. Let

$$\mathcal{P} = \{P_{i,j} : (i, j) \in [r] \times [w]\}$$

be an array of rw permutations from $\{0, 1\}^n$ to $\{0, 1\}^n$, i.e., $P_{i,j}$ is a permutation from $\{0, 1\}^n$ to $\{0, 1\}^n$ for each $i \in [r]$ and each $j \in [w]$. Also let

$$\bar{\pi} = (\pi_1, \dots, \pi_{r-1})$$

be an arbitrary sequence of $r - 1$ permutations, each from $\{0, 1\}^{wn}$ to $\{0, 1\}^{wn}$.

Given \mathcal{P} and $\mathbf{x} \in \{0, 1\}^{wn}$ we let

$$P_i(\mathbf{x})$$

denote the value in $\{0, 1\}^{wn}$ obtained by applying the permutations $P_{i,1}, \dots, P_{i,w}$ blockwise to \mathbf{x} . In other words, $P_i : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$ is defined by setting

$$P_i(\mathbf{x})[j] = P_{i,j}(\mathbf{x}[j])$$

for all $j \in [w]$. It is obvious that P_i is a permutation of $\{0, 1\}^{wn}$.

Given \mathcal{P} and $\bar{\pi}$, we define the permutation $P = P[\mathcal{P}, \bar{\pi}]$ from $\{0, 1\}^{wn}$ to $\{0, 1\}^{wn}$ as the composition

$$P[\mathcal{P}, \bar{\pi}] = P_r \circ \pi_{r-1} \circ \dots \circ P_2 \circ \pi_1 \circ P_1.$$

I.e.,

$$P[\mathcal{P}, \bar{\pi}](\mathbf{x}) = P_r(\pi_{r-1}(\dots P_2(\pi_1(P_1(\mathbf{x}))) \dots))$$

for $\mathbf{x} \in \{0, 1\}^{wn}$. We call $P[\mathcal{P}, \bar{\pi}]$ the *confusion-diffusion network* built from \mathcal{P} and $\bar{\pi}$. The permutations in \mathcal{P} are variously called the *confusion permutations* or *S-boxes*. The permutations in $\bar{\pi}$ are variously called the *diffusion permutations* or *D-boxes*.

The values n , w and r will be called the *wire length*, the *width* and the *number of rounds* respectively.

In practice, the *S-boxes* are implemented by “convoluted” or “random-like” permutations while the *D-boxes* are implemented by “easy” (typically linear) permutations that are cryptographically weak. In our indistinguishability model, described next, the *S-boxes* are modeled as random permutations while the *D-boxes* are publically fixed parameters of the network.

INDIFFERENTIABILITY. Let C be a construction making calls to an ideal set of primitives \mathcal{P} , which we notate as $C^{\mathcal{P}}$. Let \mathcal{Z} be an ideal primitive with the same interface as $C^{\mathcal{P}}$ (e.g., \mathcal{Z} is a random permutation if $C^{\mathcal{P}}$ implements a permutation). Indistinguishability is meant to capture the intuitive notion that the construction $C^{\mathcal{P}}$ is “just as good” as \mathcal{Z} , in some precise sense. The definition involves a simulator:

Definition 1. An (oracle) circuit C with access to a set of ideal primitives \mathcal{P} is (t_S, q_S, ε) -indifferentiable from an ideal primitive \mathcal{Z} if there exists a simulator S such that

$$\Pr \left[D^{C^{\mathcal{P}}, \mathcal{P}} = 1 \right] - \Pr \left[D^{\mathcal{Z}, S^{\mathcal{Z}}} \right] \leq \varepsilon$$

for every distinguisher D making at most q_0 queries to its oracles, and such that S runs in total time t_S and makes at most q_S queries to \mathcal{Z} . Here t_S , q_S and ε are functions of q_0 .

We note that in the “real world” D has oracle access to the construction $C^{\mathcal{P}}$ as well as to the primitives \mathcal{P} ; in the “ideal world” $C^{\mathcal{P}}$ is replaced by the ideal primitive \mathcal{Z} and the ideal primitives \mathcal{P} are replaced by the simulator S . Thus, S ’s job is to make \mathcal{Z} look like $C^{\mathcal{P}}$ by inventing “answers that fit” for D ’s queries to the primitives in \mathcal{P} . For this, S requires query access to \mathcal{Z} (notated as $S^{\mathcal{Z}}$); on the other hand, S does not get to see which queries D is making to \mathcal{Z} .

We emphasize that D is information-theoretic: there is no hardness assumption. What is at stake is the statistical distance between two worlds with two different underlying sources of randomness (\mathcal{P} versus \mathcal{Z}).

Informally, $C^{\mathcal{P}}$ is *indifferentiable* from \mathcal{Z} if it is (t_S, q_S, ε) -indifferentiable for “reasonable” values of (t_S, q_S, ε) . An essential composition theorem [8, 21] states that any cryptosystem that is secure when implemented with \mathcal{Z} remains secure if \mathcal{Z} is replaced with $C^{\mathcal{P}}$, if $C^{\mathcal{P}}$ is indifferentiable from \mathcal{Z} . However, the class of adversaries with respect to which the cryptosystem’s security is defined must be a class that is large enough to accommodate the simulator S from Definition 1. See, e.g., [25] for a dramatic example in which indifferentiability fails completely.

The value q_S in Definition 1 is called the *query complexity* of the simulator. Having a small value of q_S is important for “quality of composition”, i.e., to avoid an unacceptable watering-down of the security bounds when the afore-mentioned composition theorem is applied; ideally q_S should be as close to q_0 as possible.

In our setting “ \mathcal{P} will be \mathcal{P} ” (i.e., the set of ideal primitives \mathcal{P} will be the set of *wr* independent random permutations discussed in the previous subsection), while $C^{\mathcal{P}}$ will be $P[\mathcal{P}, \bar{\pi}]$. (As explained, the diffusion permutations $\bar{\pi}$ are a fixed, publically known parameter of the construction.) Consequently, \mathcal{Z} (matching $C^{\mathcal{P}}$ ’s syntax) will be a random permutation from $\{0, 1\}^{wn}$ to $\{0, 1\}^{wn}$. Like all permutation oracles, \mathcal{Z} can be queried in both forward and backward directions.

3 Attack on Two-round Confusion-Diffusion Networks

In this section we outline a simple distinguishing attack that shows confusion-diffusion networks of two rounds or less cannot be indifferentiable from a random permutation. Unfortunately we could not find a similarly general attack for networks with three rounds, which leaves open the possibility that 3- or 4-round confusion-diffusion network might already be indifferentiable.

The attack on 2-round networks requires $w \geq 2$, which is indeed a trivial requirement since if $w = 1$ then a 1-round network is already indifferentiable from a random permutation.

For concreteness we sketch the attack with $w = 2$. The confusion-diffusion network then has four S -boxes labeled $P_{i,j}$ for $(i, j) \in [2] \times [2]$ and one diffusion permutation $\pi : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$. The S -boxes in the first round are $P_{1,j}$, $j \in [2]$, the S -boxes in the second round are $P_{2,j}$, $j \in [2]$.

We will say the distinguisher “rejects” if it believes that it is in the simulated world; “accepts” if it believes it is in the real world.

The distinguishing attack is as follows:

1. The distinguisher randomly chooses $\mathbf{x} \in \{0, 1\}^{2n}$ and queries $\mathcal{Z}(\mathbf{x})$, where $\mathcal{Z} : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$ is the random permutation, obtaining $\mathbf{y} \in \{0, 1\}^{2n}$ as answer.
2. The distinguisher make the two S -box queries $P_{1,1}(\mathbf{x}[1])$ and $P_{2,1}^{-1}(\mathbf{y}[1])$ receiving answers $a \in \{0, 1\}^n$ and $b \in \{0, 1\}^n$ respectively.
3. If there exists no pair of values (c, d) such that $\pi(a||c) = (b||d)$, the distinguisher rejects.
4. If there exists a pair of values (c, d) such that $\pi(a||c) = (b||d)$, the distinguisher chooses any such pair, queries $P_{1,2}^{-1}(c)$ obtaining answer t , and accepts if and only if $\mathcal{Z}(\mathbf{x}[1]||t)[1] = \mathbf{y}[1]$.

It is clear that the distinguisher always accepts in the real world. We now argue that the simulator has negligible chance of making the distinguisher accept.

It is helpful to picture the simulator as knowing the distinguisher's attack. Moreover, we can be generous to the simulator and give both $\mathbf{x}[1]$ and $\mathbf{y}[1]$ to the simulator before requesting the answers a and b from the simulator.

By choosing a and b , the simulator knows which of options 3 and 4 the distinguisher will execute, so the simulator is essentially choosing between these options when it chooses a and b .

Obviously, case 3 is no good for the simulator; moreover, the simulator has no further information on \mathbf{x} and \mathbf{y} besides $\mathbf{x}[1]$ and $\mathbf{y}[1]$, from which it is computationally infeasible, if \mathcal{Z} is a random permutation, to locate a value t such that $\mathcal{Z}(\mathbf{x}[1]||t)[1] = \mathbf{y}[1]$, and which rules out case 4. The simulator is therefore doomed.

4 Combinatorial Definitions

In this section (re-)define the four combinatorial metrics on diffusion permutations mentioned in the introduction, these being *entry-wise randomized preimage resistance* (RPR), *entry-wise randomized collision resistance* (RCR), *conductance* and *all-but-one-conductance*.

Properties are defined unidirectionally: π might satisfy a property while π^{-1} does not.

Given $\pi : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$, a vector $\mathbf{x} \in \{0, 1\}^{wn}$ and indices $j, j' \in [w]$, we let

$$\pi_{j,j'}^{\mathbf{x}} : \{0, 1\}^n \rightarrow \{0, 1\}^n$$

be the function from $\{0, 1\}^n$ to $\{0, 1\}^n$ obtained by restricting the i -th block of input of π , $i \neq j$, to $\mathbf{x}[i]$, by replacing $\mathbf{x}[j]$ with the input $x \in \{0, 1\}^n$, and by considering only the j' -th block of output. (The value $\mathbf{x}[j]$ being, thus, immaterial to $\pi_{j,j'}^{\mathbf{x}}$, since it is replaced by the input.)

ENTRY-WISE RANDOMIZED PREIMAGE RESISTANCE AND ENTRY-WISE RANDOMIZED COLLISION RESISTANCE. The *entry-wise randomized preimage resistance* (RPR) of π is denoted $\text{MaxPreim}(\pi)$, and defined as

$$\text{MaxPreim}(\pi) = \max_{\mathbf{x}, j, h, y} |\{x \in \{0, 1\}^n : \pi_{j,h}^{\mathbf{x}}(x) = y\}|$$

while the *entry-wise randomized collision resistance* (RCR) of π is denoted $\text{MaxColl}(\pi)$, and defined as

$$\text{MaxColl}(\pi) = \max_{\mathbf{x}, \mathbf{x}', j, h} |\{x \in \{0, 1\}^n : \pi_{j,h}^{\mathbf{x}}(x) = \pi_{j,h}^{\mathbf{x}'}(x)\}|$$

where the latter maximum is taken over all tuples $\mathbf{x}, \mathbf{x}', j, h$ such that $\mathbf{x}[j'] \neq \mathbf{x}'[j']$ for some $j' \neq j$. Then by definition

$$\Pr_x[\pi_{j,h}^{\mathbf{x}}(x) = y] \leq \frac{\text{MaxPreim}(\pi)}{2^n}$$

for all $\mathbf{x} \in \{0, 1\}^{wn}$, $y \in \{0, 1\}^n$, and $j, h \in [w]$, where the probability is computed over a uniform choice of $x \in \{0, 1\}^n$, and

$$\Pr_x[\pi_{j,h}^{\mathbf{x}}(x) = \pi_{j,h}^{\mathbf{x}'}(x)] \leq \frac{\text{MaxColl}(\pi)}{2^n}$$

for all $\mathbf{x}, \mathbf{x}' \in \{0, 1\}^{wn}$, $j, h \in [w]$, such that $\mathbf{x}[j'] \neq \mathbf{x}'[j']$ for at least one $j' \neq j$.

Small values of $\text{MaxPreim}(\pi)$ and of $\text{MaxColl}(\pi)$ are better. It is easy to construct permutations with $\text{MaxPreim}(\pi) = 1$ (which is optimal): simply use a linear permutation $\pi : \text{GF}(2^n)^w \rightarrow \text{GF}(2^n)^w$ whose associated matrix (a $w \times w$ matrix with entries in $\text{GF}(2^n)$) has all nonzero entries. Constructing permutations with small values of $\text{MaxColl}(\pi)$ is a bit more involved. In Appendix D we show how to construct permutations that simultaneously achieve small RPR and small RCR (of order $O(w)$ each), and we also show unconditional lower bounds on RCR.

An interesting research direction would be to devise faster-than-linear RPR permutations and faster-than-polynomial RCR permutations. (Our construction of RCR permutations from Appendix D uses finite field operations.)

CONDUCTANCE AND ALL-BUT-ONE CONDUCTANCE. The *conductance* Cond_π of a permutation $\pi : (\{0, 1\}^n)^w \rightarrow (\{0, 1\}^n)^w$ is a function of q defined by

$$\text{Cond}_\pi(q) = \max_{\substack{U_1, \dots, U_w, V_1, \dots, V_w \subseteq \{0, 1\}^n \\ |U_1| = \dots = |V_w| = q}} |\{(\mathbf{x}, \mathbf{y}) : \mathbf{y} = \pi(\mathbf{x}), \mathbf{x}[j] \in U_j, \mathbf{y}[j] \in V_j, 1 \leq j \leq w\}|.$$

One can observe that, like for RPR and RCR, w is a hidden parameter of the conductance (e.g., the conductance of a permutation $\pi : \{0, 1\}^{1024} \rightarrow \{0, 1\}^{1024}$ is not defined on its own without knowing w) and that

$$q \leq \text{Cond}_\pi(q) \leq q^w$$

for any permutation π , $0 \leq q \leq 2^n$. Indeed, $q \leq \text{Cond}_\pi(q)$ since we can always choose q distinct vectors $\mathbf{x}_1, \dots, \mathbf{x}_q \in \{0, 1\}^{wn}$ and set

$$\begin{aligned} U_j &= \{\mathbf{x}_k[j] : 1 \leq k \leq q\} \\ V_j &= \{\mathbf{y}_k[j] : 1 \leq k \leq q\} \end{aligned}$$

where $\mathbf{y}_k = \pi(\mathbf{x}_k)$. On the other hand $\text{Cond}_\pi(q) \leq q^w$ since the number of vectors \mathbf{x} such that $\mathbf{x}[j] \in U_j$ for $1 \leq j \leq w$ is at most $|U_1 \times \dots \times U_w| = q^w$.

The notion all-but-one conductance is essentially the same as conductance, except that one coordinate position in either the input or output is ignored. The all-but-one conductance of a permutation π at q queries is denoted $\text{aboCond}_\pi(q)$.

Formally, given a permutation $\pi : (\{0, 1\}^n)^w \rightarrow (\{0, 1\}^n)^w$, we define the *all-but-one conductance* of π by

$$\text{Cond}_\pi^{h,+}(q) = \max_{\substack{U_1, \dots, U_w, V_1, \dots, V_w \subseteq \{0, 1\}^n \\ |U_1| = \dots = |V_w| = q}} |\{(\mathbf{x}, \mathbf{y}) : \mathbf{y} = \pi(\mathbf{x}), \mathbf{x}[j] \in U_j \forall j \in [w], \mathbf{y}[j] \in V_j \forall j \in [w] \setminus h\}|,$$

$$\text{Cond}_\pi^{h,-}(q) = \max_{\substack{U_1, \dots, U_w, V_1, \dots, V_w \subseteq \{0, 1\}^n \\ |U_1| = \dots = |V_w| = q}} |\{(\mathbf{x}, \mathbf{y}) : \mathbf{y} = \pi(\mathbf{x}), \mathbf{x}[j] \in U_j \forall j \in [w] \setminus h, \mathbf{y}[j] \in V_j \forall j \in [w]\}|$$

$$\text{aboCond}_\pi(q) = \max_{h \in [w]} (\max_{h \in [w]} (\text{Cond}_\pi^{h,+}(\pi, q)), \max_{h \in [w]} (\text{Cond}_\pi^{h,-}(\pi, q)))$$

(Here the first two definitions are for all $h \in [w]$, and we use \forall in postfix notation.) Thus the set V_h is immaterial in the definition of $\text{Cond}_\pi^{h,+}$, while the set U_h is immaterial in the definition of $\text{Cond}_\pi^{h,-}$.

In Appendix B we show that the conductance and all-but-one conductance of a random permutation π are both roughly qwn , which is essentially $q \log(q)$ since q is exponential in n . Constructing explicit permutations with low conductance and low all-but-one conductance remains a central open problem.

5 Network Nomenclature and Main Result

In this section we start with a syntax-oriented description of the confusion-diffusion networks for which our results are obtained. We follow with the formal statement of our main theorem.

While we mentioned six different networks in the introduction (cf. Fig. 1), our full results actually encompass two extra networks, of 6 and 10 rounds (but structurally different from the “other” 6 and 10 round networks); these two extra networks are not mentioned in Fig. 1 because they offer no structural, security or query complexity advantages over other networks. However we include them what follows since they come anyway “for free”, so that our result will subsume eight different networks. (Moreover the proof is easier to write this way, since the eight different networks happen to arise as the result of setting three independent boolean flags.)

NETWORK NOMENCLATURE. A *round* of a confusion-diffusion network refers to a round of S -boxes. More precisely, all S -box permutations $P_{i,j}$ with the same value of i lie in the same *round* of the network.

Since, say, the middle round of our 5-round confusion-diffusion network plays the same structural role (with respect to our simulator) as the middle round in our 9-round network, it makes more sense to designate rounds according to their structural role instead of by their round number (as the latter will keep changing from network to network, even while the structural purpose of the round stays the same).

For this purpose, we replace the array $\mathcal{P} = \{P_{i,j}\}$ of $r \times w$ random permutations with an array \mathcal{Q} of $12 \times w$ random permutations where each “round” (value of the index i) is designated by a different alphabet letter. Specifically, we let

$$\mathcal{Q} = \{F_j, G_j, I_j, D_j, J_j, B_j, A_j, C_j, K_j, E_j, L_j, H_j : j \in [w]\} \quad (1)$$

be a set of $12w$ random permutations, where each permutation is thus indexed by an alphabet letter from the set $\{A, \dots, L\}$ as well as by an index $j \in [w]$.

Having traded the set of indices $\{i : i \in [r]\}$ (the possible round numbers) for the set of letters $\{A, \dots, L\}$, a “round” will henceforth mean a member of the latter set, i.e., a “round” means one of the letters A, \dots, L .

Not all rounds will be used for all confusion-diffusion networks, and we describe which rounds appear in which networks below. (See also Fig. 2 below as an aid.) However two rounds always appear in the same order as they appear listed in \mathcal{Q} above (cf. (1)), when they both appear in a network.

In more detail, our eight different confusion-diffusion networks correspond to the eight different possible settings of three independent boolean flags called `XtraMiddleRnd`, `XtraOuterRnd` and `XtraUntglRnds`. The rounds that appear in each network, as a function of these boolean flags, are as follows:

$$\begin{cases} A & \text{if XtraMiddleRnd is off} \\ B, C & \text{if XtraMiddleRnd is on} \end{cases}$$

$$\begin{cases} G, H & \text{if XtraOuterRnd is off} \\ F, G, H & \text{if XtraOuterRnd is on} \\ D, E & \text{if XtraUntglRnds is off} \\ I, D, J, K, E, L & \text{if XtraUntglRnds is on} \end{cases}$$

As can be seen, toggling either of `XtraMiddleRnd` or `XtraOuterRnd` “costs” one extra round, whereas toggling `XtraUntglRnds` costs four extra rounds. Hence, and because the network has 5 rounds when all flags are off, the number of rounds in the network will be

$$5 + \text{XtraMiddleRnd} + \text{XtraOuterRnd} + 4 \cdot \text{XtraUntglRnds}$$

which spans the integers 5, 6, 6, 7, 9, 10, 10, 11. In Fig. 1 the flag bits appear in the order `XtraMiddleRnd`, `XtraOuterRnd`, `XtraUntglRnds`, hence the two “missing” combinations are those where `XtraMiddleRnd/XtraOuterRnd` are off/on. In addition, the top half of the table is missing all rows with `XtraUntglRnds = true` since toggling this flag does not (significantly) improve security or query complexity when the D -boxes can be arbitrary and, in particular, when the D -boxes can be RCR.

For example, our 11-round network consists of the rounds

$$F, G, I, D, J, B, C, K, E, L, H$$

in this order. (We refer to Fig. 2 further down.) The 10-round network with `XtraMiddleRnd = false` consists of the rounds

$$F, G, I, D, J, A, K, E, L, H$$

in this order as well. All other networks can be obtained by removing rounds from one of these two sequences. In more detail, round F is removed to un-toggle `XtraOuterRnd` and rounds I, J, K, L are removed to un-toggle `XtraUntglRnds`.

We will also rename the generic diffusion permutations $\bar{\pi} = (\pi_1, \dots, \pi_r)$ according to their structural roles in the diffusion network. Specifically, we let

$$\bar{\pi} = (\nu, \pi_G, \pi_I, \pi_J, \pi_B, \tau, \pi_C, \pi_K, \pi_L, \pi_H)$$

where each element in the sequence $\bar{\pi}$ is a permutation from $\{0, 1\}^{wn}$ to $\{0, 1\}^{wn}$. (Thus, we redefine $\bar{\pi}$ to a specific sequence of diffusion permutations.) In the 11-round confusion-diffusion network, diffusion permutations appear interleaved with the S -box rounds in the order

$$F-\nu-G-\pi_G-I-\pi_I-D-\pi_J-J-\pi_B-B-\tau-C-\pi_C-K-\pi_K-E-\pi_L-L-\pi_H-H$$

(i.e., the S -box round consisting of the parallel application of the permutations F_j is followed by the diffusion permutation ν , and so on), whereas in the 10-round network with `XtraMiddleRnd=false` the diffusion permutations appear in the order

$$F-\nu-G-\pi_G-I-\pi_I-D-\pi_J-J-\pi_B-A-\pi_C-K-\pi_K-E-\pi_L-L-\pi_H-H$$

with τ dropped. From either of these configurations one can un-toggle `XtraOuterRnd` by dropping $F-\nu-$ and one can un-toggle `XtraUntglRnds` by dropping $I-\pi_I-$, $-\pi_J-J$, $K-\pi_K-$ and $-\pi_L-L$. For example, our 9-round confusion-diffusion network has the order

$$G-\pi_G-I-\pi_I-D-\pi_J-J-\pi_B-A-\pi_C-K-\pi_K-E-\pi_L-L-\pi_H-H$$

whereas the 5-round and 6-round network with XtraMiddleRnd toggled respectively have order

$$G-\pi_G-D-\pi_B-A-\pi_C-E-\pi_H-H$$

$$G-\pi_G-D-\pi_B-B-\tau-C-\pi_C-E-\pi_H-H$$

and so on.

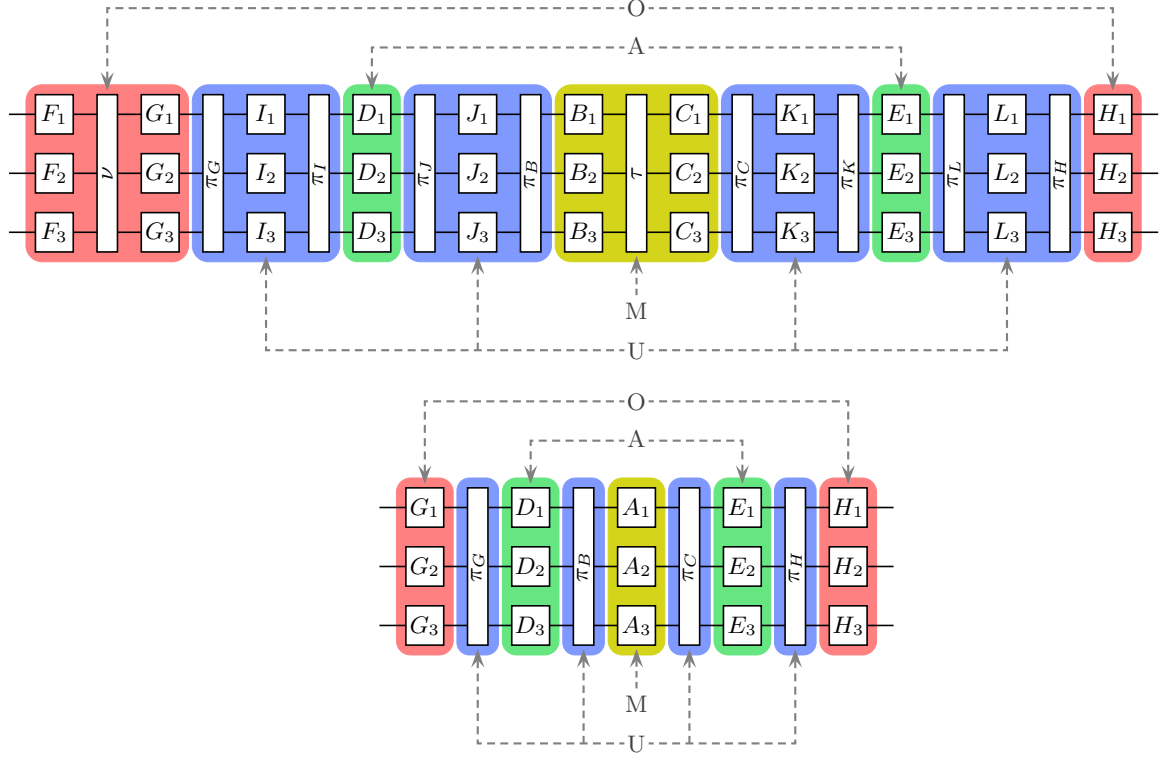


Fig. 2. Emplacement of the outer detect (O), adapt (A), middle detect (M) and untangle (U) zones for the 11- and 5-round simulators. The adapt zones always consist of rounds D and E .

In summary, the confusion-diffusion network under consideration is a function of the confusion permutations \mathcal{Q} , of the diffusion permutation vector $\bar{\pi}$ and of the three boolean flags XtraMiddleRnd, XtraOuterRnd and XtraUntglRnds. For brevity we write this network as

$$P[\mathcal{Q}, \bar{\pi}]$$

keeping the three boolean flags implicit. Depending on the value of the flags some permutations in \mathcal{Q} and/or $\bar{\pi}$ are of course unused. In particular, we assume that unused permutations in \mathcal{Q} are simply ignored for the purpose of the indistinguishability experiment (i.e., these unused permutations are not accessible as oracles).

Semantically, moreover, our simulator divides the confusion rounds and diffusion permutations into nine “zones” of four different types, to wit, one *middle detect* zone (M), left and right *outer detect* zones (O), four *untangle* zones (U) and two *adapt* zones (A). Each zone consists of one

or more contiguous rounds and/or diffusion permutations, with every round and every diffusion permutation belonging to exactly one zone. Fig. 2 shows how the zones appear in the 11- and 5-round networks, while the zoning of other networks can be interpolated from these, given that each of our zones either appears the same as in the 11-round network or the same as in the 5-round network. (For example, if XtraMiddleRnd is off, then round A exists and rounds B and C do not, so the middle detect zone consists of round A only, as in the 5-round simulator; etc.) In particular,

XtraMiddleRnd

determines the aspect of the middle detect zone (and nothing else), while

XtraOuterRnd

determines the aspect of the left outer detect zone (and nothing else), and

XtraUntglRnds

determines the aspect of the untangle zones (and nothing else).

In particular, the D -box τ appears in the middle detect zone if it appears at all, the D -box ν appears in the left outer detect zone if it appears at all, and remaining D -boxes (of the form $\pi\dots$) appear in the untangle zones. Further simulator details, including the purpose of the zones, are given in Section 6.

MAIN RESULT: TAKEAWAY POINTS. Before giving our full technical result (which can be difficult to parse the first time around), we summarize its key implications for security, query complexity, and D -box properties at a high level:

Security. The security of our simulator is essentially a function of the middle detect zone and of the conductance of τ . If XtraMiddleRnd is off, more precisely, so that the middle detect zone consists only of round A , then security is approximately q^{2w}/N where $N = 2^n$, whereas if XtraMiddleRnd is on, the security improves (essentially) to q^2/N , assuming that τ has conductance $\sim q$.

Query complexity. The simulator's query complexity is determined by the left outer detect zone and by the middle detect zone. The query complexity is approximately q^{w^2} if neither of XtraMiddleRnd or XtraOuterRnd is toggled, is approximately q^w if exactly one of these two flags is toggled, and is approximately q if both flags are toggled, where the bounds quoted in the last two cases assume that τ and/or ν both have conductance $\sim q$.

Combinatorial Requirements on the Diffusion Permutations. The diffusion permutations τ and ν must have low conductance and low all-but-one-conductance. (Near q in order to have bounds approaching those of Fig. 1.) The combinatorial requirements on the permutations $\pi\dots$ are different depending on whether XtraUntglRnds is toggled or not:

- if XtraUntglRnds is off the permutations in the untangle zones must be both RPR- and RCR-resistant (or else the above-quoted security bounds are not achieved); in particular, if $w > 2$ then these permutations *cannot* be $\text{GF}(2^n)$ -linear, because, as mentioned above, $\text{GF}(2^n)$ -linear permutations are not RCR-resistant for $w > 2$;

- if XtraUntglRnds is on then the permutations in the untangle zones need only be RPR-resistant; in particular, these permutations can now be $\text{GF}(2^n)$ -linear even for $w > 2$.

One can also summarize the situation orthogonally, according to which boolean flags have which effect. Namely:

- toggling `XtraMiddleRnd` impacts security and query complexity
- toggling `XtraOuterRnd` impacts query complexity
- toggling `XtraUntglRnds` reduces the cryptographic requirements on the permutations $\pi\dots$

It is natural to wonder whether even further extension of the middle, outer left, or untangle zones can be beneficial. The short answer to this question, that we consider in detail (but without proof) in Appendix A, is yes; namely, by adding more rounds to these zones one can further reduce the combinatorial requirements on the diffusion permutations, leading, e.g., to networks in which the permutations in the untangle zones are not even RPR-resistant. (!) We refer to Appendix A for more details.

As a final high-level comment, a potentially striking feature of our simulator is the left-right asymmetry that arises when `XtraOuterRnd` is toggled. As explained in Section 6, adding an extra round to the right-hand outer detect zone (assuming `XtraOuterRnd` is already set) does not further decrease the simulator's query complexity, nor improve security. Nonetheless, such an extra zone can be used to reduce the simulator's *space complexity* while maintaining essentially the same security and query complexity. (We refer the reader to the discussion after Lemma 52 without going into more details here.) Hence, while the benefit of adding an additional round to the right outer detect zone is somewhat technical, such an extra round does have a potential justification in terms of simulator design.

MAIN RESULT. We recall that $\bar{\pi} = (\nu, \pi_G, \pi_I, \pi_J, \pi_B, \tau, \pi_C, \pi_K, \pi_L, \pi_H)$ is the vector of diffusion permutations, not necessarily all of which are used. In order to more succinctly state the main result, we define

$$\begin{aligned} \text{MaxColl}(\bar{\pi}) &= \max(\text{MaxColl}(\pi_G), \text{MaxColl}(\pi_B^{-1}), \text{MaxColl}(\pi_C), \text{MaxColl}(\pi_H^{-1})) \\ \text{MaxPreim}(\bar{\pi}) &= \max(\text{MaxPreim}(\pi_G), \text{MaxPreim}(\pi_B^{-1}), \text{MaxPreim}(\pi_C), \text{MaxPreim}(\pi_H^{-1}), \\ &\quad \text{MaxPreim}(\pi_I), \text{MaxPreim}(\pi_J^{-1}), \text{MaxPreim}(\pi_K), \text{MaxPreim}(\pi_L^{-1})) \\ \text{MaxCoPr}(\bar{\pi}) &= \max(\text{MaxColl}(\bar{\pi}), \text{MaxPreim}(\bar{\pi})) \end{aligned}$$

where π^{-1} denotes the inverse of π .

Moreover we define

$$\begin{aligned} \alpha(q) &= \begin{cases} (2q)^w & \text{if XtraMiddleRnd is off,} \\ \text{Cond}_\tau(2q) & \text{if XtraMiddleRnd is on,} \end{cases} \\ \beta(q) &= \begin{cases} (q + \alpha(q))^w & \text{if XtraOuterRnd is off,} \\ \text{Cond}_\nu(q + \alpha(q)) & \text{if XtraOuterRnd is on.} \end{cases} \end{aligned}$$

The definitions of $\alpha(q)$ and $\beta(q)$ might seem annoyingly technical right now. In Section 6 we provide more digestible semantic explanations for $\alpha(q)$ and $\beta(q)$.

Theorem 1. *Let $N = 2^n$. The confusion-diffusion network $P[\mathcal{Q}, \bar{\pi}]$ achieves (t_S, q_S, ε) -indifferentiability from a random permutation $\mathcal{Z} : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$ for ε equal to*

$$\frac{\beta(q)(q + \alpha(q))^w}{N^w - q - \alpha(q)} + \frac{1}{N^w} + \frac{4w(q + \alpha(q))^2}{N - q - \alpha(q)}$$

$$\begin{aligned}
& + \frac{4wq \text{aboCond}_\tau(2q)}{N - 2q} && \text{if XtraMiddleRnd is on} \\
& + \frac{2w(q + \alpha(q)) \text{aboCond}_\nu(q + \alpha(q))}{N - q - \alpha(q)} && \text{if XtraOuterRnd is on} \\
& + \frac{4w\alpha(q)(q + \alpha(q)) \text{MaxCoPr}(\bar{\pi})}{N - q - \alpha(q)} && \text{if XtraUntglRnds is off} \\
& + \frac{6w(q + \alpha(q))^2 \text{MaxPreim}(\bar{\pi})}{N - q - \alpha(q)} && \text{if XtraUntglRnds is on}
\end{aligned}$$

and for $q_S = \beta(q)$, $t_S = O(w(q + \alpha(q))^w)$. Here $q = q_0(1 + rw)$ where q_0 is the number of distinguisher queries and $r \in \{5, 6, 7, 9, 10, 11\}$ is the number of rounds in the confusion-diffusion network.

The full formal proof of this result is given in Appendix E.

Parsing the security bound. In order to get a rough feel for the security bound of Theorem 1 it is helpful to make the order-of-magnitude approximations

$$\begin{aligned}
\text{MaxPreim}(\bar{\pi}) &= \text{MaxColl}(\bar{\pi}) \approx O(1) \\
\text{Cond}_\tau(2q) &= \text{aboCond}_\tau(2q) \approx q \\
\text{Cond}_\nu(q + \alpha(q)) &= \text{aboCond}_\nu(q + \alpha(q)) \approx \alpha(q).
\end{aligned}$$

With these approximations in place, and given $q \ll N$ (in fact we can assume $q \leq N^{1/2}$, since the security bound is void otherwise) it easy to verify that the largest terms in Theorem 1 are of the order

$$\frac{\alpha(q)^2}{N}$$

and which is, therefore, a first approximation to the security ε that we achieve. Since

$$\alpha(q) = \begin{cases} (2q)^w & \text{if XtraMiddleRnd is off,} \\ q & \text{if XtraMiddleRnd is on,} \end{cases}$$

under the order-of-magnitude approximations given above, we find

$$\varepsilon \approx \begin{cases} (2q)^{2w}/N & \text{if XtraMiddleRnd is off,} \\ q^2/N & \text{if XtraMiddleRnd is on} \end{cases}$$

for the security, to a first approximation. On the other hand we find

$$q_S = \beta(q) \approx \begin{cases} (2q)^{w^2} & \text{if XtraMiddleRnd/XtraOuterRnd are off/off} \\ (2q)^w & \text{if XtraMiddleRnd/XtraOuterRnd are off/on} \\ (2q)^w & \text{if XtraMiddleRnd/XtraOuterRnd are on/off} \\ q & \text{if XtraMiddleRnd/XtraOuterRnd are on/on} \end{cases}$$

for the query complexity, again to a first approximation.

Digging a little deeper into lower-order factors, if we let

$$\mu(\bar{\pi}) = \begin{cases} \text{MaxCoPr}(\bar{\pi}) & \text{if XtraUntglRnds is off} \\ \text{MaxPreim}(\bar{\pi}) & \text{if XtraUntglRnds is on} \end{cases}$$

then it is easy to verify that

$$\varepsilon = \begin{cases} O(w)\alpha(q)^2\mu(\bar{\pi})/N & \text{if XtraOuterRnd is off,} \\ O(w)\alpha(q)\text{aboCond}_\nu(q + \alpha(q))/N & \text{if XtraOuterRnd is on} \end{cases} \quad (2)$$

by keeping only the biggest term(s) in the security bound and by folding lower-order terms—as well as factors of the type $N/(N - q - \alpha(q))$, which are dominated by constants—into the big-Oh constant. The second line of (2) is typically larger than the first, and more substantially larger than the first if the all-but-one-conductance of ν is poor; indeed, security actually *degrades* when XtraUntglRnds is toggled, even as query complexity improves. (This degradation is plainly visible in the second half of Fig. 1, due to the bigger values of all-but-one-conductance for $\text{GF}(2^n)$ -linear permutations.)

As already mentioned, Appendix D shows explicit constructions of diffusion permutations with low values of MaxColl, and also unconditional lower bounds on MaxColl. More precisely that appendix shows that

$$\mu(\bar{\pi}) = \begin{cases} O(w) & \text{if XtraUntglRnds is off,} \\ 1 & \text{if XtraUntglRnds is on,} \end{cases}$$

in the best case, with these bounds being matched by explicit constructions.

