

Cryptographic Hashing From Strong One-Way Functions

Or: One-Way Product Functions and their Applications

The full version of this paper is freely available on the Cryptology ePrint Archive [1].

Justin Holmgren
MIT CSAIL
Cambridge, USA
Email: holmgren@mit.edu

Alex Lombardi
MIT CSAIL
Cambridge, USA
Email: alexjl@mit.edu

Abstract—Constructing collision-resistant hash families (CRHFs) from one-way functions is a long-standing open problem and source of frustration in theoretical cryptography. In fact, there are strong negative results: black-box separations from one-way functions that are $2^{-(1-o(1))n}$ -secure against polynomial time adversaries (Simon, EUROCRYPT '98) and even from indistinguishability obfuscation (Asharov and Segev, FOCS '15).

In this work, we formulate a mild strengthening of exponentially secure one-way functions, and we construct CRHFs from such functions. Specifically, our security notion requires that every polynomial time algorithm has at most $2^{-n} \cdot \text{negl}(n)$ probability of inverting *two independent challenges*.

More generally, we consider the problem of simultaneously inverting k functions f_1, \dots, f_k , which we say constitute a “one-way product function” (OWPF). We show that sufficiently hard OWPFs yield hash families that are multi-input correlation intractable (Canetti, Goldreich, and Halevi, STOC '98) with respect to all sparse (bounded arity) output relations. Additionally assuming indistinguishability obfuscation, we construct hash families that achieve a broader notion of correlation intractability, extending the recent work of Kalai, Rothblum, and Rothblum (CRYPTO '17). In particular, these families are sufficient to instantiate the Fiat-Shamir heuristic in the plain model for a natural class of interactive proofs.

An interesting consequence of our results is a potential new avenue for bypassing black-box separations. In particular, proving (with necessarily non-black-box techniques) that parallel repetition amplifies the hardness of specific one-way functions – for example, all one-way permutations – suffices to directly bypass Simon’s impossibility result.

I. INTRODUCTION

Cryptographically secure hash functions are a fundamental building block in cryptography. Some of their most ubiquitous applications include the construction of digital signature schemes [2], efficient CCA-secure encryption [3], succinct delegation of computation [4], and removing interaction from protocols [5]. In their most general form, hash functions can be modeled as “random oracles” [3],

in which case it is heuristically assumed that an explicitly described hash function H (possibly sampled at random from a family) behaves like a random function, as far as a computationally bounded adversary can tell.

One of the most basic properties one might desire from a hash function is *collision resistance*, which requires that a computationally bounded adversary, given an explicit (shrinking) function H , cannot find a pair of distinct inputs (x, y) such that $H(x) = H(y)$. Since their introduction [6], collision-resistant hash functions have proved extremely useful in designing cryptographic primitives and protocols. As such, the following problem has received much attention in theoretical cryptography.

Question I.1. *What are the assumptions from which collision-resistant hash functions can be built? In particular, can they be built from an arbitrary one-way function?*

The question of building CRHFs from arbitrary one-way functions is particularly intriguing because OWFs are sufficient to construct a wide class of cryptographic primitives, including: pseudorandom generators [7], pseudorandom functions [8] and secret-key encryption, universal one-way hash functions [9] and digital signatures, commitment schemes [10], zero-knowledge proofs [11], and garbled circuits [12], [13].

Unfortunately, all known constructions of CRHFs have required assumptions beyond general one-way functions, such as *structured* generic assumptions (e.g. the existence of claw-free pairs of permutations) or the hardness of specific problems (e.g. computing discrete logarithms or finding approximately short vectors on lattices). Even worse, there are strong negative results on the prospect of constructing CRHFs from arbitrary OWFs in the form of *black-box impossibility results*. The first such result is due to Simon [14].

Theorem I.2 ([14], informal). *There is an oracle relative to which no collision-resistant hash functions exist, but exponentially secure one-way permutations exist.*

In fact, CRHFs have proved to be an extremely frustrating primitive in theoretical cryptography, as they have evaded attempts to describe a hierarchy of cryptographic primitives (with “weaker” objects implied by the existence of “stronger” objects). In a stark demonstration of this problem, Asharov and Segev [15] proved that CRHFs are not even implied (in a black box¹ way) by one-way functions and the extremely powerful notion of indistinguishability obfuscation [16], [17].

Theorem I.3 ([15], informal). *There is an oracle relative to which no collision-resistant hash functions exist, but exponentially secure one-way permutations and indistinguishability obfuscation exist.*

These negative results indicate substantial barriers to building CRHFs from OWFs (or OWPs, or indeed from any of the vast array of primitives implied by IO and OWPs). Collision resistance is also just *one* desirable property of random oracles, and our question above is a special case of the following more ambitious question.

Question I.4. *Which random oracle properties can be guaranteed under standard cryptographic assumptions, and how weak can these assumptions be made?*

It is known that some random oracle properties are *not realizable* in the standard model [18], [19]. However, there has been a recent line of work [20]–[22] showing that under strong assumptions, many random oracle properties (specifically in the context of “single input correlation intractability”) *can* be realized, and Question I.4 in its full generality remains wide open.

A. Our Contributions

In this work, we make progress on all of the above questions by defining a natural strengthening of exponentially secure OWFs² that suffices for building CRHFs and more. An “uber” version of our assumption – which we state for the purpose of intuition but is quantitatively and qualitatively much stronger than what we actually require – states that for every $k = \text{poly}(n)$, there exists an injective (polynomial-time computable) function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ with the following “batch one-wayness” property: For every polynomial-size adversary \mathcal{A} , the probability that $\mathcal{A}(f(X_1), \dots, f(X_k)) = (X_1, \dots, X_k)$ for $X_1, \dots, X_k \stackrel{\text{i.i.d.}}{\leftarrow} \{0, 1\}^n$ is bounded by $2^{-kn} \cdot \text{poly}(n)$.

Based on various significant *weakenings* of this uber-assumption, we construct:

- Collision-resistant hash families whose security against polynomial-time adversaries matches that of a random oracle.

