On the Black-Box Complexity of Correlation Intractability

Nico Döttling
CISPA Helmholtz Center for Information Security, Saarbrücken, Germany

Tamer Mour[1]
Bocconi University, Milan, Italy

Abstract

Correlation intractability is an emerging cryptographic paradigm that enabled several recent breakthroughs in establishing soundness of the Fiat-Shamir transform and, consequently, basing non-interactive zero-knowledge proofs and succinct arguments on standard cryptographic assumptions. In a nutshell, a hash family is said to be *correlation intractable* for a class of relations $\mathcal{R}$ if, for any relation $R \in \mathcal{R}$, it is hard given a random hash function $h \leftarrow H$ to find an input z s.t. $(z, h(z)) \in R$, namely a correlation.

Despite substantial progress in constructing correlation intractable hash functions, all constructions known to date are based on highly-structured hardness assumptions and, further, are of complexity scaling with the circuit complexity of the target relation class.

In this work, we initiate the study of the barriers for building correlation intractability. Our main result is a lower bound on the complexity of any black-box construction of CIH from collision resistant hash (CRH), or one-way permutations (OWP), for any sufficiently expressive relation class. In particular, any such construction for a class of relations with circuit complexity t must make at least $\Omega(t)$ invocations of the underlying building block.

We see this as a first step in developing a methodology towards broader lower bounds.

2012 ACM Subject Classification Theory of computation → Computational complexity and cryptography

Keywords and phrases Correlation Intractability, Fiat-Shamir, Black-box Complexity, Black-box Separations

Digital Object Identifier 10.4230/LIPIcs.ITCS.2024.40

Related Version *Full Version*: https://eprint.iacr.org/2023/1365 [24]

Funding *Nico Döttling*: Funded by the European Union (ERC, LACONIC, 101041207). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.
Tamer Mour: Supported by the Binational Science Foundation (Grant No. 2016726), by the European Union Horizon 2020 Research and Innovation Program (ERC Projects REACT – Grant No. 756482, PROMETHEUS – 780701, and FGC – 101019547), and by the Azrieli Fellows Program.

Acknowledgements We thank Zvika Brakerski for comments and discussion, and Venkata Koppula for his helpful feedback on a prior version of this work.

1 Introduction

The *Fiat-Shamir (FS) transform* [5, 25] is a popular technique for eliminating interaction in interactive public-coin protocols. The technique was first conceived to transform 3-round identification protocols into non-interactive signature schemes [25]. Since its introduction,

[1] Work was done while at Weizmann Institute of Science.

15th Innovations in Theoretical Computer Science Conference (ITCS 2024).
Editor: Venkatesan Guruswami; Article No. 40; pp. 40:1–40:24
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

this methodology has had a substantial impact on modern cryptography through several lines of research. The concept gave rise to a number of key innovations in modern cryptography, both for achieving new theoretical feasibility results and for designing communication-efficient practical solutions. In particular, among its noticeable applications are non-interactive zero knowledge protocols [11,12,18,42,51], succinct non-interactive arguments (SNARGs) [6,7,39, 43,47], and complexity-theoretic hardness results [17,40,45].

The basic blueprint of the FS-transform, as laid out in [5], is to transform a (multi-round) public coin protocol by using a hash function H to generate the verifier's public coin messages deterministically based on the protocol transcript so-far.

While it is usually a straight-forward to show that Fiat-Shamir preserves some properties of the original interactive protocol, e.g. completeness and zero-knowledge, it is typically a lot more challenging to show that it preserves soundness using any hash function H. Intuitively, this complication arises as a malicious prover has some control over the computed challenges, e.g. it may just discard a protocol run and retry. In fact, in most constructions the soundness of FS is based on *heuristics.*

More concretely, the soundness of the transformed protocol is often established in an idealized model such as the random oracle model [5]: by modeling the hash function as a random oracle, which both parties have access to, one can prove that the FS transform is sound as long as a cheating prover does not make unreasonably many queries to the oracle. Thus, if the hash function behaves like a random function in the eyes of a bounded adversary, then the non-interactive protocol is sound. The heuristic leap occurs when the random oracle is instantiated by an "unstructured" function such as SHA-2.

Although the random oracle model provides a clean theoretical framework, it is not clear that a sound Fiat-Shamir under the random oracle is a strong enough evidence that provably sound Fiat-Shamir in the plain model exists. In fact, Goldwasser and Kalai [27] show that there exists a computationally sound protocol on which the Fiat-Shamir transform is never sound when instantiated with any actual efficient hash function, even though it is sound in the random oracle model. Further, Bitansky et al. [8] rule out the possibility of constructing a "universal" Fiat-Shamir hash function for all 3-message public-coin protocols based on standard assumptions, or even basing the soundness of Fiat-Shamir for some specific protocols on any falsifiable assumption.

This gap between the conjectured soundness of Fiat-Shamir using "sufficiently unstructured" functions and its provability under cryptographic assumptions in the plain model led Canetti, Goldreich and Halevi [14] to introduce the notion of *Correlation Intractability.* Essentially, correlation intractability captures the computational hardness needed from a Fiat-Shamir hash function in order to prove the soundness of the transform. We say that H is a correlation-intractable hash for a relation class $\mathcal{R}$ (CIH for $\mathcal{R}$) if, for any relation $R \in \mathcal{R}$, it is computationally hard given a random hash key k to find an input x such that $(x, H(k, x)) \in R$. Roughly speaking, in order to show that a Fiat-Shamir instantiation is sound for a given protocol, we require that the underlying hash function is correlation-intractable for the relation between partial protocol transcripts and "bad" verifier challenges that allow for soundness error. Based on this outline, it is known [4,13,42] that a CIH for *all sparse relations* (i.e. relations where any x is in relation with at most a negligible fraction of all y's) is sufficient for Fiat-Shamir over *any* constant-round public-coin proof system (the special case of 3-message protocols has appeared already in [23,31]).

While Canetti et al. [14] show that obtaining correlation intractability in its most general form is impossible, an extensive line of work has eventually led to CIH constructions that are useful for a wide class of protocols, including zero knowledge [12,13,42], statistical

ZAP arguments [2, 29] and, most recently, succinct argument [16, 18, 36, 40, 41]. Overall, the state-of-the-art constructions of CIH are based on advanced well-studied cryptographic primitives which are, in turn, provably secure under standard assumptions such as LWE [46,51] (through special fully-homomorphic commitments or shiftable shift-hiding functions [50]) and DDH [11, 39] (through trapdoor hash functions [21]).

Towards Understanding The Complexity of Correlation Intractability

For a complete comprehension of the notion of correlation intractability, it is fundamental to investigate not only the possibilities but also the limitations in basing correlation intractability on existing hardness notions. In this work, we focus on the relation between correlation intractability and two of the most prominent hardness notions in cryptography: *One-wayness* and *collision-resistance.*

One-way functions (OWF) [20] are functions that are easy to compute but hard to invert. OWFs constitute a central building block in modern cryptography, and were shown to be essential and sufficient for obtaining basic symmetric-key cryptographic notions (a.k.a. Impagliazzo's "Minicrypt" [38]), such as pseudorandom generators [33], pseudorandom functions [26], symmetric encryption [28], commitment schemes [48], zero knowledge [49], and more.

A collision resistant hash family (CRH) is a family of hash functions where it is hard to find collisions under a given random hash sampled from the family. CRH is one of the most widely used primitives in cryptography, with applications ranging from the most basic cryptographic tasks [19,32] to more advanced ones [3,10,43]. Despite its conceptual simplicity, it has been proven that collision resistance cannot be based on one-way functions [53] or even public-key cryptography [1, 9], at least not in a black-box manner.

While, as noted in [15], it is almost trivial that one-way functions imply restricted notions of correlation intractability, such as CIH for all relations $R_a = \{(x, h(x) + a)\}$ (where h is any arbitrary fixed function and addition is over a finite group)[2], such CIH are too weak to realize any interesting applications, in particular Fiat-Shamir for useful protocols. It is also known [34] that exponentially-secure OWF imply *output-intractability*, which is a special case of correlation-intractability for relations R where the membership $(x, y) \in R$ is determined solely by the value of y (but is more general in the sense that it considers tuples of such outputs), and has different applications. In contrast, known useful CIH constructions, for input-output relations, are either based on public-key cryptographic primitives [11, 12, 39, 46, 51], or based on (sub-)exponentially secure OWF and additionally assume the existence of indistinguishability obfuscation (iO) [34, 46].

Whereas the theoretical cryptography literature is rich with proven separations between various cryptographic notions, almost no work had been done on the limitations of reducing correlation intractable hash to other primitives, leaving our understanding of the "reduction complexity" of correlation intractability to be very lacking. The only exceptions are [30], who rule out building the strongest possible form of CIH (for all sparse relations, implying a universal Fiat-Shamir hash) on one-way functions, and [15], who ask whether we can instantiate some specific use-cases of Fiat-Shamir without (or with very weak) cryptography. What we aim for is a more general and accurate picture where we place correlation intractability among the prominent hardness notions in cryptography, specifically one-wayness and collision resistance.

[2] The hash function $H(k, x) = f(x) + h(x) + k$, where f is a OWF, is correlation intractable for $\{R_a\}$. An adversary that breaks the correlation intractability of H for some R_a inverts f at a random image y when given the random key $k = a - y$.

While it is typically beyond our field's current capabilities to rule out general reductions between different hardness notions, a useful framework, that has been developed along the past decades to facilitate reaching meaningful separation results, considers the special case of *fully-black-box* constructions [52], where

(i) the construction makes only black-box use of the underlying primitive, i.e. is oblivious in its implementation, and

(ii) the reduction is assumed to use the provided adversary against the base primitive in a black-box manner.

Such separations are insightful in particular since they already rule out most of the techniques used constructions in the cryptographic literature. The fully-black-box framework has been shown to be extremely fruitful to obtain fundamental separation results, such as separating CRH from OWFs or public-key cryptography [1, 9, 53] and separating key-agreement (and hence, public-key cryptography) from OWFs [37].

Restricting our focus to fully-black-box reductions, we propose the following question to initiate a thorough study of the complexity of correlation intractability:

What is the black-box complexity of correlation intractable hashing from CRH?

We believe collision resistance is a natural starting point in this general direction as it is a sufficiently simple and basic notion to constitute a first step towards broader research. There are two items to note here: First, as collision resistance implies OWFs (in a fully-black-box manner), any answer for the above question would immediately imply a similar statement for constructions from OWFs. Second, CRH are a special case of *multi-input correlation intractability*, which is a generalization of correlation intractability where it is hard to find a *tuple* of inputs that satisfy some relation between themselves and their images under the hash (in standard CI relations are over a single input-output pair). While general multi-input CI is clearly stronger than regular CI and implies it, what we ask above is whether multi-input CI for a specific natural (multi-input) relation can be useful to build CI for a more general class of (single-input) relations, that is – whether "multiplicity" of the relation class can be exchanged for "expressiveness".