On the other hand, it is relatively easy to see that $\text{MaxColl}(\pi) = 2^n = N$ for any linear permutation $\pi : \text{GF}(2^n)^w \rightarrow \text{GF}(2^n)^w$ as long as $w > 2$. In this case $\text{MaxCoPr}(\pi) = N$, and Theorem 1 becomes void if XtraUntglRnds is off. Thus one of the main reasons for toggling XtraUntglRnds might be to enable the use of linear diffusion permutations, or any other (fast) family of permutations that have small entry-wise randomized preimage resistance (but potentially large entry-wise randomized collision resistance).

EXTENSIONS. In Appendix A we discuss the extension of our main result to even longer networks. In particular, we show that each of our four combinatorial properties are specialized (extreme) cases of more general adversarial games involving entire segments of CD networks. We give a more general theorem (Theorem 2; albeit without proof) in terms of these properties defined on network segments. The use of more rounds, then, enables the D -boxes to become even weaker (provably so in the case of the untangle rounds, conjecturally so for the detect zones), and thus even faster. This leaves many exciting possibilities/questions for future work, as explained in Appendix A.

6 Simulator Overview

CONTEXT. We start with some very high-level description and reminder-of-purpose of our simulator. For this discussion it will be more convenient if we momentarily revert to indexing the S -boxes by coordinate pairs (i, j) where $i \in [r]$ the round number and $j \in [w]$ the layer number, with r being the number of rounds and w being the width. The diffusion permutation between the i -th and $(i + 1)$ -th rounds will again be denoted π_i as well.

The simulator is responsible for answering queries to the S -boxes, and has access to a random permutation oracle $\mathcal{Z} : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$ that is being independently accessed by the distinguisher. The simulator’s job is to keep the S -box answers compatible with \mathcal{Z} in the sense that it looks to the distinguisher as if \mathcal{Z} is implemented by the confusion-diffusion network.

For each pair $(i, j) \in [r] \times [w]$ the simulator maintains a pair of tables $P_{i,j}$ and $P_{i,j}^{-1}$, each containing 2^n entries of n bits each, in which the simulator keeps a record of “what it has already

decided” about the (i, j) -th S -box. Initially the tables are blank, meaning that $P_{i,j}(x) = P_{i,j}^{-1}(y) = \perp$ for all $x, y \in \{0, 1\}^n$. The simulator sets $P_{i,j}(x) = y$, $P_{i,j}^{-1}(y) = x$ to indicate that the (i, j) -th S -box maps x to y . The simulator never overwrites values in $P_{i,j}$ or in $P_{i,j}^{-1}$ and always keeps these two tables consistent. Hence $P_{i,j}$ encodes a partial matching (or “partial permutation”) from $\{0, 1\}^n$ to $\{0, 1\}^n$ from which edges are never subtracted. We also note that the edges in $P_{i,j}$ are a *superset* of those queries that the distinguisher has made to the (i, j) -th S -box or to its inverse (i.e., $P_{i,j}$ contains the answers to those queries, and possibly more).

By analogy with the notation of Section 2 we write

$$P_i(\mathbf{x}) = \mathbf{y} \tag{3}$$

if $\mathbf{x}, \mathbf{y} \in \{0, 1\}^{wn}$ are vectors such that $P_{i,j}(\mathbf{x}[j]) = \mathbf{y}[j]$ for all $j \in [w]$. Note that (3) is a time-dependent statement, in the sense that the tables $P_{i,j}$ keep accruing entries as the distinguishing experiment proceeds. For example, (3) is initially false for all i and all vectors \mathbf{x}, \mathbf{y} . Moreover P_i is not an actual table maintained by the simulator—i.e., (3) is “just notation”.

A sequence of vectors $(\mathbf{x}^1, \mathbf{y}^1, \dots, \mathbf{x}^r, \mathbf{y}^r)$ is called a *completed path*⁶ if $P_i(\mathbf{x}^i) = \mathbf{y}^i$ for $i = 1, \dots, r$ and if $\pi_i(\mathbf{y}^i) = \mathbf{x}^{i+1}$ for $i = 1, \dots, r - 1$. The set of completed paths is also time-dependent. The vectors \mathbf{x}^1 and \mathbf{y}^r are called the *endpoints* of the path.

We informally say that the distinguisher *completes a path* if it makes queries to the simulator that form a completed path. (There are many different possible ways to order such a set of queries, obviously.) One can picture the distinguisher as trying to complete paths in various devious ways (typically, reusing the same queries as part of different paths), and checking that the path endpoints are each time compatible with \mathcal{Z} .

The simulator’s job, in response, is to run ahead of the distinguisher and pre-emptively complete paths that it thinks the distinguisher is interested in, such as to make these paths compatible with \mathcal{Z} . The simulator’s dilemma is that it must choose under which conditions to complete a path; if it waits too long, or completes paths in only highly specialized cases, it may find itself trapped in a contradiction (typically, while trying to complete several paths at once); but if it is too trigger-happy, having a very large number of conditions under which it will choose to complete a path, the simulator runs the risk creating⁷ an out-of-control chain reaction of path completions.

Essentially the simulator must be safe, but in a smart enough way that it avoids (out-of-control) chain reactions. We will informally refer to the problem of showing that no out-of-control chain reactions occur as the problem of *simulator termination*.

SIMULATOR ZONES. As already mentioned our simulator divides the confusion rounds and diffusion permutations into nine zones of four different types, shown in Fig. 2 for the 11- and 5-round networks.

Specifically, the “middle zone” is

$$\begin{cases} A & \text{if XtraMiddleRnd is off} \\ B\text{-}\tau\text{-}C & \text{if XtraMiddleRnd is on} \end{cases}$$

⁶ This definition, made for the sake of expository convenience, is superseded further down, where we redefine “completed path” by adding the requirement that the endpoints be compatible with \mathcal{Z} , i.e., that $\mathcal{Z}(\mathbf{x}^1) = \mathbf{y}^r$.

⁷ Indeed, the simulator makes no distinction between those entries in its tables $P_{i,j}$ that are the direct result of an distinguisher query, and those which it created on its own while pre-emptively completing paths. It seems very hard to leverage such a distinction. Note for example that the distinguisher may know values in $P_{i,j}$ without having made the relevant queries, simply by virtue of knowing how the simulator works.

depending only on `XtraMiddleRnd`; the left and right outer detect zones are

$$\begin{cases} G, H & \text{if XtraOuterRnd is off} \\ F^{-\nu}G, H & \text{if XtraOuterRnd is on} \end{cases}$$

depending only on `XtraOuterRnd`; the four “untangle zones” are

$$\begin{cases} \pi_G, \pi_B, \pi_C, \pi_H & \text{if XtraUntglRnds is off} \\ \pi_{G-I}-\pi_I, \pi_{J-J}-\pi_B, \pi_{C-K}-\pi_K, \pi_{L-L}-\pi_H & \text{if XtraUntglRnds is on} \end{cases}$$

depending only on `XtraUntglRnds`; the two “adapt” zones, finally, are rounds D and E in all networks.

We observe that zones (specifically, untangle zones with `XtraUntglRnds = false`) might consist solely of diffusion permutations. This contrasts with semantically similar zoning systems for Feistel network simulators and key-alternating simulators [18, 19, 27] for which the zones always contain at least one ideal component. We also observe that in the minimalist 5-round simulator, each round and each diffusion permutation corresponds to an individual zone.

TABLE NOTATION, COLUMNS AND MATCHING QUERY SETS. We revert to identifying rounds with letters in $\{A, \dots, L\}$. Under this notation tables $P_{i,j}, P_{i,j}^{-1}$ described above become, e.g., tables A_j and A_j^{-1} . Thus for each round $T \in \{A, \dots, L\}$ that is “operational” (as will depend on the flag settings) the simulator maintains tables T_j, T_j^{-1} for $1 \leq j \leq w$, as already described above under the notation $P_{i,j}, P_{i,j}^{-1}$.

We write $T(\mathbf{x}) = \mathbf{y}$ if $T_j(\mathbf{x}[j]) = \mathbf{y}[j]$ for each $j \in [w]$, and for all $T \in \{A, \dots, L\}$. We also write $T(\mathbf{x}) = \perp$ if $T_j(\mathbf{x}[j]) = \perp$ for at least one value of j . The notation $T^{-1}(\mathbf{y})$ is analogous, with T^{-1} being the inverse of T . As emphasized above this notation is really “notation only” in the sense that the simulator does not maintain such things as tables T or T^{-1} .

A non-null entry in table T_j will be called an (S -box) *query*. More formally, an S -box query is a quadruple (T, j, x, y) where $T \in \{A, \dots, L\}$, $j \in [w]$, $x, y \in \{0, 1\}^n$, such that $T_j(x) = y$.

A set of w queries with the same value of T but with different values of j will be called a *column* or a T -*column* when we wish to emphasize the round. The (unique) vectors \mathbf{x}, \mathbf{y} such that $(T, j, \mathbf{x}[j], \mathbf{y}[j])$ is a query in the column for each $j \in [w]$ are called the *input* and *output* of the column respectively. We note that a column is uniquely determined by either its input or output.

Two columns are *adjacent* if their rounds are adjacent. (E.g., a B -column and a C -column are adjacent.) Two adjacent columns are *matching* if $\pi(\mathbf{y}) = \mathbf{x}$, where \mathbf{y} is the output of the first column, where \mathbf{x} the input of the second column, and where π is the diffusion permutation between the two rounds.

A pair of columns from the first and last round of the confusion-diffusion network are likewise *matching* if $\mathcal{Z}(\mathbf{x}) = \mathbf{y}$, where \mathbf{x} is the input to the first-round column (either an F - or G -column) and \mathbf{y} is the output of the last-round column (the H -column).

The notion of matching columns naturally extends to sequences of columns from consecutive rounds of length greater than two. (The first and last round of the network are always considered adjacent.) If a set of matching columns encompasses all rounds we call the set a *completed path*. Thus, since the first and last column of a network are considered adjacent, completed paths are compatible with \mathcal{Z} by definition.

The set of queries in a matching sequence of columns of any length is called a *matching set* of queries. We will also refer to the queries of a single column as a matching set, considering that column as a matching sequence of columns of length 1.

One can also observe that two different completed paths cannot contain the same column. Indeed, each D -box is a permutation and each round implements a partial permutation as well.

SIMULATOR OPERATION AND TERMINATION ARGUMENT. Our simulator applies a design paradigm pioneered by Yannick Seurin [27] that has also been used in other simulators for Feistel networks or key-alternating ciphers [18, 19]. As for all such simulators, our simulator completes two types of paths, where one type of path is triggered by a “middle detect zone” and the other is triggered by an “outer detect zone”.

In more detail, our simulator pre-emptively completes a path⁸ for every matching set of queries in the middle detect zone (such a matching set will consist of either one or two columns, depending on whether `XtraMiddleRnd` is toggled) and also for every matching set of queries in the two outer detect zones, considered as a single consecutive set of columns (the latter kind of matching set will thus consist of either two or three columns, depending on whether `XtraOuterRnd` is toggled).

Crucially, one can show that, unless some bad event of negligible probability occurs, each outer-triggered path completion is associated to a unique pre-existing query made by the distinguisher to \mathcal{Z} . Since the distinguisher makes only q queries in total, this means that at most q outer-triggered path completions occur, with high probability. In fact our simulator aborts if it counts⁹ more than q outer-triggered path completions, so in our case at most q outer-triggered path completions occur with probability 1.

Moreover, a middle-triggered path completion does not add any new queries to the middle detect zone. This means that all queries in the middle detect zone can be chalked up to one of two causes: (1) queries directly made by the distinguisher, (2) queries made by the simulator during outer-triggered path completions.

Cause (1) accounts for at most q queries to each S -box and, as just seen, cause (2) accounts for at most q queries as well. Hence a middle detect zone S -box is never queried at more than $2q$ points, i.e., the table T_j for $T \in \{A, B, C\}$ never has more than $2q$ non- \perp entries.

In particular, if `XtraMiddleRnd` is off, the latter implies that no more than $(2q)^w$ middle-triggered path completions occur, or one for every possible combination of an entry from each of the tables A_1, \dots, A_w . If `XtraMiddleRnd` is on, on the other hand, than at most $\text{Cond}_\tau(2q)$ middle-triggered path completions occur, as is easy to see per the definition¹⁰ of conductance. In other words, $\alpha(q)$ (cf. Section 5) is an upper bound on the number of middle-triggered path completions. In fact, $\alpha(q)$ is an upper bound for the *total* number of path completions performed (including also outer path

⁸ The phrase “completes a path” is informal at this point, as there are generally many different ways to complete a path. (E.g., where to “adapt” a path to make it compatible with \mathcal{Z} , etc.) More details follow below.

⁹ This means the simulator knows the value of q beforehand, which introduces a small amount of non-uniformity into the simulator. One could remove this non-uniformity—i.e., not tell the simulator the value of q beforehand—at the cost of a more complicated theorem statement and proof. But the fact that essentially all security games allow adversaries that know the value of q for which they want to carry out an attack makes the issue a bit moot.

¹⁰ More precisely, let $U_j = \{y \in \{0, 1\}^n : B_j^{-1}(y) \neq \perp\}$, let $V_j = \{x \in \{0, 1\}^n : C_j(x) \neq \perp\}$ at the end of the distinguishing experiment. Then $|U_j|, |V_j| \leq 2q$ for each j , and the number of middle-triggered path completions that have occurred is at most

$$\{(\mathbf{x}, \mathbf{y}) : \tau(\mathbf{x}) = \mathbf{y}, \mathbf{x} \in \prod_j U_j, \mathbf{y} \in \prod_j V_j\} \leq \text{Cond}_\tau(2q).$$

completions), since each completed path is also associated to a unique set of matching queries from the middle detect zone.

As for all S -boxes outside the middle detect zone, their queries can also be chalked up to one of two sources, namely direct distinguisher queries and path completions. There are at most q direct distinguisher queries, and each completed path contributes at most 1 query to each S -box, so each S -box outside the middle detect zone ends up with at most $q + \alpha(q)$ queries.

The simulator, moreover, only queries \mathcal{Z} in order to either complete paths or else in order to detect outer-triggered path completions. This implies the number of distinct simulator queries to \mathcal{Z} is upper bounded by the number of matching sets of queries in the left-hand outer detect zone. Indeed each completed path is obviously associated to a matching set of queries in the left-hand outer detect zone; and for the purpose of outer-triggered path detection, it is easy to see that the simulator only needs to query \mathcal{Z} at most once for each such matching set as well by maintaining a table¹¹ of queries already made to \mathcal{Z} . If `XtraOuterRnd` is off, thus, the number of simulator queries to \mathcal{Z} will be at most $(q + \alpha(q))^w$; whereas if `XtraOuterRnd` is on, the same number will be at most $\text{Cond}_\nu(q + \alpha(q))$. The simulator query complexity is thus at most $\beta(q)$ (cf. Section 5).

PROCEDURE NAMES AND ALIASES. The procedures implementing the oracles maintained by the simulator are denoted using uppercase plaintext letters, where the letter matches the round of the S -box in question. The first argument to the oracle denotes the index of the S -box within the round. Thus $A(j, x)$ denotes a query to the j -th S -box in round A in input $x \in \{0, 1\}^n$. Likewise $A^{-1}(j, y)$ denotes a query to the inverse of the j -th S -box in round A on input $y \in \{0, 1\}^n$, and so on. These are the oracles available to the distinguisher, ranging over all rounds that are in use as determined by the boolean flags.

Moreover, it will be convenient to have a single name, e.g., for “the last round of the middle detect zone” or “the first round of the left-hand outer detect zone” irrespective of the boolean flag values. For this we implement the following three “round aliases” U , V and X :

- U denotes A if `XtraMiddleRnd` is off, denotes B otherwise
- V denotes A if `XtraMiddleRnd` is off, denotes C otherwise
- X denotes G if `XtraOuterRnd` is off, denotes F otherwise

For example, the table U_j denotes the table A_j if `XtraMiddleRnd` is off, the table B_j otherwise, and so on.

FURTHER DETAILS: QUERY HANDLING. S -box queries from the adversary to the simulator, and from the simulator to itself, are handled differently according to whether or not they occur in a position and in a direction that might trigger a path completion. Queries outside the detect zones—i.e., queries in the adapt or untangle zones—cannot trigger path completions and are answered by lazy permutation sampling, if they are not already present in the relevant table T_j . Likewise, new queries to B , C^{-1} , G^{-1} , H as well as to F and F^{-1} are unlikely to trigger new path completions and are answered by lazy permutation sampling. (For example, a new query to C^{-1} is unlikely to form a new pair of matching B and C columns together with some other $2w - 1$ pre-existing queries. The degree to which this is unlikely, however, depends on the all-but-one conductance of τ , which is precisely how all-but-one conductance enters into the proof. As another example, a new query to F^{-1} or H is unlikely to trigger a new outer path completion, because it is unlikely that the random answer to this query will “meet up” with a previous query to \mathcal{Z} .)

¹¹ Thus, in particular, adding an extra round to the right-hand detect zone would not further reduce the query complexity, since the query complexity is only proportional to the number of matching query sets for the left-hand outer detect zone.

All other queries—namely queries to A , A^{-1} , B^{-1} , C , G and H^{-1} —can reasonably be expected to result in new path completions. Instead of being lazy-sampled “on the spot”, these queries are handled by a system of *requests*. In more detail, when the simulator fields a *request* to A , A^{-1} , B^{-1} , C , G or H^{-1} , the simulator makes note of the request and determines what path completions the request will trigger. Instead of answering the request, the simulator proceeds to begin these path completions in “the other direction” (i.e., the direction for which the answer to the current request isn’t needed); these tentative path completions eventually engender new *requests* of their own, which engender new path completions of their own, and so on. The recursive process eventually stops with a list of pending requests and a list of pending path completions. At this point (and we give more details below) the simulator simultaneously lazy samples all the pending requests, and simultaneously completes all the pending path completions.

Crucially for this idea, requests to U^{-1} (i.e., requests to A^{-1} or B^{-1} , depending) can only give rise to new requests to G , as is easy to see¹², and likewise requests to G can only engender new requests to U^{-1} . Symmetrically, requests to V only engender new requests to H^{-1} and vice-versa. Thus, there exist only two possible “closed systems” of recursive request calls: the G/U^{-1} system and the V/H^{-1} system. At most one of the two possible systems is triggered by any given distinguisher query. (To be precise, a G or U^{-1} query will trigger the G/U^{-1} system; a V or H^{-1} query will trigger the V/H^{-1} system; other queries trigger neither system.)

UNFOLDING OF RECURSIVE CALLS. We now describe in more detail how a recursive system of G/U^{-1} requests unfolds. The simulator maintains sets $\text{ToBeAssigned}_{G_j}$ and $\text{ToBeAssigned}_{U_j}^-$ for $1 \leq j \leq w$ that hold the requests to G and U^{-1} respectively. At the beginning of the query cycle (a *query cycle* is the portion of game execution between when the distinguisher makes a query and when that query is answered) these sets are empty.

Assume, say, that the system of recursive calls starts with a distinguisher query $G(j, x)$. At this point the simulator *requests* the value $G(j, x)$ by an internal call $\text{RequestG}(j, x)$. If $G_j(x) \neq \perp$ the request has no effect and the call returns with no value.

Otherwise, if $G_j(x) = \perp$, the simulator adds the value x to the set $\text{ToBeAssigned}_{G_j}$. The simulator then enumerates outer path completions, if any, that will be triggered by the new query (G, j, x, \cdot) . In more detail, the simulator enumerates all vectors \mathbf{x}^G such that $\mathbf{x}^G[j] = x$ and such that either $G_{j'}(\mathbf{x}^G[j']) \neq \perp$ or $\mathbf{x}^G[j'] \in \text{ToBeAssigned}_{G_{j'}}$ for all $j' \neq j$. If XtraOuterRnd is off the simulator defines $\mathbf{y}^H = \mathcal{Z}(\mathbf{x}^G)$ and checks that $H^{-1}(\mathbf{y}^H) \neq \perp$; if XtraOuterRnd is on, the simulator checks that $\perp \neq \mathbf{x}^F := F^{-1}(\nu^-(\mathbf{x}^G))$, evaluates $\mathbf{y}^H = \mathcal{Z}(\mathbf{x}^F)$, and checks again that $H^{-1}(\mathbf{y}^H) \neq \perp$. (As soon as the first check fails, the simulator moves on to the next candidate for \mathbf{x}^G . In particular, the simulator does not define any new queries via lazy sampling at this stage.)

For each \mathbf{x}^G for which all of the above checks succeed, the simulator starts completing a path backward starting from the vector $\mathbf{x}^H = H^{-1}(\mathbf{y}^H)$. For this the simulator fills in missing queries to the path (working backward through the rounds) by lazy sampling until it reaches round U ; here, instead of filling in the missing queries, it calls $\text{RequestU}^{-1}(j, \mathbf{y}^U[j])$ for $1 \leq j \leq w$, where \mathbf{y}^U is the input to U^{-1} for the path in question. It also adds the pair $(\mathbf{x}^G, \mathbf{y}^U)$ to a list of such pairs named ToBeAdapted_D (initially empty at the start of the query cycle).

A description of how requests to U^{-1} are handled will finish the description of the recursive procedure. Firstly, as for G -requests, a U^{-1} -request (j, y) returns immediately if $U_j^{-1}(y) \neq \perp$

¹² The path completion(s) triggered by the request to U^{-1} will “pass through” rounds H and F (if present) without triggering new path completions, due to the fact, argued above, that (forward) queries to H and F are unlikely to cause new outer-triggered path completions.

of if $y \in \text{ToBeAssigned}_{U_j}^-$ already. Otherwise, the value y is added to $\text{ToBeAssigned}_{U_j}^-$ and the simulator enumerates all middle path completions, if any, that will be triggered by the new request. In more detail the simulator enumerates all vectors \mathbf{y}^U such that $\mathbf{y}^U[j] = y$ and such that either $U_j(\mathbf{y}^U[j']) \neq \perp$ or $\mathbf{y}^U[j'] \in \text{ToBeAssigned}_{U_{j'}}^-$ for all $j \neq j'$; moreover if `XtraMiddleRnd` is on the simulator checks that $\perp \neq \mathbf{y}^V := C(\tau(\mathbf{y}^U))$; otherwise (if `XtraMiddleRnd` is off) we have $\mathbf{y}^V = \mathbf{y}^U$.

For each such \mathbf{y}^U for which \mathbf{y}^V is defined, the simulator starts completing a path forward from \mathbf{y}^V . The simulator fills missing queries to the path by lazy sampling until it reaches round G (after calling \mathcal{Z}^{-1}). At this point the simulator adds the pair $(\mathbf{x}^G, \mathbf{y}^U)$ for the corresponding \mathbf{x}^G to the list ToBeAdapted_D and calls $\text{RequestG}(j, \mathbf{x}^G[j])$ for all $1 \leq j \leq w$, so the recursion continues.

When the recursion terminates control reverts to the original procedure $G(j, x)$ (queried by the distinguisher) that placed the first request. So far, no requests have been answered, though all requests have been recorded, and all partially-built, yet-to-be-completed paths have been recorded as well by a pair $(\mathbf{x}^G, \mathbf{y}^U)$ in the set ToBeAdapted_D .

At this point $G(j, x)$ will call a procedure AdaptLeft which simultaneously samples all the pending requests and completes all pending paths, “adapting” each such path at round D . This procedure is discussed below.

A recursive system of G/U^{-1} calls can also be triggered by a distinguisher query to U^{-1} (A^{-1} or B^{-1} depending), but this case is exactly analogous. Moreover a recursive system of V/H^{-1} calls is analogous as well.

PATH COMPLETION AND THE PURPOSE OF THE UNTANGLE ZONES. When AdaptLeft is called it has the sets $\text{ToBeAssigned}_{G_j}$ and $\text{ToBeAssigned}_{U_j}^-$ for $1 \leq j \leq w$, as well as the list ToBeAdapted_D consisting of pairs of the form $(\mathbf{x}^G, \mathbf{y}^U)$ for paths to be completed. We will assume that ToBeAdapted_D has k elements, and write the pairs in ToBeAdapted_D as

$$(\mathbf{x}_1^G, \mathbf{y}_1^U), \dots, (\mathbf{x}_k^G, \mathbf{y}_k^U).$$

It is easy to see that $G_j(x) = \perp$ and $U_j^{-1}(y) = \perp$ for each $x \in \text{ToBeAssigned}_{G_j}$ and each $y \in \text{ToBeAssigned}_{U_j}^-$. Moreover, and as we argue in the proof, one can show that for each pair $(\mathbf{x}_i^G, \mathbf{y}_i^U)$ there is some $j \in [w]$ such that $\mathbf{x}_i^G[j] \in \text{ToBeAssigned}_{G_j}$ and some $h \in [w]$ such that $\mathbf{y}_i^U[h] \in \text{ToBeAssigned}_{U_h}^-$. Moreover, one can also show that $\mathbf{x}_i^G \neq \mathbf{x}_{i'}^G, \mathbf{y}_i^U \neq \mathbf{y}_{i'}^U$ for all $i \neq i'$.

AdaptLeft consists of three stages:

(1) Values of $G_j(x), U_j^{-1}(y)$ are (finally) lazy sampled for each $x \in \text{ToBeAssigned}_{G_j}, y \in \text{ToBeAssigned}_{U_j}^-$, for each $j \in [w]$, whence vectors $\mathbf{y}_i^G := G(\mathbf{x}_i^G), \mathbf{x}_i^U := U^{-1}(\mathbf{y}_i^U)$ become defined for $1 \leq i \leq k$.

(2) The vector \mathbf{y}_i^G is “evaluated forward” to a vector \mathbf{x}_i^D that is an input to round D for $1 \leq i \leq k$, while the vector \mathbf{x}_i^U is “evaluated backward” to a vector \mathbf{y}_i^D that is an output to round D (in a little more detail, if `XtraUntglRnds` is off, we have $\mathbf{x}_i^D = \pi_G(\mathbf{y}_i^G), \mathbf{y}_i^D = \pi_B^-(\mathbf{x}_i^U)$; if `XtraUntglRnds` is on, lazy sampling of rounds I and J is required to define $\mathbf{x}_i^D, \mathbf{y}_i^D$).

(3) AdaptLeft (tries to) set $D_j(\mathbf{x}_i^D[j]) = \mathbf{y}_i^D[j]$ for each $i \in [k]$ and each $j \in [w]$.

The only step that can run into problems is step (3). Namely, the step (3) fails if either (i) there exists values $i \in [k], j \in [w]$ such that either $D_j(\mathbf{x}_i^D[j]) \neq \perp$ or $D_j^{-1}(\mathbf{y}_i^D[j]) \neq \perp$ before the start of step (3), or (ii) there exists values $i, i' \in [k], i \neq i', j \in [w]$, such that either $\mathbf{x}_i^D[j] = \mathbf{x}_{i'}^D[j]$ or $\mathbf{y}_i^D[j] = \mathbf{y}_{i'}^D[j]$ before the start of step (3). The purpose of the untangle zones is to make these two types of bad events unlikely.

If XtraUntglRnds is off, a bit of thought reveals that small values of $\text{MaxPreim}(\pi_G)$, $\text{MaxPreim}(\pi_C^-)$ are sufficient and necessary to ensure that (i) only occurs with low probability whereas small values of $\text{MaxColl}(\pi_G)$, $\text{MaxColl}(\pi_C^-)$ are sufficient and necessary to ensure (ii) only occurs with low probability. The argument is simple (involving a union bound) but the details are notationally tedious, and we refer to the proof of Lemma 38 in Section E for the details.

If XtraUntglRnds is on, a little bit more thought¹³ reveals that small values of $\text{MaxPreim}(\pi_G)$, $\text{MaxPreim}(\pi_I)$ as well as small values of $\text{MaxPreim}(\pi_J^-)$, $\text{MaxPreim}(\pi_C^-)$ ensure that both (i) and (ii) are unlikely. Here too we refer to the proof of Lemma 38 in Section E for more details.

Interestingly, one could imagine relying on even weaker combinatorial properties of the diffusion permutations in the untangle zones, if the untangle zones had more rounds. We refer to Section A for more ideas in this direction.

THE SIMULATOR IN PSEUDOCODE. Our “official” simulator is given by the pseudocode in Figures 3–5 in Section E, more specifically by all procedures in game G_1 of that pseudocode, except for procedures Z and Z^{-1} which implement the random permutation \mathcal{Z} , and which are not part of the simulator properly speaking.

We note that following [18] our simulator uses *explicit random tapes* for its lazy sampling. More precisely, for each $T \in \{A, \dots, L\}$ and for each $j \in [w]$ there is a tape p_{T_j} (and, implicitly, an inverse $p_{T_j}^-$) that encodes a random permutation of $\{0, 1\}^n$ in the form of a table: $p_{T_j}(x)$ is the image of x for all $x \in \{0, 1\}^n$; moreover $p_{T_j}^-(y)$ is the preimage of y for all $y \in \{0, 1\}^n$. Likewise the random permutation \mathcal{Z} (implemented by the procedures Z and Z^{-1}) relies on a random permutation tape p_Z .

Of course, in a practical implementation of the simulator such random tapes would be too large to keep in memory, and an independent lazy sampling process would substitute for each tape access.

This completes the simulator overview. The full description of the simulator (and the formal security proof) is given in Appendix E.

References

1. Elena Andreeva, Andrey Bogdanov, Yevgeniy Dodis, Bart Mennink, and John P. Steinberger. On the indistinguishability of key-alternating ciphers. In *Advances in Cryptology—CRYPTO 2013* (volume 1), pages 531–550.
2. Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. On the indistinguishability of the sponge construction. In Smart [30], pages 181–197.
3. Alex Biryukov, Charles Bouillaguet, and Dmitry Khovratovich. Cryptographic schemes based on the ASASA structure: Black-box, white-box, and public-key (extended abstract). In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology and Information Security, Kaoshiung, Taiwan, R.O.C., December 7-11, 2014. Proceedings, Part I*, volume 8873 of *Lecture Notes in Computer Science*, pages 63–84. Springer, 2014.
4. D. Chakraborty and P. Sarkar. A new mode of encryption providing a tweakable strong pseudo-random permutation. In *Proceedings of FSE '06*, LNCS 4047, pp. 293–309, 2006.

¹³ Focusing on the first untangle zone, the intuition is, firstly, that small $\text{MaxPreim}(\pi_G)$ implies that all queries to I will be fresh, i.e., will be lazy sampled during AdaptLeft , even if there are some indices $i \neq i'$ for which (some of) those queries overlap; then because π_G is a permutation, for each $i \neq i'$, there is at least one $j \in [w]$ such that $\mathbf{x}_i^I[j] \neq \mathbf{x}_{i'}^I[j]$ and, hence, such that $I_j(\mathbf{x}_i^I[j])$ is independently lazy sampled from $I_j(\mathbf{x}_{i'}^I[j])$; focusing on the last of these two value to be lazy sampled, and coupled with the fact that $\text{MaxPreim}(\pi_I)$ is small, we can conclude that all entries in \mathbf{x}_i^D , $\mathbf{x}_{i'}^D$ are distinct with high probability.

5. D. Chakraborty and P. Sarkar. HCH: A new tweakable enciphering scheme using the hash-encrypt-hash approach. In *Proceedings of Indocrypt '06*, LNCS 4329, pp. 287–302, 2006.
6. Stanley Chow, Philip A. Eisen, Harold Johnson, and Paul C. van Oorschot. White-box cryptography and an AES implementation. In Kaisa Nyberg and Howard M. Heys, editors, *Selected Areas in Cryptography, 9th Annual International Workshop, SAC 2002, St. John's, Newfoundland, Canada, August 15-16, 2002. Revised Papers*, volume 2595 of *Lecture Notes in Computer Science*, pages 250–270. Springer, 2002.
7. Stanley Chow, Philip A. Eisen, Harold Johnson, and Paul C. van Oorschot. A white-box DES implementation for DRM applications. In Joan Feigenbaum, editor, *Security and Privacy in Digital Rights Management, ACM CCS-9 Workshop, DRM 2002, Washington, DC, USA, November 18, 2002, Revised Papers*, volume 2696 of *Lecture Notes in Computer Science*, pages 1–15. Springer, 2002.
8. Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. Merkle-damgård revisited: How to construct a hash function. In Victor Shoup, editor, *Advances in Cryptology—CRYPTO 2005*, volume 3621 of *LNCS*, pages 430–448. Springer-Verlag, 14–18 August 2005.
9. Jean-Sébastien Coron, Yevgeniy Dodis, Avradip Mandal, and Yannick Seurin. A domain extender for the ideal cipher. To appear in the *Theory of Cryptography Conference (TCC)*, February 2010.
10. Jean-Sébastien Coron, Jacques Patarin, and Yannick Seurin. The random oracle model and the ideal cipher model are equivalent. In Wagner [31], pages 1–20.
11. Yevgeniy Dodis, Krzysztof Pietrzak, and Prashant Puniya. A new mode of operation for block ciphers and length-preserving macs. In Smart [30], pages 198–219.
12. Yevgeniy Dodis, Leonid Reyzin, Ronald L. Rivest, and Emily Shen. Indifferentiability of permutation-based compression functions and tree-based modes of operation, with applications to md6. In Orr Dunkelman, editor, *Fast Software Encryption: 16th International Workshop, FSE 2009*, volume 5665 of *Lecture Notes in Computer Science*, pages 104–121. Springer-Verlag, 22–25 February 2009.
13. Yevgeniy Dodis, Thomas Ristenpart, and Thomas Shrimpton. Salvaging merkle-damgård for practical applications. In Antoine Joux, editor, *Advances in Cryptology - EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 371–388. Springer-Verlag, 2009.
14. Hörst Feistel. Cryptographic coding for data-bank privacy. IBM Technical Report RC-2827, March 18 1970.
15. Hörst Feistel, William A. Notz, J. Lynn Smith. Some Cryptographic Techniques for Machine-to-Machine Data Communications. IEEE proceedings, **63**(11), pages 1545–1554, 1975.
16. S.R. Fluhrer and D.A. McGrew. The extended codebook (XCB) mode of operation. Technical Report 2004/078, IACR eprint archive, 2004.
17. S. Halevi. Invertible Universal hashing and the TET Encryption Mode. In *Proceedings of CRYPTO '07*, LNCS 4622, pp. 412–429, 2007.
18. Thomas Holenstein, Robin Künzler, and Stefano Tessaro. The equivalence of the random oracle model and the ideal cipher model, revisited. In Lance Fortnow and Salil P. Vadhan, editors, *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 89–98. ACM, 2011.
19. Rodolphe Lampe and Yannick Seurin. How to construct an ideal cipher from a small set of public permutations. In Kazuo Sako and Palash Sarkar, editors, *Advances in Cryptology - ASIACRYPT 2013 - 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part I*, volume 8269 of *Lecture Notes in Computer Science*, pages 444–463. Springer, 2013.
20. M. Luby and C. Rackoff. How to construct pseudorandom permutations and pseudorandom functions. *SIAM Journal on Computing*, 17(2):373–386, April 1988.
21. Ueli M. Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In Moni Naor, editor, *First Theory of Cryptography Conference — TCC 2004*, volume 2951 of *LNCS*, pages 21–39. Springer-Verlag, February 19–21 2004.
22. Ueli M. Maurer, and Stefano Tessaro. Domain Extension of Public Random Functions: Beyond the Birthday Barrier In *Advances in Cryptology—CRYPTO 2007*, August 2007.
23. Eric Miles and Emanuele Viola. Substitution-permutation networks, pseudorandom functions, and natural proofs. In *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, pages 68–85, 2012.
24. M. Naor and O. Reingold, *On the construction of pseudorandom permutations: Luby-Rackoff revisited*, J. of Cryptology, 1999. Preliminary Version: STOC 1997.
25. Thomas Ristenpart and Hovav Shacham and Thomas Shrimpton, Careful with Composition: Limitations of the Indifferentiability Framework. EUROCRYPT 2011, LNCS 6632, 487–506.
26. Phillip Rogaway and John P. Steinberger. Constructing cryptographic hash functions from fixed-key blockciphers. In Wagner [31], pages 433–450.

27. Yannick Seurin, *Primitives et protocoles cryptographiques à sécurité prouvée*. PhD thesis, Université de Versailles Saint-Quentin-en-Yvelines, France, 2009.
28. Claude E. Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, 28(4):656–715, October 1949. Available at <http://www.cs.ucla.edu/~jkong/research/security/shannon.html> and <http://www3.edgenet.net/dcowley/docs.html>.
29. Thomas Shrimpton and Martijn Stam. Building a collision-resistant compression function from non-compressing primitives. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz, editors, *ICALP (2)*, volume 5126 of *Lecture Notes in Computer Science*, pages 643–654. Springer, July 7-11 2008.
30. Nigel P. Smart, editor. *Advances in Cryptology - EUROCRYPT 2008*, volume 4965 of *LNCS*. Springer-Verlag, 2008.
31. David Wagner, editor. *Advances in Cryptology - CRYPTO 2008*, volume 5157 of *LNCS*. Springer-Verlag, 2008.
32. P. Wang, D. Feng, and W. Wu. HCTR: A variable-input-length enciphering mode. In *Proceedings of CISC '05*, LNCS 3822, pp. 175–188, 2005.

A Extensions and Future Work

In this section we make some speculative remarks about possible extensions of our work. In so doing we also outline a broader perspective for understanding our simulator.

A *network segment* consists of an alternation of D -boxes and of S -box rounds, starting either with a D -box or with an S -box round and ending either with a D -box or with an S -box round. The network segment is *even* if it starts and ends with S -box rounds; it is *odd* if starts and ends with D -box rounds.

For example, the untangle zones in our simulator are odd network segments regardless of whether `XtraUntglRnds` is toggled. All other zones (and considering the left outer detect and right outer detect zones as separate zones) are even.

