

On the Complexity of Collision Resistant Hash Functions: New and Old Black-Box Separations

Nir Bitansky* Akshay Degwekar[†]

June 10, 2019

Abstract

The complexity of collision-resistant hash functions has been long studied in the theory of cryptography. While we often think about them as a Minicrypt primitive, black-box separations demonstrate that constructions from one-way functions are unlikely. Indeed, theoretical constructions of collision-resistant hash functions are based on rather structured assumptions.

We make two contributions to this study:

1. *A New Separation:* We show that collision-resistant hashing does not imply hard problems in the class Statistical Zero Knowledge in a black-box way.
2. *New Proofs:* We show new proofs for the results of Simon, ruling out black-box reductions of collision-resistant hashing to one-way permutations, and of Asharov and Segev, ruling out black-box reductions to indistinguishability obfuscation. The new proofs are quite different from the previous ones and are based on simple *coupling arguments*.

*Tel Aviv University. Email nirbitan@tau.ac.il. Member of the Check Point Institute of Information Security. Supported in part by Alon Young Faculty Fellowship, by Len Blavatnik and the Blavatnik Family foundation, and an ISF grant 18/484.

[†]MIT. Email: akshayd@mit.edu. This work was done in part while visiting the FACT Center in IDC Herzliya, supported in part by ISF grant 1861/16 and AFOSR Award FA9550-17-1-0069. Research supported in part by NSF Grants CNS-1413920 and CNS-1350619, and by the Defense Advanced Research Projects Agency (DARPA) and the U.S. Army Research Office under contracts W911NF-15-C-0226 and W911NF-15-C-0236.

Contents

1	Introduction	1
1.1	This Work	2
1.2	More Related Work on Black-Box Separations	2
2	Techniques	3
2.1	Collision Resistance When SZK is Easy	3
2.2	Proving Simon & Asharov-Segev : A Coupling-Based Approach	5
3	Preliminaries	7
3.1	Conventions	7
3.2	Coupling and Statistical Distance.	7
4	Separating SZK and CRHFs	8
4.1	Fully Black-Box Constructions of SZK Problems	8
4.2	The Statistical Distance Oracle	9
4.3	Insensitivity to Local Changes	9
4.4	Collision Resistance in the Presence of SDO Oracle.	11
5	New Proofs of Old Separations	14
5.1	Oracle-Aided Indistinguishability Obfuscation	14
5.2	Simon’s Collision Finding Oracle and Puncturing	17
5.3	Smoothing for the Collision Finding Oracle	19
5.4	Warm up: One Way Permutations in the presence of Coll	19
5.5	Indistinguishability Obfuscation (and OWPs) in the Presence of Coll	22
	References	30
A	Proofs from [BDV17]	34

1 Introduction

Collision-resistant hash functions (CRHFs) are perhaps one of the most studied and widely used cryptographic primitives. Their applications range from basic ones like “hash-and-sign” [Dam87, Mer89] and statistically hiding commitments [DPP93, HM96] to more advanced ones like verifiable delegation of data and computation [Ki92, BEG⁺94] and hardness results in complexity theory [MP91, KNY17].

Constructions. Collision resistance is trivially satisfied by random oracles and in common practice, to achieve it, we heuristically rely on unstructured hash functions like SHA. Accordingly, we often think of CRHFs as a creature of *Minicrypt*, the realm of symmetric key cryptography [Imp95]. However, when considering theoretical constructions with formal reductions, collision resistance is only known based on problems with some algebraic structure, like Factoring, Discrete Log, and different short vector and bounded distance decoding problems (in lattices or in binary codes) [Dam87, GGH96, PR06, LM06, AHI⁺17, YZW⁺17, BLVW19]. Generic constructions are known from claw-free permutations [Dam87, Rus95], homomorphic primitives [OK91, IKO05], and private information retrieval [IKO05], which likewise are only known from similar structured assumptions. An exception is a recent work by Holmgren and Lombardi [HL18] which constructs CRHFs from a new assumption called *one-way product functions*. These are functions where efficient adversaries succeed in inverting two random images with probability at most $2^{-n-\omega(\log n)}$. Indeed, this assumption does not explicitly require any sort algebraic structure.

Understanding the Complexity of CRHFs. In light of the above, it is natural to study what are the minimal assumptions under which CRHFs can be constructed, and whether they require any sort of special structure. Here Simon [Sim98] provided an explanation for our failure to base CRHFs on basic Minicrypt primitives like one-way functions or one-way permutations. He showed that there are no black-box reductions of CRHFs to these primitives. In fact, Asharov and Segev [AS15] demonstrated that the difficulty in constructing CRHFs from general assumptions runs far deeper. They showed that CRHFs cannot be black-box reduced even to *indistinguishability obfuscation* (and one-way permutations), and accordingly not to anyone of the many primitives it implies, like public key encryption, oblivious transfer, or functional encryption.

CRHFs and SZK. An aspect common to many CRHF constructions is that they rely on assumptions that imply hardness in the class SZK. Introduced by Goldwasser, Micali and Rackoff [GMR85], SZK is the class of *promise problems* with statistical zero-knowledge proofs. Indeed, SZK hardness is known to follow from various algebraic problems that lead to CRHFs, such as Discrete Logarithms [GK93], Quadratic Residuosity [GMR85], and Lattice Problems [GG98, MV03], as well as from generic primitives that lead to CRHFs such as homomorphic encryption [BL13], lossy functions [PVW08], and computational private information retrieval [LV16].

The formal relation between SZK and CRHFs is still not well understood. As possible evidence that SZK hardness may be sufficient to obtain collision resistance, Komargodski and Yogev [KY18] show that average-case hardness in SZK implies a relaxation of CRHFs known as *distributional* CRHFs. Applebaum and Raykov [AR16] show that CRHFs are implied by average-case hardness in a subclass of SZK of problems that have a *perfect randomized encoding*. Berman et al. [BDRV18] showed that average-case hardness of a variant of entropy approximation, a complete problem for the class of Non-Interactive SZK (NISZK), suffices to construct yet a different relaxation known as *multi-collision resistance*.

Is hardness in SZK necessary for CRHFs? Our perception of CRHFs as a Minicrypt primitive, as well as the result by Holmgren and Lombardi mentioned above, suggest that this should not be the case. However, we do not know how to prove this. Meaningfully formalizing a statement of the form “CRHFs do not require SZK hardness” requires care — it is commonly believed that SZK *does* contain hard problems, and if this is the case then formally, CRHFs (or any other assumption for that matter) imply hardness in SZK. To capture this statement we again resort to the methodology of black-box separations; that is, we aim to prove that hard problems in SZK cannot be obtained from CRHFs in a black-box way.

Recent work by Bitansky, Degwekar, and Vaikuntanathan [BDV17] showed that a host of primitives, essentially, all primitives known to follow from IO, do not lead to hard problems in SZK through black-box reductions. Their separation, however, does not imply a separation from CRHFs; indeed, CRHFs are not known to follow from IO, and in fact according to Asharov and Segev [AS15], cannot in a black-box way.

1.1 This Work

In this work, we close the above gap, proving that CRHFs do not imply hardness in SZK through black-box reductions.

Theorem. *There are no fully black-box reductions of any (even worst-case) hard problem in SZK to CRHFs.*

Here by *fully* black box we mean reductions where both the construction and the security proof are black box (in the CRHF and the attacker, respectively). This is the common type of reductions used in cryptography. We refer the reader to the technical overview in Section 2 for more details.

New proofs of Simon and Asharov and Segev. Our second contribution is new proofs for the results of Simon [Sim98], ruling out black-box reductions of CRHFs to OWPs, and of Asharov and Segev [AS15], ruling out black-box reductions of CRHFs to OWPs and IO. The new proofs draw from ideas used in [BDV17]. They are based mostly on simple *coupling arguments* and are quite different from the original proofs.

1.2 More Related Work on Black-Box Separations

Following the seminal work of Impagliazzo and Rudich [IR89], black-box separations in cryptography have been thoroughly studied (see, e.g., [Rud88, KST99, GKM⁺00, GT00, GMR01, BT03, RTV04, HR04, GGKT05, Pas06, GMM07, BM09, HH09, BKS11, DLMM11, KSS11, GKLM12, DHT12, Fis12, BBF13, Pas13, BB15, GHMM18]). Most of this study has been devoted to establishing separations between different cryptographic primitives and some of it to putting limitations on basing cryptographic primitives on NP-hardness [GG98, AGGM06, MX10, BL13, BB15, LV16].

Perhaps most relevant to our works are the works of Simon [Sim98], Asharov and Segev [AS15] and [BDV17] mentioned above, as well as the work by Haitner et al. [HHR15] who gave an alternative proof for the Simon result (and extended it to the case of statistically-hiding commitments of low round complexity).

We also note that [KNY18] claim to show that distributional CRHFs cannot be reduced to multi-collision resistant hash functions in a black box way, which given the black-box construction of distributional CRHFs from SZK hardness [KY18], would imply that SZK hardness cannot be obtained from multi-collision resistance in a black box way. However, for the time being there seems to be a gap in the proof of this claim [Per].

2 Techniques

We now give an overview of the techniques behind our results.

Ruling Out Black-Box Reductions. Most constructions in cryptography are fully black-box [RTV04], in the sense that both the construction and (security) reduction are black box. In a bit more detail, a fully black-box construction of a primitive \mathcal{P}' from another primitive \mathcal{P} consists of two algorithms: a construction C and a reduction R . The construction $C^{\mathcal{P}}$ implements \mathcal{P}' for any valid oracle \mathcal{P} . The reduction $R^{A, \mathcal{P}}$, given oracle-access to any adversary A that breaks $C^{\mathcal{P}}$, breaks the underlying \mathcal{P} . Hence, breaking the instantiation $C^{\mathcal{P}}$ of \mathcal{P}' is at least as hard as breaking \mathcal{P} itself.

A common methodology to rule out fully black black-box constructions of a primitive \mathcal{P}' from primitive \mathcal{P} (see e.g., [Sim98, HR04, HHR15]), is to demonstrate oracles (Γ, A) such that:

- relative to Γ , there exists a construction C^{Γ} realizing \mathcal{P} that is secure in the presence of A ,
- but *any* construction C^{Γ} realizing \mathcal{P}' can be broken in the presence of A .

Indeed, if such oracles (Γ, A) exist, then no efficient reduction will be able to use (as a black-box) the attacker A against \mathcal{P}' to break \mathcal{P} (as the construction of \mathcal{P} is secure in the presence of A).

We now move on to explain how each of our results is shown in this framework.

2.1 Collision Resistance When SZK is Easy

Our starting point is the work by [BDV17] who showed oracles relative to which Indistinguishability Obfuscation (IO) and One-Way Permutations (OWPs) exist and yet SZK is easy. We next recall their approach and explain why it falls short of separating CRHFs from SZK. We then explain the approach that we take in order to bridge this gap.

Black-box Constructions of SZK Problems. The [BDV17] modeling of problems in SZK follows the characterization of SZK by Sahai and Vadhan [SV03] through its complete Statistical Difference Problem (SDP). SDP is a promise problem, where given circuit samplers (C_0, C_1) , the task is to determine if the statistical distance between their respective output distributions is large ($> 2/3$) or small ($< 1/3$). Accordingly, we can model a black-box construction of a statistical distance problem SDP^{Ψ} , relative to an oracle Ψ , defined by

$$\begin{aligned} \text{SDP}_Y^{\Psi} &= \left\{ (C_0, C_1) : \text{SD}(C_0^{\Psi}, C_1^{\Psi}) \geq \frac{2}{3} \right\} , \\ \text{SDP}_N^{\Psi} &= \left\{ (C_0, C_1) : \text{SD}(C_0^{\Psi}, C_1^{\Psi}) \leq \frac{1}{3} \right\} . \end{aligned}$$

Jumping ahead, our eventual goal will be construct an oracle $\Gamma = (\Psi, A)$ such that SDP^{Ψ} is easy in the presence of A , and yet Ψ can be used to securely realize a CRHF, in the presence of A . Here we naturally choose Ψ to be a random shrinking function f , and for the SZK breaker A adopt the oracle SDO^f from [BDV17]. SDO^f is a randomized oracle that takes as input a pair of oracle-aided circuits $(C_0^{(\cdot)}, C_1^{(\cdot)})$, computes the statistical distance $s = \text{SD}(C_0^f, C_1^f)$, samples a random value $t \leftarrow (1/3, 2/3)$, and outputs:

$$\text{SDO}^f(C_0, C_1; t) := \begin{cases} N & \text{If } s < t \\ Y & \text{If } s \geq t \end{cases} .$$

This oracle is clearly sufficient to break (or rather, decide) SDP^f . The challenge is in showing that CRHFs exist in the presence of the oracle SDO^f , which may make exponentially many queries to f when computing the statistical distance.

One-Way Permutations in the Presence of SDO. Toward proving the existence of CRHFs in the presence of SDO, we first recall the argument from [BDV17] as to why one-way permutations exist relative to SDO, and then explain why it falls short of establishing the existence of CRHFs.

Consider the oracle $\Gamma = (f, \text{SDO}^f)$, where f is a random permutation. Showing that $f(x)$ is hard to invert for an adversary $A^{f, \text{SDO}^f}(f(x))$ with access to f and SDO^f relies on two key observations:

1. Inverting f requires detecting random *local changes*. Indeed, imagine an alternative experiment where we replace f with a slightly perturbed function $f_{x' \rightarrow f(x)}$, which diverts a random x' to $f(x)$. In this experiment, the attacker would not be able to distinguish x from x' and would output them with the exact same probability. Note, however, that if the attacker can invert f in the real experiment (namely, output x) with noticeable probability, then this means that the probabilities of outputting x and x' in the original experiment must noticeably differ. Indeed, in the original experiment x' is independent of the attacker's view. It is not hard to show that without access to the oracle SDO^f , such perturbations cannot be detected (this can be shown for example via a coupling argument, as we explain in more detail in Section 2.2).
2. The SDO^f oracle itself, and thus A^{f, SDO^f} , can be made oblivious to random, local changes. Hence, even given access to the SDO^f oracle, the adversary cannot invert with non-trivial probability. This is shown based on the idea of “smoothening”: any two circuits (C_0^f, C_1^f) can be transformed into new circuits that do not make any specific query x with high probability. This allows arguing that even if we perturb f at a given point, their statistical distance s does not change by much. In particular, if s is moderately far from the random threshold t , chosen by SDO, s' the statistical distance of the perturbed circuits remains on the same side of t , which means that SDO's answer will remain invariant. Indeed, such “farness” holds with overwhelming probability over SDO's choice of t .

What About Collision Resistance? The above approach is not sufficient to argue that collisions are hard to find (when f is replaced with a shrinking function). The reason is that collisions are “non-local” — they are abundant, and it is impossible to eliminate all of them in a shrinking function. In fact, as we shall show later on, a similar argument to the one above can be made to work relative to an oracle that trivially breaks CRHFs (this leads to our new proofs of the separations of CRHFs from OWPs and IO [Sim98, AS15]). Accordingly, a different approach is required.

Our Approach: Understanding What Statistical Difference Oracles Reveal. At high level, to show that collisions in f are hard to find, we would like to argue that queries to SDO^f leak no information about any $f(x)$, except for inputs x , which the adversary had already explicitly revealed by querying f itself. This would essentially reduce the argument to the standard argument showing that random oracles are collision resistant — each new query collides with any previous query with probability at most 2^{-m} , where m is f 's output length. Overall, an attacker making q queries cannot find a collision except with negligible probability $q^2 2^{-m}$.

However, showing that SDO^f reveals nothing is too good to be true. Rather, we show that this is the case with overwhelming probability. That is, with overwhelming probability on any partial execution, the value $f(x)$ of any x not explicitly queried within the execution is uniformly random.

Roughly speaking, the property that such partial executions should satisfy is that all queries to SDO^f satisfy smoothness and fairness conditions similar to those discussed above. The essential observation is that when such conditions hold the answer of SDO^f remains invariant not only to a random local change, but to *any* local change. In particular, a partial execution transcript satisfying these conditions would remain invariant if we change the value $f(x)$ for any x not explicitly queried to any particular $y \neq f(x)$.