¹“Black box” usage of IO and one-way functions is formalized through the notion of obfuscation for *oracle-aided* circuits. We refer the reader to [15] for details.

²Actually, OWFs where any *polynomial-time* algorithm can invert with only exponentially small probability

- More generally, for every k , we construct hash families \mathcal{H} that are “ k -ary output intractable” (inspired by a related definition of Zhandry [23]). Loosely speaking, given $H \leftarrow \mathcal{H}$, it is computationally hard to find distinct inputs X_1, \dots, X_k such that $(H(X_1), \dots, H(X_k))$ satisfy any fixed sparse relation R . The quantitative hardness that we achieve again matches that of a random oracle.

We are able to construct even stronger hash families if we additionally assume sub-exponentially secure indistinguishability obfuscation. This construction allows for applications including an instantiation of the Fiat-Shamir heuristic [5] for a natural class of interactive proofs.

Our main results and contributions are, in more detail, as follows.

1) *Defining OWPFs:* We introduce the notion of a family of one-way k -product functions (k -OWPFs), which is a family of k -tuples of functions (f_1, \dots, f_k) that are jointly “extremely one-way”. Such a family is most interesting when the hardness of inversion exceeds that of any individual f_i . For simplicity, suppose that each f_i is injective. In this case, we consider the assumption that no polynomial-time algorithm can recover $X_1, \dots, X_k \stackrel{\text{i.i.d.}}{\leftarrow} \{0, 1\}^n$ given $(f_1(X_1), \dots, f_k(X_k))$ with probability better than δ . Ideally, this could be true for δ as large as $2^{-(k-o(k))n}$. We call this a δ -hardness assumption of *batch inversion* for (f_1, \dots, f_k) .

The existence of such a family would follow from the following two conditions:

- A $\delta^{1/k}$ -secure injective one-way function f , and
- An optimal *parallel repetition theorem* for the hardness of f , i.e. one which states that if a function f is (s, δ) -hard to invert, then its k -wise repetition f^k is (s, δ^k) -hard to invert.

While such a dream parallel repetition property likely does not hold for *general* f [24], the counterexample presented therein does not preclude a similar result for a broad class of functions f .

In fact, the parallel repetition framework described above yields a special kind of OWPF family: one in which all k functions f_1, \dots, f_k are equal. We say that such OWPF families are *symmetric*. Another special case of interest, which we call a *one-way power family*, is a OWPF family of the form \mathcal{F}^k , meaning that the k functions f_1, \dots, f_k are sampled independently at random from a fixed family \mathcal{F} .

Our constructions (that do not require obfuscation) are based directly on symmetric injective OWPFs as a building block rather than general OWPFs. We augment these constructions by providing generic transformations between different notions of OWPFs, including constructions of (weaker) symmetric OWPFs from (stronger) general OWPFs, and constructions of *injective* k -OWPFs from arbitrary k -OWPFs (with some security loss).

One of our main contributions in this work is initiating the study of OWPFs and establishing their basic proper-

ties. We expect that OWPFs will prove useful in future work.

On Extreme Hardness Amplification: For all of our constructions without obfuscation, we actually rely on *symmetric* OWPF families. That is, we want a family $\mathcal{F} = \{\mathcal{F}_n\}$ such that if we sample $f \leftarrow \mathcal{F}_n$ and $x_1, \dots, x_k \leftarrow \{0, 1\}^n$, it is δ^k -hard to simultaneously invert $f(x_1), \dots, f(x_k)$. Clearly a necessary condition for this is that \mathcal{F} is a δ -secure one-way function family. But is this sufficient? The answer in general is no, as we discuss next.

First of all, this type of attempted hardness amplification fails for any family whose functions have short trapdoors that enable polynomial-time inversion. Given $f, f(x_1), \dots, f(x_k)$, an adversary can simply guess the trapdoor for f , succeed with some small probability that *does not depend on k* , and conditioned on guessing correctly can efficiently invert $f(x_1), \dots, f(x_k)$.

It is natural to next consider *functions* (or ensembles of functions $\{f_n : \{0, 1\}^n \rightarrow \{0, 1\}^*\}_n$ indexed only by input length) that are secure against non-uniform adversaries, and in particular do not have any trapdoors. However, [24] present an example of a single one-way function f for which it is as easy to invert $f(x_1), \dots, f(x_k)$ as it is to invert a single $f(x)$. Although their counterexample heavily relies on the fact that there are multiple permissible solutions to each instance x , there is also evidence that parallel repetition sometimes fails to increase the security of *injective* one-way functions [25].

Despite the above negative results, we emphasize that symmetric OWPFs only require direct products to amplify hardness for *specific* functions, rather than broad classes of functions. Moreover, one-way product functions may exist even if parallel repetition does not amplify the hardness of *any* function f beyond negligible. In particular, f_1, \dots, f_k may not all be the same function, and may be sampled from a joint distribution on k -tuples of functions. These observations leave us with at least two promising avenues towards constructing OWPF candidates:

- 1) Given the contrived nature of known counterexamples to one-way function parallel repetition, any “natural” δ -secure injective OWF family also serves as a candidate one-way power family with security roughly δ^k .
- 2) It may be possible to “fortify” any one-way function family \mathcal{F} into a related family \mathcal{F}' whose security *does* amplify to an extreme degree, yielding symmetric OWPFs.

Finally, we mention a concrete candidate symmetric OWPF family based on the *multiple discrete logarithm problem*. That is, in some group \mathbb{G}_n of order $|\mathbb{G}_n| \approx 2^n$, the problem is to simultaneously compute k discrete logarithms $X_1, \dots, X_k \stackrel{\text{i.i.d.}}{\leftarrow} [2^n]$ given input $(g, g^{X_1}, \dots, g^{X_k})$, where g is a generator for \mathbb{G}_n . In [26], evidence for the hardness of computing multiple discrete logarithms is given in the form of lower bounds in the generic group

model [27]. In particular, [26] show that (in our language) k -batch inversion is nearly 2^{-kn} -hard for polynomial-time generic-group algorithms.