An important element of the present “broader perspective” is the observation that all combinatorial properties of D -boxes that we have defined—i.e., conductance, all-but-one conductance, randomized preimage resistance, randomized collision resistance—are special cases of more general definitions that apply to network segments.

We start by describing these generalized definitions. Formally, each property we define is a predicate of a sequence of D -boxes $\bar{\pi} = (\pi_1, \dots, \pi_t)$. Each definition either targets even or odd network segments. In the case of an even network segment the case $t = 0$ is allowed, corresponding to a network segment that consists of a single S -box round.

In more detail, for definitions that apply to even network segments we view $\bar{\pi}$ as part of a $(t + 1)$ -round confusion-difusion network $P[\mathcal{P}, \bar{\pi}]$ where $\mathcal{P} = \{P_{i,j} : i \in [t + 1], j \in [w]\}$ and where π_i is the D -box between the S -box rounds $\{P_{i,j} : j \in [w]\}$ and $\{P_{i+1,j} : j \in [w]\}$. For definitions that apply to odd network segments we view $\bar{\pi}$ as part of a $(t - 1)$ -round confusion-diffusion network $P[\mathcal{P}, \bar{\pi}]$ where $\mathcal{P} = \{P_{i,j} : i \in \{2, \dots, t\}, j \in [w]\}$, where π_i is again the D -box between rounds $\{P_{i,j} : j \in [w]\}$ and $\{P_{i+1,j} : j \in [w]\}$. In either case the permutations $P_{i,j} \in \mathcal{P}$ will be modeled as a set of *random* independent permutations accessible as oracles to an adversary.

Conductance. The *generalized conductance* $\text{Cond}_{\bar{\pi}}^*$ of $\bar{\pi} = (\pi_1, \dots, \pi_t)$ is defined with respect to even network segments. Formally $\text{Cond}_{\bar{\pi}}^*$ is a function

$$\text{Cond}_{\bar{\pi}}^* : \mathbb{Z} \times \mathbb{Z} \rightarrow [0, 1]$$

defined by the following game. Let q and Q be integers. Let $\mathcal{P} = \{P_{i,j} : i \in [t + 1], j \in [w]\}$ be the array of random permutation oracles for the even network, as described above. We give an

adversary q queries (forward or backward) to each of the oracles in \mathcal{P} . When the adversary is done, we count the number of completed paths that can be assembled from these queries (i.e., inputs $\mathbf{x} \in \{0, 1\}^{wn}$ to $P[\mathcal{P}, \bar{\pi}]$ that can be evaluated to outputs using only the adversary's queries). We say the adversary “wins” if this number of paths is more than Q . We let

$$\text{Cond}_{\bar{\pi}}^*(q, Q)$$

be the supremum¹⁴ of the probability of winning this game, taken over all possible adversaries.

Relation to “ordinary” conductance. If $t = 1$, i.e., $\bar{\pi} = (\pi_1)$, then $\text{Cond}_{\bar{\pi}}^*$ is characterized by Cond_{π_1} in the sense that

$$\text{Cond}_{\bar{\pi}}^*(q, Q) = \begin{cases} 0 & \text{if } Q > \text{Cond}_{\pi_1}(q), \\ 1 & \text{if } Q \leq \text{Cond}_{\pi_1}(q). \end{cases}$$

This is easy to see from the definitions. (Indeed, the number of paths assembled is easily seen to be purely a function of the sets $U_j := \{y : P_{1,j}^{-1}(y) \neq \perp\}$, $1 \leq j \leq w$, and $V_j := \{x : P_{2,j}(x) \neq \perp\}$, $1 \leq j \leq w$, where we write $P_{1,j}^{-1}(y) \neq \perp$ to mean that the adversary either made the query $P_{1,j}^{-1}(y)$ or made a query $P_{1,j}(x)$ whose answer resulted in y . Hence, the best possible adversary will simply choose the sets U_1, \dots, V_w beforehand, and make the relevant inverse queries in the first round and the relevant forward queries in the second round.)

Moreover when $t = 0$ then $\bar{\pi} = ()$ where $()$ denotes an empty sequence of S -boxes, and one has

$$\text{Cond}_{\bar{\pi}}^*(q, Q) = \begin{cases} 0 & \text{if } Q > q^w, \\ 1 & \text{if } Q \leq q^w. \end{cases}$$

Thus one can observe that $\alpha(q)$, as defined in our main proof, is simply the threshold Q at which the generalized conductance $\text{Cond}_{\bar{\pi}}^*(2q, \cdot)$ of the middle round goes from 1 to 0, whether or not `XtraMiddleRnd` is on.

Likewise, $\beta(q)$ is the same threshold, only with respect to the generalized conductance $\text{Cond}_{\bar{\pi}}^*(q + \alpha(q), \dots)$ and with respect to the sequence of D -boxes $\bar{\pi}$ in the outer left detect zone.

All-but-one conductance. The *generalized all-but-one conductance* $\text{aboCond}_{\bar{\pi}}^*$ of $\bar{\pi}$ is also defined with respect to even network segments. Formally, $\text{aboCond}_{\bar{\pi}}^*$ is a function

$$\text{aboCond}_{\bar{\pi}}^* : \mathbb{Z} \rightarrow [0, 1]$$

defined with respect to the following game. Let q be an integer. An adversary is given q queries to each of the permutations in $\mathcal{P} = \{P_{i,j} : i \in [t+1], j \in [w]\}$. We say that a given adversarial query is *internal* if it has the form $P_{i,j}(x)$ with $i < t+1$ or $P_{i,j}^{-1}(y)$ with $i > 1$. Moreover we assume that the adversary does not make redundant queries. Then the adversary wins if, at any point, it makes an internal query such that a completed path can be assembled using (the result of) this internal query together with queries that were made prior to this internal query. We let

$$\text{aboCond}_{\bar{\pi}}^*(q)$$

be the maximum probability of an adversary winning this game, taken over all adversaries.

¹⁴ Given the finite amount of randomness involved, it is easy to see that the supremum is achieved by some adversary. Hence we might also say *maximum*.

Relation to “ordinary” all-but-one conductance. If $t = 1$ then one can verify that

$$\text{aboCond}_{\bar{\pi}}^*(q) \leq \frac{2wq \cdot \text{aboCond}_{\pi_1}(q)}{N - q}$$

by a straightforward union bound. On the other hand if $t = 0$ then one can verify that

$$\text{aboCond}_{\bar{\pi}}^*(q) = 0$$

regardless of q , since there is no such thing as an internal query in this case.

Entry-wise randomized preimage resistance. The *generalized entry-wise randomized preimage resistance* $\text{MaxPreim}_{\bar{\pi}}^*$ of $\bar{\pi}$ is defined with respect to odd network segments. Formally $\text{MaxPreim}_{\bar{\pi}}^*$ is a function

$$\text{MaxPreim}_{\bar{\pi}}^* : \mathbb{Z} \rightarrow [0, 1]$$

defined by the following game. Let q be an integer. An adversary is allowed q queries to each of the permutations in $\mathcal{P} = \{P_{i,j} : i \in \{2, \dots, t\}, j \in [w]\}$ in the odd network. After making its queries, the adversary names an index $j \in [w]$ and a set of values $\{x_{j'} : j' \neq j\}$ as well as an index $h \in [w]$ and a value $y_h \in \{0, 1\}^n$. A value x_j is sampled uniformly at random from $\{0, 1\}^n$ and the (odd) network is evaluated on the input vector $\mathbf{x} = (x_1, \dots, x_w)$, making additional queries to the $P_{i,j}$'s if necessary, until an output vector \mathbf{y} is obtained. (Thus \mathbf{x} is the input to the first D -box, \mathbf{y} is the output of the last D -box.) The adversary wins if $\mathbf{y}[h] = y_h$. We define

$$\text{MaxPreim}_{\bar{\pi}}^*(q)$$

to be the maximum probability of an adversary winning this game, taken over all adversaries.

Relation to “ordinary” entry-wise randomized preimage resistance. If $t = 1$ then there are no S -boxes to query, so the value q is immaterial, and it is easily verified that

$$\text{MaxPreim}_{\bar{\pi}}^*(q) = \frac{\text{MaxPreim}(\pi_1)}{N}.$$

If $t = 2$ one can also check that

$$\text{MaxPreim}_{\bar{\pi}}^*(q) \leq \frac{q \cdot \text{MaxPreim}(\pi_1)}{N} + \frac{q}{N} + \frac{\text{MaxPreim}(\pi_2)}{N}$$

by using an argument somewhat similar to the case $\text{XtraUntglRnds} = \text{true}$ of Lemma 38 in Section E.3. (A factor w is missing from the first fraction due to the fact that it is only necessary for, say, the query $P_{2,1}(\pi_1(\mathbf{x}))[1]$ to be “fresh”, among all the queries $P_{2,j}(\pi_1(\mathbf{x}))[j]$, $1 \leq j \leq w$. The fraction q/N accounts for the difference between sampling uniformly at random from $\{0, 1\}^n$ and sampling uniformly at random from a subset of $\{0, 1\}^n$ of size $N - q$; specifically, the latter kind of sampling can be implemented by sampling uniformly at random, and giving the adversary a win if the sample falls in the forbidden subset of q elements.)

Entry-wise randomized collision resistance. The *generalized entry-wise randomized preimage resistance* $\text{MaxColl}_{\bar{\pi}}^*(q)$ of $\bar{\pi}$ is also defined with respect to odd network segments. Formally

$$\text{MaxColl}_{\bar{\pi}}^* : \mathbb{Z} \rightarrow [0, 1]$$

is a function defined via the following game. Let q be an integer. An adversary is allowed q queries to each of the odd network permutations $\mathcal{P} = \{P_{i,j} : i \in \{2, \dots, t\}, j \in [w]\}$. After making its queries, the adversary names an index $j \in [w]$ and two set of values $\{x_{j'} : j' \neq j\}$, $\{x'_{j'} : j' \neq j\}$ such that $x_{j'} \neq x'_{j'}$ for at least one $j' \neq j$. A uniform random value $x_j \in \{0, 1\}^n$ is revealed, and the network is evaluated (requiring new queries if necessary) on the inputs $\mathbf{x} = (x_1, \dots, x_w)$ and $\mathbf{x}' = (x'_1, \dots, x'_w)$ where $x'_j := x_j$, resulting in outputs \mathbf{y} and \mathbf{y}' . The adversary wins if $\mathbf{y}[h] = \mathbf{y}'[h]$ for some $h \in [w]$. We define

$$\text{MaxColl}_{\bar{\pi}}^*(q)$$

to be the maximum probability of an adversary winning this game, taken over all adversaries.

Relation to “ordinary” entry-wise randomized collision resistance. If $t = 1$ then q is extraneous and it is easy to see that

$$\text{MaxColl}_{\bar{\pi}}^*(q) = \frac{w \text{MaxColl}(\pi_1)}{N}$$

whereas if $t = 2$ then one can see that

$$\text{MaxColl}_{\bar{\pi}}^*(q) \leq \frac{wq \cdot \text{MaxPreim}(\pi_1)}{N} + \frac{q+1}{N} + \frac{w \cdot \text{MaxPreim}(\pi_2)}{N}. \quad (4)$$

similarly to the argument in the case $\text{XtraUntglRnds} = \mathbf{true}$ proof of Lemma 38, Section E.3.

A MORE GENERAL THEOREM. One can easily state a more general version of our main theorem. The parameters of the network would be sequences of D -boxes $\bar{\nu}$, $\bar{\pi}_I$, $\bar{\pi}_J$, $\bar{\tau}$, $\bar{\pi}_K$ and $\bar{\pi}_L$ that appear in respectively the left outer detect zone, the first untangle zone, the second untangle zone, the middle detect zone, the third untangle zone and the fourth untangle zone. Since the untangle zones are implemented by odd network segments, each of $\bar{\pi}_I$, $\bar{\pi}_J$, $\bar{\pi}_K$ and $\bar{\pi}_L$ contain at least one D -box each.

To state such a generalized theorem, let $\eta = (\bar{\nu}, \bar{\pi}_I, \bar{\pi}_J, \bar{\tau}, \bar{\pi}_K, \bar{\pi}_L)$ (a sequence-of-sequences, not a concatenation of sequences) and let

$$\begin{aligned} \text{MaxPreim}_{\eta}^*(q) &:= \max(\text{MaxPreim}_{\bar{\pi}_I}^*(q), \text{MaxPreim}_{\bar{\pi}_J}^*(q), \text{MaxPreim}_{\bar{\pi}_K}^*(q), \text{MaxPreim}_{\bar{\pi}_L}^*(q)) \\ \text{MaxColl}_{\eta}^*(q) &:= \max(\text{MaxColl}_{\bar{\pi}_I}^*(q), \text{MaxColl}_{\bar{\pi}_J}^*(q), \text{MaxColl}_{\bar{\pi}_K}^*(q), \text{MaxColl}_{\bar{\pi}_L}^*(q)) \end{aligned}$$

where $\bar{\pi}^-$ stands for the sequence of inverses of permutations in $\bar{\pi}$, in reverse order than $\bar{\pi}$. Moreover let $|\bar{\pi}|$ be the length (number of sequences) in $\bar{\pi}$, and let

$$r_{\nu} = |\bar{\nu}| + 1, r_I = |\bar{\pi}_I| - 1, r_J = |\bar{\pi}_J| - 1, r_{\tau} = |\bar{\tau}| + 1, r_K = |\bar{\pi}_K| - 1, r_L = |\bar{\pi}_L| - 1$$

be the number of (S -box) rounds in the outer left detect zone, etc, and let

$$r = r_{\nu} + r_I + r_J + r_{\tau} + r_K + r_L + 3.$$

be the total number of rounds in the network. (The ‘3’ accounts for rounds D , E and H in our original nomenclature.) Finally let

$$\bar{\eta} = \bar{\nu} \|\bar{\pi}_I\| \bar{\pi}_J \|\bar{\tau}\| \bar{\pi}_K \|\bar{\pi}_L$$

be the concatenation of the D -boxes sequences $\bar{\nu}, \dots, \bar{\pi}_L$ into a single sequence, and let $\mathcal{Q} = \{P_{i,j} : i \in [r], j \in [w]\}$ denote a set of random permutations.

Then we have the following result (stated without proof):

Theorem 2. For all integers $Q_{\overline{D}}$ and $Q_{\overline{E}}$, the confusion-diffusion network $P[\mathcal{Q}, \overline{\eta}]$ achieves (t_S, q_S, ε) -indifferentiability from a random permutation $\mathcal{Z} : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$, for ε equal to

$$\text{Cond}_{\overline{E}}^*(2q, Q_{\overline{E}}) + \text{Cond}_{\overline{D}}^*(q + Q_{\overline{E}}, Q_{\overline{D}}) \quad (5)$$

$$+ \frac{w(q + Q_{\overline{E}})^2}{N - q - Q_{\overline{E}}} + \frac{(q + wQ_{\overline{E}})(q + Q_{\overline{E}})}{N - q - Q_{\overline{E}}} + \frac{2w(q + Q_{\overline{E}})Q_{\overline{E}}}{N - q - Q_{\overline{E}}} \quad (6)$$

$$+ \text{aboCond}_{\overline{E}}^*(2q) + \text{aboCond}_{\overline{D}}^*(q + Q_{\overline{E}}) \quad (7)$$

$$+ 2wQ_{\overline{E}} [(q + Q_{\overline{E}}) \cdot \text{MaxPreim}_{\eta}^*(q + Q_{\overline{E}}) + (q + Q_{\overline{E}})/N] \quad (8)$$

$$+ 2Q_{\overline{E}} [Q_{\overline{E}} \cdot \text{MaxColl}_{\eta}^*(q + Q_{\overline{E}}) + (q + Q_{\overline{E}})/N] \quad (9)$$

$$+ \frac{Q_{\overline{D}}(q + Q_{\overline{E}})^w}{N^w - q - Q_{\overline{E}}} + \frac{1}{N^w} \quad (10)$$

and for $q_S = Q_{\overline{E}}$, $t_S = O(w(q + Q_{\overline{E}})^w)$. Here $q = q_0(1 + rw)$ where r is the number of S -box rounds.

In this statement, $Q_{\overline{E}}$ is the (expected) maximum number of matching query sets for the middle detect zone, given $2q$ queries to each S -box in that detect zone; $Q_{\overline{D}}$ is the (expected) maximum number of matching query sets for the outer left detect zone, given $q + Q_{\overline{E}}$ queries to each S -box in that zone. The first two terms in the security bounds account for the distinguisher's probability of beating these "expected" conductances, and $Q_{\overline{E}}$, $Q_{\overline{D}}$ should be set high enough to make the first two terms negligible.

The remainder of the security bound affords the following term-by-term interpretation (we slightly abusively refer to "game G_3 ", having the natural extrapolation of our previous proof in mind):

- the first term of line (6) upper bounds the probability that a fresh query to Z or Z^{-1} in game G_3 accidentally hits a predefined value in one of the S -boxes of the last or first round, respectively (compare with Lemma 26)

- the second term of line (6) upper bounds the probability that a forward query in the last round or a backward query in the first round causes abort in game G_3 by "hitting" (the coordinate of) a pre-existing query to Z or Z^{-1} (compare with Lemma 28)

- the third term of line (6) upper bounds the probability that procedures D/D^{-1} or E/E^{-1} abort in game G_3 (compare with Lemma 31)

- the first term of line (7) upper bounds the probability that a middle path completion is ever triggered in game G_3 by a query that is neither a forward query to the last round of the middle detect zone nor a backward query in the first round of the middle detect zone (compare with Lemma 34)

- the second term of line (7) is similar, but for the left outer detect zone (compare with Lemma 29)

- line (8) upper bounds the probability that `LeftAdapt` or `RightAdapt` ever abort in game G_3 because an input or output to D (respectively E) hits a predefined value in D (respectively E); the extra additive term $(q + Q_{\overline{E}})/N$ inside the square brackets accounts for the difference between uniform sampling and lazy sampling

- line (9) upper bounds the probability that `LeftAdapt` or `RightAdapt` ever abort in game G_3 because inputs or outputs to D (respectively E) for two different adapted paths collide; the extra additive term $(q + Q_{\overline{E}})/N$ again accounts for the difference between uniform sampling and lazy sampling

- the first term of line (10) upper bounds the probability that a fresh query to Z or Z^{-1} by the simulator in game G_1 (or G_2 or G_3) accidentally hits a column of defined queries (compare with Lemma 40)

- the second term of line (10) is incurred by probability ratios in the randomness map from G_3 to G_4 , cf. Lemma 46

One can observe that MaxPreim_η^* and MaxColl_η^* appear with the same arguments and with essentially the same coefficients. Hence it makes sense to define

$$\text{MaxCoPr}_{\bar{\pi}}^*(q) := \max(\text{MaxPreim}_{\bar{\pi}}^*(q), \text{MaxColl}_{\bar{\pi}}^*(q))$$

and to characterize the quality of an untangle zone by the single function $\text{MaxCoPr}^*(\cdot)$. Typically we expect $\text{MaxCoPr}_{\bar{\pi}}^* = \text{MaxColl}_{\bar{\pi}}^*$, though it is possible to devise pathological examples for which $\text{MaxPreim}_{\bar{\pi}}^* \gg \text{MaxColl}_{\bar{\pi}}^*$.

MAKING IT FAST. Given these general observations, a basic class of questions that can be pursued is the following:

how can one build fast networks¹⁵ with low conductance/randomized preimage resistance/randomized collision resistance?

As a special case of this question, we have:

how can one build fast diffusion permutations with low conductance/randomized preimage resistance/randomized collision resistance?

As seen by the form of our main theorem (or equivalently, as evidenced by (4)), the second question can also potentially serve as a building block for the first.

Importantly, one should not necessarily equate efficiency with the smallest possible number of rounds. For example, a $\text{GF}(2^n)$ -linear diffusion permutation whose matrix has no zero entries has low randomized preimage resistance on its own. But arbitrary field operations are slow, and one could easily imagine there might exist a sequence $\bar{\pi} = (\pi_1, \pi_2)$ with good randomized preimage resistance that outperforms—that is, even taking into account the S -box round between π_1 and π_2 —a standalone linear diffusion permutation. The question is whether the use of an extra round makes it possible to sufficiently simplify the diffusion permutations.

Similarly, a network zone with low conductance might be more quickly achieved by using three S -box rounds and two “cheap” diffusion permutations than by using one S -box round and one “expensive” diffusion permutation that controls conductance on its own.

As a concrete idea, consider the D -boxes $\oplus_\uparrow : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$, $\oplus_\downarrow : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$ defined by

$$\begin{aligned} \oplus_\uparrow(\mathbf{x})[j] &= \bigoplus_{1 \leq h \leq j} \mathbf{x}[h], \\ \oplus_\downarrow(\mathbf{x})[j] &= \bigoplus_{j \leq h \leq w} \mathbf{x}[h]. \end{aligned}$$

Then \oplus_\uparrow and \oplus_\downarrow can be efficiently implemented in software or in hardware. While \oplus_\uparrow , \oplus_\downarrow both have worst-possible randomized preimage and collision resistance on their own, one can check that

$$\text{MaxPreim}_{(\oplus_\downarrow, \oplus_\uparrow)}^*(q) \leq \frac{q}{N} + \frac{1}{N - q}$$

¹⁵ By “networks” we mean “network segments” or, even more precisely, “ D -box sequence”. In this discussion we write *conductance* instead of *generalized conductance*, etc.

and moreover that

$$\begin{aligned}\text{MaxPreim}_{(\oplus_{\downarrow}, \oplus_{\uparrow}, \oplus_{\downarrow}, \oplus_{\uparrow})}^*(q) &\leq \frac{q}{N} + \frac{2q}{N-q} + \frac{1}{N-q} \\ \text{MaxColl}_{(\oplus_{\downarrow}, \oplus_{\uparrow}, \oplus_{\downarrow}, \oplus_{\uparrow})}^*(q) &\leq \frac{q}{N} + \frac{2q}{N-q} + \frac{w}{N-q}\end{aligned}$$

which suggests that the sequence $\overline{\oplus}_4 := (\oplus_{\downarrow}, \oplus_{\uparrow}, \oplus_{\downarrow}, \oplus_{\uparrow})$ could serve as a very efficient¹⁶ untangle zone. A confusion-diffusion network that used this untangle zone four times over¹⁷ and that used, say, two S -box rounds for each of the outer left and middle detect zones would have $2 \cdot 2 + 4 \cdot 3 + 2 + 1 = 19$ rounds. The main “takeaway lesson” being that by targeting “the right combinatorial property at the right time” one can potentially arrive at sequences of very simple D -boxes.

Tantalizingly, however—and as far as we know!—the above upper bounds on $\text{MaxPreim}_{\overline{\oplus}_4}^*$, $\text{MaxColl}_{\overline{\oplus}_4}^*$ crucially require the S -boxes in each round to be independent. Another potential research direction, therefore, would be to devise D -box sequences that are “ S -box reuse resistant”, or more precisely, that preserve low randomized preimage and collision resistance even when the S -boxes in a column are all the same, or even when all the S -boxes in the network segment are the same.

In parallel, one could envision extending Theorem 1 or Theorem 2 to cover the possibility of non-independent S -boxes. (Within each round, within each zone, or—most ambitiously but also most messily—within the entire network.)

Of course, and even with extensions to cover the case of non-independent S -boxes, one should bear in mind that security guarantees of the form $q^{O(1)}/N$ are meaningless for, say, $N = 2^8$, not to mention the fact that our results are obtained with respect to an ideal model, while S -boxes must be concretely implemented. Still, we tend to believe that indifferentiability represents a nontrivial and potentially important “sanity check” for the high-level design of a confusion-diffusion network, and which has been absent so far. It could therefore be worthwhile to pursue more concrete instantiations (accompanied by “real” cryptanalysis) for constructions that fall under the umbrella of Theorem 2, or some even more general extension thereof.

B Probabilistic constructions of low-conductance permutations

In this section we show the probabilistic construction of permutations $\tau : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$ with low conductance and low all-but-one conductance.

Let $U_1, \dots, U_w, V_1, \dots, V_w \subseteq \{0, 1\}^n$ such that $|U_j| = |V_j| = q$, $1 \leq j \leq w$. Let $a \in \mathbb{N}$. We start by upper bounding the probability that when a permutation $\tau : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$ is chosen uniformly at random,

$$|\{\mathbf{x} \in U_1 \times \dots \times U_w : \tau(\mathbf{x}) \in V_1 \times \dots \times V_w\}| \geq a. \quad (11)$$

¹⁶ In fact one could even use the simpler transformation

$$\oplus'_{\downarrow}(\mathbf{x})[j] = \begin{cases} \mathbf{x}[j] & \text{if } j \neq w \\ \bigoplus_{1 \leq h \leq w} \mathbf{x}[h] & \text{if } j = w. \end{cases}$$

in place of \oplus_{\downarrow} , though one’s cryptographic instincts recoil in the face of such weak diffusion. Interestingly, too, $(\oplus_{\uparrow}, \oplus'_{\downarrow})$ has bad randomized preimage resistance, even while $(\oplus'_{\downarrow}, \oplus_{\uparrow})$ has randomized preimage resistance similar to $(\oplus_{\downarrow}, \oplus_{\uparrow})$. Similarly $(\oplus_{\uparrow}, \oplus'_{\downarrow}, \oplus_{\uparrow}, \oplus'_{\downarrow})$ has bad randomized collision resistance while $(\oplus'_{\downarrow}, \oplus_{\uparrow}, \oplus'_{\downarrow}, \oplus_{\uparrow})$ has good randomized collision resistance.

¹⁷ Indeed, our bounds don’t suggest that different untangle zones should be implemented any differently apart from the obvious left-right mirror symmetry, i.e., $\overline{\pi}_J = \overline{\pi}_J^-$.

We assume that $q^w \leq N^w/2$ (where $N = 2^n$ as usual), which means that q cannot be too close to N . (In our applications we have $q \leq \sqrt{N}$, which is much smaller.)

Let $\mathbf{x}_1, \dots, \mathbf{x}_{q^w}$ be an enumeration of the elements in $U_1 \times \dots \times U_w$. If $\tau : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$ is a random permutation and the values $\tau(\mathbf{x}_1), \dots, \tau(\mathbf{x}_\ell)$ have already been revealed, then the probability that $\tau(\mathbf{x}_{\ell+1}) \in V_1 \times \dots \times V_w$ is at most

$$\frac{q^w}{N^w - \ell} \leq \frac{q^w}{N^w - q^w} \leq \frac{2q^w}{N^w}$$

by the assumption $q^w \leq N^w/2$. That is,

$$\Pr_{\tau} [\tau(\mathbf{x}_{\ell+1}) \in V_1 \times \dots \times V_w \mid \tau(\mathbf{x}_1), \dots, \tau(\mathbf{x}_\ell)] \leq \frac{2q^w}{N^w}$$

for any values of $\tau(\mathbf{x}_1), \dots, \tau(\mathbf{x}_\ell)$, $\ell \leq q^w - 1$.

Now the probability that (11) occurs is upper bounded by

$$\binom{q^w}{a} \left(\frac{2q^w}{N^w} \right)^a \leq q^{wa} \left(\frac{2q^w}{N^w} \right)^a = \left(\frac{2q^{2w}}{N^w} \right)^a.$$

By a union bound, then, the probability (taken over the choice of a random τ) that there *exist* $U_1, \dots, U_w, V_1, \dots, V_w$ such that (11) occurs is upper bounded by

$$\binom{N}{q}^{2w} \left(\frac{2q^{2w}}{N^w} \right)^a \leq N^{2qw} \left(\frac{2q^{2w}}{N^w} \right)^a. \quad (12)$$

This probability is less than 1, thus, as long as $q^{2w} \leq N^w/4$ (which essentially translates to $q < \sqrt{N}$) and as long as $a > 2nqw$. We have thus proved:

Theorem 3. *Let $q, n, w \in \mathbb{N}$ be such that $q^{2w} \leq N^w/4$, $N = 2^n$. Then there exists a permutation $\tau : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$ such that*

$$\text{Cond}(\tau) \leq 2nqw + 1.$$

The conductance of the permutation τ constructed by Theorem 3 is thus quite close to q , being only a factor $2nw = O(\log(N^w))$ away from q . One can also observe that (12) does not deliver much better conductance even for $q^{2w} \ll N^w$, since

$$N^{2qw} \left(\frac{2q^{2w}}{N^w} \right)^a \geq N^{2qw} \cdot \frac{1}{N^{wa}}$$

means that we at least need $a > 2q$ in order to strictly upper bound the union bound by 1.

All-but-one conductance. For the case of all-but-one-conductance, start by fixing sets $U_1, \dots, U_w, V_2, \dots, V_w \subseteq \{0, 1\}^n$ and a value $a \in \mathbb{N}$. This time we focus on the probability (over random choice of τ) that

$$|\{\mathbf{x} \in U_1 \times \dots \times U_w : \tau(\mathbf{x}) \in \{0, 1\}^n \times V_2 \times \dots \times V_w\}| \geq a \quad (13)$$

occurs.

As before let $\mathbf{x}_1, \dots, \mathbf{x}_{q^w}$ be an enumeration of $U_1 \times \dots \times U_w$. Then for $0 \leq \ell \leq q^w - 1$, we have

$$\Pr[\tau(\mathbf{x}_{\ell+1}) \in \{0, 1\}^n \times V_2 \times \dots \times V_w \mid \tau(\mathbf{x}_1), \dots, \tau(\mathbf{x}_\ell)] \leq \frac{Nq^{w-1}}{N - q^w} \leq \frac{Nq^{w-1}}{N^w/2} = \frac{2q^{w-1}}{N^{w-1}}$$

for any values of $\tau(\mathbf{x}_1), \dots, \tau(\mathbf{x}_\ell)$.

The probability that (13) occurs is thus upper bounded by

$$\binom{q^w}{a} \left(\frac{2q^{w-1}}{N^{w-1}} \right)^a \leq q^{wa} \left(\frac{2q^{w-1}}{N^{w-1}} \right)^a = \left(\frac{2q^{2w-1}}{N^{w-1}} \right)^a.$$

By a union bound, the probability that there exist sets $U_1, \dots, U_w, V_2, \dots, V_w$ such that (13) occurs is at most

$$\binom{N}{q}^{2w-1} \left(\frac{2q^{2w-1}}{N^{w-1}} \right)^a \leq N^{q(2w-1)} \left(\frac{2q^{2w-1}}{N^{w-1}} \right)^a.$$

Unfortunately, for this bound to give a non-void result we need

$$\frac{2q^{2w-1}}{N^{w-1}} < 1$$

which is too constraining to be useful, since in our setting we need to let q approach $N^{1/2}$. Hence, a different approach is needed.

Multiplicative Chernoff bound. As above, start by fixing sets $U_1, \dots, U_w, V_2, \dots, V_w$ and let $\mathbf{x}_1, \dots, \mathbf{x}_{q^w}$ be an enumeration of $U_1 \times \dots \times U_w$. Also assume that $q \leq \sqrt{N}$, since this is the regime which interests us. Let X_i be the indicator random variable for the event $\tau(\mathbf{x}_i) \in \{0, 1\}^n \times V_2 \times \dots \times V_w$, $1 \leq i \leq q^w$. Then

$$\Pr[X_i = 1 \mid X_1 \dots X_{i-1}] \leq \frac{Nq^{w-1}}{N^w - q^w} \leq \frac{2Nq^{w-1}}{N^w} = \frac{2q^{w-1}}{N^{w-1}} \quad (14)$$

for all outcomes $X_1 \dots X_{i-1}$. Let Y_1, \dots, Y_{q^w} be *independent* indicator random variables such that

$$\Pr[Y_i = 1] = \frac{2q^w}{N^w}$$

for all i . Let $X = \sum_i X_i$, $Y = \sum_i Y_i$. An easy coupling argument shows that

$$\Pr[X > t] \leq \Pr[Y > t]$$

for all $t > 0$, by (14).

Let $\mu = E[Y] = q^w(2q^{w-1}/N^{w-1}) = 2q^{2w-1}/N^{w-1}$. Since $N \geq q^2$ we have

$$\mu \leq 2q^{2w-1}/(q^2)^{w-1} = 2q \leq qw \ln(N) \quad (15)$$

(using $w \geq 2$) and we also have

$$\mu \leq 2qw \ln(N)/(e^2 + 1) \quad (16)$$

and

$$\mu \leq qw \ln(N) - \frac{1}{2} \ln(2w) \quad (17)$$

for sufficiently large n (e.g., $n \geq 3$, which implies $\ln(N) \geq 2$). These inequalities will play a role later.

A well-known multiplicative Chernoff bound states that

$$\Pr[Y > (1 + \delta)\mu] < \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^\mu$$

for all $\delta > 0$, since the Y_i 's are independent. As

$$\left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^\mu \leq \left(\frac{e^\delta}{\delta^\delta} \right)^\mu = e^{(1 - \ln(\delta))\delta\mu} \leq e^{-\delta\mu}$$

for $\delta \geq e^2$, we thus have

$$\Pr[X > (1 + \delta)\mu] < e^{-\delta\mu}$$

for all $\delta \geq e^2$.

Thus

$$N^{2qw} \Pr[X > (1 + \delta)\mu] < e^{2qw \ln(N) - 2\delta\mu}$$

for $\delta \geq e^2$. Note that

$$\begin{aligned} 2qw \ln(N) - 2\delta\mu &< 0 \\ \iff \delta\mu &> qw \ln(N) \\ \iff (1 + \delta)\mu &\geq 2qw \ln(N) \end{aligned}$$

where the last implication uses (15). Thus letting

$$\delta = 2qw \ln(N)/\mu - 1 \geq 2qw \ln(N) - 1$$

we have $\delta \geq e^2$ by (16) and

$$N^{2qw} \Pr[X > 2qw \ln(N)] = N^{2qw} \Pr[X > (1 + \delta)\mu] < 1$$

which implies, by a union bound over all

$$\binom{N}{q}^{2w} \leq N^{2qw}$$

possible choices of $U_1, \dots, U_w, V_2, \dots, V_w$ the existence of a permutation $\tau : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$ such that

$$\text{Cond}^{1,+}(\tau, q) \leq 2qw \ln(N).$$

In fact, moreover, a quick revisit of the computations above shows that our choice of δ not only gives

$$2qw \ln(N) - 2\delta\mu < 0$$

but more strongly gives (using (17))

$$2qw \ln(N) - 2\delta\mu < -\ln(2w)$$

which thus implies

$$N^{2qw} \Pr[X > 2qw \ln(N)] = N^{2qw} \Pr[X > (1 + \delta)\mu] < \frac{1}{2w}$$

and, by an additional union bound over $h \in [w]$ and in $\sigma \in \{+, -\}$, the existence of a permutation $\tau : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$ such that

$$\text{Cond}^{h,\sigma}(\tau, q) \leq 2qw \ln(N)$$

for all $h \in [w]$, $\sigma \in \{+, -\}$, i.e., such that

$$\text{aboCond}(\tau, q) \leq 2qw \ln(N).$$

We have proved:

Theorem 4. *Let $q, n, w \in \mathbb{N}$ be such that $n \geq 3$, $q \leq \sqrt{N}$. Then there exists a permutation $\tau : \{0, 1\}^{wn} \rightarrow \{0, 1\}^{wn}$ such that*

$$\text{aboCond}(\tau) \leq 2qw \ln(N).$$

C Lower Bounds on the Conductance of Linear Permutations

In this section we show that linear permutations cannot achieve conductance as low as random permutations. Specifically, we show that

$$\text{Cond}_\tau(q) \geq \Omega(q^{2 - \frac{1}{2w-1}})$$

for a “generic” \mathbb{F} -linear transformation $\tau : \mathbb{F}^w \rightarrow \mathbb{F}^w$, where \mathbb{F} is a finite field. This lower bound is valid for any $q = o(|\mathbb{F}|)$, and hides a constant factor dependent on w . (The constant factor is roughly 2^{-w^3} .) The linear permutation is “generic” in the sense that our lower bound holds with probability¹⁸ $1 - o(1)$ over the choice of a random matrix $A \in \mathbb{F}^{w \times w}$. We conjecture that our lower bound should carry over to non-generic (i.e., arbitrary) linear permutations as well, but our analysis does not cover this case.

We start by describing the result for $w = 2$, as the general case is a bit unwieldy. When $w = 2$ the linear transformation τ takes the form of a 2×2 matrix

$$A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}$$

where $\tau(\mathbf{x}) = A\mathbf{x}$. Or, if $\mathbf{x} = (x_1, x_2)$, $\tau(\mathbf{x}) = \mathbf{y} = (y_1, y_2)$,

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 &= y_1, \\ a_{21}x_1 + a_{22}x_2 &= y_2. \end{aligned}$$

The object is to build sets $U_1, U_2, V_1, V_2 \subseteq \mathbb{F}$ such that $\tau(\mathbf{x}) \in V_1 \times V_2$ for many $\mathbf{x} \in U_1 \times U_2$.

¹⁸ The underlying asymptotic parameter is $|\mathbb{F}|$. I.e., the result has the form “for all functions $q(|\mathbb{F}|) = o(|\mathbb{F}|)$, ...”.

To begin with, define

$$\begin{aligned} U_1^* &= \{ia_{12} + ja_{22} : 1 \leq i, j \leq q^{1/3}\}, \\ U_2^* &= \{ka_{11} + ha_{21} : 1 \leq k, h \leq q^{1/3}\}. \end{aligned}$$

(These sets only have size $q^{2/3}$, but later we will extend them to sets of size q .) Note that if

$$ia_{12} + ja_{22} = i'a_{12} + j'a_{22}$$

has a nontrivial solution for some coefficients i, j, i', j' between 1 and $q^{1/3}$ then

$$ia_{12} + ja_{22} = 0$$

has a nontrivial solution for some i and j between $-q^{1/3}$ and $q^{1/3}$. When a_{12} and a_{22} are chosen randomly, the probability of such a solution occurring is upper bounded by $4q^{2/3}/|\mathbb{F}| \ll 1$, since $q = o(|\mathbb{F}|)$. Hence the ‘‘syntactically distinct’’ elements of U_1^* are also distinct elements of \mathbb{F} with high probability (and likewise for U_2^*) with high probability over the choice of a random matrix A .