A Note on Leakage from Random Oracles. Our approach is in part inspired by the works of Unruh [Unr07] and Coretti et al. [CDGS18] on *random oracles with auxiliary information*. They show that revealing short auxiliary information about f (so called leakage), essentially has the effect of fixing f on a small set of values, while the rest of f remains hidden. This does not suffice for us, because it does not restrict in any way which values are fixed. We need to ensure that *all* values not explicitly queried remain hidden even under the leakage from the oracle SDO . (Our argument is restricted though to the specific oracle SDO and does not say anything about arbitrary leakage.)

2.2 Proving Simon & Asharov-Segev : A Coupling-Based Approach

Next, we sketch the main ideas underlying the new proofs of Simon’s result that OWPs do not imply CRHFs through fully black-box constructions, and the extended result by Asharov, which consider not only OWPs, but also IO. In this overview, we focus on the simpler result by Simon. We refer the reader to Section 5.4 for a complete proof of this result and to Section 5.5 for the extension to IO.

Simon’s Collision Finding Oracle. The oracle $\Gamma = (f, \text{Coll}^f)$ introduced by Simon consists of a random permutation f and a collision finding oracle Coll^f . The oracle Coll^f given a circuit C^f returns a random w along with a random element that collides with w ; namely a random w' in the preimage of $y = C^f(w)$. In particular, if the circuit C is compressing, then the oracle will output a collision $w \neq w'$ with high probability, meaning that CRHFs cannot exist in its presence.

Our Proof. To prove that Coll does not help inverting f , Simon used careful conditional probability arguments, whereas Haitner et al. [HHS15], and then Asharov and Segev [AS15] adding also IO to the picture, relied on a *compression and reconstruction argument*, originally due to Gennaro and Trevisan [GT00] (see Section 5.5.2 for a comparison). Our proof is inspired by the [BDV17] proof that the statistical distance oracle SDO does not help inverting permutations (discussed above). At high level, we would like to argue that the collision-finding oracle Coll , like the oracle SDO , is oblivious to random local changes. Following the intuition outlined for SDO , an attacker that fails to detect random local changes will also fail in inverting random permutations.

Punctured Collision Finders. To fulfil this plan, we consider a *punctured* version PColl of the oracle Coll , where the function f can be erased at a given set of values S . Roughly speaking, PColl will allow us to argue that Coll is not particularly sensitive to the value $f(x)$ of almost any x . To define PColl , we first give a more concrete description of Coll and then explain how we change it.

The oracle Coll , for any circuit $C : \{0, 1\}^k \rightarrow \{0, 1\}^*$, assigns a random input $w \in \{0, 1\}^k$ and a random permutation π of $\{0, 1\}^k \simeq [2^k]$. It then returns (w, w') , where w' is the first among $\pi(1), \pi(2), \dots$ such that $C^f(w) = C^f(w')$. The oracle PColl_S^f is parameterized by a set of punctured inputs $S \subseteq \{0, 1\}^n$. Like Coll , for any C , it samples a random input w and a permutation π . Differently from Coll , if $C^f(w)$ queries any $x \in S$, the oracle returns \perp . Else, it iterates over the inputs $\{0, 1\}^k$ according to π and finds the first value w' such that (1) $C^f(w')$ makes no queries to

any $x \in S$, and (2) $C^f(w) = C^f(w')$. The oracle outputs the collision (w, w') .

The PColl oracle satisfies the following essential property. Let τ be a transcript generated by the attacker A^{f, Coll^f} and assume that for all Coll answers (w, w') in τ , neither $C^f(w)$ nor $C^f(w')$ query any $x \in S$. Then A^{f, PColl_S^f} generates the exact same transcript τ . Indeed, this follows directly from the definition of the punctured oracle PColl.

Proving Hardness of Inversion by Smoothing and Coupling. Equipped with the punctured oracle, we now explain how it can be used to argue the hardness of inversion. We first consider a smoothing process analogous to the one considered in the statistical distance separation discussed above. That is, we make sure that (with overwhelming probability) all queries C made to Coll are smooth in the sense that $C^f(w)$ does not query any specific input with high probability when w is chosen at random. We then make a few small perturbations to our oracles, and argue that they are undetectable by a coupling argument. Finally, we deduce univertability.

Step 1: Let x be the preimage that $A^{f, \text{Coll}^f}(f(x))$ aims to find. We first consider, instead of Coll, the punctured oracle $\text{PColl}_{\{x\}}^f$. Due to smoothness, almost every transcript produced by $A^{f, \text{Coll}^f}(f(x))$ is such that x is not queried by $C^f(w), C^f(w')$ for any query C and answer (w, w') returned by Coll. Any transcript satisfying the latter can be coupled with an identical transcript generated by $A^{f, \text{PColl}_{\{x\}}^f}(f(x))$, and deduce that the probability of inversion (outputting x) in this new experiment E_1 is close to the probability in the original experiment E_0 .

Step 2: We perturb the oracle again. We sample a random $x' \leftarrow \{0, 1\}^n$ and make the following two changes: (1) we change the oracle f to $f_{x' \rightarrow f(x)}$, which diverts x' to $f(x)$, and (2) we puncture at x' , namely, we consider $\text{PColl}_{\{x, x'\}}^f$.

We next observe that in this new experiment E_2 , x and x' are symmetric. Accordingly, x and x' are output with the same probability in the experiment E_2 . To complete the proof, we apply a coupling argument to show that x and x' are output with *almost* the same probability also in the previous experiment E_1 . This is enough as in E_1 the view of the attacker is independent of x' , which will allow us to deduce that the probability of inversion is negligible overall.

Let us describe the coupling argument more explicitly. Both experiments E_1 and E_2 are determined by the choice of f, x, x' and randomness $R = \{w, \pi\}$ for Coll. We can look at the events $X_1 = X_1(f, x, x', R)$ and $X_2 = X_2(f, x, x', R)$, where X_1 occurs when the attacker outputs x in the experiment E_1 and X_2 occurs when it outputs x in E_2 . Similarly, we can look at X'_1 and X'_2 , which describe the events that x' is output in each of the experiments. Then by coupling, we know that

$$\left| \Pr[X_1] - \Pr[X_2] \right| \leq \Pr_{f, x, x', R} [I_{X_1} \neq I_{X_2}] ,$$

where I_{X_1}, I_{X_2} are the corresponding indicators. The same holds for X'_1, X'_2 . Thus, we can bound:

$$\begin{aligned} \left| \Pr[X_1] - \Pr[X'_1] \right| &\leq \left| \Pr[X_1] - \Pr[X_2] \right| + \left| \Pr[X_2] - \Pr[X'_2] \right| + \left| \Pr[X'_1] - \Pr[X'_2] \right| \\ &\leq \Pr_{f, x, x', R} [I_{X_1} \neq I_{X_2}] + 0 + \Pr_{f, x, x', R} [I_{X'_1} \neq I_{X'_2}] . \end{aligned}$$

It is left to see that when fixing f, x, R the outputs in the two experiments E_1, E_2 (and thus also X_1, X_2 and X'_1, X'_2) are identical as long as x' does not coincide with any of the queries to f , nor with any of the queries induced by any $\text{PColl}_{\{x\}}$ answer (w, w') . Since the number of such queries is bounded and x' is chosen independently at random, this will almost surely be the case.

3 Preliminaries

In this section, we introduce the basic definitions and notation used throughout the paper.

3.1 Conventions

For a distribution D , we denote the process of sampling from D by $x \leftarrow D$. A function $\text{negl} : \mathbb{N} \rightarrow \mathbb{R}^+$ is negligible if for every constant c , there exists a constant n_c such that for all $n > n_c$ $\text{negl}(n) < n^{-c}$.

Randomized Algorithms. As usual, for a random algorithm A , we denote by $A(x)$ the corresponding output distribution. When we want to be explicit about the algorithm using randomness r , we shall denote the corresponding output by $A(x; r)$. We refer to uniform probabilistic polynomial-time algorithms as PPT algorithms.

Oracles. We consider *oracle-aided algorithms (or circuits)* that make repeated calls to an oracle Γ . Throughout, we will consider deterministic oracles Γ that are a-priori sampled from a distribution Γ on oracles. More generally, we consider infinite oracle ensembles $\Gamma = \{\Gamma_n\}_{n \in \mathbb{N}}$, one distribution Γ_n for each security parameter $n \in \mathbb{N}$ (each defined over a finite support). For example, we may consider an ensemble $f = \{f_n\}$ where each $f_n : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a random function. For such an ensemble Γ and an oracle aided algorithm (or circuit) A with finite running time, we will often abuse notation and denote by $A^\Gamma(x)$ and execution of A on input x where each of (finite number of) oracle calls that A makes is associated with a security parameter n and is answered by the corresponding oracle Γ_n . When we write $A_1^\Gamma, \dots, A_k^\Gamma$ for k algorithms, we mean that they all access the same realization of Γ .

3.2 Coupling and Statistical Distance.

Definition 3.1 (Coupling). *Given two random variables X, Y over \mathcal{X}, \mathcal{Y} , a coupling of X, Y is defined to be any distribution $P_{X,Y}$ on $\mathcal{X} \times \mathcal{Y}$ such that, the marginals of $P_{X,Y}$ on \mathcal{X} and \mathcal{Y} are the distributions X, Y respectively.*

Denote by \mathcal{P}_{XY} the set of all couplings of X, Y .

Lemma 3.2. *Given any two distributions X, Y supported on \mathcal{X} ,*

$$\text{SD}(X, Y) = \inf_{P_{X,Y} \in \mathcal{P}_{XY}} \Pr_{(x,y) \leftarrow P_{X,Y}} [x \neq y].$$

Furthermore, for distributions over a discrete domain \mathcal{X} the infimum is attained: that is, there exists a coupling P_{XY} such that $\text{SD}(X, Y) = \Pr_{(x,y) \leftarrow P_{XY}} [x \neq y]$.

The lemma allows us to bound the statistical distance between two random variables (hybrid experiments in our case) by setting up a coupling between two experiments and bounding the probability of them giving a different outcome. Looking ahead, in Lemma 5.9, we describe an explicit coupling for the Simon's collision finder oracle, of the form above that allows us to bound the statistical distance between hybrids.

4 Separating SZK and CRHFs

4.1 Fully Black-Box Constructions of SZK Problems

The class of problems with Statistical Zero Knowledge Proofs (SZK) [GMR85, Vad99] can be characterized by complete promise problems [SV03], particularly statistical difference, and the transformation is black-box. In order to consider black-box constructions of hard problems in SZK, we start by defining statistical difference problem relative to oracles. This modelling follows [BDV17].

In the following definition, for an oracle-aided (sampler) circuit $C^{(\cdot)}$ with n -bit input and an oracle Ψ , we denote by \mathbf{C}^Ψ the output distribution $C^\Psi(r)$ where $r \leftarrow \{0, 1\}^n$. We denote statistical distance by SD: for two distributions X and Y $SD(X, Y) = \frac{1}{2} \sum_x |\Pr[X = x] - \Pr[Y = x]|$.

Definition 4.1 (Statistical Difference Problem relative to oracles). *For an oracle Ψ , the statistical difference promise problem relative to Ψ , denoted as $SDP^\Psi = (SDP_Y^\Psi, SDP_N^\Psi)$, is given by*

$$SDP_Y^\Psi = \left\{ (C_0, C_1) : SD(\mathbf{C}_0^\Psi, \mathbf{C}_1^\Psi) \geq \frac{2}{3} \right\},$$

$$SDP_N^\Psi = \left\{ (C_0, C_1) : SD(\mathbf{C}_0^\Psi, \mathbf{C}_1^\Psi) \leq \frac{1}{3} \right\}.$$

Next, we formally define fully black-box reductions from CRHFs to SZK.

Definition 4.2 (Black-Box Construction of SZK-hard Problems). *A fully black-box construction of a hard statistical distance problems (SDP) from CRHFs consists of*

- **Black-box construction:** *A collection of oracle-aided circuit pairs $\Pi^{(\cdot)} = \left\{ \Pi_n^{(\cdot)} \right\}_{n \in \mathbb{N}}$ where $\Pi_n = \left\{ (C_0^{(\cdot)}, C_1^{(\cdot)}) \in \{0, 1\}^{n \times 2} \right\}$ such that each (C_0, C_1) defines an SDP instance.*
- **Black-box security proof:** *A probabilistic oracle-aided reduction R with functions $q_R(\cdot), \epsilon_R(\cdot)$ such that the following holds: Let f be any distribution on functions. For any probabilistic oracle-aided A that decides Π in the worst-case, namely, for all $n \in \mathbb{N}$,*

$$\Pr \left[A^f(C_0, C_1) = B \quad \text{for all} \quad \begin{array}{l} (C_0, C_1) \in \Pi_n, B \in \{Y, N\} \\ \text{such that } (C_0, C_1) \in SDP_B^f \end{array} \right] = 1$$

the reduction breaks collision resistance of f , namely, for infinitely many $n \in \mathbb{N}$,

$$\Pr_f \left[f_n(x) = f_n(x') \text{ where } (x, x') \leftarrow R^{f, A} \right] \geq \epsilon_R(n),$$

where R makes at most $q_R(n)$ queries to any of its oracles (A, f) where each query to A consists of circuits C_0, C_1 each of which makes at most $q_R(n)$ queries to f .

Next, we state the main result of this section: that any fully black-box construction of SDP problems from CRHFs has to either run in time exponential in the security parameter or suffer exponential security loss.

Theorem 4.3. *For any fully black-box construction $(\Pi, R, q_R, \epsilon_R)$ of SDPs from CRHFs, the following holds:*

1. (The reduction runs in exponential time.) $q_{\mathbb{R}}(n) \geq 2^{n/10}$. Or,
2. (Reduction succeeds with exponentially small probability.) $\varepsilon_{\mathbb{R}}(n) \leq 2^{-n/10}$.

We prove the theorem by describing an oracle $\Gamma = (f, A)$ such that, A solves SDP^f but f is a CRHF relative to Γ . The rest of the section is devoted to describing this oracle and proving the theorem. We start by describing the adversary that breaks SDP: the statistical distance oracle.

4.2 The Statistical Distance Oracle

Next we describe the statistical distance oracle SDO from [BDV17] that solves SZK instances.

Definition 4.4 (Oracle SDO^Ψ). *The oracle consists of $t = \{t_n\}_{n \in \mathbb{N}}$ where $t_n : \{0, 1\}^{2n} \rightarrow (\frac{1}{3}, \frac{2}{3})$ is a uniformly random function. Given n -bit descriptions of oracle-aided circuits $C_0, C_1 \in \{0, 1\}^n$, let $t^* = t_n(C_0, C_1)$, and let $s = \text{SD}(\mathbf{C}_0^\Psi, \mathbf{C}_1^\Psi)$, return*

$$\text{SDO}^\Psi(C_0, C_1; t) := \begin{cases} 0 & \text{If } s < t^* \\ 1 & \text{If } s \geq t^* \end{cases}$$

It is immediate to see that SDO^Ψ decides SDP^Ψ in the worst-case.

Claim 4.4.1. *For any oracle Ψ ,*

$$\text{SDP}^\Psi \in \mathsf{P}^{\Psi, \text{SDO}^\Psi} .$$

Remark 4.5 (On the Oracle Used). Our separation is sensitive to the oracle used. Subsequent to [BDV17], [KY18] observed that the Simon’s collision finding oracle Coll can be used to decide SZK. Clearly, no separation between CRHFs and SZK holds relative to the Simon’s oracle. It turns out that Simon’s oracle can be used to estimate a different measure of distance between distributions, the Triangular Discrimination,¹ which like statistical distance also gives an SZK-complete promise problem [BDRV19]. Our separation does hold with a variant of Coll and SDO that measures triangular discrimination, but does not output a collision.

4.3 Insensitivity to Local Changes

Next, we recall the notions of smoothness and fairness from [BDV17] that are used to argue that the SDO^Ψ oracle is insensitive to local changes. Roughly speaking *fairness* says that the random threshold t used for a query (C_0, C_1) to SDO^Ψ is “far” from the actual statistical distance. [BDV17] show that with high probability over the choice of random threshold t , fairness holds for all queries (C_0, C_1) made to SDO^Ψ by any (relatively) efficient adversary. This intuitively means that changing the distributions $(\mathbf{C}_0^\Psi, \mathbf{C}_1^\Psi)$, on sets of small density, will not change the oracle’s answer. The proofs are included in Appendix A for completeness.

