

MIT Open Access Articles

The Magic of ELFs

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

As Published: <https://doi.org/10.1007/s00145-018-9289-9>

Publisher: Springer US

Persistent URL: <https://hdl.handle.net/1721.1/131500>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



The Magic of ELFs

MARK ZHANDRY
MIT & Princeton
mzhandry@princeton.edu

Abstract

We introduce the notion of an *Extremely Lossy Function* (ELF). An ELF is a family of functions with an image size that is tunable anywhere from injective to having a polynomial-sized image. Moreover, for any efficient adversary, for a sufficiently large polynomial r (necessarily chosen to be larger than the running time of the adversary), the adversary cannot distinguish the injective case from the case of image size r .

We develop a handful of techniques for using ELFs, and show that such extreme lossiness is useful for instantiating random oracles in several settings. In particular, we show how to use ELFs to build secure point function obfuscation with auxiliary input, as well as polynomially-many hardcore bits for any one-way function. Such applications were previously known from strong knowledge assumptions — for example polynomially-many hardcore bits were only known from differing inputs obfuscation, a notion whose plausibility has been seriously challenged. We also use ELFs to build a simple hash function with *output intractability*, a new notion we define that may be useful for generating common reference strings.

Next, we give a construction of ELFs relying on the *exponential* hardness of the decisional Diffie-Hellman problem, which is plausible in elliptic curve groups. Combining with the applications above, our work gives several practical constructions relying on qualitatively different — and arguably better — assumptions than prior works.

1 Introduction

Hash functions are a ubiquitous tool in cryptography: they are used for password verification, proofs of work, and are central to a variety of cryptographic algorithms including efficient digital signatures and encryption schemes.

Unfortunately, formal justifications of many of the uses of hash functions have been elusive. The trouble stems from the difficulty of even defining what security properties a hash function should satisfy. On one extreme, a hash function can be assumed to have standard security properties such as one-wayness or collision resistance, which are useless for most of the applications above. On the other extreme, a hash function can be modeled as a truly random function, where it is assumed that an adversary only has black-box access. In the so-called random oracle model (ROM) [BR93], all of the above applications are secure. However, random oracles clearly do not exist and moreover provably cannot be replaced by any concrete hash function [CGH98]. In this light, it is natural to ask:

*What are useful properties of random oracles
that can be realized by real-world hash functions.*

Some attempts have been made to answer this question; however, many such attempts have serious limitations. For example Canetti, Goldreich, and Halevi [CGH98] propose the notion of *correlation intractability* as a specific feature of random oracles that could potentially have a standard model instantiation. However, they show that for some parameter settings such standard model hash functions cannot exist. The only known positive example [CCR16] relies on extremely strong cryptographic assumptions such as general-purpose program obfuscation. For another example, Bellare, Hoang, and Keelveedhi [BHK13] define a security property for hash functions called Universal Computational Extractors (UCE), and show that hash functions with UCE security suffice for several uses of the random oracle model, such as deterministic public key encryption, polynomially-many hardcore bits for any one-way function, point obfuscation, and more. While UCE’s present an important step toward understanding which hash function properties might be achievable and which are not, UCE’s have several limitations. For example, the formal definition of a UCE is somewhat complicated to even define. Moreover, UCE is not a single property, but a family or “framework” of assumptions. The most general form of UCE is trivially unattainable, and some of the natural restricted classes of UCE which have been used for applications have been challenged [BFM14, BST16]. Therefore, it is unclear which versions of UCE should be trusted and which untrusted.

Similar weaknesses have been shown for other strong assumptions that can be cast as families of assumptions or as knowledge/extracting assumptions, such as extractable one-way functions (eOWFs) [BCPR14] and differing inputs obfuscation (diO) [BCP14, ABG⁺13, GGHW14]. Roughly, UCEs, eOWFs, and diO are assumptions of the form “for any efficient adversary A solving a task from a certain task family, there exists an efficient adversary B solving a related task from a different task family,” without giving an explicit means to convert A to B . Without an explicit conversion procedure that works for every single task, it is not clear whether such assumptions should be possible at all. In contrast, standard assumptions in cryptography are of the form “there do not exist any efficient adversaries solving a one particular task.”

The weakness of UCE, eOWFs, and diO are part of a general pattern for strong assumptions that are not specified by a cryptographic game. In particular, these assumptions do not meet standard notions of falsifiability ([Nao03, GW11]), and are not *complexity assumptions* in the sense of Goldwasser and Kalai [GK16]. It is difficult to analyze and gain confidence in such non-standard assumptions, and such assumptions often turn out to be implausible [BFM14, BST16, BCPR14, GGHW14]. We stress that knowledge/extracting/framework assumptions are highly desirable as security *properties*; however, making the *assumption* that these properties hold without formal justification is risky. Instead, in order to trust that the property actually holds, it should be derived from a trusted assumption. Therefore, an important question in this space is the following:

Are there primitives based on “nice” (e.g. simple, well-established, game-based, falsifiable, complexity assumption, etc) assumptions that can be used to build hash functions suitable for instantiating random oracles for many protocols.

1.1 Our Work

Our Random Oracle Targets. In this work, we aim to base several applications of random oracles on concrete, “nice” assumptions with relatively simple instantiations.

- **Boosting selective to adaptive security.** A trivial application of random oracles is to boost selective to adaptive security in the case of signatures and identity-based encryption. This is

done by first hashing the message/identity with the random oracle before signing/generating secret keys. There has been no standard-model security notion for hash functions that allows for this conversion to be secure, though in the case of signatures, chameleon hash functions [KR00] achieve this conversion with a small tweak to the hash-and-sign approach.

