

Distributional Collision Resistance Beyond One-Way Functions

Nir Bitansky^{*} Iftach Haitner[†] Ilan Komargodski[‡] Eylon Yogev[§]

Abstract

Distributional collision resistance is a relaxation of collision resistance that only requires that it is hard to sample a collision (x, y) where x is uniformly random and y is uniformly random conditioned on colliding with x . The notion lies between one-wayness and collision resistance, but its exact power is still not well-understood. On one hand, distributional collision resistant hash functions cannot be built from one-way functions in a black-box way, which may suggest that they are stronger. On the other hand, so far, they have not yielded any applications beyond one-way functions.

Assuming distributional collision resistant hash functions, we construct *constant-round* statistically hiding commitment scheme. Such commitments are not known based on one-way functions and are impossible to obtain from one-way functions in a black-box way. Our construction relies on the reduction from inaccessible entropy generators to statistically hiding commitments by Haitner et al. (STOC '09). In the converse direction, we show that two-message statistically hiding commitments imply distributional collision resistance, thereby establishing a loose equivalence between the two notions.

A corollary of the first result is that constant-round statistically hiding commitments are implied by average-case hardness in the class SZK (which is known to imply distributional collision resistance). This implication seems to be folklore, but to the best of our knowledge has not been proven explicitly. We provide yet another proof of this implication, which is arguably more direct than the one going through distributional collision resistance.

^{*}School of Computer Science, Tel Aviv University. Email: nirbitan@tau.ac.il. Member of the Check Point Institute of Information Security. Supported by ISF grant 18/484, the Alon Young Faculty Fellowship, and by Len Blavatnik and the Blavatnik Family foundation.

[†]School of Computer Science, Tel Aviv University. Email: iftachh@cs.tau.ac.il. Member of the Check Point Institute for Information Security. Research supported by ERC starting grant 638121.

[‡]Cornell Tech, New York, NY. Email: komargodski@cornell.edu. Supported in part by an AFOSR grant FA9550-15-1-0262.

[§]Department of Computer Science, Technion. Email: eylony@gmail.com. Supported by the European Union's Horizon 2020 research and innovation program under grant agreement No. 742754.

1 Introduction

Distributional collision resistant hashing (dCRH), introduced by Dubrov and Ishai [DI06], is a relaxation of the notion of collision resistance. In (plain) collision resistance, it is guaranteed that no efficient adversary can find *any* collision given a random hash function in the family. In dCRH, it is only guaranteed that no efficient adversary can sample *a random* collision given a random hash function in the family. More precisely, given a random hash function h from the family, it is computationally hard to sample a pair (x, y) such that x is uniform and y is uniform in the preimage set $h^{-1}(x) = \{z: h(x) = h(z)\}$. This hardness is captured by requiring that the adversary cannot get statistically-close to this distribution over collisions.¹

The power of dCRH. Intuitively, the notion of dCRH seems quite weak. The adversary may even be able to sample collisions from the set of *all* collisions, but only from a skewed distribution, far from the random one. Komargodski and Yagev [KY18] show that dCRH can be constructed assuming average-case hardness in the complexity class *statistical zero-knowledge* (SZK), whereas a similar implication is not known for multi-collision resistance.² (let alone plain collision resistance). This can be seen as evidence suggesting that dCRH may be weaker than collision resistance, or even multi-collision resistance [KNY17, BDRV18, BKP18, KNY18].

Furthermore, dCRH has not led to the same cryptographic applications as collision resistance, or even multi-collision resistance. In fact, dCRH has no known applications beyond those implied by one-way functions.

At the same time, dCRH is not known to follow from one-way functions, and actually, cannot follow based on black-box reductions [Sim98]. In fact, it can even be separated from indistinguishability obfuscation (and one-way functions) [AS16]. Overall, we are left with a significant gap in our understanding of the power of dCRH:

Does the power of dCRH go beyond one-way functions?

1.1 Our Results

We present the first application of dCRH that is not known from one-way functions and is provably unachievable from one-way functions in a black-box way.

Theorem 1. *dCRH implies constant-round statistically hiding commitment scheme.*

Such commitment schemes cannot be constructed from one-way functions (or even permutations) in a black-box way due to a result of Haitner, Hoch, Reingold and Segev [HHRS15]. They show that the number of rounds in such commitments must grow quasi-linearly in the security parameter.

The heart of Theorem 1 is a construction of an inaccessible-entropy generator [HRVW09, HRVW18] from dCRH.

¹There are some subtleties in defining this precisely. The definition we use differs from previous ones [DI06, HN10, KY18]. We elaborate on the exact definition and the difference in the technical overview below and in Section 3.4.

²Multi-collision resistance is another relaxation of collision resistance, where it is only hard to find multiple elements that all map to the same image. Multi-collision resistance does not imply dCRH in a black-box way [KNY18], but Komargodski and Yagev [KY18] give a non-black-box construction.

An implication of the above result is that constant-round statistically hiding commitments can be constructed from average-case hardness in SZK. Indeed, it is known that such hardness implies the existence of a dCRH [KY18].

Corollary 1. *A Hard-on-average problem in SZK implies a constant-round statistically hiding commitment scheme.*

The statement of Corollary 1 has been treated as known in several previous works (c.f. [HRVW09, DGRV11, BDV17]), but a proof of this statement has so far not been published or (to the best of our knowledge) been publicly available. We also provide an alternative proof of this statement (and in particular, a different commitment scheme) that does not go through a construction of a dCRH, and is arguably more direct.

A limit on the power of dCRH. We also show a converse connection between dCRH and statistically hiding commitments. Specifically, we show that *any* two-message statistically hiding commitment implies a dCRH function family.

Theorem 2. *Any two-message statistically hiding commitment scheme implies dCRH.*

This establishes a loose equivalence between dCRH and statistically hiding commitments. Indeed, the commitments we construct from dCRH require more than two messages. Interestingly, we can even show that such commitments imply a stronger notion of dCRH where the adversary’s output distribution is not only noticeably far from the random collision distribution, but is $(1 - \text{negl}(n))$ -far.

1.2 Related Work on Statistically Hiding Commitments

Commitment schemes, the digital analog of sealed envelopes, are central to cryptography. More precisely, a commitment scheme is a two-stage interactive protocol between a sender S and a receiver R . After the commit stage, S is bound to (at most) one value, which stays hidden from R , and in the reveal stage R learns this value. The immediate question arising is what it means to be “bound to” and to be “hidden”. Each of these security properties can come in two main flavors, either *computational security*, where a polynomial-time adversary cannot violate the property except with negligible probability, or the stronger notion of *statistical security*, where even an unbounded adversary cannot violate the property except with negligible probability. However, it is known that there do *not* exist commitment schemes that are simultaneously statistically hiding and statistically binding.

There exists a one-message (i.e., non-interactive) statistically binding commitment schemes assuming one-way permutations (Blum [Blu81]). From one-way functions, such commitments can be achieved by a two-message protocol (Naor [Nao91] and Håstad, Impagliazzo, Levin and Luby [HILL99]).

Statistically hiding commitments schemes have proven to be somewhat more difficult to construct. Naor, Ostrovsky, Venkatesan and Yung [NOVY92] gave a statistically hiding commitment scheme protocol based on one-way permutations, whose linear number of rounds matched the lower bound of [HHRS15] mentioned above. After many years, this result was improved by Haitner, Nguyen, Ong, Reingold and Vadhan [HNO⁺09] constructing such commitment based on the minimal hardness assumption that one-way functions exist. The reduction of [HNO⁺09] was later

simplified and made more efficient by Haitner, Reingold, Vadhan and Wee [HRVW09, HRVW18] to match, in some settings, the round complexity lower bound of [HHR15]. Constant-round statistically hiding commitment protocols are known to exist based on families of collision resistant hash functions [NY89, DPP93, HM96]. Recently, Berman, Degwekar, Rothblum and Vasudevan [BDRV18] and Komargodski, Naor and Yogev [KNY18] constructed constant-round statistically hiding commitment protocols assuming the existence of *multi-collision* resistant hash functions.

Constant-round statistically hiding commitments are a basic building block in many fundamental applications. Two prominent examples are constructions of *constant-round* zero-knowledge proofs for all NP (Goldreich and Kahan [GK96]) and *constant-round* public-coin statistical zero-knowledge arguments for NP (Barak [Bar01], Pass and Rosen [PR08]).

Statistically hiding commitment are also known to be tightly related to the hardness of the class of problems that possess a statistical zero-knowledge protocol, i.e., the class SZK. Ong and Vadhan [OV08] showed that a language in NP has a zero-knowledge protocol if and only if the language has an “instance-dependent” commitment scheme. An instance-dependent commitment scheme for a given language is a commitment scheme that can depend on an instance of the language, and where the hiding and binding properties are required to hold only on the YES and NO instances of the language, respectively.

1.3 Directions for Future Work

The security notions of variants of collision resistance, including plain collision resistance and multi-collision resistance, can be phrased in the language of entropy. For example, plain collision resistance requires that once a hash value y is fixed the (max) entropy of preimages that any efficient adversary can find is zero. In multi-collision resistance, it may be larger than zero, even for every y , but still bounded by the size of allowed multi collisions. In distributional collision resistance, the (Shannon) entropy is close to maximal.

Yet, the range of applications of collision resistance (or even multi-collision resistance) is significantly larger than those of distributional collision resistance. Perhaps the most basic such application is *succinct* commitment protocols which are known from plain/multi-collision resistance but not from distributional collision resistance (by *succinct* we mean that the total communication is shorter than the string being committed to). Thus, with the above entropy perspective in mind, a natural question is to characterize the full range or parameters between distributional and plain collision resistance and understand for each of them what are the applications implied. A more concrete question is to find the minimal notion of security for collision resistance that implies succinct commitments.

A different line of questions concerns understanding better the notion of distributional collision resistance and constructing it from more assumptions. Komargodski and Yogev constructed it from multi-collision resistance and from the average-case hardness of SZK. Can we construct it, for example, from the multivariate quadratic (MQ) assumption [MI88] or can we show an attack for random degree 2 mappings? Indeed, we know that random degree 2 mappings cannot be used for plain collision resistant hashing [AHI⁺17, Theorem 5.3].