Definition 4.6 (Fairness). *The oracles (Ψ, SDO^Ψ) satisfy ε -fairness with respect to oracle-aided circuits $(C_0, C_1) \in \{0, 1\}^n$ if the statistical difference $s = \text{SD}(\mathbf{C}_0^\Psi, \mathbf{C}_1^\Psi)$ and the threshold $t^* = t_n(C_0, C_1)$ sampled by SDO are ε -far:*

$$|s - t^*| \geq \varepsilon .$$

¹The triangular discrimination is defined as $\text{TD}(X, Y) = \frac{1}{2} \sum_x \frac{(\Pr[X=x] - \Pr[Y=x])^2}{(\Pr[X=x] + \Pr[Y=x])}$. This measure also lies in the interval $[0, 1]$ and is a metric.

For an adversary A , we denote by $\text{farness}(A, \Psi, \varepsilon)$ the event that $\Gamma = (\Psi, \text{SDO}^\Psi)$ satisfies ε -farness for all queries (C_0, C_1) made by A to SDO^Ψ .

Lemma 4.7 ([BDV17](Claim 3.7)). *Fix any Ψ and any oracle-aided adversary A such that $A^{\Psi, \text{SDO}^\Psi}$ makes at most q queries to SDO^Ψ . Then*

$$\Pr_t[\text{farness}(A, \Psi, \varepsilon)] \geq 1 - 6q\varepsilon ,$$

where the probability is over the choice t of random thresholds by SDO .

We now turn to define the notion of *smoothness*. Roughly speaking we will say that an oracle-aided circuit C is smooth with respect to some oracle Ψ if any specific oracle query is only made with small probability. In particular, for a pair of smooth circuits (C_0, C_1) , local changes to the oracle Ψ should not change significantly the statistical distance $s = \text{SD}(C_0^\Psi, C_1^\Psi)$.

Definition 4.8 ((Ψ, ε) -Smoothness). *A circuit $C^{(\cdot)}$ is ε -smooth w.r.t oracle Ψ , if every location $x \in \{0, 1\}^*$ is queried with probability at most ε . That is,*

$$\max_x \Pr_w[C^\Psi(w) \text{ queries } \Psi \text{ at } x] < \varepsilon .$$

Lemma 4.9 ([BDV17](Claim 3.9)). *Let Ψ, Ψ' be oracles that differ on at most c values in the domain. Let C_0 and C_1 be ε -smooth w.r.t Ψ . Let $s = \text{SD}(C_0^\Psi, C_1^\Psi)$ and $s' = \text{SD}(C_0^{\Psi'}, C_1^{\Psi'})$ then $|s - s'| \leq 2c\varepsilon$.*

The above roughly means that (under the likely event that farness holds) making smooth queries should not help the adversary detect local changes in the oracle Ψ . [BDV17] show that we can always “smoothen” the adversary’s circuit at the expense of making (a few) more queries to Ψ , which intuitively deems the statistical difference oracle SDO^Ψ useless altogether for detecting local changes in Ψ .

In what follows, a (q', q) -query algorithm A makes at most q' queries to the oracle Ψ and q queries to SDO^Ψ such that for each query (C_0, C_1) to SDO , the circuits C_0, C_1 themselves make at most q queries to Ψ on any input.

Lemma 4.10 (Smoothing Lemma for SDO [BDV17](Lemma 3.10)). *For any (q, q) -query algorithm A and $\beta \in \mathbb{N}$, there exists a $(q + 2\beta q^2, q)$ -query algorithm S such that for any input $z \in \{0, 1\}^*$ and oracles Ψ, SDO^Ψ :*

1. $S^{\Psi, \text{SDO}^\Psi}(z)$ perfectly simulates the output of $A^{\Psi, \text{SDO}^\Psi}(z)$,
2. $S^{\Psi, \text{SDO}^\Psi}(z)$ only makes queries (C_0, C_1) where both C_0, C_1 are (Ψ, ε) -smooth queries to SDO^Ψ with probability:

$$\Pr_S[\text{smooth}(S, \Psi, \varepsilon)] \geq 1 - 2^{-\varepsilon\beta + \log(2q^2/\varepsilon)} ,$$

over its own random coin tosses.

4.4 Collision Resistance in the Presence of SDO Oracle.

In this section, we prove the oracle separation between collision resistant hash functions and SZK.

Let \mathcal{F}_n be the set of all functions from $\{0, 1\}^n$ to $\{0, 1\}^{m(n)}$ where $m(n) < n$ is a shrinking function. Let $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ denote the family of these sets of functions. Let $\mathcal{T} = \{\mathcal{T}_n\}_{n \in \mathbb{N}}$ where \mathcal{T}_n denotes the set of threshold functions $t : \{0, 1\}^n \rightarrow (1/3, 2/3)$.²

Definition 4.11 (The Oracle f). *The oracle $f = \{f_n\}_{n \in \mathbb{N}}$ on input $x \in \{0, 1\}^n$ returns $f_n(x)$ where $f_n : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is a random function from \mathcal{F}_n .*

The oracle we consider is $\Gamma = (f, \text{SDO}^f)$. It is easy to see that all $\text{SDP}^f \in \text{Pf}^{f, \text{SDO}^f}$. What remains to show is that f is still collision resistant in the presence of the SDO^f oracle. We do so next.

Theorem 4.12. *Let A be a (q, q) query adversary for $q = O(2^{m/10})$. Then,*

$$\Pr \left[f_n(x) = f_n(x') \text{ where } (x, x') \leftarrow A^{f, \text{SDO}^f}(1^n) \right] \leq 2^{-m/10}.$$

Proof. Fix oracle $f_{-n} = \{f_k\}_{k \neq n}$ arbitrarily. Consider the $(q + 2\beta q^2, q)$ query smooth version S , of A given by Lemma 4.10 for $\beta = 2^{m/5} \cdot m$ and $\varepsilon = 2^{-m/5}$. We assume w.l.o.g that S makes no repeated oracle queries and that whenever S outputs a collision (x, x') , x is its last oracle query and x' is a previous query (both to the f oracle).

The first assumption is w.l.o.g because S may store a table of previously made queries and answers. The second is w.l.o.g because S may halt once its f -queries include a collision and output that collision; also, if one, or both, outputs x, x' have not been queried, S can query it at the end (and if needed change the order of the output so that x is queries last). The latter costs at most two additional queries, and does not affect the smoothness of S .

Next, we define some notation about transcripts generated in the process.

Transcripts. A transcript π consists of all queries asked and answers received by S to the oracle (f, SDO^f) . Let x_i denote the i -th query to the f -oracle. We say that $x \notin \pi$ if the location x is not among the queries explicitly made in π .

The Underlying Joint Distribution. The proof infers properties of the joint distribution (f, t, π) consisting of the oracle f , the SDO oracle's random thresholds t and the transcript generated by S . The distribution is generated as follows: $f \leftarrow \mathcal{F}$ and $t \leftarrow \mathcal{T}$ and $\pi \leftarrow S^{f, \text{SDO}^{f;t}}$ where $\text{SDO}^{f;t}$ denotes running the SDO oracle with random thresholds t . Denote this distribution by $P_{FT\Pi}$.

Note that given f, t , the transcript π is generated in a deterministic manner as S is deterministic and the oracle's behavior is completely specified. Furthermore, we also consider partial transcripts obtained by running S and stopping after i queries. This transcript is denoted by $\pi_{<i}, x_i$: that is the $\pi_{<i}$ consists of queries and responses received and x_i is the next query to the oracle f . Note that x_i is a deterministic function of $\pi_{<i}$. Given the distribution $P_{FT\Pi}$, the conditional distributions $P_{FT|\Pi=\pi}$ or $P_{FT|\Pi=\pi_{<i}}$ are well defined: these consist of uniform distribution on pairs (f, t) that when run using S result in the transcript being π (or $\pi_{<i}$).

The Good Event. We define the concept of Good transcripts. Roughly speaking, these are transcripts π that satisfy sufficient smoothness and farness so to guarantee that the value $f(x)$ at any $x \notin \pi$ is completely hidden.

²While we describe the threshold function as a real valued function, it can be safely discretized because statistical distance for any pair of circuits $C_0, C_1 : \{0, 1\}^m \rightarrow \{0, 1\}^*$, takes values that are multiples of $2^{-(m+1)}$. We omit the details here.

Definition 4.13 (Good). A tuple $(f, t, \pi, x, \varepsilon)$ is good, denoted by $\text{good}(f, t, \pi, x, \varepsilon)$ if the following hold:

1. $\pi = \mathbf{S}^{f_{x \rightarrow \perp}, \text{SDO}^{f_{x \rightarrow \perp} \cdot t}}(1^n)$, where $f_{x \rightarrow \perp}$ is the function equal to f everywhere except at x where it takes the value \perp .
2. (x is not explicitly queried.) $x \notin \pi$.
3. (Transcript is smooth.) For every query made by $\mathbf{S}^{f_{x \rightarrow \perp}, \text{SDO}^{f_{x \rightarrow \perp} \cdot t}}(1^n)$, the event $\text{smooth}(f_{x \rightarrow \perp}, 2\varepsilon)$ holds. Denote this event by $\text{smooth}(f_{x \rightarrow \perp}, t, \pi, 2\varepsilon)$.
4. (Transcript is far.) For every query by $\mathbf{S}^{f_{x \rightarrow \perp}, \text{SDO}^{f_{x \rightarrow \perp} \cdot t}}(1^n)$, event $\text{far}(f_{x \rightarrow \perp}, 6\varepsilon)$ holds. Denote this by $\text{far}(f, t, \pi, 12\varepsilon)$.

The key reason for using $f_{x \rightarrow \perp}$ instead of f in the definition is that when an execution of $\mathbf{S}^{f_{x \rightarrow \perp}, \text{SDO}^{f_{x \rightarrow \perp} \cdot t}}$ generates a transcript π while making only smooth and far queries, all executions of $\mathbf{S}^{f_{x \rightarrow z}, \text{SDO}^{f_{x \rightarrow z} \cdot t}}$ for all z , also generate π while not necessarily being smooth or far themselves.

A tuple (f, t, π, ε) is good if for all $x \notin \pi$, $\text{good}(f, t, \pi, x, \varepsilon)$ holds.

Lemma 4.14. Let $P_{FT\Pi}$ as defined above. Then,

$$\Pr_{(f,t,\pi) \leftarrow P_{FT\Pi}} [\text{good}(f, t, \pi, \varepsilon)] \geq 1 - 16q\varepsilon - 2^{-\beta\varepsilon + \log(2q^2/\varepsilon)}$$

The same holds for i -length partial transcripts generated as well, for all i .

Lemma 4.15. For any transcript π and query $x \notin \pi$ such that $\Pr_{(f,t,\pi) \leftarrow P_{FT\Pi}} [\text{good}(f, t, \pi, x, \varepsilon)] > 0$, it holds that,

$$\{f(x) : (f, t) \leftarrow P_{FT|\Pi=\pi, \text{good}(f,t,\pi,x,\varepsilon)}\} \equiv U_m.$$

Next, we prove Theorem 4.12 assuming Lemmas 4.14 and 4.15. Then, we prove the two lemmas. Let $\text{hit}(\pi)$ denote the event that π contains two queries x, x' such that $f_n(x) = f_n(x')$. Then,

$$\begin{aligned} \Pr_{f,t} [f_n(x) = f_n(x') \wedge (x, x') = \mathbf{S}^{f, \text{SDO}^{f \cdot t}}(1^n)] &= \Pr_{f,t,\pi} [\text{hit}(\pi)] \\ &\leq \Pr_{f,t,\pi} [\text{hit}(\pi) \wedge \text{good}(f, t, \pi, \varepsilon)] + \Pr_{f,t,\pi} [\overline{\text{good}(f, t, \pi, \varepsilon)}] \end{aligned}$$

We will bound the two terms separately. The first term will involve using Lemma 4.15 while the second term is bound using Lemmas 4.7 and 4.10. We begin by bounding the first term. This is done by decomposing the probability of hitting a collision by the first query that hits a collision:

$$\begin{aligned} &\Pr_{f,t} [\text{hit}(\pi) \wedge \text{good}(f, t, \pi, \varepsilon)] \\ &\leq \sum_i \Pr_{f,t} [\text{hit}(\pi_{\leq i}) \wedge \overline{\text{hit}(\pi_{< i})} \wedge \text{good}(f, t, \pi_{< i}, \varepsilon)] \\ &= \sum_i \Pr_{f,t} [f(x_i) \in \text{hitSet}(\pi_{< i}) \wedge \overline{\text{hit}(\pi_{< i})} \wedge \text{good}(f, t, \pi_{< i}, \varepsilon)] \end{aligned}$$

where $x_i \notin \pi$ denotes the i -th f query made by \mathbf{S} and $\text{hitSet}(\pi_{< i})$ denotes the answers to f -queries in $\pi_{< i}$,

$$= \sum_i \sum_{\pi_{< i}, x_i} \Pr_{f,t} [(\pi_{< i}, x_i) = \mathbf{S}^{f, \text{SDO}^{f \cdot t}}(1^n) \wedge \text{good}(f, t, \pi_{< i}, x_i, \varepsilon)] \cdot \Pr_{f,t \leftarrow P_{FT|\Pi=\pi_{< i}, \text{good}}} [f(x_i) \in \text{hitSet}(\pi_{< i})]$$

The last equality follows from the definition of conditional probability. At this point, we can use Lemma 4.15 to argue that $\Pr_{f,t \leftarrow P_{FT} | \Pi = \pi_{<i}, \text{good}(f,t,\pi_{<i},x_i,\varepsilon)} [f(x_i) \in \text{hitSet}(\pi_{<i})] \leq \frac{i}{2^m}$ because $f(x_i)$ is uniformly random and $|\text{hitSet}(\pi_{<i})| \leq i$. Hence, we get that,

$$\begin{aligned} &\leq \sum_i \frac{i}{2^m} \cdot \sum_{\pi_{<i}, x_i} \Pr_{f,t} \left[(\pi_{<i}, x_i) = \mathcal{S}^{f, \text{SDO}^{f,t}}(1^n) \wedge \text{good}(f, t, \pi_{<i}, x_i, \varepsilon) \right] \\ &\leq \sum_{i=1}^{q'} \frac{i}{2^m} \leq \frac{q'^2}{2^m}. \end{aligned}$$

where $q' = q + 2\beta q^2 + 2$, the number queries that \mathcal{S} makes to f .

Hence, by Lemma 4.14, the algorithm's success probability is bounded by

$$\begin{aligned} \Pr_{f,t} \left[f_n(x) = f_n(x') \wedge (x, x') = \mathcal{S}^{f, \text{SDO}^{f,t}}(1^n) \right] &\leq \Pr_{f,t} \left[\text{hit}(\pi) \wedge \text{good}(f, t, \pi) \right] + \Pr_{f,t} \left[\overline{\text{good}(f, t, \pi)} \right] \\ &\leq \frac{(q + 2\beta q^2 + 2)^2}{2^m} + 16q\varepsilon + 2^{-\beta\varepsilon + \log(2q^2/\varepsilon)} \\ &\leq O(q^4 \beta^2 2^{-m} + 16q\varepsilon + q^2/\varepsilon 2^{-\varepsilon\beta}) \\ &\leq O(2^{-m/10}). \end{aligned}$$

when substituting $\varepsilon = 2^{-m/5}$, $\beta = 2^{m/5} \cdot m$, and $q \leq 2^{m/10}$.

Proof of Lemma 4.14. The proof follows from the observation if $\mathcal{S}^{f, \text{SDO}^f}$ outputs π with all the queries being both smooth, and far, then, the same holds for $\mathcal{S}^{f_{x \rightarrow \perp}, \text{SDO}^{f_{x \rightarrow \perp}}}$ with slightly degraded parameters. That is,

$$\begin{aligned} \Pr_{(f,t,\pi) \leftarrow P_{FT\Pi}} \left[\text{good}(f, t, \pi, \varepsilon) \right] &= \Pr_{f,t,\pi} \left[\bigwedge_{x \notin \pi} \text{good}(f, t, \pi, x, \varepsilon) \right] \\ &\geq \Pr_{f,t,\pi} \left[\text{smooth}(f, t, \pi, \varepsilon) \wedge \text{farness}(f, t, \pi, 8\varepsilon) \right] \\ &\geq 1 - 16\varepsilon q - 2^{-\beta\varepsilon + \log(2q^2/\varepsilon)} \end{aligned}$$

Hence, to complete the proof, we need to show that, for any (f, t) if $\mathcal{S}^{f, \text{SDO}^f}(1^n)$ outputs π with all the queries being (f, ε) -smooth, and $(f, 16\varepsilon)$ -far, then, $\mathcal{S}^{f_{x \rightarrow \perp}, \text{SDO}^{f_{x \rightarrow \perp}}}(1^n)$ generates π with all the queries being $(f, 2\varepsilon)$ -smooth and $(f, 12\varepsilon)$ -far.

