Lossy Functions Do Not Amplify Well

Krzysztof Pietrzak¹, Alon Rosen², and Gil Segev³

¹ IST Austria

`pietrzak@ist.ac.at`

² Efi Arazi School of Computer Science, IDC Herzliya, Israel

`alon.rosen@idc.ac.il`

³ Microsoft Research, Mountain View, CA 94043, USA

`gil.segev@microsoft.com`

Abstract. We consider the problem of amplifying the “lossiness” of functions. We say that an oracle circuit $C^* : \{0, 1\}^m \rightarrow \{0, 1\}^*$ amplifies relative lossiness from ℓ/n to L/m if for every function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ it holds that

1. If f is injective then so is C^f .
2. If f has image size of at most $2^{n-\ell}$, then C^f has image size at most 2^{m-L} .

The question is whether such C^* exists for $L/m \gg \ell/n$. This problem arises naturally in the context of cryptographic “lossy functions,” where the relative lossiness is the key parameter.

We show that for every circuit C^* that makes at most t queries to f , the relative lossiness of C^f is at most $L/m \leq \ell/n + O(\log t)/n$. In particular, no black-box method making a polynomial $t = \text{poly}(n)$ number of queries can amplify relative lossiness by more than an $O(\log n)/n$ additive term. We show that this is tight by giving a simple construction (cascading with some randomization) that achieves such amplification.

1 Introduction

Lossy trapdoor functions, introduced by Peikert and Waters [14], are a powerful cryptographic primitive. Soon after their introduction, they were found to be useful for realizing new constructions of traditional cryptographic concepts, as well as for demonstrating the feasibility of new ones. Their wide applicability, simple definition, and realizability under a variety of cryptographic assumptions make them a clear candidate for induction into the “pantheon” of cryptographic primitives.

1.1 Lossy Trapdoor Functions

A collection of lossy trapdoor functions consists of two families of functions. Functions in the first family are injective (and can be inverted using a trapdoor), whereas functions in the second are “lossy,” meaning that the size of their image is significantly smaller than the size of their domain. The security requirement

is that the description of a function sampled from the injective family is computationally indistinguishable from the description of a function sampled from the lossy family.

As demonstrated by Peikert and Waters, lossy trapdoor functions imply primitives such as trapdoor functions, collision-resistant hash functions, and oblivious transfer [14]. Amongst “higher level” applications, we can find chosen-ciphertext secure public-key encryption [14], deterministic public-key encryption [4], OAEP-based public-key encryption [10], “hedged” public-key encryption for protecting against bad randomness [2], security against selective opening attacks [3], and non-interactive universally-composable string commitments [13].¹

1.2 Relative Lossiness

A key parameter in all the applications of lossy trapdoor functions is the amount of lossiness guaranteed in case that a lossy function was sampled. We say that a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is (n, ℓ) -lossy if its image size is at most $2^{n-\ell}$. Intuitively, this means that an application of f on an input $x \in \{0, 1\}^n$ loses at least ℓ bits of information, on average, about x . We refer to ℓ as the *absolute* lossiness of the function and to ℓ/n as the *relative* lossiness of the function.

Peikert and Waters [14] showed how to obtain chosen ciphertext secure encryption assuming relative lossiness $\ell/n = \Omega(1)$. This was subsequently improved by Mol and Yilek [12] who, building on work by Rosen and Segev [16], demonstrated how to obtain the same result assuming relative lossiness of only $1/\text{poly}(n)$. One-way functions and similarly trapdoor functions and oblivious transfer, can be constructed assuming relative lossiness of $1/\text{poly}(n)$. Collision resistant hashing requires relative lossiness of at least $1/2 + 1/\text{poly}(n)$. All other known applications of lossy trapdoor functions currently assume relative lossiness that is at least as large as $1 - o(1)$.

Currently, relative lossiness of $1 - o(1)$ seems to be necessary for most “non-traditional” applications of lossy trapdoor functions. While some of the known instantiations are able to guarantee such a high rate of lossiness, some other constructions fall short. Most notably, the lattice-based construction of Peikert and Waters [14], which is the only one based on a worst-case assumption and the only one for which no sub-exponential attack is known, only guarantees relative lossiness of $\Omega(1)$.

High relative lossiness is also relevant for applications that do not necessitate it. This is because the lossiness rate typically has a pronounced effect on the efficiency of the resulting construction. Specifically, higher lossiness rate enables the use of a smaller security parameter, and in many applications also enables the extraction of a larger number of “information theoretic” hard-core bits from the underlying function. This is useful, for example, for efficiently handling long messages.

¹ We note that for some of these constructions (e.g., collision-resistant hashing) the existence of a trapdoor is not required.

1.3 Lossiness Amplification

All of the above leads to the question of whether, given a specific construction of lossy trapdoor functions, it is possible to apply an efficient transformation that would result in a construction with significantly higher lossiness. It can be easily seen that parallel evaluation of t independent copies of an (n, ℓ) -lossy function amplifies the absolute lossiness from ℓ to $t\ell$. Specifically, given an (n, ℓ) -lossy function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ the function $g : \{0, 1\}^{tn} \rightarrow \{0, 1\}^{tn}$, defined as

$$g(x_1, \dots, x_t) = (f(x_1), \dots, f(x_t))$$

is $(tn, t\ell)$ -lossy. However, this comes at the cost of blowing up the input size by a factor of t and hence leaves the relative lossiness ℓ/n unchanged. What we are really looking for is a construction of a (m, L) -lossy function $h : \{0, 1\}^m \rightarrow \{0, 1\}^m$ where $L/m \gg \ell/n$. A natural candidate is sequential evaluation (also known as “cascading”), defined as

$$h(x) = \underbrace{f(f(\dots, f(f(x)) \dots))}_{t \text{ times}}$$

Unfortunately, in general h might not be more lossy than f . In particular, this is the case when f is injective on its own range. One can do a bit better though. By shuffling the outputs in-between every invocation, using randomly chosen $r_1, \dots, r_t$, one obtains the function

$$h_{r_1, \dots, r_t}(x) = f(f(\dots, f(f(x) \oplus r_1) \oplus r_2) \dots \oplus r_t),$$

for which it is possible to show that, if f is say $(n, 1)$ -lossy, then with overwhelming probability over the choice of $r_1, \dots, r_t$, the function $h_{r_1, \dots, r_t}$ has relative lossiness of $\Omega(\log t)/n$.

While already not entirely trivial, relative lossiness of $\Omega(\log t)/n$ is a fairly modest improvement over $\Omega(1)/n$, and would certainly not be considered sufficient for most applications. Still, it is not a-priori inconceivable that there exists more sophisticated ways to manipulate f so that the relative lossiness is amplified in a more significant manner. In this paper, we show that an additive gain of $O(\log n)/n$ is actually the best one can hope for, at least with respect to black-box constructions.

1.4 Our Results

We show that no efficient black-box amplification method can additively improve the relative lossiness of a given function f by more than $O(\log n)/n$. To this end, we consider a circuit $C^* : \{0, 1\}^m \rightarrow \{0, 1\}^*$ with oracle access to a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ such that the following hold:

1. If f is injective then so is C^f .
2. If f has image size of at most $2^{n-\ell}$, then C^f has image size at most 2^{m-L} .