Note that

$$\begin{aligned} &\{a_{11}x_1 + a_{12}x_2 : \mathbf{x} \in U_1^* \times U_2^*\} \\ &= \{a_{11}a_{12}(i+k) + a_{11}a_{22}j + a_{12}a_{21}h : 1 \leq i, j, k, h \leq q^{1/3}\} \\ &\supseteq \{a_{11}a_{12}i + a_{11}a_{22}j + a_{12}a_{21}h : 1 \leq i, j, h \leq q^{1/3}\}. \end{aligned}$$

Likewise

$$\begin{aligned} &\{a_{21}x_1 + a_{22}x_2 : \mathbf{x} \in U_1^* \times U_2^*\} \\ &= \{a_{21}a_{22}(j+h) + a_{21}a_{12}i + a_{11}a_{22}k : 1 \leq i, j, k, h \leq q^{1/3}\} \\ &\supseteq \{a_{21}a_{22}j + a_{21}a_{12}i + a_{11}a_{22}k : 1 \leq i, j, k \leq q^{1/3}\}. \end{aligned}$$

This motivates defining

$$\begin{aligned} V_1 &= \{a_{11}a_{12}i + a_{11}a_{22}j + a_{12}a_{21}h : 1 \leq i, j, h \leq q^{1/3}\}, \\ V_2 &= \{a_{21}a_{22}j + a_{21}a_{12}i + a_{11}a_{22}k : 1 \leq i, j, k \leq q^{1/3}\}. \end{aligned}$$

Then if $x_1 = ia_{12} + ja_{22} \in U_1^*$, $x_2 = ka_{11} + ha_{21} \in U_2^*$, we have

$$(a_{11}x_1 + a_{12}x_2, a_{21}x_1 + a_{22}x_2) \in V_1 \times V_2$$

if $i+k \leq q^{1/3}$ and $j+h \leq q^{1/3}$. The latter conditions are both implied if $i, j, k, h \leq (q^{1/3})/2$, so

$$|\{\mathbf{x} \in U_1^* \times U_2^* : \tau(\mathbf{x}) \in V_1 \times V_2\}| \geq (q^{1/3}/2)^4 = \Omega(q^{4/3}).$$

I.e., $\text{Cond}_\tau(q) \geq \Omega(q^{4/3})$.

As already pointed out, however, this construction has $|U_1^*| = |U_2^*| = q^{2/3} \ll q$, which suggests that even higher conductance could conceivably be shown. For this, our idea is to replace U_1^* and U_2^* by sets of the form $U_1^* + R$, $U_2^* + S$ where R and S are arithmetic progressions of length $q^{1/3}$. The step sizes for R and S will be different. Details follow.

Let $\mathbf{b} = (b_1, b_2)$ be the first column of A 's inverse. In other words

$$\begin{aligned} a_{11}b_1 + a_{12}b_2 &= 1, \\ a_{21}b_1 + a_{22}b_2 &= 0. \end{aligned}$$

We let

$$\begin{aligned} U_1 &= \{rb_1a_{11}a_{22} + ia_{12} + ja_{22} : 1 \leq r, i, j \leq q^{1/3}\}, \\ U_2 &= \{rb_2a_{11}a_{22} + ka_{11} + ha_{21} : 1 \leq r, k, h \leq q^{1/3}\}. \end{aligned}$$

A similar argument as above shows (using $q = o(|\mathbb{F}|)$) that syntactically distinct elements of U_i are also distinct as field elements with high probability over the choice of A .

Let

$$r, i, j, k, h$$

be arbitrary values between 1 and $(q^{1/3})/2$ and set

$$\begin{aligned} x_1 &= rb_1a_{11}a_{22} + ia_{12} + ja_{22} \\ x_2 &= rb_2a_{11}a_{22} + ia_{11} + ka_{21} \end{aligned}$$

so that $\mathbf{x} = (x_1, x_2) \in U_1 \times U_2$. (Note that we are using the “same r ” for both x_1 and x_2 ; hence, only a negligible fraction of vectors in $U_1 \times U_2$ have this form.)

If we let $\tau(\mathbf{x}) = \mathbf{y} = (y_1, y_2)$ then it is easy to see from the fact that \mathbf{b} is the first column of A 's inverse that

$$\begin{aligned} y_1 &= a_{11}a_{12}(i + k) + a_{11}a_{22}(j + r) + a_{12}a_{21}h, \\ y_2 &= a_{21}a_{22}(j + h) + a_{21}a_{12}i + a_{11}a_{22}k. \end{aligned}$$

Hence $\mathbf{y} \in V_1 \times V_2$ (the definition of the sets V_1, V_2 hasn't changed), by our stipulation that $1 \leq r, i, j, k, h \leq (q^{1/3})/2$. Hence

$$\text{Cond}_\tau(q) \geq ((q^{1/3})/2)^5 = \Omega(q^{5/3}) = \Omega(q^{2-1/(2w-1)})$$

as claimed.

We note that the constant $a_{11}a_{22}$, appearing in the definition of U_1, U_2 , is somewhat arbitrary and could be replaced with $a_{12}a_{21}$ or even $a_{11}a_{12}$ or $a_{21}a_{22}$ without (much) damaging the bound. (In more detail, choosing $a_{11}a_{12}$ or $a_{21}a_{22}$ means that r, i, j, k, h should be restricted to be less than $(q^{1/3})/3$ instead of $(q^{1/3})/2$, which is obviously a small loss.)

In the general case, the constant $a_{11}a_{22}$ is replaced by a collection of constants each of the form $a_{1j}a_{xy}$ with $x \neq 1$ and $y \neq j$, with the arithmetic progression $\{rb_ia_{11}a_{22} : 1 \leq r \leq q^{1/3}\}$ being replaced by a higher-dimensional arithmetic progression, with each such constant used for one dimension of the arithmetic progression.

General case. Let $\tau(\mathbf{x}) = A\mathbf{x}$ where $A \in \mathbb{F}^{w \times w}$ is a $w \times w$ matrix with entries in \mathbb{F} , and let a_{ij} be the (i, j) -th entry of A .

Let $m = w(w-1)^2 + w(w-1)/2$. Let B_1, \dots, B_w be subsets of $[w] \setminus \{1\} \times [w]$ such that $B_j \subseteq [w] \setminus \{1\} \times [w] \setminus \{j\}$ for each $1 \leq j \leq w$ and such that

$$\sum_j |B_j| = m - w(w-1) = w(w-1)^2 - w(w-1)/2.$$

We write $xy \in B_j$ as a shorthand for $(x, y) \in B_j$. [For $w = 2$ we had $B_1 = (2, 2)$, $B_2 = \emptyset$.]

Let $\mathbf{b} = (b_1, \dots, b_w)$ be the first column of A 's inverse.

We define

$$U_i = \left\{ \sum_j \sum_{xy \in B_j} r_{jxy} b_i a_{1j} a_{xy} + \sum_{i'} \sum_{j \neq i} h_{i'j} a_{i'j} : 1 \leq r_{jxy} \leq q^{1/m}, 1 \leq h_{i'j} \leq q^{1/m} \right\}$$

for $1 \leq i \leq w$. Then $|U_i| \leq q$. To argue distinctness of elements in U_i , note that the equation

$$\sum_j \sum_{xy \in B_j} r_{jxy} b_i a_{1j} a_{xy} + \sum_{i'} \sum_{j \neq i} h_{i'j} a_{i'j} = \sum_j \sum_{xy \in B_j} r_{jxy}^* b_i a_{1j} a_{xy} + \sum_{i'} \sum_{j \neq i} h_{i'j}^* a_{i'j}$$

has a nontrivial solution with coefficients $r_{jxy}, r_{jxy}^*, h_{i'j}, h_{i'j}^*$ between 1 and $q^{1/m}$ if and only if the equation

$$\sum_j \sum_{xy \in B_j} r_{jxy} b_i a_{1j} a_{xy} + \sum_{i'} \sum_{j \neq i} h_{i',j} a_{i'j} = 0$$

has a nontrivial solution with coefficients between $-q^{1/m}$ and $q^{1/m}$. By a union bound, the probability of such a solution existing for a randomly chosen matrix A is at most $(2q^{1/m})^{(m-w(w-1))+w(w-1)}/|\mathbb{F}| = 2^m q / |\mathbb{F}| = o(1)$ if $q = o(|\mathbb{F}|)$. Hence $|U_i| = q$ for all i with high probability over A .

We also define

$$V_i = \left\{ \sum_{j < j'} g_{jj'} a_{ij} a_{i'j'} + \sum_{i' \neq i} \sum_{j \neq j'} h_{i'jj'} a_{ij} a_{i'j'} : 1 \leq g_{jj'}, h_{i'jj'} \leq q^{1/m} \right\}.$$

Since the number of distinct choices for the $g_{jj'}$'s is $(q^{1/m})^{w(w-1)/2}$ and the number of distinct choices for the $h_{i'jj'}$'s is $(q^{1/m})^{(w-1)w(w-1)}$, we have $|V_i| \leq q$ by the fact that $m = w(w-1)/2 + w(w-1)^2$.

Let

$$\{r_{zxy} : xy \in B_z\}, \{h_{i'j}^i : j \neq i\}$$

be values such that $1 \leq r_{zxy} \leq (q^{1/m})/2$ for all z, x, y such that $xy \in B_z$ and such that $1 \leq h_{i'j}^i \leq (q^{1/m})/2$ for all i, i', j such that $j \neq i$. We note there are

$$\begin{aligned} (q^{1/m}/2)^{m-w(w-1)} \cdot (q^{1/m}/2)^{w^2(w-1)} &= \Omega(1) \cdot q^{[m-w(w-1)+w^2(w-1)]/m} \\ &= \Omega(1) \cdot q^{1+w(w-1)^2/m} \\ &= \Omega(1) \cdot q^{1+[m-(m-w(w-1)^2)]/m} \\ &= \Omega(1) \cdot q^{2-w(w-1)/2m} \\ &= \Omega(1) \cdot q^{2-1/(2w-1)} \end{aligned}$$

such tuples. Moreover, each such tuple naturally corresponds to a distinct element \mathbf{x} in the direct product $U_1 \times \dots \times U_w$, which is namely the vector \mathbf{x} such that

$$\mathbf{x}[i] = \sum_z \sum_{xy \in B_z} r_{zxy} b_i a_{1z} a_{xy} + \sum_{i'} \sum_{j \neq i} h_{i'j}^i a_{i'j}.$$

Then

$$\begin{aligned}
\tau(\mathbf{x})[i] &= (A\mathbf{x})[i] = \sum_k a_{ik}\mathbf{x}[k] \\
&= \sum_k a_{ik} \left(\sum_z \sum_{xy \in B_z} r_{zxy} b_k a_{1z} a_{xy} + \sum_{i'} \sum_{j \neq k} h_{i'j}^k a_{i'j} \right) \\
&= \sum_z \sum_{xy \in B_z} r_{zxy} a_{1z} a_{xy} \sum_k a_{ik} b_k + \sum_{i'} \sum_{j \neq k} h_{i'j}^k a_{ik} a_{i'j}
\end{aligned}$$

(Note the symbol $\sum_{j \neq k}$ in the second line does not mean the same as the symbol $\sum_{j \neq k}$ in the third line, since in the former sum k is already quantified, and the sum only takes place over j , whereas in the latter sum summation takes place over both j and k .) Here one can note that for each z , $xy \in B_z$,

$$r_{zxy} a_{1z} a_{xy} \sum_k a_{ik} b_k = \begin{cases} r_{zxy} a_{1z} a_{xy} & \text{if } i = 1 \\ 0 & \text{otherwise} \end{cases}$$

since $\mathbf{b} = (b_1, \dots, b_w)$ is the first column of A 's inverse. Thus we can write $\tau(\mathbf{x})[i]$ as a sum

$$\tau(\mathbf{x})[i] = \sum_{i' \neq i} \sum_{j \neq k} c_{ij'k} a_{ik} a_{i'j} + \sum_{k < j} c_{ikj} a_{ik} a_{ij}$$

for some coefficients $c_{iki'j}$, $i' \neq i$, $j \neq k$, as well as some coefficients c_{ikj} , $k < j$. More precisely, one has

$$c_{ikj} = h_{ij}^k + h_{ik}^j$$

for all $k < j$, and

$$c_{iki'j} = \begin{cases} h_{i'j}^k & \text{if } i \neq 1 \\ h_{i'j}^k & \text{if } i = 1, i'j \notin B_k \\ h_{i'j}^k + r_{ki'j} & \text{if } i = 1, i'j \in B_k \end{cases}$$

for all $i' \neq i$ and all $j \neq k$.

Since all the indices r_{zxy} and $h_{i'j}^k$ are between 1 and $(q^{1/m})/2$, we thus find that all the coefficients $c_{iki'j}$ as well as c_{ikj} are between 1 and $q^{1/m}$, and hence that $\tau(\mathbf{x})[i] \in V_i$ and that $\tau(\mathbf{x}) \in V_1 \times \dots \times V_w$. The claim on the conductance of τ then follows from the fact (argued above) that there are $\Omega(1) \cdot q^{2-1/(2w-1)}$ such \mathbf{x} 's.

All-but-one conductance. In the same vein, we could also show that all-but-one conductance is at least $q^{2+(w-2)/(w-1)(2w-3)}$ for generic linear permutations. For the sake of brevity, however, and since all-but-one conductance does not seem as intrinsically interesting as conductance, we skip this construction.

D Explicit constructions of permutations with low values of MaxColl

In this appendix we show an explicit construction of a permutation with small entry-wise random collision resistance of $O(w)$. Moreover we show this is essentially tight by proving entry-wise random collision resistance cannot be at least $\Omega(w)$ for any permutation π . (For the latter see Theorem 7 below.)

More specifically, our construction consists of an explicit permutation $\pi : \mathbb{F}^w \rightarrow \mathbb{F}^w$ such that $\text{MaxColl}(\pi) \leq 2w$, for a sufficiently large finite field \mathbb{F} . (The “sufficiently large” is needed in order for a matrix with certain generic features to exist. In fact a field of size $O(w^2)$ is already sufficient.) Note that if $\pi : \mathbb{F}^w \rightarrow \mathbb{F}^w$ is a linear permutation whose associated matrix has no zero entries, then $\text{MaxPreim}(\pi) = 1$, which is optimal. Hence, achieving small entry-wise random preimage resistance is trivial. Nonetheless, since we are generally interested in permutations that *simultaneously* achieve small entry-wise random preimage resistance and small entry-wise random collision resistance, we provide an analysis of the entry-wise random preimage resistance as well.

The Construction. Let A be a $w \times w$ invertible matrix over a finite field \mathbb{F} such that A has no zero entries and such that A^{-1} has no zero entries in its first column. By Cramer’s rule for the computation of the inverse, the latter condition is equivalent to $\det(A_{1k}) \neq 0$ for all $k \in [w]$, where A_{1k} denote the $(w-1) \times (w-1)$ matrix obtained by deleting the first row and k -th column of A . Clearly a random $w \times w$ matrix A over \mathbb{F} satisfies these requirements with high probability if \mathbb{F} is sufficiently large.

We also let

$$D = (d_2, \dots, d_w)$$

be an increasing sequence of positive integers, none of which is divisible by the characteristic p of \mathbb{F} , and such that $d_2 \geq 2$. (E.g., $(d_2, \dots, d_w) = (3, 5, \dots, 2w-1)$ if $p = 2$.) The matrix A and the sequence D will be the parameters of the construction.¹⁹

Let $\sigma : \mathbb{F}^w \rightarrow \mathbb{F}^w$ be defined by

$$\sigma(\mathbf{x}) = A\mathbf{x}$$

for $\mathbf{x} = (x_1, \dots, x_w) \in \mathbb{F}^w$. Also let $\eta : \mathbb{F}^2 \rightarrow \mathbb{F}^w$ be the Feistel-round-like transformation given by

$$\eta(\mathbf{x})[i] = \begin{cases} x_i + \sum_{j=2}^w x_j^{d_j} & \text{if } i = 1, \\ x_i & \text{if } i \neq 1. \end{cases}$$

We note that η has a straightforward inverse:

$$\eta^{-1}(\mathbf{x})[i] = \begin{cases} x_i - \sum_{j=2}^w x_j^{d_j} & \text{if } i = 1, \\ x_i & \text{if } i \neq 1. \end{cases}$$

We define the permutation $\pi[A, D] : \mathbb{F}^w \rightarrow \mathbb{F}^w$ as

$$\pi[A, D] = \sigma^{-1} \circ \eta \circ \sigma$$

with inverse

$$\pi[A, D]^{-1} = \sigma^{-1} \circ \eta^{-1} \circ \sigma.$$

¹⁹ We believe that the restriction $d_2 \geq 2$ can be replaced with $d_2 \geq 1$ with some extra technical (but annoying) legwork. If this restriction were lifted one could use $(d_2, \dots, d_w) = (1, 3, \dots, 2w-3)$ for $p = 2$.

Let

$$z(\mathbf{x}) = \sum_{j=2}^w x_j^{d_j}$$

be the polynomial appearing in the definition of η , and let $Z : \mathbb{F}^w \rightarrow \mathbb{F}^w$ be defined by

$$Z(\mathbf{x}) = \eta(\mathbf{x}) - \mathbf{x}.$$

Then $Z(\mathbf{x})[1] = z(\mathbf{x})$ and $Z(\mathbf{x})[i] = 0$ for $i \geq 2$. Moreover we can write $\eta(\mathbf{x})$ as $Z(\mathbf{x}) + \mathbf{x}$. Thus

$$(\eta \circ \sigma)(\mathbf{x}) = A\mathbf{x} + Z(A\mathbf{x}).$$

Thus if \mathbf{a}_i^T denotes the i -th row of A and \mathbf{b}_i^T denotes the i -th row of A^{-1} , we have

$$\begin{aligned} (\sigma^{-1} \circ \eta \circ \sigma)(\mathbf{x})[i] &= (A^{-1}A\mathbf{x})[i] + A^{-1}Z(A\mathbf{x})[i] \\ &= x_i + \mathbf{b}_i^T Z(A\mathbf{x}) \\ &= x_i + \mathbf{b}_i[1]z(A\mathbf{x}) \\ &= x_i + \mathbf{b}_i[1] \sum_{j=2}^w (\mathbf{a}_j^T \mathbf{x})^{d_j} \end{aligned}$$

for all $1 \leq i \leq w$.

The analysis. We start with entry-wise random collision resistance (the more interesting of the two properties), and follow with entry-wise random preimage resistance.

Theorem 5. $\text{MaxColl}_{\pi[A,D]} \leq d_w$ for A, D and $\pi[A, D]$ as described above.

Proof. Fix values $k, \ell \in [w]$. Let

$$\begin{aligned} \mathbf{x}' &= (x'_1, \dots, x'_{k-1}, x'_{k+1}, \dots, x'_k) \\ \mathbf{x}'' &= (x''_1, \dots, x''_{k-1}, x''_{k+1}, \dots, x''_k) \end{aligned}$$

be two arbitrary $(w-1)$ -dimensional vectors such that $\mathbf{x}' \neq \mathbf{x}''$. We will also let \mathbf{x}, \mathbf{x}' stand for functions from \mathbb{F} to \mathbb{F}^w by setting

$$\begin{aligned} \mathbf{x}'(x_k) &= (x'_1, \dots, x'_{k-1}, x_k, x'_{k+1}, \dots, x'_k) \\ \mathbf{x}''(x_k) &= (x''_1, \dots, x''_{k-1}, x_k, x''_{k+1}, \dots, x''_k) \end{aligned}$$

for $x_k \in \mathbb{F}$. Then

$$\begin{aligned} P(x_k) &:= (\sigma^{-1} \circ \eta \circ \sigma)(\mathbf{x}'(x_k))[\ell] - (\sigma^{-1} \circ \eta \circ \sigma)(\mathbf{x}''(x_k))[\ell] \\ &= \begin{cases} x'_\ell - x''_\ell + \mathbf{b}_\ell[1](z(A\mathbf{x}'(x_k)) - z(A\mathbf{x}''(x_k))) & \text{if } \ell \neq k \\ \mathbf{b}_\ell[1](z(A\mathbf{x}'(x_k)) - z(A\mathbf{x}''(x_k))) & \text{if } \ell = k \end{cases} \end{aligned}$$

is a polynomial of degree (at most) d_w in x_k .

We assume next by contradiction that $P(x_k) = 0$ is the zero polynomial. We consider the cases $\ell = k$ and $\ell \neq k$ separately.

For $\ell \neq k$ we have

$$P(x_k) = x'_\ell - x''_\ell + \mathbf{b}_\ell[1] \left(\sum_{j=2}^w (\mathbf{a}_j^T \mathbf{x}'(x_k))^{d_j} - \sum_{j=2}^w (\mathbf{a}_j^T \mathbf{x}''(x_k))^{d_j} \right).$$

We recall that $\mathbf{b}_i[1] \neq 0$ for all i by our initial assumptions on A . Writing $\mathbf{a}_{j \setminus k}^T$ for the j -th row of A with the k -th entry removed we can write

$$P(x_k) = x'_\ell - x''_\ell + \mathbf{b}_\ell[1] \left(\sum_{j=2}^w (\mathbf{a}_{j \setminus k}^T \mathbf{x}' + \mathbf{a}_j[k]x_k)^{d_j} - \sum_{j=2}^w (\mathbf{a}_{j \setminus k}^T \mathbf{x}'' + \mathbf{a}_j[k]x_k)^{d_j} \right).$$

Since $P(x_k) = 0$ by assumption, the coefficient of $x_k^{d_w-1}$ in $P(x_k)$ must be zero, i.e.,

$$d_w(\mathbf{a}_{w \setminus k}^T \mathbf{x}') \mathbf{a}_w[k]^{d_w-1} = d_w(\mathbf{a}_{w \setminus k}^T \mathbf{x}'') \mathbf{a}_w[k]^{d_w-1}.$$

Since the characteristic of \mathbb{F} doesn't divide d_w , and since $\mathbf{a}_w[k] \neq 0$, this implies

$$\mathbf{a}_{w \setminus k}^T \mathbf{x}' = \mathbf{a}_{w \setminus k}^T \mathbf{x}''$$

whence

$$P(x_k) = x'_\ell - x''_\ell + \mathbf{b}_\ell[1] \left(\sum_{j=2}^{w-1} (\mathbf{a}_{j \setminus k}^T \mathbf{x}' + \mathbf{a}_j[k]x_k)^{d_j} - \sum_{j=2}^{w-1} (\mathbf{a}_{j \setminus k}^T \mathbf{x}'' + \mathbf{a}_j[k]x_k)^{d_j} \right).$$

Since the coefficient of $x_k^{d_{w-1}-1}$ in $P(x_k)$ must also be zero, we then find

$$d_{w-1}(\mathbf{a}_{w-1 \setminus k}^T \mathbf{x}') \mathbf{a}_{w-1}[k]^{d_{w-1}-1} = d_{w-1}(\mathbf{a}_{w-1 \setminus k}^T \mathbf{x}'') \mathbf{a}_{w-1}[k]^{d_{w-1}-1}$$

and (since $p \nmid d_{w-1}$, $\mathbf{a}_{w-1}[k] \neq 0$)

$$\mathbf{a}_{w-1 \setminus k}^T \mathbf{x}' = \mathbf{a}_{w-1 \setminus k}^T \mathbf{x}''$$

whence

$$P(x_k) = x'_\ell - x''_\ell + \mathbf{b}_\ell[1] \left(\sum_{j=2}^{w-2} (\mathbf{a}_{j \setminus k}^T \mathbf{x}' + \mathbf{a}_j[k]x_k)^{d_j} - \sum_{j=2}^{w-2} (\mathbf{a}_{j \setminus k}^T \mathbf{x}'' + \mathbf{a}_j[k]x_k)^{d_j} \right).$$

Continuing like this, we eventually find that

$$P(x_k) = x'_\ell - x''_\ell + \mathbf{b}_1[\ell] \left((\mathbf{a}_{2 \setminus k}^T \mathbf{x}' + \mathbf{a}_2[k]x_k)^{d_2} - (\mathbf{a}_{2 \setminus k}^T \mathbf{x}'' + \mathbf{a}_2[k]x_k)^{d_2} \right)$$

and that

$$\mathbf{a}_{2 \setminus k}^T \mathbf{x}' = \mathbf{a}_{2 \setminus k}^T \mathbf{x}''$$

by considering the term $x_k^{d_2-1}$ in $P(x_k)$. (Here we use the fact that $d_2 > 1$.) Hence

$$\mathbf{a}_{j \setminus k}^T \mathbf{x}' = \mathbf{a}_{j \setminus k}^T \mathbf{x}'' \tag{18}$$

for all $2 \leq j \leq w$, i.e.,

$$A_{1k}\mathbf{x}' = A_{1k}\mathbf{x}'',$$

but this implies $\mathbf{x}' = \mathbf{x}''$, a contradiction, since $\det(A_{1k}) \neq 0$.

To establish a contradiction for the case $\ell = k$ one can reason analogously. In fact this case is simpler, since the only thing that changes is that $P(x_k)$ doesn't contain the term $x'_\ell - x''_\ell$. (As a technical comment, one can observe that for the case $\ell = k$ it suffices to have $d_2 \geq 1$ instead of $d_2 \geq 2$. For in the absence of the constant term $x'_\ell - x''_\ell$ one arrives at equation (18) even if $d_2 - 1 = 0$.)

Having obtained a contradiction in both cases (being $\ell = k$ and $\ell \neq k$), we can conclude that $P(x_k)$ is a nonzero polynomial. Since $P(x_k)$ has degree at most d_w , however, this implies that

$$(\sigma^{-1} \circ \eta \circ \sigma)(\mathbf{x}'(x_k)) = (\sigma^{-1} \circ \eta \circ \sigma)(\mathbf{x}''(x_k))$$

for at most d_w values of $x_k \in \mathbb{F}$. This precisely implies that $\text{MaxColl}_{\pi[A,D]} \leq d_w$ since k, ℓ as well as \mathbf{x}' and \mathbf{x}'' were arbitrary. \square

Theorem 6. $\text{MaxPreim}_{\pi[A,D]} \leq d_w$ for A, D and $\pi[A, D]$ as described above.

Proof. Fix values $k, \ell \in [w]$ as well as a vector

$$\mathbf{x}' = (x'_1, \dots, x'_{k-1}, x'_{k+1}, \dots, x'_w) \in \mathbb{F}^{w-1}.$$

We also set

$$\mathbf{x}'(x_k) = (x'_1, \dots, x'_{k-1}, x_k, x'_{k+1}, \dots, x'_w) \in \mathbb{F}^{w-1}$$

as in the proof of Theorem 5. Then

$$\pi[A, D](\mathbf{x}'(x_k))[\ell] = \begin{cases} x'_\ell + \mathbf{b}_\ell[1]z(\mathbf{A}\mathbf{x}'(x_k)) & \text{if } \ell \neq k, \\ x_k + \mathbf{b}_k[1]z(\mathbf{A}\mathbf{x}'(x_k)) & \text{if } \ell = k. \end{cases}$$

In either case (whether $\ell = k$ or $\ell \neq k$) $\pi[A, D](\mathbf{x}'(x_k))[\ell]$ is a nonzero polynomial in x_k of degree d_w . (The nonzeroness can be seen by focusing on the term of degree d_w , which has nonzero coefficient of $\mathbf{a}_w^T[k]^{d_w}$.) Hence the probability that $\pi[A, D](\mathbf{x}'(x_k))$ equals any particular value, over random choice of x_k , is at most d_w . \square

Optimality. We can also prove that every permutation π has $\text{MaxColl}(\pi) = \Omega(w)$. Specifically, we can prove the following theorem:

Theorem 7. For all sufficiently large $|\mathbb{F}|$, any permutation $\pi : |\mathbb{F}|^w \rightarrow |\mathbb{F}|^w$ has $\text{MaxColl}(\pi) \geq (w-1)/2$.

For this universal lower bound, the algebraic structure of \mathbb{F} obviously doesn't matter. Here \mathbb{F} can be considered as a placeholder for an arbitrary finite set.

In fact we will prove something a bit stronger than Theorem 7, being namely that

$$\text{MaxColl}_{j,h}(\pi) \geq (w-1)/2$$

for all $j, h \in [w]$ (and sufficiently large $|\mathbb{F}|$), where

$$\text{MaxColl}_{j,h}(\pi) := \max_{\mathbf{x}, \mathbf{x}'} |\{x \in \mathbb{F} : \pi_{j,h}^{\mathbf{x}}(x) = \pi_{j,h}^{\mathbf{x}'}(x)\}|$$

where $\pi_{j,h}^{\mathbf{x}}(x)$ is as defined in Section 2. We can obviously assume that $j = h = 1$.

Let $q = |\mathbb{F}|$. We will analyze the family of functions $\{\pi_{1,1}^{\mathbf{x}} : \mathbf{x} \in \mathbb{F}^w, \mathbf{x}[1] = 0\}$. This family is indexed by q^{w-1} different \mathbf{x} 's, which we will denote $\mathbf{x}_1, \dots, \mathbf{x}_{q^{w-1}}$.

Each function $\pi_{1,1}^{\mathbf{x}_i} : \mathbb{F} \rightarrow \mathbb{F}$ can be encoded as a $q \times q$ 0-1 matrix β_i by setting

$$\beta_i[x, y] = \begin{cases} 1 & \text{if } \pi_{1,1}^{\mathbf{x}_i}(x) = y, \\ 0 & \text{otherwise.} \end{cases}$$

Then $\text{MaxColl}_{1,1}(\pi)$ can be equivalently defined as the maximum dot product (over \mathbb{R}) between any two (distinct) β_i 's. Thus, Theorem 7 will follow by the following proposition (which further relaxes²⁰ the problem by removing the requirement that the β_i 's encode functions):

Proposition 1. *Let $B = \{\beta_1, \dots, \beta_{q^{w-1}}\}$ be distinct 0-1 matrices of size $q \times q$. Then if q is sufficiently large compared to w there exists some $i \neq i'$ such that $\langle \beta_i, \beta_{i'} \rangle \geq (w-1)/2$ where $\langle x, y \rangle$ denotes the real-valued dot product of x and y .*

Proof. A family B that minimizes the maximum value of $\langle \beta_i, \beta_j \rangle$, $i \neq j$ can obviously be chosen such that deleting any entry of '1' from any vector β_i produces another vector in B (or else replace β_i with the new vector). We will thus assume that B has this form. Then

$$\max_{i \neq j} \langle \beta_i, \beta_j \rangle \geq \max_i \|\beta_i\|_1 - 1$$

where $\|\cdot\|_1$ denotes the 1-norm. (In fact this inequality holds with equality, but we only need the inequality.)

Let $d = \max_i \|\beta_i\|_1$. Each element in B can then be specified by a sequence of d choices, where each choice has arity $q^2 + 1$ (where to put the next '1', and if it should be a '1' at all), so

$$q^{w-1} = |B| \leq (1 + q^2)^d$$

which obviously implies that $d \geq (w-1)/2$ for sufficiently large q , for fixed w . □

E Proof of Theorem 1

In this section we provide the proof of Theorem 1. More precisely, we show that the simulator discussed in Section 6 and formally given by the pseudocode in Figures 3–7 fulfills the claims of Theorem 1.

On a high level, the formal structure of our indifferentiability proof(s) follows [1] rather closely. On the other hand our simulator shares more in common with the simulator of [18] (which also employs Seurin's termination paradigm) so our proof is also descended from [18]. Moreover our proof utilizes (with some refinements) the *randomness mapping* technique pioneered by [18], and described in more detail below.

²⁰ Well, not quite: Proposition 1 introduces the new requirement that β_i 's be distinct, but this is obviously without loss of generality since otherwise $\text{MaxColl}_{1,1}(\pi) = \text{MaxColl}(\pi) = |\mathbb{F}|$.

E.1 Overview

An indistinguishability theorem encompasses at least three claims: an upper bound on the simulator running time, an upper bound on the simulator query complexity, and lastly the indistinguishability of the real and simulated worlds. Here (in the proof overview) we focus on the latter aspect. Lemmas 51, 52 in subsection E.5 establish upper bounds on the simulator efficiency and query complexity respectively.

The indistinguishability argument uses a sequence of five games G_1, G_2, \dots, G_5 . The games are specified by the pseudocode in Figs. 3–9. More precisely, the pseudocode in Figs. 3–7 collectively describes games G_1, G_2 and G_3 (the differences between the various games being highlighted in red, and also by comment lines), whereas Fig. 8–9 depicts games G_4 and G_5 . Game G_1 is the simulated world and G_5 is the real world.

Following [1] the simulator explicitly aborts (**‘abort’**) if runs into an unexpected state. The distinguisher is consequently notified and, in such a case, knows that it is in the simulated world (since game G_5 , the real world, doesn’t abort).

As stated, game G_1 is the simulated world. The interface available to the distinguisher consists of the **public** procedures, of which the simulator controls all but Z and Z^{-1} (which implement \mathcal{Z} and \mathcal{Z}^{-1}). The procedures Z/Z^{-1} maintain a partial permutation table P_Z that records which queries have already been made to these procedures either by the adversary or by the simulator. In game G_1 the table P_Z serves no effective purpose, but this changes in later games. Moreover Z and Z^{-1} are the only procedures to access the permutation table P_Z in game G_1 (as should be, since the simulator shouldn’t know which queries to Z/Z^{-1} the distinguisher has already made).

Game G_2 contains only one change from game G_1 , in the procedure `CheckZ` used by the simulator (see Fig. 3). In G_1 `CheckZ(x, y)` calls $Z(x)$ to see if $Z(x) = y$. In game G_2 , however, `CheckZ(x, y)` answers instead according to the table P_Z of values already queried to Z/Z^{-1} by either the distinguisher or the simulator, and returns **false** if the relevant table entry is blank, without calling Z . Thus `CheckZ` may return a “false negative” in game G_2 . The G_1 – G_2 game transition is inspired by a similar transition in [18]. Our transition is substantially simpler, however, since (and following [1]) we keep Z as a permutation instead of replacing it by a “two-way random function” [18].

The fact that G_1 and G_2 are hard to distinguish for a q -query distinguisher is proven in Lemma 40, Section E.4. The proof is relatively straightforward but relies on having an upper bound on the number of times that `CheckZ` is called in G_1 and G_2 . This upper bound is given by Lemma 23 at the end of Section E.2.

Game G_3 adds a number of **abort** conditions over game G_2 . (These abort conditions could not be present in G_2 because they involve the simulator reading values in P_Z , or because they involve abort conditions placed in Z/Z^{-1} , which should of course not abort in the real simulated world.) In fact, one doesn’t need to upper bound the distinguishability of G_2 from G_3 : one simply observes that the pair (G_3, G_5) is strictly *easier* to distinguish than the pair (G_2, G_5) (as abortions never occur in G_5 —we refer to Lemma 41 for a formal statement and proof). Hence it suffices to show that G_3 and G_5 are hard to distinguish.

The transition from G_3 to G_5 is analyzed via G_4 . In G_4 the procedures Z and Z^{-1} are changed: instead of using the random tape p_Z these procedures now use (some subset of) the random tapes p_{A_1}, \dots, p_{L_w} and the “real” substitution-permutation network; see Figs. 8–9. There is another minor technicality: the random tape p_{T_j} is renamed q_{T_j} in G_4 , in order to clarify subsequent simultaneous discussion of G_3 and G_4 .

The proof’s key transition is the transition from G_3 to G_4 . For this a *randomness mapping* argument is used, following [18]. We borrow the idea of *footprints* from [1], but we manage to eschew their “execution trees”.

The randomness mapping argument is sketched in some detail after Lemma 41 in Section E.4. For completeness, however, we provide a briefer sketch here.

A *footprint* is that subset of the random tapes (plus its contents) which is actually accessed during an execution²¹. If we change portions of a random tape outside of its footprint, the execution doesn’t change. The set of possible footprints is a function of the distinguisher (which is presumed fixed and deterministic, wlog). A moment’s thought reveals that two distinct footprints (with respect to the same distinguisher) are never compatible: they always overlap in at least one conflicting entry.

A footprint is said to be *good* if the execution doesn’t **abort** for that footprint. Essentially, the randomness mapping argument maps good footprints of G_3 to good footprints of G_4 by the following rule: the content of the table T_j at the end of the G_3 -execution becomes the G_4 -footprint for p_{T_j} (actually, q_{T_j} in G_4). One must prove that points in the image of this mapping are really footprints for G_4 , that the mapping is injective (maps distinct footprints to distinct footprints), and that a G_3 -execution with a good footprint looks identical from the distinguisher’s perspective to a G_4 -execution run with the image footprint. For all this to be useful one must also show that the “probability mass” of good footprints in G_3 (i.e., the probability of obtaining a good footprints, taken over the uniform choice of random tapes) is high, and that a good G_3 footprint has roughly the same probability of occurring as its image in G_4 . Once all these elements in place it is easily seen that G_3 and G_4 are hard to distinguish.

We note that the probability of obtaining a good footprint in G_3 is exactly the probability of not aborting in G_3 . Hence for the randomness mapping argument to be effective one must, in particular, show that G_3 aborts with low probability. Section E.3 is devoted to upper bounding the latter probability. For more details we refer to the afore-mentioned material in Section E.4.