First observe that by Lemma 4.9, since 16ε -farness and ε -smoothness hold, answers by $\text{SDO}^{f_{x \rightarrow \perp}}$ are identical to those by SDO^f . Accordingly, the transcript $\pi = \mathcal{S}^{f_{x \rightarrow \perp}, \text{SDO}^{f_{x \rightarrow \perp}}}(1^n)$.

Next, we show that 2ε -smoothness holds with respect to $\text{SDO}^{f_{x \rightarrow \perp}}$. Indeed, any SDO-query $(C_0^{(\cdot)}, C_1^{(\cdot)})$ is ε -smooth with respect to f , accordingly the probability that either circuit C_b queries any individual z is bounded by

$$\Pr \left[C_b^{f_{x \rightarrow \perp}} \text{ queries } z \right] \leq \Pr \left[C_b^f \text{ queries } x \right] + \Pr \left[C_b^f \text{ queries } z \right] \leq 2\varepsilon.$$

Finally, to conclude the proof, we show that 12ε -farness holds with respect to $f_{x \rightarrow \perp}$. Indeed, for any query (C_0, C_1) , let $s = \text{SD}(C_0^f, C_1^f)$ be the statistical distance with respect to f , then by ε -smoothness with respect to f , the statistical distance $s^x = \text{SD}(C_0^{f_{x \rightarrow \perp}}, C_1^{f_{x \rightarrow \perp}})$ with respect to $f_{x \rightarrow \perp}$

is at most 2ε -far from s . Letting $t = t(C_0, C_1)$ be the threshold chosen by SDO, we know by 16ε -farness that $|s - t| \geq 16\varepsilon$ and thus $|s^x - t| \geq 12\varepsilon$, which implies the require farness with respect to $f_{x \rightarrow \perp}$.

The above argument holds unaltered for partial transcripts output by S as well. Even there, when a partial transcript is output by S^{f, SDO^f} with all queries being (f, ε) -smooth and $(f, 16\varepsilon)$ -far, then, $S^{f_{x \rightarrow \perp}, \text{SDO}^{f_{x \rightarrow \perp}}}(1^n)$ generates the same partial transcript with all the queries being $(f, 2\varepsilon)$ -smooth and $(f, 12\varepsilon)$ -far. \square

Proof of Lemma 4.15. Given $\pi, x \notin \pi$, for any y

$$\Pr_{f, t \leftarrow P_{FT} | \Pi = \pi, \text{good}(f, t, \pi, x, \varepsilon)} [f(x) = y] = \frac{\Pr_{f, t} [\pi = S^{f, \text{SDO}^{f, t}}(1^n) \wedge f(x) = y \wedge \text{good}(f, t, \pi, x, \varepsilon)]}{\Pr_{f, t} [\pi = S^{f, \text{SDO}^t}(1^n) \wedge \text{good}(f, t, \pi, x, \varepsilon)]}$$

In order to show that, the distribution $\{f(x) : f \leftarrow P_{F | \Pi = \pi, \text{good}}\}$ is uniform, it suffices to show that for all $y_1, y_2 \in \{0, 1\}^m$,

$$\begin{aligned} & \Pr_{f, t} [\pi = S^{f, \text{SDO}^{f, t}}(1^n) \wedge f(x) = y_1 \wedge \text{good}(f, t, \pi, x, \varepsilon)] \\ &= \Pr_{f, t} [\pi = S^{f, \text{SDO}^{f, t}}(1^n) \wedge f(x) = y_2 \wedge \text{good}(f, t, \pi, x, \varepsilon)] \end{aligned}$$

To prove this, it suffices to show that that for every (f, t) where $f(x) = y_1$,

$$\pi = S^{f, \text{SDO}^{f, t}}(1^n) \wedge \text{good}(f, t, \pi, x, \varepsilon) = 1 \iff \pi = S^{f_{x \rightarrow y_2}, \text{SDO}^{f_{x \rightarrow y_2}, t}}(1^n) \wedge \text{good}(f_{x \rightarrow y_2}, t, \pi, x, \varepsilon)$$

This follows because as $\text{good}(f, t, \pi, x, \varepsilon)$ holds, $\pi = S^{f_{x \rightarrow \perp}, \text{SDO}^{f_{x \rightarrow \perp}, t}}(1^n)$ and every query made to $\text{SDO}^{f_{x \rightarrow \perp}, t}$ is both 12ε -far and 2ε -smooth. Hence, when we change the oracle to $(f_{x \rightarrow y_2}, \text{SDO}^{f_{x \rightarrow y_2}, t})$, each query is answered identically to $f_{x \rightarrow \perp}, \text{SDO}^{f_{x \rightarrow \perp}, t}$. Indeed, for any query (C_0, C_1) , let $s = \text{SD}(C_0^{f_{x \rightarrow \perp}}, C_1^{f_{x \rightarrow \perp}})$ be their statistical distance with respect to $f_{x \rightarrow \perp}$, then by 2ε -smoothness with respect to $f_{x \rightarrow \perp}$, the statistical distance $s' = \text{SD}(C_0^{f_{x \rightarrow y_2}}, C_1^{f_{x \rightarrow y_2}})$ is at most 4ε -far from s . As the threshold $t = t(C_0, C_1)$ is more than 12ε far by farness, the answer will be unchanged to this query.

Hence, $S^{(f_{x \rightarrow y_2}, \text{SDO}^{f_{x \rightarrow y_2}, t})}$ will also return π as the answer. Also, by definition, $\text{good}(f_{x \rightarrow y_2}, t, \pi, x, \varepsilon)$ will hold because $\pi = S^{f_{x \rightarrow \perp}, \text{SDO}^{f_{x \rightarrow \perp}, t}}(1^n)$ and every query made to $\text{SDO}^{f_{x \rightarrow \perp}, t}$ is both 12ε -far and 2ε -smooth. Hence, the claim follows. \square

This completes the proof of Theorem 4.12. \square

5 New Proofs of Old Separations

In this section, we give new proofs of results by Simon [Sim98] and Asharov-Segev [AS16].

5.1 Oracle-Aided Indistinguishability Obfuscation

The notion of *indistinguishability obfuscation* (IO) was introduced by Barak et al. [BGI⁺01] and the first candidate construction was demonstrated in the work of Garg et al. [GGH⁺13]. Since then, IO

has given rise to a plethora of applications in cryptography and beyond. Nevertheless, Asharov and Segev [AS15, AS16] demonstrated that IO is insufficient to achieve some cryptographic tasks, most notably (domain-invariant) one-way permutations, collision-resistant hashing. To formally show such a statement, they introduced the framework of indistinguishability obfuscation for oracle-aided circuits. We give a different proof for their result that IO does not imply collision resistance, for which we follow their framework.

We begin by recalling the notion of two oracle-aided circuits being equivalent, and move on to defining IO relative to oracles.

Definition 5.1. Let C_0 and C_1 be two oracle-aided circuits and let f be a function. C_0 and C_1 are said to be functionally equivalent relative to f , denoted as $C_0^f \equiv C_1^f$, if for every input x , $C_0^f(x) = C_1^f(x)$.

Definition 5.2. Let $\mathcal{C} = \{C_n\}_{n \in \mathbb{N}}$ be a class of oracle-aided circuits, where each $C \in \mathcal{C}_n$ is of size n .³ A PPT algorithm O is an indistinguishability obfuscator for \mathcal{C} relative to an oracle distribution ensemble $\Gamma = \{\Gamma_n\}_{n \in \mathbb{N}}$ if the following conditions are met:

1. **Functionality.** For all $n \in \mathbb{N}$ and for all $C \in \mathcal{C}_n$ it holds that

$$\Pr_{\Gamma, O} \left[C^\Gamma \equiv \widehat{C}^\Gamma \mid \widehat{C} \leftarrow O^\Gamma(1^n, C) \right] = 1 .$$

2. **Indistinguishability.** For any non-uniform PPT distinguisher $D = (D_1, D_2)$ there exists a negligible function negl such that for all $n \in \mathbb{N}$

$$\text{Adv}_{\Gamma, O, \mathcal{C}, D}^{\text{IO}}(n) = \left| \Pr \left[\text{Expt}_{\Gamma, O, \mathcal{C}, D}^{\text{IO}}(n) = 1 \right] - \frac{1}{2} \right| \leq \text{negl}(n)$$

where the random variable $\text{Expt}_{\Gamma, O, \mathcal{C}, D}^{\text{IO}}(n)$ is defined via the following experiment:

- (a) $b \leftarrow \{0, 1\}$.
- (b) $(C_0, C_1, \text{state}) \leftarrow D_1^\Gamma(1^n)$ where $C_0, C_1 \in \mathcal{C}_n$ and $C_0^\Gamma \equiv C_1^\Gamma$.
- (c) $\widehat{C} \leftarrow O^\Gamma(1^n, C_b)$.
- (d) $b' = D_2^\Gamma(\text{state}, \widehat{C})$.
- (e) If $b = b'$ output 1 else output 0.

We further say that O satisfies δ -indistinguishability if the above negligible advantage is at most δ .

5.1.1 Fully Black Box Constructions of CRHFs from IO and OWPs

We begin by defining oracle-aided constructions of CRHFs and then specialize it to the setting of IO (and OWPs).

Definition 5.3 (Oracle-Aided Collision-Resistant Function Families). A pair of polynomial-time oracle-aided algorithms $(\text{Gen}, \text{Eval})$ is a collision-resistant function family relative to an oracle Γ if it satisfies the following properties:

³As in [AS15], we assume throughout that the size of the obfuscated circuits equals the security parameter. This is only for simplicity of notation, and is without loss of generality as the circuits can be padded up if they are too small, and the security parameter can be polynomially increased if the circuits are too large.

- The index-generation algorithm Gen is a probabilistic algorithm that on input 1^n and oracle access to Γ outputs a function index $\sigma \in \{0, 1\}^{m(n)}$.
- The evaluation algorithm Eval is a deterministic algorithm that takes as input a function index $\sigma \in \{0, 1\}^{m(n)}$ and a string $x \in \{0, 1\}^n$, has oracle access to Γ , and outputs a string $y = \text{Eval}^\Gamma(\sigma, x) \in \{0, 1\}^{n-1}$.

Definition 5.4 (Black-Box Construction of CRHFs from OWPs). A fully black-box construction of a Collision Resistant Hash Functions (CRHFs) from One-Way Permutations consists of a pair of PPT oracle-aided algorithms $(\text{Gen}, \text{Eval})$, an oracle-reduction R along with functions $q_{\text{R}}(n)$, $\varepsilon_{\text{R}}(n)$ such that the following two conditions hold:

- **Correctness:** For any $n \in \mathbb{N}$, for any permutation f , and for any function index σ produced by $\text{Gen}^f(1^n)$, it holds that $\text{Eval}^f(\sigma, \cdot) : \{0, 1\}^n \rightarrow \{0, 1\}^{n-1}$.
- **Black-box security proof:** For any permutation f and probabilistic oracle-aided algorithm A , if

$$\Pr[\text{Eval}^f(\sigma, x) = \text{Eval}^f(\sigma, x') \wedge x \neq x'] \geq \frac{1}{2}$$

where the experiment is $\sigma \leftarrow \text{Gen}^f(1^n)$ and $(x, x') \leftarrow \text{A}^f(1^n, \sigma)$, for infinitely many n , then the reduction breaks f , namely, for infinitely many $n \in \mathbb{N}$ either

$$\Pr_{\substack{x \leftarrow \{0, 1\}^n \\ f, \text{A}}}[\text{R}^{\text{A}, f}(f_n(x)) = x] \geq \varepsilon_{\text{R}}(n),$$

for infinitely many values of n where R makes at most $q_{\text{R}}(n)$ queries to the oracles A , f and for every circuit $D^{(\cdot)}$ queried to A makes at most $q_{\text{R}}(n)$ queries to f on any input.

Definition 5.5 (Black-Box Construction of CRHFs from IO and OWPs). A fully black-box construction of a Collision Resistant Hash Functions (CRHFs) from OWPs and Indistinguishability Obfuscation for a class $\mathcal{C} = \{C_n\}_{n \in \mathbb{N}}$ of all polynomial-sized oracle-circuits consists of a pair of PPT oracle-aided algorithms $(\text{Gen}, \text{Eval})$, an oracle-reduction R along with functions $q_{\text{R}}(n)$, $\varepsilon_{\text{R}}(n)$ such that the following two conditions hold:

- **Correctness:** For any $n \in \mathbb{N}$, for any permutation f and any function $i\mathcal{O}$ such that $i\mathcal{O}(C; r)^f \equiv C$ for all $C \in \mathcal{C}$ and $r \in \{0, 1\}^*$, and for any function index σ produced by $\text{Gen}^{f, i\mathcal{O}}(1^n)$, it holds that $\text{Eval}^{f, i\mathcal{O}}(\sigma, \cdot) : \{0, 1\}^n \rightarrow \{0, 1\}^{n-1}$.
- **Black-box security proof:** For any permutation f and any function $i\mathcal{O}$ such that $i\mathcal{O}(C; r)^f \equiv C$ for all $C \in \mathcal{C}$ and $r \in \{0, 1\}^*$, for any probabilistic oracle-aided algorithm A , if

$$\Pr[\text{Eval}^{f, i\mathcal{O}}(\sigma, x) = \text{Eval}^{f, i\mathcal{O}}(\sigma, x') \wedge x \neq x'] \geq \frac{1}{2}$$

where the experiment is $\sigma \leftarrow \text{Gen}^{f, i\mathcal{O}}(1^n)$ and $(x, x') \leftarrow \text{A}^{f, i\mathcal{O}}(1^n, \sigma)$, for infinitely many n , then the reduction breaks either f or $i\mathcal{O}$, namely, for infinitely many $n \in \mathbb{N}$ either

$$\Pr_{\substack{x \leftarrow \{0, 1\}^n \\ f, i\mathcal{O}, \text{A}}}[\text{R}^{\text{A}, f, i\mathcal{O}}(f_n(x)) = x] \geq \varepsilon_{\text{R}}(n),$$

or

$$\left| \Pr \left[\text{Expt}_{(f, i\mathcal{O}), i\mathcal{O}, \text{CRA}}^{\text{IO}}(n) = 1 \right] - \frac{1}{2} \right| \geq \varepsilon_{\text{R}}(n) ,$$

for infinitely many values of n . The random variable $\text{Expt}_{(f, i\mathcal{O}), i\mathcal{O}, \text{CRA}}^{\text{IO}}$ represents the reductions winning probability in the IO security game (Definition 5.2) relative to $(f, i\mathcal{O})$ where R makes at most $q_{\text{R}}(n)$ queries to the oracles A , f , $i\mathcal{O}$ and for every oracle-aided circuit query $C^{(\cdot)}$ to $i\mathcal{O}$, C itself makes at most $q_{\text{R}}(n)$ queries to f . Furthermore, every query to A consisting of a circuit $D^{(\cdot)}$ follows the conditions above: that is $q_{\text{R}}(n)$ queries in total to the oracles $f, i\mathcal{O}$, with each $i\mathcal{O}$ query making at most $q_{\text{R}}(n)$ queries to the f oracle.

We remark that ruling out black-box reductions as defined in Definition 5.5 where the reduction has to break the OWP or IO given an adversary that breaks CRHFs w.p. over 1/2 only makes our result stronger. In the standard setting, the reduction has to break IO (or OWP) given an adversary that succeeds with any noticeable probability.