Our main result is that, if $\ell < n - \omega(\log n)$, then for every C^* that makes at most t queries to f , the relative lossiness, L/m , of C^f is at most $(\ell + O(\log t))/n$. The impossibility result holds regardless of whether the injective mode of f has a trapdoor, and rules out even probabilistic constructions C^* (i.e., ones which amplify lossiness only with high probability over the choice of some randomness). In Section 2 we provide a high-level overview of our approach, and in Section 3 we formally present our proof. We then show (in Section 4) how to extend the above result to a “full fledged” cryptographic setting, in which one does not simply get black-box access to a single lossy or injective function f . In this setting, lossy functions are defined by a triple of algorithms $\{g_0, g_1, f\}$, where one requires that a function f_k is injective if the key is sampled by $k \leftarrow g_1$, and lossy if the key is sampled by $k \leftarrow g_0$. Moreover, the distributions generated by the injective and lossy key generation algorithms g_0, g_1 must be computationally indistinguishable.

1.5 Relation to the Collision Problem

Closely related to our setting is the *collision problem*, in which one is given black-box access to a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and is required to distinguish between the case that f is injective and the case that it is 2^ℓ -to-1. A simple argument shows that any (randomized) classical algorithm that tries to distinguish between the cases must make $\Omega(2^{(n-\ell)/2})$ calls to f . Kutin [11], extending work of Aaronson and Shi [1], proves an analogous bound of $\Omega(2^{(n-\ell)/3})$ in the quantum setting.

Lower bounds on the collision problem can be seen to directly imply a weak version of our results. Specifically, if non-trivial lossiness amplification were possible then one could have applied it, and then invoked known upper bounds for the collision problem (either $O(2^{(n-\ell)/2})$ randomized classical or $O(2^{(n-\ell)/3})$ quantum), resulting in a violation of the corresponding lower bounds. However, this approach will only work if the amplification circuit does not blow up f ’s input size (specifically, only if $m < n + (L - \ell)$). In contrast, our results also hold with respect to arbitrary input blow-up.

1.6 Related Work

Several instantiations of lossy trapdoor functions guarantee relative lossiness of $1 - o(1)$. Peikert and Waters present constructions based on the Decisional Diffie-Hellman assumption [14]. These are further simplified by Freeman et al, who also present a generalization based on the d -linear assumption [6]. Boldyreva et al. [4], and independently Freeman et al. [6], present a direct construction based on Paillier’s Composite Residuosity assumption.

Hemenway and Ostrovsky [7] generalize the approach of Peikert and Waters, and obtain relative lossiness of $1 - o(1)$ from any homomorphic hash proof system (a natural variant of hash proof systems [5]). In turn, this implies a unified construction based on either Decisional Diffie Hellman, Quadratic Residuosity, or Paillier’s Composite Residuosity assumptions.

Constructions with relative lossiness $\Omega(1)$ are known based on the hardness of the “learning with errors” problem, which is implied by the worst case hardness of various lattice problems [14]. Kiltz et al. argue that RSA with exponent e satisfies relative lossiness $(\log e)/n$ under the phi-hiding assumption, and that use of multi-prime RSA increases relative lossiness up to $(m \log e)/n$ where m is the number of prime factors of the modulus [10]. Finally, Freeman et al. [6] propose an instantiation based on the Quadratic Residuosity assumption with relative lossiness of $\Omega(1/n)$.

1.7 On Black-Box Separations

The use of black-box separations between cryptographic primitives was pioneered by Impagliazzo and Rudich [9], who proved that there is no black-box construction of a key-exchange protocol from a one-way permutation. Since then, black-box separations have become the standard tool for demonstrating such assertions. We note that our main result is “unconditional”, in the sense that it holds regardless of any cryptographic assumption. Our “cryptographic” result, in contrast, is more standard in that it relies on the indistinguishability property of lossy functions (see the work of Reingold et al. [15] for an extensive discussion on black-box separations).

Strictly speaking, it is not clear whether black-box separations should be interpreted as strong impossibility results. Certainly not as long as non-black-box techniques are still conceivable. Nevertheless, since as far as we know any of the primitives could exist unconditionally (cf. [8]), it is currently not clear how else one could have gone about proving cryptographic lower bounds. In addition, most of the known construction and reductions in cryptography are black-box. Knowing that no such technique can be used to establish an implication serves as a good guideline when searching for a solution. Indeed, it would be extremely interesting to see if non-black box techniques are applicable in the context of lossy function amplification.

2 Overview of Our Approach

We say that a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^{n'}$ is (n, ℓ) -lossy if its image $\{f(x) : x \in \{0, 1\}^n\}$ has size at most $2^{n-\ell}$. We refer to ℓ as the *absolute lossiness*, and ℓ/n as the *relative lossiness* of f . An (n, ℓ) -lossy function f is *balanced* if $f(x)$ has exactly 2^ℓ preimages for every $x \in \{0, 1\}^{n'}$, i.e. $|\{z : f(z) = f(x)\}| = 2^\ell$. We denote with $\mathcal{F}_{n,\ell}$ the set of all balanced (n, ℓ) -lossy functions.

Definition 2.1 (Lossiness amplification). *We say that an oracle circuit $C^* : \{0, 1\}^m \rightarrow \{0, 1\}^{m'}$ amplifies the relative lossiness from ℓ/n to L/m if*

1. *for every injective function f_0 over $\{0, 1\}^n$, C^{f_0} is injective.*
2. *for every $f_1 : \{0, 1\}^n \rightarrow \{0, 1\}^n$ with image size $2^{n-\ell}$, the image of C^{f_1} has size at most 2^{m-L} .*

We say C^* weakly amplifies if C^* is probabilistic and the second item above only holds with probability ≥ 0.9 over the choice of C^* 's randomness.

Remark 2.2 (Permutations vs. injective functions). *In order to make our negative result as strong and general as possible, we require the oracle to be length preserving (and thus the injective f_0 is a permutation), whereas the input and output domain of C^* can be arbitrary.*

For concreteness, in this proof sketch we only consider the case $\ell = 1$. We will also assume that $m = nk$ is an integer multiple of n . The basic idea of our proof is to show that for any C^* , property 1. of Definition 2.1 implies that C^{f_1} has very low collision probability if $f_1 \in \mathcal{F}_{n,1}$ is a randomly chosen 2-1 function. More concretely, let t denote the number of oracle gates in C^* and assume we could prove that

$$\Pr_{X,Y \in \{0,1\}^m} [C^{f_1}(X) = C^{f_1}(Y)] \leq 2^{-k \cdot n + O(k \log t)} \quad (1)$$

Such a low collision probability implies that C^{f_1} must have a large range and thus cannot be too lossy. In particular, Eq. (1) implies that the absolute lossiness of C^{f_1} is at most $O(k \log t)$, or equivalently, the relative lossiness is $O(k \log t)/kn = O(\log t)/n$, which matches (ignoring the constant hidden in the big-oh) the lossiness of the construction $h_{r_1, \dots, r_t}$ from Section 1.3. Unfortunately Eq. (1) is not quite true. For example consider a circuit $\tilde{C}^* : \{0,1\}^{kn} \rightarrow \{0,1\}^{kn}$ which makes only $t = 2$ queries to its oracle and is defined as

$$\tilde{C}^f(x_1, x_2, \dots, x_k) \stackrel{\text{def}}{=} \begin{cases} 0^{kn} & \text{if } f(x_1) = f(x_2) \text{ and } x_1 \neq x_2 \\ (x_1, x_2, \dots, x_k) & \text{otherwise} \end{cases}$$