2) *Constructions from OWPFs:* Our first application of OWPFs is a construction of a collision-resistant hash family from suitably secure symmetric 2-OWPFs. Informally, we prove

Theorem I.5. *Suppose that there exist symmetric injective 2-OWPFs with security $2^{-n-\omega(\log n)}$. Then, there exists a collision-resistant hash family.*

This type of OWPF does not follow in a black-box way from even exponentially-hard one-way permutations; this is how we avoid the [14], [15] impossibility results.

Through one of our generic transformations of OWPFs, we also obtain a construction that does not assume injectivity:

Theorem I.6. *Suppose that there exist symmetric 2-OWPFs with security $2^{-(1.6+\epsilon)n}$. Then, there exists a collision-resistant hash family.*

Optimality and Implications of Theorem I.5: While we have explained how our result is not captured by the [14], [15] framework, one could question the necessity of this new OWPF assumption. For example, [15] only rules out black-box constructions of CRHFs from $2^{-\epsilon n}$ -secure IO and one-way permutations (for $\epsilon = \frac{1}{50}$ in particular), and [14] proves a quantitatively similar impossibility. What about assuming only $2^{-n/2}$ -secure OWPs, which are weaker and more standard than our symmetric OWPFs? As a complementary result, we show that these are insufficient – we strengthen the Asharov-Segev analysis to rule out black box constructions from IO and even 2^{-n} -secure one-way permutations.

Theorem I.7 (Extension of [15] Theorem 1.1, informal). *There is no black-box construction of CRHFs from sub-exponentially secure IO, sub-exponentially secure OWPs, and OWPs that ppt algorithms \mathcal{A} can invert with probability at most $\text{size}(\mathcal{A})^c \cdot 2^{-n}$ for some absolute constant c .*

Theorem I.7 indicates a sharp limit on directly improving Theorem I.5; in the latter, we show that injective 2-OWPFs that are $2^{-n} \cdot \text{negl}(n)$ -hard to invert suffice for constructing CRHFs from IO, while the former result says that improving the $2^{-n} \cdot \text{negl}(n)$ to $\frac{2^{-n}}{\text{negl}(n)}$ is impossible for black-box constructions. In particular, for black-box constructions, exponentially secure one-way permutations (in the usual sense) are insufficient.

Extension to Output Intractability: Theorem I.5 can be substantially generalized beyond collision-resistance. In particular, given a $2k$ -ary relation, we consider the problem of finding X_1, \dots, X_k such that $(X_1, \dots, X_k, H(X_1), \dots, H(X_k)) \in R$ for $H \leftarrow \mathcal{H}_n$. If this problem is hard, then \mathcal{H} is said to be multi-input correlation intractable for R , a notion due to [18].

Collision-resistance is the special case when $k = 2$ and

$$R = \{(x_1, x_2, y_1, y_2) : (x_1 \neq x_2) \wedge (y_1 = y_2)\}.$$

Random oracles are correlation intractable for any *sparse* relation R – that is, as long as for every $\mathbf{x} = (x_1, \dots, x_k)$, $\Pr_{\mathbf{Y} \leftarrow \{0,1\}^{n-1}} [(\mathbf{x}, \mathbf{Y}) \in R] \leq \text{negl}(n)$. In many applications, this correlation-intractability is the crucial property of a random oracle, and a fundamental theoretical question is whether it can be achieved by *concrete* hash families.

Despite the initial negative result of [18], which ruled out correlation intractability for arbitrary (e.g., unbounded-arity) relations, there has been substantial work on constructing hash families that are correlation intractable for “bounded” single-input/output relations [20]–[22] as well as hash families that are “output intractable” [23], that is, correlation intractable with respect to relations of the form “ $(x_i \neq x_j \text{ for all } i \neq j) \wedge R(y_1, \dots, y_k) = 1$.”³

Using suitably secure k -OWPFs, we construct hash families that are output intractable for *all* sparse output relations (with known bounded arity). The quantitative intractability that we prove depends on the sparsity of the relation, similarly to the situation for a true random oracle. Equivalently, we rely on weaker assumptions to show correlation-intractability of sparser relations.

A simplified version of our result is as follows.

Theorem I.8 (informal). *Suppose that there exists a family of symmetric injective k -OWPFs with security $(s + \text{poly}(n), \delta)$, let $m = m(n)$ denote any output length, and let $p = p(n)$ denote any sparsity. Then, there exists a hash family $\mathcal{H} = \{\mathcal{H}_{n,m(n)}\}$ that is output intractable, with security $(s, \delta \cdot p \cdot 2^{kn})$, with respect to all k -ary relations of sparsity p .*

In particular, if the k -OWPF family has optimal (2^{-kn}) security, then the hash family constructed in Theorem I.8 has output intractability matching that of a random oracle.

As an interesting special case, we note that Theorem I.8 gives a construction of k -multi-collision resistant hash functions (formally introduced in [28] and further studied in [29]–[31]) from symmetric injective k -OWPFs with security $2^{-n-k \log(k)} \cdot \text{negl}(n)$, an assumption that (up to a lower order term in the exponent) becomes weaker as k increases from 2 to any $o(\frac{n}{\log(n)})$. As any multi-collision-resistant hash family implies the existence of constant round statistically hiding commitments [29], [31], this yields constant round statistically hiding commitments from $2^{-n} \cdot \text{negl}(n)$ -secure (injective and symmetric) k -OWPFs for any $k = o(\frac{n}{\log(n)})$. Unlike the assumptions required for collision resistance, this assumption would follow from optimal parallel repetition for *any polynomially secure (injective) one-way function*.

³ [23] considers a slightly different notion of output intractability. We elaborate on this in the full version of this paper [1].

3) *Combining OWPFs with Indistinguishability Obfuscation*: Our results above, Theorem I.5 and Theorem I.8, are constructions of cryptographic hash families from (symmetric) OWPFs alone, and hence (partially) address the question of what hash families can be constructed from assumptions in the realm of one-wayness.