5.2 Simon's Collision Finding Oracle and Puncturing

Recall that the Simon's collision finding oracle is defined as follows:

Definition 5.6 (Simon's Oracle Coll^{Ψ}). *Given any description of a circuit C with m -bit inputs, the oracle's randomness contains a random input $w_{\text{C}} \in \{0, 1\}^m$ and a random permutation $\pi_{\text{C}} : \{0, 1\}^m \rightarrow \{0, 1\}^m$. The Coll^{Ψ} oracle returns the following:*

$$\text{Coll}^{\Psi}(C) := (w_{\text{C}}, w'_{\text{C}}) \text{ where } w'_{\text{C}} = \pi_{\text{C}}(i) \text{ for the smallest } i \text{ such that } C^{\Psi}(w_{\text{C}}) = C^{\Psi}(\pi_{\text{C}}(i)).$$

W.l.o.g, along with $(w_{\text{C}}, w'_{\text{C}})$, let Coll also return the queries made to Ψ , and their answers, when evaluating $C^{\Psi}(w_{\text{C}})$ and $C^{\Psi}(w'_{\text{C}})$.

The collision-finding oracle breaks any oracle-aided collision resistant hash function.

Lemma 5.7 ([Sim98]). *Let $\Gamma = (\Psi, \text{Coll}^{\Psi})$. Let $C^{(\cdot)} : \{0, 1\}^n \rightarrow \{0, 1\}^{n-1}$ be any candidate construction of CRHFs. Then,*

$$\Pr[C^{\Psi}(w) = C^{\Psi}(w') \wedge w \neq w' \text{ where, } (w, w') \leftarrow \text{Coll}^{\Psi}(C)] \geq \frac{1}{2}$$

where the randomness is over the randomness of Coll .

Proof. Fix Ψ and omit it from the notation. For any string $y \in \{0, 1\}^{n-1}$, let $a_y = |\{x : C(x) = y\}|$. Then,

$$\begin{aligned} \Pr[w \neq w'] &= \sum_{y \in \text{Supp}(C)} \Pr_{w, w' \leftarrow C^{-1}(y)} [w \neq w'] \cdot \Pr_w [C(w) = y] \\ &= \sum_{y \in \text{Supp}(C)} \frac{a_y - 1}{a_y} \cdot \frac{a_y}{2^n} \\ &= \sum_{y \in \text{Supp}(C)} \frac{a_y}{2^n} - \sum_{y \in \text{Supp}(C)} \frac{1}{2^n} \geq 1 - \frac{2^{n-1}}{2^n}, \end{aligned}$$

where the second inequality follows from the fact that $\Pr_{w, w' \leftarrow C^{-1}(y)} [w \neq w'] = \Pr_{w' \leftarrow C^{-1}(y)} [w' \neq w] = \frac{a_y - 1}{a_y}$. \square

Next we define a variant of the Simon's oracle, dubbed as the punctured Simon's oracle. This collision finding oracle allows Ψ to be punctured, that is, a set of values in Ψ are erased. As we will show later, this oracle returns the same answers as Coll^Ψ most of the time, and we can characterize when it does not.

Definition 5.8 (Punctured Simon's Oracle PColl_S^Ψ). *Let $\Psi : \{0, 1\}^* \rightarrow \{0, 1\}^*$ be an oracle. Let $S \subseteq \{0, 1\}^*$ be a subset of inputs. The oracle PColl_S^Ψ 's randomness contains for any circuit C with m -bit inputs, a random input $w_C \in \{0, 1\}^m$ and a random permutation $\pi_C : \{0, 1\}^m \rightarrow \{0, 1\}^m$. The PColl_S^Ψ oracle returns the following:*

$$\text{PColl}_S^\Psi(C) = \perp, \text{ if } C^\Psi(w_C) \text{ queries any } x \in S.$$

Else,

$$\text{PColl}_S^\Psi(C) := (w_C, w'_C)$$

where $w'_C = \pi_C(i)$ for the smallest i such that $C^\Psi(w_C) = C^\Psi(\pi_C(i))$ and $C^\Psi(\pi_C(i))$ does not query any $x \in S$. Along with (w_C, w'_C) , let it also return the queries made to Ψ when evaluating $C^\Psi(w_C)$ and $C^\Psi(w'_C)$. We refer to these queries as Ψ queries induced by the Coll oracle.

There are two key properties of the punctured oracle: (1) The answers of PColl_S^Ψ are independent of the values of the oracle Ψ on all of S ; and (2) there is a natural coupling between Coll^Ψ and PColl_S^Ψ such that, as long as there is no *explicit* query $x \in S$ to Ψ , the two oracles return identical answers. This is captured by the following lemma.

Lemma 5.9. *Let $\Psi : \{0, 1\}^* \rightarrow \{0, 1\}^*$ be an oracle, let $S \subseteq \{0, 1\}^*$. Consider the coupling of Coll^Ψ and PColl_S^Ψ that instantiates the two oracles with identical randomness. Let A be any deterministic oracle-aided algorithm. Let τ be the transcript generated by $A^{\Psi, \text{Coll}^\Psi}$. Then,*

$$A^{\Psi, \text{PColl}_S^\Psi} = \tau \text{ if and only if, } \Psi\text{-set}(\tau) \cap S = \emptyset,$$

where $\Psi\text{-set}(\tau)$ is the set of all queries made to Ψ in the execution. This includes the queries to Ψ returned by the Coll oracle.

Proof. Every direct query to Ψ by A is returned identically in both the executions. Furthermore, in any transcript τ , such that $\Psi\text{-set}(\tau) \cap S = \emptyset$, all queries to Coll^Ψ and PColl_S^Ψ are answered identically. This follows from the definition of PColl because for every query C to Coll and response (w_C, w'_C) , all the queries made to Ψ when evaluating $C^\Psi(w_C)$ and $C^\Psi(w'_C)$ are explicitly made directly to Ψ , and are thus in $\Psi\text{-set}$. In more detail, for any query C^Ψ made to Coll^Ψ with answer (w_C, w'_C) , $C^\Psi(w_C)$ does not make any queries in S , and thus PColl , will also return w_C . In addition, any w'' that is lexicographically prior to w'_C will not be returned because it either induces queries in S , or if it does then it is such that $C^\Psi(w'') \neq C^\Psi(w_C)$. In contrast, $C(w'_C)$ does not make any queries to S , and is such that $C(w'_C) = C(w_C)$. Hence w'_C will also be returned by PColl (and likewise the queries to Ψ induced by w_C, w'_C).

□

A Word of Caution. In Lemma 4.15, we showed that the distribution $f(x)$ when conditioned on a transcript τ is close to uniformly random.⁴

$$\{f(x) : f \leftarrow P_{F|\Pi=\pi, \text{good}}\} \equiv U_m$$

⁴We are using τ for transcript here to avoid the ambiguity with the Coll oracle randomness π .

Lemma 5.9 seems to suggest the same for the collision finding oracle. That is, the oracle reveals no information about $f(x)$ for any location x not explicitly queried in τ . Unfortunately, we do not know how to show this. The key reason for this is that the probability of seeing this transcript τ could itself depend on the value of $f(x)$. This issue is not new: it also comes up with the SDO oracle. We are able to remedy this issue in the case of the SDO oracle in part because of its short output: it allows us to define the notion of farness which shows that the SDO oracle is robust to *any* small changes to the SDO oracle. Puncturing only allows us to erase a value, and not set it to a different one.

5.3 Smoothing for the Collision Finding Oracle

Similar to Lemma 4.10, we can show that any algorithm $A^{\Psi, \text{Coll}^\Psi}$ can be transformed to a smoothed algorithm $S^{\Psi, \text{Coll}^\Psi}$ that with high probability makes only smooth queries to the Coll^Ψ oracle.

A (q', q) -query algorithm A makes at most q' queries to the oracle f and q queries to Coll^f such that each for each query C to Coll , the circuit C makes at most q queries to f on any input.

Lemma 5.10 (Smoothing Lemma for Coll). *For any (q, q) -query algorithm A and $\beta \in \mathbb{N}$, there exists a $(q + \beta q^2, q)$ -query algorithm S such that for any input $z \in \{0, 1\}^*$ and oracles Ψ, Coll^Ψ :*

1. $S^{\Psi, \text{Coll}^\Psi}(z)$ perfectly simulates the output of $A^{\Psi, \text{Coll}^\Psi}(z)$,
2. $S^{\Psi, \text{Coll}^\Psi}(z)$ only makes queries that are (Ψ, ε) -smooth queries to Coll^Ψ with probability:

$$\Pr_S[\text{smooth}(S, \Psi, \varepsilon)] \geq 1 - 2^{-\varepsilon\beta + \log(q^2/\varepsilon)},$$

over its own random coin tosses.

The proof of the lemma is identical to that of Lemma 4.10, the bound differs in a factor of 2: $(q + \beta q^2)$ instead of $(q + 2\beta q^2)$ in case of Lemma 4.10 because Coll oracle takes only one circuit as input.

5.4 Warm up: One Way Permutations in the presence of Coll

In this section, we show that CRHFs cannot be constructed from OWPs in a black-box manner (Definition 5.4). That is, we show,

Theorem 5.11. *Let $(\text{Gen}, \text{Eval}, R, q_R, \varepsilon_R)$ be a fully black-box construction of CRHFs from OWPs. Then, either*

1. (Large Running Time) R makes at least $q_R(n) \geq 2^{n/6}$ queries. Or,
2. (Large Security Loss) $\varepsilon_R(n) \leq 2^{-n/6}$.

To prove the theorem, we consider the oracle $\Gamma = (f, \text{Coll}^f)$ where f is a random permutation. We show that a random permutation f is hard to invert even given access to Coll^f . We start by defining the oracle. In what follows, \mathcal{P}_n denotes the set of permutations of $\{0, 1\}^n$.

Definition 5.12 (The Oracle f). $f = \{f_n\}_{n \in \mathbb{N}}$ on input $x \in \{0, 1\}^n$ answers with $f_n(x)$ where f_n is a random permutation $f_n \leftarrow \mathcal{P}_n$.

It is clear that Coll^f breaks any potential CRHF construction with probability at least $1/2$. Our main result states that f cannot be inverted, except with exponentially small probability, even given an exponential number of oracle queries to f and Coll^f . Here, consistently with the previous subsection, we say that an adversary A is q -query if A^{f, Coll^f} makes at most q queries to f and q queries to Coll^f , and any query made to Coll^f consists of oracle-aided circuit C that makes at most q queries to f , on any specific input.

Theorem 5.13. *Let $q \leq O(2^{n/6})$. Then for any (q, q) -query adversary A ,*

$$\Pr_{f, \text{Coll}, x} \left[A^{f, \text{Coll}^f}(f(x)) = x \right] \leq O(2^{-n/6}) .$$

Proof. We, in fact, prove a stronger statement: the above holds when fixing the oracles $f_{-n} := \{f_k\}_{k \neq n}$. Let $\varepsilon = 2^{-n/3}$ and $\beta = 2^{n/3} \cdot n$. Fix a q -query adversary A and let S be its smooth $(q + \beta q^2 + 2q^2, q)$ query simulator given by Lemma 4.10. The extra $2q^2$ queries are incurred by the fact that along with each collision w, w' from $\text{Coll}^f(C)$, the queries made to f in computing $C^f(w)$ and $C^f(w')$ are also returned. Since S perfectly emulates A , it is enough to bound the probability that S successfully inverts. To bound S 's inversion probability, we consider six hybrid experiments $\{\mathbf{H}_i\}_{i \in [6]}$ given in Table 1. Throughout, for a permutation $f \in \mathcal{P}_n$ and $x, y \in \{0, 1\}^n$, we denote by $f_{x \rightarrow y}$ the function that maps x to y and is identical to f on all other inputs (in particular, $f_{x \rightarrow y}$ is no longer a permutation when $x \neq f^{-1}(y)$).

Hybrid	\mathbf{H}_1 (Real)	\mathbf{H}_2	\mathbf{H}_3	\mathbf{H}_4	\mathbf{H}_5	\mathbf{H}_6 (Ideal)
Permutation	$f_n \leftarrow \mathcal{P}_n$					
Preimage	$x \leftarrow \{0, 1\}^n$					
2nd Preimage	$z \leftarrow \{0, 1\}^n$					
Planted Image	$y \leftarrow \{0, 1\}^n$					
Challenge	$f(x)$			y		
Oracle	f, Coll^f	$f, \text{PColl}_{\{x\}}^f$	$f_{z \rightarrow f(x)}, \text{PColl}_{\{x, z\}}^f$	$f_{x \rightarrow y}, \text{PColl}_{\{f^{-1}(y), x\}}^f$	$f, \text{PColl}_{\{f^{-1}(y)\}}^f$	f, Coll^f
Winning Condition	Find x					

Table 1: The hybrid experiments.

Hybrid \mathbf{H}_1 is identical to the real world where S wins if it successfully inverts the permutation at a random output. We show that the probability that the adversary wins in any of the experiments is roughly the same, and that in hybrid \mathbf{H}_6 the probability that S wins is tiny.

Claim 5.13.1. $|\Pr[S \text{ wins in } \mathbf{H}_1] - \Pr[S \text{ wins in } \mathbf{H}_2]| \leq O(2^{-n/6})$

Proof. The difference between the two hybrids is in the collision finding oracle: in \mathbf{H}_1 , S gets the standard Coll^f oracle, while in \mathbf{H}_2 , punctured oracle $\text{PColl}_{\{x\}}^f$, punctured at x . Note that by coupling the two experiments, we can bound the statistical distance (and hence the winning probabilities) in \mathbf{H}_1 and \mathbf{H}_2 as follows:

$$\left| \Pr[S \text{ wins in } \mathbf{H}_1] - \Pr[S \text{ wins in } \mathbf{H}_2] \right| \leq \Pr_{f, x, z}^{\text{Coll}} \left[S^{f, \text{Coll}^f}(f(x)) \neq S^{f, \text{PColl}_{\{x\}}^f}(f(x)) \right]$$

Let $\text{smooth} = \text{smooth}(S(f(x)), f, \varepsilon)$ be the event that all Coll -queries made by $S^{f, \text{Coll}^f}(f(x))$ are (f, ε) -smooth (Definition 4.8). And let $\text{collHit} = \text{collHit}(S, f, x, z)$ denote the event that the collision finder oracle Coll^f for some query C returns an answer (w, w') such that $C^f(w)$ or $C^f(w')$ queries x during the evaluation. Note that collHit does not occur when f is queried at x by S , but only when its indirectly queried by Coll^f .

Observe that by Lemma 5.9, as long as punctured set $\{x\}$ is not queried by a collision returned, that is as long as collHit event does not occur, the two oracles Coll^f and $\text{PColl}_{\{x\}}^f$ would return identical answers. Hence,

$$\Pr_{f,x,z}^{\text{Coll}} \left[S^{f, \text{Coll}^f}(f(x)) \neq S^{f, \text{PColl}_{\{x\}}^f}(f(x)) \right] \leq \Pr_{f,x,z}^{\text{Coll}} [\text{collHit}]$$

We bound the probability of collHit as:

$$\Pr[\text{collHit}] \leq \Pr[\overline{\text{smooth}}] + \Pr[\text{smooth} \wedge \text{collHit}]$$

By the smoothness lemma, 5.10,

$$\Pr[\overline{\text{smooth}}] \leq 2^{-\varepsilon\beta + \log(2q^2/\varepsilon)},$$

and, when smooth holds, we can bound the probability of a collHit .

$$\Pr[\text{smooth} \wedge \text{collHit}] \leq 2q\varepsilon$$

This follows from the fact that for any (f, ε) -smooth circuit C , and any x , the following holds:

$$\Pr_r [C^f(r) \text{ queries } x] \leq \varepsilon$$

Hence, as the marginal of each coordinate of a collision returned by the Coll oracle is uniformly random, by a union bound, the probability of collHit occurring for this particular Coll query C is at most $2 \cdot \varepsilon$. Hence the total probability is bounded by $q \cdot (2\varepsilon)$ as desired.

Hence, we can bound the difference between \mathbf{H}_1 and \mathbf{H}_2 by

$$2^{-\varepsilon\beta + \log(2q^2/\varepsilon)} + 2q\varepsilon \leq O(2^{-n/6})$$

when setting $\varepsilon = 2^{-n/3}$, $\beta = 2^{n/3} \cdot n$ and recalling that $q \leq O(2^{n/6})$. \square

Claim 5.13.2. $|\Pr[S \text{ wins in } \mathbf{H}_2] - \Pr[S \text{ wins in } \mathbf{H}_3]| \leq O(2^{-n/6})$.