If $f_0 : \{0,1\}^n \rightarrow \{0,1\}^n$ is a permutation, so is $\tilde{C}^{f_0}$ (in fact, it's the identity function), thus property 1. holds. On the other hand, for any $(n,1)$ -lossy f_1 we have $f_1(x_1) = f_1(x_2)$ and $x_1 \neq x_2$ with probability 2^{-n} for uniform x_1, x_2 . Thus the probability that $\tilde{C}^{f_1}$ outputs 0^{kn} on a random input is also 2^{-n} , which implies

$$\begin{aligned} \Pr_{X,Y \in \{0,1\}^m} [\tilde{C}^{f_1}(X) = \tilde{C}^{f_1}(Y)] &\geq \Pr_{X,Y \in \{0,1\}^m} [\tilde{C}^{f_1}(X) = \tilde{C}^{f_1}(Y) = 0^{2k}] \\ &\geq 2^{-2n} \end{aligned}$$

contradicting Eq. (1) for $k > 2$.

The idea behind the counterexample $\tilde{C}^f$ is to query f on two random inputs and check if f collides on these inputs. If this is the case, C^f “knows” that f is not a permutation and so it must not be a permutation itself as required by property 1, in this case mapping to some fixed output. Although Eq. (1) is wrong, we can prove a slightly weaker statement, where we exclude inputs X where the evaluation of C^f on X involves two invocations of f on inputs $x \neq x'$ where $f(x) = f(x')$ (we will call such bad inputs “burned”). As with high probability, for a random $(n,1)$ -lossy f , most inputs are not burned, already this weaker statement implies that C^f has large range.

The cryptographic setting. In a cryptographic setting, one usually does not simply get black-box access to a single lossy or injective function f , but lossy functions are defined by a collection (indexed by a security parameter λ) of triples of algorithms $\{g_0, g_1, f\}_{\{\lambda \in \mathbb{N}\}}$, where one requires that $f(k, \cdot)$ is injective if the key is sampled by $k \leftarrow g_1$, and lossy if the key is sampled by $k \leftarrow g_0$. Moreover the distributions generated by the injective and lossy key generation algorithms g_0, g_1 must be computationally indistinguishable.

In this setting one can potentially do more sophisticated amplification than what is captured by Definition 2.1, e.g. by somehow using the key-generation algorithms g_0, g_1 . In Section 4 we prove that black-box lossiness amplification is not possible in this setting either.

In a nutshell, we show that constructions which amplify collections of lossy functions can be classified in two classes depending on whether the lossiness of the construction depends only on the lossiness of the oracle (we call such amplifiers “non-communicating”) or if the property of being lossy is somehow encoded into the key. In the first case, the proof goes along the lines of the proof of Theorem 3.1 (in particular, amplifiers as in Definition 2.1 are “non-communicating” as there’s not even a key). In the second case, where the construction is “communicating”, we show that the output of the key-generation algorithms (of the amplified construction) will not always be indistinguishable. This proof borrows ideas from the work of Impagliazzo and Rudich [9] who show that one cannot construct a key-agreement from one-way permutations. Their proof shows that for any two parties Alice and Bob who can communicate over a public channel and who have access to random oracle $\mathcal{R}$, there exists an adversary Eve who can with high probability make all queries to $\mathcal{R}$ that both, Alice and Bob, made. As a consequence, Alice and Bob cannot use $\mathcal{R}$ to “secretly” communicate. In a similar vein we show that the lossy key-generation algorithm cannot “communicate” the fact that the key it outputs is lossy to the evaluation function or we can catch it, and thus distinguish lossy from injective keys.

3 An Upper Bound on Black-Box Lossiness Amplification

We now state our main theorem, asserting that simple sequential composition is basically the best black-box amplification that can be achieved.

Theorem 3.1 (Impossibility of Black-Box Amplification). *Consider any $n, \ell, t \in \mathbb{N}$ where*

$$n \geq \ell + 2 \log t + 2 \tag{2}$$

and any oracle aided circuit $C^ : \{0, 1\}^m \rightarrow \{0, 1\}^{m'}$ which makes t oracle queries per invocation, then the following holds: If C^* weakly amplifies relative lossiness from ℓ/n to L/n ,² then $L \leq \ell + 3 \log t + 4$. More concretely, for a random $f \in \mathcal{F}_{n, \ell}$, the construction C^f will have relative lossiness less than $(\ell + 3 \log t + 4)/n$ with probability at least $1/2$.*

² Note that we denote the relative lossiness of C^* by L/n , not L/m like in the previous sections. In particular, the absolute lossiness of C^* is Lm/n (not L).

Remark 3.2. The bound $n \geq \ell + 2 \log t + 2$ is basically tight, as for $n = \ell + 2 \log t - O(1)$ one can with constant advantage p distinguish any (n, ℓ) -lossy function from an injective one by simply making t random queries and looking for a collision. The exact value of p depends on the $O(1)$ term, in particular, replacing the $O(1)$ with a sufficiently large constant we get a $p \geq .9$ as required by Definition 2.1. Then $C^f(x)$ which outputs x if no such collision is found, and some fixed value (say $0^{m'}$) otherwise is a weak amplifier as in Definition 2.1.

Remark 3.3 (Probabilistic C^* vs. random f). Instead of considering a probabilistic C^* and constructing a particular lossy f such that C^f is not too lossy with high probability over C^* 's randomness (as required by Definition 2.1), we consider a deterministic C^* and show that C^f fails to be lossy with high probability for a randomly chosen f . As f is sampled independently of (the description of) C^* , the latter implies the former.

Below we formally define what we mean by an input being burned as already outlined.

Definition 3.4 (Burned input). For $X \in \{0, 1\}^m$, we denote with $\text{in}(X)$ and $\text{out}(X)$ the inputs and outputs of the t invocations of f in an evaluation of $C^f(X)$. Consider an input $X \in \{0, 1\}^m$ and let $\{x_1, \dots, x_t\} \leftarrow \text{in}(X)$, we say that X is burned if for some $1 \leq i < j \leq t$, $x_i \neq x_j$ and $f(x_i) = f(x_j)$. $\phi(X)$ denotes the event that X is burned.

Below is the main technical Lemma which we will use to prove Theorem 3.1 (recall that $m = nk$).