- **Password hashing.** Another common use of hash functions is to securely store a password in “encrypted” form, allowing for the password to be verified, but hiding the actual password in case the password database is compromised. This use case is a special instance of *point obfuscation* (PO). For password distributions with polynomial min-entropy, point obfuscation can be constructed from general primitives such as regular collision resistant hashing [CMR98]. Sometimes, we would like even lower-entropy distributions, optimally as low as super-logarithmic. In this case, Wee [Wee05] shows a construct based on a strong one-way permutation assumption. In the case that there may be side information about the password, we have the notion of *auxiliary input* point obfuscation (AIPO). The only prior constructions of AIPO [Can97, BP12] rely on very strong knowledge assumptions. The first is Canetti’s [Can97] strong knowledge variant of the decisional Diffie Hellman (DDH) assumption, which is an assumption family and therefore shares similar weaknesses to the strong assumptions discussed above. In particular, the plausibility of this assumption has been called into question by a recent work showing it is incompatible with the existence of certain strong forms of obfuscation [BST16]. The second is a strong knowledge assumption about one-way permutations due to Bitansky and Paneth [BP12], which is a strengthening of Wee’s assumption and again is a family of knowledge assumptions. To the best of our knowledge, the only currently known ways to instantiate the [BP12] assumption is to make the tautological assumption that a particular one-way permutation is secure. For reasons mentioned above, such tautological knowledge assumptions are generally considered undesirable in cryptography.
- **Generating system parameters.** A natural use case of hash functions is for generating common random strings (crs) in a trusted manner. More specifically, suppose a (potentially untrusted) authority is generating a crs for some protocol. Unfortunately, such a crs may admit a “trapdoor” that allows for breaking whatever protocol is using it (Dual_EC_DRBG is a prominent example of this). In order to ensure to untrusting parties that no trapdoor is known, the authority will generate the crs as an output of the hash function on some input. The authority may have some flexibility in choosing the input; we wish to guarantee that it cannot find an input such that it also knows a trapdoor for the corresponding output. In the random oracle model, this methodology is sound: the authority cannot choose an input so that it knows the trapdoor for the output. However, standard notions of security for hash functions do not preclude the possibility of the adversary choosing a strange input to the hash function so that it knows a trapdoor for the output. We propose (Section 5) the notion of *output intractability* as a standard-model security notion that captures this use case. Output intractability is related to, but incomparable with, the notion of correlation intractability mentioned above. As an assumption, our notion of output intractability takes the form of a knowledge assumption on hash functions; no construction based on “nice” assumptions is currently known.
- **Hardcore bits for any one-way function.** A random oracle serves as a good way to extract many hardcore bits for any one-way function. This fact gives rise to a simple public-key encryption scheme from trapdoor permutations. While it is known how to extract many

hardcore bits for specific functions [SS90, PS98, AGV09], extracting many bits for general one-way functions may be useful in settings where we cannot freely choose the function, such as if the function is required to be a trapdoor permutation. Unfortunately, for general one-way functions, the only known way to extract more than a logarithmic number of hardcore bits is to use very strong (and questionable [GGHW14]) knowledge assumptions: differing inputs obfuscation [BST14] (plus one-way functions) or extractable witness PRFs [Zha16]. In the case of *injective* one-way functions, Bellare, Stepanovs, and Tessaro [BST14] show that the weaker assumption of *indistinguishability* obfuscation (iO) (plus one-way functions) suffices. While weaker than diO, iO is still one of the strongest assumptions made in cryptography. Either way, the forms of obfuscation required are also completely impractical [AHKM14]. Another limitation of prior constructions is that randomness used to sample the hardcore function needs to be kept secret.

- **Instantiating Full Domain Hash (FDH) signatures.** Finally, we consider using random oracles to instantiate the Full Domain Hash (FDH) protocol transforming trapdoor permutations into signatures. Hohenberger, Sahai, and Waters [HSW14] show that (indistinguishability) obfuscating a (puncturable) pseudorandom function *composed with the permutation* is sufficient for FDH signatures. However, their proof has two important limitations. First, the resulting signature scheme is only selectively secure. Second, the instantiation depends on the particular trapdoor permutation used, as well as the public key of the signer. Thus, each signer needs a separate hash function, which needs to be appended to the signer’s public keys. To use their protocol, everyone will therefore need to publish new keys, even if they already have published keys for the trapdoor permutation.

Our approach. We take a novel approach to addressing the questions above. We isolate a (generally ignored) property of random oracles, namely that random oracles are indistinguishable from functions that are extremely lossy. More precisely, the following is possible in the random oracle model. Given any polynomial time oracle adversary \mathcal{A} and an inverse polynomial δ , we can choose the oracle such that (1) the image size of the oracle is a polynomial r (even for domain/range sizes where a truly random oracle will have exponential image size w.h.p.), and (2) \mathcal{A} cannot tell the difference between such a lossy oracle and a truly random oracle, except with advantage smaller than δ . Note that the tuning of the image size must be done with knowledge of the adversary’s running time — an adversary running in time $O(\sqrt{r})$ can with high probability find a collision, thereby distinguishing the lossy function from a truly random oracle. However, by setting \sqrt{r} to be much larger than the adversary’s running time, the probability of finding a collision diminishes. We stress that any protocol would still use a truly random oracle and hence not depend on the adversary; the image size tuning would only appear in the security proof. Our observation of this property is inspired by prior works of Boneh and Zhandry [Zha12, BZ13], who use it for the entirely different goal of giving security proofs in the so-called *quantum* random oracle model (random oracle instantiation was not a goal nor accomplishment of these prior works).

We next propose the notion of an *Extremely Lossy Function (ELF)* as a standard-model primitive that captures this tunable image size property. The definition is related to the notion of a lossy *trapdoor* function due to Peikert and Waters [PW08], with two important differences: we do not need any trapdoor, giving hope that ELFs could be constructed from symmetric primitives. On the other hand, we need the functions to be much, much more lossy, as standard lossy functions still have exponential image size.

On the surface, extreme lossiness without a trapdoor does not appear incredibly useful, since many interesting applications of standard lossy functions (e.g. (CCA-secure) public key encryption) require a trapdoor. Indeed, using an ELF as a hash function directly does not solve most of the tasks outlined above. Perhaps surprisingly, we show that this extremely lossy property, in conjunction with other tools — usually pairwise independence — can in fact be quite powerful, and we use this power to give new solutions to each of the tasks above. Our results are as follows:

- (Section 3) We give a practical construction of ELFs assuming the *exponential* hardness of the decisional Diffie-Hellman (DDH) problem: roughly, that the best attack on DDH for groups of order p takes time $O(p^c)$ for some constant c . Our construction is based on the lossy trapdoor functions due to Peikert and Waters [PW08] and Freeman et al. [FGK⁺10], though we do not need the trapdoor from those works. Our construction starts from a trapdoor-less version of the DDH-based construction of [FGK⁺10], and iterates it many times at different security levels, together with pairwise independent hashing to keep the output length from growing too large. Having many different security levels allows us to do the following: when switching the function to be lossy, we can do so at a security level that is just high enough to prevent the particular adversary from detecting the switch. Using the exponential DDH assumption, we show that the security level can be set low enough so that (1) the adversary cannot detect the switch, and (2) so that the resulting function has polynomial image size.