2 Technical Overview

In this section, we give an overview of our techniques. We start with a more precise statement of the definition of dCRH and a comparison with previous versions of its definition.

A dCRH is a family of functions $\mathcal{H}_n = \{h: \{0, 1\}^n \rightarrow \{0, 1\}^m\}$. (The functions are not necessarily compressing.) The security guarantee is that there exists a universal polynomial $p(\cdot)$ such that for every efficient adversary A it holds that

$$\Delta((h, A(1^n, h)), (h, \text{Col}(h))) \geq \frac{1}{p(n)},$$

where Δ denotes statistical distance, $h \leftarrow \mathcal{H}_n$ is chosen uniformly at random, and Col is a random variable that is sampled in the following way: Given h , first sample $x_1 \leftarrow \{0, 1\}^n$ uniformly at random and then sample x_2 uniformly at random from the set of all preimages of x_1 relative to h (namely, from the set $\{x: h(x) = h(x_1)\}$). Note that Col may not be efficiently samplable and intuitively, the hardness of dCRH says that there is no efficient way to sample from Col , even approximately.

Our definition is stronger than previous definitions of dCRH [DI06, HN10, KY18] by that we require the existence of a universal polynomial $p(\cdot)$, whereas previous definitions allow a different polynomial per adversary. Our modification seems necessary to get non-trivial applications of dCRH, as the previous definitions are not known to imply one-way functions. In contrast, our notion of dCRH implies distributional one-way functions which, in turn, imply one-way functions [IL89] (indeed, the definition of distributional one-way functions requires a universal polynomial rather than one per adversary).³ We note that previous constructions of dCRH (from multi-collision resistance and SZK-hardness) [KY18] apply to our stronger notion as well.

2.1 Commitments from dCRH and Back

We now describe our construction of constant-round statistically hiding commitments from dCRH. To understand the difficulty, let us recall the standard approach to constructing statistically hiding commitments from (fully) collision resistant hash functions [NY89, DPP93, HM96]. Here to commit to a bit b , we hash a random string x , and output $(h(x), s, b \oplus \text{Ext}_s(x))$, where s is a seed for a strong randomness extractor Ext and b is padded with a (close to) random bit extracted from x . When h is collision resistant, x is computationally fixed and thus so is the bit b . However, for a dCRH h , this is far from being the case: for any y , the sender might potentially be able to sample preimages from the set of all preimages.

The hash $h(x)$, however, does yield a weak binding guarantee. For simplicity of exposition, let us assume that any $y \in \{0, 1\}^m$ has exactly 2^k preimages under h in $\{0, 1\}^n$. Then, for a noticeable fraction of commitments y , the adversary cannot open y to a uniform x in the preimage set $h^{-1}(y)$. In particular, the adversary must choose between two types of *entropy losses*: it either outputs a commitment y of entropy m' noticeably smaller than m , or after the commitment, it can only open to a value x of entropy k' noticeably smaller than k . One way or the other, in total $m' + k'$ must be noticeably smaller than $n = m + k$. This naturally leads us to the notion of *inaccessible entropy* defined by Haitner, Reingold, Vadhan and Wee [HRVW09, HRVW18].

³The previous definition is known to imply a weaker notion of distributional one-way functions (with a different polynomial bound per each adversary) [HN10], which is not known to imply one-way functions.

Let us briefly recall what inaccessible entropy is (see Section 4.1 for a precise definition). The entropy of a random variable X is a measure of “the amount of randomness” that X contains. The notion of (in)accessible entropy measures the feasibility of sampling high-entropy strings that are *consistent* with a given random process. Consider the two-block generator (algorithm) G that samples $x \leftarrow \{0, 1\}^n$, and then outputs $y = h(x)$ and x . The *real entropy* of G is defined as the entropy of the generator’s (total) output in a random execution, and is clearly equal to n , the length of x . The *accessible entropy* of G measures the entropy of these output blocks from the point of view of an efficient G -consistent generator, which might act arbitrarily, but still outputs a value in the support of G .

Assume for instance that h had been (fully) collision resistant. Then from the point of view of any efficient G -consistent generator \tilde{G} , conditioned on its first block y , and its internal randomness, its second output block is fixed (otherwise, G can be used for finding a collision). In other words, while the value of x given y may have entropy $k = n - m$, this entropy is completely *inaccessible* for an efficient G -consistent generator. (Note that we do not measure here the entropy of the output blocks of \tilde{G} , which clearly can be as high as the real entropy of G by taking $\tilde{G} = G$. Rather, we measure the entropy of the block from \tilde{G} ’s *point of view*, and in particular, the entropy of its second block given the randomness used for generating the first block.). Haitner et al. show that any noticeable gap between the real entropy and the inaccessible entropy of such an efficient generator can be leveraged for constructing statistically hiding commitments, with a number of rounds that is linear in the number of blocks.

Going back to dCRH, we have already argued that in the simple case that h is regular and onto $\{0, 1\}^m$, we get a noticeable gap between the real entropy $n = m + k$ and the accessible entropy $m' + k' \leq m + k - 1/\text{poly}(n)$. We prove that this is, in fact, true for any dCRH:

Lemma 1. *dCRH implies a two-block inaccessible entropy generator.*

The block generator itself is the simple generator described above:

output $h(x)$ and then x , for $x \leftarrow \{0, 1\}^n$.

The proof, however, is more involved than in the case of collision resistance. In particular, it is sensitive to the exact notion of entropy used. Collision resistant hash functions satisfy a very clean and simple guarantee — the *maximum entropy*, capturing the support size, is always at most $m < n$. In contrast, for dCRH (compressing or not), the maximum entropy could be as large as n , which goes back to the fact that the adversary may be able to sample from the set of *all* collisions (albeit from a skewed distribution). Still, we show a gap with respect to average (a.k.a Shannon) accessible entropy, which suffices for constructing statistically hiding commitments [HRVW18].

From commitments back to dCRH. We show that any two-message statistically hiding commitment implies a dCRH function family. Let $(\mathcal{S}, \mathcal{R})$ be the sender and receiver of a statistically hiding bit commitment. The first message sent by the receiver is the description of the hash function: $h \leftarrow \mathcal{R}(1^n)$. The sender’s commitment to a bit b , using randomness r , is the hash of $x = (b, r)$. That is, $h(x) = \mathcal{S}(h, b; r)$.

To argue that this is a dCRH, we show that any attacker that can sample collisions that are close to the random collision distribution Col can also break the binding of the commitment scheme. For this, it suffices to show that a collision $(b, r), (b', r')$ sampled from Col , translates to equivocation — the corresponding commitment can be opened to two distinct bits $b \neq b'$. Roughly speaking, this

is because statistical hiding implies that a random collision to a random bit b (corresponding to a random hash value) is statistically independent of the underlying committed bit. In particular, a random preimage of such a commitment will consist of a different bit b' with probability roughly $1/2$. See details in Section 4.3.

2.2 Commitments from SZK Hardness

We now give an overview of our construction of statistically hiding commitments directly from average-case hardness in SZK. Our starting point is a result of Ong and Vadhan [OV08] showing that any promise problem in SZK has an *instance-dependent commitment*. These are commitments that are also parameterized by an instance x , such that if x is a *yes instance*, they are statistically hiding and if x is a *no instance*, they are statistically binding. We construct statistically hiding commitments from instance-dependent commitments for a hard-on-average problem $\Pi = (\Pi_N, \Pi_Y)$ in SZK.

A first attempt: using zero-knowledge proofs. To convey the basic idea behind the construction, let us first assume that Π satisfies a strong form of average-case hardness where we can efficiently sample no-instances from Π_N and yes-instances from Π_Y so that the two distributions are computationally indistinguishable. Then a natural protocol for committing to a message m is the following: The receiver \mathcal{R} would sample a yes-instance $x \leftarrow \Pi_Y$, and send it to the sender \mathcal{S} along with zero-knowledge proof [GMR89] that x is indeed a yes-instance. The sender \mathcal{S} would then commit to m using an x -dependent commitment.

To see that the scheme is statistically hiding, we rely on the soundness of the proof which guarantees that x is indeed a yes-instance, and then on the hiding of the instance-dependent scheme. To prove (computational) binding, we rely on zero knowledge property and the hardness of Π . Specifically, by zero knowledge, instead of sampling x from Π_Y , we can sample it from any computationally indistinguishable distribution, without changing the probability that an efficient malicious sender breaks binding. In particular, by the assumed hardness of Π , we can sample x from Π_N . Now, however, the instance-dependent commitment guarantees binding, implying that the malicious sender will not be able to equivocate.

The main problem with this construction is that constant-round zero-knowledge proofs (with a negligible soundness error) are only known assuming constant-round statistically hiding commitments [GK96], which is exactly what we are trying to construct.

A second attempt: using witness-indistinguishable proofs. Instead of relying on zero-knowledge proofs, we rely on the weaker notion of witness-indistinguishable proofs and use the *independent-witnesses paradigm* of Feige and Shamir [FS90]. (Indeed such proofs are known for all of NP, based average-case hardness in SZK [GMW87, Nao91, OW93], see Section 5 for details.) We change the previous scheme as follows: the receiver \mathcal{R} will now sample *two* instances x_0 and x_1 and provide a witness-indistinguishable proof that at least one of them is a yes-instance. The sender, will secret share the message m into two random messages m_0, m_1 such that $m = m_0 \oplus m_1$, and return two instance-dependent commitments to m_0 and m_1 relative to x_0 and x_1 , respectively.

Statistical hiding follows quite similarly to the previous protocol — by the soundness of the proof one of the instances x_b is a yes-instance, and by the hiding of the x_b -dependent commitment, the corresponding share m_b is statistically hidden, and thus so is m . To prove binding, we first

note that by witness indistinguishability, to prove its statement, the receiver could use x_b for either $b \in \{0, 1\}$. Then, relying on the hardness of Π , we can sample x_{1-b} to be a no-instance instead of a yes-instance. If b is chosen at random, the sender cannot predict b better than guessing. At the same time, in order to break binding, the sender must equivocate with respect to at least one of the instance-dependent commitments, and since it cannot equivocate with respect to the no-instance x_{1-b} , it cannot break binding unless it can get an advantage in predicting b .