Proof. The difference between the two hybrids is that in \mathbf{H}_2 , S receives the normal f oracle, while in \mathbf{H}_3 , it receives the planted oracle $f_{z \rightarrow f(x)}$. And it receives $\text{PColl}_{\{x\}}^f$ in \mathbf{H}_2 while receiving $\text{PColl}_{\{x,z\}}^f$ in \mathbf{H}_3 . In what follows, we denote by $\text{zHit} = \text{zHit}(S, f, x, z)$ the event that $S^{f, \text{PColl}_{\{x\}}^f}(f(x))$ queries f on z , either directly or indirectly through a collision returned.

Consider the execution of $S^{f, \text{PColl}_{\{x\}}^f}$ in \mathbf{H}_2 , every query S makes to the oracle is answered identically in \mathbf{H}_3 , unless the event zHit occurs. This follows because the f oracle itself differs only at z in the two hybrids, and the PColl oracle returns the same value by Lemma 5.9 unless zHit occurs. Hence, as S receives the same answers and hence asks the same questions in both hybrids, it would

have the same output, unless zHit occurs. As z is picked uniformly at random, independent of everything else in \mathbf{H}_2 ,

$$\Pr[\text{zHit}] \leq 2^{-n} \cdot |\text{total } f\text{-queries made by S}| \leq 2^{-n} \cdot (q + \beta q^2 + 2q^2) \leq O(2^{-n/6})$$

when setting $\varepsilon = 2^{-n/3}$, $\beta = 2^{n/3} \cdot n$ and recalling that $q \leq O(2^{n/6})$. \square

Claim 5.13.3. $\Pr[\text{S wins in } \mathbf{H}_3] = \Pr[\text{S wins in } \mathbf{H}_4]$.

Proof. First, by symmetry, observe that in \mathbf{H}_3 , the probability of S outputting x is the same as that of S outputting z , because they are completely symmetrical in this hybrid. Then observe that these two hybrids \mathbf{H}_3 and \mathbf{H}_4 are relabellings of each other: $z \leftrightarrow x$, $f(x) \leftrightarrow y$ and $x \leftrightarrow f^{-1}(y)$. This implies that the probability of the probability of S outputting z in \mathbf{H}_3 is the same as that of S outputting x in \mathbf{H}_4 . This completes the argument. \square

Claim 5.13.4. $|\Pr[\text{S wins in } \mathbf{H}_4] - \Pr[\text{S wins in } \mathbf{H}_5]| \leq O(2^{-n/6})$.

The difference between the two hybrids is two fold: the f and PColl oracles differs at x and are identical otherwise. Note that x is independent of the adversary's view in \mathbf{H}_5 . The proof of this claim is identical to that of Claim 5.13.2 and is omitted.

Claim 5.13.5. $|\Pr[\text{S wins in } \mathbf{H}_5] - \Pr[\text{S wins in } \mathbf{H}_6]| \leq O(2^{-n/6})$.

The only difference between the two hybrids is that the Coll oracle from \mathbf{H}_6 is punctured at $f^{-1}(y)$ in \mathbf{H}_5 . The proof of this claim is identical to that of Claim 5.13.1, relies on smoothness, and is omitted.

To conclude the proof of Theorem 5.13, we observe that

Claim 5.13.6. $\Pr[\text{S wins in } \mathbf{H}_6] \leq 2^{-n}$.

Proof. The view of S in this hybrid is completely independent of the random choice of x . \square

This completes the proof of Theorem 5.13. \square

5.5 Indistinguishability Obfuscation (and OWPs) in the Presence of Coll

In this section, we prove the main theorem of this section:

Theorem 5.14. *Let $(\text{Gen}, \text{Eval}, \mathbf{R}, q_{\mathbf{R}}, \varepsilon_{\mathbf{R}})$ be a fully black-box construction of CRHFs from IO and OWPs. Then, one of the following conditions holds:*

1. (Large Running Time) \mathbf{R} makes at least $q_{\mathbf{R}}(n) \geq 2^{n/12}$ queries. Or,
2. (Large Security Loss) $\varepsilon_{\mathbf{R}}(n) \leq 2^{-n/6}$.

To prove the theorem, we consider an oracle Ψ that realizes both indistinguishability obfuscation (IO) and one-way permutations (OWPs) and show that neither break in the presence of the collision finding oracle Coll^{Ψ} . We start by defining the oracle Ψ . In a nutshell, the oracle realizes OWPs through a random permutation oracle.

In what follows, \mathcal{P}_n denotes the set of permutations of $\{0, 1\}^n$, \mathcal{F}_n^m denotes the set of functions mapping $\{0, 1\}^n$ to $\{0, 1\}^m$, and \mathcal{I}_n^m denotes the set of injective functions mapping $\{0, 1\}^n$ to $\{0, 1\}^m$.

Definition 5.15 (The Oracle Ψ). The oracle $\Psi = (f, O, \text{Eval}^{f,O})$ consists of three parts:

- $f = \{f_n\}_{n \in \mathbb{N}}$ on input $x \in \{0, 1\}^n$ answers with $f_n(x)$, where f_n is a random permutation $f_n \leftarrow \mathcal{P}_n$.
- $O = \{O_n\}_{n \in \mathbb{N}}$ on input $(C, r) \in \{0, 1\}^n \times \{0, 1\}^n$ answers with $\widehat{C} := O_n(C, r)$ where O_n is a random injective function $O_n \leftarrow \mathcal{I}_{2n}^{5n}$ into $\{0, 1\}^{5n}$.
- $\text{Eval}^{f,O}$ given $\widehat{C} \in \{0, 1\}^{5n}$, $x \in \{0, 1\}^*$ computes $(C, r) = O_n^{-1}(\widehat{C})$, interprets C as an oracle-aided circuit, and returns $C^f(x)$. If \widehat{C} does not have a unique preimage, or the input size of C is inconsistent with $|x|$, the oracle returns \perp .

In the next two subsections, we show that the oracle Ψ securely realizes OWPs and IO in the presence of the Simon's collision finding oracle Coll^Ψ . Throughout, we address adversaries with oracles $\Psi = (f, O, \text{Eval}^{f,O})$ and Coll^Ψ . We will say that such an adversary is (q', q) -query if they make only:

1. q' queries to f, O, Eval in total, where each circuit query C to O makes at most q queries to f , and for each query \widehat{C} to Eval and any circuit C such that $\widehat{C} \in O(C, *)$, C makes at most q queries to f .
2. q queries to Coll^Ψ , and any circuit query D to Coll^Ψ , makes at most q -queries to Ψ with each O or Eval query being a circuit that itself makes at most q queries to f on any input.

5.5.1 One-Way Permutations

We show that f cannot be inverted, except with exponentially small probability even given an exponential number of oracle queries to $\Psi = (f, O, \text{Eval}^{f,O})$ and Coll^Ψ .

Theorem 5.16. Let $q(n) \leq O(2^{n/12})$. Then for any (q, q) query adversary A

$$\Pr_{\substack{\Psi=(f,O,\text{Eval}) \\ \text{Coll},x}} \left[A^{\Psi, \text{Coll}^\Psi}(f(x)) = x \right] \leq O(2^{-n/6}) .$$

Proof. We will, in fact, prove a stronger statement: the above holds when fixing the oracles $f_{-n} := \{f_k\}_{k \neq n}$, $O = \{O_n\}_{n \in \mathbb{N}}$. We prove the theorem by a reduction to the case that Ψ only consists of the permutation f (and does not include O, Eval). Concretely, fix any (q, q) -query adversary A that inverts the random permutation f_n given access to $\Psi = (f, O, \text{Eval})$ and Coll^Ψ , we show how to reduce it to a $(O(q^2), O(q^2))$ -query adversary $B^{f, \text{Coll}^f}(f_n(x))$ that inverts f_n for a random $x \leftarrow \{0, 1\}^n$ with the same probability as A , given oracle access to (f, Coll^f) . The proof then follows from Theorem 5.13. The new adversary $B^{f, \text{Coll}^f}(f_n(x))$ emulates $A^{\Psi, \text{Coll}^\Psi}(f_n(x))$ answering Ψ -queries as follows:

- **f queries:** answered according to B 's oracle f . This translates to at most q queries to f .
- **O queries:** answered according to the fixed oracle O . This does not add any calls to f .
- **$\text{Eval}^{f,O}$ queries:** given query (\widehat{C}, x) to Eval , invert the fixed oracle O to find $(C, r) = O^{-1}(\widehat{C})$. If no such preimage exists, return \perp . If a preimage does exist, using the f -oracle, compute $C^f(x)$ and return the result. This translates to at most q^2 queries to f : q queries by C , for each of the q queries \widehat{C} to Eval .

- **Coll^Ψ queries:** given query (C_0, C_1) , where C_b makes Ψ -queries translate to D_0, D_1 that only make f -queries, where each query to $\Psi = (f, O, \text{Eval})$ is translated to a query to f according to the previous three items. The resulting oracle-aided (D_0, D_1) may thus make up to $q + q^2$ queries f : q corresponding to the first item, and q^2 corresponding to the third.⁵

Overall B^f is $(O(q^2), O(q^2))$ -query and perfectly emulates the view of A^Ψ . The theorem now follows from Theorem 5.13. \square

5.5.2 Indistinguishability Obfuscation

We now turn to show that Ψ also realizes an indistinguishability obfuscator that does not break in the presence of Coll^Ψ .

Construction 5.17 (The Obfuscator O^Ψ). Let $\Psi = (f, O, \text{Eval}^{f,O})$. Given an oracle-aided circuit $C \in \{0, 1\}^n$, $O^\Psi(1^n, C)$ samples a random $r \leftarrow \{0, 1\}^n$, computes $\widehat{C} = O(C, r)$, and returns an oracle aided circuit $E_{\widehat{C}}$ that given input x , computes $\text{Eval}^{f,O}(\widehat{C}, x)$.

It is easy to see that $O^{f,O,\text{Eval}}$ satisfies the functionality requirement of Definition 5.2 for the class \mathcal{C} of f -aided circuits; indeed, this follows by the fact that O is injective, and by the definition of O and the oracles O, Eval . We now show that it also satisfies indistinguishability, with an exponentially small distinguishing gap, even given an exponential number of oracle queries to $\Psi = (f, O, \text{Eval}^{f,O})$ and the statistical difference oracle Coll^Ψ .

Theorem 5.18. Let $q(n) \leq O(2^{n/6})$. Then for any (q, q) -query adversary A

$$\left| \Pr \left[\text{Expt}_{\Psi, \text{Coll}, O, \mathcal{C}, A}^{\text{IO}}(n) = 1 \right] - \frac{1}{2} \right| \leq O(2^{-n/6})$$

where the random variable $\text{Expt}_{T, O, \mathcal{C}, A}^{\text{IO}}(n)$ denotes the adversary's winning probability in the IO security game (Definition 5.2) relative to $\Psi = (f, O, \text{Eval}^{f,O})$ and Coll^Ψ .

At a very high-level, the proof of the theorem follows a similar rationale to the proof of Theorem 5.13 showing that one-way permutations do not break in the presence of the collision finding oracle. Roughly speaking, we show that in order to break the above construction of IO, the adversary must be able to detect local changes in the oracles realizing it, whereas the collision finding oracle is insensitive of these changes. At a technical level, the case of IO requires somewhat more care than the case of one-way permutations. For once, it has a more elaborate interface consisting not only of a hard to invert mapping O , but also of the evaluation oracle $\text{Eval}^{f,O}$. In particular, a single change to O may introduce many changes to $\text{Eval}^{f,O}$, which could potentially be detected by the statistical difference oracle. Another aspect that complicates the proof is that the IO game is more interactive in its nature. In particular, we need to deal with the fact that the actual circuits of the IO challenge are chosen adaptively, after the adversary had already interacted with all the oracles. We now turn to the actual proof.

⁵We note that while there is a bound on the number of queries that they make, we do not put any restrictions on their size, which allows to hardwire the fixed O and f_{-n} as required in the previous three items. Indeed, Theorem 5.13 does not put any restriction on the size of these circuits.

Proof. We prove a stronger statement: the above holds when fixing the oracles f and $O_{-n} = \{O_k\}_{k \neq n}$. We begin by refining the notion of smoothening for the Eval oracle.

Smoothening and Puncturing for Eval Oracle. As the Eval oracle takes as input a pair (\widehat{C}, z) , where \widehat{C} is an obfuscation and z is an input to evaluate the obfuscated circuit on, we need to refine the notion of smoothness and puncturing for the Eval oracle. For both smoothening and puncturing, we want the definition to consider the obfuscation \widehat{C} while disregarding the specific input it is queried on.

Definition 5.19. Let $\Psi = (f, O, \text{Eval}^{f,O})$. A query D^Ψ to oracle Coll^Ψ is (O, ε) -smooth if: ⁶

1. $\max_{(C,r) \in \{0,1\}^{2n}} \Pr_w [D^\Psi \text{ queries } O \text{ at } (C, r)] \leq \varepsilon.$
2. $\max_{\widehat{C} \in \{0,1\}^{5n}} \Pr_w [D^\Psi \text{ queries } \text{Eval}^{f,O} \text{ at } (\widehat{C}, z) \text{ for some } z] \leq \varepsilon.$

The Smoothening lemma (Lemma 5.10) applies to this definition as well with an essentially identical proof, which is omitted.

Fix a (q, q) -query adversary $A = (A_1, A_2)$ ⁷ and let $S = (S_1, S_2)$ be its smooth $(q + \beta q^2 + 2q^2, q)$ -query simulator given by Lemma 5.10, where the additional $2q^2$ queries are induced by the Coll oracle (see Definition 5.6). That is, the algorithm makes at most $q + 2\beta q^2 + 2q^2$ queries in total to the O and Eval oracles and every query $C^{(\cdot)}$ it makes to $\text{Coll}^{f,O,\text{Eval}}$ in turn makes at most q queries to the O or Eval oracle.

We will specify β later. Since S perfectly emulates A , it suffices to prove the theorem for S .

To bound S 's advantage in breaking O , we consider six hybrid experiments $\{\mathbf{H}_i\}_{i \in [6]}$ given in Table 2. We introduce some notation that will be useful to describe the hybrids:

- We use regular expressions to describe sets. In particular, the $*$ expression. We denote by $(*, r) = \{(C, r) : C \in \{0,1\}^* \text{ and } |C| = |r|\}$ and $O(*, r) = \{O(C, r) : C \in \{0,1\}^* \text{ and } |C| = |r|\}$.
- For a function $O = \{O_k : \{0,1\}^{2k} \rightarrow \{0,1\}^{5k}\}_{k \in \mathbb{N}}$, a pair $(C, r) \in \{0,1\}^{n \times 2}$, and $\widehat{C} \in \{0,1\}^{5n}$, we denote by $O_{(C,r) \rightarrow \widehat{C}}$ the function that maps (C, r) to \widehat{C} and is otherwise identical to O . If $\widehat{C} = \perp$, then $\text{Eval}^{f,O}(\widehat{C}, z)$ returns \perp for all inputs z . This operation naturally extends to sets.
- For a function $O = \{O_k : \{0,1\}^{2k} \rightarrow \{0,1\}^{5k}\}_{k \in \mathbb{N}}$, we denote by $\Gamma(f, O)$ the oracle

$$\Gamma(f, O) := f, O, \text{Eval}^{f,O}, \text{Coll}^{f,O,\text{Eval}^{f,O}},$$

and by $\Gamma_T(f, O)$ the oracle

$$\Gamma_T(f, O) = f, O, \text{Eval}^{f,O}, \text{PColl}_T^{f,O,\text{Eval}^{f,O}},$$

where PColl is the punctured oracle given by Definition 5.8.

⁶Note that the f oracle is fixed. We do not consider it here because it is completely available to the adversary.

⁷In this section, we completely ignore all the queries to f . The adversary has complete access to f and can hard code it in the circuits. We do not limit the size of the circuits, but the number of queries they make.