We note that a couple prior works [BH12, BHJ⁺13, DS15] have used a similar technique of combining several “bounded adversary” instances at multiple security levels, and invoking the security of the instance with “just high enough” security. The main difference is that in prior works, “bounded adversary” refers to bounded queries, and the security parameter itself is kept constant across instances; in our work, “bounded adversary” refers to bounding the running time of the adversary, and the security parameter is what is varied across instances.

Our iteration at multiple security levels is somewhat generic and would potentially apply to other constructions of lossy functions, such as those based on LWE. However, LWE-based constructions of lossy functions are not quite lossy enough for our needs since even “exponentially secure” LWE can be solved in time sub-exponential in the length of the secret.

The exponential hardness of DDH is plausible on elliptic curve groups — despite over a decade of wide-spread use and cryptanalysis attempts, there are virtually no non-trivial attacks on most elliptic curve groups and the current best attacks on DDH take time $\Omega(p^{1/2})$. In fact, the parameter settings for most real-world uses of the Diffie-Hellman problem are set assuming the Diffie-Hellman problem takes exponential time to solve. If our assumption turns out to be false, it would have significant ramifications as it would suggest that parameters for many cryptosystems in practice are set too aggressively. It would therefore be quite surprising if DDH turned out to *not* be exponentially hard on elliptic curves. While not a true falsifiable assumption in the sense of Naor [Nao03] or Gentry and Wichs [GW11] due to the adversary being allowed to run in exponential time, the exponential DDH assumption is falsifiable in spirit and naturally fits within the complexity assumption framework of Goldwasser and Kalai [GK16].

While our ELFs are built from public key tools, we believe such tools are unnecessary and we leave as an interesting open question the problem of constructing ELFs from symmetric or generic tools.

We observe that our construction achieves a public coin notion, which is useful for obtaining public coin hash functions in applications¹.

- We give several different hash function instantiations based on ELF’s ranging in complexity and the additional tools used. In doing so, we give new solutions to each of the problems above. Each construction uses the ELF’s in different ways, and we develop new techniques for the analysis of these constructions. Thus we give an initial set of tools for using ELF’s that we hope to be useful outside the immediate scope of this work.

- The simplest instantiation is just to use an ELF itself as a hash function. Such a function can be used to generically boost selective security to adaptive security in signatures and identity-based encryption by first hashing the message/user identity (more details below).
- (Section 4) The next simplest instantiation is to pre-compose the ELF with a pairwise independent hash function. This function gives rise to a simple (public coin) point function obfuscation (PO). Proving this uses a slight generalization of the “crooked leftover hash lemma” [DS05].
- (Section 5) A slightly more complicated instantiation is given by *post*-composing an ELF with a k -wise independent function. We show that this construction satisfies our notion of *output intractability*. It is moreover public coin.
- (Section 6) We then give an even more complicated construction, though still using ELF’s as the only underlying source of cryptographic hardness. The construction roughly follows a common paradigm used in leakage resilience [DP08]: apply a computational step (in our case, involving ELF’s), compress with pairwise independence, and repeat. We note however that the details of the construction and analysis are new to this work.

We demonstrate that our construction is a pseudorandom generator attaining a very strong notion of leakage resilience for the seed. This property generalizes the strong variant of one-wayness of Bitansky and Paneth [BP12]. Our construction therefore shows how to instantiate the knowledge properties conjectured in their work using a more traditional-style assumption.

An immediate consequence of our generator requirement is a (public coin) point function obfuscation that is secure even in the presence of auxiliary information (AIPO), which was previously known from either *permutations* satisfying [BP12]’s one-wayness requirement (our function is *not* a permutation), or from Canetti’s strong knowledge variant of DDH [Can97, BP12]². Our AIPO construction is qualitatively very different from these existing constructions, and when plugging in our ELF construction, again relies on just exponential DDH.

Our generator also immediately gives a family of (public coin) hardcore functions of arbitrary stretch for any one-way function. Unlike the previous obfuscation-based solutions, ours is practical, and public coin, and ultimately based on a well-studied game-based assumption.

¹The construction of [FGK⁺10] can also be made public coin by tweaking the generation procedure. However, this necessarily loses the trapdoor, as having a trapdoor and being public coin are incompatible. To the best of our knowledge, however, we are the first to observe this public coin feature.

²One drawback of our construction — which is shared with some of the prior constructions — is that we achieve a relaxed notion of correctness where for some sparse “bad” choices of the obfuscation randomness, the outputted program may compute the wrong function.

Our analysis also demonstrates that our ELF-based function can be used in a standard random oracle public key encryption protocol [BR93].

- (Section 7) Finally, we give an instantiation useful for Full Domain Hash signatures which involves obfuscating the composition of an ELF and a (puncturable) pseudorandom function (but not any trapdoor permutation) using an indistinguishability obfuscator. Since we use obfuscation as in Hohenberger, Sahai, and Waters’ [HSW14] scheme, this construction is still completely impractical and therefore currently only of theoretical interest. We show that our construction can be used in the FDH protocol, solving some of the limitations in [HSW14]. In particular, by composing with an ELF, we immediately get adaptive security as observed above. Our construction is moreover independent of the permutation (except for the size of the circuit computing it), and is also independent of the signer’s public key. Thus, our instantiation is universal and one instantiation can be used by any signer, even using existing keys. Similar to [HSW14], this construction is still required to be secret coin, even if the underlying ELF is public coin.

Note that most of the above applications can be attained using an appropriate UCE [BHK13]. However, as noted above, UCE’s require assuming a large family of non-standard assumptions, some members of which have been challenged [BFM14, BST16]. Our ELF-based constructions give an alternative instantiation that requires only making a single standard-style assumption, namely exponential DDH, or perhaps a few assumptions in the case of Full Domain Hash signatures.

Warm up: generically boosting selective to adaptive security. To give a sense for our techniques, we show how ELFs can be used to generically boost selective to adaptive security in signatures and identity-based encryption. We demonstrate the case for signatures; the case for identity based encryption is almost identical.

Recall that in selective security for signatures, the adversary commits to a message m^* at the beginning of the experiment before seeing the public key. Then the adversary makes a polynomial number q of adaptive signing queries on messages m_1, \dots, m_q , receiving signatures $\sigma_1, \dots, \sigma_q$. It is required that $m_i \neq m^*$ for any i . Then, the adversary produces a forged signature σ^* on m^* , and security states that σ^* is with overwhelming probability invalid for any efficient adversary. Adaptive security, in contrast, allows the adversary to choose m^* potentially *after* the q adaptive queries.