We additionally consider which hash families can be constructed in the plain model under stronger assumptions. Namely, we combine OWPFs with the powerful notion of indistinguishability obfuscation [16], [17]. This line of reasoning yields another construction of CRHFs, and more generally a construction of multi-input correlation intractable hash functions for a broader class of relations than achieved by Theorem I.8. In our IO-based construction, we are able to handle relations R which depend on both the input variables \mathbf{x} and the output variables \mathbf{y} , as long as the relation R is efficiently *locally* samplable. Informally, we need to be able to efficiently sample a random output \mathbf{Y} such that $(\mathbf{x}, \mathbf{Y}) \in R$ such that each output Y_i is sampled only knowing the corresponding input x_i (with arbitrary preprocessed shared randomness “between the variables”).

Moreover, our construction is extremely simple and confirms typical intuition about obfuscation: our hash family is an obfuscated (puncturable) PRF $\mathcal{O}(F_s(\cdot))$. We only require the existence of suitably secure OWPFs in the security proof; they are not needed in the construction. This result extends the framework of [20], [21] on constructing strong hash functions from obfuscation (and additional assumptions).

Our main result utilizing obfuscation is proved by viewing OWPFs themselves as a (weak) form of obfuscation: an injective k -OWPF (f_1, \dots, f_k) allows us to obfuscate *multi-point functions*, i.e., programs of the form

$$P_{x_1, \dots, x_k}(x) = \begin{cases} i & x = x_i \text{ for some } i \\ 0 & \text{otherwise.} \end{cases}$$

Since this construction is oblivious to whether or not the OWPF family \mathcal{F} is symmetric, this yields a construction of correlation intractable hash families (and in particular, of CRHFs) relying on weaker OWPF assumptions, at the cost of additionally assuming IO. That is, the assumptions on asymmetric OWPFs required here are quantitatively (and even qualitatively) weaker than those required without obfuscation, as we avoid the cost of converting asymmetric OWPFs into symmetric OWPFs.

As an interesting special case, the notion of correlation intractability that we achieve is powerful enough to capture nontrivial cases of the Fiat-Shamir paradigm for converting (constant round, public-coin) interactive proof systems into non-interactive argument systems. We elaborate on one such formal result in the full version of this paper, but the main intuition is that we can instantiate the Fiat-Shamir transform for any proof system that has the property that a malicious prover can efficiently determine which verifier messages he can cheat on.

This captures protocols that follow the “commit-challenge-response” framework. This approach yields a construction of NIZK argument schemes (in the common reference string model) through the Fiat-Shamir transform whose security relies on IO and the existence of exponentially secure one-way functions – no OWPF assumptions are needed in this case.

B. Related Work

a) Multi-Instance Security: There are a few other cryptographic constructions in the literature that are secure assuming a strong form of hardness amplification for one-way functions, or more generally some notion of multi-instance security. Several notable examples, although not a comprehensive listing, are as follows.

- In the context of password-based cryptography, [32] study the multi-instance security of encryption schemes and key derivation functions. Their work is motivated by the common practice of “salting”, which is intended to insure that the running time required for an adversary to compromise k users scales linearly with k .
- In the context of chosen ciphertext security, [33] consider the problem of simultaneously inverting $(f(x_1), \dots, f(x_k))$ where (x_1, \dots, x_k) are sampled from a joint distribution (rather than i.i.d.). In contrast to our work, they only ask that the inversion probability should be $\text{negl}(\lambda)$; that is, they do not ask for hardness to amplify. They show that *trapdoor functions* satisfying certain security properties of this flavor suffice to construct CCA-secure public key encryption.
- Inspired by Merkle puzzles, [34] construct a public-key encryption scheme that allows for adversaries that run in time at most quadratically larger than that of the honest parties. They prove the security of their scheme under the assumption that there is an injective one-way function f , a polynomial $k = k(n)$, a constant $0 < \delta < \frac{1}{2}$, and a (randomized) “multi-source hard-core predicate” H such that for random $x_1, \dots, x_k \leftarrow \{0, 1\}^n$, every algorithm running in time $2^{(1-\delta)n}$ on input $(f(x_1), \dots, f(x_k), r)$ successfully guesses $H(x_1, \dots, x_k, r)$ with advantage at most $2^{-\omega(n)}$.
- In concurrent and independent work, Bitansky and Lin [35] introduce the notion of an *amplifiable one-way function*. Roughly speaking, a one-way function f is (sub-exponentially) amplifiable if for all $k = \text{poly}(n)$ there exists a hard-core predicate hcb for f and an efficiently computable *combiner* C such that given $(y_1 = f(x_1), \dots, y_k = f(x_k))$ it is 2^{-k^ϵ} -hard (for 2^{n^ϵ} -time algorithms) to predict the combined hard-core bit $C(\text{hcb}(x_1), \dots, \text{hcb}(x_k))$. The work [35] shows that such a one-way function is useful in the construction of a one message non-malleable commitment scheme.

b) Extremely Lossy Functions: [23] introduces the notion of an extremely lossy function (ELF). In [23], ELFs are used as a central building block to construct several hash families with strong security properties. In particular, they can be used to construct hash functions satisfying a notion of output intractability that is incomparable to we achieve. Informally, [23] considers the more general setting of $k + 1$ -ary relations $R(y_1, \dots, y_k, w)$ with the property that for random (y_1, \dots, y_k) , it is computationally hard to find a witness w for which $R(y_1, \dots, y_k, w) = 1$ (where our notion would correspond to the case that for random (y_1, \dots, y_k) , *no such witness exists*), and constructs hash functions that are correlation intractable for such relations R that are efficiently decidable.