For the transition from G_4 to G_5 , we first argue that queries in G_4 are all answered by the value found in the underlying random tape (i.e., query $T(j, x)$ is answered by $q_{T_j}(x)$, and so on) as long as **abort** does not occur. (See Lemma 49.) Hence the only effective difference between G_4 and G_5 is the possibility that G_4 aborts. The probability that G_4 aborts can be upper bounded from the probability that G_3 aborts and from the randomness map. This is the argument adopted by [1]. However, we note that a more careful additive accounting of the transition probabilities causes the probability of abortion in G_4 to cancel out entirely from the final distinguishing upper bound (see the computation in the proof of Lemma 50) which makes the latter argument superfluous, and which also saves us a factor of two²² in the final indistinguishability bound over [1, 18, 19]. This is another (small) technical contribution of our proof.

PROOF ORGANIZATION. A prerequisite for much of the analysis is to have upper bounds on the operations of various types (e.g., path completions, queries) carried out by the simulator. The proof starts by subsection E.2, where we establish the most important such upper bounds.

Another crucial component of the analysis is to upper bound the probability of bad executions (i.e., executions in which the simulator aborts) in game G_3 . This is the topic of subsection E.3.

²¹ An *execution* consists of one run of the distinguishing experiment from beginning to end.

²² In fact, Andreeva et al. [1], thanks to their use of footprints, save a factor of two in the transition from G_3 to G_4 over the traditional approach of [18, 19]. But Andreeva et al. lose this factor again in the transition from G_4 to G_5 by double-counting abort probabilities. In our approach the saved factor of two is conserved all the way through.

The high-level assembly of these prior low-level results occurs in subsection E.4, where we separately analyze each of the game transitions G_1 – G_2 , G_3 – G_4 and G_4 – G_5 and, in particular, carry out the randomness mapping argument (for the transition from G_3 to G_4).

E.2 Simulator Efficiency

LAST MINUTE REMINDERS. We have $\alpha(q) = (2q)^w$ if XtraMiddleRnd is off, $\alpha(q) = \text{Cond}_\tau(2q)$ if XtraMiddleRnd is on; $\beta(q) = (q + \alpha(q))^w$ if XtraOuterRnd is off, $\beta(q) = \text{Cond}_\nu(q + \alpha(q))$ if XtraOuterRnd is on. We recall that p_T, p_T^- are a pair of random tapes encoding a random permutation from $\{0, 1\}^n$ to $\{0, 1\}^n$ for all $T \in \{A, \dots, L\}$ whereas p_Z, p_Z^- are a pair of random tapes encoding a random permutation from $\{0, 1\}^{wn}$ to $\{0, 1\}^{wn}$.

An *execution* refers to the entire run of the distinguishing experiment; a *q-query execution* is an execution in which the distinguisher makes at most q queries; a *query cycle* is that time portion of an execution between when the distinguisher makes a query and when that query is answered.

For $T \in \{A, \dots, L\}$ we write $T(\mathbf{x})$ to denote the vector whose j -th entry is $T_j(\mathbf{x}[j])$, if those w values are all defined, where T_j refers to the table maintained by the simulator; $T(\mathbf{x}) = \perp$ if $T_j(\mathbf{x}[j]) = \perp$ for any j . Similar conventions hold for $T^{-1}(\mathbf{y})$ with respect to the tables $T_j^{-1}, j \in [w]$.

For the time being, q is the number of distinguisher queries. The distinction between q and q_0 is elucidated in subsection E.4.

Lemma 1. *Each pair of tables $T_j, T_j^{-1}, T \in \{F, G, I, D, J, B, A, C, K, E, L, H\}, 1 \leq j \leq w$, represents a partial permutation at all points in the execution of G_1, G_2 and G_3 outside of calls to SetTable. Moreover, values in these tables are never overwritten.*

Proof. The lemma transparently follows from the fact that these tables are initialized to \perp and from the fact that the only function to write to the tables T_j, T_j^{-1} is SetTable, which is set to abort the game rather than overwrite. \square

Lemma 2. *The pair of tables P_Z, P_Z^{-1} represents a partial permutation at all points in the execution of G_1, G_2 and G_3 outside of calls to SetTable. Moreover P_Z remains consistent with p_Z .*

Proof. The lemma directly follows from the fact that table P_Z, P_Z^{-1} are presumed initialized to \perp and from the fact that the tables P_Z, P_Z^{-1} are only modified by the ReadTape calls in Z and Z^{-1} . \square

Definition 2. *For G_1, G_2 and G_3 we define*

$$\begin{aligned} \text{Index}_{G_j} &= \text{ToBeAssigned}_{G_j} \cup \{x : G_j(x) \neq \perp\}, \\ \text{Index}_{H_j}^- &= \text{ToBeAssigned}_{H_j}^- \cup \{y : H_j^{-1}(y) \neq \perp\}, \\ \text{Index}_{A_j} &= \text{ToBeAssigned}_{A_j} \cup \{x : A_j(x) \neq \perp\}, \\ \text{Index}_{A_j}^- &= \text{ToBeAssigned}_{A_j}^- \cup \{y : A_j^{-1}(y) \neq \perp\}, \\ \text{Index}_{B_j}^- &= \text{ToBeAssigned}_{B_j}^- \cup \{y : B_j^{-1}(y) \neq \perp\}, \\ \text{Index}_{C_j} &= \text{ToBeAssigned}_{C_j} \cup \{x : C_j(x) \neq \perp\}, \end{aligned}$$

at all time points in the execution.

As per the aliases introduced in Section 6 $\text{Index}_{U_j}^-$ will stand for $\text{Index}_{A_j}^-$ if XtraOuterRnd is off, for $\text{Index}_{B_j}^-$ otherwise, and Index_{V_j} will stand for Index_{A_j} if XtraOuterRnd is off, for Index_{C_j} otherwise.

Lemma 3. *The sets $\text{Index}_{T_j}^{(-)}$ are monotone increasing in G_1 , G_2 and G_3 : elements are never removed from these sets.*

Proof. The set $\{x : G_j(x) \neq \perp\}$ is monotone increasing by Lemma 1. Moreover x is only removed from $\text{ToBeAssigned}_{G_j}$ in AdaptLeft after $G_j(x)$ has been read from p_{G_j} (see Fig. 6). Similar statements apply to $\text{Index}_{H_j}^-$, $\text{Index}_{U_j}^-$, and Index_{V_j} . \square

We will refer to the procedures RequestG , RequestU^{-1} , AdaptLeft as well as the public procedures G and U^{-1} (A^{-1} or B^{-1} depending on XtraMiddleRnd) as *left-hand* procedures and to the procedures RequestH^{-1} , RequestV , AdaptRight as well as the public procedures H^{-1} and V (A or C depending on XtraMiddleRnd) as *right-hand* procedures. Moreover, the data structures $\text{ToBeAssigned}_{G_j}$, $\text{ToBeAssigned}_{U_j}^-$, ToBeAdapted_D are *left-hand* (data structures), while $\text{ToBeAssigned}_{V_j}$, $\text{ToBeAssigned}_{H_j}^-$, ToBeAdapted_E are *right-hand* (data structures).

A query made by the distinguisher is *left-hand* if it is a query to G or U^{-1} it is *right-hand* if it is a query to H^{-1} or V .²³ Other adversarial queries are *neutral*. Finally a left-hand distinguisher query gives rise to a “left-hand” query cycle, a right-hand distinguisher query gives rise to a “right-hand” query cycle, etc.

The following lemma holds for G_1 , G_2 and G_3 .

Lemma 4. *Left-hand data structures are only modified by left-hand procedures, right-hand data structures are only modified by right-hand procedures. Right-hand procedures are never evaluated during left-hand or neutral query cycles and left-hand procedures are never evaluated during right-hand or neutral query cycles.*

Proof. These claims are all syntactically clear from the code. In particular, one can observe that the only closed systems or recursive calls that occurs are between G or U^{-1} , which is a left-hand closed system, and between H^{-1} or V , which is a right-hand closed system. \square

Lemma 5. *Outside a left-hand query cycle, the left-hand data sets $\text{ToBeAssigned}_{G_j}$, $\text{ToBeAssigned}_{U_j}^-$, ToBeAdapted_D are empty in the execution of G_1 , G_2 and G_3 . A symmetric statement holds for right-hand data.*

Proof. $\text{ToBeAssigned}_{G_j}$, $\text{ToBeAssigned}_{U_j}^-$, ToBeAdapted_D are empty when AdaptLeft returns without abort. Moreover, these sets are initially empty and never modified in a right-hand or neutral query cycle by Lemma 4. \square

Lemma 6. *For any $j \in [w]$, the sets $\{y : G_j(y) \neq \perp\}$ and $\text{ToBeAssigned}_{G_j}$ are disjoint at all points in the execution of G_1 , G_2 and G_3 except for inside the first **forall** loop in AdaptLeft .*

Similar statements hold for $\text{ToBeAssigned}_{U_j}^-$, $\text{ToBeAssigned}_{V_j}$ and $\text{ToBeAssigned}_{H_j}^-$ (for the latter two, with respect to AdaptRight).

²³ We observe one last time that U^{-1} and V aren't actual procedures, but stand for A^{-1} or B^{-1} in the first case and for A or C in the second case. On the other hand, RequestU^{-1} , RequestV as well as $\text{BlockRequestU}^{-1}$, BlockRequestV , BlockU^{-1} and BlockV are not aliases, but the actual names of pseudocode procedures.

Proof. We focus on the table G_j . Other tables are treated similarly.

Outside a left-hand query cycle $\text{ToBeAssigned}_{G_j}$ is empty by Lemma 5.

Inside a left-hand query, before AdaptLeft is called, G_j is unmodified and x is added to $\text{ToBeAssigned}_{U_j}$ only if $G_j(x) = \perp$.

Finally, $\text{ToBeAssigned}_{G_j}$ is empty after the first **forall** loop in AdaptLeft . \square

Lemma 7. *For any q -query execution of G_1, G_2 or G_3 , the number of times lines 17–18 of RequestG are executed plus the number of times lines 17–18 of RequestH^{-1} are executed is at most q . Moreover, the number of times line 20 of RequestG is executed plus the number of times line 20 of RequestH^{-1} is executed is at most q . (Nb: line 20 is the last line of $\text{RequestG}, \text{RequestH}^{-1}$.)*

Proof. Both statements easily from the fact that when line 16 of RequestG^{-1} or RequestH is executed, NumOuter is increased by 1 and from the fact that the simulator aborts when NumOuter exceeds q . \square

Lemma 8. *For any $j \in [w]$ and for any q -query execution of G_1, G_2 or G_3 the number of calls made to $\text{RequestU}^{-1}(j, \cdot)$ plus the number of calls made to $B(j, \cdot)$ is at most $2q$ if XtraMiddleRnd is on. A symmetric statement holds for $C^{-1}(j, \cdot)$ and $\text{RequestV}(j, \cdot)$.*

If XtraMiddleRnd is off, the number of calls made to $\text{RequestU}^{-1}(j, \cdot)$ plus the number of calls made to $\text{RequestV}(j, \cdot)$ is at most $2q$.

Proof. Note that $\text{RequestU}^{-1}(j, \cdot)$ is called only by $U^{-1}(j, \cdot)$ or by $\text{BlockRequestU}^{-1}$. Moreover $U^{-1}(j, \cdot)$ (which is $A^{-1}(j, \cdot)$ or $B^{-1}(j, \cdot)$, depending) is only called by the distinguisher.

If XtraMiddleRnd is on: The number of calls to $B(j, \cdot)$ made by the distinguisher plus the number of calls to $B^{-1}(j, \cdot)$ is at most q (as just noted, $B^{-1}(j, \cdot)$ is only called by the distinguisher). Hence, the number of calls to $B(j, \cdot)$ made by the distinguisher plus the number of times $\text{RequestU}^{-1}(j, \cdot)$ is called by $B^{-1}(j, \cdot)$ is at most q .

It remains to account for non-distinguisher calls to $B(j, \cdot)$ as well as for the number of times $\text{BlockRequestU}^{-1}$ calls $\text{RequestU}^{-1}(j, \cdot)$, i.e., the number of times $\text{BlockRequestU}^{-1}$ is called. However non-distinguisher calls to $B(j, \cdot)$ only occur in BlockB , called only by the function LeftToMiddle , itself called only on line 17 of RequestH^{-1} . Moreover $\text{BlockRequestU}^{-1}$ is only called on line 20 of RequestG . Hence the sum of these two types of calls is at most q by Lemma 7.

If XtraMiddleRnd is off: The number of times $\text{RequestU}^{-1}(j, \cdot)$ is called by $A^{-1}(j, \cdot)$ plus the number of times $\text{RequestV}(j, \cdot)$ is called by $A(j, \cdot)$ is at most q , the number of distinguisher queries. Indeed, A and A^{-1} are not called by the simulator.

It remains again to account for calls to $\text{RequestU}^{-1}(j, \cdot)$ made by $\text{BlockRequestU}^{-1}$ and for calls to $\text{RequestV}(j, \cdot)$ made by BlockRequestV . But this number of calls is at most q by Lemma 7, as $\text{BlockRequestU}^{-1}$ only appears on line 20 of RequestG and BlockRequestV only appears on line 20 of RequestV . \square

Lemma 9. *For any $j \in \{1, \dots, w\}$ and for any q -query execution of G_1, G_2 or G_3 the size of $\text{Index}_{U_j}^{-1}$ is at most $2q$ and the size of U_j is at most $2q$.*

The same holds for Index_{V_j} and V_j .

Proof. The size of U_j is no greater than the size of $\text{Index}_{U_j}^{-1}$, so it suffices to prove the statement for Index_{U_j} .

If XtraMiddleRnd is on: The only functions to possibly increase the size of $\text{Index}_{U_j}^{-1}$ are $B(j, \cdot)$ and $\text{RequestU}^{-1}(j, \cdot)$. Moreover the size of $\text{Index}_{U_j}^{-1}$ increases by at most 1 during calls to either of these functions. The conclusion thus follows by Lemma 8.

If XtraMiddleRnd is off: The only functions to possibly increase the size of $\text{Index}_{A_j}^{-1} = \text{Index}_{U_j}^{-1}$ are $\text{RequestU}^{-1}(j, \cdot)$ and AdaptRight . (Indeed, the latter adds entries to A_j^{-1} ; AdaptLeft also adds entries to A_j^{-1} , but the y 's for which AdaptLeft sets $A_j^{-1}(y)$ are already in $\text{ToBeAssigned}_{U_j}^-$ at the start of AdaptLeft .) Moreover, for each value x such that $A_j(x)$ is assigned in AdaptRight there obviously exists a unique corresponding call to $\text{RequestV}(j, \cdot)$, namely the call which added x to $\text{ToBeAssigned}_{V_j}$. Hence the conclusion also follows by Lemma 8. \square

Lemma 10. *In G_1, G_2 and G_3 , when a pair $(\mathbf{x}^G, \mathbf{y}^U)$ is added to ToBeAdapted_D in RequestU^{-1} , $\mathbf{y}^U[k] \in \text{Index}_{U_k}^-$ for every $k \in [w]$, and there exists $j \in [w]$ such that $\mathbf{y}^U[j] \in \text{ToBeAssigned}_{G_j}^-$.*

Moreover, if XtraMiddleRnd is on, let $\mathbf{x}^V = \tau(\mathbf{y}^U)$. Then $\mathbf{x}^V[k] \in \text{Index}_{V_k}$ for every $k \in [w]$ when $(\mathbf{x}^G, \mathbf{y}^U)$ is added to ToBeAdapted_D .

A symmetric statement holds for RequestV .

Proof. If $(\mathbf{x}^G, \mathbf{y}^U)$ is added to ToBeAdapted_D in RequestU^{-1} then $\text{BlockDefined}(U, \mathbf{y}^U, -) = \mathbf{true}$, so $\mathbf{y}^U[k] \in \text{Index}_{U_k}^-$ for all $k \in [w]$. Moreover if j and y are the arguments to RequestU^{-1} then it's clear that $\mathbf{y}^U[j] = y$ and that $y \in \text{ToBeAssigned}_{U_j}^-$. (See the pseudocode of RequestU^{-1} .)

If XtraMiddleRnd is on, moreover, then it's necessary (see the pseudocode again) that $\text{BlockDefined}(C, \tau(\mathbf{y}^U), \mathbf{true})$ in order for a pair $(\mathbf{x}^G, \mathbf{y}^U)$ to be added to ToBeAdapted_D in RequestU^{-1} . So $\mathbf{x}^V[k] \in \text{Index}_{V_k}$ for all $k \in [w]$. \square

Lemma 11. *For every pair $(\mathbf{x}^G, \mathbf{y}^U)$ ever added to ToBeAdapted_D ,*

$$\mathbf{y}^U \in \text{Index}_{U_1}^- \times \dots \times \text{Index}_{U_w}^-$$

Moreover, if XtraMiddleRnd is on, then

$$\tau(\mathbf{y}^U) \in \text{Index}_{V_1} \times \dots \times \text{Index}_{V_w}.$$

A symmetric statement holds for every pair $(\mathbf{x}^V, \mathbf{y}^H)$ added to ToBeAdapted_E . (This for G_1, G_2 and G_3 .)

Proof. This is a simple corollary of Lemma 10. \square

Lemma 12. *The same pair $(\mathbf{x}^G, \mathbf{y}^U)$ is never added twice to ToBeAdapted_D . In fact, if $(\mathbf{x}^G, \mathbf{y}^U)$ and $(\mathbf{x}'^G, \mathbf{y}'^U)$ are two pairs added to ToBeAdapted_D at different points in time, then both $\mathbf{x}^G \neq \mathbf{x}'^G$ and $\mathbf{y}^U \neq \mathbf{y}'^U$. A symmetric statement holds for pairs $(\mathbf{x}^V, \mathbf{y}^H)$ and ToBeAdapted_E . (This for G_1, G_2 and G_3 .)*

Proof. We start by noting that the second statement easily follows from the first: when $(\mathbf{x}^G, \mathbf{y}^U)$ is added to ToBeAdapted_D in RequestU^{-1} , \mathbf{y}^U is associated to a unique vector \mathbf{y}^H , which association remains unchanged for the rest of the execution by virtue of the fact that the tables P_{T_j} are never overwritten (Lemma 1); moreover the mapping from \mathbf{x}^X to \mathbf{y}^H is also invariant given that values in P_Z are never overwritten either (Lemma 2); finally the mapping from \mathbf{x}^X to \mathbf{x}^G is invariant as P_{F_j} is never overwritten. Hence \mathbf{y}^U “non-malleably” determines \mathbf{x}^G and vice-versa.

We now argue that the same pair $(\mathbf{x}^G, \mathbf{y}^U)$ is never added twice to ToBeAdapted_D in RequestU^{-1} . For the following argument it helps to picture a modified (but equivalent) version of the code in which the last **forall** loop of RequestU^{-1} is detached from the main body of that function, and executed separately by whatever function called RequestU^{-1} in the first place, right after the call to RequestU^{-1} . With this (cosmetic) change in place the execution stack only ever contains one active copy of RequestU^{-1} . (In other words, RequestU^{-1} is never called while another instance of RequestU^{-1} is still waiting to terminate.)

Firstly, now, a given pair $(\mathbf{x}^G, \mathbf{y}^U)$ cannot be added twice to ToBeAdapted_D by the same call to RequestU^{-1} because the **forall** loop of RequestU^{-1} only considers each value of \mathbf{y}^U once. Thus if a pair $(\mathbf{x}^G, \mathbf{y}^U)$ is added twice to \mathbf{y}^U it is added during two distinct calls to RequestU^{-1} .

Let the first such call be $\text{RequestU}^{-1}(j_1, y_1)$ and the second call be $\text{RequestU}^{-1}(j_2, y_2)$. (We do not assume $(j_1, y_1) \neq (j_2, y_2)$.) By the remarks above we can assume—via our mental experiment—that $\text{RequestU}^{-1}(j_1, y_1)$ has finished executing when $\text{RequestU}^{-1}(j_2, y_2)$ is called. Also $\mathbf{y}^U[j_1] = y_1$ and $\mathbf{y}^U[j_2] = y_2$ since the **forall** loop of $\text{RequestU}^{-1}(j, y)$ only iterates over vectors \mathbf{y}^U such that $\mathbf{y}^U[j] = y$. For $\text{RequestU}^{-1}(j_2, y_2)$ to not immediately return we must have $U_{j_2}^{-1}(y_2) = \perp$ and $y_2 \notin \text{ToBeAssigned}_{U_{j_2}^{-1}}$ when the call $\text{RequestU}^{-1}(j_2, y_2)$ occurs. But this implies $\mathbf{y}^U[j_1] \notin \text{Index}_{U_{j_1}}$ before the second call, which contradicts the fact that $\mathbf{y}^U[j_1] \in \text{Index}_{U_{j_1}}$ after the first call by Lemma 3. \square

Lemma 13. *Let $(\mathbf{x}^G, \mathbf{y}^U)$ be a pair added to ToBeAdapted_D at some point and let $(\mathbf{x}^V, \mathbf{y}^H)$ be a pair added to ToBeAdapted_E at some other point of an execution of G_1, G_2 or G_3 . Then $\tau(\mathbf{y}^U) \neq \mathbf{x}^V$ if XtraMiddleRnd is on. Moreover if XtraMiddleRnd is off, $A(\mathbf{x}^V) \neq \mathbf{y}^U$ (and $A^{-1}(\mathbf{y}^U) \neq \mathbf{x}^V$) for all $j \in [w]$ if $A(\mathbf{x}^V)$ (or $A^{-1}(\mathbf{y}^U)$) ever becomes defined.*

Proof. Without loss of generality, we can assume the addition of $(\mathbf{x}^V, \mathbf{y}^H)$ to ToBeAdapted_E occurs after the addition of $(\mathbf{x}^G, \mathbf{y}^U)$ to ToBeAdapted_D . Note moreover these additions will occur in distinct query cycles, since RequestU^{-1} is a left-hand procedure and RequestV is a right-hand procedure. For the second query cycle to occur at all, naturally, the first query cycle must be successful (i.e., not abort).

If XtraMiddleRnd is on: Lemma 11 implies that $\tau(\mathbf{y}^U) \in \text{Index}_{V_1} \times \dots \times \text{Index}_{V_w}$ when $(\mathbf{x}^G, \mathbf{y}^U)$ is added to ToBeAdapted_D . Hence any subsequent call $\text{RequestV}(j, x)$ that doesn't immediately return must have $x \neq \tau(\mathbf{y}^U)[j]$. (More precisely, $C_j(\tau(\mathbf{y}^U)[j])$ will be defined for all $j \in [w]$ if the query cycle that adds $(\mathbf{x}^G, \mathbf{y}^U)$ to ToBeAdapted_D completes without aborting.) So any pair $(\mathbf{x}^V, \mathbf{y}^H)$ added to ToBeAdapted_E by RequestV must have $\mathbf{x}^V \neq \tau(\mathbf{y}^U)$.

If XtraMiddleRnd is off: Lemma 11 implies that $\mathbf{y}^U \in \text{Index}_{U_1}^- \times \dots \times \text{Index}_{U_w}^-$ when $(\mathbf{x}^G, \mathbf{y}^U)$ is added to ToBeAdapted_D , and Lemma 3 implies that $A^{-1}(\mathbf{y}^U) \neq \perp$ after that query cycle completes (since $\text{ToBeAssigned}_{U_j}^- = \emptyset$ for all j at the end of a non-aborted query cycle). Thus any subsequent call $\text{RequestV}(j, x)$ such that $\mathbf{x}^G[j] = x$ will return immediately. So any pair $(\mathbf{x}^V, \mathbf{y}^H)$ added later to ToBeAdapted_E by RequestV must have $\mathbf{x}^V \neq A^{-1}(\mathbf{y}^U)$ or, equivalently (cf. Lemma 1) $A(\mathbf{x}^V) \neq \mathbf{y}^U$. \square

For the remaining lemmas of this section, we define

$$\begin{aligned} \text{Index}_{U_j}^- \wedge \text{Index}_{V_j} &= \{(x, y) : A_j(x) = y\} \cup \\ &\quad \{(x, \perp) : A_j(x) = \perp, x \in \text{Index}_{V_j}\} \cup \\ &\quad \{(\perp, y) : A_j^-(y) = \perp, y \in \text{Index}_{U_j}^-\} \end{aligned}$$

if XtraMiddleRnd is off. We note that since XtraMiddleRnd is off, and by definition of $\text{Index}_{U_j}^-$ and Index_{V_j} , this definition can also be written

$$\begin{aligned} \text{Index}_{U_j}^- \wedge \text{Index}_{V_j} &= \{(x, y) : A_j(x) = y\} \cup \\ &\quad \{(x, \perp) : A_j(x) = \perp, x \in \text{ToBeAssigned}_{A_j}\} \cup \\ &\quad \{(\perp, y) : A_j^-(y) = \perp, y \in \text{ToBeAssigned}_{A_j}^-\}. \end{aligned}$$

We note that by definition, a pair (x, y) cannot exist in $\text{Index}_{U_j}^- \wedge \text{Index}_{V_j}$ if either of the pairs (x, \perp) or (\perp, y) is in $\text{Index}_{U_j}^- \wedge \text{Index}_{V_j}$. Moreover pairs of the form (x, \perp) can only exist during right-hand query cycles, and pairs of the form (\perp, y) can only exist during left-hand query cycles, as is clear from the definition and from Lemma 5.

Each pair in $\text{Index}_{U_j}^- \wedge \text{Index}_{V_j}$ will be called a *slot*. A slot of the form (x, y) is permanent by Lemma 1. Thus, two distinct slots cannot have the same non-null first coordinate, nor can they share the same non-null second coordinate. (Case analysis: if (x, y) and (x, y') are two slots then $A_j(x) = y$ and $A_j(x) = y'$, so $y = y'$ by Lemma 1; secondly, by the definition of $\text{Index}_{U_j}^- \wedge \text{Index}_{V_j}$, (x, y) and (x, \perp) cannot both be slots, as noted above.)

Lemma 14. *If XtraMiddleRnd is off, $\text{Index}_{U_j}^- \wedge \text{Index}_{V_j}$ never contains more than $2q$ elements at any point in an execution of G_1, G_2 or G_3 .*

Proof. AdaptLeft and AdaptRight do not add entries to $\text{Index}_{U_j}^- \wedge \text{Index}_{V_j}$, since they simply turns pairs of the form (\perp, y) into pairs of the form (x, y) (in the case of AdaptLeft) or turn pairs of the form (x, \perp) into pairs of the form (x, y) (in the case of AdaptRight).

Thus, the only functions to add pairs to $\text{Index}_{U_j}^- \wedge \text{Index}_{V_j}$ are calls to $\text{RequestU}^{-1}(j, \cdot)$ and calls to $\text{RequestV}(j, \cdot)$. However each such call adds at most one pair (since these call add at most one element to $\text{ToBeAssigned}_{A_j}^-$ and $\text{ToBeAssigned}_{A_j}$, respectively), and the total number of such calls is at most $2q$ by Lemma 8. \square

We note that when XtraMiddleRnd is off, the value of the function ϕ —defined in the following lemma—on a given pair $(\mathbf{x}^G, \mathbf{y}^U)$ or $(\mathbf{x}^V, \mathbf{y}^H)$ is time-dependent, as the table A_j is time-dependent.

Lemma 15. *If XtraMiddleRnd is on: Let $\phi(\mathbf{x}^G, \mathbf{y}^U) = \mathbf{y}^U$ for each pair $(\mathbf{x}^G, \mathbf{y}^U)$ added to ToBeAdapted_D , and let $\phi(\mathbf{x}^V, \mathbf{y}^H) = \tau^{-1}(\mathbf{x}^V)$ for each pair $(\mathbf{x}^V, \mathbf{y}^H)$ added to ToBeAdapted_E . Then ϕ is injective, and the range of ϕ always lies in the set*

$$\{\mathbf{y} : \mathbf{y} \in \text{Index}_{U_1}^- \times \cdots \times \text{Index}_{U_w}^- \wedge \tau(\mathbf{y}) \in \text{Index}_{V_1} \times \cdots \times \text{Index}_{V_w}\}.$$

If XtraMiddleRnd is off: Let $\phi(\mathbf{x}^G, \mathbf{y}^U)$ be the w -tuple whose j -th element is the pair $(A_j^-(\mathbf{y}^U[j]), \mathbf{y}^U[j])$ for each j , for each $(\mathbf{x}^G, \mathbf{y}^U)$ added to ToBeAdapted_D , and let $\phi(\mathbf{x}^V, \mathbf{y}^H)$ be the w -tuple whose j -th element is the pair $(\mathbf{x}^V[j], A_j(\mathbf{x}^V[j]))$ for each j , for each $(\mathbf{x}^V, \mathbf{y}^H)$ added to ToBeAdapted_E . Then ϕ is injective, and the range of ϕ remains in the set $(\text{Index}_{U_1}^- \wedge \text{Index}_{V_1}) \times \cdots \times (\text{Index}_{U_w}^- \wedge \text{Index}_{V_w})$.

Proof. Whether XtraMiddleRnd is on or off ϕ has the claimed range by Lemma 10 and is injective by lemmas 12 and 13. \square

Lemma 16. *At the end of each successful query cycle in G_1 , G_2 and G_3 , every pair $(\mathbf{y}^U, \mathbf{x}^V)$ in*

$$\Gamma = \begin{cases} \left\{ (\mathbf{y}^B, \mathbf{x}^C) : \tau(\mathbf{y}^B) = \mathbf{x}^C, \forall j \mathbf{y}^B[j] \in \text{Index}_{B_j}^- \wedge \mathbf{x}^C[j] \in \text{Index}_{C_j} \right\} & \text{if XtraMiddleRnd is on} \\ \left\{ (\mathbf{y}^A, \mathbf{x}^A) : \forall j A_j(\mathbf{x}^A[j]) = \mathbf{y}^A[j] \right\} & \text{if XtraMiddleRnd is off} \end{cases}$$

is in a completed path.

Proof. If XtraMiddleRnd is on: For any such pair $(\mathbf{y}^U, \mathbf{x}^V)$, consider the query cycle during which $(\mathbf{y}^U, \mathbf{x}^V)$ becomes contained by the above set. Without loss of generality, has that $j' \in [w]$ has the property that is $\mathbf{y}^U[j']$ added to $\text{Index}_{B_{j'}}^-$ after $\mathbf{y}^U[j]$ is added to $\text{Index}_{B_j}^-$ for all $j \neq j'$, and also after $\mathbf{x}^V[j]$ is added to Index_{C_j} for all j .

As already noted in Lemma 9, a value y is added to $\text{Index}_{B_{j'}}^-$ only by calls to $B(j', \cdot)$ or to $\text{RequestU}^{-1}(j', y)$. If $\mathbf{y}^U[j']$ is added to $\text{Index}_{B_{j'}}^{-1}$ via B , then the query cycle is neutral or a right-hand query cycle, so $B_{j'}^-(\mathbf{y}^U[j]) \neq \perp$ for all $j \neq j'$ and $\mathbf{x}^V[j] \in \text{Index}_{V_j}$ for all $j \in [w]$ by assumption, so the simulator will abort inside of B . (See the pseudocode for procedure $B(j, x)$.)

Therefore we can assume that $\mathbf{y}^U[j']$ is added to $\text{Index}_{B_{j'}}^-$ by a function call of the form $\text{RequestU}^{-1}(j', y)$, which can only occur in a left-hand query cycle. By assumption, then $\mathbf{y}^U[j] \in \text{Index}_{U_j}^{-1}$ and $\mathbf{x}^V[j] \in \text{Index}_{V_j}$ for all $1 \leq j \leq w$ after $\mathbf{y}^U[j']$ is added to $\text{ToBeAssigned}_{B_{j'}}^{-1}$ inside of $\text{RequestU}^{-1}(j', \mathbf{y}^U[j'])$. The **forall** loop of $\text{RequestU}^{-1}(j', \mathbf{y}^U[j'])$ will iterate with \mathbf{y}^U and both $\text{BlockDefined}(U, \mathbf{y}^U, -)$ and $\text{BlockDefined}(V, \mathbf{x}^V = \tau(\mathbf{y}^U), +)$ will return **true**. A corresponding pair $(\mathbf{x}^G, \mathbf{y}^U)$ is then added to ToBeAdapted_D and, assuming successful completion of the query cycle, a completed chain containing $\mathbf{y}^U, \mathbf{x}^V$ is created by AdaptLeft .

If XtraMiddleRnd is off: We assume without loss of generality that $(\mathbf{y}^A, \mathbf{x}^A)$ enters Γ for the first time during a left-hand query cycle, as neutral query cycles don't modify the tables A_j . Then there exists some $j' \in [w]$ such that $\mathbf{y}^A[j']$ is added to $\text{ToBeAssigned}_{A_{j'}}^-$ during that query cycle. We can moreover define j' such that $\mathbf{y}^A[j']$ is added to $\text{ToBeAssigned}_{A_{j'}}^-$ after any other value $\mathbf{y}^A[j]$ is added to $\text{ToBeAssigned}_{A_j}^-$ during that query cycle. Then we have $\mathbf{y}^A[j] \in \text{Index}_{A_j}^-$ for all j when $\mathbf{y}^A[j']$ is added to $\text{ToBeAssigned}_{A_{j'}}^-$, since $\mathbf{y}^A[j] \in \text{Index}_{A_j}^-$ for all j at the end of the query cycle (and being added to $\text{Index}_{A_j}^-$ can only happen by being added to $\text{ToBeAssigned}_{A_j}^-$, for a left-hand query cycle). Thus the call to RequestU^{-1} that adds $\mathbf{y}^A[j']$ to $\text{ToBeAssigned}_{A_{j'}}^-$ will add a pair $(\mathbf{x}^G, \mathbf{y}^U)$ to ToBeAdapted_D with $\mathbf{y}^U = \mathbf{y}^A$. Assuming successful completion of the query cycle, a completed chain containing \mathbf{y}^A is thus created by AdaptLeft . \square

Lemma 17. *The total number of pairs $(\mathbf{x}^G, \mathbf{y}^U)$ ever added to ToBeAdapted_D plus the total number of pairs $(\mathbf{x}^V, \mathbf{y}^H)$ ever added to ToBeAdapted_E in G_1 , G_2 or G_3 is at most $(2q)^w$ if XtraMiddleRnd is off and is at most $\text{Cond}_\tau(2q)$ if XtraMiddleRnd is on.*

Proof. When XtraMiddleRnd is off this is a direct corollary of lemmas 14 and 15.

When XtraMiddleRnd is on this follows from Lemma 15 and from the fact that

$$\left\{ \mathbf{y}^U : \mathbf{y}^U \in \text{Index}_{U_1}^- \times \dots \times \text{Index}_{U_w}^- \wedge \tau(\mathbf{y}^U) \in \text{Index}_{V_1} \times \dots \times \text{Index}_{V_w} \right\}$$

has size at most $\text{Cond}_\tau(2q)$ by the definition of conductance. \square

In other words, the total number of pairs $(\mathbf{x}^G, \mathbf{y}^U)$ added to ToBeAdapted_D plus the total number of pairs $(\mathbf{x}^V, \mathbf{y}^H)$ added to ToBeAdapted_E is at most $\alpha(q)$.

Lemma 18. *The number of times lines 8–11 of RequestU^{-1} plus the number of times lines 8–9 of RequestV are executed is at most $\alpha(q)$.*

Proof. This is a corollary of lemmas 17 and 12, since lines 8–11 of RequestU^{-1} and lines 8–9 of RequestV are executed each time before an element is added respectively to ToBeAdapted_D or ToBeAdapted_E . \square

For the next two lemmas, let AllAdapted be the set of all pairs $(\mathbf{x}^G, \mathbf{y}^U)$ ever added to ToBeAdapted_D unioned with the set of all pairs $(\mathbf{x}^V, \mathbf{y}^H)$ ever added to ToBeAdapted_E .

Lemma 19. *For every $T \in \{D, \dots, L\}$ and for every $j \in [w]$ there is an injection from the set of pairs (x, y) such that $T_j(x) = y$ to $\text{AllAdapted} \cup [q]$, at any point in time of any execution of G_1 , G_2 or G_3 .*

Proof. If the distinguisher previously queried $T(j, x)$ or $T^{-1}(y)$ with (say) his h -th query, we map the pair (x, y) to the value $h \in [q]$.

Otherwise the distinguisher did not query $T(j, x)$ or $T^{-1}(y)$, and the entry was added to T_j either by line 17 of RequestG , by line 17 of RequestH^{-1} , by line 8 of RequestU^{-1} , by line 8 of RequestV , by line 11 of RequestU^{-1} , by line 11 of RequestV , or else inside the second **forall** loop of AdaptLeft or AdaptRight . (Recall that $T \neq A, B, C$.) However, for each of these cases it is straightforward (by direct examination of the pseudocode) to associate either some pair $(\mathbf{x}^G, \mathbf{y}^U)$ added to ToBeAdapted_D or some pair $(\mathbf{x}^V, \mathbf{y}^H)$ added to ToBeAdapted_E . Moreover it is equally easy to see that under this natural mapping, no element of AllAdapted is the image of more than one pair (x, y) , for a given T and j (in a nutshell, because the query $T_j(x) = y$ lies on the path corresponding to the element in AllAdapted , which limits x and y to one possible combination). \square

Lemma 20. *Let $T \in \{D, \dots, L\}$, $j \in [w]$. Then the number of x for which $T_j(x) \neq \perp$ is at most $q + \alpha(q)$ at any point of any execution of G_1 , G_2 and G_3 . Moreover if $T \in \{A, B, C\}$ then the same number is at most $2q$.*

Proof. The first statement follows by Lemma 19 coupled with Lemma 17. The second statement is a restatement of Lemma 9. \square

Lemma 21. *In any execution of G_1 , G_2 or G_3 the total number of distinct²⁴ calls from RequestU^{-1} to BlockRequestG and from RequestV to $\text{BlockRequestH}^{-1}$ is at most $\alpha(q)$.*

Proof. This follows from Lemma 17 since for each call $\text{BlockRequestG}(\mathbf{x}^G)$ there is a pair $(\mathbf{x}^G, \mathbf{y}^U) \in \text{ToBeAdapted}_D$ for some \mathbf{y}^U and for each call $\text{BlockRequestH}^{-1}(\mathbf{y}^H)$ there is a pair $(\mathbf{x}^V, \mathbf{y}^H) \in \text{ToBeAdapted}_E$ for some \mathbf{x}^V . \square

Lemma 22. *For any $j \in [w]$ and at any point in an execution of G_1 , G_2 or G_3 , Index_{G_j} and $\text{Index}_{H_j}^-$ have size at most $q + \alpha(q)$.*

²⁴ Distinct means with distinct arguments. The total number of calls may be greater.