We now convert selective to adaptive security using ELFs: first hash the message using the ELF, and then sign. Adaptive security is proved through a sequence of hybrids. The first is the standard adaptive security game above. Toward contradiction, suppose that the adversary runs in polynomial time t and succeeds in forging a signature on m^* with non-negligible probability ϵ . Let δ be an inverse polynomial that lower bounds ϵ infinitely often. In the second hybrid, the ELF is selected to have polynomial image size r , where $r \geq 2q$ is chosen, say, so that no t -time adversary can distinguish between this ELF and an injective ELF, except with probability at most $\delta/2$. Thus, in this hybrid, the adversary still successfully forges with probability $\epsilon - \delta/2$. This is lower bounded by $\delta/2$ infinitely often, and is therefore non-negligible.

In the next hybrid, at the beginning of the experiment, one of the r image points of the ELF, y^* , is chosen at random³. Then we abort the experiment if the adversary’s chosen m^* does not hash to

³The ability to sample a random image point does not follow immediately from our basic ELF definition, though this can be done in our construction. If the ELF is regular, then this can be accomplished by sampling a random

y^* ; with probability $1/r$, we do not abort⁴. This abort condition is independent of the adversary’s view, meaning that we do not abort, and the adversary successfully forges, with probability at least $(\epsilon - \delta/2)/r$, which again is non-negligible. Notice now that y^* can be chosen at the beginning of the experiment. This is sufficient for obtaining an adversary for the selective security of the original signature scheme.

1.2 Complexity Absorption

Assumptions come in a spectrum, long standing assumptions that are widely accepted at one end, and new suspect assumptions at the other. It may therefore be more reasonable to assume the sub-exponential or exponential hardness of an existing well-studied problem than to assume such hardness for new and untested problems. Moreover, there might be implementation issues (such as having to re-publish longer keys, see Section 7 for a setting where this could happen) that make assuming the sub-exponential hardness of certain primitives undesirable.

The selective-to-adaptive security result above can be seen as an instance of a more general task of *complexity absorption*, where an extra complexity-absorbing primitive (in our case, and ELF) is introduced into the protocol. The original building blocks of the protocol (the underlying signature/identity-based encryption in this case) can be reduced from (sub)exponential security to polynomial security. Meanwhile, the complexity-absorbing primitive may still require exponential hardness as in our case, but hopefully such hardness is a reasonable assumption. Our hardcore function with arbitrary span can also be seen in this light: it is straightforward to extend Goldreich-Levin [GL89] to a hardcore function of polynomial span for exponentially-secure one-way functions. By introducing an ELF into the hardcore function, the ELF can absorb the complexity required of the one-way function, yielding a hardcore function for *any* one-way function, even one-way functions that are only polynomially secure. Similarly, our random oracle instantiation for Full Domain Hash can also be seen as an instance of complexity absorption.

Thus, our work can be seen as providing an initial set of tools and techniques for the task of complexity absorption that may be useful in other settings where some form of sub-exponential hardness is difficult or impossible to avoid. For example, Rao [Rao14] argues that any proof of adaptive security for multiparty non-interactive key exchange (NIKE) will likely incur an exponential loss. As all current multiparty NIKE protocols are built from multilinear maps or obfuscation, which in turn rely on new, untested (and in many cases broken) hardness assumptions, assuming the sub-exponential security of the underlying primitives to attain adaptive security is undesirable. Hofheinz et al. [HJK⁺14] give a construction in the random oracle model that only has a polynomial loss; our work gives hope that a standard model construction based on ELFs may be possible where the ELF is the only primitive that needs stronger than polynomial hardness.

For a more ambitious example, Garg, Gentry, Sahai, and Waters [GGSW13] argue that any proof of security of witness encryption and obfuscation will also likely incur an exponential loss. This suggests that at least one of the underlying assumptions (those on multilinear maps) will need to be sub-exponentially secure. An intriguing possibility is to use ELFs (or some other (sub)exponentially secure primitive with better-understood security) to absorb the complexity, yielding an obfuscation

input to the ELF and then applying the ELF. More generally, if it is possible to efficiently enumerate all the image points, then randomly sampling an image point is easy. Of course, enumerating all the image points will take time at least r , which is larger than the running time of the adversary, but can still potentially be done efficiently.

⁴We also need to abort if any of the m_i do hash to y_i . It is straightforward to show that we still do not abort with probability at least $\frac{1}{2r}$.

construction from a constant number of assumptions *without* assuming the sub-exponential hardness of multilinear maps.

1.3 Non-black box simulation

Our proofs inherently require knowledge of the adversary’s running time (and success probability). Thus, they do not make black box use of the adversary. Yet, this is the only non-black box part of our proofs — the reduction does not need to know the description or internal workings of the adversary. This is similar to the case of Goldreich-Levin [GL89], where only the adversary’s success probability is needed. Thus our reductions are very nearly black box, while potentially giving means to circumvent black-box impossibility results. For example, proving the security of AIPO is known to require non-black box access to the adversary [Wee05, BP12], and yet our reduction proves the security of AIPO knowing only the adversary’s running time and success probability. We leave it as an interesting open question to see if your techniques can be used to similarly circumvent other black box impossibility results.

1.4 On the minimal assumptions needed to build ELF

We show how to construct extremely lossy functions from a specific assumption on elliptic curve groups. One could also hope for generic constructions of ELF based on existing well-studied primitives. Unfortunately, this appears to be a difficult task, and there are several relevant black-box barriers to constructing ELF. For example, lossy functions (even standard ones) readily imply collision resistance [PW08], which cannot be built from one-way functions in a black-box fashion [Sim98]. Rosen and Segev [RS09] show a similar separation from functions that are secure under correlated products. Pietrzak, Rosen, and Segev [PRS12] show that efficiently amplifying lossiness in a black box way is impossible — this suggests that building ELF from standard lossy functions will be difficult, if not impossible.