1.1 Our Results

In this work, we explore inherent limitations in constructing correlation intractable hash functions and initiate the study of the black-box complexity of correlation intractability. We draw the following connection between the complexity of any fully-black-box construction of CIH from CRH or OWP and the complexity of the relations we get correlation intractability for.

▶ **Theorem 1** (Black-box Complexity of CIH from CRH or OWP; Informal). *Any fully-black-box construction of correlation intractable hash for any t-wise independent class of relations from collision-resistant hash, or one-way permutations, must make at least $O(t)$ calls to the underlying base primitive(s).*

A t-wise independent class of relations $\mathcal{R}$ is class of relations where for any t pairs (z_1, w_1), $\ldots$, (z_t, w_t), the events $\{(z_i, w_i) \in R\}$ for a random relation $R \leftarrow \mathcal{R}$ are all independent. One example is the class of all relations searchable by degree t polynomials, i.e. any relation consisting of all pairs $(z, p(z))$ for a degree t polyonomial p specified by the relation. Consequently, as polynomials of degree t can be computed by (arithmetic) circuits of size t, we get that the class of relations searchable by t-bounded circuits is $\Omega(t)$-wise independent. Hence, the degree of independence provides a meaningful proxy for the complexity of a class of relations. To give some sense, CIH suitable for cryptographic applications, such as NIZKs or succinct non-interactive arguments, as far as we know requires intractability for relations with complexity proportional to the security parameter.

Our result carries a couple of caveats. First, it holds only for fully-black-box constructions. Although this is already insightful and captures many of the existing and imaginable techniques, there might exist non fully-black-box constructions that circumvent this impossibility. Second, this is not an absolute separation in the sense that it does not entirely rule out building one primitive from another, rather it only sets a lower bound on the efficiency of such constructions. While such a result is partial in nature, we believe that the analysis underlying the proof provides many insights regarding the essence of correlation intractability and its complexity, potentially leading to future work advancing our understanding further, through stronger separations and even new constructions. We elaborate below.

1.2 Discussion and Open Questions

We view our result as initiating the research on the complexity of correlation intractable hashing. While our bottom-line yields a lower bound that is far from what is known or even believed to be possible, our hope is that the techniques and observations introduced in our analysis will eventually lead to a better understanding of the notion of correlation intractability.

For instance, it may not be unlikely that, with some additional effort and insights, our proof can be extended to achieve a similar impossibility for CIH from public-key cryptography; In the work of [9] by which our initial ideas were inspired, they are able to show separation of CRH not only from OWPs but also from the combination of OWPs with iO, which, in particular, implies a separation from public-key encryption (PKE) [3]. While extending our result in an analogous manner is doomed to fail due to the existence of (fully black-box) CIH based on OWP and iO, demonstrated in [46], we expect that a more careful adaptation of the developed ideas has the potential to yield a separation from PKE.

A more intriguing direction is to investigate the gap between our limited separation result, that does not entirely rule out constructions of CIH from CRH or OWP, and the state-of-the-art CIH constructions which are known from building blocks that are much more complex. We see it is important to understand whether it is merely an artifact of our proof technique that we were not able to extend it to rule out constructions for *any* (non-trivial) class of relations or whether there is an inherent barrier in proving such a separation. In particular, one may ask

Is it indeed impossible to build non-trivial CIH in Minicrypt [38] (or Hashomania [44])?
Which relation classes can we get CIH for, based on OWFs or CRH?

1.3 Technical Overview

We will now discuss the ideas behind our main result, Theorem 1, which states that any fully-black-box construction of CIH from CRH, or OWP for relations of complexity t (more accurately, that are t-wise independent) must make $\Omega(t)$ invocations of the underlying base primitive(s).

The starting point of our proof is the work of Bitansky and Degwekar [9], which provides a separation of collision-resistant hash functions (CRH) from one-way permutations (OWP). We generalize their framework to the correlation intractability setting and further extend

[3] This, in fact, was first established in [1]

it to capture the separation from CRH. Along the way, we introduce a new notion which facilitates establishing hardness under oracles (e.g. of inversion or finding collisions) which we call *differential indistinguishability.* Proving oracle-relative hardness is always at the core of separations of this theme since, typically, the underlying cryptographic primitive, which is accessible only in a black-box manner, is modelled as an oracle that provably satisfies the corresponding intractability property. Interestingly, through differential indistinguishability, we show how to use techniques resembling those from the differential privacy literature in order to obtain traditional cryptographic hardness relative to an oracle.

For the sake of this overview, we outline the lower bound on CIH constructions from one-way permutations and then briefly discuss how the underlying techniques can be further expanded to obtain the lower bound on constructions from CRH.

Let us first recall the fully black-box separation framework which we follow in this work.

Fully Black-box Separations and How to Prove Them

We say that a construction P of a cryptographic primitive $\mathbf{P}$ from a different primitive $\mathbf{Q}$ is *fully black-box* [52] if

(i) the construction makes only black-box use of $\mathbf{Q}$ (that is, any *instantiation* of $\mathbf{Q}$, independently of its implementation) and, further,

(ii) there is a black-box security reduction $\mathcal{M}$ which, given a black-box access to *any* adversary $\mathcal{A}$ that breaks P, breaks the underlying instantiation of $\mathbf{Q}$.

A *fully black-box separation* of $\mathbf{P}$ from $\mathbf{Q}$ simply means that fully black-box constructions of $\mathbf{P}$ from $\mathbf{Q}$ are impossible. In many cases, such as ours, conditioned separations are considered, namely, where it is only argued that fully-black-box constructions that satisfy certain constraints (e.g. efficiency) are impossible.

Similarly to prior work on fully-black-box separations, we follow the "Two-Oracle Methodology" [1,9,35,53] where, to show that is it impossible to build correlation intractable hash from another primitive $\mathbf{Q}$, e.g. OWP or CRH – again, possibly assuming certain efficiency constraints – it is shown that there exists an oracle Q, which models an idealized implementation of $\mathbf{Q}$, and an oracle that models an adversary against correlation intractability, namely, a *correlation finder* CF, such that

(i) CF breaks any black-box construction of CIH from Q that satisfies the presumed constraints, yet,

(ii) Q is still secure, as per the security definition of $\mathbf{Q}$, in the presence of CF.

Given that such oracles Q and CF exist, any fully black-box reduction $\mathcal{M}$ fails in breaking Q using CF and, hence, no fully black-box construction of CIH from $\mathbf{Q}$ exists.

We first focus on separating CIH constructions from OWP, as this captures many of the key concepts in the extended result, and only later discuss how to further derive a separation from CRH.

The Challenge in Designing a Correlation Finder

We model our "ideal" OWP via a random permutation oracle $f : \{0,1\}^\lambda \to \{0,1\}^\lambda$. While it is straight-forward to show that it is infeasible to invert a random permutation at a random image given bounded black-box access, our goal is to show this is still infeasible even given access to a successful correlation finder CF. The correlation finder CF takes as input a circuit $C \in H$ describing a hash function with oracle access to f. Here we can think of the set H as abstracting away from the keys in a hash construction. In essence the set H limits the adversary's choice.

One natural way to show that any bounded reduction $\mathcal{M}$ still cannot invert f given CF is to show that $\mathcal{M}$ is able to simulate, with little extra cost, any useful information it receives from CF by itself alone, making the correlation finder redundant and using the one-wayness of a random permutation to complete the proof. This approach has been successful to separate, in particular, CRH from OWP [9,53]. An elementary reason is that, under any (sufficiently shrinking) CRH, the marginal distribution of any "half" of a uniformly random collision, is almost uniform. Let $C \in H$ be a circuit with oracle-access to f describing a hash function. Thus, when letting the collision finder, on input such a circuit C^f, simply output a random collision (z, z') (s.t. $C^f(z) = C^f(z')$ and $z \neq z'$), the reduction can simulate the marginals of each of z and z' without the help of the collision-finder. Roughly speaking, as has been shown particularly in [9], the marginals capture the only "useful" information the reduction can obtain for inverting f.

Things are not that simple, however, when the goal of the oracle is to return a correlation. Here, the correlation finder CF_R^f depends on a relation R as well as f. We will omit R and f when the context is clear. On input a circuit $C \in H$ the correlation finder $\mathsf{CF}_R^f(C)$ should produce an input z such that $(z, C^f(z)) \in R$. In this setting, a reduction $\mathcal{M}$ may produce a query to CF where all possible correlations under a chosen relation R, i.e. all "correct" answers that CF may possibly return, coincide with a set of inputs that is most useful for inverting f at any given point. For example, we may think of an $\mathcal{M}$ that, given a challenge y, calls $\mathsf{CF}(C)$ where C is the hash circuit that on any input z, outputs a w s.t. $(z, w) \in R$ if and only if $f(z) = y$ (it is reasonable to assume that such a w is efficiently computable)[4]. For this C, there is only one such z satisfying $(z, C^f(z)) \in R$, an hence $CF(C)$ must return this z.

Picky Correlation Finder

Given the inherent tension between correctness of the correlation finder and its usefulness for inverting f, we propose the following way out. We design CF to be *imperfect*, that is, to return a correct answer, say a uniformly random correlation, for most inputs while rejecting to do so for the others. The distinction between functions on which CF may "cooporate" and functions on which CF must reject is made possible by the fact that, in order to break correlation intractability, CF must succeed on some relation R only for an *average-case* hash $\tilde{C} \leftarrow H$. Thus, in the CIH game, which one can think of as the "honest" case, the circuit $\tilde{C}$ that computes the hash function is independent of R and should not exhibit any extraordinary behavior w.r.t. correlations under R. On the other hand, if $\mathcal{M}$ attempts to abuse CF to invert f, then it must produce a "malicious" circuit C which is specifically tailored to be useful for inversion and, therefore, as we argue in our proof, must highly "depend" on R.

Thus, we need to construct a correlation finder CF that is able to tell when a circuit C is likely to be malicious, yet does not overshoot as it still needs to answer for an honest C. To that end, we articulate a measure of "extraordinariness" that captures "usefulness" for inversion, which, roughly speaking, happens to be tightly related to the Rényi divergence of infinite order (this can be thought of as an analog of KL-divergence for min-entropy) between what useful information is obtained from CF and what useful information can be simulatable without CF. We let our CF reject any circuit C that is extraordinary w.r.t. R to obtain a *picky correlation finder*, namely a correlation finder that is successful only with high probability.