Our actual scheme. The only gap remaining between the scheme just described and our actual scheme is our assumption regarding the strong form of average-case hardness of Π . In contrast, the standard form of average-case hardness only implies a single samplable distribution D , such that given a sample x from D it is hard to tell whether x is a yes-instance or a no-instance better than guessing.

This requires the following changes to the protocol. First, lacking a samplable distribution on yes-instances, we consider instead the product distribution D^n , as a way to sample *weak yes instances* — n -tuples of instances where at least one is a yes-instance in Π_Y . Unlike before, where everything in the support of the yes-instance sampler was guaranteed to be a yes-instance, now we are only guaranteed that a random tuple is a weak yes instance with overwhelming probability. To deal with this weak guarantee, we add a *coin-tossing into the well* phase [GMW87], where the randomness for sampling an instance from D^n is chosen together by the receiver and sender. We refer the reader to Section 5 for more details.

3 Preliminaries

Unless stated otherwise, the logarithms in this paper are base 2. For a distribution \mathcal{D} we denote by $x \leftarrow \mathcal{D}$ an element chosen from \mathcal{D} uniformly at random. For an integer $n \in \mathbb{N}$ we denote by $[n]$ the set $\{1, \dots, n\}$. We denote by U_n the uniform distribution over n -bit strings. We denote by \circ the string concatenation operation. A function $\text{negl}: \mathbb{N} \rightarrow \mathbb{R}^+$ is *negligible* if for every constant $c > 0$, there exists an integer N_c such that $\text{negl}(n) < n^{-c}$ for all $n > N_c$.

3.1 Cryptographic Primitives

A function f , with input length $m_1(n)$ and outputs length $m_2(n)$, specifies for every $n \in \mathbb{N}$ a function $f_n: \{0, 1\}^{m_1(n)} \rightarrow \{0, 1\}^{m_2(n)}$. We only consider functions with polynomial input lengths (in n) and occasionally abuse notation and write $f(x)$ rather than $f_n(x)$ for simplicity. The function f is computable in polynomial time (efficiently computable) if there exists a probabilistic machine that for any $x \in \{0, 1\}^{m_1(n)}$ outputs $f_n(x)$ and runs in time polynomial in n .

A function family ensemble is an infinite set of function families, whose elements (families) are indexed by the set of integers. Let $\mathcal{F} = \{\mathcal{F}_n: \mathcal{D}_n \rightarrow \mathcal{R}_n\}_{n \in \mathbb{N}}$ stand for an ensemble of function families, where each $f \in \mathcal{F}_n$ has domain \mathcal{D}_n and range \mathcal{R}_n . An efficient function family ensemble is one that has an efficient sampling and evaluation algorithms.

Definition 1 (Efficient function family ensemble). *A function family ensemble $\mathcal{F} = \{\mathcal{F}_n: \mathcal{D}_n \rightarrow \mathcal{R}_n\}_{n \in \mathbb{N}}$ is efficient if:*

- \mathcal{F} is samplable in polynomial time: there exists a probabilistic polynomial-time machine that given 1^n , outputs (the description of) a uniform element in \mathcal{F}_n .

- There exists a deterministic algorithm that given $x \in \mathcal{D}_n$ and (a description of) $f \in \mathcal{F}_n$, runs in time $\text{poly}(n, |x|)$ and outputs $f(x)$.

3.2 Distance and Entropy Measures

Definition 2 (Statistical distance). *The statistical distance between two random variables X, Y over a finite domain Ω , is defined by*

$$\Delta(X, Y) \triangleq \frac{1}{2} \cdot \sum_{x \in \Omega} |\Pr[X = x] - \Pr[Y = x]|.$$

We say that X and Y are δ -close (resp. -far) if $\Delta(X, Y) \leq \delta$ (resp. $\Delta(X, Y) \geq \delta$).

Entropy. Let X be a random variable. For any $x \in \text{supp}(X)$, the sample-entropy of x with respect to X is

$$H_X(x) = \log \left(\frac{1}{\Pr[X = x]} \right).$$

The Shannon entropy of X is defined as:

$$H(X) = \mathbf{E}_{x \leftarrow X} [H_X(x)].$$

Conditional entropy. Let (X, Y) be a jointly distributed random variable.

- For any $(x, y) \in \text{supp}(X, Y)$, the conditional sample-entropy to be

$$H_{X|Y}(x | y) = \log \left(\frac{1}{\Pr[X = x | Y = y]} \right).$$

- The conditional Shannon entropy is

$$H(X | Y) = \mathbf{E}_{(x,y) \leftarrow (X,Y)} [H_{X|Y}(x | y)] = \mathbf{E}_{y \leftarrow Y} [H(X|Y=y)] = H(X, Y) - H(Y).$$

Relative entropy. We also use basic facts about relative entropy (also known as , Kullback-Leibler divergence).

Definition 3 (Relative entropy). *Let X and Y be two random variables over a finite domain Ω . The relative entropy is*

$$\mathbf{D}_{\text{KL}}(X \| Y) = \sum_{x \in \Omega} \Pr[X = x] \cdot \log \left(\frac{\Pr[X = x]}{\Pr[Y = x]} \right).$$

Proposition 1 (Chain rule). *Let (X_1, X_2) and (Y_1, Y_2) be random variables. It holds that*

$$\mathbf{D}_{\text{KL}}((X_1, X_2) \| (Y_1, Y_2)) = \mathbf{D}_{\text{KL}}(X_1 \| Y_1) + \mathbf{E}_{x \leftarrow X_1} [\mathbf{D}_{\text{KL}}(X_2 |_{X_1=x} \| Y_2 |_{Y_1=x})].$$

A well-known relation between statistical distance and relative entropy is given by Pinsker's inequality.

Proposition 2 (Pinsker's inequality). *For any two random variables X and Y over a finite domain it holds that*

$$\Delta(X, Y) \leq \sqrt{\frac{\ln 2}{2} \cdot \mathbf{D}_{\text{KL}}(X \| Y)}.$$

Another useful inequality is Jensen's inequality.

Proposition 3 (Jensen's inequality). *If X is a random variable and f is concave, then*

$$\mathbf{E}[f(X)] \leq f(\mathbf{E}[X]).$$

3.3 Commitment Schemes

A commitment scheme is a two-stage interactive protocol between a sender \mathcal{S} and a receiver \mathcal{R} . The goal of such a scheme is that after the first stage of the protocol, called the commit protocol, the sender is bound to at most one value. In the second stage, called the opening protocol, the sender opens its committed value to the receiver. Here, we are interested in statistically hiding and computationally binding commitments. Also, for simplicity, we restrict our attention to protocols that can be used to commit to bits (i.e., strings of length 1).

In more detail, a commitment scheme is defined via a pair of probabilistic polynomial-time algorithms $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ such that:

- The commit protocol: \mathcal{S} receives as input the security parameter 1^n and a bit $b \in \{0, 1\}$. \mathcal{R} receives as input the security parameter 1^n . At the end of this stage, \mathcal{S} outputs `decom` (the decommitment) and \mathcal{R} outputs `com` (the commitment).
- The verification: \mathcal{V} receives as input the security parameter 1^n , a commitment `com`, a decommitment `decom`, and outputs either a bit b or \perp .

A commitment scheme is *public coin* if all messages sent by the receiver are independent random coins.

Denote by $(\text{decom}, \text{com}) \leftarrow \langle \mathcal{S}(1^n, b), \mathcal{R} \rangle$ the experiment in which \mathcal{S} and \mathcal{R} interact with the given inputs and uniformly random coins, and eventually \mathcal{S} outputs a decommitment string and \mathcal{R} outputs a commitment. The completeness of the protocol says that for all $n \in \mathbb{N}$, every $b \in \{0, 1\}$, and every tuple $(\text{decom}, \text{com})$ in the support of $\langle \mathcal{S}(1^n, b), \mathcal{R} \rangle$, it holds that $\mathcal{V}(\text{decom}, \text{com}) = b$. Unless otherwise stated, \mathcal{V} is the canonical verifier that receives the sender's coins as part of the decommitment and checks their consistency with the transcript.

Below we define two security properties one can require from a commitment scheme. The properties we list are *statistical-hiding* and *computational-binding*. These roughly say that after the commit stage, the sender is *bound* to a specific value but the receiver cannot know this value.

Definition 4 (binding). *A commitment scheme $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ is binding if for every probabilistic polynomial-time adversary \mathcal{S}^* there exists a negligible function $\text{negl}(n)$ such that*

$$\Pr \left[\begin{array}{l} \mathcal{V}(\text{decom}, \text{com}) = 0 \text{ and} \\ \mathcal{V}(\text{decom}', \text{com}) = 1 \end{array} : (\text{decom}, \text{decom}', \text{com}) \leftarrow \langle \mathcal{S}^*(1^n), \mathcal{R} \rangle \right] \leq \text{negl}(n)$$

for all $n \in \mathbb{N}$, where the probability is taken over the random coins of both \mathcal{S}^* and \mathcal{R} .

Given a commitment scheme $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ and an adversary \mathcal{R}^* , we denote by $\text{view}_{\langle \mathcal{S}(b), \mathcal{R}^* \rangle}(n)$ the distribution on the view of \mathcal{R}^* when interacting with $\mathcal{S}(1^n, b)$. The view consists of \mathcal{R}^* 's random coins and the sequence of messages it received from \mathcal{S} . The distribution is taken over the random coins of both \mathcal{S} and \mathcal{R} . Without loss of generality, whenever \mathcal{R}^* has no computational restrictions, we can assume it is deterministic.

Definition 5 (hiding). *A commitment scheme $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ is statistically hiding if there exists a negligible function $\text{negl}(n)$ such that for every (deterministic) adversary \mathcal{R}^* it holds that*

$$\Delta(\{\text{view}_{\langle \mathcal{S}(0), \mathcal{R}^* \rangle}(n)\}, \{\text{view}_{\langle \mathcal{S}(1), \mathcal{R}^* \rangle}(n)\}) \leq \text{negl}(n)$$

for all $n \in \mathbb{N}$.

3.4 Distributional Collision Resistant Hash Functions

Roughly speaking, a distributional collision resistant hash function [DI06] guarantees that no efficient adversary can sample a uniformly random collision. We start by defining more precisely what we mean by a random collision throughout the paper, and then move to the actual definition.