Lemma 3.5. For a random balanced (n, ℓ) -lossy function f , and two random inputs X, Y , the probability that X, Y are colliding inputs for C^f and at the same time both are not burned can be upper bounded as

$$\Pr_{\substack{f \in \mathcal{F}_{n, \ell} \\ X, Y \in \{0, 1\}^m}} [(C^f(X) = C^f(Y)) \wedge \neg \phi(X) \wedge \neg \phi(Y)] \leq 2^{-kn+k(3 \log t + \ell)} \quad (3)$$

We postpone the proof of this Lemma to Section 3.1. The following simple claim upper bounds the probability (over the choice of $f \in \mathcal{F}_{n, \ell}$) that an input x to C^f is burned

Claim 3.6. For any $x \in \{0, 1\}^m$

$$\Pr_{f \in \mathcal{F}_{n, \ell}} [\phi(x)] \leq \frac{2^\ell t^2}{2^n} \quad (4)$$

Proof. For $i \in \{1, \dots, t\}$, the probability that the i th query to f made during the evaluation of $C^f(x)$ provides a collision for f (assuming there's been no collision so far) is at most $\frac{(i-1)(2^\ell-1)}{2^n-i-1}$. To see this, note that as f is balanced, there are exactly $(i-1)(2^\ell-1)$ possible inputs which will lead to a collision as each of the $(i-1)$ queries we did so far has $2^\ell-1$ other preimages. As f is random, the probability the i th query (for which there are 2^n-i-1 choices)

will hit one of these values is $\frac{(i-1)(2^\ell-1)}{2^n-i-1}$. The claim follows by taking the union bound over all i

$$\Pr_{f \in \mathcal{F}_{n,\ell}} [\phi(x)] \leq \sum_{i=1}^t \frac{(i-1)(2^\ell-1)}{2^n-i-1} \leq \frac{2^\ell t^2}{2^n}$$

The second step above used $t \leq 2^{n/2}$ which is implied by Eq. (2). $\blacksquare$

Proof of Theorem 3.1. Consider a C^* as in the statement of the theorem and a random $f \in \mathcal{F}_{n,\ell}$. Let $\Phi \stackrel{\text{def}}{=} \{x \in \{0,1\}^m : \phi(x)\}$ denote the set of inputs which are burned (cf. Definition 3.4) and $\overline{\Phi} = \{0,1\}^m \setminus \Phi$. Using the chain rule, we can state eq.(3) as

$$\Pr_{\substack{f \in \mathcal{F}_{n,\ell} \\ X, Y \in \overline{\Phi}}} [C^f(X) = C^f(Y)] \leq \frac{2^{-kn+k(3 \log t + \ell)}}{\Pr_{\substack{f \in \mathcal{F}_{n,\ell} \\ X, Y \in \{0,1\}^m}} [X, Y \in \overline{\Phi}]} \quad (5)$$

Using eq.(4) we can bound the expected size (over the choice of $f \in \mathcal{F}_{n,\ell}$) of Φ as

$$\mathbb{E}[|\Phi|] = |\{0,1\}^n| \cdot \Pr_{\substack{f \in \mathcal{F}_{n,\ell} \\ X \in \{0,1\}^m}} [\phi(X)] \leq 2^n \cdot \frac{2^\ell t^2}{2^n} = 2^\ell t^2$$

Using the Markov inequality and eq.(2), this implies that Φ is not too big, say at most half of the domain $\{0,1\}^m$, with probability $1/2$

$$\begin{aligned} \Pr_{f \in \mathcal{F}_{n,\ell}} [|\Phi| \geq 2^{n-1}] &= \Pr_{f \in \mathcal{F}_{n,\ell}} [|\Phi| \geq 2^{n-1-\ell-2 \log t} \mathbb{E}[|\Phi|]] \\ &\leq 1/2^{n-1-\ell-2 \log t} \\ &\stackrel{(2)}{\leq} 1/2 \end{aligned}$$

By the above equation, $|\overline{\Phi}| > 2^{n-1}$ with probability $\geq 1/2$ over the choice of f , and for such a “good” f , two random X, Y are in $\overline{\Phi}$ with probability at least $(1/2)^2 = 1/4$. Thus the denominator on the right side of eq.(5) is at least $1/8$, replacing the denominator in eq.(5) with $2^{-3} = 1/8$ we get

$$\Pr_{\substack{f \in \mathcal{F}_{n,\ell} \\ X, Y \in \{0,1\}^m}} [C^f(X) = C^f(Y)] \leq 2^{-kn+k(3 \log t + \ell)+3} \quad (6)$$

Again using Markov, this means that for a randomly chosen $f \in \mathcal{F}_{n,\ell}$, with probability at least $1/2$

$$\Pr_{X, Y \in \{0,1\}^m} [C^f(X) = C^f(Y)] \leq 2^{-kn+k(3 \log t + \ell)+4} \quad (7)$$

As two values sampled independently from a distribution with support of size u collide with probability at least $1/u$ (this is tight if the distribution is flat), eq.(7) implies that the range of C^f must be at least of size $2^{kn-k(3 \log t + \ell)-4}$, thus the relative lossiness (recall that $m = nk$) is $(k\ell + k3 \log t + 4)/kn \leq (\ell + 3 \log t + 4)/n$. $\blacksquare$

3.1 Proof of Lemma 3.5

We consider a random experiment denoted Γ where $X_0, Y_0 \in \{0, 1\}^m$ and $f \in \mathcal{F}_{n, \ell}$ are chosen at random, and then $C^f(X_0)$ and $C^f(Y_0)$ are evaluated. This evaluations result in $2t$ invocations of f . Let $\{x_1, \dots, x_t\} \leftarrow \text{in}(X_0)$ and $\{X_1, \dots, X_t\} \leftarrow \text{out}(X_0)$ denote the inputs and outputs of f in the evaluation of $C^f(X_0)$. Analogously we define values y_i, Y_i occurring in the evaluation of $C^f(Y_0)$. For $I \subseteq \{1, \dots, t\}$, we define an event E_I which holds if for every $i \in I$ (and only for such i), there exists a j such that $y_i \neq x_j$ and $f(y_i) = f(x_j)$ and $y_i \neq y_k$ for all $k < i$ (i.e. we have a fresh, non-trivial collision). Γ defines a “transcript”

$$v_{f, X_0, Y_0} \stackrel{\text{def}}{=} \{X_0, Y_0, x_1, \dots, x_t, f(x_1), \dots, f(x_t), y_1, \dots, y_t, f(y_1), \dots, f(y_t)\}$$

The values x_i and y_i in the transcript are redundant, i.e., they can be computed from values $X_0, Y_0, f(x_i)$ and $f(y_i)$, and only are added for convenience. For $I \subseteq \{1, \dots, t\}$ we define V_I as all transcripts where (1) both inputs are not burned (2) we have a collision and (3) E_I holds, i.e.

$$V_I \stackrel{\text{def}}{=} \{v_{f, X_0, Y_0} : \neg\phi(X_0) \wedge \neg\phi(Y_0) \wedge (C^f(X_0) = C^f(Y_0)) \wedge E_I\}$$

V_{col} is the union of all V_I , i.e.

$$V_{\text{col}} = \cup_I V_I = \{v_{f, X_0, Y_0} : \neg\phi(X_0) \wedge \neg\phi(Y_0) \wedge C^f(X_0) = C^f(Y_0)\} \quad (8)$$

For a set of transcripts V , we denote with $\Pr_\Gamma[V]$ the probability that the transcript generated by Γ is in V . It is not hard to see³ that $\Pr_\Gamma[V_\emptyset] \leq 2^{-nk}$, we prove that this bound (up to a factor 2) holds for any V_I .