[4] We are implicitly assuming that the input spaces for f and C are the same. When this is not the case, the reduction can use any arbitrary 1-1 mapping that maps any C-input w to a corresponding f-input x_w and the implication still holds.

Detecting Malicious CF-Queries

In the following let $\mathsf{Corr}^f_{R,C}$ denote the set of all z which satisfy $(z, C^f(z)) \in R$. To identify what "useful information" is w.r.t. inverting one-way permutations, we take inspiration from [9], where they implicitly show that to invert an oracle f at a random image y^*, it is necessary for the reduction $\mathcal{M}$ to be able to distinguish between black-box access to f and black-box access to a different permutation $f' = f_{x^* \leftrightarrow x'}$ that is obtained from f by swapping the solution pre-image $x^* = f^{-1}(y^*)$ with a uniformly random x'. Given this, a malicious query to CF is then a circuit C for which the distribution of a non-rejecting $\mathsf{CF}^f(C)$, namely the correlation finder's answer to C under f can be distinguished from a $\mathsf{CF}^{f'}(C)$, namely its answer under f'. Only using such malicious queries, the reduction $\mathcal{M}$ can use CF to distinguish between f and f' and, thus, invert f. We observe that C can induce such two distinguishable distributions under functions f and f' only if the swap $x^* \leftrightarrow x'$ significantly affects the set of correlations, from which CF samples its answer. This may occur only if, given a random correlation $z \leftarrow \mathsf{Corr}^f_{R,C}$, the hash function $C^f(z)$ calls any of x^*, x' with noticeable probability or, in other words, only if any of x^*, x' are *heavy among correlations*. Hence, our correlation finder should, in particular, look at the weight of any worst-case x w.r.t. the given query C and the chosen relation R (it is crucial to note that CF has no knowledge of x^*, x' as they exist only in the inversion game and its analysis), which we define as

$$\omega^f_{R,C}(x) = \Pr_{z \leftarrow \mathsf{Corr}^f_{R,C}}[C^f(z) \rightharpoondown x], \tag{1}$$

where $C^f(z) \rightharpoondown x$ denotes the event that the computation $C^f(z)$ calls f at x. This alone is not sufficient, however, since it may be the case that there are heavy inputs also under an honest query C, that does not depend on the relation R. For instance, consider a CIH construction C that *always* calls f at some fixed x_0, regardless of its input being a correlation or not. Then, in such case we have that $\max_x \omega(x)$ takes its maximal value 1 and the correlation finder always rejects and is, therefore, never successful. It is clear that such a query to CF cannot possibly be helpful to invert f since, intuitively speaking, any information that $\mathcal{M}$ may extract from $\mathsf{CF}^f(C)$ regarding the image of the heavy input x_0 it could already extract without calling CF by calling C at random inputs (that are not necessarily a correlation). Keeping our initial outline in mind, we are interested in the *relative* "usefulness" of information obtained form CF compared to information simulatable without CF's help. We refine our basic idea to consider the *relative* weight of any worst-case x among correlations compared to its weight in the entire input space. For that, we define the *scale* of any input x as

$$\sigma^f_C(x) = \Pr_{z \leftarrow \{0,1\}^m}[C^f(z) \rightharpoondown x], \tag{2}$$

and look at the *amplification* in the likelihood of observing x in an execution $C^f(z)$ due to restricting z to be a random correlation compared to being a random input (i.e., a random CF answer compared to a random input which is simulatable without CF), that is,

$$\alpha^f_{R,C}(x) = \omega^f_{R,C}(x) / \sigma^f_C(x). \tag{3}$$

We let CF reject only if there exists an x for which $\alpha^f(x) \gg 1$. This gives us a CF that is successful in the honest case since one can easily see that $\alpha^f(x)$ has an average of 1 when the relation R is sampled independently in C (further, a sufficient tail bound for worst-case $\alpha^f(x)$ can be derived already when R is pairwise independent). On the other hand, CF rejects whenever x^* or x' are too heavy among correlations since $\sigma^f_C(x^*)$ and $\sigma^f_C(x')$ can be assumed to be small: x' is a random input that is sampled in the analysis and is independent in the reduction's choice of C, while x^* is the solution pre-image and, had it been heavy among random inputs, the reduction would have been able to observe it by sampling random inputs to C without the help of CF.

As already mentioned, it turns out that $\max_x \alpha^f(x)$ is precisely the Rényi divergence of order infinity between the distributions over the f-input space induced by the PDFs ω and σ. In Figure 1 we visualize the distinction between honest queries, which give low divergence between ω and σ, and malicious queries.

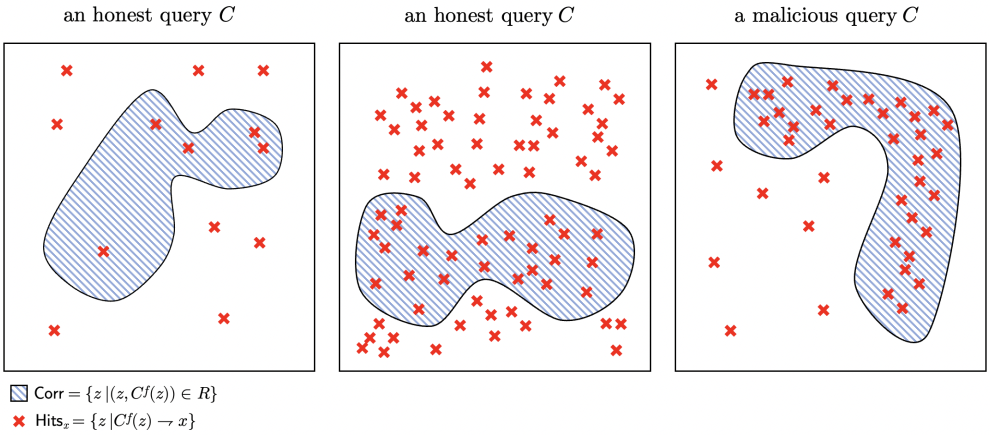

an honest query C

an honest query C

a malicious query C

▨ $\mathsf{Corr} = \{z \,|(z, C^f(z)) \in R\}$
✖ $\mathsf{Hits}_x = \{z \,|C^f(z) \to x\}$

Figure 1 An illustration of the domain of the circuit C when C is an honest query vs. when it is a malicious one. Notice that C is malicious only if some x is observed with much higher likelihood in the correlation set, compared with the entire space.

While our efforts so far already constitute a major step towards a separation, a new problem pops up: the decision for rejection, namely whether CF rejects or not, might itself give information that is useful for inverting f! This is not merely a hypothetical scenario; One can show, in fact, a concrete "attack" against such CF, namely a reduction $\mathcal{M}$ that, while unable to learn anything from the non-rejection answers, can learn the pre-image entirely based on the rejection decisions of CF (note each such decision conveys at most a single bit of information).

Differential Indistinguishability

To solve the issue raised above, we encapsulate the "usefulness" (or "uselessness") of a correlation finder in inverting f via an novel oracle-relative hardness notion, we call *differential indistinguishability*. At a high level, we say that CF is differentially indistinguishable if its answers and, in particular, its *rejection decision*, do not substantially change when the function f is modified locally (specifically, when two inputs are swapped under the permutation as described above). As already mentioned, the notion of differential indistinguishability is implicitly used in the work of [9] and is inspired by their proof. However, the straight-forward collision finder, that returns a random collision, already satisfies differential indistinguishability and, therefore, the main effort in their proof goes to show that differential indistinguishability implies "uselessness" for inversion. One of the contributions in our work is formalizing then extending the implicit framework from [9] to capture correlation finders (in fact, any oracle) and additionally show a similar implication regarding "uselessness" for finding collisions, through which we derive the separation from CRH.

A Differentially Indistinguishable Rejection Policy

Given that differential indistinguishability implies that f is hard to invert given CF, it remains to design a differentially indistinguishable CF. To that end, we further refine the picky correlation finder from above and design a "soft" rejection policy that is robust against local changes of the function f. Unsurprisingly, our mechanism for achieving this looks, in retrospect, as if taken from the world of differential privacy (DP) [22]. In more details, we consider a rejection policy that takes into account not only extraordinary behaviour, namely large $\alpha^f(x)$, w.r.t. the given function f, but also w.r.t. functions f' that are in the "neighborhood" of f, namely functions that can be obtained by performing a limited series of swaps on f. By weighing in the impact of any *extraordinary behaviour* in a way that quickly vanishes with increasing distance from f, e.g. an weight function which drops exponentially with the distance from f, we are able to derive differential indistinguishability. Overall, our rejection policy, namely the probability that CF rejects at some circuit query C, is computed by a function similar to the following

$$\rho_R^f(C) = \epsilon \cdot \max_{f',x}[e^{-\Delta(f,f')/d} \cdot \alpha_{R,C}^{f'}(x)], \tag{4}$$

where $\Delta(f,f')$ denotes the *swap distance* [5] between f and f', and $\epsilon \ll 1$ and $d \gg 1$ are carefully chosen normalization parameters; The larger d and ϵ are, the stronger is the differential indistinguishability guarantee, yet the harder it is to show that CF is successful in the honest case. For the sake of this overview, d can be thought of as superpolynomial in the security parameter and ϵ as inverse-superpolynomial.

Having safeguarded intractability of inverting f via differential indistinguishability, we need to make sure we have not broken the subtle balance with the necessary correctness requirement on CF. While it is almost immediate that any circuit C (modelling a random hash function $\tilde{C} \leftarrow H$ sampled from the CIH candidate) behaves "nicely" under f w.r.t. an independently chosen relation $R \leftarrow \mathcal{R}$, even when relations in $\mathcal{R}$ are only pairwise-independent, it is not clear that this holds under *all* functions f' in the close neighborhoods of f. The straight-forward attempt to apply a union bound over all possible functions in the neighborhood inherently requires that relations in the class have a description of exponential size, which would dramatically weaken our result. Instead, we exploit the fact that the candidate CIH is bounded to make t invocations of f and observe strong dependencies between the $\alpha(x)$ values under different functions across the neighborhood of f. More specifically, since the behavior of C on any input is determined by the images of at most t points in f, we notice that we can represent $\alpha^{f'}(x)$, under any f', as the average of a collection of values $\{\hat{\alpha}^{f'}(x)\}$ where each $\hat{\alpha}^{f'}(x)$ depends solely on $t \ll d$ points in f'. This allows us to apply a much more benign union bound to establish that none of the possible $\hat{\alpha}^{f'}(x)$ is too large. Consequently, we are able to argue that if the relation class is sufficiently "expressive", namely $\Omega(t)$-wise independent, then when the relation R is chosen at random and independently in the hash circuit C, C is indeed behaving well under *any* f' that is sufficiently close to f. This implies that an honest query C is rejected with low probability and, therefore, CF is correct and we have a separation of CIH from OWP.