Perhaps an even stronger barrier to realizing ELF from standard assumptions is the following. Our assumption, unfortunately, is about *exponential*-time adversaries, as opposed to typical assumptions about polynomial-time adversaries. One could hope for basing ELF on standard polynomial assumptions, such as polynomial DDH. However, we now argue that this would require major breakthroughs in complexity theory. Indeed, lossy and injective modes of an ELF can be distinguished very efficiently using a super-logarithmic amount of non-determinism as follows. Let $D = [2^{\omega(\log m)}]$ where m is the number of input bits to the ELF. In the injective mode, there will be no collisions when the domain is restricted to D . However, in the lossy mode for *any* polynomial image size $r = r(m)$, there is guaranteed to be a collision in D . Points in D can be specified by $\omega(\log m)$ bits. Therefore, we can distinguish the two modes by non-deterministically guessing two inputs in D (using $\omega(\log m)$ bits of non-determinism) and checking that they form a collision. Therefore, if NP restricted to some super-logarithmic amount of non-determinism was solvable in polynomial time, then this algorithm could be made efficient while removing all non-determinism. Such an algorithm would violate ELF security.

Theorem 1.1. *If ELF exist, then for any super-logarithmic function t , NP with t bits of non-determinism is not solvable in polynomial time.*

Therefore, it seems implausible to base ELF on any polynomially-secure primitive, since it is consistent with our current knowledge that NP with, say, \log^2 bits of non-determinism is solvable in

polynomial time, but polynomially-secure cryptographic primitives exist.

The above may seem to suggest that ELF's are too strong of a starting point for our applications; to the contrary, we argue that for most of our applications — point functions⁵ (Section 4), output intractability (Section 5), and polynomially-many hardcore bits for any one-way function (Section 6) — similar barriers exist to realizing those applications. Therefore, this limitation of ELF's is actually inherent to any primitive strong enough to realize the applications.

Therefore, instead of starting from standard polynomially-secure primitives, we may hope to build ELF's generically from, say, an exponentially secure primitive which has a similar limitation. Can we build ELF's from exponentially secure (injective) one-way functions? Exponentially-secure collision resistant hash functions? To what extent do the black-box barriers above extend into the regime of exponential hardness? We leave these as interesting open questions for future work.

2 Preliminaries

Given integers $a, b, a \leq b$, we define $[a, b]$ as the set $\{a, a + 1, \dots, b - 1, b\}$. For integer $a > 1$, we define $[a] = [1, a]$. Given a distribution \mathcal{D} over a set \mathcal{X} , define the support of \mathcal{D} , $\text{Supp}(\mathcal{D})$, to be the set of points in \mathcal{X} that occur with non-zero probability. For any $x \in \mathcal{X}$, let $\Pr[\mathcal{D} = x]$ be the probability that \mathcal{D} selects x . For any set \mathcal{X} , define $U_{\mathcal{X}}$ to be the uniform distribution on \mathcal{X} . Define the collision probability of \mathcal{D} to be

$$CP(\mathcal{D}) = \Pr[x_1 = x_2 : x_1, x_2 \leftarrow \mathcal{D}] = \sum_{x \in \mathcal{X}} \Pr[\mathcal{D} = x]^2 .$$

Given two distributions $\mathcal{D}_1, \mathcal{D}_2$, define the statistical distance between \mathcal{D}_1 and \mathcal{D}_2 to be

$$\Delta(\mathcal{D}_1, \mathcal{D}_2) = \frac{1}{2} \sum_{x \in \mathcal{X}} |\Pr[\mathcal{D}_1 = x] - \Pr[\mathcal{D}_2 = x]| .$$

Suppose $\text{Supp}(\mathcal{D}_1) \subseteq \text{Supp}(\mathcal{D}_2)$. Define the Rényi Divergence between \mathcal{D}_1 and \mathcal{D}_2 to be

$$RD(\mathcal{D}_1, \mathcal{D}_2) = \sum_{x \in \text{sup}(\mathcal{D}_1)} \frac{\Pr[\mathcal{D}_1 = x]^2}{\Pr[\mathcal{D}_2 = x]} \quad \text{6} .$$

The Rényi divergence is related to the statistical distance via the following lemma:

Lemma 2.1. *For any distributions $\mathcal{D}_1, \mathcal{D}_2$ over a set \mathcal{Z} such that $\text{Supp}(\mathcal{D}_1) \subseteq \text{Supp}(\mathcal{D}_2)$,*

$$\Delta(\mathcal{D}_1, \mathcal{D}_2) \leq \frac{1}{2} \sqrt{RD(\mathcal{D}_1, \mathcal{D}_2) - 1} .$$

Consider a distribution \mathcal{H} over the set of functions $h : \mathcal{X} \rightarrow \mathcal{Y}$. We say that \mathcal{H} is *pairwise independent* if, for any $x_1 \neq x_2 \in \mathcal{X}$, the random variables $\mathcal{H}(x_1)$ and $\mathcal{H}(x_2)$ are independent and identically distributed, though not necessarily uniform. Similarly define k -wise independence. We say that \mathcal{H} has *output distribution* \mathcal{D} if for all x , the random variable $\mathcal{H}(x)$ is identical to \mathcal{D} . Finally,

⁵The case of point functions is more or less equivalent to a similar result of Wee [Wee05].

⁶Often, the Rényi Divergence is defined to be proportional to the logarithm of this quantity. For our purposes, this representation of the divergence will be more convenient.

we say that \mathcal{H} is *uniform* if it has output distribution $U_{\mathcal{Y}}$ ⁷. We will sometimes abuse notation and say that a function h is a pairwise independent function (resp. uniform) if h is drawn from a pairwise independent (resp. uniform) distribution of functions. We will need the following lemma:

Lemma 2.2. *Let H be a pairwise independent and uniform family of hash functions with domain \mathcal{X} and codomain \mathcal{Y}_{i+1} . Let S be a subset of \mathcal{X} . Consider sampling an h from H . Then except with probability $|S|^2/2|\mathcal{X}|$, there is no pair of distinct inputs $x_0, x_1 \in S, x_0 \neq x_1$ such that $h(x_0) = h(x_1)$. In other words, the restriction of h to S is injective. In this case, we say that h is injective on S .*

Proof. For any pair $x_0, x_1 \in S$ of distinct inputs, the probability $h(x_0) = h(x_1)$ is $1/|\mathcal{Y}|$, by the pairwise independence and uniformity of H . The number of such pairs is $|S|(|S| - 1)/2 \leq |S|^2/2$. By the union bound, the probability that there exists a pair $x_0 \neq x_1$ with $h(x_0) = h(x_1)$ is at most $|S|^2/2|\mathcal{X}|$. \square

We will say that a (potentially probabilistic) algorithm \mathcal{A} outputting a bit b distinguishes two distributions $\mathcal{D}_0, \mathcal{D}_1$ with advantage ϵ if $\Pr[\mathcal{A}(\mathcal{D}_b) : b \leftarrow \{0, 1\}] \in [\frac{1}{2} - \epsilon, \frac{1}{2} + \epsilon]$. This is equivalent to the random variables $\mathcal{A}(\mathcal{D}_0)$ and $\mathcal{A}(\mathcal{D}_1)$ have 2ϵ statistical distance.