- For a function $O = \{O_k : \{0, 1\}^{2k} \rightarrow \{0, 1\}^{5k}\}_{k \in \mathbb{N}}$, a string $\widehat{C} \in \{0, 1\}^{5n}$, and a circuit C , we denote by $\Gamma(f, O, \widehat{C}, C)$ the oracle

$$\Gamma(f, O, \widehat{C}, C) := f, O, \text{Eval}_{\widehat{C}, C}^{f, O}, \text{Coll}^{f, O, \text{Eval}_{\widehat{C}, C}^{f, O}},$$

where $\text{Eval}_{\widehat{C}, C}^{f, O}$ is an oracle that

- Given (\widehat{D}, x) where $\widehat{D} \neq \widehat{C}$, acts like $\text{Eval}^{f, O}(\widehat{D}, x)$. Namely, it computes $(D, r) = O^{-1}(\widehat{D})$, and returns $D(x)$, or \perp in case there is no unique preimage or the size of x does not match the input size of D .
- Given (\widehat{C}, x) returns $C(x)$, or \perp in case $C = \perp$, or the size of x does not match the input size of C .

(In both oracles $\Gamma(f, O)$ and $\Gamma(f, O, \widehat{C}, C)$, the fixed O_{-n}, f are also a part of Γ and the oracles of Eval and Coll. These are omitted from our notation for the sake of simplicity.)

Hybrid	\mathbf{H}_1 (Real)	\mathbf{H}_2	\mathbf{H}_3	\mathbf{H}_4	\mathbf{H}_5	\mathbf{H}_6 (Ideal)
Obfuscator Function	$O_n \leftarrow \mathcal{F}_{2n}^{5n}$		$O_n \leftarrow \mathcal{F}_{2n}^{5n}$			$O_n \leftarrow \mathcal{F}_{2n}^{5n}$
Challenger Randomness	$b \leftarrow \{0, 1\}, r \leftarrow \{0, 1\}^n$					
Planted Obfuscation	$\widehat{C} \leftarrow \{0, 1\}^{5n}$					
Prechallenge Punctured Set T	$(*, r) \cup O(*, r)$					
Prechallenge Oracle	$\Gamma(f, O)$	$\Gamma_T(f, O_{(*, r) \rightarrow \perp})$				$\Gamma(f, O)$
Challenge Obfuscation	$O(C_b, r)$			\widehat{C}		
Postchallenge Punctured Set U		$\{(C_b, r)\} \cup (O(C_b, r), *)$		$\{(C_b, r)\} \cup (\widehat{C}, *)$		
Postchallenge Oracle	$\Gamma(f, O)$	$\Gamma_U(f, O)$	$\Gamma_U(f, O_{(C_b, r) \rightarrow \widehat{C}})$		$\Gamma(f, O, \widehat{C}, C_0)$	
Winning Condition	Guess b					

Table 2: The hybrid experiments.

Hybrid \mathbf{H}_1 is identical to the real world where S wins if it produces functionally equivalent C_0, C_1 , and it successfully guesses the bit b . We show that the probability that the simulator wins in any of the experiments is roughly the same, and that in hybrid \mathbf{H}_6 the probability that S wins is $1/2$. Throughout the hybrids, we denote by (C_0, C_1) the circuits output by S_1 in the prechallenge phase. These are w.l.o.g included in the state that S_1 outputs for the postchallenge phase.

Claim 5.19.1. $|\Pr[S \text{ wins in } \mathbf{H}_1] - \Pr[S \text{ wins in } \mathbf{H}_2]| \leq O(2^{-n/6})$

Proof. There are two differences between the two hybrids. The first is in the oracle that S_1 is given in the prechallenge phase: $\Gamma(f, O)$ in \mathbf{H}_1 , and its tweaked version $\Gamma_T(f, O_{(*,r) \rightarrow \perp})$, in \mathbf{H}_2 , where $T = (*, r) \cup O(*, r)$. The second difference between the two hybrids is in the postchallenge oracle that S_2 is given: the normal Coll oracle in \mathbf{H}_1 and the punctured PColl $_U$ in \mathbf{H}_2 where $U = \{(C_b, r)\} \cup (O(C_b, r), *)$.

We can bound the difference between the winning probabilities in \mathbf{H}_1 and \mathbf{H}_2 as follows:

$$\begin{aligned} & \left| \Pr[\text{S wins in } \mathbf{H}_1] - \Pr[\text{S wins in } \mathbf{H}_2] \right| \\ & \leq \Pr_{\substack{S_1, O \\ b, r, \Gamma}} \left[S_1^{\Gamma(f, O)}(1^n) \neq S_1^{\Gamma_T(f, O_{(*,r) \rightarrow \perp})}(1^n) \right] \\ & + \Pr_{\substack{S_2, O \\ b, r, \Gamma}} \left[S_2^{\Gamma(f, O)}(\text{state}, O(C_b, r)) \neq S_2^{\Gamma_U(f, O)}(\text{state}, O(C_b, r)) \mid \text{state} = S_1^{\Gamma(f, O)}(1^n) \right], \end{aligned}$$

We will, in fact, show that the above is bounded for any fixed $b \in \{0, 1\}$. Indeed, for the rest of the claim, fix $b \in \{0, 1\}$.

In what follows, we denote by $\text{hit} = \text{hit}(O, S_1, r)$ the event that O is queried at (C, r) for any $C \in \{0, 1\}^n$ or Eval is queried at $(O(C, r), z)$, for any C, z . The event hit occurs when the queries are made either directly by S_1 or induced through one of the Coll queries: that is, by $D^{f, O, \text{Eval}}(w)$ or $D^{f, O, \text{Eval}}(w')$ where D is a query to Coll with answer (w, w') .

Lemma 5.9 implies that, as long as the hit event does not occur, the two hybrids have identical results in the prechallenge phase. Hence,

$$\Pr_{\substack{S_1, O \\ r, b, \Gamma}} \left[S_1^{\Gamma(f, O)}(1^n) \neq S_1^{\Gamma_T(f, O_{(C_b, r) \rightarrow \perp})}(1^n) \right] \leq \Pr_r[\text{hit}] \leq \frac{(\beta + 2)q^2 + q}{2^n} \leq O(2^{-n/6}).$$

Above we use the fact that r is independent of S_1 's execution in \mathbf{H}_1 and is drawn at random from a set of size 2^n . For any query (C', r') to O , the probability of $r' = r$, that is a hit, is at most 2^{-n} . Similarly, using the fact that O is injective, for any Eval query (\widehat{C}, z) , the probability that $\widehat{C} \in O(*, r)$ is at most 2^{-n} . We then use a union, considering all queries to O or Eval—at most q direct queries, βq^2 queries due to smoothening, and $2q^2$ induced by answers from Coll.

In what follows, let $\text{collHit} = \text{collHit}(S_2, O, U)$ denote the event that one of the queries induces by the collision finder oracle Coll $^{f, O, \text{Eval}}$ hits U . That is, some query Coll $^{f, O, \text{Eval}}(D)$ returns an answer (w, w') such that $D^{f, O, \text{Eval}}(w)$ or $D^{f, O, \text{Eval}}(w')$ queries the oracle at a point inside set $U = \{(C_b, r), (O(C_b, r), *)\}$. Note that, a collHit event does not occur when S directly queries O, Eval inside U , but only via an induced Coll query.

Then, we claim that,

Claim 5.19.2.

$$\Pr_{\substack{S_2, O \\ r, \Gamma}} \left[S_2^{\Gamma(f, O)}(\text{state}, O(C_b, r)) \neq S_2^{\Gamma_U(f, O)}(\text{state}, O(C_b, r)) \mid \text{state} = S_1^{\Gamma(f, O)}(1^n) \right] \leq \Pr[\text{collHit}]$$

Proof. We argue that whenever the complement $\overline{\text{collHit}}$ occurs,

$$S_2^{\Gamma(f, O)}(\text{state}, O(C_b, r)) = S_2^{\Gamma_U(f, O)}(\text{state}, O(C_b, r))$$

The two oracles in the two executions differ only on the Coll oracle where one of them gets Coll oracle while the other gets the oracle punctured at PColl $_U$. By Lemma 5.9, as long as punctured set U is not queried by a collision returned, that is as long as collHit event does not occur, the two oracles Coll and PColl $_U$ return identical answers. The claim follows. \square

We bound the probability of collHit. Let smooth = smooth(S_2, f, O, ε) be the event that all Coll-queries made by $S_2^{\Gamma(f,O)}$ are (O, ε) -smooth (Definition 5.19). Then

$$\Pr[\text{collHit}] \leq \Pr[\overline{\text{smooth}}] + \Pr[\text{smooth} \wedge \text{collHit}]$$

By the smoothness lemma, 5.10,

$$\Pr[\overline{\text{smooth}}] \leq 2^{-\varepsilon\beta + \log(2q^2/\varepsilon)}.$$

To bound $\Pr[\text{smooth} \wedge \text{collHit}]$, observe that,

$$\begin{aligned} & \Pr[\text{smooth} \wedge \text{collHit}] \\ &= \Pr[\text{smooth} \wedge (D^\Psi(w) \text{ or } D^\Psi(w') \text{ queries } O(C_b, r) \text{ for some query } D)] \\ &+ \Pr[\text{smooth} \wedge (D^\Psi(w) \text{ or } D^\Psi(w') \text{ queries } \text{Eval}^{f,O}(O(C_b, r), z) \text{ for some } z, \text{ for some query } D)] \end{aligned}$$

Next, for each (O, ε) -smooth circuit D , the following holds: For a uniformly random w ,

$$\begin{aligned} & \Pr_w [D^{f,O,\text{Eval}}(w) \text{ queries } O \text{ at } (C_b, r)] \leq \varepsilon \\ & \Pr_w [D^{f,O,\text{Eval}}(w) \text{ queries } \text{Eval}^{f,O} \text{ at } (O(C_b, r), z) \text{ for some } z] \leq \varepsilon \end{aligned}$$

Hence, as w, w' answered by Coll each have uniform marginals, by a union bound,

$$\begin{aligned} & \Pr_{w,w' \leftarrow \text{Coll}^\Psi(D)} [D^{f,O,\text{Eval}}(w) \text{ or } D^{f,O,\text{Eval}}(w') \text{ queries } O \text{ at } (C_b, r)] \leq 2\varepsilon \\ & \Pr_{w,w' \leftarrow \text{Coll}^\Psi(D)} [D^{f,O,\text{Eval}}(w) \text{ or } D^{f,O,\text{Eval}}(w') \text{ queries } \text{Eval}^{f,O} \text{ at } (O(C_b, r), z) \text{ for some } z] \leq 2\varepsilon \end{aligned}$$

The bound now follows from a union bound over all the q queries to Coll. Hence,

$$\Pr[\text{smooth} \wedge \text{collHit}] \leq 4q\varepsilon.$$

Overall, we can bound the difference between \mathbf{H}_1 and \mathbf{H}_2 by

$$2^{-n}(q + (\beta + 2) + 2^{-\varepsilon\beta + \log(2q^2/\varepsilon)}q^2) + 4q\varepsilon \leq O(2^{-n/6})$$

when setting $\varepsilon = 2^{-n/3}$, $\beta = 2^{n/3} \cdot n$ and recalling that $q \leq O(2^{n/6})$. \square

Claim 5.19.3. $|\Pr[\text{S wins in } \mathbf{H}_2] - \Pr[\text{S wins in } \mathbf{H}_3]| \leq 2^{-n}$

Proof. The difference between the two hybrids is in the choice of the oracle O : a random injective function in the first, and a random function in the second. Thus,

$$\left| \Pr[\text{S wins in } \mathbf{H}_2] - \Pr[\text{S wins in } \mathbf{H}_3] \right| \leq \Pr_{O \leftarrow \mathcal{F}_{2n}^{5n}} [O \text{ is not injective}] \leq 2^{-n}.$$

\square

Claim 5.19.4. $\Pr[S \text{ wins in } \mathbf{H}_3] = \Pr[S \text{ wins in } \mathbf{H}_4]$.

Proof. The difference between \mathbf{H}_3 and \mathbf{H}_4 is that in \mathbf{H}_4 , in the challenge and post challenge phases, the value $O(C_b, r)$ is re-sampled uniformly at random, namely, it is replaced everywhere by $\widehat{C} \leftarrow \{0, 1\}^{5n}$. We claim that this induces exactly the same distribution on S 's view as in \mathbf{H}_3 . Indeed, in \mathbf{H}_3 and \mathbf{H}_4 , the view of S in pre-challenge phase is completely independent the outputs $O(C, r)$ for any circuit C . □

Claim 5.19.5. $|\Pr[S \text{ wins in } \mathbf{H}_4] - \Pr[S \text{ wins in } \mathbf{H}_5]| \leq 2^{-n}$

Proof. The difference between the two hybrids is in the choice of the oracle O : a random injective function in the first, and a random function in the second. The proof is thus identical to the proof Claim 5.19.3. □

Claim 5.19.6. $|\Pr[S \text{ wins in } \mathbf{H}_5] - \Pr[S \text{ wins in } \mathbf{H}_6]| \leq O(2^{-n/6})$.

Proof. There are two differences between the hybrids. The first is in the oracle that S_1 is given before the challenge phase: $\Gamma(f, O)$ in \mathbf{H}_6 , and its tweaked version $\Gamma_T(f, O_{(*,r) \rightarrow \perp})$ in \mathbf{H}_5 . The second is in the oracle that S_2 is given after the challenge phase: $\Gamma_U(f, O, \widehat{C}, C_0)$ in \mathbf{H}_6 , and $\Gamma(f, O_{(C_b, r) \rightarrow \widehat{C}}, \widehat{C}, C)$ in \mathbf{H}_5 . We can thus bound the difference between the winning probabilities in \mathbf{H}_5 and \mathbf{H}_6 as follows:

$$\begin{aligned} & \left| \Pr[S \text{ wins in } \mathbf{H}_6] - \Pr[S \text{ wins in } \mathbf{H}_5] \right| \\ & \leq \Pr_{\substack{S_1, O \\ r, \Gamma}} \left[\text{state} := S_1^{\Gamma(f, O)}(1^n) \neq S_1^{\Gamma_T(f, O_{(*,r) \rightarrow \perp})}(1^n) \right] \\ & + \Pr_{\substack{S_2, O \\ r, \Gamma}} \left[S_2^{\Gamma(f, O, \widehat{C}, C_0)}(\text{state}, \widehat{C}) \neq S_2^{\Gamma_U(f, O_{(C_b, r) \rightarrow \widehat{C}}, \widehat{C}, C)}(\text{state}, \widehat{C}) \mid \text{state} = S_1^{\Gamma(f, O)}(1^n) \right], \end{aligned}$$

where the probabilities are over the coins of $S = (S_1, S_2)$ and Γ (specifically, Coll) and the choice of $r \leftarrow \{0, 1\}^n$, and $O \leftarrow \mathcal{I}_{2n}^{5n}$, and $b \leftarrow \{0, 1\}$.

The proof of this claim is essentially identical to that of Claim 5.13.1. To show equivalence, we start with the execution in \mathbf{H}_6 , which is independent of (b, r) and show that the execution in \mathbf{H}_5 would result in an identical outcome. There are two differences between \mathbf{H}_5 and \mathbf{H}_6 , like Claim 5.13.1. In the prechallenge phase, the two oracles differ at $T = (*, r) \cup O(*, r)$. In the post challenge phase, they differ at $\text{Eval}(\widehat{C})$, in \mathbf{H}_6 , Eval answers according to C_0 , while in \mathbf{H}_5 it does so according to C_b . However, by the functional equivalence of (C_0, C_1) , the two Eval oracles answer identically. They also differ on the set U . As long U is not directly or indirectly queried in \mathbf{H}_6 , the two hybrids return identical outcomes. And the probability of them diverging is low because r is picked independent of everything else in \mathbf{H}_6 .

Claim 5.19.7. $\Pr[S \text{ wins in } \mathbf{H}_6] = \frac{1}{2}$.

The claim follows because in this hybrid, the challenge is completely independent of the expected answer b . □

This complete the proof of Theorem 5.18. □

Our Approach and the Compression Argument Approach. The proof of Asharov and Segev is based on a compression argument developed in [GGKT05, Wee07, HHRS15, AS15]. At a high level, the compression argument shows that a OWP inverter or an IO breaker that succeeds with non-negligible probability, implies a succinct description of the underlying oracle (e.g., concretely of a random function or permutation), which is information theoretically impossible.

Our proof has the form of common indistinguishability based arguments in cryptography. We show, via a sequence of hybrids, that the adversary’s behavior in the real experiment is indistinguishable from that in an ideal experiment where the security is obvious. Proving that the hybrids are indistinguishable is done via coupling arguments that demonstrate insensitivity to local changes. This rather direct approach leads to a relatively short and simple proof.