[5] We define the swap distance between permutations f and f' as the minimal number of times we need to swap outputs between input pairs x_1 and x_2 in order to transform f into f'

Extending to Constructions from Collision Resistance

We will now discuss how these results can be extended to capture constructions from CRH as well. At the core of our extended result is the observation that the task of finding collisions in an oracle can be conceptually thought of as an "adaptive" inversion task; To find a collision in an oracle f, an adversary must invert an image y for which he had already seen a pre-image under f. The difference from breaking the one-wayness of f, namely inverting f at a random image, is clear: In the collision-finding game, the adversary, in some sense, chooses the images he aims to invert. Hence the "adaptiveness".

Recall that in order to establish a separation of CIH from OWP, we

(i) define a notion of differential indistinguishability and prove a random f is hard to invert even given a differentially indistinguishable correlation finder, and

(ii) construct a differentially indistinguishable correlation finder (that is successful in the honest case).

Based on the above insight, we propose a notion of *adaptive differential indistinguishably*, then prove that it is sufficient to imply collision-resistance of f under the correlation finder were the latter to satisfy it. Lastly, we show how to generalize our construction of correlation finder from above to satisfy the new adaptive notion and, by this, finish. We elaborate below.

Re-randomizing Siblings

The reason that differential indistinguishability is sufficient to imply hardness of inversion is that it allows us to re-randomize the target pre-image (that is, swap the original x^* whose image is given as a challenge with a random x') without the adversary noticing that he is given access to a different function $f' = f_{x^* \leftrightarrow x'}$. Thus, the probability that the adversary returns x^* under f is equal to the probability he returns x^* under f', which carries no information about x^*, and therefore he cannot do better than guessing. To adapt this idea to an adversary that is trying to find collisions, we re-randomize, as hinted above, the *siblings* of any input x at which the adversary calls f in his executions. The siblings of any x under f are all x' that make a collision with x, i.e. all $x' \neq x$ such that $f(x) = f(x')$. Roughly speaking, since we may assume w.l.o.g. that the adversary calls f at a collision the moment he finds it, then we can assume that a successful adversary must call a sibling of a previously queried input. It would not be sufficient, however, to re-randomize, at every step of the execution, only siblings from previous queries since, hypothetically, the adversary's strategy might be to collect information about "future" siblings, namely siblings of some x before actually making the query to f at x. We must therefore re-randomize *all* siblings induced by the execution at *any* of its steps.

Adaptive Differential Indistinguishability

An inherent difference from the OWP case then arises: In proving intractability of inverting an OWP f, we re-randomize a pre-image x^* which is fixed apriori to the execution of the adversary. In contrast, when re-randomizing siblings, specifically "future" siblings, we are re-randomizing pre-images that are implicitly determined by the adversary's execution. This difference motivates us to define an adaptive analog of differential indistinguishability, where, in a high level, we require that the answers of CF do not change when the function f is swapped even at points chosen adaptively in the answers themselves (essentially, the answers are what constitute the view of the adversary on which he bases his choice of siblings). The new adaptive notion introduces various non-trivial subtleties. For instance, unlike its non-adaptive counterpart, adaptive differential indistinguishability against any general choice of swap sets is impossible to realize while preserving correctness of the correlation finder. To see this, consider an adaptive choice that, given an answer on some query C, namely a correlation $z \leftarrow \mathsf{CF}^f(C)$, chooses to swap the function f at an input x that is called by $C^f(z)$. Swapping x possibly changes the outcome of the computation $C^f(z)$ making z no longer a correlation under the modified function f' and, therefore, no longer a "correct" answer for z. A successful CF will most likely not output z in such a case, practically implying that such a modification of f must cause CF to answer differently with high probability.

Fortunately, we are able to show that, unless our adaptive choice is that "targeted" (that is, chooses to swap inputs that specifically appear in the execution of the query circuit C on CF's answer), then adaptive differential indistinguishability is achievable. On the other hand, we prove that the choice to swap the set of siblings is never such a "targeted" choice under one condition in particular: that the execution of C on the answer $z \leftarrow \mathsf{CF}^f(C)$ does not observe a collision, namely does not make two f-queries that collide, with high probability over CF's randomness. Through these observations, we are able to generalize our correlation finder from above to satisfy the required adaptive notion against any choice of siblings. In particular, our new correlation finder looks at an analog of the amplification values $\alpha = \omega/\sigma$ that we define for *pairs* of inputs and, further, for the "soft" rejection considers the neighborhood of functions that are obtained by swapping between *sets* of inputs (rather than individual points). Overall, we get a correlation finder that is adaptively differentially indistinguishable against siblings, implying a similar separation of correlation intractable hash from collision resistance.

1.4 Technical Notation

We introduce basic notation to be used in the paper. For a distribution $\mathcal{X}$, we write $x \in \mathcal{X}$ to say that x is in the support of $\mathcal{X}$, and $x \leftarrow \mathcal{X}$ to denote that x is sampled from the distribution $\mathcal{X}$. We overload the notation for sets and write $x \leftarrow X$ when x is sampled uniformly at random from a set X. We use $\mathbb{P}(X)$ to denote the power set of X. For an event E, we use $\mathbb{1}(\mathsf{E})$ to denote the binary value which takes 1 if and only if E occurs. $\mathbf{SD}(\mathcal{X}, \mathcal{Y})$ denotes statistical distance between distributions X and Y. For an oracle-aided algorithm $\mathcal{A}$, an oracle Ψ, and a Ψ-input z, we denote by $\mathcal{A}^\Psi(x) \xrightarrow{\Psi} Q$ the event where $\mathcal{A}$, on input x, calls the oracle Ψ at Q. We extend this notation for tuples of Ψ-inputs: $\mathcal{A}^\Psi(x) \xrightarrow{\Psi} Q_1, \ldots, Q_n$ if $\mathcal{A}^\Psi(x)$ calls Q_i for *all* i.

1.5 Paper Organization

Due to space limitation, in what follows in this version of the paper, we focus only on the separation result of CIH from OWP and present a series of arguments and abstractions through which we establish the result. The full version of the paper [24] contains all necessary proofs, as well as the presentation of the stronger separation result from CRH, which relies on a generalization of many of the ideas presented hereby.

In Section 2 we define the notions of fully black-box constructions and separations then formally state our results. In Section 3 we present a generic framework for proving bounds on constructions of correlation intractability via the notion of differential indistinguishability and, in Section 4 we build a differentially indistinguishable correlation finder that satisfies our requirements, allowing us to derive the main theorem separating CIH from OWP.

2 Our Results: Statement of Main Theorems

In this section we formally state our separation results of correlation intractability from CRH and OWPs. Let us first define a fully black-box construction of CIH from OWP.

▶ **Definition 2** (Fully Black-box Construction of CIH from OWP). *Let $m := m(n)$ be a length parameter and let $\mathcal{R}$ be a class of relations. A (t, q, ϵ)-*fully black-box construction of correlation intractable hash (CIH) for $\mathcal{R}$ with input length m from one-way permutations (OWP)*, for $t := t(n)$, $q := q(\lambda)$ and $\epsilon := \epsilon(\lambda)$, is an ensemble of distributions $H = \{H_n\}$ where, for any $n \in \mathbb{N}$, H_n is a distribution over functions mapping m-bit inputs to n-bit outputs, and an oracle-aided reduction $\mathcal{M}$ satisfying the following properties:*

- ***Construction Efficiency:*** *For any $n \in \mathbb{N}$ and any $h \in H_n$, h^f makes at most $t(n)$ queries to f on any input.*
- ***Black-box Security Reduction:*** *For any oracle $f = \{f_\lambda : \{0,1\}^\lambda \to \{0,1\}^\lambda\}$ and any probabilistic oracle-aided adversary $\mathcal{A}$, if there exists a relation $R \in \mathcal{R}$ such that*

$$\mathbf{Adv}^{H}_{\mathbf{CI}}(n, R, \mathcal{A}) := \Pr_{\substack{h \leftarrow H_n \\ z \leftarrow \mathcal{A}^f(1^n, h)}} [(z, h^f(z)) \in R] > \frac{1}{2}$$

for infinitely many $n \in \mathbb{N}$, then,

$$\mathbf{Adv}^{f}_{\mathbf{OWP}}(\lambda, \mathcal{M}^{\mathcal{A}}) := \Pr_{x \leftarrow \{0,1\}^\lambda} [\mathcal{M}^{f,\mathcal{A}}(f_\lambda(x)) = x] \geq \epsilon(\lambda)$$

for infinitely many $\lambda \in \mathbb{N}$.

- ***Reduction Efficiency:*** *For any $\lambda \in \mathbb{N}$ and $y \in \{0,1\}^\lambda$, $\mathcal{M}^{f,\mathcal{A}}(y)$ makes at most $q(\lambda)$ queries to the oracles f and $\mathcal{A}$, and for every $\mathcal{A}$-query $(1^n, h)$ made by $\mathcal{M}(y)$, it holds that $n < 2^{\lambda/24}$ and $h^f(\cdot)$ makes at most $q(\lambda)$ queries to f on any input.*

A *fully black-box construction of CIH from CRH* is defined along similar lines, except that the reduction aims at finding collisions in a shrinking oracle f. A full definition is given in the full version [24].

Lastly, we define a fully black-box α-separation to embody the impossibility of any fully black-box construction abiding a trade-off (parameterized by α) between the complexity of the underlying reduction and its success probability. A larger value of α gives stronger separation and, in particular, superpolynomial α indicates the impossibility of a reduction that is both polynomial time and has non-negligible advantage, as typically required in the traditional cryptographic setting.

▶ **Definition 3** (Black-box Separation of CIH from OWP (or CRH)). *Let $m := m(n)$ be a length parameter and let $\mathcal{R}$ be a class of relations. We say that t-*bounded CIH functions for $\mathcal{R}$ (with input length m) are $\alpha(\lambda)$-fully black-box separated *from OWP (or CRH), for $t := t(n)$ and $\alpha(\lambda) > 1$, if for any (t, q, ϵ)-fully black-box construction of such CIH from OWP (or, resp., CRH), it holds that either $q(\lambda) > O(\alpha(\lambda))$, or $\epsilon(\lambda) \leq O(1/\alpha(\lambda))$.*

Our impossibility results rule out any construction of a CIH for relation classes that, roughly speaking, constitute complexity greater than the black-box complexity of the construction. To articulate the complexity of a given relation class $\mathcal{R}$, we refer to the degree of "unpredictability" induced by a random relation; We say that $\mathcal{R}$ is k-wise universal if, in particular, the likelihood of any (z, w) to be in a random relation $R \leftarrow \mathcal{R}$ does not change even given the membership (or non-membership) in R of any $k - 1$ pairs.