Definition 6 (Ideal collision finder). *Let Col be the random function that given a (description) of a function $h: \{0, 1\}^n \rightarrow \{0, 1\}^m$ as input, returns a collision (x_1, x_2) with respect to h as follows: it samples a uniformly random element, $x_1 \leftarrow \{0, 1\}^n$, and then samples a uniformly random element that collides with x_1 under h , $x_2 \leftarrow \{x \in \{0, 1\}^n: h(x) = h(x_1)\}$. (Note that possibly, $x_1 = x_2$.)*

Definition 7 (Distributional collision resistant hashing). *Let $\mathcal{H} = \{\mathcal{H}_n: \{0, 1\}^n \rightarrow \{0, 1\}^{m(n)}\}_{n \in \mathbb{N}}$ be an efficient function family ensemble. We say that \mathcal{H} is a secure distributional collision resistant hash (dCRH) function family if there exists a polynomial $p(\cdot)$ such that for any probabilistic polynomial-time algorithm A , it holds that*

$$\Delta((h, A(1^n, h)), (h, \text{Col}(h))) \geq \frac{1}{p(n)},$$

for $h \leftarrow \mathcal{H}_n$ and large enough $n \in \mathbb{N}$.

Comparison with the previous definition. Our definition deviates from the previous definition of distributional collision resistance considered in [DI06, HN10, KY18]. The definition in the above-mentioned works is equivalent to requiring that for any efficient adversary A , there exists a polynomial p_A , such that the collision output by A is $\frac{1}{p_A(n)}$ -far from a random collision on average (over h). Our definition switches the order of quantifiers, requiring that there is one such polynomial $p(\cdot)$ for all adversaries A .

We note that the previous definition is, in fact, not even known to imply one-way functions. In contrast, the definition presented here strengthens that of *distributional one-way functions*, which in turn implies one-way functions [IL89]. Additionally, note that both constructions of distributional collision resistance in [KY18] (from multi-collision resistance and from SZK hardness) satisfy our stronger notion of security (with a similar proof).

On compression. As opposed to classical notions of collision resistance (such as plain collision resistance or multi-collision resistance), it makes sense to require distributional collision resistance even for *non-compressing* functions. So we do not put a restriction on the order between n and $m(n)$. As a matter of fact, by padding, the input, arbitrary polynomial compression can be assumed without loss of generality.

4 From dCRH to Statistically Hiding Commitments and Back

We show distributional collision resistant hash functions imply constant-round statistically hiding commitments.

Theorem 3. *Assume the existence of a distributional collision resistant hash function family. Then, there exists a constant-round statistically hiding and computationally binding commitment scheme.*

Our proof relies on the transformation of Haitner et al. [HRVW09, HRVW18], translating inaccessible-entropy generators to statistically hiding commitments. Concretely, we construct appropriate inaccessible-entropy generators from distributional collision resistant hash functions. In Section 4.1, we recall the necessary definitions and the result of [HRVW18], and then in Section 4.2, we prove Theorem 3.

We complement the above result by showing a loose converse to Theorem 3, namely that two message statistically hiding commitments (with possibly large communication) imply the existence of distributional collision resistance hashing.

Theorem 4. *Assume the existence of a binding and statistically hiding two-message commitment scheme. Then, there exists a dCRH function family.*

This proof of Theorem 4 appears in Section 4.3.

4.1 Preliminaries on Inaccessible Entropy Generators

The following definitions of real and accessible entropy of protocols are taken from [HRVW18].

Definition 8 (Block generators). *Let n be a security parameter, and let $c = c(n)$, $s = s(n)$ and $m = m(n)$. An m -block generator is a function $G: \{0, 1\}^c \times \{0, 1\}^s \mapsto (\{0, 1\}^*)^m$. It is efficient if its running time on input of length $c(n) + s(n)$ is polynomial in n .*

We call parameter n the security parameter, c the public parameter length, s the seed length, m the number of blocks, and $\ell(n) = \max_{(z,x) \in \{0,1\}^{c(n)} \times \{0,1\}^{s(n)}, i \in [m(n)]} |G(z, x)_i|$ the maximal block length of G .

Definition 9 (Real sample-entropy). *Let G be an m -block generator over $\{0, 1\}^c \times \{0, 1\}^s$, let $n \in \mathbb{N}$, let Z_n and X_n be uniformly distributed over $\{0, 1\}^{c(n)}$ and $\{0, 1\}^{s(n)}$, respectively, and let $\mathbf{Y}_n = (Y_1, \dots, Y_m) = G(Z_n, X_n)$. For $n \in \mathbb{N}$ and $i \in [m(n)]$, define the real sample-entropy of $\mathbf{y} \in \text{Supp}(Y_1, \dots, Y_i)$ given $z \in \text{Supp}(Z_n)$ as*

$$\text{RealH}_{G,n}(\mathbf{y}|z) = \sum_{j=1}^i H_{Y_j|Z_n, Y_{<j}}(\mathbf{y}_j|z, \mathbf{y}_{<j}).$$

We omit the security parameter from the above notation when clear from the context.

Definition 10 (Real entropy). *Let G be an m -block generator, and let Z_n and \mathbf{Y}_n be as in Definition 9. Generator G has real entropy at least $k = k(n)$, if*

$$\mathbf{E}_{(z, \mathbf{y}) \leftarrow (Z_n, \mathbf{Y}_n)} [\text{RealH}_{G,n}(\mathbf{y}|z)] \geq k(n)$$

for every $n \in \mathbb{N}$.

The generator G has real min-entropy at least $k(n)$ in its i 'th block for some $i = i(n) \in [m(n)]$, if

$$\Pr_{(z, \mathbf{y}) \leftarrow (Z_n, \mathbf{Y}_n)} [\text{H}_{Y_i|Z_n, Y_{<i}}(\mathbf{y}_i|z, \mathbf{y}_{<i}) < k(n)] = \text{negl}(n).$$

We say the above bounds are invariant to the public parameter if they hold for any fixing of the public parameter Z_n .⁴

It is known that the real Shannon entropy amounts to measuring the standard conditional Shannon entropy of G 's output blocks.

Lemma 2 ([HRVW18, Lemma 3.4]). *Let G , Z_n and \mathbf{Y}_n be as in definition 9 for some $n \in \mathbb{N}$, then*

$$\mathbf{E}_{(z, \mathbf{y}) \leftarrow (Z_n, \mathbf{Y}_n)} [\text{RealH}_{G,n}(\mathbf{y}|z)] = \text{H}(\mathbf{Y}_n|Z_n).$$

Toward the definition of *inaccessible entropy*, we first define *online block-generators* which are a special type of block generators that toss fresh random coins before outputting each new block.

Definition 11 (Online block generator). *Let n be a security parameter, and let $c = c(n)$ and $m = m(n)$. An m -block online generator is a function $\tilde{G}: \{0, 1\}^c \times (\{0, 1\}^v)^m \mapsto (\{0, 1\}^*)^m$ for some $v = v(n)$, such that the i 'th output block of \tilde{G} is a function of (only) its first i input blocks. We denote the transcript of \tilde{G} over random input by $T_{\tilde{G}}(1^n) = (Z, R_1, Y_1, \dots, R_m, Y_m)$, for $Z \leftarrow \{0, 1\}^c$, $(R_1, \dots, R_m) \leftarrow (\{0, 1\}^v)^m$ and $(Y_1, \dots, Y_m) = \tilde{G}(Z, R_1, \dots, R_i)$.*

That is, an online block generator is a special type of block generator that tosses fresh random coins before outputting each new block. In the following, we let $\tilde{G}(z, r_1, \dots, r_i)_i$ stand for $\tilde{G}(z, r_1, \dots, r_i, x^*)_i$ for arbitrary $x^* \in (\{0, 1\}^v)^{m-i}$ (note that the choice of x^* has no effect on the value of $\tilde{G}(z, r_1, \dots, r_i, x^*)_i$).

Definition 12 (Accessible sample-entropy). *Let n be a security parameter, and let \tilde{G} be an online $m = m(n)$ -block online generator. The accessible sample-entropy of $\mathbf{t} = (z, r_1, y_1, \dots, r_m, y_m) \in \text{Supp}(Z, R_1, Y_1, \dots, R_m, Y_m) = T_{\tilde{G}}(1^n)$ is defined by*

$$\text{AccH}_{\tilde{G},n}(\mathbf{t}) = \sum_{i=1}^m \text{H}_{Y_i|Z, R_{<i}}(y_i|z, r_{<i}).$$

Again, we omit the security parameter from the above notation when clear from the context.

As in the case of real entropy, the expected accessible entropy of a random transcript can be expressed in terms of the standard conditional Shannon entropy.

⁴In particular, this is the case when there is no public parameter, i.e., $c = 0$.

Lemma 3 ([HRVW18, Lemma 3.7]). *Let \tilde{G} be an online m -block generator and let $(Z, R_1, Y_1, \dots, R_m, Y_m) = T_{\tilde{G}}(1^n)$ be its transcript. Then,*

$$\mathbf{E}_{\mathbf{t} \leftarrow T_{\tilde{G}}(Z, 1^n)} [\text{AccH}_{\tilde{G}}(\mathbf{t})] = \sum_{i \in [m]} \mathbf{H}(Y_i | Z, R_{<i}).$$

We focus on efficient generators that are consistent with respect to G . That is, the support of their output is contained in that of G .

Definition 13 (Consistent generators). *Let G be a block generator over $\{0, 1\}^{c(n)} \times \{0, 1\}^{s(n)}$. A block (possibly online) generator G' over $\{0, 1\}^{c(n)} \times \{0, 1\}^{s'(n)}$ is G consistent if, for every $n \in \mathbb{N}$, it holds that $\text{Supp}(G'(U_{c(n)}, U_{s'(n)})) \subseteq \text{Supp}(G(U_{c(n)}, U_{s(n)}))$.*

Definition 14 (Accessible entropy). *A block generator G has accessible entropy at most $k = k(n)$ if, for every efficient G -consistent, online generator \tilde{G} and all large enough n ,*

$$\mathbf{E}_{\mathbf{t} \leftarrow T_{\tilde{G}}(1^n)} [\text{AccH}_{\tilde{G}}(\mathbf{t})] \leq k.$$

We call a generator whose real entropy is noticeably higher than its accessible entropy an inaccessible entropy generator.