Lemma 3.7. *For any $I \subseteq \{1, \dots, t\}$ we have (recall that $m = nk$)*

$$\Pr_\Gamma[V_I] \leq 2^{-nk+1}$$

We postpone the proof of this main technical lemma and first prove how it implies Theorem 3.1. But let us here give some intuition as to why Lemma 3.7 holds. The experiment Γ generates a transcript in V_I if (besides $C^f(X_0) = C^f(Y_0)$ colliding and X_0, Y_0 not being burnt) for every $i \in I$, the i th invocation of f during the evaluation of $C^f(Y_0)$ produces a fresh collision. Now, conditioned on such a collision happening, the probability of actually getting a collision $C^f(X_0) = C^f(Y_0)$ can potentially raise significantly (by something like $2^{n-\ell}$) as this is a rare event, but then, the probability of having such a collision is also around $2^{n-\ell}$, and if this collision does not occur, we definitely will not end up

³ We have $\Pr_\Gamma[V_\emptyset] \leq \Pr_\Gamma[X_0 = Y_0] = 2^{-nk}$. The second step follows as $X_0, Y_0 \in \{0, 1\}^{nk}$ are uniformly random. The first step follows as $\neg\phi(X_0), \neg\phi(Y_0)$ and $E_\emptyset$ together imply that there are no collisions in the $2t$ invocations of f , and thus f is “consistent” with being a permutation. But in this case, $C^f(X_0) = C^f(Y_0)$ implies $X_0 = Y_0$.

with a transcript in V_I . These two probabilities even out, and we end up with roughly the same probability for a transcript V_I as we had for $V_\emptyset$.

Before we can prove the theorem we need one more lemma, which bounds the probability of Γ generating a transcript with lots (k or more) collisions.

Lemma 3.8

$$\sum_{I:|I|\geq k} \Pr_F[V_I] \leq \sum_{I:|I|\geq k} \Pr_F[E_I] \leq 2^{k(\ell+2\log t-(n-1))} \quad (9)$$

Proof. The first step of Eq. (9) follows as V_I implies E_I . Let E_I^+ denote the event which holds if $E_{I'}$ holds for any $I' \supseteq I$. We have

$$\Pr_F[E_I^+] \leq \left(\frac{(2^\ell - 1)t}{2^n - 2t} \right)^{|I|} \leq \left(\frac{2^\ell t}{2^{n-1}} \right)^{|I|} \quad (10)$$

To see this, note that to get E_I^+ , in every step $i \in I$, x_i must be fresh, and then $f(y_i)$ must “hit” one of the at most t distinct $f(x_i)$. As f is a random 2^ℓ -1 function evaluated on at most $2t$ inputs, this probability can be upper bounded by $(2^\ell - 1)t/(2^n - 2t)$ as at most $(2^\ell - 1)t$ of the at least $2^n - 2t$ fresh inputs can “hit” as described above. The probability that we have such a “hit” for all $i \in I$ is the $|I|$ ’th power of this probability. The number of different I where $|I| = k$ can be upper bounded by $2^{k \log t}$, using this and Eq. (10) we get

$$\begin{aligned} \sum_{I:|I|\geq k} \Pr_F[E_I] &\leq \sum_{I:|I|=k} \Pr_F[E_I^+] \\ &\leq 2^{k \log t} \left(\frac{2^{\ell k} t^k}{2^{(n-1)k}} \right) \\ &= 2^{k(\ell+2\log t-(n-1))} \end{aligned}$$

■

Proof of Lemma 3.5. Lemma 3.5 states that $\Pr_F[V_{col}] \leq 2^{-kn+O(k \log(t))}$, which we can write as

$$\Pr_F[V_{col}] \stackrel{Eq.(8)}{=} \sum_{I:|I|<k} \Pr_F[V_I] + \sum_{I:|I|\geq k} \Pr_F[V_I]$$

Using Lemma 3.7 and 3.8 and the fact that there are $\binom{t}{k-1} < t^k$ different I ’s with $|I| < k$, we get

$$\begin{aligned} \sum_{I:|I|<k} \Pr_F[V_I] + \sum_{I:|I|\geq k} \Pr_F[V_I] &\leq t^k \cdot 2^{-nk+1} + 2^{k(\ell+2\log t-(n-1))} \\ &\leq 2^{1+k(\ell+2\log t-(n-1))} \\ &< 2^{-nk+k(3\log t+\ell)} \end{aligned}$$

■

Proof of Lemma 3.7. For any I , we consider a new random experiment Γ_I . This experiment will define a distribution $X'_t \in \{0, 1\}^m, Y'_t \in \{0, 1\}^m \cup \perp$. We'll show that

$$\Pr_{\Gamma_I}[X'_t = Y'_t] \leq 2^{-m} \quad (11)$$

and

$$\Pr_{\Gamma_I}[V_I] \leq 2 \cdot \Pr_{\Gamma_I}[X'_t = Y'_t] \quad (12)$$

Note that the two equations above imply Lemma 3.7. The experiment Γ_I is defined as follows

1. We sample random $X'_0, Y'_0 \in \{0, 1\}^m$ and a random *permutation* g over $\{0, 1\}^n$.
2. Let $x'_1, \dots, x'_t$ be the inputs to g in the evaluation of $C^g(X'_0)$. Let $X'_t \stackrel{\text{def}}{=} C^g(X'_0)$.
3. Now evaluate $C^g(Y'_0)$ in steps (one invocation of g per step), where for any $i \in I$ do the following:
 - if y'_i is “fresh” (that is $y'_i \neq x'_j$ for any $1 \leq j \leq t$ and $y'_i \neq y'_j$ for any $1 \leq j < i$). we change the value of $g(y'_i)$ and set it to some uniformly random value $z_i \in_U \{0, 1\}^n$ (note that g is no longer a permutation).
 - If y'_i is no fresh set $Y'_t = \perp$ and stop.

We will first prove Eq. (11). Let's consider a new random experiment Γ_I^* which will define outputs $X''_t, Y''_t \in \{0, 1\}^m$. This experiment is defined exactly as the experiment Γ_I defining X'_t, Y'_t , but when y'_i is not fresh we nonetheless redefine $g(y'_i)$ to a random z_i (instead of setting $Y'_t = \perp$ and aborting). As the two experiments only differ when $Y'_t = \perp$, but X'_t cannot be $\perp$, we have.

$$\Pr_{\Gamma_I}[X'_t = Y'_t] \leq \Pr_{\Gamma_I^*}[X''_t = Y''_t]$$

Moreover $X''_t = C^g(X''_0)$ is uniformly random (as X''_0 is uniform and C^g is a permutation) and Y''_t is *independent* of X''_t (the reason we consider the experiment Γ_I^* is because in Γ_I we don't have this independence), thus

$$\Pr_{\Gamma_I^*}[Y''_t = X''_t] = 2^{-m}$$

The two equations above imply Eq. (11). Now we show Eq. (12), i.e.