Unless otherwise stated, all cryptographic protocols will implicitly take a security parameter λ as input. Moreover, any sets (such as message spaces, ciphertext spaces, etc) will be implicitly indexed by λ , unless otherwise stated. In this context, when we say that an adversary is efficient, we mean its running time is polynomial in λ . A non-negative function ϵ is *negligible* if it is smaller than any inverse polynomial. When discussing cryptographic protocols, we say that a probability of an event or advantage of an adversary is negligible if it is negligible in λ . Two distributions $\mathcal{D}_0, \mathcal{D}_1$ (implicitly parameterized by λ) are computationally indistinguishable if any efficient algorithm has only negligible distinguishing advantage, and are statistically indistinguishable if the distributions have negligible statistical distance. In the statistical setting, we also sometimes say that $\mathcal{D}_0, \mathcal{D}_1$ are statistically close.

2.1 The Crooked Leftover Hash Lemma

Here we prove a slight generalization of the “crooked Leftover Hash Lemma” of Dodis and Smith [DS05].

Lemma 2.3. *Let H be a distribution on functions $h : \mathcal{X} \rightarrow \mathcal{Y}$ that is pairwise independent with output distribution \mathcal{E} , for some distribution \mathcal{E} that is possibly non-uniform. Let \mathcal{D} be an arbitrary distribution over \mathcal{X} . Then we have that*

$$\Delta((H, H(\mathcal{D})), (H, \mathcal{E})) \leq \frac{1}{2} \sqrt{CP(\mathcal{D})(|\text{Supp}(\mathcal{E})| - 1)}$$

Proof. Our proof will proceed analogously to the proof of the original Leftover Hash Lemma [ILL89]. [ILL89] proceeds by bounding the collision probability of $(H, H(\mathcal{D}))$, and then showing that distributions with low collision probability must be close to uniform. In general, the collision probability of a distribution \mathcal{F} is equivalent to the Rényi divergence $RD(\mathcal{F}, U)$ between \mathcal{F} and the uniform distribution, up to an overall scaling factor. Therefore, [ILL89] can be equivalently proved by bounding the Rényi divergence, and then relating the divergence to statistical distance using Lemma 2.1. This is the view we use for our proof.

⁷Note that the typical use of *pairwise independence* is equivalent to our notion of pairwise independence *plus* uniformity. For our purposes, it will be convenient to separate out the two properties.

Given the following claim and Lemma 2.1, Lemma 2.3 follows:

Claim 2.4. $RD((H, H(\mathcal{D})) , (H, \mathcal{E})) = 1 + CP(\mathcal{D})(|\text{Supp}(\mathcal{E})| - 1)$

We now prove the claim. First,

$$RD((H, H(\mathcal{D})) , (H, \mathcal{E})) = \sum_{h,y} \frac{\Pr[H = h \wedge h(\mathcal{D}) = y]^2}{\Pr[H = h]\Pr[\mathcal{E} = y]} = \sum_y \frac{\sum_h \Pr[H = h] \Pr[h(\mathcal{D}) = y]^2}{\Pr[\mathcal{E} = y]} .$$

Now, notice that $\Pr[h(\mathcal{D}) = y]^2 = \Pr[h(x_1) = y \wedge h(x_2) = y : x_1, x_2 \leftarrow \mathcal{D}]$, so that

$$\sum_h \Pr[H = h] \Pr[h(\mathcal{D}) = y]^2 = \Pr[H(x_1) = y \wedge H(x_2) = y : x_1, x_2 \leftarrow \mathcal{D}] .$$

Also, notice that $\Pr[E = y] = \Pr[H(x_1) = y]$ for any $x_1 \in \mathcal{X}$. Therefore, we can write

$$RD((H, H(\mathcal{D})) , (H, \mathcal{E})) = \sum_{x_1, x_2} \Pr[X = x_1] \Pr[X = x_2] \sum_y \Pr[H(x_1) = y | H(x_2) = y] .$$

Now we divide into two cases. If $x_1 = x_2$, then $\Pr[H(x_1) = y | H(x_2) = y] = 1$, so $\sum_y \Pr[H(x_1) = y | H(x_2) = y] = |\text{Supp}(\mathcal{E})|$. If $x_1 \neq x_2$, then by independence $\Pr[H(x_1) = y | H(x_2) = y] = \Pr[H(x_1) = y] = \Pr[\mathcal{E} = y]$, so $\sum_y \Pr[H(x_1) = y | H(x_2) = y] = \sum_y \Pr[\mathcal{E} = y] = 1$. Therefore,

$$\begin{aligned} RD((H, H(\mathcal{D})) , (H, \mathcal{E})) &= \sum_x \Pr[\mathcal{D} = x]^2 |\text{Supp}(\mathcal{E})| + \sum_{x_1 \neq x_2} \Pr[\mathcal{D} = x_1] \Pr[\mathcal{D} = x_2] \\ &= CP(\mathcal{D})|\text{Supp}(\mathcal{E})| + (1 - CP(\mathcal{D})) = 1 + CP(\mathcal{D})(|\text{Supp}(\mathcal{E})| - 1) . \end{aligned}$$

This completes the proof of the claim, completing the proof of Lemma 2.3. \square

3 Extremely Lossy Functions

Here, we define our notion of *extremely lossy functions*, or ELFs. A standard lossy function [PW08] is intuitively a function family with two modes: an injective mode where the function is injective, and a lossy mode where the image size of the function is much smaller than the domain. The standard security requirement is that no polynomial-time adversary can distinguish the two modes⁸.

An ELF is a lossy function with a much stronger security requirement. In the lossy mode, the image size can be taken to be a polynomial r . Clearly, such a lossy mode can be distinguished from injective by an adversary running in time $O(\sqrt{r})$ that simply evaluates the function on \sqrt{r} inputs, looking for a collision. Therefore, we cannot have security against arbitrary polynomial-time attackers. Instead, we require security against r^c -time attackers, for some $c \leq 1/2$. Moreover, we require that r is actually tunable, and can be chosen based on the adversary in question. This means that for *any* polynomial time attacker, we can set the lossy function to have domain r for some polynomial r , and the lossy function will be indistinguishable from injective to that particular attacker (note that the honest protocol will always use the injective mode, and therefore will not depend on the adversary in any way).