The only current construction of ELFs relies on an exponentially strong DDH assumption. An interesting open question is whether OWPFs imply the existence of ELFs, or even ordinary (i.e. moderately) lossy one-way functions.

c) CRHFs from Extremely Strong LPN: Two recent works [36], [37] give constructions of CRHFs from the Learning Parity with Noise (LPN) problem in parameter settings that resemble an exponential hardness assumption. We note that one of the same works [37] proves that these particular LPN assumptions imply hardness in the complexity class BPP^{SZK} , placing this construction on similar complexity-theoretic ground as prior constructions from discrete logarithm and SIS. The LPN-based CRHFs are also provably broken in quasi-polynomial time, while our CRHF is plausibly as collision-resistant as a random oracle.

d) Single-Input Correlation Intractability: Correlation intractability [18] is a clean but powerful property of random oracles that has drawn considerable interest, particularly for its relevance to the Fiat-Shamir transform [3], [5]. Circumventing the negative results of [18], [19], [38], there has been a recent line of work [20]–[22] on constructing (single input) correlation intractable hash functions and instantiating the Fiat-Shamir heuristic in the standard model, under strong assumptions. We build on this line of work, particularly the work of [21], to achieve results for special cases of *multi-input* correlation intractability under weaker or incomparable assumptions than are required in these previous works.

e) CRHFs from IO and SZK-hardness: [39] constructs CRHFs from indistinguishability obfuscation and any average-case hard problem in the complexity class $\text{SZK}^{0,1}$. We consider SZK-hardness to be a “structured assumption” which makes it different from (even very strong) assumptions on injective one-way functions; indeed, the same work proves an Asharov-Segev-like impossibility result for constructing (even worst-case) hard SZK instances from IO and OWPFs. A fascinating open question is whether OWPFs (with or without IO) imply SZK-hardness of any form.

C. Technical Overview

We now outline some of our constructions in more detail. In order to clearly demonstrate the power of OWPFs and our techniques, we focus on the following two special cases: constructing CRHFs from symmetric 2-OWPFs, and constructing CRHFs from IO and (asymmetric) injective 2-OWPFs.

1) *Construction of CHRFS:* For simplicity, we first assume that we have an ensemble of one-way permutations $\{f_n : \{0, 1\}^n \rightarrow \{0, 1\}^n\}$, where for every constant $c > 0$, double inversion is $2^{-n} \cdot n^{-c}$ hard for size- n^c adversaries. In this case, we construct a particularly simple CRHF: to sample a collision-resistant $H : \{0, 1\}^n \rightarrow \{0, 1\}^{n-1}$, first sample $P : \{0, 1\}^n \rightarrow \{0, 1\}^{n-1}$ from a pairwise independent hash family \mathcal{P}^4 . $H = P \circ f_n$. This and similar constructions have proved very useful in prior works [2], [23], [40].

We now sketch the proof of security. Assume for contradiction that some poly-size algorithm \mathcal{A} finds collisions in H with probability $\epsilon = \epsilon(n)$. We show how to use \mathcal{A} to simultaneously find $X_1^* = f_n^{-1}(Y_1^*)$ and $X_2^* = f_n^{-1}(Y_2^*)$ with probability roughly $\epsilon \cdot 2^{-n}$, given uniformly random $Y_1^*, Y_2^* \stackrel{\text{i.i.d.}}{\leftarrow} \{0, 1\}^n$. Specifically, we will invoke \mathcal{A} not on a uniformly sampled $H = P \circ f_n$, but on a differently defined $H = P_{\text{plant}} \circ f_n$, where P_{plant} is sampled from \mathcal{P} conditioned on $P_{\text{plant}}(Y_1^*) = P_{\text{plant}}(Y_2^*)$.

Intuitively, we now argue (by a purely statistical argument) that (X_1^*, X_2^*) looks sufficiently like a *uniformly random* collision of H that \mathcal{A} must output that exact collision with probability roughly $\epsilon \cdot 2^{-n}$. To make this intuition rigorous, suppose first that we ignore Y_1^* and Y_2^* , and simply invoke \mathcal{A} on a randomly sampled $H = P \circ f_n$. Then with probability ϵ , \mathcal{A} will find a collision (X_1, X_2) in H . Conditioned on this event, (X_1, X_2) will be *equal* to (X_1^*, X_2^*) with probability 2^{-2n} , for a total probability of $\epsilon \cdot 2^{-2n}$ that both events occur. But (X_1^*, X_2^*) is a collision in H with probability only $2^{-(n-1)}$. Thus, conditioning on this event (i.e., sampling $H = P_{\text{plant}} \circ f_n$ instead of $H = P \circ f_n$) boosts the probability that \mathcal{A} outputs (X_1^*, X_2^*) to $\epsilon \cdot 2^{-2n} \cdot 2^{n-1} = \epsilon \cdot 2^{-n-1}$.

Therefore, the CRHF we constructed satisfies the standard notion of security: every polynomial-size adversary finds collisions with probability that is negligible in n . From stronger hardness assumptions on $\{f_n\}$, i.e. that double-inversion is $\delta(n)$ -hard for size- $s(n)$ adversaries, one obtains a correspondingly more secure CRHF.

The above argument actually does not rely in any way on f_n being a permutation. It is, however, important that f_n is injective, so that all collisions in $P \circ f_n$ are due to P , and thus in some sense are randomly distributed.

We also show that the injectivity requirement can be traded off against a stronger hardness assumption. In fact,

⁴We also require that the hash family is *programmable* at any two points, meaning that it is possible to sample a uniformly random $p \leftarrow \mathcal{P}$ subject to the condition that $p(y_1) = z_1$ and $p(y_2) = z_2$.

if $\{f_n\}$ is extremely secure to begin with, we can construct a family of functions which is statistically injective, and still nearly as secure.

For simplicity, we illustrate this transformation for *one-way functions*. Suppose that $\{f_n\}$ is $\delta(n)$ -hard to invert for polynomial-time adversaries (think of $\delta(n) = 2^{-(1-o(1))n}$, although such extreme parameters are not necessary). We first observe that $\{f_n\}$ cannot be “extremely” non-injective; if one independently samples $X_1 \leftarrow \{0, 1\}^n$ and $X_2 \leftarrow \{0, 1\}^n$, then the probability that $f_n(X_1) = f_n(X_2)$ must be at most δ (otherwise one could break the security of f_n by random guessing). This can be leveraged to obtain a fully injective function (with some small error probability), as follows.