We use the following reduction from inaccessible entropy generators to constant round statistically hiding commitment.

Theorem 5 ([HRVW18, Thm. 6.24]). *Let G be an efficient block generator with constant number of blocks. Assume G 's real Shannon entropy is at least $k(n)$ for some efficiently computable function k , and that its accessible entropy is bounded by $k(n) - 1/p(n)$ for some $p \in \text{poly}$. Then there exists a constant-round statistically hiding and computationally binding commitment scheme. Furthermore, if the bound on the real entropy is invariant to the public parameter, then the commitment is receiver public-coin.*

Remark 1 (Inaccessible max/average entropy). *Our result relies on the reduction from inaccessible Shannon entropy generators to statistically hiding commitments, given in [HRVW18]. The proof of this reduction follows closely the proof in previous versions [HV17, HRVW09], where the reduction was from inaccessible max entropy generators. The extension to Shannon entropy generators is essential for our result.*

4.2 From dCRH to Inaccessible Entropy Generators – Proof of Theorem 3

In this section we show that there is a block generator with two blocks in which there is a gap between the real entropy and the accessible entropy. Let $\mathcal{H} = \{\mathcal{H}_n: \{0, 1\}^n \rightarrow \{0, 1\}^m\}_{n \in \mathbb{N}}$ be a dCRH for $m = m(n)$ and assume that each $h \in \mathcal{H}_n$ requires $c = c(n)$ bits to describe. By Definition 7, there exists a polynomial $p(\cdot)$ such that for any probabilistic polynomial-time algorithm A , it holds that

$$\Delta((h, A(1^n, h)), (h, \text{Col}(h))) = \mathbf{E}_{h \leftarrow \mathcal{H}_n} [\Delta(A(1^n, h), \text{Col}(h))] \geq \frac{1}{p(n)}$$

for large enough $n \in \mathbb{N}$, where $h \leftarrow \mathcal{H}_n$.

The generator $G: \{0, 1\}^c \times \{0, 1\}^n \rightarrow \{0, 1\}^m \times \{0, 1\}^n$ is defined by

$$G(h, x) = (h(x), x).$$

The public parameter length is c (this is the description size of h), the generator consists of two blocks, and the maximal block length is $\max\{n, m\}$. Since the random coins of G define x and x is completely revealed, the real Shannon entropy of G is n . That is,

$$\mathbf{E}_{y \leftarrow G(U_c, U_n)} [\text{RealH}_G(y)] = n.$$

Our goal in the remaining of this section is to show a non-trivial upper bound on the accessible entropy of G . We prove the following lemma.

Lemma 4. *There exists a polynomial $q(\cdot)$ such that for every G -consistent online generator \tilde{G} , it holds that*

$$\mathbf{E}_{t \leftarrow T_{\tilde{G}}(Z, 1^n)} [\text{AccH}_{\tilde{G}}(t)] \leq n - \frac{1}{q(n)}$$

for all large enough $n \in \mathbb{N}$.

Proof. Fix a G -consistent online generator \tilde{G} . Let us denote by Y a random variable that corresponds to the first part of G 's output (i.e., the first m bits) and by X the second part (i.e., the last n bits). Denote by R the randomness used by the adversary to sample Y . Denote by Z the random variable that corresponds to the description of the hash function h . Fix $q(n) \triangleq 4 \cdot p(n)^2$. Assume towards contradiction that for infinitely many n 's it holds that

$$\mathbf{E}_{t \leftarrow T_{\tilde{G}}(Z, 1^n)} [\text{AccH}_{\tilde{G}}(t)] > n - \frac{1}{q(n)}.$$

By Lemma 3, this means that

$$\text{H}(Y | Z) + \text{H}(X | Y, Z, R) > n - \frac{1}{q(n)} \tag{1}$$

We show how to construct an adversary A that can break the security of the dCRH. The algorithm A , given a hash function $h \leftarrow \mathcal{H}$, does the following:

1. Sample r and let $y = \tilde{G}(h, r)_1$
2. Sample r_1, r_2 and output $x_1 = \tilde{G}(h, r, r_1)_2$ and $x_2 = \tilde{G}(h, r, r_2)_2$.

In other words, A tries to create a collision by running G to get the first block, y , and then running it twice (by rewinding) to get two inputs x_1, x_2 that are mapped to y . Indeed, A runs in polynomial-time and if \tilde{G} is G -consistent, then x_1 and x_2 collide relative to h . Denote by Y^A , X_1^A , and X_2^A be random variables that correspond to the output of the emulated \tilde{G} . Furthermore, denote by $(X_1^{\text{Col}}, X_2^{\text{Col}})$ a random collision that $\text{Col}(h)$ samples. To finish the proof it remains to show that

$$\mathbf{E}_{h \leftarrow \mathcal{H}_n} \left[\Delta((X_1^A, X_2^A), (X_1^{\text{Col}}, X_2^{\text{Col}})) \right] \leq \frac{1}{p(n)}$$

which is a contradiction.

By Pinsker's inequality (Proposition 2) and the chain rule from Proposition 1, it holds that

$$\begin{aligned} \Delta\left(\left(X_1^A, X_2^A\right), \left(X_1^{\text{Col}}, X_2^{\text{Col}}\right)\right) &\leq \sqrt{\frac{\ln(2)}{2} \cdot \mathbf{D}_{\text{KL}}\left(X_1^A, X_2^A \parallel X_1^{\text{Col}}, X_2^{\text{Col}}\right)} \\ &= \sqrt{\mathbf{D}_{\text{KL}}\left(X_1^A \parallel X_1^{\text{Col}}\right) + \mathbf{E}_{x_1 \leftarrow X_1^A} \left[\mathbf{D}_{\text{KL}}\left(X_2^A \mid X_1^A = x_1 \parallel X_2^{\text{Col}} \mid X_1^{\text{Col}} = x_1\right) \right]} \\ &\leq \sqrt{\mathbf{D}_{\text{KL}}\left(X_1^A \parallel X_1^{\text{Col}}\right) + \sqrt{\mathbf{E}_{x_1 \leftarrow X_1^A} \left[\mathbf{D}_{\text{KL}}\left(X_2^A \mid X_1^A = x_1 \parallel X_2^{\text{Col}} \mid X_1^{\text{Col}} = x_1\right) \right]}}. \end{aligned}$$

Hence, by Jensen's inequality (Proposition 3), it holds that

$$\begin{aligned} \mathbf{E}_{h \leftarrow \mathcal{H}_n} \left[\Delta\left(\left(X_1^A, X_2^A\right), \left(X_1^{\text{Col}}, X_2^{\text{Col}}\right)\right) \right] &\leq \sqrt{\mathbf{E}_{h \leftarrow \mathcal{H}_n} \left[\mathbf{D}_{\text{KL}}\left(X_1^A \parallel X_1^{\text{Col}}\right) \right]} + \\ &\quad \sqrt{\mathbf{E}_{h \leftarrow \mathcal{H}_n} \left[\mathbf{D}_{\text{KL}}\left(X_2^A \mid X_1^A = x_1 \parallel X_2^{\text{Col}} \mid X_1^{\text{Col}} = x_1\right) \right]}. \end{aligned}$$

We complete the proof using the following claims.

Claim 1. *It holds that*

$$\mathbf{E}_{h \leftarrow \mathcal{H}_n} \left[\mathbf{D}_{\text{KL}}\left(X_1^A \parallel X_1^{\text{Col}}\right) \right] \leq \frac{1}{p(n)^2}.$$

Claim 2. *It holds that*

$$\mathbf{E}_{\substack{h \leftarrow \mathcal{H}_n \\ x_1 \leftarrow X_1^A}} \left[\mathbf{D}_{\text{KL}}\left(X_2^A \mid X_1^A = x_1 \parallel X_2^{\text{Col}} \mid X_1^{\text{Col}} = x_1\right) \right] \leq \frac{1}{p(n)^2}.$$

Proof of Claim 1. Recall that X_1^{Col} is the *uniform* distribution over the inputs of the hash function and thus

$$\mathbf{D}_{\text{KL}}\left(X_1^A \parallel X_1^{\text{Col}}\right) = \sum_x \Pr\left[X_1^A = x\right] \cdot \log \frac{\Pr\left[X_1^A = x\right]}{2^{-n}} = n - \mathbf{H}\left(X_1^A\right).$$

To sample X_1^A , the algorithm A first runs $\tilde{G}(r)_1$ to get y and then runs $G(r, r_1)$ to get x_1 . Thus, by Equation (1), it holds that

$$\mathbf{E}_{h \leftarrow \mathcal{H}_n} \left[\mathbf{H}\left(X_1^A\right) \right] = \mathbf{E}_{h \leftarrow \mathcal{H}_n} \left[\mathbf{H}(X) \right] = \mathbf{H}(X, Y \mid Z) = \mathbf{H}(Y \mid Z) + \mathbf{H}(X \mid Y, Z, R) \geq n - \frac{1}{q(n)},$$

where the second equality follows since \tilde{G} is G -consistent and thus X fully determines Y . This implies that

$$\mathbf{E}_{h \leftarrow \mathcal{H}_n} \left[\mathbf{D}_{\text{KL}}\left(X_1^A \parallel X_1^{\text{Col}}\right) \right] \leq \frac{1}{q(n)} = \frac{1}{p(n)^2},$$

as required. □

Proof of Claim 2. For $x_1 \in \text{supp}(X_1^A)$, it holds that

$$\begin{aligned} \mathbf{D}_{\text{KL}}(X_2^A |_{X_1^A=x_1} \| X_2^{\text{Col}} |_{X_1^{\text{Col}}=x_1}) &= \sum_x \mathbf{Pr}[X_2^A = x |_{X_1^A=x_1}] \cdot \log \frac{\mathbf{Pr}[X_2^A = x |_{X_1^A=x_1}]}{|h^{-1}(h(x_1))|^{-1}} \\ &= \log |h^{-1}(h(x_1))| - \mathbf{H}(X_2^A |_{X_1^A=x_1}). \end{aligned}$$