⁸[PW08] additionally require that, in the injective mode, there is a trapdoor that allows inverting the function. We will not need any such trapdoor

Definition 3.1. An *extremely lossy function* (ELF) consists of an algorithm ELF.Gen . ELF.Gen which takes as input integers M and $r \in [M]$. There is no security parameter here; instead, $\log M$ acts as the security parameter. ELF.Gen outputs the description of a function $f : [M] \rightarrow [N]$ such that:

- f is computable in time polynomial in the bit-length of its input, namely $\log M$. The running time is independent of r .
- If $r = M$, then f is injective with overwhelming probability (in $\log M$).
- For all $r \in [M]$, $|f([M])| \leq r$ with overwhelming probability. That is, the function f has image size at most r .
- For any polynomial p and inverse polynomial function δ (in $\log M$), there is a polynomial q such that: for any adversary \mathcal{A} running in time at most p , and any $r \in [q(\log M), M]$, we have that \mathcal{A} distinguishes $\text{ELF.Gen}(M, M)$ from $\text{ELF.Gen}(M, r)$ with advantage less than δ . Intuitively, no polynomial-time adversary \mathcal{A} can distinguish an injective from polynomial image size (where the polynomial size depends on the adversary's running time.).

For some applications, we will need an additional requirement for ELFs:

Definition 3.2. An ELF has an *efficiently enumerable* image space if, there is a (potentially randomized) procedure running in time polynomial in r and $\log M$ that, given $f \leftarrow \text{ELF.Gen}(M, r)$ and r , outputs a polynomial-sized set S of points in $[N]$ such that, with overwhelming probability over the choice of f and the randomness of the procedure, S contains the entire image set $f([M])$. If $S = f([M])$ with overwhelming probability, we say that the ELF has a strongly efficiently enumerable image space, and otherwise, we call the image space weakly efficiently enumerable.

Definition 3.3. An ELF has a weakly *efficiently sampleable* image space if there is a polynomial s and a (potentially randomized) polynomial time procedure running (where “polynomial” means polynomial in r and $\log M$) such that the following holds. Given $f \leftarrow \text{ELF.Gen}(M, r)$ and r , the procedure outputs a point $y \in [N]$ such that with overwhelming probability over the choice of f , the point y has a distribution that places weight at least $1/s$ on each image point in $f([M])$. We say that the ELF is properly efficiently sampleable if additionally $y \in f([M])$ with overwhelming probability (over the choice of f and the randomness of the algorithm). We say that the ELF is strongly efficiently sampleable if, with overwhelming probability over the choice of f , the distribution on y is statistically close to uniform on $f([M])$.

Lemma 3.4. *An ELF's image space is (strongly/weakly) efficiently sampleable if and only if it is (strongly/weakly) efficiently enumerable. Moreover, proper sampleability is equivalent to strong sampleability.*

Proof. In one direction, we just sample a random element from the polynomial-sized list S - if $S = f([M])$, this samples an almost uniform random image point, and if $S \supseteq f([M])$ then this will hit each image point with probability at least $1/|S|$. In the other direction, by sampling $\lambda s \log s$ points independently at random, except with negligible probability in λ , the set of sampled points will contain every one of the r image points. If the points are guaranteed to be in $f([M])$ (regardless of the distribution), then the list S will actually be equal to $f([M])$. This shows that proper or strong sampling implies strong enumeration. \square

The following property will be useful for attaining ELF's with efficiently enumerable/sampleable image spaces:

Definition 3.5. An ELF is strongly regular if, for all polynomial r , with overwhelming probability over the choice of $f \leftarrow \text{ELF.Gen}(M, r)$, the distribution $f(x)$ for a uniform $x \leftarrow [M]$ is statistically close to uniform over $f([M])$. The ELF is weakly regular if there is some polynomial s in $\log M$ and r such that the distribution $f(x)$ hits each image point with probability at least $1/s$.

Lemma 3.6. *If an ELF is weakly regular, then it is strongly efficiently sampleable/enumerable*

Proof. To sample, just apply the ELF to a random domain point. □

The final ELF property we define is *public coin*.

Definition 3.7. An ELF is *public coin* if the description of an injective mode f outputted by $\text{ELF.Gen}(M, M)$ is simply the random coins used by $\text{ELF.Gen}(M, M)$. The descriptions of lossy mode f 's outputted by $\text{ELF.Gen}(M, r), r < M$ may (and in fact, must) be a more complicated function of the random coins.

3.1 Constructing ELF's

We now show how to construct ELF's. Our construction will have two steps: first, we will show that ELF's can be constructed from a weaker primitive called a *bounded adversary* ELF, which is basically an ELF that is only secure against a priori bounded adversaries. Then we show essentially that the DDH-based lossy function of [FGK⁺10], when the group size is taken to be polynomial, satisfies our notion of a bounded-adversary ELF.

3.1.1 Bounded Adversary ELF's

Definition 3.8. A *bounded adversary* extremely lossy function (ELF) consists of an algorithm $\text{ELF.Gen}'$. $\text{ELF.Gen}'$ which takes as input integers $M, r \in [M]$, and $b \in \{0, 1\}$. Here, b will indicate whether the function should be lossy, and r will specify the lossiness. Similar to regular ELF's, there is no security parameter here; instead, $\log M$ acts as the security parameter. $\text{ELF.Gen}'$ outputs the description of a function $f : [M] \rightarrow [N]$ such that:

- f is computable in time polynomial in the bit-length of its input, namely $\log M$. The running time is independent of r .
- If $b = 0$, then f is injective with overwhelming probability (in $\log M$).
- For all $r \in [M]$, if $b = 1$, then $|f([M])| \leq r$ with overwhelming probability. That is, the function f has image size at most r .
- For any polynomial p and inverse polynomial function δ (in $\log M$), there is a polynomial q such that: for any adversary \mathcal{A} running in time at most p , and any $r \in [q(\log M), M]$, we have that \mathcal{A} distinguishes $\text{ELF.Gen}'(M, r, 0)$ from $\text{ELF.Gen}'(M, r, 1)$ with advantage less than δ .

Intuitively, the difference between a regular ELF and a bounded adversary ELF is that in a regular ELF, r can be chosen dynamically based on the adversary, whereas in a bounded adversary ELF, r must be chosen first, and then security only applies to adversaries whose running time is

sufficiently small. In a bounded adversary ELF, security holds even if the adversary sees r ; we will therefore assume without loss of generality that r is included in the output of $\text{ELF.Gen}'_k(M, r, b)$.