We note that while the proofs are different there are some analogous aspects to them, such as smoothening, which effectively occurs in both proofs (there by the name of “direct hit elimination”).

References

- [AGGM06] Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz. On Basing One-way Functions on NP-hardness. In *STOC*, 2006.
- [AHI⁺17] Benny Applebaum, Naama Haramaty, Yuval Ishai, Eyal Kushilevitz, and Vinod Vaikuntanathan. Low-Complexity Cryptographic Hash Functions. In *ITCS*, 2017.
- [AR16] Benny Applebaum and Pavel Raykov. On the Relationship Between Statistical Zero-Knowledge and Statistical Randomized Encodings. In *CRYPTO*, 2016.
- [AS15] Gilad Asharov and Gil Segev. Limits on the Power of Indistinguishability Obfuscation and Functional Encryption. In *FOCS*, 2015.
- [AS16] Gilad Asharov and Gil Segev. On constructing one-way permutations from indistinguishability obfuscation. In *TCC*. Springer, 2016.
- [BB15] Andrej Bogdanov and Christina Brzuska. On Basing Size-Verifiable One-Way Functions on NP-Hardness. In *TCC*, 2015.
- [BBF13] Paul Baecher, Christina Brzuska, and Marc Fischlin. Notions of black-box reductions, revisited. In *ASIACRYPT*, 2013.
- [BDRV18] Itay Berman, Akshay Degwekar, Ron D. Rothblum, and Prashant Nalini Vasudevan. Multi-Collision Resistant Hash Functions and Their Applications. In *EUROCRYPT*, 2018.
- [BDRV19] Itay Berman, Akshay Degwekar, Ron D. Rothblum, and Prashant Nalini Vasudevan. Statistical Difference Beyond the Polarizing Regime. In *(ECCC)*, 2019.
- [BDV17] Nir Bitansky, Akshay Degwekar, and Vinod Vaikuntanathan. Structure vs. Hardness through the Obfuscation Lens. In *CRYPTO*, 2017.

- [BEG⁺94] Manuel Blum, William S. Evans, Peter Gemmell, Sampath Kannan, and Moni Naor. Checking the correctness of memories. *Algorithmica*, 12(2/3):225–244, 1994.
- [BGI⁺01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (Im)possibility of Obfuscating Programs. In *CRYPTO*, 2001.
- [BKSY11] Zvika Brakerski, Jonathan Katz, Gil Segev, and Arkady Yerukhimovich. Limits on the Power of Zero-Knowledge Proofs in Cryptographic Constructions. In *TCC*, 2011.
- [BL13] Andrej Bogdanov and Chin Ho Lee. Limits of provable security for homomorphic encryption. In *CRYPTO*, 2013.
- [BLVW19] Zvika Brakerski, Vadim Lyubashevsky, Vinod Vaikuntanathan, and Daniel Wichs. Worst-case hardness for LPN and cryptographic hashing via code smoothing. In *EUROCRYPT*, 2019.
- [BM09] Boaz Barak and Mohammad Mahmoody-Ghidary. Merkle Puzzles Are Optimal - An $O(n^2)$ -Query Attack on Any Key Exchange from a Random Oracle. In *CRYPTO*, 2009.
- [BT03] Andrej Bogdanov and Luca Trevisan. On Worst-Case to Average-Case Reductions for NP Problems. In *FOCS*, 2003.
- [CDGS18] Sandro Coretti, Yevgeniy Dodis, Siyao Guo, and John Steinberger. Random Oracles and Non-Uniformity. In *EUROCRYPT*, 2018.
- [Dam87] Ivan Damgård. Collision free hash functions and public key signature schemes. In *EUROCRYPT*, 1987.
- [DHT12] Yevgeniy Dodis, Iftach Haitner, and Aris Tentes. On the instantiability of hash-and-sign RSA signatures. In *TCC*, 2012.
- [DLMM11] Dana Dachman-Soled, Yehuda Lindell, Mohammad Mahmoody, and Tal Malkin. On the black-box complexity of optimally-fair coin tossing. In *TCC*, 2011.
- [DPP93] Ivan Damgård, Torben P. Pedersen, and Birgit Pfitzmann. On the existence of statistically hiding bit commitment schemes and fail-stop signatures. In *CRYPTO*, 1993.
- [Fis12] Marc Fischlin. Black-box reductions and separations in cryptography. In *AFRICACRYPT*, 2012.
- [GG98] Oded Goldreich and Shafi Goldwasser. On the Possibility of basing Cryptography on the assumption that $P \neq NP$. *IACR Cryptology ePrint Archive*, 1998.
- [GGH96] Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Collision-free hashing from lattice problems. *IACR Cryptology ePrint Archive*, 1996:9, 1996.
- [GGH⁺13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate Indistinguishability Obfuscation and Functional Encryption for all Circuits. In *FOCS*, 2013.

- [GGKT05] Rosario Gennaro, Yael Gertner, Jonathan Katz, and Luca Trevisan. Bounds on the efficiency of generic cryptographic constructions. *SIAM J. Comput.*, 35(1):217–246, 2005.
- [GHMM18] Sanjam Garg, Mohammad Hajiabadi, Mohammad Mahmoody, and Ameer Mohammed. Limits on the power of garbling techniques for public-key encryption. In *CRYPTO*, 2018.
- [GK93] Oded Goldreich and Eyal Kushilevitz. A perfect zero-knowledge proof system for a problem equivalent to the discrete logarithm. *Journal of Cryptology*, 6(2):97–116, 1993.
- [GKLM12] Vipul Goyal, Virendra Kumar, Satyanarayana V. Lokam, and Mohammad Mahmoody. On black-box reductions between predicate encryption schemes. In *TCC*, 2012.
- [GKM⁺00] Yael Gertner, Sampath Kannan, Tal Malkin, Omer Reingold, and Mahesh Viswanathan. The relationship between public key encryption and oblivious transfer. In *FOCS*, 2000.
- [GMM07] Yael Gertner, Tal Malkin, and Steven Myers. Towards a Separation of Semantic and CCA Security for Public Key Encryption. In *TCC*, 2007.
- [GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The Knowledge Complexity of Interactive Proof-Systems (Extended Abstract). In *STOC*, 1985.
- [GMR01] Yael Gertner, Tal Malkin, and Omer Reingold. On the impossibility of basing trapdoor functions on trapdoor predicates. In *FOCS*, 2001.
- [GT00] Rosario Gennaro and Luca Trevisan. Lower bounds on the efficiency of generic cryptographic constructions. In *FOCS*, 2000.
- [HH09] Iftach Haitner and Thomas Holenstein. On the (im)possibility of key dependent encryption. In *TCC*, 2009.
- [HHRS15] Iftach Haitner, Jonathan J Hoch, Omer Reingold, and Gil Segev. Finding collisions in interactive protocols—tight lower bounds on the round and communication complexities of statistically hiding commitments. *SIAM Journal on Computing*, 44(1):193–242, 2015.
- [HL18] Justin Holmgren and Alex Lombardi. Cryptographic Hashing from Strong One-Way Functions (Or: One-Way Product Functions and Their Applications). In *FOCS*, 2018.
- [HM96] Shai Halevi and Silvio Micali. Practical and provably-secure commitment schemes from collision-free hashing. In *CRYPTO*, 1996.
- [HR04] Chun-Yuan Hsiao and Leonid Reyzin. Finding collisions on a public road, or do secure hash functions need secret coins? In *CRYPTO*, 2004.
- [IKO05] Yuval Ishai, Eyal Kushilevitz, and Rafail Ostrovsky. Sufficient conditions for collision-resistant hashing. In *TCC*, 2005.
- [Imp95] Russell Impagliazzo. A Personal View of Average-Case Complexity. In *CCC*, 1995.

- [IR89] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *STOC*, 1989.
- [Kil92] Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *STOC*, 1992.
- [KNY17] Ilan Komargodski, Moni Naor, and Eylon Yogev. White-box vs. black-box complexity of search problems: Ramsey and graph property testing. In *FOCS*, 2017.
- [KNY18] Ilan Komargodski, Moni Naor, and Eylon Yogev. Collision resistant hashing for paranoids: Dealing with multiple collisions. In *EUROCRYPT*, 2018.
- [KSS11] Jeff Kahn, Michael E. Saks, and Clifford D. Smyth. The dual BKR inequality and rudich’s conjecture. *Combinatorics, Probability & Computing*, 20(2):257–266, 2011.
- [KST99] Jeong Han Kim, Daniel R. Simon, and Prasad Tetali. Limits on the efficiency of one-way permutation-based hash functions. In *FOCS*, 1999.
- [KY18] Ilan Komargodski and Eylon Yogev. On distributional collision resistant hashing. In *CRYPTO*, 2018.
- [LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized Compact Knapsacks Are Collision Resistant. In *ICALP*, 2006.
- [LV16] Tianren Liu and Vinod Vaikuntanathan. On Basing Private Information Retrieval on NP-Hardness. In *TCC*, 2016.
- [Mer89] Ralph C. Merkle. A Certified Digital Signature. In *CRYPTO*, 1989.
- [MP91] Nimrod Megiddo and Christos H. Papadimitriou. On total functions, existence theorems and computational complexity. *Theor. Comput. Sci.*, 1991.
- [MV03] Daniele Micciancio and Salil P. Vadhan. Statistical Zero-Knowledge Proofs with Efficient Provers: Lattice Problems and More. In *CRYPTO*, 2003.
- [MX10] Mohammad Mahmoody and David Xiao. On the power of randomized reductions and the checkability of SAT. In *CCC*, 2010.
- [OK91] Wakaha Ogata and Kaoru Kurosawa. On claw free families. In *ASIACRYPT*, 1991.
- [Pas06] Rafael Pass. Parallel Repetition of Zero-Knowledge Proofs and the Possibility of Basing Cryptography on NP-Hardness. In *CCC*, 2006.
- [Pas13] Rafael Pass. Unprovable security of perfect NIZK and non-interactive non-malleable commitments. In *TCC*, pages 334–354, 2013.
- [Per] Personal communication with the authors of [KNY18].
- [PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, 2006.
- [PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In *CRYPTO*, pages 554–571, 2008.

- [RTV04] Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Notions of reducibility between cryptographic primitives. In *TCC*, 2004.
- [Rud88] Steven Rudich. *Limits on the Provable Consequences of One-Way Functions*. PhD thesis, University of California, Berkeley, 1988.
- [Rus95] Alexander Russell. Necessary and sufficient conditions for collision-free hashing. *J. Cryptology*, 8(2):87–100, 1995.
- [Sim98] Daniel R Simon. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In *EUROCRYPT*, 1998.
- [SV03] Amit Sahai and Salil Vadhan. A complete problem for statistical zero knowledge. *Journal of the ACM (JACM)*, 50(2):196–249, 2003.
- [Unr07] Dominique Unruh. Random Oracles and Auxiliary Input. In *CRYPTO*, 2007.
- [Vad99] Salil Pravin Vadhan. *A study of statistical zero-knowledge proofs*. PhD thesis, Massachusetts Institute of Technology, 1999.
- [Wee07] Hoeteck Wee. One-way Permutations, Interactive Hashing and Statistically Hiding Commitments. In *TCC*, 2007.
- [YZW⁺17] Yu Yu, Jiang Zhang, Jian Weng, Chun Guo, and Xiangxue Li. Collision resistant hashing from learning parity with noise. *IACR Cryptology ePrint Archive*, 2017:1260, 2017.

A Proofs from [BDV17]

In this appendix, we include the omitted proofs of lemmas from [BDV17].

Lemma 4.7 ([BDV17](Claim 3.7)). *Fix any Ψ and any oracle-aided adversary A such that $A^{\Psi, \text{SDO}^\Psi}$ makes at most q queries to SDO^Ψ . Then*

$$\Pr_t[\text{farness}(A, \Psi, \varepsilon)] \geq 1 - 6q\varepsilon ,$$

where the probability is over the choice t of random thresholds by SDO .

Proof. This follows from the fact that, for any query (C_0, C_1) to SDO^Ψ with $s = \text{SD}(C_0^\Psi, C_1^\Psi)$, ε -farness does not hold only if the threshold $t^* = t(C_0, C_1)$, chosen at random for this query, happens to be in the interval $(s - \varepsilon, s + \varepsilon)$, which occurs with probability at most $|(s - \varepsilon, s + \varepsilon)| / |(\frac{1}{3}, \frac{2}{3})| = 6\varepsilon$. The lemma then follows by a union bound over at most q queries. \square

Lemma 4.9 ([BDV17](Claim 3.9)). *Let Ψ, Ψ' be oracles that differ on at most c values in the domain. Let C_0 and C_1 be ε -smooth w.r.t Ψ . Let $s = \text{SD}(C_0^\Psi, C_1^\Psi)$ and $s' = \text{SD}(C_0^{\Psi'}, C_1^{\Psi'})$ then $|s - s'| \leq 2c\varepsilon$.*

Proof. For either $b \in \{0, 1\}$,

$$\begin{aligned} & \text{SD}(C_b^\Psi, C_b^{\Psi'}) \leq \\ & \Pr_r[C_b^\Psi(r) \neq C_b^{\Psi'}(r)] \leq \\ & \Pr_r[C_b^\Psi(r) \text{ queries } \Psi \text{ at } x \text{ where } \Psi(x) \neq \Psi'(x)] \leq \\ & \sum_{x: \Psi(x) \neq \Psi'(x)} \Pr_r[C_b^\Psi(r) \text{ queries } \Psi \text{ at } x] \leq c \cdot \varepsilon . \end{aligned}$$

The claim then follows by the fact that

$$|s - s'| := |\text{SD}(C_0^\Psi, C_1^\Psi) - \text{SD}(C_0^{\Psi'}, C_1^{\Psi'})| \leq \text{SD}(C_0^\Psi, C_0^{\Psi'}) + \text{SD}(C_1^\Psi, C_1^{\Psi'}) \leq 2c\varepsilon .$$

□

Lemma 4.10 (Smoothing Lemma for SDO [BDV17](Lemma 3.10)). *For any (q, q) -query algorithm A and $\beta \in \mathbb{N}$, there exists a $(q + 2\beta q^2, q)$ -query algorithm S such that for any input $z \in \{0, 1\}^*$ and oracles Ψ, SDO^Ψ :*

1. $S^{\Psi, \text{SDO}^\Psi}(z)$ perfectly simulates the output of $A^{\Psi, \text{SDO}^\Psi}(z)$,
2. $S^{\Psi, \text{SDO}^\Psi}(z)$ only makes queries (C_0, C_1) where both C_0, C_1 are (Ψ, ε) -smooth queries to SDO^Ψ with probability:

$$\Pr_S[\text{smooth}(S, \Psi, \varepsilon)] \geq 1 - 2^{-\varepsilon\beta + \log(2q^2/\varepsilon)} ,$$

over its own random coin tosses.

Proof. The simulator S emulates A and whenever A makes a query (C_0, C_1) to Ψ , the simulator S first evaluates each of the two circuits C_0^Ψ, C_1^Ψ on β random inputs and stores all the queries they make to Ψ along with their answers in a table T . It then generates a new query consisting of circuits (C'_0, C'_1) that have the table T hardwired in them. Each C'_b emulates C_b , but whenever the emulated C_b makes an oracle query to Ψ , C'_b first tries to answer using the table T , and only if the answer is not there turns to the oracle Ψ .

This process incurs $2\beta q$ queries to Ψ for each $\text{SDO}^\Psi(C_0, C_1)$ query for smoothing. As there are q such queries, we incur at most $2\beta q^2$ additional Ψ -queries.

By construction, S perfectly emulates the view of A . We now bound the probability that S generates a circuit that is not (Ψ, ε) -smooth. Fix any query (C_0, C_1) and let x be a *heavy query* in the sense that it is queried with probability larger than ε by one of the two circuits. Then the query x will be put in the table T except with probability $(1 - \varepsilon)^\beta \leq 2^{-\varepsilon\beta}$. Furthermore, each one of the two circuits makes at most q oracle queries and thus each has at most q/ε inputs x as above. The claim now follows by a union bound over at most q queries (C_0, C_1) and at most q/ε heavy inputs that each of the two has. □