Set n to be any function of n' (think of $n(n') = 3n'$). Then define the ensemble of function families $\mathcal{F} = \{\mathcal{F}_{n'}\}$ as follows. To sample a function $f \leftarrow \mathcal{F}_{n'}$, sample $P : \{0, 1\}^{n'} \rightarrow \{0, 1\}^n$ from a pairwise independent hash family, and define $\tilde{f}_{n'} = f_n \circ P$. A simple pairwise independence argument shows that \mathcal{F} is statistically injective, with failure probability at most $2^{2n'} \cdot \delta(n)$ (with the suggested parameters in mind, this is $2^{-(1-o(1))n'}$).

Security of \mathcal{F} follows from observing that if an adversary cannot invert $f_n(X)$ with probability better than δ when sampling $X \leftarrow \{0, 1\}^n$, then for any subset $\mathcal{X} \subseteq \{0, 1\}^n$, the adversary cannot invert $f_n(X')$ with probability better than $\delta \cdot \frac{2^n}{|\mathcal{X}|}$ when sampling $X' \leftarrow \mathcal{X}$. With good probability $(1 - 2^{2n'-n})$, or with our suggested parameters $1 - 2^{-n'}$, it holds that $P : \{0, 1\}^{n'} \rightarrow \{0, 1\}^n$ is actually injective, so that inverting $f_n \circ P$ corresponds to inverting f_n when inputs are drawn from the uniform distribution on $\text{Img}(P)$. The above discussion shows that this is $\delta \cdot 2^{n-n'}$ -hard (or with our suggested parameters $2^{-(1-o(1))n'}$ -hard) even for adversaries that are given arbitrary advice about P .

While the above description refers to the case of one-way functions (i.e. 1-OWPFs), similar arguments can be made for arbitrary OWPFs (with different quantitative tradeoffs).

2) *Constructions Using Obfuscation:* We now outline our general proof strategy – which we informally refer to as the *planting technique* – for all of our constructions based on IO, using collision resistance as an example. The planting technique is inspired by the recent work of Kalai, Rothblum, and Rothblum [21] on instantiating the Fiat-Shamir heuristic using obfuscation.

For simplicity, we focus on hash functions that shrink by a single bit. Our construction is then simply an obfuscation $H \stackrel{\text{def}}{=} \mathcal{O}(F_S)$ of a puncturable pseudorandom function $F_S : \{0, 1\}^n \rightarrow \{0, 1\}^{n-1}$, where \mathcal{O} is an indistinguishability obfuscator. Recall that we also assume the existence of an injective but *not necessarily symmetric* 2-OWPF that cannot be inverted in polynomial time with probability better than $2^{-n-\omega(\log n)}$.

The proof of security then proceeds as follows. Assume for contradiction that some ppt algorithm \mathcal{A} finds a colli-

sion (X_1, X_2) of H with non-negligible⁵ probability ϵ . We then consider the behavior of \mathcal{A} on an obfuscation of a *different* program H_{plant} which overrides the functionality of F_S with a hard-coded planted collision $H_{\text{plant}}(X_1^*) = H_{\text{plant}}(X_2^*) = Y^*$, for independent and uniformly random X_1^*, X_2^* , and Y^* . That is, the functionality of H_{plant} is

$$H_{\text{plant}}(x) \stackrel{\text{def}}{=} \begin{cases} Y^* & \text{if } x = X_1^* \text{ or } x = X_2^* \\ F_S(x) & \text{otherwise.} \end{cases}$$

We then prove two contradictory claims.

- 1) The probability that \mathcal{A} outputs (X_1^*, X_2^*) is approximately $\epsilon \cdot 2^{-n-1}$, i.e. $2^{-n-O(\log n)}$.

This claim is argued as follows.

- a) If \mathcal{A} is given an obfuscation of a program H_{punc} that (in contrast to H_{plant}) overrides F_S with hard-coded mappings $X_1^* \mapsto Y_1^*$ and $X_2^* \mapsto Y_2^*$ for *independent* uniform $Y_1^*, Y_2^* \leftarrow \{0, 1\}^{n-1}$, then the probability that \mathcal{A} successfully produces a collision *and that collision is* (X_1^*, X_2^*) is very nearly $\epsilon \cdot 2^{-2n}$ by the security of \mathcal{O} and F_S .
 - b) (X_1^*, X_2^*) is only a valid collision of H_{punc} when $Y_1^* = Y_2^*$, so the probability that \mathcal{A} outputs (X_1^*, X_2^*) conditioned on $Y_1^* = Y_2^*$ is approximately $\epsilon \cdot 2^{-2n} \cdot 2^{n-1} = \epsilon \cdot 2^{-n-1}$. But the distribution of H_{punc} conditioned on $Y_1^* = Y_2^*$ is exactly the distribution of H_{plant} .
- 2) The probability that \mathcal{A} outputs (X_1^*, X_2^*) is $2^{-n-\omega(\log n)}$.

Since IO is the “best-possible” obfuscation [41], it suffices for there to exist *some* obfuscation of H_{plant} that hides (X_1^*, X_2^*) . This would follow from a “special-purpose” obfuscator \mathcal{O}' for membership testing in two-element sets (in our case $\{X_1^*, X_2^*\}$). The security property we need is that every ppt algorithm recovers (X_1^*, X_2^*) from $\mathcal{O}'(\{X_1^*, X_2^*\})$ with probability bounded by $2^{-n-\omega(\log n)}$.

This is a variant of “point function obfuscation”, a notion which was studied by [42]–[44]. Our variant (with uniformly random X_1^*, X_2^*) admits a particularly easy construction from injective 2-OWPFs – the obfuscation is $(W_1^* = f_1(X_1^*), W_2^* = f_2(X_2^*))$, and is evaluated on an input x as

$$\begin{cases} 1 & \text{if } f_1(x) = W_1^* \text{ or } f_2(x) = W_2^* \\ 0 & \text{otherwise.} \end{cases}$$

There are conceivably other ways to obtain this point function obfuscation, but for this particular construction, security is equivalent to the hardness of batch inverting (f_1, f_2) .