▶ **Definition 4** (k-wise Universal Relations). *Let $k : \mathbb{N} \to \mathbb{N}$ and $p : \mathbb{N} \to (0,1)$. We say that a relation class $\mathcal{R} = \{\mathcal{R}_n \subseteq \mathbb{P}(\{0,1\}^m \times \{0,1\}^n)\}$, for $m := m(n)$ is* k-wise p-universal *if, for any $n \in \mathbb{N}$, there exists a distribution over relations in $\mathcal{R}_n$ (which we ambiguously denote by $\mathcal{R}_n$) such that for any $k' \leq k$ and any distinct $(z_1, w_1), \ldots, (z_{k'}, w_{k'}) \in \{0,1\}^m \times \{0,1\}^n$, it holds that*

$$\Pr_{R \leftarrow \mathcal{R}_n}[(z_i, w_i) \in R \quad \forall i \in [k']] = p(n)^{k'(n)}.$$

We now formally state our main separation theorems: In any fully-black-box construction of CIH against a k-wise independent relation class from CRH (or OWP) with non-trivial security, the hash function must invoke the underlying CRH (or OWP) at least $\Omega(k)$ times in its computation. We provide the formal theorems with accurate quantitative details below.

▶ **Theorem 5** (Black-box Separation of CIH from OWP). *Let $m := m(n)$ and $p : \mathbb{N} \to [0,1]$ be such that $p(n) \geq 4n^2 2^{-m(n)}$ and let $k, t : \mathbb{N} \to \mathbb{N}$ be such that $k(n) > 20 \cdot t(n)$ for all $n \in \mathbb{N}$. Then, t-bounded CIH functions, with input length m, for any class of k-wise p-universal relations are $2^{\lambda/10}$-fully black-box separated from OWP.*

▶ **Theorem 6** (Black-box Separation of CIH from CRH). *Let $m := m(n)$ and $p : \mathbb{N} \to [0,1]$ be such that $p(n) \geq 4n^2 2^{-m(n)}$ and let $k, t : \mathbb{N} \to \mathbb{N}$ be such that $k(n) > 25 \cdot t(n)$ for all $n \in \mathbb{N}$. Then, t-bounded CIH functions, with input length m, for any class of k-wise p-universal relations are $2^{\lambda/25}$-fully black-box separated from CRH mapping $\lambda + O(1)$ bits to λ bits.*

Note that our separation result from collision-resistance considers CRH that shrinks its input only by a constant number of bits. We stress, however, that our proof technique results in equally-merited separations from any CRH with constant multiplicative shrinkage smaller than $\frac{1}{2}$, where we still require $k = \Omega(t)$ and obtain a $2^{\Omega(\lambda)}$-separation. Since such a CRH implies CRH with any polynomial shrinkage via a logarithmic number of sequential invocations, one may derive more general separation results with corresponding parameters.

3 A Generic Framework: One-Wayness under Differentially Indistinguishable Correlation Finder

We introduce a generic framework for showing barriers on CIH constructions. Our approach builds on the "Two-Oracle Methodology" [1, 9, 35, 53] where, in order to obtain bounds on cryptographic constructions, one creates an idealized (oracle-relative) world under which such constructions are impossible. In our case, such a world would consist mainly of an ideal oracle representation of a cryptographic primitive (be it a CRH or OWP) and a correlation finder that should be able to break any construction of CIH from the ideal primitive that satisfies certain constraints, e.g. query complexity, yet is useless for breaking the intractability of the underlying oracle (that is, inverting it or finding induced collision).

Among our contributions is the formulation of a somewhat unified hardness notion, namely differential indistinguishability, and show that any correlation finder that satisfies it is indeed useless breaking the ideal OWP or CRH. We believe that our approach is sufficiently modular to allow for adaptation in different settings. For this outline, we focus on separating CIH from OWP. In the full version [24], we give a generalization of the framework that enables the separation result from CRH in Theorem 6.

Setting and Notation

Our proof will be centered around two “computational games”: In the first, a correlation finder aims to break the correlation intractability of a candidate CIH that maps $m := m(n)$ bits to n bits (n can be thought of as the “security parameter” in this game). In the second game, an adversary is given access to the correlation finder and aims to break the intractability of an idealized OWP that is given as a permutation oracle over λ-bit inputs (here, λ is the security parameter). We now list the main playing parts in this settings:

- A *relation class* $\mathcal{R} = \{\mathcal{R}_n\}$ where, for any $n \in \mathbb{N}$, $\mathcal{R}_n$ is a class of relations over $\{0,1\}^{m(n)} \times \{0,1\}^n$. This will denote a relation class which we seek to build (actually, rule out) correlation intractability for.
- A (random) oracle $\mathcal{F} = \{\mathcal{F}_\lambda\}$ where for any $\lambda \in \mathbb{N}$, $\mathcal{F}_\lambda$ is the uniform distribution of permutations over λ-bit inputs. We typically use f to denote a function chosen from $\mathcal{F}$. An algorithm is f-aided if it is given access to an oracle with the syntax of $f \in \mathcal{F}$.
- The family of f-aided circuits $\mathcal{C} = \{\mathcal{C}_n\}$ where, for any $n \in \mathbb{N}$, $\mathcal{C}_n$ is the set of all f-aided circuits mapping $m(n)$-bit inputs to n-bit inputs. In particular, a *CIH candidate for* $\mathcal{R}$ (from OWP) is an ensemble of f-aided circuits $C = \{C_n\}$ where $C_n \in \mathcal{C}_n$ for all $n \in \mathbb{N}$.
- A *correlation finder* $\mathcal{O} = \{\mathcal{O}_R\}$ where, for any $R \in \mathcal{R}$, $\mathcal{O}_R$ is a distribution over f-aided oracles that on input 1^n and an f-aided circuit $C \in \mathcal{C}_n$ (for any $n \in \mathbb{N}$) outputs a C-input of length $m(n)$ bits (which should be a correlation w.r.t. R if successful). We often omit the input 1^n as it is clearly determined by C and sometimes omit the relation R when it is irrelevant in the context.
- An *oracle-aided adversary* $\mathcal{A}$ that is given access to an oracle $f \in \mathcal{F}$ and a correlation finder $\mathsf{O} \in \mathcal{O}$ (which in turn has access to f). In particular, we will be interested in adversaries against OWP, which take as input $y \in \{0,1\}^\lambda$ and return some $x \in \{0,1\}^\lambda$.

The Two-Oracle Methodology

Following the two-oracle methodology developed in prior separation results [1, 9, 35, 53], we formulate, in Lemma 8 below, sufficient conditions for separating correlation intractability from OWP. The proof of the Lemma may be found in the full version [24].

▶ **Definition 7** ((q, q', q'')-Bounded Adversary). *Let $q, q', q'' : \mathbb{N} \to \mathbb{N}$. We say that an oracle-aided adversary $\mathcal{A}$ is* (q, q', q'')-bounded *if, for any fixed correlation finder O and any $f = \{f_\lambda : \{0,1\}^{\ell(\lambda)} \to \{0,1\}^\lambda\}$, $\mathcal{A}^{f,\mathsf{O}^f}$ on any input with security parameter λ makes at most $q(\lambda)$ queries to f_λ and $q'(\lambda)$ queries to O, where each O-query C makes at most $q''(\lambda)$ queries to f_λ on any input.*

▶ **Lemma 8** (Separation via Correlation-Finding Oracle). *Let $\kappa : \mathbb{N} \to [0,1]$ and $c \in \mathbb{N}$ be a constant. Let $\mathcal{R}$ be a class of relations. Assume there exists a correlation finder $\mathcal{O} = \{\mathcal{O}_R\}_{R\in\mathcal{R}}$ (see* Setting and Notation *above), such that*

- ***$\mathcal{O}$ breaks all CIH (Correctness):*** *For any CIH candidate $H = \{H_n\}$ with query complexity bounded by $t(n)$, it holds that*

$$\mathbb{E}_{\substack{R\leftarrow\mathcal{R}\\ \mathsf{O}\leftarrow\mathcal{O}_R}}[\mathbf{Adv}^{H}_{\mathbf{CI}}(n, R, \mathsf{O})(n)] = \Pr_{\substack{f\leftarrow\mathcal{F}, h\leftarrow H_n\\ R\leftarrow\mathcal{R}, \mathsf{O}\leftarrow\mathcal{O}_R}}[z \leftarrow \mathsf{O}^f(1^n, h);\ (z, h^f(z)) \in R] > 1 - \frac{1}{2n^2}$$

  for infinitely many $n \in \mathbb{N}$.

- $\mathcal{F}$ ***is one-way under*** $\mathcal{O}$ ***(Security):*** *For any* (q,q,q)*-bounded* $\mathcal{A}$ *that on security parameter* λ *calls* O *only with queries* $(1^n, h)$ *s.t.* $n < 2^{\lambda/24}$*, any* $\lambda \in \mathbb{N}$*, and any* $R \in \mathcal{R}$*,*

$$\mathbb{E}_{\substack{f \leftarrow \mathcal{F} \\ \mathsf{O} \leftarrow \mathcal{O}_R}} [\mathbf{Adv}^f_{\mathbf{OWP}}(\lambda, \mathcal{A}^{\mathsf{O}})] = O(q(\lambda)^c \cdot \kappa(\lambda)).$$

Then, t*-bounded CIH functions for* $\mathcal{R}$ *with input length* m *are* $(\kappa^{-1/(1+c)})$*-fully black-box separated from OWP.*

Generic Assumptions on the Correlation Finder and the Adversary

To facilitate our proof, we will make few assumptions over the structure and behavior of both the correlation finder and the oracle-aided adversary attacking the underlying OWP. The first assumption we make is that the correlation finder answers each of its queries using independent randomness. This will be immediately satisfied by our construction later on.

▶ **Definition 9** (Query-Independent Oracle). *We say that a distribution* $\mathcal{O} : \mathcal{Q} \to \mathcal{Z}$ *over oracles is* query-independent *if the answers of a random oracle* $\mathsf{O} \leftarrow \mathcal{O}$ *to different queries are independent or, more formally, if*

$$(\mathsf{O}(Q))_{Q \in \mathcal{Q}} \equiv (\mathsf{O}_Q(Q))_{Q \in \mathcal{Q}}$$

where O *and* $\{\mathsf{O}_Q\}_{Q \in \mathcal{Q}}$ *are all sampled independently at random from* $\mathcal{O}$*.*

Next, we will assume that the adversary is *canonical*, namely that it follows the natural structure of a competent adversary.