$$\Pr_{\Gamma_I}[V_I] \leq 2 \cdot \Pr_{\Gamma_I}[X'_t = Y'_t] \quad (13)$$

We will show a stronger statement, namely that for every transcript $\hat{v} \in V_I$ we have

$$\Pr_{\Gamma_I}[\hat{v}] \leq 2 \cdot \Pr_{\Gamma_I}[X'_t = Y'_t] \quad (14)$$

This implies (13) as

$$\Pr_{\Gamma_I}[V_I] = \sum_{\hat{v} \in V_I} \Pr_{\Gamma_I}[\hat{v}] \leq 2 \cdot \sum_{\hat{v} \in V_I} \Pr_{\Gamma_I}[\hat{v}] = 2 \cdot \Pr_{\Gamma_I}[V_I] \leq 2 \cdot \Pr_{\Gamma_I}[X'_t = Y'_t]$$

We'll use the following notation for the transcript $\hat{v}$ and the transcripts generated by Γ and Γ_I , respectively.

$$\hat{v} \stackrel{\text{def}}{=} \{\hat{X}_0, \hat{Y}_0, \hat{x}_1, \dots, \hat{x}_t, a_1, \dots, a_t, \hat{y}_1, \dots, \hat{y}_t, b_1, \dots, b_t\}$$

$$v \stackrel{\text{def}}{=} \{X_0, Y_0, x_1, \dots, x_t, f(x_1), \dots, f(x_t), y_1, \dots, y_t, f(y_1), \dots, f(y_t)\}$$

$$v' \stackrel{\text{def}}{=} \{X'_0, Y'_0, x'_1, \dots, x'_t, g(x'_1), \dots, g(x'_t), y'_1, \dots, y'_t, g(y'_1), \dots, g(y'_t)\}$$

As $\hat{X}_0, \hat{Y}_0, X_0, Y_0, X'_0, Y'_0$ are uniformly random, we have

$$\Pr[(\hat{X}_0, \hat{Y}_0) = (X_0, Y_0)] = \Pr[(\hat{X}_0, \hat{Y}_0) = (X'_0, Y'_0)] = 2^{-2m}$$

Further

$$\begin{aligned} \Pr_{\Gamma}[(\hat{x}_1, \dots, \hat{x}_t, a_1, \dots, a_t) = (x_1, \dots, x_t, f(x_1), \dots, f(x_t)) \mid (\hat{X}_0, \hat{Y}_0) = (X_0, Y_0)] &\leq \\ \Pr_{\Gamma_I}[(\hat{x}_1, \dots, \hat{x}_t, a_1, \dots, a_t) = (x'_1, \dots, x'_t, g(x'_1), \dots, g(x'_t)) \mid (\hat{X}_0, \hat{Y}_0) = (X'_0, Y'_0)] & \end{aligned}$$

Using the chain rule, the above is implied by

$$\prod_{i=1}^t \Pr_{\Gamma}[a_i = f(x_i) \mid \dots] \leq \prod_{i=1}^t \Pr_{\Gamma_I}[a_i = g(x'_i) \mid \dots] \quad (15)$$

where here and below we use the convention that “...” always means that the transcript defined up to this point is consistent with the transcript $\hat{v}$. E.g. on the left side of eq.(15) the “...” stands for

$$(\hat{X}_0, \hat{Y}_0) = (X_0, Y_0) \quad , \quad \forall j = 1 \dots i-1 : f(x_j) = a_j \quad (16)$$

Note that we don't have to explicitly require $\forall j = 1 \dots i-1 : x_j = \hat{x}_j$ as this is already implied by (16).⁴

For $i = 1, \dots, 2t$ we will denote with $q_i \leq i$ the number of *distinct* elements that appeared as inputs to f in the first i queries. I.e., for $i \leq t$ $q_i = |\{\hat{x}_1, \dots, \hat{x}_i\}|$ and for $t < i \leq 2t$, $q_i = |\{\hat{x}_1, \dots, \hat{x}_t, \hat{y}_1, \dots, \hat{y}_{i-t}\}|$.

To see that Eq. (15) holds, note that for any i where $\hat{x}_i$ is not fresh (i.e. $x_i = x_j$ for some $j < i$) we have

$$\Pr[a_i = f(x_i) \mid \dots] = \Pr[a_i = g(x'_i) \mid \dots] = 1$$

For i 's where x_i is fresh, let q_i denote the number of distinct elements in $\hat{x}_1, \dots, \hat{x}_{i-1}$. A g is a random permutation and $a_i \neq g(x'_j)$ for $j < i$ because $\neg \phi(X_0)$, we have

$$\Pr_{\Gamma_I}[g(x'_i) = a_i \mid \dots] = \frac{1}{2^n - q_i}$$

⁴ If the inputs $(\hat{X}_0, \hat{Y}_0) = (X_0, Y_0)$ are identical, and all the oracle queries so far gave the same outputs, also all intermediate values (including the next oracle query) will be the same.

On the other hand

$$\Pr_{\Gamma}[f(x_i) = a_i | \dots] \leq \frac{1}{2^n - q_i}$$

To see this note that $\Pr_{\Gamma}[f(x_i) = a_i | \dots]$ is exactly $\frac{1}{2^n - q_i}$ if one additionally conditions on the fact that $f(x_i) \neq a_j$ for all $j < i$. Not conditioning on this event can only decrease the probability as $a_i \neq a_j$ for $j < i$ as $\neg\phi(X_0)$.

Now we come to the second part of the transcript. Here we will show that

$$\begin{aligned} & \Pr_{\Gamma}[(\hat{y}_1, \dots, \hat{y}_t, a_1, \dots, a_t) = (y_1, \dots, y_t, f(y_1), \dots, f(y_t)) | \dots] \leq \\ & 2 \cdot \Pr_{\Gamma}[(\hat{y}_1, \dots, \hat{y}_t, a_1, \dots, a_t) = (y'_1, \dots, y'_t, g(y'_1), \dots, g(y'_t)) | \dots] \end{aligned}$$

The proof is almost identical as for the first part, except that now for fresh y'_i we have a slightly smaller probability

$$\Pr_{\Gamma}[g(y'_i) = b_i | \dots] = 2^{-n}$$

that g maps to the right value b_i in the experiment Γ_I , as by definition of Γ_I the output of g is assigned a uniformly random value in this case. Using the fact that $t \leq 2^{-n/4}$ this difference is covered by the extra factor 2. ■

4 Extension to Collections of Lossy Functions

By Theorem 3.1 no circuit C^* (of size polynomial in n) can amplify relative lossiness better than sequential composition. That is, if C^f is injective for any permutation $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$, then there exists an (n, ℓ) -lossy f (i.e. it has relative lossiness ℓ/n) such that C^f has relative lossiness only $(\ell + O(\log n))/n$. In fact, a random (n, ℓ) -lossy f will have this property with very high probability. In a cryptographic setting, lossy functions are not given as a single function, but by a collection of triple of algorithms as defined below.