Proof. The size of Index_{G_j} is upper bounded by the number of distinguisher calls to G and G^{-1} plus the number of distinct values x for which a call $\text{RequestG}(j, x)$ occurs plus the number of times BlockG^{-1} is called. The number of distinguisher calls to G and G^{-1} is at most q in total. Moreover, since RequestG is otherwise (that is, when it isn't being called by G because of a direct distinguisher query to G) only called by BlockRequestG and the end of RequestU^{-1} , and since BlockG^{-1} is only called from within MiddleToLeft from within RequestV , the conclusion follows by Lemma 21. \square

Lemma 23. *In any execution of G_1 , G_2 or G_3 , the total number of distinct pairs $(\mathbf{x}^X, \mathbf{y}^H)$ such that a call $\text{CheckZ}(\mathbf{x}^X, \mathbf{y}^H)$ occurs is at most $\beta(q)(q + \alpha(q))^w$.* \square

Proof. CheckZ is only called from within on arguments of the type $(\mathbf{x}^X, \mathbf{y}^H)$, where \mathbf{y}^H satisfies $\text{BlockDefined}(H, \mathbf{y}^H, -) = \mathbf{true}$. There are at most $(q + \alpha(q))^w$ distinct such \mathbf{x}^X over the course of an execution by Lemma 20.

Moreover, if XtraOuterRnd is off then $\text{BlockDefined}(G, \mathbf{x}^X, +) = \mathbf{true}$ and there are at most $\beta(q) = (q + \alpha(q))^w$ distinct such \mathbf{x}^X over the course of an execution by the same lemma.

If XtraOuterRnd is on, finally, CheckZ is only called on arguments of then \mathbf{x}^X satisfies $\text{BlockDefined}(F, \mathbf{x}^X, +) = \mathbf{true}$ and $\text{BlockDefined}(G, \nu(\text{BlockF}(\mathbf{x}^X)), +) = \mathbf{true}$. Obviously there are at most $\beta(q) = \text{Cond}_\nu(\alpha(q) + q)$ distinct such \mathbf{x}^X by Lemma 20 and by the definition of conductance. \square

E.3 Bounding the Abort Probability in G_3

Lemma 24. *An execution of G_3 never aborts inside ReadTape unless the function calling ReadTape is D , D^{-1} , E or E^{-1} .*

Proof. Note that values in $A_j, B_j, C_j, F_j, G_j, H_j, I_j, J_j, K_j, L_j$ and P_Z are always set via ReadTape ; hence, ReadTape never aborts except the second argument to ReadTape is D_j or E_j , and such calls to ReadTape only occur in procedures eD , D^{-1} , E and E^{-1} . \square

Lemma 25. *The number of times Z/Z^{-1} is called during an execution of G_2 or G_3 is at most $q + \alpha(q)$.*

Proof. Note that Z is no longer called by CheckZ in game G_2 and G_3 . Hence, the only calls to Z are those which occur in RequestU^{-1} and RequestV , or else are distinguisher calls. The number of former calls is at most $\alpha(q)$ by Lemma 17 while the latter is at most q . \square

Lemma 26. *The probability that G_2 or G_3 aborts inside of Z or Z^{-1} is at most $w(q + \alpha(q))^2 / (N - q - \alpha(q))$.*

Proof. Note that values in P_Z are always set via ReadTape and never via SetTable ; hence, calls to ReadTape never abort in Z or Z^{-1} .

The number of entries in any of the tables X_j, Y_j is at most $q + \alpha(q)$ by Lemma 20. Moreover, a lookup in p_Z produces a value $z_1 \dots z_w \in \{0, 1\}^{wn}$ where each z_i is uniform in a set of size at least $N - q - \alpha(q)$ by Lemma 25. Thus, by a union bound over the entries in X_j (for queries to Z^{-1}) or over the entries in Y_j (for queries to Z) as well as over the value of j the probability that a query to Z or Z^{-1} causes an abort is at most $w(q + \alpha(q)) / (N - q - \alpha(q))$. The total number of queries to Z/Z^{-1} being at most $q + \alpha(q)$, the lemma follows by a union bound over all these queries. \square

Lemma 27. *The procedure CheckZ is never called twice on the same input pair (\mathbf{x}, \mathbf{y}) over the course of an execution of G_3 .*

Proof. Calls to CheckZ occur only in RequestG and RequestH⁻¹. Obviously, the given call to RequestG or RequestH⁻¹ will not call CheckZ twice on the same pair $(\mathbf{x}^X, \mathbf{y}^H)$, as is clear by inspection of the pseudocode. Therefore if CheckZ is called twice on the same pair $(\mathbf{x}^X, \mathbf{y}^H)$, the two calls must occur from within distinct instances of RequestG and/or RequestH⁻¹. For concreteness, assume to start with that the two calls are issued by distinct instances of RequestG. (This is in fact the more interesting case, since if one call occurs within RequestG and the other occurs within RequestH⁻¹, the calls are happening in different query cycles.)

Let the two calls to RequestG in question be RequestG(j_1, x_1) and RequestG(j_2, x_2), where RequestG(j_1, x_1) starts executing before RequestG(j_2, x_2); we do not assume $(j_1, x_1) \neq (j_2, x_2)$.

Since RequestG(j_1, x_1) does not start making recursive calls until its last **forall** loop, the call CheckZ($\mathbf{x}^X, \mathbf{y}^H$) that occurs in RequestG(j_1, x_1) occurs before RequestG(j_2, x_2) starts executing. At the point when the call CheckZ($\mathbf{x}^X, \mathbf{y}^H$) is made (or even when RequestG(j_1, x_1) is called) \mathbf{x}^X is associated to a unique \mathbf{x}^G , which is namely $\mathbf{x}^G = \nu(F(\mathbf{x}^X))$ if XtraOuterRnd is on and which is $\mathbf{x}^G = \mathbf{x}^X$ if XtraOuterRnd is off. Moreover $\mathbf{x}^G[j_1] = x_1$ by design of RequestG, and $\mathbf{x}^G[j_2] = x_2$ because RequestG(j_2, x_2) also makes the call CheckZ($\mathbf{x}^X, \mathbf{y}^H$) by assumption. But since $\mathbf{x}^G[j] \in \text{Index}_{G_j}$ for all j when RequestG(j_1, x_1) calls CheckZ($\mathbf{x}^X, \mathbf{y}^H$), this implies that $x_2 \in \text{Index}_{G_{j_2}}$ when RequestG(j_2, x_2) is called, since the latter call occurs after RequestG(j_1, x_1) calls CheckZ($\mathbf{x}^X, \mathbf{y}^H$), and hence RequestG(j_2, x_2) would return immediately, a contradiction.

Other cases (e.g., when one call to CheckZ is caused by RequestG, and the other by RequestH⁻¹) can be handled by a similar analysis. \square

Lemma 28. *The probability that G_3 aborts inside of ForwardOutside or BackwardOutside is at most $(q + w\alpha(q))(q + \alpha(q))/(N - q - \alpha(q))$.*

Proof. BackwardOutside(T, j, y) is called with $T \in \{F, G\}$ while ForwardOutside(T, j, x) is called with $T = H$. By Lemma 20, values read from $p_{F_j}^-$, $p_{G_j}^-$ or p_{H_j} come uniformly at random from a set of size at least $N - q - \alpha(q)$. Since P_Z contains at most $q + \alpha(q)$ entries by Lemma 25, the abort probability in a query to ForwardOutside or BackwardOutside is at most $(q + \alpha(q))/(N - q - \alpha(q))$.

The total number of queries to ForwardOutside and BackwardOutside is no more than $q + w\alpha(q)$, including no more than q queries directly issued by the distinguisher and, by Lemma 18, no more than $w\alpha(q)$ queries made by MiddleToLeft and MiddleToRight. The lemma thus follows by a union bound. \square

Lemma 29. *When XtraOuterRnd is on, the probability that G_3 aborts inside of F or G⁻¹ is at most*

$$\frac{2w(q + \alpha(q))\text{aboCond}_\nu(q + \alpha(q))}{N - q - \alpha(q)}. \quad (19)$$

Proof. We focus on a call F(j, x). Calling F(j, x) causes abort if and only if $F_j(x) = \perp$ and if once we set $F_j(x) = y := p_{F_j}(x)$, there exist $\mathbf{y}^F, \mathbf{x}^G$ such that $\mathbf{y}^F[j] = y$, $\mathbf{x}^G = \nu(\mathbf{y}^F)$ and $\text{BlockDefined}(F, \mathbf{y}^F, -) = \text{BlockDefined}(G, \mathbf{x}^G, +) = \mathbf{true}$. Since there is no such thing as a set $\text{ToBeAssigned}_{F_j}^-$, the set of values y that will cause an abort is

$$\{y \in \{0, 1\}^n : \exists \mathbf{y}^F, \mathbf{x}^G \text{ s.t. } \mathbf{y}^F[j] = y, \nu(\mathbf{y}^F) = \mathbf{x}^G, \forall k \neq j \ F_k^-(\mathbf{y}^F[k]) \neq \perp, \forall k \ \mathbf{x}^G[k] \in \text{Index}_{G_k}\}.$$

The size of this set can be upper bounded by the size of

$$\{(\mathbf{y}^F, \mathbf{x}^G) : \nu(\mathbf{y}^F) = \mathbf{x}^G, \forall k \neq j F_k^-(\mathbf{y}^F[k]) \neq \perp, \forall k \mathbf{x}^G[k] \in \text{Index}_{G_k}\}$$

which is upper bounded by $\text{aboCond}_\nu(q + \alpha(q))$ by the definition of all-but-one conductance and by lemmas 20, 22. Moreover y comes uniformly at random from a set of size at least $N - q - \alpha(q)$ by Lemma 20. Thus the probability of abort for a call to $F(j, \cdot)$ is no more than $\text{Cond}_\nu(q + \alpha(q))/(N - q - \alpha(q))$, and it is easy to see that the same bound holds for a call to $G^{-1}(j, \cdot)$. Multiplying by $q + \alpha(q)$ to account for all calls to $F(j, \cdot)$ and by a further factor of $2w$ to account for all values of j as well as for calls to G^{-1} yields the final bound. \square

Lemma 30. *The probability that NumOuter exceeds q in G_3 is zero.*

Proof. Each increment of NumOuter is associated to a unique pair $(\mathbf{x}^X, \mathbf{y}^H)$ such that $\text{CheckZ}(\mathbf{x}^X, \mathbf{y}^H) = \text{true}$. Moreover a given pair $(\mathbf{x}^X, \mathbf{y}^H)$ only contributes at most once to the increase in NumOuter by Lemma 27. We claim that before the call $\text{CheckZ}(\mathbf{x}^X, \mathbf{y}^H)$ occurs the distinguisher has already made the call $Z(\mathbf{x}^X) \rightarrow \mathbf{y}^H$ or $Z^{-1}(\mathbf{y}^H) \rightarrow \mathbf{x}^X$. This would complete the proof since the distinguisher has at most q queries.

To prove the latter claim, consider wlog the case where NumOuter is incremented within a left-hand query cycle, so that the call to CheckZ occurs from within RequestG. Assume by contradiction that the distinguisher did not make the call $Z(\mathbf{x}^X) \rightarrow \mathbf{y}^H$ or $Z^{-1}(\mathbf{y}^H) \rightarrow \mathbf{x}^X$, so that the entry $P_Z(\mathbf{x}^X) = \mathbf{y}^H$ which results in $\text{CheckZ}(\mathbf{x}^X, \mathbf{y}^H)$ returning true (see the pseudocode of procedure CheckZ in G_3 , Fig. 3) was created by a call to Z^{-1} in RequestU⁻¹ or by a call to Z in RequestV, as no other simulator functions call Z/Z⁻¹ in G_3 . In the latter case the call to Z must have occurred during a previous query cycle because RequestV is a right-hand function while RequestG is a left-hand function and, thus, the pair $(\mathbf{x}^X, \mathbf{y}^H)$ was in a completed path already at the beginning of the current query cycle, a contradiction to the first line of RequestG since values are never removed from the tables T_j and since the vector \mathbf{x}^X would have already been associated to a unique vector \mathbf{x}^G such that $G(\mathbf{x}^G) \neq \perp$ at the beginning of the query cycle. (I.e., the vector \mathbf{x}^G on the same completed path as \mathbf{x}^X .) In the other case, when the entry $P_Z(\mathbf{x}^X) = \mathbf{y}^H$ was created by a call to Z^{-1} in RequestU⁻¹, then either that query to Z^{-1} occurred in a previous query cycle, in which case the same argument applies, or else $(\mathbf{x}^X, \mathbf{y}^U) \in \text{ToBeAdapted}_D$ for some vector \mathbf{y}^U , contradicting the fact that RequestG made it to the increment of NumOuter. This completes the proof. \square

Lemma 31. *The probability that an execution of G_3 aborts inside either of D, D^{-1}, E or E^{-1} is, altogether, at most $2w\alpha(q)(q + \alpha(q))/(N - q - \alpha(q))$.*

Proof. We fix a value of j and focus on calls $D(j, \cdot), D^{-1}(j, \cdot)$, and later multiply by $2w$ (union bound) to obtain the final bound.

Note that at most $\alpha(q)$ entries of D_j are fixed by SetTable, since at most $\alpha(q)$ paths are completed by Lemma 17. Moreover, every lookup in the table p_{D_j} or $p_{D_j}^{-1}$ gives a value drawn uniformly at random from a set of size at least $N - q - \alpha(q)$ by Lemma 20. Hence the probability that an abort occurs because a drawn value “hits” a predefined value within D or D^{-1} is at most $\alpha(q)/(N - q - \alpha(q))$ per query. Furthermore the number of queries to $D(j, \cdot)$ or $D^{-1}(j, \cdot)$ that result in a call ReadTable is at most $q + \alpha(q)$ again by Lemma 20. Multiplying by $2w$ to account for the different values of j and for queries to D gives the union bound its final form. \square

Lemma 32. *When `AdaptLeft` is called, for every pair $(\mathbf{x}^G, \mathbf{y}^U) \in \text{ToBeAdapted}_D$ and for every $1 \leq j \leq w$, either $\mathbf{x}^G[j] \in \text{ToBeAssigned}_{G_j}$ or $G_j(\mathbf{x}^1[j]) \neq \perp$, but not both, and, similarly, either $\mathbf{y}^U[j] \in \text{ToBeAssigned}_{U_j}^{-1}$ or $U_j^{-1}(\mathbf{y}^U[j]) \neq \perp$, but not both. A symmetric statement holds for calls to `AdaptRight` and pairs $(\mathbf{x}^V, \mathbf{y}^H) \in \text{ToBeAdapted}_E$.*

Proof. The “but not both” clauses follows from Lemma 6. Subsequently, the statement is obvious for \mathbf{y}^U , given that the **forall** loop of `RequestU-1` doesn’t proceed unless `BlockDefined(U, \mathbf{y}^U , -) = \mathbf{true}`. Moreover, once a pair $(\mathbf{x}^G, \mathbf{y}^U)$ is added to `ToBeAdaptedD` the call `BlockRequestG(\mathbf{x}^G)` occurs, which implies (by direct inspection of `BlockRequestG` and of `RequestG`) that the statement also hold for \mathbf{x}^G . \square

Lemma 33. *The probability that an execution of G_3 aborts inside of `BlockU-1`, `BlockV`, `BlockG` or `BlockH-1` is zero.*

Proof. First we consider the call to `BlockG` in `RequestH-1` via `LeftToMiddle`. Since `RequestH-1` is only called during a right-hand query cycle, the tables `ToBeAssignedGj` remain empty during a right-hand query cycle. Hence if `BlockDefined(G, \mathbf{x}^G , +) = \mathbf{true}` in `RequestH-1` it implies that $G_j(\mathbf{x}^G[j]) \neq \perp$ for $1 \leq j \leq w$. This implies that the call to `BlockG` doesn’t abort within `RequestH-1`.

When `XtraMiddleRnd` is on, `BlockU-1` can be called in `RequestV` via `MiddleToLeft`. Since `RequestV` is only called during a right-hand query cycle, the tables `ToBeAssignedUj`⁻¹ remain empty during a right-hand query cycle. Hence if `BlockDefined(B, \mathbf{y}^B , -) = \mathbf{true}` in `RequestV` it implies that $B_j(\mathbf{y}^B[j]) \neq \perp$ for $1 \leq j \leq w$. This implies that the call to `BlockU-1` never aborts within `RequestV`.

Otherwise, `BlockU-1`, `BlockG` are only called in `AdaptLeft`. The claim thus follows by Lemma 32 and by direct inspection of `AdaptLeft`.

Calls to `BlockV`, `BlockH-1` are symmetrically analyzed. \square

Lemma 34. *When `XtraMiddleRnd` is on, the probability that G_3 aborts inside of `B` or `C-1` is at most*

$$\frac{4wq \cdot \text{aboCond}_\tau(2q)}{N - 2q}.$$

Proof. By Lemma 24, `ReadTape` does not abort when called by `B` or `C-1`. By Lemma 9, values read from the tape p_{B_j} or from the tape $p_{C_j}^{-1}$ come uniformly at random from a set of size at least $N - 2q$.

Calling `B(j, x)` causes abort if and only $B_j(x) = \perp$ and if once we set $B_j(x) = y := p_{B_j}(x)$, there exists a pair $\mathbf{y}^B, \mathbf{x}^C$ such that $\mathbf{y}^B[j] = y$, $\mathbf{x}^C = \tau(\mathbf{y}^B)$ and `BlockDefined(B, \mathbf{y}^B , -) = \mathbf{true}` and `BlockDefined(C, \mathbf{x}^C , +) = \mathbf{true}`. As in Lemma 29, the number of such y ’s is upper bounded by the size of the set

$$\left\{ (\mathbf{y}^B, \mathbf{x}^C) : \tau(\mathbf{y}^B) = \mathbf{x}^C, \forall k \neq j \mathbf{y}^B[k] \in \text{Index}_{B_k}^{-1}, \forall k \mathbf{x}^C[k] \in \text{Index}_{C_k} \right\}$$

which is upper bounded by `aboCondτ(2q)` by the definition of all-but-one conductance and by Lemma 9. Hence the probability of abort for a single call to `B` is no more than `aboCondτ(2q)/(N - 2q)`. Moreover at most $2q$ distinct queries to `B(j, ·)` occur by Lemma 9 again, so the final bound follows by symmetry between `B` and `C-1` and by a straightforward union bound. \square

Lemma 35. *During a left-hand query cycle, entries are not added to tables G_j or U_j until `AdaptLeft` is called. Moreover, the only entries added to H_j are via forward calls to H . A symmetric statement holds for right-hand query cycles.*

Proof. Focusing on left-hand query cycles, the first statement follows by inspection of the procedures `RequestG` and `RequestU-1`. The second statement follows by inspection of the same procedures and by noting that the procedure `BlockH-1` (called in `RequestG` via `RightToMiddle`) does not actually call H^{-1} , but performs look-ups in the tables H_j^{-1} instead. \square

Lemma 36. *At the end of each non-aborted query cycle in G_3 , every pair $(\mathbf{x}^G, \mathbf{y}^H)$ such that*

1. $G_j(\mathbf{x}^G[j]) \neq \perp$ and $H_j^{-1}(\mathbf{y}^H[j]) \neq \perp$ for all $1 \leq j \leq w$;
2. If `XtraOuterRnd` is on and $\mathbf{y}^F = \nu^{-1}(\mathbf{x}^G)$ then $F_j^{-1}(\mathbf{y}^F[j]) \neq \perp$ for all $1 \leq j \leq w$;
3. $P_Z(\mathbf{x}^X) = \mathbf{y}^H$, where $\mathbf{x}^X = F^{-1}(\mathbf{y}^F)$ if `XtraOuterRnd` is on, $\mathbf{x}^X = \mathbf{x}^G$ otherwise

is in a completed path.

Proof. We prove the statement by induction on the number of distinguisher queries. We refer to the lemma’s assertion as “the invariant” in the proof below. Obviously, the invariant holds before the distinguisher starts making queries.

Distinguisher queries to Z and Z^{-1} do not affect the invariant, thanks to the abort conditions installed in those functions in game G_3 . Nor do distinguisher queries to G^{-1} , H and F , F^{-1} thanks to similar abort conditions installed in those functions.

Other neutral calls obviously leave the lemma’s invariant unaffected as well, since such calls don’t affect any of the tables on which the invariant depends.

For the rest we focus on a left-hand query (i.e., to G or U^{-1}), the case of a right-hand query being symmetric.

Let $(\mathbf{x}_*^G, \mathbf{y}_*^H)$ be a putative pair that satisfies the lemma’s constraints at the end of the query cycle, but such that \mathbf{x}_*^G and \mathbf{y}_*^H don’t lie in a (common) completed path at the end of the query cycle.

At the end of a non-aborted query cycle the condition $G_j(\mathbf{x}_*^G[j]) \neq \perp$ is equivalent to $\mathbf{x}_*^G[j] \in \text{Index}_{G_j}$. So \mathbf{x}_*^G and \mathbf{y}_*^H satisfy

1. $\mathbf{x}_*^G[j] \in \text{Index}_{G_j}$ for all $j \in [w]$;
2. $H_j^{-1}(\mathbf{y}_*^H[j]) \neq \perp$ for all $j \in [w]$;
3. if `XtraOuterRnd` is on then $F^{-1}(\mathbf{y}_*^F) \neq \perp$ where and $\mathbf{y}_*^F = \nu^{-1}(\mathbf{x}_*^G)$;
4. $P_Z(\mathbf{x}_*^X) = \mathbf{y}_*^H$, where $\mathbf{x}_*^X = F^{-1}(\mathbf{y}_*^F)$ if `XtraOuterRnd` is on, and where $\mathbf{x}_*^X = \mathbf{x}_*^G$ otherwise

at the end of the query cycle. Moreover, by the hypothesis that the invariant holds at the start of the query cycle, one or more of these constraints doesn’t hold at the start of the query cycle.

Consider the first time the pair $(\mathbf{x}_*^G, \mathbf{y}_*^H)$ satisfies constraints 1–4. Note this cannot occur in H , F , F^{-1} , Z and Z^{-1} because of the abort conditions in those functions. (As far as the last claim is concerned, in fact, constraint 1 could be replaced with “ $G_j(\mathbf{x}_*^G[j]) \neq \perp$ ” as well, which would only make the constraints harder to satisfy. One can also note that F^{-1} is never actually called during a left-hand query cycle.)

As argued in Lemma 35, H^{-1} and $\text{Request}H^{-1}$ are never called during a left-hand query cycle. Moreover AdaptLeft leaves Index_{G_j} and H_j , F_j and P_Z invariant, so $(\mathbf{x}_*^G, \mathbf{y}_*^H)$ must satisfy the constraints 1–4 before AdaptLeft is called.

We have ruled out all the functions that affect the tables F_j , H_j and P_Z (being respectively F , H and Z/Z^{-1}) during a left-hand query cycle as being the point where constraints 1–4 become satisfied for the first time, and we know the constraints become satisfied before AdaptLeft . The only remaining possibility is that the constraints become satisfied when $\mathbf{x}_*^G[k]$ is added to $\text{ToBeAssigned}_{G_k}$ for some k . But when $\mathbf{x}_*^G[k]$ is added to $\text{ToBeAssigned}_{G_k}$ inside $\text{Request}G(k, \mathbf{x}_*^G[k])$, and since by assumption all the constraints 1–4 are satisfied at that moment, the pair $(\mathbf{x}_*^G, \mathbf{y}_*^U)$ would be added to ToBeAdapted_D for some \mathbf{y}_*^U and, if abort doesn't occur, \mathbf{x}_*^G must be in a completed path at the end of the query cycle, which would perforce include \mathbf{y}_*^H . \square

Lemma 37. *When AdaptLeft is called in G_3 , for every pair $(\mathbf{x}^G, \mathbf{y}^U) \in \text{ToBeAdapted}_D$ there exists a $j \in [w]$ such that $\mathbf{x}^G[j] \in \text{ToBeAssigned}_{G_j}$ and there exists $k \in [w]$ such that $\mathbf{y}^U[k] \in \text{ToBeAssigned}_{U_k}^{-1}$. A symmetric statement holds for calls to AdaptRight and pairs $(\mathbf{x}^V, \mathbf{y}^H) \in \text{ToBeAdapted}_E$.*

Proof. The claim is clear for vector \mathbf{y}^U by inspection of $\text{Request}U^{-1}$, where pairs are added to ToBeAdapted_D . Indeed every addition of a pair $(\mathbf{x}^G, \mathbf{y}^U)$ to ToBeAdapted_D is preceded by adding $\mathbf{y}^U[j]$ to $\text{ToBeAssigned}_{U_j}^{-1}$ for some j , and $\mathbf{y}^U[j]$ is not removed from $\text{ToBeAssigned}_{U_j}^{-1}$ until AdaptLeft is called. We also note that in this case the pair $(\mathbf{x}^G, \mathbf{y}^U)$ is added to ToBeAdapted_D during the call $\text{Request}U^{-1}(j, \mathbf{y}^U[j])$.

In order to argue a similar conclusion for the vector \mathbf{x}^G of the pair $(\mathbf{x}^G, \mathbf{y}^U)$ in question we can consider the point in time after \mathbf{x}^G has been computed in $\text{Request}U^{-1}(j, \mathbf{y}^U[j])$ but before $\text{BlockRequest}G(\mathbf{x}^G)$ is called.

If at this point $G_h(\mathbf{x}^G[h]) = \perp$ for some $h \in [w]$ then $\text{BlockRequest}G(\mathbf{x}^G)$ will add $\mathbf{x}^G[h]$ to $\text{ToBeAssigned}_{G_h}$, and the conclusion follows as well for \mathbf{x}^G . Therefore we can assume that $G_h(\mathbf{x}^G[h]) \neq \perp$ for all $h \in [w]$ at this point, i.e., that $G(\mathbf{x}^G) \neq \perp$ at this point. Since new entries aren't added to the tables G_h before AdaptLeft is called, this also implies that $G(\mathbf{x}^G) \neq \perp$ already at the start of the query cycle.

However, $\text{Request}U^{-1}(j, \mathbf{y}^U[j])$ has computed a (forward) path going all the way from \mathbf{y}^U to \mathbf{x}^G via a vector \mathbf{y}^H (computed in MiddleToRight), and one can check that the pair $(\mathbf{x}^G, \mathbf{y}^H)$ now meets the conditions 1–4 listed in the previous lemma. But following the same argument as in the proof of that lemma (i.e., seeking the first point in time at which $(\mathbf{x}^G, \mathbf{y}^H)$ meets conditions 1–4) leads us to a contradiction, for we concluded there the existence of some value k such that $\mathbf{x}^G[k]$ is added to $\text{ToBeAssigned}_{G_k}$ during the query cycle, which cannot be the case here if $G(\mathbf{x}^G) \neq \perp$ at the start of the query cycle. \square

Lemma 38. *The probability that abort occurs in AdaptLeft or in AdaptRight in an execution of G_3 is no more than*

$$\frac{2w\alpha(q)(q + \alpha(q)) \cdot \text{MaxPreim}(\bar{\pi})}{N - q - \alpha(q)} + \frac{2w\alpha(q)^2 \cdot \text{MaxCoPr}(\bar{\pi})}{N - q - \alpha(q)}$$

if XtraUntglRnds is off, and no more than

$$\frac{(4w\alpha(q) + 2w\alpha(q)^2) \cdot \text{MaxPreim}(\bar{\pi})(q + \alpha(q))}{N - q - \alpha(q)}$$

if `XtraUntglRnds` is on.

Proof. As all values of G_j, U_j, I_j, J_j are set by `ReadTape`, the calls to `ReadTape` that occur in `AdaptLeft` cannot cause abort (Lemma 24).

Moreover, by Lemma 33, the calls to `BlockG` and `BlockU`⁻¹ that occur in `AdaptLeft` cannot cause abort either. Therefore, the only calls that might cause `AdaptLeft` to abort are the calls to `SetTable`. Specifically, a call `SetTable(D_j, x, y)` will abort if either the x - or y -argument already appears in D_j . Arguing by symmetry, we will bound the probability that $D_j(x) \neq \perp$.

For concreteness, assume that `ToBeAdaptedD` contains k pairs

$$(\mathbf{x}_1^G, \mathbf{y}_1^U), \dots, (\mathbf{x}_k^G, \mathbf{y}_k^U)$$

when `AdaptLeft` is called. Let

$$\mathcal{S}_i = \{j : 1 \leq j \leq w, G_j(\mathbf{x}_i^G[j]) = \perp\}$$

Note that $\mathcal{S}_i \neq \emptyset$ for $1 \leq i \leq k$ by Lemma 37.

Let $\mathcal{D} \subseteq \{0, 1\}^n$ be the domain of D_j when `AdaptLeft` is called and, if `XtraUntglRnds` is on, let $\mathcal{I} \subseteq \{0, 1\}^n$ be the domain of I_j when `AdaptLeft` is called. In other words,

$$\begin{aligned} \mathcal{D}_j &= \{x \in \{0, 1\}^n : D_j(x) \neq \perp\} \\ \mathcal{I}_j &= \{x \in \{0, 1\}^n : I_j(x) \neq \perp\} \end{aligned}$$

with these “snapshots” of D_j and I_j being taken when `AdaptLeft` is called. (Thus subsequent modifications to D_j, I_j do not affect $\mathcal{D}_j, \mathcal{I}_j$.)

Define

$$\begin{aligned} \mathbf{y}_i^G &= G(\mathbf{x}_i^G) \\ \mathbf{x}_i^D &= \pi_G(\mathbf{y}_i^G) && \text{if XtraUntglRnds is off} \\ \mathbf{x}_i^I &= \pi_G(\mathbf{y}_i^G) && \text{if XtraUntglRnds is on} \\ \mathbf{y}_i^I &= I(\mathbf{x}_i^I) && \text{if XtraUntglRnds is on} \\ \mathbf{x}_i^D &= \pi_I(\mathbf{y}_i^I) && \text{if XtraUntglRnds is on} \end{aligned}$$

If `XtraUntglRnds` is off, \mathbf{y}_i^G and \mathbf{x}_i^D become defined after the first **forall** loop of `AdaptLeft` has completed. If `XtraUntglRnds` is on, \mathbf{y}_i^G and \mathbf{x}_i^I become defined after the **forall** loop has completed, while \mathbf{y}_i^I and \mathbf{x}_i^D only become defined within the second (outermost) **forall** loop of `AdaptLeft`.

In either case, the underlying probability space for the following argument is all the coins thrown during the `AdaptLeft` call. One can think of $\mathbf{y}_i^G, \mathbf{x}_i^I, \mathbf{y}_i^I, \mathbf{x}_i^D$ as random variables over this probability space. Moreover, for the sake of having these random values always be well-defined, we will assume that `SetTable(D_j, ·, ·)` doesn’t actually stop the execution of `AdaptLeft` if **abort** occurs, but throws some kind of uncaught exception instead. This way both **forall** loops fully complete and all afore-mentioned values are defined, modulo the flag `XtraUntglRnds`.

We define the following bad events:

- `PreExistingi,j` occurs if and only if $\mathbf{x}_i^D[j] \in \mathcal{D}_j$ ($i \in [k], j \in [w]$)
- `Collidingi,i',j` occurs if and only if $\mathbf{x}_i^D[j] = \mathbf{x}_{i'}^D[j]$ ($1 \leq i < i' \leq k, j \in [w]$)

It is easy to see that some call $\text{SetTable}(D_j, x, y)$ aborts because $D_j(x) \neq \perp$ if and only if an event from these two families occurs.

First we bound the probability of one of these events occurring when XtraUntglRnds is off. We have $|\mathcal{D}_j| \leq q + \alpha(q)$ by Lemma 20. Note that $\mathbf{x}_i^D[j]$ only becomes defined (and for all j at once) once $\text{ReadTape}(G_j, \mathbf{x}_i^G[j], p_{G_j})$ has been called for all $j \in \mathcal{S}_i$. The last such call to ReadTape , for a given value of i , will be called the *defining* call for \mathbf{x}_i^D .

As in Section 2, let

$$\pi_{j,j'}^{\mathbf{x}}$$

be the function from $\{0, 1\}^n$ to $\{0, 1\}^n$ obtained by restricting the i -th block of input of π , $i \neq j$, to $\mathbf{x}[i]$, and by considering only the j' -th block of input. Recall that

$$\text{MaxPreim}(\pi) = \max_{\mathbf{x}, j, h, y} |\{x \in \{0, 1\}^n : \pi_{j,h}^{\mathbf{x}}(x) = y\}|.$$

Since the random value accessed by ReadTape during the defining call for \mathbf{x}_i^D comes uniformly at random from a set of size at least $N - q - \alpha(q)$ by Lemma 20, the defining call to ReadTape for \mathbf{x}_i^D has probability at most

$$\frac{\text{MaxPreim}(\pi_G)|\mathcal{D}_j|}{N - q - \alpha(q)} \quad (20)$$

of causing $\text{PreExisting}_{i,j}$ to occur for each $1 \leq j \leq w$, by a union bound over \mathcal{D}_j . Thus, given that $|\mathcal{D}_j| \leq q + \alpha(q)$,

$$\Pr \left[\bigvee_{1 \leq i \leq k, 1 \leq j \leq w} \text{PreExisting}_{i,j} \right] \leq \frac{wk \cdot \text{MaxPreim}(\pi_G)(q + \alpha(q))}{N - q - \alpha(q)}. \quad (21)$$

If the defining calls for vectors $\mathbf{x}_i^D, \mathbf{x}_{i'}^D$ are distinct calls, moreover, then

$$\Pr [\text{Colliding}_{i,i',j}] \leq \frac{\text{MaxPreim}(\pi_G)}{N - q - \alpha(q)}.$$

(The right-hand side is the same as (20), but with $|\mathcal{D}_j|$ replaced by 1, as there is only one bad value for the second of the two defining calls to land on.) If $\mathbf{x}_i^D, \mathbf{x}_{i'}^D$ share the same defining call, however, then

$$\Pr [\text{Colliding}_{i,i',j}] \leq \frac{\text{MaxColl}(\pi_G)}{N - q - \alpha(q)}$$

where we recall that

$$\text{MaxColl}(\pi) = \max_{\mathbf{x} \neq \mathbf{x}', j, h} |\{x \in \{0, 1\}^n : \pi_{j,h}^{\mathbf{x}}(x) = \pi_{j,h}^{\mathbf{x}'}(x)\}|.$$

Thus,

$$\Pr \left[\bigvee_{i \neq i', j} \text{Colliding}_{i,i',j} \right] \leq \frac{wk^2 \cdot \text{MaxCoPr}(\pi_G)}{N - q - \alpha(q)} \quad (22)$$

where $\text{MaxCoPr}(\pi) = \max(\text{MaxPreim}(\pi), \text{MaxColl}(\pi))$. Thus, the probability that AdaptLeft aborts during this call because $\text{SetTable}(D_j, x, y)$ is called with $D_j(x) \neq \perp$ is at most

$$\frac{wk \cdot \text{MaxPreim}(\pi_G)(q + \alpha(q))}{N - q - \alpha(q)} + \frac{wk^2 \cdot \text{MaxCoPr}(\pi_G)}{N - q - \alpha(q)} \quad (23)$$

and, by symmetry and by a union bound, the overall probability that `AdaptLeft` aborts (thus also including bad events of type $D_j^-(y) \neq \perp$) is upper bounded by

$$\frac{2wk \cdot \text{MaxPreim}(\bar{\pi})(q + \alpha(q))}{N - q - \alpha(q)} + \frac{2wk^2 \cdot \text{MaxCoPr}(\bar{\pi})}{N - q - \alpha(q)} \quad (24)$$

as per the definitions of $\text{MaxPreim}(\bar{\pi})$, $\text{MaxCoPr}(\bar{\pi})$ given in Section 4.

By symmetry, (24) holds equally well for a call to `AdaptRight`, provided `ToBeAdaptedE` has k pairs when `AdaptRight` is called.

If the distinguisher's i -th query is a left-hand query let k_i be the size of `ToBeAdaptedD` when `AdaptLeft` is called; if the distinguisher's i -th query is a right-hand query, let k_i be the size of `ToBeAdaptedE` when `AdaptRight` is called; and if the distinguisher's i -th query is a neutral query let $k_i = 0$. Then k_1, \dots, k_q are random variables (functions of the execution) such that

$$k_1 + \dots + k_q \leq \alpha(q) \quad (25)$$

with probability 1 by Lemma 17 and, by extension, such that

$$k_1^2 + \dots + k_q^2 \leq \alpha^2(q) \quad (26)$$

with probability 1. Even while the k_i 's are random variables, a little thought reveals that the total probability of aborting in `AdaptLeft` or `AdaptRight` is upper bounded by replacing respectively k and k^2 in (24) by the upper bounds $\alpha(q)$ and $\alpha^2(q)$ on these sums, cf. (25), (26). The total probability that G_3 aborts in `AdaptLeft` or `AdaptRight` is thus upper bounded by

$$\frac{2w\alpha(q)(q + \alpha(q)) \cdot \text{MaxPreim}(\bar{\pi})}{N - q - \alpha(q)} + \frac{2w\alpha(q)^2 \cdot \text{MaxCoPr}(\bar{\pi})}{N - q - \alpha(q)} \quad (27)$$

when `XtraUntglRnds` is off.