▶ **Definition 10** (Canonical Adversary). *We say that an oracle-aided adversary* $\mathcal{A}$ *is* canonical *if it satisfies the following three properties for any* $\mathsf{O} = \{\mathsf{O}_n : \mathcal{C}_{m(n),n} \to \{0,1\}^n\}$ *and any* f*:*

(i) $\mathcal{A}^{f,\mathsf{O}^f}$ *never makes the same oracle query twice.*

(ii) *After any* O*-query* C *that* $\mathcal{A}^{f,\mathsf{O}^f}$ *makes,* $\mathcal{A}^{f,\mathsf{O}^f}$ *immediately calls* f *at any* x *such that* $C^f(\mathsf{O}(C)) \rightharpoondown x$*.*

(iii) $\mathcal{A}^{f,\mathsf{O}^f}$ *immediately halts and outputs answer if found (i.e. when* $\mathcal{A}^{f,\mathsf{O}^f}(y^*)$ *calls* f *at* $x = f^{-1}(y^*)$*).*

(iv) $\mathcal{A}^{f,\mathsf{O}^f}$ *always calls* f *at its final output(s) (i.e.* $\mathcal{A}^{f,\mathsf{O}^f}(y^*) = x$ *implies* $\mathcal{A}^{f,\mathsf{O}^f}(y^*) \xrightarrow{f} x$*).*

Crucial to our analysis is one more assumption over the adversary, namely that it is *smooth*. Conceptually, a smooth adversary never calls the correlation finder with queries that already convey sufficient amount of information for succeeding in its task (without the help of the correlation finder). Specifically, these are queries C where the transcript of $C^f(z)$ for a random input z may contain a correct solution with noticeable probability. That is, we require that any pre-image x is observed by $C^f(z)$ with negligible probability[6].

▶ **Definition 11** (Smooth Inputs). *Fix an oracle* f *and an* f*-aided circuit* $C \in \mathcal{C}_n$ *and let* $\gamma : \mathbb{N} \to [0,1]$*. We define the set of* γ-smooth inputs of C under f *as follows*

$$\mathsf{Smooth}^f_\gamma(C) = \{x \in \{0,1\}^\lambda \mid \lambda \in \mathbb{N}, \Pr_{z \leftarrow \{0,1\}^m}[C^f(z) \rightharpoondown x] \le \gamma(\lambda)\}.$$

[6] In fact, we require this only for x that has not been already observed by $\mathcal{A}$ since, by canonicality, if an x has already been observed then it is not the solution. Notice that, otherwise, the correlation finder has no knowledge about the identity of the targeted pre-image and, therefore, we quantify over all such x's.

▶ **Definition 12** (Smooth Adversary). *Let $\tau, \gamma : \mathbb{N} \to [0,1]$. We say that an oracle-aided adversary is* (τ, γ)-smooth *if, for any f, any fixed correlation finder O, any $\lambda \in \mathbb{N}$, any $\mathcal{A}$-input a, and any $i \in \mathbb{N}$, letting C_i be the (random) i^{th} O-query made by $\mathcal{A}^{f,\mathsf{O}}(a)$ and $X_{<i} \subset \{0,1\}^{\ell(\lambda)}$ be the set of all f_λ-queries made by $\mathcal{A}$ prior to C_i, it holds that*

$$\Pr_{\mathcal{A}}[\overline{X_{<C}} \cap \{0,1\}^\lambda \subseteq \mathsf{Smooth}^f_\gamma(C)] > \tau(\lambda)$$

We note that the notion of smooth adversaries was considered already in [9] with the difference that we require smoothness only for unobserved inputs – this simplifies the adaption to correlation intractability and is w.l.o.g. assuming canonical adversaries.[7]

In the lemma below, we argue that canonical and smooth adversaries are complete, in the sense that it would be sufficient to show one-wayness against them if we can tolerate a small cost in complexity. A proof of a more general lemma, that shows smoothening of adversaries against CRH as well, is given in the full version [24].

▶ **Lemma 13** (The Smoothening Lemma). *For any (q, q', q'')-bounded adversary $\mathcal{A}$ and any $\beta := \beta(\lambda)$, there exists a canonical $(q + \beta q' q'', q', q'')$-bounded adversary $\mathcal{B}$ such that the following two properties hold:*

- ***Correctness:*** *for any fixed correlation finder O, any f, and any $\lambda \in \mathbb{N}$,*

$$\mathbf{Adv}^f_{\mathbf{OWP}}(\lambda, \mathcal{B}^{\mathsf{O}}) \geq \mathbf{Adv}^f_{\mathbf{OWP}}(\lambda, \mathcal{A}^{\mathsf{O}}).$$

- ***Smoothness:*** *$\mathcal{B}$ is a $(1 - 2^{\log(q''/\gamma) - \gamma\beta}, \gamma)$-smooth adversary, for all $\gamma > 0$.*

3.1 One-wayness from Differential Indistinguishability

In this section, we formalize the notion of differential indistinguishability for correlation finders and show that it is hard to invert f under any differentially indistinguishable correlation finder. This allows us to focus our design on obtaining differential indistinguishability to establish separation from OWP (via Lemma 8). We begin by defining this new notion.

▶ **Definition 14** (Differential Indistinguishability). *Let $\delta, \gamma : \mathbb{N} \to \mathbb{R}^+$ and $q : \mathbb{N} \to \mathbb{N}$. We say that a correlation finder $\mathcal{O} = \{\mathcal{O}_R\}$ is* (non-adaptively) differentially (q, γ, δ)-indistinguishable for $\mathcal{F}$ *if for any $R \in \mathcal{R}$, any $\lambda \in \mathbb{N}$, any $f \in \mathcal{F}$, any f-aided circuit $C \in \mathcal{C}$ which makes at most $q(\lambda)$ queries to f_λ on any input, and any $x^* \in \mathsf{Smooth}^f_\gamma(C)$ of length $\ell(\lambda)$, it holds that*

$$\mathbf{SD}(\mathsf{O}^f(C),\ \mathsf{O}^{f'}(C)) \leq \delta(\lambda),$$

where $\mathsf{O} \leftarrow \mathcal{O}_R$ and $f' = f_{x^* \leftrightarrow x'}$ for a uniformly random $x' \leftarrow \{0,1\}^\ell$.

In the lemma below, we argue that it is hard to invert a random permutation f, even when given access to a differential indistinguishable correlation finder. The lemma is followed by a proof sketch whereas the full proof can be found in the full version [24].

[7] More specifically, in [9], they make any adversary smooth by modifying his O-queries. Their collision finder is oblivious to these modifications since the functionality of the queries (as f-aided circuits) is preserved. This does not hold for our correlation finder and, therefore, their smoothening method does not preserve advantage in our settings. Hence, we slightly modify the definition, w.l.o.g., to allow for generic smoothening under any oracle, in particular for our correlation finder.

▶ **Lemma 15.** *Let $\mathcal{F} = \{\mathcal{F}_\lambda\}$ be the distribution of random permutations over $\{0,1\}^\lambda$ (i.e. $\ell(\lambda) = \lambda$). Let $\delta : \mathbb{N} \to [0,1]$ and let $\mathcal{A}$ be a (q, q', q'')-bounded* Inv*-adversary that is canonical and (τ, γ)-smooth (see Definitions 10 and 12). Let $\mathcal{O}$ be a correlation finder that is query-independent (see Definition 9) and differentially (q'', γ, δ)-indistinguishable for $\mathcal{F}$. Then, it holds for any $\lambda \in \mathbb{N}$ that*

$$\mathbb{E}_{f,\mathsf{O}}[\mathbf{Adv}^{f}_{\mathbf{OWP}}(\lambda, \mathcal{A}^{\mathsf{O}})] \leq O((1 - \tau(\lambda)) + q(\lambda) \cdot 2^{-\lambda} + q'(\lambda) \cdot \delta(\lambda)),$$

where $f \leftarrow \mathcal{F}$ and $\mathsf{O} \leftarrow \mathcal{O}$.

Proof Sketch. To show hardness of inversion, we switch the OWP security game to an experiment where the adversary (and the oracle O) is given access to an oracle f' that statistically hides any information about the pre-image of the given challenge x^*, which deems non-trivial success in computing a successful answer virtually impossible. More specifically, we swap f at x^* with a random $x' \leftarrow \{0,1\}^\lambda$, essentially "randomizing" the pre-image of the given challenge under the given function $f' = f_{x^* \leftrightarrow x'}$. By the presumed differential indistinguishability of $\mathcal{O}$ (in particular) we are able to show that such a swap does not affect the view of $\mathcal{A}$ except with a negligible probability. Note that the notion of differential indistinguishability from Definition 14, that considers swapping of smooth inputs only, is already sufficient by the canonicality and the smoothness of the adversary. Lastly, by the symmetry between f and f' given $y^* = f(x^*)$, we may then conclude that f' hides any information about x^* and, hence, inversion is impossible. ◀

4 The Correlation Finder

In this section, we build a correlation finder that is correct and differentially indistinguishable. Given the framework from Section 3, this suffices to derive the desired separation from OWP. We stress that the presented correlation finder satisfies a more general notion of differential indistinguishability (namely, adaptive differential indistinguishability) which is crucial for the separation from CRH. We refer the reader to the full version [24] for more details.

Strategy: Picky Correlation Finder

Our correlation finder follows a natural structure of a *picky correlation finder*, namely a correlation finder that given any input, rejects with some probability (outputs $\perp$[8]) and, otherwise, simply outputs a uniformly random correlation under the target relation.

First, we define the *set of correlations* between a circuit and a relation.

▶ **Definition 16** (Set of Correlations). *Let $f : \{0,1\}^* \to \{0,1\}^*$ be any oracle function and $C^f : \{0,1\}^m \to \{0,1\}^n$ be an f-aided circuit. Let R be a relation. The* set of (R, C^f)-correlations *is defined as*

$$\mathsf{Corr}^{f}_{R,C} = \{z \mid (z, C^f(z)) \in R\}.$$

We sometimes omit f, R and C from notation when clear by context.

We now present our generic picky correlation finder, which we will later instantiate with a proper rejection policy.

[8] Although our definition for a correlation finder allows only for outputs in $\{0,1\}^m$, we can always dedicate some $z \in \{0,1\}^m$ (e.g. the all-zeros input) to correspond to $\perp$.