⁵In fact, our approach readily generalizes to obtain exponentially-secure CRHFs, at the cost of quantitatively stronger computational assumptions.

D. Conclusions and Questions

In this work, we have introduced a new family of computational assumptions – namely, the existence of various flavors of one-way product functions (OWPFs). We find these assumptions to be clean, plausible, and useful.

In terms of power, OWPFs allow the construction of hash families that achieve several elusive random oracle-like properties. In particular, our black-box construction of CRHFs shows that OWPFs are more powerful than *black box usage* of exponentially-secure one-way functions.

OWPFs are also extremely plausible. Depending on s, δ , and k , we view (s, δ) -secure k -OWPFs as somewhere between standard and exponentially-secure one-way functions. The plausibility is supported by a concrete candidate instantiation – the discrete log problem, which is probably a nearly optimal OWPF in the generic group model.

Indeed, this particular combination of plausibility and usefulness gives us some hope that CRHFs can be constructed solely based on exponentially strong one-way functions. More generally, our results suggest a possible blueprint for circumventing black-box impossibility results from OWPFs:

- 1) Build OWPFs from OWFs (using necessarily non-black-box techniques).
- 2) Build primitives in a black-box way from OWPFs.

One bonus of this approach is that it could result in constructions that are non-black-box only *in the security proof*, and thus has the potential for practical efficiency.

Independently, OWPFs satisfy several desirable properties for a cryptographic assumption. For example, for any family \mathcal{F} , the assumption “ \mathcal{F} is a k -OWPF” is a *search complexity assumption* [45]: for some efficiently sampleable distribution \mathcal{D} and efficiently checkable relation \mathcal{R} , the assumption is equivalent to requiring that on input $x \sim \mathcal{D}$, every bounded-time algorithm has bounded probability of finding y such that $(x, y) \in \mathcal{R}$.

There remain many intriguing questions about the precise power of OWPFs. In particular:

- What are the complexity-theoretic implications of OWPFs? For example, do they imply hardness in SZK? We emphasize that all prior constructions of CRHFs have been from assumptions that imply (average-case) SZK hardness, but CRHFs themselves are not known to imply any sort of SZK hardness.
- What implies OWPFs? Is it possible to construct non-trivial k -OWPFs from previously studied cryptographic assumptions? Above we outlined an approach to *generically* constructing OWPFs, but it is also possible that OWPFs can be based on concrete, structured assumptions.

ACKNOWLEDGEMENTS

We thank Zvika Brakerski, Yael Kalai, Omer Paneth, and Vinod Vaikuntanathan for helpful discussions and

comments. We also thank the anonymous FOCS reviewers for their useful feedback.

This work was done in part while the authors were visiting the Weizmann Institute of Science in January 2018, supported by the Binational Science Foundation (Grants No. 2016726, 2014276) and European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482) and Project PROMETHEUS (Grant 780701).

JH was supported in part by NSF Grant CNS-1413920. AL was supported in part by NSF Grants CNS-1350619 and CNS-1414119, and by the Defense Advanced Research Projects Agency (DARPA) and the U.S. Army Research Office under contracts W911NF-15-C-0226 and W911NF-15-C-0236.