When `XtraUntglRnds` is on, we define an extra bad event:

- `PreHiti,j` occurs if and only if $\mathbf{x}_i^I[j] \in \mathcal{I}_j$ ($1 \leq i \leq k$, $1 \leq j \leq w$)

Similarly to (21), we have

$$\Pr \left[\bigvee_{1 \leq i \leq k, 1 \leq j \leq w} \text{PreHit}_{i,j} \right] \leq \frac{wk \cdot \text{MaxPreim}(\pi_G)(q + \alpha(q))}{N - q - \alpha(q)}. \quad (28)$$

If none of the events `PreHiti,j` occur each $\mathbf{y}_i^I[j]$ is randomly chosen from a set of at least $N - q - \alpha(q)$ values (Lemma 20) during the second `forall` loop of `AdaptLeft`. We note this does not preclude the possibility that $\mathbf{y}_i^I[j] = \mathbf{y}_{i'}^I[j]$ for certain pairs $i \neq i'$ and certain values of j .

We proceed to upper bound the probability that some event `PreExistingi,j` or `Collidingi,i',j` occurs given that no event `PreHiti,j` occurs. For the remainder of the argument the underlying randomness is only the randomness remaining in `I`. (We note, indeed, that this randomness isn't skewed by assuming none of the events `PreHiti,j` occurs, since these events depend only the randomness in `G`.)

We can define the “defining call” of \mathbf{x}_i^D similarly to the case `XtraUntglRnds = false`; the only difference is that we are now conditioning on `G`-values, and that the defining call is a call to `I`.

Once again, the random value accessed by ReadTape during the defining call for \mathbf{x}_i^D comes uniformly at random from a set of size at least $N - q - \alpha(q)$ by Lemma 20. Thus, the defining call for \mathbf{x}_i^D has probability at most

$$\frac{\text{MaxPreim}(\pi_I)|\mathcal{D}_j|}{N - q - \alpha(q)} \quad (29)$$

of causing $\text{PreExisting}_{i,j}$ to occur for each $1 \leq j \leq w$, by a union bound over \mathcal{D}_j . Thus, given that $|\mathcal{D}_j| \leq q + \alpha(q)$,

$$\Pr \left[\bigvee_{1 \leq i \leq k, 1 \leq j \leq w} \text{PreExisting}_{i,j} \right] \leq \frac{wk \cdot \text{MaxPreim}(\pi_I)(q + \alpha(q))}{N - q - \alpha(q)}. \quad (30)$$

If the defining calls for vectors $\mathbf{x}_i^D, \mathbf{x}_{i'}^D$ are distinct calls, moreover, then

$$\Pr [\text{Colliding}_{i,i',j}] \leq \frac{\text{MaxPreim}(\pi_I)}{N - q - \alpha(q)} \quad (31)$$

as before. Note that $\mathbf{x}_i^G \neq \mathbf{x}_{i'}^G$ implies $\mathbf{x}_i^I \neq \mathbf{x}_{i'}^I$, so there exists $1 \leq j' \leq w$ that $\mathbf{x}_i^I[j'] \neq \mathbf{x}_{i'}^I[j']$. In fact we can assume without loss of generality that $\text{ReadTape}(I_{j'}, \mathbf{x}_{i'}^I[j'], p_{I_{j'}})$ is the defining call of \mathbf{x}_i^D . (This is because changing the order in which values are sampled does not change the sampled values; indeed, the underlying random tapes are the same.) Thus

$$\Pr \left[\bigvee_{i \neq i', j} \text{Colliding}_{i,i',j} \right] \leq \frac{wk^2 \cdot \text{MaxPreim}(\pi_I)}{N - q - \alpha(q)}. \quad (32)$$

because (31) holds for each i, i', j individually.

The probability that AdaptLeft aborts during this call because $\text{SetTable}(D_j, x, y)$ is called with $D_j(x) \neq \perp$ is thus at most

$$\frac{wk \cdot \text{MaxPreim}(\pi_G)(q + \alpha(q))}{N - q - \alpha(q)} + \frac{wk \cdot \text{MaxPreim}(\pi_I)(q + \alpha(q))}{N - q - \alpha(q)} + \frac{wk^2 \cdot \text{MaxPreim}(\pi_I)}{N - q - \alpha(q)}. \quad (33)$$

as obtained by summing (28), (30), (32). The overall probability that AdaptLeft aborts during this call is thus upper bounded by

$$\frac{(4wk(q + \alpha(q)) + 2wk^2) \cdot \text{MaxPreim}(\bar{\pi})}{N - q - \alpha(q)} \quad (34)$$

by taking bad events of the type $D_j^-(y) \neq \perp$ into account. The same upper bound also holds for AdaptRight by symmetry.

Similarly to when XtraUntglRnds is off, one can then argue that the total probability that G_3 aborts in AdaptLeft or AdaptRight is upper bounded by

$$\frac{(4w\alpha(q)(q + \alpha(q)) + 2w\alpha(q)^2) \cdot \text{MaxPreim}(\bar{\pi})}{N - q - \alpha(q)} \quad (35)$$

as obtained by replacing k with $\alpha(q)$ in (34). \square

Lemma 39. *The probability that G_3 aborts for a q -query distinguisher is at most*

$$\begin{aligned}
& \frac{w(q + \alpha(q))^2}{N - q - \alpha(q)} + \frac{(q + w\alpha(q))(q + \alpha(q))}{N - q - \alpha(q)} + \frac{2w\alpha(q)(q + \alpha(q))}{N - q - \alpha(q)} \\
& + \frac{4wq \cdot \text{aboCond}_\tau(2q)}{N - 2q} && \text{if XtraMiddleRnd is on} \\
& + \frac{4w(q + \alpha(q))(\text{aboCond}_\nu(q + \alpha(q)))}{N - q - \alpha(q)} && \text{if XtraOuterRnd is on} \\
& + \frac{2w\alpha(q)(q + \alpha(q)) \cdot \text{MaxPreim}(\bar{\pi})}{N - q - \alpha(q)} + \frac{2w\alpha(q)^2 \cdot \text{MaxCoPr}(\bar{\pi})}{N - q - \alpha(q)} && \text{if XtraUntglRnds is off} \\
& + \frac{(4w\alpha(q)(q + \alpha(q)) + 2w\alpha(q)^2) \cdot \text{MaxPreim}(\bar{\pi})}{N - q - \alpha(q)} && \text{if XtraUntglRnds is on}
\end{aligned}$$

Proof. This follows by summing the bounds of Lemmas 24, 26, 28, 29, 30, 31, 33, 34 and 38. \square

E.4 Proof of Indifferentiability

We say that a distinguisher D *completes all paths* if, for every query $Z(\mathbf{x}^X) \leftarrow \mathbf{y}^H$ or query $Z^{-1}(\mathbf{y}^H) \leftarrow \mathbf{x}^X$ made by D , D subsequently and eventually makes the (possibly redundant) rwq queries

$$\begin{aligned}
& F(1, \mathbf{x}^F[1]) \rightarrow \mathbf{y}^F[1], \dots, F(w, \mathbf{x}^F[w]) \rightarrow \mathbf{y}^F[w], \quad \text{if XtraOuterRnd is on} \\
& G(1, \mathbf{x}^G[1]) \rightarrow \mathbf{y}^G[1], \dots, G(w, \mathbf{x}^G[w]) \rightarrow \mathbf{y}^G[w], \\
& I(1, \mathbf{x}^I[1]) \rightarrow \mathbf{y}^I[1], \dots, I(w, \mathbf{x}^I[w]) \rightarrow \mathbf{y}^I[w], \quad \text{if XtraMiddleRnd is on} \\
& \quad \quad \quad \vdots \quad \quad \quad \ddots \quad \quad \quad \vdots \\
& H(1, \mathbf{x}^H[1]) \rightarrow \mathbf{y}^H[1], \dots, H(w, \mathbf{x}^H[w]) \rightarrow \mathbf{y}^H[w]
\end{aligned}$$

unless the game aborts, where $\mathbf{x}^F = \mathbf{x}^X$, $\mathbf{x}^G = \nu(\mathbf{y}^F)$ if XtraOuterRnd is on, $\mathbf{x}^G = \mathbf{x}^X$ otherwise, etc. (Here we don't presume $\mathbf{y}^H = \mathbf{y}^H$, though obviously D will be able to distinguish the simulated world from the real world if $\mathbf{y}^H \neq \mathbf{y}^H$.)

For this section we presume a fixed q -query (information-theoretic) distinguisher D that completes all paths. The object is to upper bound

$$\Delta_D(G_1, G_5) = \Pr[D^{G_1} = 1] - \Pr[D^{G_5} = 1]. \quad (36)$$

The upper bounds we obtain for D will imply appropriately modified upper bounds for arbitrary distinguishers by noting that if D' is an arbitrary q_0 -query distinguisher there exists a $(q_0 + rwq_0)$ -query distinguisher D^* that completes all paths and that achieves advantage at least as good as D' .

Since D is information-theoretic we can assume without loss of generality that D is deterministic. Moreover since G_5 never aborts we can assume that if the environment aborts then D outputs 1, which can only serve to maximize (36).

Lemma 40. *Let D be as described above. Then $\Delta_D(G_1, G_2) \leq \beta(q)(q + \alpha(q))^w / (N^w - q - \alpha(q))$.*

Proof. The only difference between G_1 and G_2 is in the procedure CheckZ. Specifically, CheckZ will call Z in G_1 whereas CheckZ consults P_Z in G_2 . Hence CheckZ is answered according to p_Z in G_1 and according to P_Z in G_2 . Also, a call to CheckZ might modify P_Z in G_1 whereas a call to

CheckZ never modifies P_Z in G_2 . (We note, however, that P_Z always remains consistent with p_Z in either G_1 or G_2 : P_Z is a subset of p_Z .)

Consider two concurrent executions of G_1 and G_2 , executed on the same random tapes $p_{A_1}, \dots, p_{L_w}, p_Z$. We say the two executions *diverge* if CheckZ is ever called with arguments (\mathbf{x}, \mathbf{y}) such that $p_Z(\mathbf{x}) = \mathbf{y}$ and such that $P_Z(\mathbf{x}) = \perp$ in the G_2 -execution. We claim that as long as the two executions don't diverge in this sense, every function call is answered the same in each execution. Indeed, calls to CheckZ are answered the same (by the definition of divergence) and calls of the form $Z(\mathbf{x}), Z^{-1}(\mathbf{y})$ are also answered the same, being answered by $p_Z(\mathbf{x})$ and $p_Z^{-1}(\mathbf{y})$ (respectively) in either game regardless of the state of P_Z . Furthermore, while P_Z may contain more defined entries in G_1 than in G_2 , CheckZ, Z and Z^{-1} are the only functions to use P_Z in either game, so the differences in P_Z don't propagate to elsewhere, as we have just seen that CheckZ, Z and Z^{-1} answer function calls the same.

Next, to upper bound the probability of divergence, we start by observing that divergence is well-defined for an execution of G_2 on its own. Indeed we can say that an execution of G_2 “diverges” if CheckZ is ever called with a pair (\mathbf{x}, \mathbf{y}) such that $p_Z(\mathbf{x}) = \mathbf{y}$ and $P_Z(\mathbf{x}) = \perp$; then G_2 “diverges” on a given set of random tapes if and only if the concurrent executions of G_1 and G_2 “diverge” in the sense above on the same set of random tapes. We can therefore focus on the probability that an execution of G_2 diverges, without reference to G_1 .

To upper bound the probability of divergence in G_2 note that if CheckZ is called on arguments (\mathbf{x}, \mathbf{y}) such that $P_Z(\mathbf{x}) = \perp$ then $p_Z(\mathbf{x})$ hasn't been read yet. Indeed $p_Z(\mathbf{x})$ is only read in Z and Z^{-1} , at which point P_Z is updated with the read value. Hence, the value $p_Z(\mathbf{x})$ comes uniformly at random from a set of size at least $N^w - q - \alpha(q)$ by Lemma 25. (A subtlety here is that we are *not* assuming that divergence hasn't already occurred. Indeed such an assumption would skew the conditional distribution of the *unread* values in p_Z , as the truth of this assumption depends on unread values. Instead we are simply doing a union bound over all calls to CheckZ.) As CheckZ is called with at most $\beta(q)(q + \alpha(q))^w$ distinct arguments pairs during a G_2 -execution by Lemma 23, the probability of divergence in G_2 (and hence between G_1 and G_2) is at most $\beta(q)(q + \alpha(q))^w / (N^w - q - \alpha(q))$. \square

Lemma 41. *Let D be as described above. Then $\Delta_D(G_2, G_5) \leq \Delta_D(G_3, G_5)$.*

Proof. Since G_5 never aborts we can assume that D outputs 1 when the environment aborts. However the only difference between G_2 and G_3 is that more abort conditions are added in G_3 (specifically, two executions with the same random tapes will unfold exactly the same in G_2 and G_3 except for the eventuality that G_3 might abort while G_2 does not) so D 's probability of outputting 1 is at least as large as D 's probability of outputting 1 in G_1 , which proves the claim. \square

To upper bound $\Delta_D(G_3, G_5)$ we upper bound $\Delta_D(G_3, G_4)$ and $\Delta_D(G_4, G_5)$ separately and apply the triangle inequality. The crucial transition is the transition from G_3 to G_4 , for which we apply a randomness mapping argument. This requires some extra definitions.

Given an execution of G_3 the *footprint* of that execution is that subset of the random tapes actually read during that execution. More precisely, identify a permutation $p : \{0, 1\}^n \rightarrow \{0, 1\}^n$ with the set of pairs $(x, y) \in \{0, 1\}^n \times \{0, 1\}^n$ such that $p(x) = y$. A *partial permutation* is obtained by dropping pairs from p . Then the *footprint* of an execution of G_3 (with our fixed distinguisher D) on random tapes $p_Z, p_{A_1}, \dots, p_{L_w}$ is the set of *partial* permutations $\tilde{p}_Z, \tilde{p}_{A_1}, \dots, \tilde{p}_{L_w}$ where

$$\tilde{p}_Z \subseteq p_Z, \tilde{p}_{A_1} \subseteq p_{A_1}, \dots, \tilde{p}_{L_w} \subseteq p_{L_w}$$

and where \tilde{p}_Z and \tilde{p}_{T_j} contain only those entries of respectively p_Z and p_{T_j} that were actually read from (respectively again) p_Z and p_{T_j} during the execution. (Here an entry (x, y) is “read” from p if and only if either $p(x)$ or $p^{-1}(y)$ was accessed during the execution. In G_3 such reads occur only in the function `ReadTape`.)

Note that since D is deterministic, a footprint contains all the information necessary to re-create a particular execution (and in particular D 's output). In particular if we run D on a footprint $\tilde{p}_Z, \tilde{p}_{A_1}, \dots, \tilde{p}_{L_w}$ then (i) the execution will never “leave” that footprint (i.e., request values from the random tables that are outside the footprint), and (ii) all values included in the footprint will be read at some point during the execution. Also note that two distinct footprints are always incompatible in the sense that they don't have a common extension $p_Z, p_{A_1}, \dots, p_{L_w}$. This follows by points (i) and (ii) just mentioned. We also emphasize that the set of possible footprints is a function of D .

We say that a footprint is *accepting* for G_3 if D outputs 1 for the execution corresponding to that footprint. Moreover we say that a footprint is *non-aborting* for G_3 if the corresponding G_3 execution doesn't abort.

Letting $|\tilde{p}|$ denote the number of input-output pairs in the partial permutation \tilde{p} , the probability of obtaining a given footprint $\tilde{p}_Z, \tilde{p}_{A_1}, \dots, \tilde{p}_{L_w}$, taken over the random choice of $p_Z, p_{A_1}, \dots, p_{L_w}$, is obviously

$$\left(\prod_{\ell=0}^{|\tilde{p}_Z|-1} \frac{1}{N^w - \ell} \right) \left(\prod_T \prod_{j=1}^w \prod_{\ell=0}^{|\tilde{p}_{T_j}|-1} \frac{1}{N - \ell} \right) \quad (37)$$

since this is the probability that a randomly chosen $p_Z, p_{A_1}, \dots, p_{L_w}$ is an extension of $\tilde{p}_Z, \tilde{p}_{A_1}, \dots, \tilde{p}_{L_w}$. (Of course, this presumes that $\tilde{p}_Z, \tilde{p}_{A_1}, \dots, \tilde{p}_{L_w}$ is a real footprint and not just some set of arbitrarily restricted permutations.) In particular, for instance, the probability that D accepts is the sum of (37) taken over all accepting footprints.

We likewise define the footprint of an execution in G_4 with respect to the random tapes q_{A_1}, \dots, q_{L_w} as the set of entries read from those tapes. Thus a footprint in G_4 is again a set of partially defined permutations. The probability of obtaining a given footprint $\tilde{q}_{A_1}, \dots, \tilde{q}_{L_w}$ in G_4 is

$$\prod_T \prod_{j=1}^w \prod_{\ell=0}^{|\tilde{q}_{T_j}|-1} \frac{1}{N - \ell}. \quad (38)$$

since this is the probability that randomly chosen tapes q_{A_1}, \dots, q_{L_w} are an extension of $\tilde{q}_{A_1}, \dots, \tilde{q}_{L_w}$. We likewise talk of *accepting* and *non-aborting* footprints in G_4 .

For $i = 3, 4$ let FP_i denote the set of all G_i footprints and $\text{FP}_i^* \subseteq \text{FP}_i$ the set of all non-aborting G_i footprints. We write

$$\Pr_{G_i}[\omega]$$

for the probability of obtaining a given footprint $\omega \in \text{FP}_i$, $i = 3, 4$ (cf. (37), (38)), and

$$\Pr_{G_i}[\mathcal{A}] = \sum_{\omega \in \mathcal{A}} \Pr_{G_i}[\omega]$$

for any $\mathcal{A} \subseteq \text{FP}_i$. Note that $\Pr_{G_i}[\text{FP}_i \setminus \text{FP}_i^*]$ is the probability of abortion in game G_i for $i = 3, 4$ since FP_i^* is the set of non-aborting footprints.

We also let $\mathcal{A}_i \subseteq \text{FP}_i$ be the set of accepting footprints in game G_i and let $\mathcal{A}_i^* = \mathcal{A}_i \cap \text{FP}_i^*$ be the set of accepting footprints that are also non-aborting. Note that $\text{FP}_i \setminus \text{FP}_i^* \subseteq \mathcal{A}_i$ by our assumption that D accepts whenever the game aborts and, by the same token, $\text{FP}_i \setminus \text{FP}_i^* = \mathcal{A}_i \setminus \mathcal{A}_i^*$. (In other words, the set of non-good accepting footprints is also the set of all non-good footprints.)

The key to the randomness mapping argument is an injection $\zeta : \text{FP}_3^* \leftarrow \text{FP}_4^*$ such that $\omega \in \text{FP}_3^*$ is accepting if and only if $\zeta(\omega) \in \text{FP}_4^*$ is accepting and such that

$$\Pr_{G_4}[\zeta(\omega)] \geq \Pr_{G_3}[\omega] (1 - 1/N^w) \quad (39)$$

for all $\omega = (\tilde{p}_Z, \tilde{p}_{A_1}, \dots, \tilde{p}_{L_w}) \in \text{FP}_3^*$. Then

$$\begin{aligned} \Delta_D(G_3, G_4) &= \Pr[D^{G_3} \leftarrow 1] - \Pr[D^{G_4} \leftarrow 1] \\ &= \Pr_{G_3}[\mathcal{A}_3] - \Pr_{G_4}[\mathcal{A}_4], \\ &= \Pr_{G_3}[\mathcal{A}_3 \setminus \mathcal{A}_3^*] + \Pr_{G_3}[\mathcal{A}_3^*] - \Pr_{G_4}[\mathcal{A}_4 \setminus \mathcal{A}_4^*] - \Pr_{G_4}[\mathcal{A}_4^*] \\ &= \Pr_{G_3}[\text{FP}_3 \setminus \text{FP}_3^*] + \Pr_{G_3}[\mathcal{A}_3^*] - \Pr_{G_4}[\text{FP}_4 \setminus \text{FP}_4^*] - \Pr_{G_4}[\mathcal{A}_4^*] \\ &\leq \Pr_{G_3}[\text{FP}_3 \setminus \text{FP}_3^*] + \frac{1}{N^w} - \Pr_{G_4}[\text{FP}_4 \setminus \text{FP}_4^*] \end{aligned} \quad (40)$$

where the inequality holds by (39) (recall that (39) holds for all $\omega \in \mathcal{A}_3^*$ and that $\zeta : \mathcal{A}_3^* \leftarrow \mathcal{A}_4^*$ is injective—see Lemma 47 below for more details).

The definition of $\zeta : \text{FP}_3^* \rightarrow \text{FP}_4^*$ is fairly simple: for $\omega = (\tilde{p}_Z, \tilde{p}_{A_1}, \dots, \tilde{p}_{L_w}) \in \text{FP}_3^*$ we define $\zeta(\omega) = (\tilde{q}_{A_1}, \dots, \tilde{q}_{L_w})$ by

$$\tilde{q}_{T_j} = \{(x, y) \in \{0, 1\}^n \times \{0, 1\}^n : T_j(x) = y\}$$

where T_j refers to the table T_j as it stands at the end of the G_3 -execution. (In a nutshell, thus, $\tilde{q}_{T_j} = T_j$.)

The map ζ is obviously well-defined because ω is a footprint. But, to recapitulate, the following properties of ζ must still be proved:

- that $\zeta(\omega) \in \text{FP}_4^*$ for all $\omega \in \text{FP}_3^*$ (i.e., that the image of a good footprint is a good footprint)
- that ζ is injective
- that $\zeta(\omega)$ is accepting in G_4 if and only if ω is accepting for G_3
- that $\Pr_{G_4}[\zeta(\omega)] \geq \Pr_{G_3}[\omega](1 - 1/N^w)$ for all $\omega \in \text{FP}_3^*$

We build up to these four points in a series of lemmas. We recall the terminology of *T-columns*, *matching columns* and *completed paths* from Section 6. We note that Lemma 1 implies that both completed paths and *T-columns* are “persistent”: once a column or completed path appears, it remains unchanged until the end of the execution. Moreover, because the elements of $\overline{\pi}$ are permutations, each column is part of at most one completed path, and matching columns are in the same (if any) completed path.

Lemma 42. *At the end of each non-aborted query cycle in G_3 (and, in particular, at the end of each non-aborted G_3 -execution) every pair of matching B- and C-columns is part of a completed path if XtraMiddleRnd is on, every A-column is part of a completed path if XtraMiddleRnd is off.*

Proof. This is a corollary of Lemma 16. □

Lemma 43. *At the end of each non-aborted execution of G_3 there exists a completed path with endpoints (\mathbf{x}, \mathbf{y}) for each pair (\mathbf{x}, \mathbf{y}) such that either of the calls $Z(\mathbf{x}) \leftarrow \mathbf{y}$ or $Z^{-1}(\mathbf{y}) \leftarrow \mathbf{x}$ has occurred at some point in the execution.*

Proof. For calls to Z, Z^{-1} made by the simulator this follows from the fact that the simulator only calls Z or Z^{-1} when it has already decided to complete a path in G_3 (cf. RequestG, RequestU $^{-1}$ and AdaptLeft, AdaptRight). For calls to Z, Z^{-1} made by the distinguisher this follows from the assumption that the distinguisher completes all paths and from Lemma 42 (since the “path” built by the distinguisher will in particular go through the middle detect zone and hence, by Lemma 42, have endpoints that are compatible with Z). \square

Lemma 44. *Let the map ζ be defined on FP_3^* as above. Then $\zeta(\omega) \in \text{FP}_4^*$ for each footprint $\omega \in \text{FP}_3^*$ and $\zeta(\omega)$ is accepting for G_4 if and only if ω is accepting for G_3 . Finally, ζ is injective.*

Proof. Let $\omega = (\tilde{p}_Z, \tilde{p}_{A_1}, \dots, \tilde{p}_{L_w}) \in \text{FP}_3^*$ and let $\zeta(\omega) = (\tilde{q}_{A_1}, \dots, \tilde{q}_{L_w})$.

For syntactical convenience, let $\lambda = (q_{A_1}, \dots, q_{L_w})$ be some extension of $\zeta(\omega)$; we will first prove that executing G_4 on random tape λ gives a non-aborting execution, and then argue that this execution’s footprint is in fact $\zeta(\omega)$.

Consider concurrent executions of G_3 and G_4 with random tapes ω and λ respectively. (Note that ω is a partial random tape for G_3 whereas λ is a complete random tape G_4 , but since ω is a G_3 -footprint by assumption this poses no problem.)

Note that only two differences occur between G_3 and G_4 : (i) the internals of the functions Z, Z^{-1} ; (ii) the fact that G_4 reads from the random tapes q_{A_1}, \dots, q_{L_w} , whereas G_3 reads from the tapes $p_Z, p_{A_1}, \dots, p_{F_w}$. Moreover, while calls to Z/Z^{-1} affect the table P_Z —and only this table— P_Z is not actually used anywhere outside of Z, Z^{-1} . Consequently we say that the two executions *diverge* if (a) there either occurs a call Z to Z^{-1} that is not answered the same in either game (where the call could be made by the simulator or by the distinguisher either) or if (b) a call to $\text{ReadTape}(T_j, x, p_{T_j})$ or $\text{ReadTape}(T_j^{-1}, y, p_{T_j}^{-1})$ occurs in G_3 such that, respectively, $\tilde{p}_{T_j}(x) \neq q_{T_j}(x)$ or $\tilde{p}_{T_j}^{-1}(y) \neq q_{T_j}^{-1}(y)$. Thus if divergence doesn’t occur in this technical sense, the two executions proceed identically up to the internals of Z, Z^{-1} .

Firstly divergence of type (b) cannot occur because every time a value is read from p_{T_j} in G_3 this value ends up in T_j and hence, by definition of ζ , in \tilde{q}_{T_j} (which is a subset of q_{T_j}). Secondly divergence of type (a) cannot occur by Lemma 43. Indeed this lemma implies that for each call $Z(\mathbf{x}) \rightarrow \mathbf{y}$ or $Z^{-1}(\mathbf{y}) \rightarrow \mathbf{x}$ occurring in G_3 there eventually exists a completed path with endpoints (\mathbf{x}, \mathbf{y}) in the tables T_j ; by definition of ζ the relevant entries of T_j appear in \tilde{q}_{T_j} .

Since the two executions don’t diverge it follows that G_4 does not abort on λ (as G_3 doesn’t abort on ω), and that G_4 accepts if and only if G_3 accepts.

We next show that $\zeta(\omega)$ is a G_4 -footprint. To see that G_4 never reads values *outside* of $\zeta(\omega)$ note that all values read by G_4 in the tables q_{T_j} were either: (i) simultaneously read by G_3 in p_{T_j} or (ii) read as part of a call to Z or Z^{-1} ; by Lemma 42 it follows that the corresponding entry was in T_j at the end of the G_3 -execution. Hence, G_4 never reads outside of $\zeta(\omega)$ when executed on a random tape λ . Then to see that G_4 reads *every* value in $\zeta(\omega)$ note that each entry entered into a T_j table in G_3 is either entered by a call to ReadTape (in which case the same entry is read from $\zeta(\omega)$ in G_4 , as the two executions don’t diverge) or else, in the case of D_j and E_j , placed there while completing a path. But each completed path is completed only after a corresponding call to Z/Z^{-1} , and the matching entry from q_{D_j} or q_{E_j} will be read during the corresponding call to Z/Z^{-1} that occurs in G_4 .

Finally we show that ζ is injective. For this we build an inverse map. Given $\zeta(\omega) \in \text{FP}_4^*$, we define a footprint ω' for G_3 like so: we run G_3 with “blank” random tapes concurrently to G_4 with random tape $\zeta(\omega)$; when a call $\text{ReadTape}(T_j^\pm, \cdot, \cdot)$ occurs we set the corresponding entry of p_{T_j} according to the same entry in \tilde{q}_{T_j} , and when a call to Z or Z^{-1} occurs such that p_Z is undefined at that point, we set the corresponding entry of p_Z according to the answer of that call in G_4 . Since the execution of G_4 is “the same” as when G_4 and G_3 were executed concurrently with random tapes λ and ω respectively, it’s easy to see that the partial random tape obtained for G_3 by this inverse map must be compatible²⁵ with ω and hence must be ω (since two distinct footprints are never compatible), and hence that ζ is injective. \square

Lemma 45. *Let $\omega = (\tilde{p}_Z, \tilde{p}_{A_1}, \dots, \tilde{p}_{L_w}) \in \text{FP}_3^*$ and let $\zeta(\omega) = (\tilde{q}_{A_1}, \dots, \tilde{q}_{L_w}) \in \text{FP}_4^*$. Then $|\tilde{q}_{T_j}| = |\tilde{p}_{T_j}|$ for $T \notin \{D, E\}$ and there exist nonnegative integers a, b such that $a + b = |\tilde{p}_Z|$ and such that $|\tilde{q}_{D_j}| = |\tilde{p}_{D_j}| + a$, $|\tilde{q}_{E_j}| = |\tilde{p}_{E_j}| + b$ for $1 \leq j \leq w$.*

Proof. Recall that $|\tilde{q}_{T_j}| = |T_j|$ where $|T_j|$ is the number of entries in T_j at the end of the G_3 -execution on tape (or footprint) ω . Thus, the fact that $|\tilde{q}_{T_j}| = |\tilde{p}_{T_j}|$ for $T \notin \{D, E\}$ follows from the fact that $\text{SetTable}(T_j, \cdot, \cdot)$ is never called for $T \notin \{D, E\}$ and, hence, that the number of entries in T_j is the same as the number of calls to $\text{ReadTape}(T_j^{\pm 1}, \cdot, \cdot)$ for $T \notin \{D, E\}$.

To prove the second part of the claim let a be the number of times that a pair $(\mathbf{x}^G, \mathbf{y}^U)$ is added to ToBeAdapted_D in RequestU^{-1} and let b be the number of times that a pair $(\mathbf{x}^V, \mathbf{y}^H)$ is added to ToBeAdapted_E in RequestV , in the G_3 -execution on random tape ω . Note there is a bijection between the final set of completed paths and the entries in p_Z (in more detail, each A -column or each matching pair of B - and C -columns causes the simulator to complete a path in which it calls Z or Z^{-1} , and each call to Z or Z^{-1} has an associated completed path (containing a unique A -column or a unique matching pair of B - and C -columns) because the distinguisher completes all paths). Moreover, as noted in the proof of Lemma 42, there is a bijection between all A -columns (if XtraMiddleRnd is off) or all matching pairs of B - and C -columns (if XtraMiddleRnd is on) and the set of all pairs $(\mathbf{x}^G, \mathbf{y}^U)$ or $(\mathbf{x}^V, \mathbf{y}^H)$ added respectively to ToBeAdapted_D and ToBeAdapted_E ; hence $a + b = |p_Z|$. But a is also the number of times $\text{SetTable}(D_j, \cdot, \cdot)$ is called for each $1 \leq j \leq w$, while b is the number of times $\text{SetTable}(E_j, \cdot, \cdot)$ is called for each $1 \leq j \leq w$. Since SetTable and ReadTape both abort rather than overwrite an entry, this completes the proof. \square

Lemma 46. *We have $\Pr_{G_4}[\omega] \geq \Pr_{G_3}[\zeta(\omega)](1 - 1/N^w)$ for every $\omega \in \text{FP}_3^*$.*

Proof. Let a and b be as in Lemma 45. Then by (37), (38), and keeping other notations of Lemma 45,

$$\begin{aligned} \Pr_{G_4}[\zeta(\omega)] / \Pr_{G_3}[\omega] &= \prod_{j=0}^w \left(\prod_{h=0}^{a-1} \frac{1}{N - |\tilde{p}_{D_j}| - h} \cdot \prod_{h=0}^{b-1} \frac{1}{N - |\tilde{p}_{E_j}| - h} \right) / \prod_{\ell=0}^{a+b-1} \frac{1}{N^w - \ell} \\ &\geq \prod_{j=0}^w \left(\prod_{h=0}^{a-1} \frac{1}{N - h} \cdot \prod_{h=0}^{b-1} \frac{1}{N - h} \right) / \prod_{\ell=0}^{a+b-1} \frac{1}{N^w - \ell} \\ &= \prod_{h=0}^{a-1} \frac{N^w}{(N - h)^w} \cdot \prod_{h=0}^{b-1} \frac{N^w}{(N - h)^w} \cdot \prod_{\ell=0}^{a+b-1} \frac{N^w - \ell}{N^w} \end{aligned} \quad (41)$$

²⁵ Two partial random tapes are “compatible” if they have a common extension.

As $a + b = |p_Z|$, we have

$$a + b \leq q + \alpha(q) \leq N^{w-1}$$

by Lemma 25, where the second inequality holds without loss of generality. (Indeed, the security bound of Theorem 1 is void otherwise.) Thus

$$(N - 1)^w = N^w(1 - 1/N)^w \leq N^w(1 - 1/N) = N^w - N^{w-1} \leq N - a - b$$

and, more generally,

$$(N - 1)^w \leq N - \ell \tag{42}$$

for all $0 \leq \ell \leq a + b$. In particular, we can assume $a, b > 0$ since if $a = 0$ or $b = 0$ we have that (41) is lower bounded by 1, by (42). Thus, continuing from (41),

$$\begin{aligned} \Pr_{G_4}[\zeta(\omega)] / \Pr_{G_3}[\omega] &\geq \frac{N^w - 1}{N^w} \cdot \prod_{h=1}^{a-1} \frac{N^w}{(N-h)^w} \cdot \prod_{h=1}^{b-1} \frac{N^w}{(N-h)^w} \cdot \prod_{\ell=2}^{a+b-1} \frac{N^w - \ell}{N^w} \\ &\geq \frac{N^w - 1}{N^w} \cdot \prod_{h=1}^{a-1} \frac{N^w}{(N-h)^w} \cdot \prod_{h=1}^{b-1} \frac{N^w}{(N-h)^w} \cdot \prod_{\ell=2}^{a+b-1} \frac{(N-1)^w}{N^w} \\ &= \frac{N^w - 1}{N^w} \prod_{h=1}^{a-1} \frac{(N-1)^w}{(N-h)^w} \cdot \prod_{h=1}^{b-1} \frac{(N-1)^w}{(N-h)^w} \\ &\geq \frac{N^w - 1}{N^w} = 1 - 1/N^w \end{aligned} \tag{43}$$

as claimed. \square

Lemma 47. *We have*

$$\Pr_{G_3}[\mathcal{A}_3^*] - \Pr_{G_4}[\mathcal{A}_4^*] \leq \frac{1}{N^w}$$

where $\mathcal{A}_i^* \subseteq \text{FP}_i^*$ denotes the set of accepting, non-aborting footprints in G_i .

Proof. This follows directly by Lemma 46 since ζ maps elements of \mathcal{A}_3^* to elements of \mathcal{A}_4^* , and is injective (Lemma 44). More precisely, the fact that

$$\Pr_{G_4}[\mathcal{A}_4^*] \geq \Pr_{G_3}[\mathcal{A}_3^*] \left(1 - \frac{1}{N^w}\right)$$

implies

$$\Pr_{G_3}[\mathcal{A}_3^*] - \Pr_{G_4}[\mathcal{A}_4^*] \leq \frac{1}{N^w} \Pr_{G_3}[\mathcal{A}_3^*] \leq \frac{1}{N^w}.$$

\square

Lemma 48. *We have*

$$\Delta_D(G_3, G_4) \leq \Pr_{G_3}[\text{FP}_3 \setminus \text{FP}_3^*] + \frac{1}{N^w} - \Pr_{G_4}[\text{FP}_4 \setminus \text{FP}_4^*]$$

for any q -query distinguisher D that completes all paths.

Proof. This was shown in (40), above, where we used the fact that $\Pr_{G_3}[\mathcal{A}_3^*] - \Pr_{G_4}[\mathcal{A}_4^*] \leq \frac{1}{N^w}$, as shown in Lemma 47. \square

Lemma 49. *We have*

$$\Delta_D(G_4, G_5) \leq \Pr_{G_4}[\text{FP}_4 \setminus \text{FP}_4^*]$$

for any q -query distinguisher D that completes all paths.

Proof. Consider a non-aborting execution of G_4 . In such an execution queries to Z/Z^{-1} are answered entirely according to the random tapes q_{A_1}, \dots, q_{L_w} , as in G_5 . Also, since entries in T_j for $T \notin \{D, E\}$ are never adapted, queries to T^\pm are answered according to the random tapes q_{T_j} . Lastly we argue that queries to $D^\pm(\cdot, j)$ and $E^\pm(\cdot, j)$ are answered according to the tables q_{D_j} and q_{E_j} . This is equivalent to showing that entries written in D_j and E_j always agree with q_{D_j} and q_{E_j} , respectively. We argue the latter by induction (on the number of entries added to D_j and E_j). For entries in D_j, E_j that are directly read from the random tapes the claim is obvious. For an entry in (say) D_j that is adapted, the claim follows from the fact that every adapted value in D_j is part of a completed path, from the induction hypothesis, from the fact that queries to Z/Z^{-1} are answered according to the random tapes q_{T_j} , and from the fact that when the D_j query is adapted all queries in columns $T \neq D$ from the completed path are already made.

It follows that queries in both G_4 and G_5 are answered according the tables q_{T_j} , with the only difference that G_4 might abort. So $\Delta_D(G_4, G_5) \leq \Pr[\text{FP}_4 \setminus \text{FP}_4^*]$. \square

(One can, in fact, argue that $\Delta_D(G_4, G_5) = \Pr[\text{FP}_4 \setminus \text{FP}_4^*]$ from the fact that D outputs 1 if abort occurs. But we only need the inequality.)