Definition 4.1 (Collection of Lossy Functions). Let $\lambda \in \mathbb{N}$ denote a security parameter and $n = n(\lambda)$, $n' = n'(\lambda)$, $\ell = \ell(\lambda)$ be functions of λ . A collection of (n, n', ℓ) -lossy function is a sequence (indexed by λ) of functions $\pi = \{g_0, g_1, f\}_{\lambda \in \mathbb{N}}$ where g_0, g_1 are probabilistic key-generation functions, such that

1. **Evaluation of lossy functions:** For every function index $\sigma \leftarrow g_0(1^\lambda)$, $f(\sigma, \cdot)$ is a function $f_\sigma : \{0, 1\}^n \rightarrow \{0, 1\}^{n'}$ whose image is of size at most $2^{n-\ell}$.
2. **Evaluation of injective functions:** For every function index $\sigma \leftarrow g_1(1^\lambda)$, the function $f(\sigma, \cdot)$ computes an injective function $f_\sigma : \{0, 1\}^n \rightarrow \{0, 1\}^{n'}$.
3. **Security:** The ensembles $\{\sigma : \sigma \leftarrow g_0(1^\lambda)\}_{\lambda \in \mathbb{N}}$ and $\{\sigma : \sigma \leftarrow g_1(1^\lambda)\}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable.

We refer to ℓ as the absolute lossiness of π , and to ℓ/n as the relative lossiness of π .

Definition 4.2 (Black-Box Amplification of Lossy Collection). *A triple of probabilistic polynomial-time oracle algorithms $\Pi^* = \{G_0^*, G_1^*, F^*\}$ is a black-box amplification for relative lossiness from $\alpha = \alpha(\lambda)$ to $\beta = \beta(\lambda)$ ($\beta > \alpha$) if for every oracle $\pi = \{g_0, g_1, f\}_{\lambda \in \mathbb{N}}$ that implements a $(n, n, \alpha n)$ -lossy collection, Π^π is a $(m, m', \beta m)$ -lossy collection (where $m = m(\lambda)$, $m' = m'(\lambda)$).*

Note that if π is efficient (i.e. can be implemented by polynomial time algorithms), so is Π^π . We will prove the following theorem.

Theorem 4.3 (Impossibility of Black-Box Amplification). *Let t, ℓ, n be functions of λ such that $n(\lambda) \leq \ell(\lambda) + 2 \log(t(\lambda)) + \omega(\lambda)$. If each of the algorithms in $\Pi^* = \{G_0^*, G_1^*, F^*\}$ makes at most $t = t(\lambda)$ oracle queries per invocation and Π^* amplifies relative lossiness from $\alpha(\lambda) = \ell/n$ to $\beta(\lambda) = L/n$ then $L = \ell + O(\log t)$.*

To save on notation, we will identify the security parameter λ with the domain size n of the lossy-function we try to amplify (which will be given as an oracle).

To prove Theorem 4.3, we will show that for any construction Π^* , if we choose a random (n, ℓ) -lossy $\pi_n = \{g_0, g_1, f\}$ (“random” to be defined in Section 4.1), then with overwhelming probability either the outputs of $G_0^{\pi_n}$ and $G_1^{\pi_n}$ can be distinguished relative to π_n , or for a random lossy key $k \leftarrow G_0^{\pi_n}$, the function $F^{\pi_n}(k, \cdot)$ has very small collision probability and thus cannot be too lossy.

4.1 The Random $\pi = \{g_0, g_1, f\}$

For $n, \ell \in \mathbb{N}$ let $\mathcal{L}_{n, \ell}$ denote the set of triples of functions $g_0, g_1 : \{0, 1\}^{n-1} \rightarrow \{0, 1\}^n$, $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ where the range of g_0, g_1 covers all of $\{0, 1\}^n$ (note this means that the range of g_0 and g_1 are disjoint) and (with $\mathcal{F}_{n, \ell}$ as defined in the first paragraph of Section 2)

$$\forall x \in \{0, 1\}^{n-1} : f(g_0(x), \cdot) \in \mathcal{F}_{n, \ell} \quad \text{and} \quad f(g_1(x), \cdot) \in \mathcal{F}_{n, 0}$$

Claim 4.4. *For $\ell(n) \leq n - \omega(n)$, let $\pi = \{\pi_n\}_{n \in \mathbb{N}}$ where $\pi_n = \{g_0, g_1, f\}$ is chosen uniformly in $\mathcal{L}_{n, \ell}$ (for every $n \in \mathbb{N}$.) Then with overwhelming probability π is (n, ℓ) -lossy even relative to an EXPTIME-complete oracle.*

4.2 (Non-)Communicating Π^*

Consider a $\Pi^* = \{G_0^*, G_1^*, F^*\}$ as in Definition 4.2. We will classify such Π^* in two classes, depending on whether Π^* is close to being “non-communicating” or not. Intuitively, we say Π^* is non-communicating if the lossiness of Π^π comes entirely from the lossiness of π , that is, if π is not lossy, then also Π^π will not be lossy.

Definition 4.5 ((close to) non-communicating). *Π^* is non-communicating if for every $n \in \mathbb{N}$ and $\pi_n \in \mathcal{L}_{n, 0}$ the function computed by $F^\pi(k, \cdot)$ is injective for every $k \leftarrow G_0^\pi(1^n)$. In addition, Π^* is close to being non-communicating if*

for all but finitely many $n \in \mathbb{N}$, with probability $1/2$ over the choice of a random $\pi_n \in \mathcal{L}_{n,0}$, for at least $1/2$ of the keys $k \leftarrow \mathcal{G}_0^{\pi_n}$, there's a subset $\mathcal{M}_k \subseteq \{0,1\}^m$ of size at least $2^m/2$ such $F^{\pi_n}(k, x)$ is injective on $\mathcal{M}_k$ (i.e. for $x, x' \in \mathcal{M}_k$, $F^{\pi_n}(k, x) = F^{\pi_n}(k, x')$ implies $x = x'$).

In order to prove that Theorem 4.3 holds for some particular construction Π^* , we will use a different argument depending on whether Π^* is close to being non-communicating or not. The proof for the first case is almost identical to that of Theorem 3.1, where we rely on the fact that C^* is injective for any injective key k . The proof for the second case relies on the indistinguishability of injective and lossy functions, and requires new ideas. More specifically, in this case we prove the following lemma:

Lemma 4.6. *If Π^* (as in the statement of Theorem 4.3) is far (i.e. not close) from being non-communicating, then for infinitely many $n \in \mathbb{N}$ the following holds. For a random $\pi_n \in \mathcal{L}_{n,\ell}$ the outputs of $\mathcal{G}_0^{\pi_n}$ and $\mathcal{G}_1^{\pi_n}$ can be distinguished with constant advantage making $\text{poly}(t, n)$ oracle queries to π_n (and one query to an EXPTIME oracle).*

Due to space limitations in the remainder of this section we describe a high-level outline for the proof of Lemma 4.6, and refer the reader to the full version for the formal proof.