Hence,

$$\mathbf{E}_{\substack{h \leftarrow \mathcal{H}_n \\ x_1 \leftarrow X_1^A}} \left[\mathbf{D}_{\text{KL}}(X_2^A |_{X_1^A=x_1} \| X_2^{\text{Col}} |_{X_1^{\text{Col}}=x_1}) \right] = \mathbf{E}_{\substack{h \leftarrow \mathcal{H}_n \\ x_1 \leftarrow X_1^A}} \left[\log |h^{-1}(h(x_1))| - \mathbf{H}(X_2^A |_{X_1^A=x_1}) \right].$$

Notice that the distribution of X_2^A only depends on $y = h(x_1)$, that is, $X_2^A |_{X_1^A=x_1}$ is distributed exactly as $X_2^A |_{X_1^A=x'_1}$ for every x_1 and x'_1 that such that $y = h(x_1) = h(x'_1)$. Thus, we have that $X_2^A |_{X_1^A=x_1}$ is distributed exactly as $X |_{Y=y}$ and the distribution of $h(X_1)$ is distributed as Y . Namely,

$$\begin{aligned} \mathbf{E}_{\substack{h \leftarrow \mathcal{H}_n \\ x_1 \leftarrow X_1^A}} \left[\mathbf{D}_{\text{KL}}(X_2^A |_{X_1^A=x_1} \| X_2^{\text{Col}} |_{X_1^{\text{Col}}=x_1}) \right] &= \mathbf{E}_{\substack{h \leftarrow \mathcal{H}_n \\ x_1 \leftarrow X_1^A}} [\log |h^{-1}(y)|] - \mathbf{E}_{h \leftarrow \mathcal{H}_n} [\mathbf{H}(X | Y, R)] \\ &= \mathbf{E}_{\substack{h \leftarrow \mathcal{H}_n \\ x_1 \leftarrow X_1^A}} [\log |h^{-1}(y)|] - \mathbf{H}(X | Y, Z, R) \\ &\leq \mathbf{E}_{\substack{h \leftarrow \mathcal{H}_n \\ x_1 \leftarrow X_1^A}} [\log |h^{-1}(y)|] + \mathbf{H}(Y | Z) - n + \frac{1}{q(n)} \\ &= \frac{1}{q(n)}, \end{aligned}$$

where the first inequality follows by Equation (1) and the second follows since

$$\begin{aligned} \mathbf{E}_{\substack{h \leftarrow \mathcal{H}_n \\ y \leftarrow Y}} [\log |h^{-1}(y)|] + \mathbf{H}(Y | Z) &= \mathbf{E}_{\substack{h \leftarrow \mathcal{H}_n \\ y \leftarrow Y}} [\log |h^{-1}(y)| + \mathbf{H}_Y(y)] \\ &= \mathbf{E}_{\substack{h \leftarrow \mathcal{H}_n \\ y \leftarrow Y}} \left[\log \frac{|h^{-1}(y)|}{\mathbf{Pr}[Y = y]} \right] \\ &\leq \log \mathbf{E}_{\substack{h \leftarrow \mathcal{H}_n \\ y \leftarrow Y}} \left[\frac{|h^{-1}(y)|}{\mathbf{Pr}[Y = y]} \right] = n, \end{aligned}$$

where the inequality is by Jensen's inequality (Proposition 3). Thus, overall

$$\mathbf{E}_{\substack{h \leftarrow \mathcal{H}_n \\ x_1 \leftarrow X_1^A}} \left[\mathbf{D}_{\text{KL}}(X_2^A |_{X_1^A=x_1} \| X_2^{\text{Col}} |_{X_1^{\text{Col}}=x_1}) \right] \leq \frac{1}{q(n)} = \frac{1}{p(n)^2},$$

as required. □

□

4.3 From Statistically Hiding Commitments to dCRH– Proof of Theorem 4

Let $\pi = (\mathcal{S}, \mathcal{R}, \mathcal{V})$ be a binding and statistically hiding two-message commitment scheme. We show that there exists a dCRH family \mathcal{H} .

To sample a hash function in the family with security parameter n , we use the receiver’s first message of the protocol. Namely, we set the hash function as $h \leftarrow \mathcal{R}(1^n)$. Then, to evaluate h on input x we first parse x as $x = (b, r)$, where b is a bit, and output a commitment to the bit b using randomness r , with respect to the receiver message h . That is, we set

$$h(x) = \mathcal{S}(h, b; r).$$

Since π is efficient, then sampling and evaluating h are polynomial-time procedures. This concludes the definition of our family \mathcal{H} of hash functions. (Note that the functions in the family are not necessarily compressing.)

We next argue security. Suppose toward contradiction that \mathcal{H} is not a dCRH according to Definition 7. Then, for any $\delta(n) = n^{-O(1)}$ there exists an adversary A , such that

$$\Delta((h, A(1^n, h)), (h, \text{Col}(h))) \leq \delta, \tag{2}$$

for infinitely many n ’s. From hereon, we fix δ to be any function such that $n^{-O(1)} < \delta < \frac{1}{2} - n^{-O(1)}$.

We show how to use A to break the binding property of the commitment scheme. Our cheating receiver \mathcal{R}^* is defined as follows: On input h , \mathcal{R}^* runs $A(h)$ to get x and x' , interprets $x = (b, r)$ and $x' = (b', r')$ and outputs b and b' along with their openings r and r' , respectively. Our goal is to show that $x = (b, r)$ and $x' = (b', r')$ are two valid distinct openings to the commitment scheme.

By Equation (2), it suffices to analyze the success probability when the pair (x, x') is sampled according to the distribution Col_h , and show that it is at least $1/2 - \text{negl}(n)$. From the definition of Col_h , we have that $h(x) = h(x')$ and thus $\mathcal{S}(h, b; r) = \mathcal{S}(h, b'; r') := y$. In other words, the second message of the protocol for b with randomness r and b' with randomness r' are the same, and thus both pass as valid openings in the reveal stage of the protocol: $\mathcal{V}(h, y, b, r) = 1$ and $\mathcal{V}(h, y, b', r') = 1$.

We are left to show that these are two *distinct* openings for the commitment, namely, $b \neq b'$. To show this, we use the statistically hiding property of the commitment scheme. The following claim concludes the proof.

Claim 1. *Fix any h . Then for $((b, r), (b', r')) \leftarrow \text{Col}(h)$ it holds that $\Pr[b \neq b'] \geq 1/2 - \text{negl}(n)$.*

Proof. Let B be the uniform distribution on bits and R the uniform distribution on commitment randomness. For every commitment c , let B_c be the distribution on bits given by sampling $(b, r) \leftarrow (B, R)$ conditioned on $\mathcal{S}(h, b; r) = c$. Let C be the distribution on random commitments to a random bit.

By the statistical hiding property of the commitment scheme,

$$\Delta((\mathcal{S}(h, B, R), B), (\mathcal{S}(h, B', R), B)) \leq \epsilon,$$

where B' is an independent copy of B , and $\epsilon = \text{negl}(n)$ is a negligible function. Furthermore,

$$\Delta((\mathcal{S}(h, B, R), B), (\mathcal{S}(h, B', R), B)) = \Delta((C, B_C), (C, B)) = \mathbf{E}_{c \leftarrow C}[\Delta(B_c, B)].$$

By Markov's inequality, it holds that

$$\Pr_{c \leftarrow C} [\Delta(B_c, B) \geq \sqrt{\varepsilon}] \leq \sqrt{\varepsilon} .$$

To conclude the proof note that

$$\begin{aligned} \Pr [b = b' : (b, r), (b', r') \leftarrow \text{Col}_h] &= \Pr \left[b = b' : \begin{array}{l} (b, r) \leftarrow (B, R) \\ c = \mathcal{S}(h, b; r) \\ b' \leftarrow B_c \end{array} \right] \leq \\ \Pr \left[\begin{array}{l} (b, r) \leftarrow (B, R) \\ c = \mathcal{S}(h, b; r) \\ b' \leftarrow B_c \\ \Delta(B_c, B) \leq \sqrt{\varepsilon} \end{array} \right] &+ \Pr_{c \leftarrow C} [\Delta(B_c, B) \geq \sqrt{\varepsilon}] \leq \\ \left(\frac{1}{2} + \sqrt{\varepsilon} \right) + \sqrt{\varepsilon} &= \frac{1}{2} + \text{negl}(n) . \end{aligned}$$

□

Overall, the success probability of **A** is at least $1/2 - \text{negl}(n) - \delta \geq n^{-O(1)}$.

Using string commitments. The above proof constructs dCRH from statistically hiding *bit* commitment schemes. For schemes that support commitments to *strings*, following the above proof gives a stronger notion of dCRH, where the adversary's output distribution is $(1 - \text{negl}(n))$ -far from a random collision distribution.

Technically, the change in the proof is to interpret b in $x = (b, r)$ as a string of length n , rather than as a single bit. The proof remains the same except that the probability that $b = b'$ is (negligibly close to) 2^{-n} instead of $1/2$. Thus, overall the success probability of **A** is at least $1 - \text{negl}(n) - \delta$. To ensure a polynomial success probability we can allow any $\delta = 1 - n^{-O(1)}$.

5 From SZK-Hardness to Statistically Hiding Commitments

In this section, we give a direct construction of a constant-round statistically hiding commitment from average-case hardness in SZK. This gives an alternative proof to Corollary 1.

5.1 Hard on Average Promise Problems

Definition 15. A *promise problem* (Π_Y, Π_N) consists of two disjoint sets of yes instances Π_Y and no instances Π_N .