▶ **Construction 17** (Picky Correlation Finder). *Let $\mathcal{R}$ be a relation class and $\mathcal{F}$ be a class of oracles. Let $\rho := \{\rho_R^f : \mathcal{C} \to [0,1]\}$ be a an ensemble of functions (namely, a* rejection policy*) that take as input a description of an f-aided circuit and outputs a real value in $[0,1]$ w.r.t. fixed $R \in \mathcal{R}$ and $f \in \mathcal{F}$. We define our* picky correlation finder with rejection policy *ρ, which we denote by $\mathcal{CF}[\rho] = \{\mathcal{CF}_R[\rho_R]\}_{R \in \mathcal{R}}$, such that, for every $R \in \mathcal{R}$, $\mathcal{CF}_R[\rho_R]$ is the distribution over deterministic oracles where, for any C, letting $\mathsf{CF} \leftarrow \mathcal{CF}_R[\rho_R]$, $\mathsf{CF}_R^f(C)$ is an independent random variable equal to the random output of the following algorithm* [9]

$\underline{\mathsf{CF}_R^f(C):}$

***Reject** with probability $\rho_R^f(C)$ and, otherwise, output a uniformly random **correlation** $z_C \leftarrow \mathsf{Corr}_{R,C}^f$ (if $\mathsf{Corr} = \emptyset$, set $z_C = \perp$).*

We point out that $\mathcal{CF}[\rho]$ is by construction query-independent (as by Definition 9).

Given this framework, it remains only to specify the *rejection policy* of our correlation finder, that is, the probability with which he rejects for any given input. As a first step, we specify in the lemma below a list of conditions on the rejection policy (which is simply a function from the query space $\mathcal{C}$ to real values in $[0,1]$) that are sufficient to obtain a correlation finder that is both correct and differentially indistinguishable (via Construction 17). The proof of the lemma is deferred to the full version of the paper [24].

▶ **Lemma 18.** *Let $t, q, N : \mathbb{N} \to \mathbb{N}$ and $\epsilon, \gamma : \mathbb{N} \to [0,1]$. Let $\mathcal{R}$ be a relation class and let $\rho := \{\rho_R^f : \mathcal{C}^f \to [0,1]\}_{R \in \mathcal{R}}$ that satisfies the following properties:*

- ***Correctness:** For any $f \in \mathcal{F}$ and any circuit $C = \{C_n \in \mathcal{C}_n\}$ with query complexity bounded by $t(n)$, it holds, for infinitely many $n \in \mathbb{N}$, that*

$$\mathbb{E}_{R \leftarrow \mathcal{R}}[\rho_R^f(C_n)] < \frac{1}{2n^2}.$$

- ***Soundness:** For any $R \in \mathcal{R}$, any $f \in \mathcal{F}$ and any circuit $C \in \mathcal{C}$, if $\rho_R^f(C) < 1$ then, for any $\lambda \in \mathbb{N}$ and any $x \in \{0,1\}^\lambda$, it holds that*

$$\Pr_{z \leftarrow \mathsf{Corr}^f}[C^f(z) \rightharpoonup x] < \Pr_{z \leftarrow \{0,1\}^m}[C^f(z) \rightharpoonup x]/\epsilon(\lambda).$$

- ***Differential Indistinguishability (of Rejection)**[10]**:** For any $R \in \mathcal{R}$, any $f \in \mathcal{F}$, any circuit $C \in \mathcal{C}$, any $\lambda \in \mathbb{N}$ and any $x^*, x' \in \{0,1\}^\lambda$ such that $x^*, x' \in \mathsf{Smooth}_\gamma^f(C)$,*

$$|\rho_R^{f'}(C) - \rho_R^f(C)| < \delta(\lambda), \quad \text{where } f' = f_{x^* \leftrightarrow x'}.$$

Then, $\mathcal{CF}[\rho] = \{\mathcal{CF}_R[\rho_R]\}$ from Construction 17 satisfies

➤ ***Correctness** as required by Lemma 8, for $t(n)$-bounded CIH candidates,*

➤ ***Differential (q'', γ, δ')-Indistinguishability** , where $\delta' = O(\delta + (q''/\gamma) \cdot 2^{-\lambda} + q'' \cdot \gamma/\epsilon)$.*

[9] Although the described algorithm has access to f, we think of its random coins as being sampled obliviously of f (w.l.o.g.), and therefore $\mathcal{CF}_R[\rho_R]$ is well-defined prior to setting f. One way to sample such an oracle is proposed in the proof of Lemma 18.

[10] We note that a much weaker notion of differential indistinguishability is sufficient for the rejection policy for the separation result from OWP. The presented stronger variant of the notion (in fact, a more general one) is required for the separation from CRH and is anyway achieved by our construction.

The Rejection Policy

We are now prepared to define our rejection policy. As outlined in the overview in Section 1.3, our strategy is to have the rejection probability be proportional to the worst-case amplification $\alpha^f_{R,C}(x)$ (see (3)) to obtain soundness as required by Lemma 18 and, further, to blur out the difference between adjacent functions to obtain differential indistinguishability (of rejection), by "spreading out" large amplification factors corresponding to some "bad" f over its neighborhood in the function space.

More specifically, since it is sufficient to guarantee differential indistinguishability only w.r.t. swaps between smooth inputs (see Lemma 18), we define the *d-Neighborhood* of f (for smoothness parameter γ) as

$$\mathcal{N}^f_{d,\gamma}(\lambda, C) = \bigcup_{0 \le d' \le d} \left\{ f_{\mathbb{X}} \mid \mathbb{X} = ((x^*_1, x'_1), \ldots, (x^*_{d'}, x'_{d'})) : \ \forall\, i,\ x^*_i, x'_i \in \mathsf{Smooth}_\gamma^{f_{\mathbb{X}_{i-1}}}(C) \right\},$$

where $f_{\mathbb{X}}$ is the function obtained from f by swapping $x^*_i \leftrightarrow x'_i$ for all $(x^*_i, x'_i) \in \mathbb{X}$, and $\mathbb{X}_\ell$ denotes the first ℓ swaps in $\mathbb{X}$. Further, we define the *distance* between functions f and f' as the smallest number of swaps required between smooth inputs in $\{0,1\}^\lambda$ to obtain f' from f (and is set to ∞ if such a transformation is not possible). That is,

$$\Delta_{\gamma,\lambda,C}(f, f') = \min_{d} f' \in \mathcal{N}^f_{d,\gamma}(\lambda, C).$$

We subsequently define our rejection policy in Figure 2.

$\rho = \{\rho_R : \mathcal{C} \to [0,1]\}$ **Parameters:** $\epsilon, \gamma : \mathbb{N} \to [0,1]$ (such that $\epsilon \gg \gamma$).

For any $C \in \mathcal{C}$, letting $d = (\epsilon/\gamma t)^2$, where t is the largest number of queries that C makes on any input, we define

$$\rho^f_R(C) = \max_{\substack{\lambda \in \lambda(C), \mathbf{x} \in \mathbf{X}_\lambda \\ f' \in \mathcal{N}^f_{d,\gamma}(\lambda, C)}} [e^{-\Delta_{\gamma,\lambda,C}(f,f')/\sqrt{d}} \cdot \int(\epsilon(\lambda) \cdot \alpha^{f'}_{R,C}(\mathbf{x}))],$$

where $\int(\chi) = \min(1, \chi)$ and

$$\lambda(C) = \{\lambda \in \mathbb{N} \mid \exists z \in \{0,1\}^m, x \in \{0,1\}^{\ell(\lambda)} : \ C^f(z) \rightharpoondown x\}.$$

Figure 2 The Rejection Policy.

Deriving Theorem 5

Lastly, we may derive the proof of the separation result stated in Theorem 5 by utilizing our correlation finder construction with carefully chosen parameters to obtain them via the differential indistinguishability framework from Section 3. We hereby provide a brief outline and refer the reader to the full version [24] for more details.

By Lemma 8, a separation can be derived by a correlation finder that satisfies both the correctness and security conditions in the lemma. We choose our correlation finder to be the picky correlation finder from Construction 17 with the rejection policy from Figure 2, instantiated with parameters $\gamma = 2^{-2\lambda/3}$ and $\epsilon = 2^{-\lambda/3}$ for the construction. By the smoothening lemma (Lemma 13), it suffices to establish security of the OWP against smooth adversaries which, in turn, is guaranteed when the correlation finder is differentially indistinguishable, as demonstrated by Lemma 15. Consequently, to prove separation, it is sufficient to show that our proposed rejection policy satisfies the conditions set by Lemma 18, namely correctness, soundness, and differential indistinguishability.

While one can prove, almost immediately by our construction, that the rejection policy satisfies soundness $\epsilon = 2^{-\lambda/3}$ and differential indistinguishability $O(q\gamma/\epsilon) = O(q2^{-\lambda/3})$, proving correctness demands much more effort. We refer the reader to the technical overview in Section 1.3 for an intuitive outline of our approach and to the full version [24] for a full formal proof.

Plugging in our correlation finder into Lemma 8, with $c = 2$ and $\kappa = \lambda 2^{-\lambda/3}$, we obtain the desired $2^{\lambda/10}$-separation.

References

1 Gilad Asharov and Gil Segev. Limits on the power of indistinguishability obfuscation and functional encryption. *SIAM J. Comput.*, 45(6):2117–2176, 2016. doi:10.1137/15M1034064.

2 Saikrishna Badrinarayanan, Rex Fernando, Aayush Jain, Dakshita Khurana, and Amit Sahai. Statistical zap arguments. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020*, pages 642–667, Cham, 2020. Springer International Publishing.

3 Boaz Barak and Oded Goldreich. Universal arguments and their applications. In *Proceedings of the 17th Annual IEEE Conference on Computational Complexity, Montréal, Québec, Canada, May 21-24, 2002*, pages 194–203. IEEE Computer Society, 2002. doi:10.1109/CCC.2002.1004355.

4 Boaz Barak, Yehuda Lindell, and Salil Vadhan. Lower bounds for non-black-box zero knowledge. *J. Comput. Syst. Sci.*, 72(2):321–391, March 2006. doi:10.1016/j.jcss.2005.06.010.

5 Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Douglas R. Stinson, editor, *Advances in Cryptology – CRYPTO' 93*, pages 232–249, Berlin, Heidelberg, 1994. Springer Berlin Heidelberg.

6 Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable zero knowledge with no trusted setup. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019*, pages 701–732, Cham, 2019. Springer International Publishing.

7 Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. Interactive oracle proofs. In Martin Hirt and Adam Smith, editors, *Theory of Cryptography*, pages 31–60, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.