Lemma 50. *We have*

$$\begin{aligned} & \Delta_D(G_1, G_5) \\ & \leq \frac{\beta(q)(q + \alpha(q))^w}{N^w - q - \alpha(q)} + \frac{1}{N^w} \\ & \quad + \frac{w(q + \alpha(q))^2}{N - q - \alpha(q)} + \frac{(q + w\alpha(q))(q + \alpha(q))}{N - q - \alpha(q)} + \frac{2w\alpha(q)(q + \alpha(q))}{N - q - \alpha(q)} \\ & \quad + \frac{4wq \cdot \text{aboCond}_\tau(2q)}{N - 2q} && \text{if XtraMiddleRnd is on} \\ & \quad + \frac{2w(q + \alpha(q))(\text{aboCond}_\nu(q + \alpha(q)))}{N - q - \alpha(q)} && \text{if XtraOuterRnd is on} \\ & \quad + \frac{2w\alpha(q)(q + \alpha(q)) \cdot \text{MaxPreim}(\bar{\pi})}{N - q - \alpha(q)} + \frac{2w\alpha(q)^2 \cdot \text{MaxCoPr}(\bar{\pi})}{N - q - \alpha(q)} && \text{if XtraUntglRnds is off} \\ & \quad + \frac{(4w\alpha(q)(q + \alpha(q)) + 2w\alpha(q)^2) \cdot \text{MaxPreim}(\bar{\pi})}{N - q - \alpha(q)} && \text{if XtraUntglRnds is on} \end{aligned}$$

for any q -query distinguisher D that complete all paths.

Proof. We have, using lemmas 40, 41, 48 and 49,

$$\begin{aligned}
\Delta_D(G_1, G_5) &= \Delta_D(G_1, G_2) + \Delta_D(G_2, G_5) \\
&\leq \Delta_D(G_1, G_2) + \Delta_D(G_3, G_5) \\
&= \Delta_D(G_1, G_2) + \Delta_D(G_3, G_4) + \Delta_D(G_4, G_5) \\
&\leq \frac{\beta(q)(q + \alpha(q))^w}{N^w - q - \alpha(q)} + \left(\Pr_{G_3}[\text{FP}_3 \setminus \text{FP}_3^*] + \frac{1}{N^w} - \Pr_{G_4}[\text{FP}_4 \setminus \text{FP}_4^*] \right) + \Pr_{G_4}[\text{FP}_4 \setminus \text{FP}_4^*] \\
&= \frac{\beta(q)(q + \alpha(q))^w}{N^w - q - \alpha(q)} + \Pr_{G_3}[\text{FP}_3 \setminus \text{FP}_3^*] + \frac{1}{N^w}.
\end{aligned}$$

The bound then follows by Lemma 39, since $\Pr_{G_3}[\text{FP}_3 \setminus \text{FP}_3^*]$ is the probability of abortion in G_3 . \square

E.5 Leftovers

The security bound in Theorem 1 is easily seen to be the bound of Lemma 50 modulo some simplifications to make the bound more compact. In order to prove Theorem 1 it remains to argue the simulator's query complexity and running time.

Lemma 51. *The simulator described in game G_1 of Figs. 3–7 can be implemented in running time $O(w(q + \alpha(q))^w)$.*

Proof. Once glaring inefficiencies are removed from the simulator (many of which are caused by our efforts to avert nested loops), one can see that the simulator's running time is dominated by steps in which it “mix-and-matches” all queries within a given column, either to check for abort conditions or (more crucially) for path completion triggers.

A little more precisely, such “mix-and-matching” can be arranged to always take place such that at least one query of the column is fixed (being in fact the last query created for that column); this means that each mix-and-match enumeration takes time $O((q + \alpha(q))^{w-1})$ in the worst case by Lemma 20. Moreover there will be at most $5w(q + \alpha(q))$ mix-and-match enumerations over the whole course of the experiment, given that at most 5 columns (F, G, B, C and H) trigger mix-and-match enumerations. This leads to the stated bound. \square

Lemma 52. *The simulator described in game G_1 of Figs. 3–7 can be implemented to have query complexity at most $\beta(q)$.*

Proof. Note the simulator in game G_1 of Figs. 3–7 queries Z/Z^{-1} without checking to see if it possibly knows the answers to these queries from its previous queries to Z/Z^{-1} . For the following argument, we imagine that the simulator maintains a global table of its queries to Z , together with their inverses, such as to never make a redundant query to Z . (This obviously does not change anything from the distinguisher's perspective, nor for the simulator, except for the change itself.)

Each query to Z or Z^{-1} either occurs in CheckZ or else because the simulator is completing a middle-triggered path. In either case, the query can thus uniquely be associated with a set of matching queries in the left-hand outer detect zone. But there are at most $\beta(q)$ such matching query sets by Lemma 20. \square

Further remarks. Space complexity of the simulator. The table mentioned in the proof of Lemma 52 is in fact our simulator's dominating space requirement; the space complexity of our simulator is thus

$$O(\beta(q))$$

if we assume that some efficient data structure is used to implement the table. (We omit a factor of nw for simplicity.)

If we remove the table, the space complexity drops down to $O(q)$ (factor of nwr omitted), i.e., the space needed to store the S -box permutation tables. On the other hand one can check that the query complexity climbs to

$$2(\beta(q) + (2q)^w) \tag{44}$$

because one can establish a 2-to-1 surjection from the set of queries to the union of (i) matching query sets in the left-hand detect zone and (ii) matching query sets in the right-hand detect zone. (The factor of 2 occurs because the same query can occur once in RequestF and once in RequestU⁻¹, during a left-hand query cycle, or once in RequestG⁻¹ and once in RequestV, during a right-hand query cycle.)

However, if we change the design of the network by keeping the left- and right-hand detect zones symmetric the term $(2q)^w$ in (44) is replaced by $\beta(q)$, making the query complexity

$$4\beta(q)$$

instead of $2(\beta(q) + (2q)^w)$.

In summary, one can make potentially significant savings in space complexity (from $O(\beta(q))$ to $O(q)$) at the cost of only moderately worse query complexity (from $\beta(q)$ to $4\beta(q)$) by maintaining symmetry between the two left- and right-hand outer detect zones. On the other hand, maintaining symmetry does not improve the security of our simulator (and obviously, costs an extra round).

<p>Game G_1 G_2, G_3</p> <p>random tapes: $p_{A_1}, \dots, p_{L_w}, p_Z$</p> <p>global static: XtraOuterRnd, XtraMiddleRnd, XtraUntglRnds</p> <p>if $Table(x) \neq \perp$ then abort if $Table^{-1}(y) \neq \perp$ then abort $Table(x) \leftarrow y$ $Table^{-1}(y) \leftarrow x$</p> <p>private procedure ReadTape($Table, x, p$) $y \leftarrow p(x)$ SetTable($Table, x, y$) return y</p> <p>public procedure Z(x) if $P_Z(x) = \perp$ then $y \leftarrow$ ReadTape(P_Z, x, p_Z) if $\exists j, Y_j^{-1}(y[j]) \neq \perp$ then abort // G_3 return $P_Z(x)$</p> <p>public procedure $Z^{-1}(y)$ if $P_Z^{-1}(y) = \perp$ then $x \leftarrow$ ReadTape($P_Z^{-1}, y, p_{J^{-1}}$) if $\exists j, X_j(x[j]) \neq \perp$ then abort // G_3 return $P_Z^{-1}(y)$</p> <p>private procedure CheckZ(x, y) return $P_Z(x) = y$ // G_2, G_3 return $Z(x) = y$</p> <p>private procedure BlockDefined(T, z, σ) forall $k \in \{1, \dots, w\}$ do if $(T_k^\sigma(z[k]) = \perp \wedge z[k] \notin \text{ToBeAssigned}_{T_k}^\sigma)$ return false return true</p>	<p>public procedure D(j, x) if $D_j(x) = \perp$ then ReadTape(D_j, x, p_{D_j}) return $D_j(x)$</p> <p>public procedure $D^{-1}(j, y)$ if $D_j^{-1}(y) = \perp$ then ReadTape($D_j^{-1}, y, p_{D_j^{-1}}$) return $D_j^{-1}(y)$</p> <p>public procedure E(j, x) if $E_j(x) = \perp$ then ReadTape(E_j, x, p_{E_j}) return $E_j(x)$</p> <p>public procedure $E^{-1}(j, y)$ if $E_j^{-1}(y) = \perp$ then ReadTape($E_j^{-1}, y, p_{E_j^{-1}}$) return $E_j^{-1}(y)$</p> <p>public procedure I(j, x) if $I_j(x) = \perp$ then ReadTape(I_j, x, p_{I_j}) return $I_j(x)$</p> <p>public procedure $I^{-1}(j, y)$ if $I_j^{-1}(y) = \perp$ then ReadTape($I_j^{-1}, y, p_{I_j^{-1}}$) return $I_j^{-1}(y)$</p> <p>public procedure J(j, x) if $J_j(x) = \perp$ then ReadTape(J_j, x, p_{J_j}) return $J_j(x)$</p> <p>public procedure $J^{-1}(j, y)$ if $J_j^{-1}(y) = \perp$ then ReadTape($J_j^{-1}, y, p_{J_j^{-1}}$) return $J_j^{-1}(y)$</p> <p>public procedure K(j, x) if $K_j(x) = \perp$ then ReadTape(K_j, x, p_{K_j}) return $K_j(x)$</p> <p>public procedure $K^{-1}(j, y)$ if $K_j^{-1}(y) = \perp$ then ReadTape($K_j^{-1}, y, p_{K_j^{-1}}$) return $K_j^{-1}(y)$</p> <p>public procedure L(j, x) if $L_j(x) = \perp$ then ReadTape(L_j, x, p_{L_j}) return $L_j(x)$</p> <p>public procedure $L^{-1}(j, y)$ if $L_j^{-1}(y) = \perp$ then ReadTape($L_j^{-1}, y, p_{L_j^{-1}}$) return $L_j^{-1}(y)$</p>
--	---

Fig. 3. Games G_1 through G_3 , first set of procedures. Statements commented by // G_2 or // G_2, G_3 (in red) appear only in G_2 or in G_2 and G_3 , respectively. The simulator is implemented by game G_1 , excepted procedures Z and Z^{-1} (to which the simulator only has oracle access).

<pre> private procedure BackwardOutside(T, j, y) $x \leftarrow \text{ReadTape}(T_j^{-1}, y, p_{T_j}^{-1})$ if ($\exists \mathbf{x}$ s.t. $\mathbf{x}[j] = x \wedge$ $P_Z(\mathbf{x}) \neq \perp$) then abort // G_3 public procedure $F^{-1}(j, y)$ if $F_j^{-1}(y) = \perp$ then BackwardOutside(F, j, y) return $F_j^{-1}(y)$ public procedure $F(j, x)$ if $F_j(x) = \perp$ then $y \leftarrow \text{ReadTape}(F_j, x, p_{F_j})$ if ($\exists \mathbf{y}$ s.t. $\mathbf{y}[j] = y \wedge$ BlockDefined($F, \mathbf{y}, -$) \wedge BlockDefined($G, \nu(\mathbf{y}), +$)) then abort return $F_j(x)$ public procedure $G^{-1}(j, y)$ if $G_j^{-1}(y) = \perp$ then if XtraOuterRnd then $x \leftarrow \text{ReadTape}(G_j^{-1}, y, p_{G_j}^{-1})$ if ($\exists \mathbf{x}$ s.t. $\mathbf{x}[j] = x \wedge$ BlockDefined($G, \mathbf{x}, +$) \wedge BlockDefined($F, \nu^-(\mathbf{x}), -$)) then abort else BackwardOutside(G, j, y) return $G_j^{-1}(y)$ public procedure $G(j, x)$ RequestG(j, x) AdaptLeft() return $G_j(x)$ public procedure $B(j, x)$ if $B_j(x) = \perp$ then $y \leftarrow \text{ReadTape}(B_j, x, p_{B_j})$ if ($\exists \mathbf{y}$ s.t. $\mathbf{y}[j] = y \wedge$ BlockDefined($B, \mathbf{y}, -$) \wedge BlockDefined($C, \tau(\mathbf{y}), +$)) then abort return $B_j(x)$ public procedure $B^{-1}(j, y)$ RequestU$^{-1}(j, y)$ AdaptLeft() return $B_j^{-1}(y)$ public procedure $A^{-1}(j, y)$ RequestU$^{-1}(j, y)$ AdaptLeft() return $A_j^{-1}(y)$ </pre>	<pre> private procedure ForwardOutside(T, j, x) $y \leftarrow \text{ReadTape}(T_j, x, p_{T_j})$ if ($\exists \mathbf{y}$ s.t. $\mathbf{y}[j] = y \wedge$ $P_Z^{-1}(\mathbf{y}) \neq \perp$) then abort // G_3 public procedure $H(j, x)$ if $H(x) = \perp$ then ForwardOutside(H, j, x) return $H_j(x)$ public procedure $H^{-1}(j, y)$ RequestH$^{-1}(j, y)$ AdaptRight() return $H_j^{-1}(y)$ public procedure $C^{-1}(j, y)$ if $C_j^{-1}(y) = \perp$ then $x \leftarrow \text{ReadTape}(C_j^{-1}, y, p_{C_j}^{-1})$ if ($\exists \mathbf{x}$ s.t. $\mathbf{x}[j] = x \wedge$ BlockDefined($C, \mathbf{x}, +$) \wedge BlockDefined($B, \tau^-(\mathbf{x}), -$)) then abort return $C_j^{-1}(y)$ public procedure $C(j, x)$ RequestV(j, x) AdaptRight() return $C_j(x)$ public procedure $A(j, x)$ RequestV(j, x) AdaptRight() return $A_j(x)$ </pre>
---	---

Fig. 4. Games G_1 through G_3 , second set of procedures.

<pre> private procedure RequestG(j, x) if $G_j(x) \neq \perp$ then return if $x \in \text{ToBeAssigned}_{G_j}$ then return $\text{ToBeAssigned}_{G_j} \leftarrow \text{ToBeAssigned}_{G_j} \cup \{x\}$ $\text{NewYUs} \leftarrow \emptyset$ // local var forall $\mathbf{x}^G, \mathbf{y}^H$ s.t. $\mathbf{x}^G[j] = x$ do if $\exists \mathbf{y}^U$ s.t. $(\mathbf{x}^G, \mathbf{y}^U) \in \text{ToBeAdapted}_D$ continue if $(\text{BlockDefined}(G, \mathbf{x}^G, +) = \text{false})$ continue if $(\text{BlockDefined}(H, \mathbf{y}^H, -) = \text{false})$ continue if XtraOuterRnd then $\mathbf{y}^F \leftarrow \nu^-(\mathbf{x}^G)$ if $(\text{BlockDefined}(F, \mathbf{y}^F, -) = \text{false})$ continue $\mathbf{x}^X \leftarrow \text{BlockF}^{-1}(\mathbf{y}^F)$ else $\mathbf{x}^X \leftarrow \mathbf{x}^G$ if $(\text{CheckZ}(\mathbf{x}^X, \mathbf{y}^H) = \text{false})$ continue if $(++\text{NumOuter} \geq q + 1)$ then abort $\mathbf{y}^U \leftarrow \text{RightToMiddle}(\mathbf{y}^H)$ $\text{NewYUs} \leftarrow \text{NewYUs} \cup \mathbf{y}^U$ forall $\mathbf{y}^U \in \text{NewYUs}$ do $\text{BlockRequestU}^{-1}(\mathbf{y}^U)$ private procedure RequestU⁻¹(j, y) if $U_j^{-1}(y) \neq \perp$ then return if $y \in \text{ToBeAssigned}_{U_j}^-$ then return $\text{ToBeAssigned}_{U_j}^- \leftarrow \text{ToBeAssigned}_{U_j}^- \cup \{y\}$ forall \mathbf{y}^U s.t. $\mathbf{y}^U[j] = y$ do if $(\text{BlockDefined}(U, \mathbf{y}^U, -) = \text{false})$ continue if XtraMiddleRnd then if $(\text{BlockDefined}(C, \tau(\mathbf{y}^U), +) = \text{false})$ continue $\mathbf{y}^H \leftarrow \text{MiddleToRight}(\mathbf{y}^U)$ $\mathbf{x}^G \leftarrow \mathbf{x}^X \leftarrow Z^{-1}(\mathbf{y}^H)$ if XtraOuterRnd then $\mathbf{x}^G \leftarrow \nu(\text{BlockF}(\mathbf{x}^X))$ $\text{ToBeAdapted}_D \leftarrow \text{ToBeAdapted}_D \cup (\mathbf{x}^G, \mathbf{y}^U)$ forall $(\mathbf{x}^G, \mathbf{y}^U) \in \text{ToBeAdapted}_D$ do $\text{BlockRequestG}(\mathbf{x}^G)$ private procedure RightToMiddle(\mathbf{y}^H) $\mathbf{x}^H \leftarrow \text{BlockH}^{-1}(\mathbf{y}^H)$ if XtraUntglRnds then $\mathbf{x}^E \leftarrow \text{BlockE}^{-1}(\pi_L^-(\text{BlockL}^{-1}(\pi_H^-(\mathbf{x}^H))))$ $\mathbf{y}^V \leftarrow \pi_C^-(\text{BlockK}^{-1}(\pi_K^-(\mathbf{x}^E)))$ else $\mathbf{y}^V \leftarrow \pi_C^-(\text{BlockE}^{-1}(\pi_H^-(\mathbf{x}^H)))$ if XtraMiddleRnd then return $\tau^-(\text{BlockC}^{-1}(\mathbf{y}^V))$ else return \mathbf{y}^V </pre>	<pre> private procedure RequestH⁻¹(j, y) if $H_j^{-1}(y) \neq \perp$ then return if $y \in \text{ToBeAssigned}_{H_j}^-$ then return $\text{ToBeAssigned}_{H_j}^- \leftarrow \text{ToBeAssigned}_{H_j}^- \cup \{y\}$ $\text{NewXVs} \leftarrow \emptyset$ // local var forall $\mathbf{y}^H, \mathbf{x}^G$ s.t. $\mathbf{y}^H[j] = y$ do if $\exists \mathbf{x}^V$ s.t. $(\mathbf{x}^V, \mathbf{y}^H) \in \text{ToBeAdapted}_E$ continue if $(\text{BlockDefined}(H, \mathbf{y}^H, -) = \text{false})$ continue if $(\text{BlockDefined}(G, \mathbf{x}^G, +) = \text{false})$ continue if XtraOuterRnd then $\mathbf{y}^F \leftarrow \nu^-(\mathbf{x}^G)$ if $(\text{BlockDefined}(F, \mathbf{y}^F, -) = \text{false})$ continue $\mathbf{x}^X \leftarrow \text{BlockF}^{-1}(\mathbf{y}^F)$ else $\mathbf{x}^X \leftarrow \mathbf{x}^G$ if $(\text{CheckZ}(\mathbf{x}^X, \mathbf{y}^H) = \text{false})$ continue if $(++\text{NumOuter} \geq q + 1)$ then abort $\mathbf{x}^V \leftarrow \text{LeftToMiddle}(\mathbf{x}^X)$ $\text{NewXVs} \leftarrow \text{NewXVs} \cup \mathbf{x}^V$ forall $\mathbf{x}^V \in \text{NewXVs}$ do $\text{BlockRequestV}(\mathbf{x}^V)$ private procedure RequestV(j, x) if $V_j(x) \neq \perp$ then return if $x \in \text{ToBeAssigned}_{V_j}$ then return $\text{ToBeAssigned}_{V_j} \leftarrow \text{ToBeAssigned}_{V_j} \cup \{x\}$ forall \mathbf{x}^V s.t. $\mathbf{x}^V[j] = x$ do if $(\text{BlockDefined}(V, \mathbf{x}^V, +) = \text{false})$ continue if XtraMiddleRnd then if $(\text{BlockDefined}(B, \tau^-(\mathbf{x}^V), -) = \text{false})$ continue $\mathbf{x}^X \leftarrow \text{MiddleToLeft}(\mathbf{x}^V)$ $\mathbf{y}^H \leftarrow Z(\mathbf{x}^X)$ $\text{ToBeAdapted}_E \leftarrow \text{ToBeAdapted}_E \cup (\mathbf{x}^V, \mathbf{y}^H)$ forall $(\mathbf{x}^V, \mathbf{y}^H) \in \text{ToBeAdapted}_E$ do $\text{BlockRequestH}^{-1}(\mathbf{y}^H)$ private procedure LeftToMiddle(\mathbf{x}^X) if XtraOuterRnd then $\mathbf{y}^G \leftarrow \text{BlockG}(\nu(\text{BlockF}(\mathbf{x}^X)))$ else $\mathbf{y}^G \leftarrow \text{BlockG}(\mathbf{x}^X)$ if XtraUntglRnds then $\mathbf{y}^D \leftarrow \text{BlockD}(\pi_I(\text{BlockI}(\pi_G(\mathbf{y}^G))))$ $\mathbf{x}^U \leftarrow \pi_B(\text{BlockJ}(\pi_J(\mathbf{y}^D)))$ else $\mathbf{x}^U \leftarrow \pi_C(\text{BlockE}(\pi_G(\mathbf{y}^G)))$ if XtraMiddleRnd then return $\tau(\text{BlockB}(\mathbf{x}^U))$ else return \mathbf{x}^U </pre>
---	--

Fig. 5. Games G_1 through G_3 , third set of procedures.

<pre> private procedure MiddleToRight(y^U) if XtraMiddleRnd then $y^V \leftarrow \text{BlockV}(\tau(y^U))$ else $y^V \leftarrow y^U$ if XtraUntglRnds then $y^E \leftarrow \text{BlockE}(\pi_K(\text{BlockK}(\pi_C(y^V))))$ $x^H \leftarrow \pi_H(\text{BlockL}(\pi_L(y^E)))$ else $x^H \leftarrow \pi_H(\text{BlockE}(\pi_C(y^V)))$ return BlockH(x^H) private procedure AdaptLeft() forall $j \in \{1, \dots, w\}$ do forall $x \in \text{ToBeAssigned}_{G_j}$ do ReadTape(G_j, x, p_{G_j}) forall $y \in \text{ToBeAssigned}_{U_j^-}$ do ReadTape($U_j^{-1}, y, p_{U_j^-}$) $\text{ToBeAssigned}_{G_j} \leftarrow \emptyset$ $\text{ToBeAssigned}_{U_j^-} \leftarrow \emptyset$ forall $(x^G, y^U) \in \text{ToBeAdapted}_D$ if XtraUntglRnds then $x^D \leftarrow \pi_I(\text{BlockI}(\pi_G(\text{BlockG}(x^G))))$ $y^D \leftarrow \pi_J(\text{BlockJ}^{-1}(\pi_B^-(\text{BlockU}^{-1}(y^U))))$ else $x^D \leftarrow \pi_G(\text{BlockG}(x^G))$ $y^D \leftarrow \pi_B^-(\text{BlockU}^{-1}(y^U))$ forall $j \in \{1, \dots, w\}$ do SetTable($D_j, x^D[j], y^D[j]$) $\text{ToBeAdapted}_D \leftarrow \emptyset$ private procedure BlockU⁻¹(y) forall $j \in \{1, \dots, w\}$ do if $U_j^{-1}(y[j]) = \perp$ abort $x[j] \leftarrow U_j^{-1}(y[j])$ return x private procedure BlockG(x) forall $j \in \{1, \dots, w\}$ do if $G_j(x[j]) = \perp$ abort $y[j] \leftarrow G_j(x[j])$ return y private procedure BlockRequestG(x) forall $j \in \{1, \dots, w\}$ do RequestG($x[j], j$) private procedure BlockRequestU⁻¹(x) forall $j \in \{1, \dots, w\}$ do RequestU⁻¹($x[j], j$) </pre>	<pre> private procedure MiddleToLeft(x^V) if XtraMiddleRnd then $x^U \leftarrow \text{BlockU}^{-1}(\tau^-(x^V))$ else $x^U \leftarrow x^V$ if XtraUntglRnds then $x^D \leftarrow \text{BlockD}^{-1}(\pi_J^-(\text{BlockJ}^{-1}(\pi_B^-(x^U))))$ $y^G \leftarrow \pi_G^-(\text{BlockI}^{-1}(\pi_I^-(x^D)))$ else $y^G \leftarrow \pi_G^-(\text{BlockD}(\pi_B^-(x^U)))$ if XtraOuterRnd then return BlockF⁻¹($\nu^-(\text{BlockH}^{-1}(y^G))$) else return BlockH⁻¹(y^G) private procedure AdaptRight() forall $j \in \{1, \dots, w\}$ do forall $y \in \text{ToBeAssigned}_{H_j^-}$ do ReadTape($H_j^{-1}, y, p_{H_j^-}$) forall $x \in \text{ToBeAssigned}_{V_j}$ do ReadTape(V_j, x, p_{V_j}) $\text{ToBeAssigned}_{H_j^-} \leftarrow \emptyset$ $\text{ToBeAssigned}_{V_j} \leftarrow \emptyset$ forall $(x^V, y^H) \in \text{ToBeAdapted}_E$ if XtraUntglRnds then $y^E \leftarrow \pi_L^-(\text{BlockL}^{-1}(\pi_H^-(\text{BlockH}^{-1}(y^H))))$ $x^E \leftarrow \pi_K(\text{BlockK}(\pi_C(\text{BlockV}(x^V))))$ else $y^E \leftarrow \pi_H^-(\text{BlockH}^{-1}(y^H))$ $x^E \leftarrow \pi_C(\text{BlockV}(x^V))$ forall $j \in \{1, \dots, w\}$ do SetTable($E_j, x^E[j], y^E[j]$) $\text{ToBeAdapted}_E \leftarrow \emptyset$ private procedure BlockV(x) forall $j \in \{1, \dots, w\}$ do if $V_j(x[j]) = \perp$ abort $y[j] \leftarrow V_j^{-1}(x[j])$ return y private procedure BlockH⁻¹(y) forall $j \in \{1, \dots, w\}$ do if $H_j^{-1}(y[j]) = \perp$ abort $x[j] \leftarrow H_j^{-1}(y[j])$ return x private procedure BlockRequestH⁻¹(y) forall $j \in \{1, \dots, w\}$ do RequestH⁻¹($x[j], j$) private procedure BlockRequestV(y) forall $j \in \{1, \dots, w\}$ do RequestV($x[j], j$) </pre>
---	---

Fig. 6. Games G_1 through G_3 , fourth set of procedures.

<pre> private procedure BlockB(x) forall j ∈ {1, ..., w} do y[j] ← B(x[j], j) return y private procedure BlockG⁻¹(y) forall j ∈ {1, ..., w} do x[j] ← G⁻¹(y[j], j) return x private procedure BlockD(x) forall j ∈ {1, ..., w} do y[j] ← D(x[j], j) return y private procedure BlockD⁻¹(y) forall j ∈ {1, ..., w} do x[j] ← D⁻¹(y[j], j) return x private procedure BlockF(x) forall j ∈ {1, ..., w} do y[j] ← F(x[j], j) return y private procedure BlockF⁻¹(y) forall j ∈ {1, ..., w} do x[j] ← F⁻¹(y[j], j) return x private procedure BlockI(x) forall j ∈ {1, ..., w} do y[j] ← I(x[j], j) return y private procedure BlockI⁻¹(y) forall j ∈ {1, ..., w} do x[j] ← I⁻¹(y[j], j) return x private procedure BlockL(x) forall j ∈ {1, ..., w} do y[j] ← L(x[j], j) return y private procedure BlockL⁻¹(y) forall j ∈ {1, ..., w} do x[j] ← L⁻¹(y[j], j) return x </pre>	<pre> private procedure BlockC⁻¹(y) forall j ∈ {1, ..., w} do x[j] ← C⁻¹(y[j], j) return x private procedure BlockH(x) forall j ∈ {1, ..., w} do y[j] ← H(x[j], j) return y private procedure BlockE(x) forall j ∈ {1, ..., w} do y[j] ← E(x[j], j) return y private procedure BlockE⁻¹(y) forall j ∈ {1, ..., w} do x[j] ← E⁻¹(y[j], j) return x private procedure BlockK(x) forall j ∈ {1, ..., w} do y[j] ← K(x[j], j) return y private procedure BlockK⁻¹(y) forall j ∈ {1, ..., w} do x[j] ← K⁻¹(y[j], j) return x private procedure BlockL(x) forall j ∈ {1, ..., w} do y[j] ← L(x[j], j) return y private procedure BlockL⁻¹(y) forall j ∈ {1, ..., w} do x[j] ← L⁻¹(y[j], j) return x </pre>
---	---

Fig. 7. Games G_1 , G_2 , G_3 , continuation and end.

<p>Game G_4</p> <p>random tapes: q_{A_1}, \dots, q_{L_w}</p> <p>global static: XtraOuterRnd, XtraMiddleRnd, XtraUntglRnds</p> <pre> private procedure Fwd(T, \mathbf{x}) forall $j \in \{1, \dots, w\}$ do $\mathbf{y}[j] \leftarrow q_{T_j}(\mathbf{x}[j])$ return \mathbf{y} private procedure EvaluateForward(\mathbf{x}^X) if XtraOuterRnd then $\mathbf{y}^G \leftarrow \text{Fwd}(G, \nu(\text{Fwd}(F, \mathbf{x}^X)))$ else $\mathbf{y}^G \leftarrow \text{Fwd}(\mathbf{x}^X)$ if XtraUntglRnds then $\mathbf{y}^D \leftarrow \text{Fwd}(D, \pi_I(\text{Fwd}(I, \pi_G(\mathbf{y}^G))))$ $\mathbf{x}^U \leftarrow \pi_B(\text{Fwd}(I, \pi_I(\mathbf{y}^D)))$ else $\mathbf{x}^U \leftarrow \pi_C(\text{Fwd}(E, \pi_G(\mathbf{y}^G)))$ if XtraMiddleRnd then $\mathbf{y}^V \leftarrow \text{Fwd}(C, \tau(\text{Fwd}(B, \mathbf{x}^U)))$ else $\mathbf{y}^V \leftarrow \text{Fwd}(A, \mathbf{x}^U)$ if XtraUntglRnds $\mathbf{y}^E \leftarrow \text{Fwd}(E, \pi_K(\text{Fwd}(K, \pi_C(\mathbf{y}^V))))$ $\mathbf{x}^H \leftarrow \pi_H(\text{Fwd}(L, \pi_L(\mathbf{y}^E)))$ else $\mathbf{x}^H \leftarrow \pi_H(\text{Fwd}(E, \pi_C(\mathbf{y}^V)))$ return $\text{Fwd}(H, \mathbf{x}^H)$ public procedure Z(\mathbf{x}) if $P_Z(\mathbf{x}) = \perp$ then $\mathbf{y} \leftarrow \text{EvaluateForward}(\mathbf{x})$ SetTable($P_Z, \mathbf{x}, \mathbf{y}$) if $\exists j, H_j^{-1}(\mathbf{y}[j]) \neq \perp$ then abort return $P_Z(\mathbf{x})$ </pre>	<pre> private procedure Bwd(T, \mathbf{y}) forall $j \in \{1, \dots, w\}$ do $\mathbf{x}[j] \leftarrow q_{T_j}^{-1}(\mathbf{y}[j])$ return \mathbf{x} private procedure EvaluateBackward(\mathbf{y}^H) $\mathbf{x}^H \leftarrow \text{Bwd}(H, \mathbf{y}^H)$ if XtraUntglRnds then $\mathbf{x}^E \leftarrow \text{Bwd}(E, \pi_L^{-1}(\text{Bwd}(L, \pi_H^{-1}(\mathbf{x}^H))))$ $\mathbf{y}^V \leftarrow \pi_C^{-1}(\text{Bwd}(K, \pi_K^{-1}(\mathbf{x}^E)))$ else $\mathbf{y}^V \leftarrow \pi_C^{-1}(\text{Bwd}(E, \pi_H^{-1}(\mathbf{x}^H)))$ if XtraMiddleRnd then $\mathbf{x}^U \leftarrow \text{Bwd}(B, \tau^{-1}(\text{Bwd}(C, \mathbf{y}^V)))$ else $\mathbf{x}^U \leftarrow \text{Bwd}(A, \mathbf{y}^V)$ if XtraUntglRnds $\mathbf{x}^D \leftarrow \text{Bwd}(D, \pi_I^{-1}(\text{Bwd}(I, \pi_B^{-1}(\mathbf{x}^U))))$ $\mathbf{y}^G \leftarrow \pi_G^{-1}(\text{Bwd}(I, \pi_I^{-1}(\mathbf{x}^D)))$ else $\mathbf{y}^G \leftarrow \pi_G^{-1}(\text{Bwd}(E, \pi_B^{-1}(\mathbf{x}^U)))$ if XtraOuterRnd then return $\text{Bwd}(F, \nu^{-1}(\text{Bwd}(G, \mathbf{y}^G)))$ else return $\text{Bwd}(G, \mathbf{y}^G)$ public procedure $Z^{-1}(\mathbf{y})$ if $P_Z^{-1}(\mathbf{y}) = \perp$ then $\mathbf{x} \leftarrow \text{EvaluateBackward}(\mathbf{y})$ SetTable($P_Z, \mathbf{x}, \mathbf{y}$) if $\exists j, X_j(\mathbf{x}[j]) \neq \perp$ then abort return $P_Z^{-1}(\mathbf{y})$ </pre>
--	--

Fig. 8. Game G_4 with random tapes and public procedures Z, Z^{-1} . Other public procedures are implemented as in G_3 , substituting G_4 's random tape q_{T_j} for G_3 's random tape p_{T_j} , $T \in \{A, B, \dots, L, L\}$.

<p>Game G_5</p> <p>random tapes: q_{A_1}, \dots, q_{M_w}</p> <p>global static: XtraOuterRnd, XtraMiddleRnd, XtraUntglRnds</p> <p>private procedure Fwd(T, \mathbf{x}) forall $j \in \{1, \dots, w\}$ do $\mathbf{y}[j] \leftarrow q_{T_j}(\mathbf{x}[j])$ return \mathbf{y}</p> <p>private procedure EvaluateForward(\mathbf{x}^X) if XtraOuterRnd then $\mathbf{y}^G \leftarrow \text{Fwd}(G, \nu(\text{Fwd}(F, \mathbf{x}^X)))$ else $\mathbf{y}^G \leftarrow \text{Fwd}(\mathbf{x}^X)$ if XtraUntglRnds then $\mathbf{y}^D \leftarrow \text{Fwd}(D, \pi_I(\text{Fwd}(I, \pi_G(\mathbf{y}^G))))$ $\mathbf{x}^U \leftarrow \pi_B(\text{Fwd}(I, \pi_I(\mathbf{y}^D)))$ else $\mathbf{x}^U \leftarrow \pi_C(\text{Fwd}(E, \pi_G(\mathbf{y}^G)))$ if XtraMiddleRnd then $\mathbf{y}^V \leftarrow \text{Fwd}(C, \tau(\text{Fwd}(B, \mathbf{x}^U)))$ else $\mathbf{y}^V \leftarrow \text{Fwd}(A, \mathbf{x}^U)$ if XtraUntglRnds $\mathbf{y}^E \leftarrow \text{Fwd}(E, \pi_K(\text{Fwd}(K, \pi_C(\mathbf{y}^V))))$ $\mathbf{x}^H \leftarrow \pi_H(\text{Fwd}(L, \pi_L(\mathbf{y}^E)))$ else $\mathbf{x}^H \leftarrow \pi_H(\text{Fwd}(E, \pi_C(\mathbf{y}^V)))$ return $\text{Fwd}(H, \mathbf{x}^H)$</p> <p>public procedure Z(\mathbf{x}) return EvaluateForward(\mathbf{x})</p> <p>public procedure T(x, j) return $q_{T_j}(x)$</p>	<p>private procedure Bwd(T, \mathbf{y}) forall $j \in \{1, \dots, w\}$ do $\mathbf{x}[j] \leftarrow q_{T_j}^-(\mathbf{y}[j])$ return \mathbf{x}</p> <p>private procedure EvaluateBackward(\mathbf{y}^H) $\mathbf{x}^H \leftarrow \text{Bwd}(H, \mathbf{y}^H)$ if XtraUntglRnds then $\mathbf{x}^E \leftarrow \text{Bwd}(E, \pi_L^-(\text{Bwd}(L, \pi_H^-(\mathbf{x}^H))))$ $\mathbf{y}^V \leftarrow \pi_C^-(\text{Bwd}(K, \pi_K^-(\mathbf{x}^E)))$ else $\mathbf{y}^V \leftarrow \pi_C^-(\text{Bwd}(E, \pi_H^-(\mathbf{x}^H)))$ if XtraMiddleRnd then $\mathbf{x}^U \leftarrow \text{Bwd}(B, \tau^-(\text{Bwd}(C, \mathbf{y}^V)))$ else $\mathbf{x}^U \leftarrow \text{Bwd}(A, \mathbf{y}^V)$ if XtraUntglRnds $\mathbf{x}^D \leftarrow \text{Bwd}(D, \pi_I^-(\text{Bwd}(I, \pi_B^-(\mathbf{x}^U))))$ $\mathbf{y}^G \leftarrow \pi_G^-(\text{Bwd}(I, \pi_I^-(\mathbf{x}^D)))$ else $\mathbf{y}^G \leftarrow \pi_G^-(\text{Bwd}(E, \pi_B^-(\mathbf{x}^U)))$ if XtraOuterRnd then return $\text{Bwd}(F, \nu^-(\text{Bwd}(G, \mathbf{y}^G)))$ else return $\text{Bwd}(G, \mathbf{y}^G)$</p> <p>public procedure Z⁻¹(\mathbf{y}) return EvaluateBackward(\mathbf{y})</p> <p>public procedure T⁻¹(y, j) return $q_{T_j}^-(y)$</p>
---	--

Fig. 9. Game G_5 . The procedures EvaluateForward and EvaluateBackward are as in game G_4 . The procedures T and T⁻¹ in G_5 are templates, to be instantiated with each $T \in \{A, B, \dots, K, L\}$.