Definition 16. A *promise problem* (Π_Y, Π_N) is *hard on average* if there exists a probabilistic polynomial-time sampler Π with support $\Pi_Y \cup \Pi_N$, such that for any probabilistic polynomial-time decider D , there exists a negligible function $\text{negl}(n)$, such that

$$\Pr_{r \leftarrow \{0,1\}^n} [x \in \Pi_{D(x)} \mid x \leftarrow \Pi(r)] \leq \frac{1}{2} + \text{negl}(n) .$$

5.2 Instance-Dependent Commitments

Definition 17 ([OV08]). *An instance-dependent commitment scheme \mathcal{IDC} for a promise problem (Π_Y, Π_N) is a commitment scheme where all algorithms get as auxiliary input an instance $x \in \{0, 1\}^*$. The induced family of schemes $\{\mathcal{IDC}_x\}_{x \in \{0, 1\}^*}$ is*

- *statistically binding when $x \in \Pi_N$,*
- *statistically hiding when $x \in \Pi_Y$.*

Theorem 6 ([OV08]). *Any promise problem $(\Pi_Y, \Pi_N) \in \text{SZK}$ has a constant-round instance-dependent commitment.*

5.3 Witness-Indistinguishable Proofs

Definition 18. *A proof system \mathcal{WI} for an NP relation R is witness indistinguishable if for any x, w_0, w_1 such that $(x, w_0), (x, w_1) \in R$, the verifier's view given a proof using w_0 is computationally indistinguishable from its view given a proof using w_1 .*

Constant-round \mathcal{WI} proofs systems are known from any constant-round statistically-binding commitments [GMW87]. Statistically-binding commitments can be constructed from one-way functions [Nao91], and thus can also be obtained from average-case hardness in SZK [OW93].

Theorem 7 ([GMW87, Nao91, OW93]). *Assuming hard-on-average problems in SZK, there exist constant-round witness-indistinguishable proof systems.*

5.4 The Commitment Protocol

Here, we give the details of our protocol. Our protocol uses the following ingredients and notation:

- A \mathcal{WI} proof for NP.
- A hard-on average SZK problem (Π_Y, Π_N) with sampler Π .
- An instance-dependent commitment scheme \mathcal{IDC} for Π .

We describe the commitment scheme in Figure 1.

5.5 Analysis

Proposition 4. *Protocol 1 is computationally binding.*

Proof. Let \mathcal{S}^* be any probabilistic polynomial-time sender that breaks binding in Protocol 1 with probability ε . We use \mathcal{S}^* to construct a probabilistic polynomial-time decider D for the SZK problem Π with advantage $\varepsilon/4n - \text{negl}(n)$.

Given an instance $x \leftarrow \Pi$, the decider D proceeds as follows:

- It samples at random $i^* \in [n]$ and $b^* \in \{0, 1\}$.
- It executes the protocol $(\mathcal{S}^*, \mathcal{R})$ with the following exceptions:
 - The instance x_{i^*, b^*} , generated by \mathcal{R} , is replaced with the instance x , given to D as input.

Protocol 1

Sender input: a bit $m \in \{0, 1\}$.

Common input: security parameter 1^n .

Coin tossing into the well

- \mathcal{R} samples $2n$ independent random strings $\rho_{i,b} \leftarrow \{0, 1\}^n$, for $i \in [n], b \in \{0, 1\}$.
- The parties then execute (in parallel) $2n$ statistically-binding commitment protocols \mathcal{SBC} in which \mathcal{R} commits to each of the strings $\rho_{i,b}$. We denote the transcript of each such commitment by $C_{i,b}$.
- \mathcal{S} samples $2n$ independent random strings $\sigma_{i,b} \leftarrow \{0, 1\}^n$, and sends them to \mathcal{R} .
- \mathcal{R} sets $r_{i,b} = \rho_{i,b} \oplus \sigma_{i,b}$.

Generating hard instances

- \mathcal{R} generates $2n$ instances $x_{i,b} \leftarrow \Pi(r_{i,b})$, using the strings $r_{i,b}$ as randomness, and sends the instances to \mathcal{S} .
- The parties then execute a \mathcal{WI} protocol in which \mathcal{R} proves to \mathcal{S} that there exists a $b \in \{0, 1\}$ such that for all $i \in [n]$, $x_{i,b}$ was generated consistently. That is, there exist strings $\{\rho_{i,b}\}_{i \in [n]}$ that are consistent with the receiver's commitments $\{C_{i,b}\}_{i \in [n]}$, and $x_{i,b} = \Pi(\rho_{i,b} \oplus \sigma_{i,b})$.

As the witness, \mathcal{R} uses $b = 0$ and the strings $\{\rho_{i,0}\}_{i \in [n]}$ sampled earlier in the protocol.

Instance-binding commitment

- The sender samples $2n$ random bits $m_{i,b}$ subject to $m = \bigoplus_{i,b} m_{i,b}$.
- The parties then execute (in parallel) $2n$ instance-dependent commitment protocols $\mathcal{IDC}_{x_{i,b}}$ in which \mathcal{S} commits to each bit $m_{i,b}$ using the instance $x_{i,b}$.

Figure 1: A constant round statistically hiding commitment from SZK hardness.

- In the \mathcal{WI} protocol, as the witness we use $1 \oplus b^*$ and the strings $\{\rho_{i,1 \oplus b^*}\}_{i \in [n]}$ (instead of 0 and the strings $\{\rho_{i,0}\}_{i \in [n]}$).
- Then, at the opening phase, if \mathcal{S}^* equivocally opens the (i^*, b^*) -th instance-dependent commitment, D declares that $x \in \Pi_Y$. Otherwise, it declares that $x \in \Pi_\beta$ for a random $\beta \in \{Y, N\}$.

Analyzing D 's advantage. Denote by E the event that in the above experiment \mathcal{S}^* equivocally opens the (i^*, b^*) -th instance-dependent commitment. We first observe that the advantage of D in deciding Π is at least as large as the probability that E occurs.

Claim 3. $\Pr[x \in \Pi_{D(x)}] \geq \frac{1+\Pr[E]}{2} - \text{negl}(n)$.

Proof. By the definition of D ,

$$\Pr[x \in \Pi_{D(x)} | E] = \Pr[x \in \Pi_Y | E] = 1 - \Pr[x \in \Pi_N | E] \geq 1 - \frac{\Pr[E | x \in \Pi_N]}{\Pr[E]},$$

$$\Pr[x \in \Pi_{D(x)} | \bar{E}] = \frac{1}{2}.$$

Furthermore, if $x \in \Pi_N$ (namely, it is a no instance), then \mathcal{IDC}_x is binding, and thus

$$\Pr[E | x \in \Pi_N] = \text{negl}(n).$$

Claim 3 now follows by the law of total probability. \square

From hereon, we focus on showing that E occurs with high probability.

Claim 4. $\Pr[E] \geq \frac{\varepsilon}{2n} - \text{negl}(n)$.

Proof. To prove the claim, we consider hybrid experiments $\mathcal{H}_0, \dots, \mathcal{H}_4$, and show that that the view of the sender \mathcal{S}^* changes in a computationally indistinguishable manner throughout the hybrids. We then bound the probability that E occurs in the last hybrid experiment.

\mathcal{H}_0 : In this experiment, we consider an execution of $D(x)$ as specified above.

\mathcal{H}_1 : Here x is not sampled ahead of time, but rather first the value σ_{i^*, b^*} is obtained from \mathcal{S}^* , then a random value $\rho' \leftarrow \{0, 1\}^n$ is sampled, and x is sampled using randomness $r_{i^*, b^*} = \sigma_{i^*, b^*} \oplus \rho'$. Since ρ' is sampled independently of the rest of the experiment, the sender's view in \mathcal{H}_1 is identically distributed to its view in \mathcal{H}_0 .

\mathcal{H}_2 : Here the (i^*, b^*) -th commitment to ρ_{i^*, b^*} is replaced with a commitment to ρ' . By the (computational) hiding of the commitment \mathcal{SBC} , the sender's view in \mathcal{H}_2 is computationally indistinguishable from its view in \mathcal{H}_1 .

\mathcal{H}_3 : Here, in the \mathcal{WI} protocol, instead of using as the witness $1 \oplus b^*$ and the strings $\{\rho_{i, 1 \oplus b^*}\}_i$, we use 0 and the strings $\{\rho_{i, 0}\}_i$. By the (computational) witness-indistinguishability of the protocol, the sender's view in \mathcal{H}_3 is computationally indistinguishable from its view in \mathcal{H}_2 .

\mathcal{H}_4 : In this experiment, we consider a standard execution of the protocol between \mathcal{S}^* and \mathcal{R} (without any exceptions). The sender's view in this hybrid is identical to its view in \mathcal{H}_3 (by renaming $\rho' = \rho_{i^*, b^*}$ and $x = x_{i^*, b^*}$).

It is left to bound from below the probability that E occurs in \mathcal{H}_4 . That is, when we consider a standard execution of $(\mathcal{S}^*, \mathcal{R})$ and sample (i^*, b^*) independently at random.

Indeed, note that since the plaintext bit m is uniquely determined by the bits $\{m_{i, b}\}_{i, b}$. Whenever \mathcal{S}^* equivocally opens the commitment to two distinct bits, there exists (at least one) (i, b) such that \mathcal{S}^* equivocally opens the (i, b) -th instance-dependent commitment. Since in a standard execution \mathcal{S}^* equivocally opens the commitment with probability at least ε , and (i^*, b^*) is sampled independently, E occurs in this experiment with probability at least $\frac{\varepsilon}{2n}$.

Claim 4 follows. \square

This completes the proof that the scheme is binding. \square

Proposition 5. *Protocol 1 is statistically hiding.*

Proof. Let \mathcal{R}^* be any (computationally unbounded) receiver. We show that the view of \mathcal{R}^* given a commitment to $m = 0$ is statistically indistinguishable from its view given a commitment to $m = 1$.

For this purpose, consider the view of the receiver \mathcal{R}^* after the coin tossing and instance-generation phase (and before the instance-dependent commitment phase). We shall refer to this as the *preamble view*. We say that the preamble view is *admissible*, if either of the following occurs:

- Let $\{x_{i,b}\}_{i,b}$ be the instances sent by \mathcal{R}^* . Then there exists i^*, b^* such that $x_{i^*, b^*} \in \Pi_Y$.
- The sender \mathcal{S} rejects the \mathcal{WI} proof that $\{x_{i,b}\}_{i,b}$ were properly generated.

To complete the proof, we show that the preamble view is admissible with overwhelming probability, and that conditioned on any admissible preamble view, the view of \mathcal{R}^* given a commitment to $m = 0$ is statistically indistinguishable from its view given a commitment to $m = 1$. Since the preamble view is completely independent of m , the above two conditions are sufficient to establish statistical indistinguishability of the total views.