REFERENCES

- [1] J. Holmgren and A. Lombardi, “Cryptographic hashing from strong one-way functions,” 2018. [Online]. Available: <https://eprint.iacr.org/2018/385>
- [2] M. Naor and M. Yung, “Universal one-way hash functions and their cryptographic applications,” in *Proceedings of the twenty-first annual ACM symposium on Theory of computing*. ACM, 1989, pp. 33–43.
- [3] M. Bellare and P. Rogaway, “Random oracles are practical: A paradigm for designing efficient protocols,” in *Proceedings of the 1st ACM conference on Computer and communications security*. ACM, 1993, pp. 62–73.
- [4] J. Kilian, “On the complexity of bounded-interaction and non-interactive zero-knowledge proofs,” in *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on*. IEEE, 1994, pp. 466–477.
- [5] A. Fiat and A. Shamir, “How to prove yourself: Practical solutions to identification and signature problems,” in *Conference on the Theory and Application of Cryptographic Techniques*. Springer, 1986, pp. 186–194.
- [6] I. Damgård, “Collision free hash functions and public key signature schemes,” in *EUROCRYPT*, ser. Lecture Notes in Computer Science, vol. 304. Springer, 1987, pp. 203–216.
- [7] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby, “A pseudorandom generator from any one-way function,” *SIAM Journal on Computing*, vol. 28, no. 4, pp. 1364–1396, 1999.
- [8] O. Goldreich, S. Goldwasser, and S. Micali, “How to construct random functions,” *J. ACM*, vol. 33, no. 4, pp. 792–807, 1986.
- [9] J. Rompel, “One-way functions are necessary and sufficient for secure signatures,” in *STOC*. ACM, 1990, pp. 387–394.
- [10] M. Naor, “Bit commitment using pseudorandomness,” *Journal of cryptology*, vol. 4, no. 2, pp. 151–158, 1991.
- [11] O. Goldreich, S. Micali, and A. Wigderson, “Proofs that yield nothing but their validity or all languages in np have zero-knowledge proof systems,” *Journal of the ACM (JACM)*, vol. 38, no. 3, pp. 690–728, 1991.
- [12] A. C.-C. Yao, “How to generate and exchange secrets,” in *Foundations of Computer Science, 1986., 27th Annual Symposium on*. IEEE, 1986, pp. 162–167.
- [13] Y. Lindell and B. Pinkas, “A proof of security of Yao’s protocol for two-party computation,” *Journal of Cryptology*, vol. 22, no. 2, pp. 161–188, 2009.
- [14] D. R. Simon, “Finding collisions on a one-way street: Can secure hash functions be based on general assumptions?” in *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 1998, pp. 334–345.
- [15] G. Asharov and G. Segev, “Limits on the Power of Indistinguishability Obfuscation and Functional Encryption,” in *2015 IEEE 56th Annual Symposium on Foundations of Computer Science, 2015*. [Online]. Available: <https://eprint.iacr.org/2015/341.pdf>
- [16] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. Vadhan, and K. Yang, “On the (im) possibility of obfuscating programs,” in *Annual International Cryptology Conference – CRYPTO 2001*. Springer, 2001, pp. 1–18, journal version appears in JACM 2012.
- [17] S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters, “Candidate Indistinguishability Obfuscation and Functional Encryption for all Circuits,” in *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, 2013, pp. 40–49. [Online]. Available: <https://eprint.iacr.org/2013/451.pdf>
- [18] R. Canetti, O. Goldreich, and S. Halevi, “The random oracle methodology, revisited,” *Journal of the ACM (JACM)*, vol. 51, no. 4, pp. 557–594, 2004.
- [19] S. Goldwasser and Y. T. Kalai, “On the (in) security of the fiat-shamir paradigm,” in *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*. IEEE, 2003, pp. 102–113.
- [20] R. Canetti, Y. Chen, and L. Reyzin, “On the correlation intractability of obfuscated pseudorandom functions,” in *Theory of Cryptography Conference*. Springer, 2016, pp. 389–415.
- [21] Y. T. Kalai, G. Rothblum, and R. Rothblum, “From Obfuscation to the Security of Fiat-Shamir for Proofs,” in *Advances in Cryptology – CRYPTO 2017*, 2016. [Online]. Available: <https://eprint.iacr.org/2016/303.pdf>
- [22] R. Canetti, Y. Chen, L. Reyzin, and R. Rothblum, “Fiat-shamir and correlation intractability from strong kdm-secure encryption,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques – EUROCRYPT 2018*. Springer, 2018.
- [23] M. Zhandry, “The magic of elfs,” in *Proceedings, Part I, of the 36th Annual International Cryptology Conference on Advances in Cryptology—CRYPTO 2016-Volume 9814*. Springer-Verlag New York, Inc., 2016, pp. 479–508.
- [24] Y. Dodis, A. Jain, T. Moran, and D. Wichs, “Counterexamples to hardness amplification beyond negligible,” in *Theory of Cryptography Conference*. Springer, 2012, pp. 476–493.
- [25] D. Wichs, personal communication, April 2018.
- [26] H. Corrigan-Gibbs and D. Kogan, “The discrete-logarithm problem with preprocessing,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques – EUROCRYPT 2018*, 2018.
- [27] V. Shoup, “Lower bounds for discrete logarithms and related problems,” in *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 1997, pp. 256–266.
- [28] I. Komargodski, M. Naor, and E. Yogev, “White-box vs. black-box complexity of search problems: Ramsey and graph property testing,” in *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*. IEEE, 2017, pp. 622–632.
- [29] I. Berman, A. Degwekar, R. D. Rothblum, and P. N. Vasudevan, “Multi collision resistant hash functions and their applications,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques – EUROCRYPT 2018*. Springer, 2018.
- [30] N. Bitansky, Y. T. Kalai, and O. Paneth, “Multi-collision resistance: A paradigm for keyless hash functions,” in *STOC*, 2018.
- [31] I. Komargodski, M. Naor, and E. Yogev, “Collision resistant hashing for paranoids: Dealing with multiple collisions,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques – EUROCRYPT 2018*. Springer, 2018.
- [32] M. Bellare, T. Ristenpart, and S. Tessaro, “Multi-instance security and its application to password-based cryptography,” in *CRYPTO*, ser. Lecture Notes in Computer Science, vol. 7417. Springer, 2012, pp. 312–329.
- [33] A. Rosen and G. Segev, “Chosen-ciphertext security via correlated products,” in *Theory of Cryptography Conference*. Springer, 2009, pp. 419–436.
- [34] E. Biham, Y. J. Goren, and Y. Ishai, “Basing weak public-key cryptography on strong one-way functions,” in *TCC*, ser. Lecture Notes in Computer Science, vol. 4948. Springer, 2008, pp. 55–72.

- [35] N. Bitansky and H. Lin, “One-message zero knowledge and non-malleable commitments,” 2018. [Online]. Available: <https://eprint.iacr.org/2018/613>
- [36] Y. Yu, J. Zhang, J. Weng, C. Guo, and X. Li, “Learning parity with noise implies collision resistant hashing,” 2017, <https://eprint.iacr.org/2017/1260.pdf>.
- [37] Z. Brakerski, V. Lyubashevsky, V. Vaikuntanathan, and D. Wichs, “Cryptographic hashing and worst-case hardness for lpn via code smoothing,” 2018. [Online]. Available: <https://eprint.iacr.org/2018/279>
- [38] N. Bitansky, D. Dachman-Soled, S. Garg, A. Jain, Y. T. Kalai, A. López-Alt, and D. Wichs, “Why Fiat-shamir for proof-sâ€š lacks a proof,” in *Theory of Cryptography*. Springer, 2013, pp. 182–201.
- [39] N. Bitansky, A. Degwekar, and V. Vaikuntanathan, “Structure vs. hardness through the obfuscation lens,” in *Annual International Cryptology Conference – CRYPTO 2017*. Springer, 2017, pp. 696–723.
- [40] C. Peikert and B. Waters, “Lossy trapdoor functions and their applications,” *SIAM Journal on Computing*, vol. 40, no. 6, pp. 1803–1844, 2011.
- [41] S. Goldwasser and G. N. Rothblum, “On best-possible obfuscation,” in *Theory of Cryptography Conference*. Springer, 2007, pp. 194–213.
- [42] R. Canetti, “Towards realizing random oracles: Hash functions that hide all partial information,” *IACR Cryptology ePrint Archive*, vol. 1997, p. 7, 1997.
- [43] R. Canetti, D. Micciancio, and O. Reingold, “Perfectly one-way probabilistic hash functions (preliminary version),” in *STOC*. ACM, 1998, pp. 131–140.
- [44] H. Wee, “On obfuscating point functions,” in *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*. ACM, 2005, pp. 523–532.
- [45] S. Goldwasser and Y. T. Kalai, “Cryptographic assumptions: A position paper,” in *Theory of Cryptography Conference*. Springer, 2016, pp. 505–522.