8 Nir Bitansky, Dana Dachman-Soled, Sanjam Garg, Abhishek Jain, Yael Tauman Kalai, Adriana López-Alt, and Daniel Wichs. Why "fiat-shamir for proofs" lacks a proof. In Amit Sahai, editor, *Theory of Cryptography*, pages 182–201, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

9 Nir Bitansky and Akshay Degwekar. On the complexity of collision resistant hash functions: New and old black-box separations. In Dennis Hofheinz and Alon Rosen, editors, *Theory of Cryptography*, pages 422–450, Cham, 2019. Springer International Publishing.

10 Manuel Blum, William S. Evans, Peter Gemmell, Sampath Kannan, and Moni Naor. Checking the correctness of memories. In *32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991*, pages 90–99. IEEE Computer Society, 1991. doi:10.1109/SFCS.1991.185352.

11 Zvika Brakerski, Venkata Koppula, and Tamer Mour. Nizk from lpn and trapdoor hash via correlation intractability for approximable relations. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020*, pages 738–767, Cham, 2020. Springer International Publishing.

12 Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, Ron D. Rothblum, and Daniel Wichs. Fiat-shamir: From practice to theory. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, pages 1082–1090, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3313276.3316380.

13 Ran Canetti, Yilei Chen, and Leonid Reyzin. On the correlation intractability of obfuscated pseudorandom functions. In *Proceedings, Part I, of the 13th International Conference on Theory of Cryptography – Volume 9562*, TCC 2016-A, pages 389–415, Berlin, Heidelberg, 2016. Springer-Verlag. doi:10.1007/978-3-662-49096-9_17.

14 Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited. *J. ACM*, 51(4):557–594, July 2004. doi:10.1145/1008731.1008734.

15 Yilei Chen, Alex Lombardi, Fermi Ma, and Willy Quach. Does fiat-shamir require a cryptographic hash function? Cryptology ePrint Archive, Report 2020/915, 2020. URL: https://eprint.iacr.org/2020/915.

16 Arka Rai Choudhuri, Sanjam Garg, Abhishek Jain, Zhengzhong Jin, and Jiaheng Zhang. Correlation intractability and snargs from sub-exponential DDH. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023 – 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20-24, 2023, Proceedings, Part IV*, volume 14084 of *Lecture Notes in Computer Science*, pages 635–668. Springer, 2023. doi:10.1007/978-3-031-38551-3_20.

17 Arka Rai Choudhuri, Pavel Hubácek, Chethan Kamath, Krzysztof Pietrzak, Alon Rosen, and Guy N. Rothblum. Finding a nash equilibrium is no easier than breaking fiat-shamir. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, pages 1103–1114, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3313276.3316400.

18 Arka Rai Choudhuri, Abhishek Jain, and Zhengzhong Jin. Snargs for $\mathcal{P}$ from LWE. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 68–79. IEEE, 2021. doi:10.1109/FOCS52979.2021.00016.

19 Ivan Damgård. Collision free hash functions and public key signature schemes. In David Chaum and Wyn L. Price, editors, *Advances in Cryptology – EUROCRYPT '87, Workshop on the Theory and Application of of Cryptographic Techniques, Amsterdam, The Netherlands, April 13-15, 1987, Proceedings*, volume 304 of *Lecture Notes in Computer Science*, pages 203–216. Springer, 1987. doi:10.1007/3-540-39118-5_19.

20 Whitfield Diffie and Martin Hellman. New directions in cryptography. *IEEE transactions on Information Theory*, 22(6):644–654, 1976.

21 Nico Döttling, Sanjam Garg, Yuval Ishai, Giulio Malavolta, Tamer Mour, and Rafail Ostrovsky. Trapdoor hash functions and their applications. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019*, pages 3–32, Cham, 2019. Springer International Publishing.

22 Cynthia Dwork. Differential privacy. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II*, volume 4052 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2006. doi:10.1007/11787006_1.

23 Cynthia Dwork, Moni Naor, Omer Reingold, and Larry J. Stockmeyer. Magic functions. *J. ACM*, 50(6):852–921, 2003. doi:10.1145/950620.950623.

24 Nico Döttling and Tamer Mour. On the black-box complexity of correlation intractability. Cryptology ePrint Archive, Paper 2023/1365, 2023. URL: https://eprint.iacr.org/2023/1365.

25 Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *Advances in Cryptology — CRYPTO' 86*, pages 186–194, Berlin, Heidelberg, 1987. Springer Berlin Heidelberg.

26 Oded Goldreich, Shafi Goldwasser, and Silvio Micali. On the cryptographic applications of random functions. In *Proceedings of CRYPTO 84 on Advances in Cryptology*, pages 276–288, Berlin, Heidelberg, 1985. Springer-Verlag.

27 S. Goldwasser and Y. T. Kalai. On the (in)security of the fiat-shamir paradigm. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pages 102–113, 2003. doi:10.1109/SFCS.2003.1238185.

28 Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984. doi:10.1016/0022-0000(84)90070-9.

29 Vipul Goyal, Abhishek Jain, Zhengzhong Jin, and Giulio Malavolta. Statistical zaps and new oblivious transfer protocols. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020*, pages 668–699, Cham, 2020. Springer International Publishing.

30 Satoshi Hada and Toshiaki Tanaka. A relationship between one-wayness and correlation intractability. In *Public Key Cryptography*, pages 82–96, Berlin, Heidelberg, 1999. Springer Berlin Heidelberg.

31 Satoshi Hada and Toshiaki Tanaka. Zero-knowledge and correlation intractability. *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.*, E89-A(10):2894–2905, October 2006. doi:10.1093/ietfec/e89-a.10.2894.

32 Shai Halevi and Silvio Micali. Practical and provably-secure commitment schemes from collision-free hashing. In Neal Koblitz, editor, *Advances in Cryptology – CRYPTO '96, 16th Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 1996, Proceedings*, volume 1109 of *Lecture Notes in Computer Science*, pages 201–215. Springer, 1996. doi:10.1007/3-540-68697-5_16.

33 Johan Hastad, Russell Impagliazzo, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28, February 1999. doi:10.1137/S0097539793244708.

34 J. Holmgren and A. Lombardi. Cryptographic hashing from strong one-way functions (or: One-way product functions and their applications). In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 850–858, 2018. doi:10.1109/FOCS.2018.00085.

35 Chun-Yuan Hsiao and Leonid Reyzin. Finding collisions on a public road, or do secure hash functions need secret coins? In Matt Franklin, editor, *Advances in Cryptology – CRYPTO 2004*, pages 92–105, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.

36 James Hulett, Ruta Jawale, Dakshita Khurana, and Akshayaram Srinivasan. Snargs for P from sub-exponential DDH and QR. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022 – 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30 – June 3, 2022, Proceedings, Part II*, volume 13276 of *Lecture Notes in Computer Science*, pages 520–549. Springer, 2022. doi:10.1007/978-3-031-07085-3_18.

37 R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*, STOC '89, pages 44–61, New York, NY, USA, 1989. Association for Computing Machinery. doi:10.1145/73007.73012.

38 Russell Impagliazzo. Personal view of average-case complexity. In *Proceedings of the IEEE Annual Structure in Complexity Theory Conference*, pages 134–147, July 1995. doi:10.1109/SCT.1995.514853.

39 Abhishek Jain and Zhengzhong Jin. Non-interactive zero knowledge from sub-exponential DDH. In Anne Canteaut and Francois-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021 – 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part I*, volume 12696 of *Lecture Notes in Computer Science*, pages 3–32. Springer, 2021. doi:10.1007/978-3-030-77870-5_1.

40 Ruta Jawale, Yael Tauman Kalai, Dakshita Khurana, and Rachel Zhang. Snargs for bounded depth computations and ppad hardness from sub-exponential lwe. *IACR Cryptol. ePrint Arch*, 2020:980, 2020.

41 Yael Tauman Kalai, Alex Lombardi, and Vinod Vaikuntanathan. Snargs and PPAD hardness from the decisional diffie-hellman assumption. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023 – 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part II*, volume 14005 of *Lecture Notes in Computer Science*, pages 470–498. Springer, 2023. doi:10.1007/978-3-031-30617-4_16.

42 Yael Tauman Kalai, Guy N. Rothblum, and Ron D. Rothblum. From obfuscation to the security of fiat-shamir for proofs. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017*, pages 224–251, Cham, 2017. Springer International Publishing.

43 Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing*, STOC '92, pages 723–732, New York, NY, USA, 1992. Association for Computing Machinery. doi:10.1145/129712.129782.

44 Ilan Komargodski, Moni Naor, and Eylon Yogev. Collision resistant hashing for paranoids: Dealing with multiple collisions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018 – 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 – May 3, 2018 Proceedings, Part II*, volume 10821 of *Lecture Notes in Computer Science*, pages 162–194. Springer, 2018. doi:10.1007/978-3-319-78375-8_6.

45 Alex Lombardi and Vinod Vaikuntanathan. Fiat-shamir for repeated squaring with applications to ppad-hardness and vdfs. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020 – 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part III*, volume 12172 of *Lecture Notes in Computer Science*, pages 632–651. Springer, 2020. doi:10.1007/978-3-030-56877-1_22.

46 Alex Lombardi and Vinod Vaikuntanathan. Correlation-intractable hash functions via shift-hiding. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 – February 3, 2022, Berkeley, CA, USA*, volume 215 of *LIPIcs*, pages 102:1–102:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.ITCS.2022.102.

47 Silvio Micali. Computationally sound proofs. *SIAM J. Comput.*, 30(4):1253–1298, October 2000. doi:10.1137/S0097539795284959.

48 Moni Naor. Bit commitment using pseudorandomness. *Journal of cryptology*, 4(2):151–158, 1991.

49 R. Ostrovsky and A. Wigderson. One-way functions are essential for non-trivial zero-knowledge. In *[1993] The 2nd Israel Symposium on Theory and Computing Systems*, pages 3–17, 1993. doi:10.1109/ISTCS.1993.253489.

50 Chris Peikert and Sina Shiehian. Privately constraining and programming prfs, the lwe way. In Michel Abdalla and Ricardo Dahab, editors, *Public-Key Cryptography – PKC 2018*, pages 675–701, Cham, 2018. Springer International Publishing.

51 Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for np from (plain) learning with errors. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019*, pages 89–114, Cham, 2019. Springer International Publishing.

52 Omer Reingold, Luca Trevisan, and Salil Vadhan. Notions of reducibility between cryptographic primitives. In Moni Naor, editor, *Theory of Cryptography*, pages 1–20, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.

53 Daniel R. Simon. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In Kaisa Nyberg, editor, *Advances in Cryptology – EUROCRYPT'98*, pages 334–345, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.