Claim 5. *The probability that the preamble view is not admissible is negligible.*

Proof. Let A be the event that the \mathcal{WI} proof is accepted and let Y be the event that for some (i, b) , $x_{i,b}$ is a yes instance. To show that the preamble view is not admissible with negligible probability, we would like to prove that

$$\Pr[A \wedge \overline{Y}] \leq \text{negl}(n) .$$

Let T be the event that the statement proven by \mathcal{R}^* in the \mathcal{WI} protocol is true. Namely, there exists $b \in \{0, 1\}$ such that all $\{x_{i,b}\}_i$ are generated consistently with the coin-tossing phase (and in particular where the coin-tossing phase consists of valid commitments $\{C_{i,b}\}_i$).

First, note that by the soundness of the \mathcal{WI} system, the probability that the preamble is admissible, and in particular the proof is accepted, when the statement is false, is negligible:

$$\Pr[A \wedge \overline{T}] \leq \text{negl}(n) .$$

We now show:

$$\Pr[\overline{Y} \wedge T] \leq \text{negl}(n) .$$

For this purpose, fix any \mathcal{SBC} commitments $\{C_{i,b}\}_{i,b}$. Let $F = F[\{C_{i,b}\}_{i,b}]$ be the event, over the sender randomness $\{\sigma_{i,b}\}_{i,b}$, that there exists $\beta \in \{0, 1\}$ such that $\{C_{i,\beta}\}_i$ are valid commitments to strings $\{\rho_{i,\beta}\}_i$ and for all i , $\Pi(\rho_{i,\beta} \oplus \sigma_{i,\beta}) = x_{i,\beta} \in \Pi_N$. We show

$$\Pr[F] \leq 2^{-\Omega(n)} .$$

This is sufficient since

$$\Pr[\overline{Y} \wedge T] \leq \max_{\substack{C_{1,0} \dots C_{n,0} \\ C_{1,1} \dots C_{n,1}}} \Pr[F] \leq 2^{-\Omega(n)} .$$

To bound the probability that F occurs, fix any β and commitments $\{C_{i,\beta}\}_i$ to strings $\{\rho_{i,\beta}\}_i$. Then the strings $\rho_{i,\beta} \oplus \sigma_{i,\beta}$ are distributed uniformly and independently at random. Since $\Pi \in \Pi_Y$ with probability at least 0.49, and taking a union bound over both $\beta \in \{0, 1\}$, the bound follows.

This concludes the proof of Claim 5. \square

Claim 6. Fix any admissible preamble view V . Then, conditioned on V the view of \mathcal{R}^* when given a commitment to $m = 0$ is statistically indistinguishable from its view when given a commitment to $m = 1$.

Proof. If V is such that the WI proof is rejected then \mathcal{S} aborts and the view of \mathcal{R}^* remains independent of m . Thus, from hereon, we assume that the instances corresponding to V include an instance $x_{i^*,b^*} \in \Pi_Y$. In particular, the corresponding instance-dependent commitment $\mathcal{IDC}_{x_{i^*,b^*}}$ is statistically hiding.

It is left to note that in any execution $(\mathcal{S}, \mathcal{R}^*)$, with either $m \in \{0, 1\}$, the bits $M_{-i} := \{m_{i,b}\}_{(i,b) \neq (i^*,b^*)}$ are distributed uniformly and independently at random. Conditioned on V and M_{-i} , only the bit

$$m_{i^*,b^*} = m \bigoplus_{m' \in M_{-i}} m'$$

depends on m . By the statistical hiding of $\mathcal{IDC}_{x_{i^*,b^*}}$ a commitment to $0 \bigoplus_{m' \in M_{-i}} m'$ is statistically indistinguishable from a commitment to $1 \bigoplus_{m' \in M_{-i}} m'$.

This concludes the proof of Claim 6. □

□

References

- [AHI⁺17] Benny Applebaum, Naama Haramaty, Yuval Ishai, Eyal Kushilevitz, and Vinod Vaikuntanathan. Low-complexity cryptographic hash functions. In *8th Innovations in Theoretical Computer Science Conference, ITCS*, pages 7:1–7:31, 2017.
- [AS16] Gilad Asharov and Gil Segev. Limits on the power of indistinguishability obfuscation and functional encryption. *SIAM J. Comput.*, 45(6):2117–2176, 2016.
- [Bar01] Boaz Barak. How to go beyond the black-box simulation barrier. In *42nd IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 106–115, 2001.
- [BDRV18] Itay Berman, Akshay Degwekar, Ron D. Rothblum, and Prashant Nalini Vasudevan. Multi-collision resistant hash functions and their applications. In *Advances in Cryptology - EUROCRYPT 2018*, pages 133–161, 2018.
- [BDV17] Nir Bitansky, Akshay Degwekar, and Vinod Vaikuntanathan. Structure vs. hardness through the obfuscation lens. In *Advances in Cryptology - CRYPTO*, pages 696–723, 2017.
- [BKP18] Nir Bitansky, Yael Tauman Kalai, and Omer Paneth. Multi-collision resistance: a paradigm for keyless hash functions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC*, pages 671–684, 2018.
- [Blu81] Manuel Blum. Coin flipping by telephone. In *Advances in Cryptology - CRYPTO*, pages 11–15, 1981.

- [DGRV11] Zeev Dvir, Dan Gutfreund, Guy N. Rothblum, and Salil P. Vadhan. On approximating the entropy of polynomial mappings. In *Innovations in Computer Science - ICS*, pages 460–475, 2011.
- [DI06] Bella Dubrov and Yuval Ishai. On the randomness complexity of efficient sampling. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 711–720, 2006.
- [DPP93] Ivan Damgård, Torben P. Pedersen, and Birgit Pfitzmann. On the existence of statistically hiding bit commitment schemes and fail-stop signatures. In *Advances in Cryptology - CRYPTO '93, 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22-26, 1993, Proceedings*, pages 250–265, 1993.
- [FS90] Uriel Feige and Adi Shamir. Witness indistinguishable and witness hiding protocols. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, STOC*, pages 416–426, 1990.
- [GK96] Oded Goldreich and Ariel Kahan. How to construct constant-round zero-knowledge proof systems for NP. *J. Cryptology*, 9(3):167–190, 1996.
- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, STOC*, pages 218–229, 1987.
- [HHR15] 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 J. Comput.*, 44(1):193–242, 2015.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudo-random generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- [HM96] Shai Halevi and Silvio Micali. Practical and provably-secure commitment schemes from collision-free hashing. In *Advances in Cryptology - CRYPTO '96, 16th Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 1996, Proceedings*, pages 201–215, 1996.
- [HN10] Danny Harnik and Moni Naor. On the compressibility of NP instances and cryptographic applications. *SIAM J. Comput.*, 39(5):1667–1713, 2010.
- [HNO⁺09] Iftach Haitner, Minh-Huyen Nguyen, Shien Jin Ong, Omer Reingold, and Salil P. Vadhan. Statistically hiding commitments and statistical zero-knowledge arguments from any one-way function. *SIAM J. Comput.*, 39(3):1153–1218, 2009.
- [HRVW09] Iftach Haitner, Omer Reingold, Salil P. Vadhan, and Hoeteck Wee. Inaccessible entropy. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC*, pages 611–620, 2009.

- [HRVW18] Iftach Haitner, Omer Reingold, Salil Vadhan, and Hoeteck Wee. Inaccessible entropy I: Inaccessible entropy generators and statistically hiding commitments from one-way functions. www.cs.tau.ac.il/~iftachh/papers/AccessibleEntropy/IE1.pdf, 2018. Preliminary version, named Inaccessible Entropy, appeared in STOC 2009.
- [HV17] Iftach Haitner and Salil Vadhan. *The Many Entropies in One-Way Functions*, pages 159–217. Springer, 2017.
- [IL89] Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography (extended abstract). In *30th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 230–235, 1989.
- [KNY17] Ilan Komargodski, Moni Naor, and Eylon Yogev. White-box vs. black-box complexity of search problems: Ramsey and graph property testing. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 622–632, 2017.
- [KNY18] Ilan Komargodski, Moni Naor, and Eylon Yogev. Collision resistant hashing for paranoids: Dealing with multiple collisions. In *Advances in Cryptology - EUROCRYPT 2018*, pages 162–194, 2018.
- [KY18] Ilan Komargodski and Eylon Yogev. On distributional collision resistant hashing. In *Advances in Cryptology - CRYPTO*, pages 303–327, 2018.
- [MI88] Tsutomu Matsumoto and Hideki Imai. Public quadratic polynomial-tuples for efficient signature-verification and message-encryption. In *Advances in Cryptology - EUROCRYPT '88, Workshop on the Theory and Application of Cryptographic Techniques, Davos, Switzerland, May 25-27, 1988, Proceedings*, pages 419–453, 1988.
- [Nao91] Moni Naor. Bit commitment using pseudorandomness. *J. Cryptology*, 4(2):151–158, 1991.
- [NOVY92] Moni Naor, Rafail Ostrovsky, Ramarathnam Venkatesan, and Moti Yung. Perfect zero-knowledge arguments for NP can be based on general complexity assumptions (extended abstract). In *Advances in Cryptology - CRYPTO*, pages 196–214, 1992.
- [NY89] Moni Naor and Moti Yung. Universal one-way hash functions and their cryptographic applications. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 33–43. ACM, 1989.
- [OV08] Shien Jin Ong and Salil P. Vadhan. An equivalence between zero knowledge and commitments. In *Theory of Cryptography, Fifth Theory of Cryptography Conference, Theory of Cryptography - TCC*, pages 482–500, 2008.
- [OW93] Rafail Ostrovsky and Avi Wigderson. One-way functions are essential for non-trivial zero-knowledge. In *Second Israel Symposium on Theory of Computing Systems, ISTCS*, pages 3–17. IEEE Computer Society, 1993.
- [PR08] Rafael Pass and Alon Rosen. Concurrent nonmalleable commitments. *SIAM J. Comput.*, 37(6):1891–1925, 2008.

- [Sim98] Daniel R. Simon. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In *Advances in Cryptology - EUROCRYPT*, pages 334–345, 1998.