Proof outline. For $b \in \{0, 1\}$ consider a key $k \leftarrow \mathcal{G}_b^{\pi_n}(R)$ and let $\mathcal{Q}_k$ denote all the queries that $\mathcal{G}_b^{\pi_n}(R)$ made to its oracle π_n during sampling this key using randomness R . Now consider a $(n, 0)$ -lossy $\hat{\pi}_n \in \mathcal{F}_{n,0}$ which is sampled at random except that we require it to be consistent with the queries in $\mathcal{Q}_k$. As $\hat{\pi}_n$ is consistent with π_n on $\mathcal{Q}_k$, we have $\mathcal{G}_b^{\pi_n}(R) = \mathcal{G}_b^{\hat{\pi}_n}(R) = k$. Thus if

- $b = 1$, then k is a valid injective key relative to $\hat{\pi}_n$ and thus $F^{\hat{\pi}_n}(k, \cdot)$ has image size 2^m .
- $b = 0$, then k is a valid lossy key relative to $\hat{\pi}_n$. As Π^* is far from being non-communicating, with constant probability $F^{\hat{\pi}_n}(k, \cdot)$ will have an image size of $\leq 2^{m-1}$ despite the fact that $\hat{\pi}_n$ is not lossy at all.

Using the above two observations, here's a way to distinguish the case $b = 0$ from $b = 1$ (i.e. lossy from injective keys) with constant advantage given $\mathcal{Q}_k$ and access to an EXPTIME oracle: query the oracle on input k , $\mathcal{Q}_k$ and ask for the image size of $F^{\hat{\pi}_n}(k, \cdot)$ for a $\hat{\pi}_n$ randomly sampled as described above. If the image size is $\leq 2^{m-1}$, guess $b = 0$, guess $b = 1$ otherwise.

Unfortunately we are only given the key k , but not $\mathcal{Q}_k$. What we'll do is consider a random $\bar{\pi}_n$ which is consistent with π_n on a set of inputs/outputs $\mathcal{Q}_{k,q}^{\text{sam}}$ to π_n which is sampled by invoking $F^{\pi_n}(k, \cdot)$ on $q = \text{poly}(n, t)$ random inputs (i.e. $\mathcal{Q}_{k,q}^{\text{sam}}$ contains all inputs/outputs to π_n made during these q invocations).

We will prove that for such a $\bar{\pi}_n$ the image size of $F(k, \cdot)^{\bar{\pi}_n}$ is still close to 2^m if $b = 1$, but with constant probability $\ll 2^m$ if $b = 0$, so we can use our EXPTIME oracle to distinguish these cases by sending $k, \mathcal{Q}_{k,q}^{\text{sam}}$ (which, unlike

$\mathcal{Q}_k$, we do have) to the EXPTIME oracle asking for the image size of $F(k, \cdot)^{\bar{\pi}_n}$ when $\bar{\pi}_n \in \mathcal{F}_{n,0}$ is chosen at random but consistent with $\mathcal{Q}_{k,q}^{sam}$.

The reason it is good enough to consider a $\bar{\pi}_n$ that is consistent with $\mathcal{Q}_{k,q}^{sam}$ and not $\mathcal{Q}_k$, is that for sufficiently many samples $q = \text{poly}(n, t)$, $\mathcal{Q}_{k,q}^{sam}$ will with high probability contain all “heavy” queries in $\mathcal{Q}_k$, where we say a query is heavy if there’s a good probability that $F^{\pi_n}(k, \cdot)$ will make that query if invoked on a random input.

So for most inputs x , $F^{\bar{\pi}_n}(k, x)$ will not query $\bar{\pi}_n$ on a query which is in $\mathcal{Q}_k$ (i.e. which was made during key-generation), but is not in $\mathcal{Q}_{k,q}^{sam}$. As a consequence, $F^{\bar{\pi}_n}(k, \cdot)$ “behaves” differently from what we would get by using $\hat{\pi}_n$ (which is consistent with all of $\mathcal{Q}_k$) instead of $\bar{\pi}_n$ only for a small fraction of the inputs. In particular, the image size is close to what we would have gotten by using $\hat{\pi}_n$.

Acknowledgements. We would like to thank Oded Goldreich and Omer Reingold for discussions at an early stage of this project, and Scott Aaronson for clarifications regarding the collision problem.

References

1. Aaronson, S., Shi, Y.: Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM* 51(4), 595–605 (2004)
2. Bellare, M., Brakerski, Z., Naor, M., Ristenpart, T., Segev, G., Shacham, H., Yilek, S.: Hedged Public-Key Encryption: How to Protect against Bad Randomness. In: Matsui, M. (ed.) *ASIACRYPT 2009*. LNCS, vol. 5912, pp. 232–249. Springer, Heidelberg (2009)
3. Bellare, M., Hofheinz, D., Yilek, S.: Possibility and Impossibility Results for Encryption and Commitment Secure under Selective Opening. In: Joux, A. (ed.) *EUROCRYPT 2009*. LNCS, vol. 5479, pp. 1–35. Springer, Heidelberg (2009)
4. Boldyreva, A., Fehr, S., O’Neill, A.: On Notions of Security for Deterministic Encryption, and Efficient Constructions without Random Oracles. In: Wagner, D. (ed.) *CRYPTO 2008*. LNCS, vol. 5157, pp. 335–359. Springer, Heidelberg (2008)
5. Cramer, R., Shoup, V.: Universal Hash Proofs and a Paradigm for Adaptive Chosen Ciphertext Secure Public-Key Encryption. In: Knudsen, L.R. (ed.) *EUROCRYPT 2002*. LNCS, vol. 2332, pp. 45–64. Springer, Heidelberg (2002)
6. Freeman, D.M., Goldreich, O., Kiltz, E., Rosen, A., Segev, G.: More Constructions of Lossy and Correlation-Secure Trapdoor Functions. In: Nguyen, P.Q., Pointcheval, D. (eds.) *PKC 2010*. LNCS, vol. 6056, pp. 279–295. Springer, Heidelberg (2010)
7. Hemenway, B., Ostrovsky, R.: Lossy trapdoor functions from smooth homomorphic hash proof systems. *Electronic Colloquium on Computational Complexity*, Report TR09-127 (2009)
8. Impagliazzo, R.: A personal view of average-case complexity. In: *Structure in Complexity Theory Conference*, pp. 134–147 (1995)
9. Impagliazzo, R., Rudich, S.: Limits on the provable consequences of one-way permutations. In: *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pp. 44–61 (1989)

10. Kiltz, E., O'Neill, A., Smith, A.: Instantiability of RSA-OAEP under Chosen-Plaintext Attack. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 295–313. Springer, Heidelberg (2010)
11. Kutin, S.: Quantum lower bound for the collision problem with small range. *Theory of Computing* 1(1), 29–36 (2005)
12. Mol, P., Yilek, S.: Chosen-Ciphertext Security from Slightly Lossy Trapdoor Functions. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 296–311. Springer, Heidelberg (2010)
13. Nishimaki, R., Fujisaki, E., Tanaka, K.: Efficient Non-interactive Universally Composable String-Commitment Schemes. In: Pieprzyk, J., Zhang, F. (eds.) ProvSec 2009. LNCS, vol. 5848, pp. 3–18. Springer, Heidelberg (2009)
14. Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. In: Proceedings of the 40th Annual ACM Symposium on Theory of Computing, pp. 187–196 (2008)
15. Reingold, O., Trevisan, L., Vadhan, S.P.: Notions of Reducibility between Cryptographic Primitives. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 1–20. Springer, Heidelberg (2004)
16. Rosen, A., Segev, G.: Chosen-ciphertext security via correlated products. *SIAM Journal on Computing* 39(7), 3058–3088 (2010)