We now show that bounded adversary ELFs are sufficient for constructing full ELFs. Our construction of ELF.Gen from $\text{ELF.Gen}'$ is the following. On input M, r , ELF.Gen does:

- Let k be the smallest integer such that $M \leq 2^k$. Let $M' = 2^{3k}$, and $[N]$ be the co-domain of $\text{ELF.Gen}'$ on domain $[M']$.
- Let i^* be the integer such that $2^{i^*} \in (r/2, r]$. Set $b_{i^*} = 1$ and $b_i = 0$ for $i \neq i^*$.
- For $i = 1, \dots, k-1$, let $f_i \leftarrow \text{ELF.Gen}'(M', 2^i, b_i)$.
- For $i = 2, \dots, k$, choose a pairwise independent random function $h_i : [N'] \rightarrow [M']$.
- Choose a pairwise independent random $h_1 : [M] \rightarrow [M']$.
- Output the function $f = h_k \circ f_{k-1} \circ h_{k-1} \circ f_{k-2} \circ \dots \circ f_1 \circ h_1$.




Figure 1: An example instantiation for $k = 3$.

Theorem 3.9. *If $\text{ELF.Gen}'$ is a bounded-adversary ELF, then ELF.Gen is a (standard) ELF. If $\text{ELF.Gen}'$ is public coin, then so is ELF.Gen . If $\text{ELF.Gen}'$ is weakly enumerable, then so is ELF.Gen . If $\text{ELF.Gen}'$ is weakly/strongly regular, then so is ELF.Gen .*

Proof. First, if $r = M$, then $i^* = k$, and so each of the b_i will be 0. Thus each of the f_i will be injective with overwhelming probability. Fix $h_1, f_1, \dots, h_{i-1}, f_{i-1}$, and let S_i be the image of $f_{i-1} \circ h_{i-1} \circ f_{i-2} \circ \dots \circ f_1 \circ h_1$. Since each of the functions h_i have co-domain of size $M' \geq M^3$, by pairwise independence and Lemma 2.2, h_i will be injective on S_i with overwhelming probability. Thus, with overwhelming probability, the entire evaluation of f will be injective.

Second, notice that for any r , we have that $2^{i^*} \leq r$. Moreover, if $r < M$, the function f_{i^*} is set to be lossy with image size $2^{i^*} \leq r$. Thus, f will have image size at most r .

Third, we need to argue security. Let p be a polynomial and σ be an inverse polynomial (in $\log M$). Let p' be p plus the running time of ELF.Gen . We can think of p', σ as being functions of $\log M' = 3 \log M$. Let q be the polynomial guaranteed by $\text{ELF.Gen}'$ for p' and σ . Then we can consider q to be a polynomial in $\log M$. Consider any adversary A for ELF.Gen running in time at most p . Let $r = r(M)$ be a computable function of M such that $r \in (q(\log M), M]$. Our goal is to show that A distinguishes $\text{ELF.Gen}(M, M)$ from $\text{ELF.Gen}(M, r)$ with advantage less than δ .

Toward that goal, let i^* be such that $2^{i^*} \in (r/2, r]$. We construct an adversary A' for $\text{ELF.Gen}'$ with $r = 2^{i^*}$: let f_{i^*} be the f that A' receives, where f_{i^*} is either $\text{ELF.Gen}(M, 2^{i^*}, 0)$ or $\text{ELF.Gen}(M, 2^{i^*}, 1)$. As part of its input, A' obtains 2^{i^*} , from which it learns i^* . Then A' simulates the rest of f for itself, setting $b_i = 0, f_i \leftarrow \text{ELF.Gen}'(M, 2^i, b_i)$ for $i \neq i^*$ as well as generating the h_i . A' then runs A on

the simulated f . Notice that A' runs in time at most p' . Thus by the bounded-adversary security of $\text{ELF.Gen}'$, A' cannot distinguish injective or lossy mode, except with advantage σ . Moreover, if f_{i^*} is generated as $\text{ELF.Gen}(M, 2^{i^*}, 0)$, then this corresponds to the injective mode of ELF.Gen , and if f_{i^*} is generated as $\text{ELF.Gen}(M, 2^{i^*}, 1)$, then this corresponds to $\text{ELF.Gen}(M, r)$. Thus, A' and A have the same distinguishing advantage, and therefore A cannot distinguish the two cases except with probability less than σ .

It remains to show that ELF.Gen inherits some of the properties of $\text{ELF.Gen}'$. If $\text{ELF.Gen}'$ is public coin, then ELF.Gen is trivially public coin. To get a weak sampler for ELF.Gen , apply the weak sampler for $\text{ELF.Gen}'$ to the instance f_{i^*} that is lossy, obtaining point $y^{(i^*)}$. Then compute $y = h_k \circ f_{k-1} \circ h_{k-1} \circ f_{k-2} \circ \dots \circ f_{i^*+1} \circ h_{i^*+1}(y^{(i^*)})$. Since any image of f is necessarily computed as $h_k \circ f_{k-1} \circ h_{k-1} \circ f_{k-2} \circ \dots \circ f_{i^*+1} \circ h_{i^*+1}(y^{(i^*)})$ for *some* i^* in the image of f_{i^*} , each of the f_i for $i \neq i^*$ are injective (whp), each of the h_i are injective (whp) over the image of f_{i-1} , and since $y^{(i^*)}$ will hit each point in the image space of f_{i^*} frequently, the result $x^{(k+1)}$ will hit each image point of f frequently as well.

Finally, we show that regularity is inherited. First, fix all of the components of f except h_{i^*} . By similar analysis to the above, the map $f_{i^*-1} \circ h_{i^*-1} \circ \dots \circ f_1 \circ h_1$ will be injective, and so will $h_k \circ f_{k-1} \circ h_{k-1} \circ \dots \circ f_{i^*+1} \circ h_{i^*+1}$. This means the function $f_{i^*} \circ h_{i^*} \circ \dots \circ f_1 h_1$ is pairwise independent (by the pairwise independence of h_{i^*}), and recall that its output size is at most r . Let $y^{(i^*)}$ be the output of this truncated function. Using Lemma 2.3, we have that $h_{i^*}, y^{(i^*)}$ is negligibly-close (that is, $\frac{1}{2}\sqrt{(r-1)/M}$ -close) to $h_{i^*}, f_{i^*}(z^{(i^*)})$ for a random $z^{(i^*)}$. By the weak/strong regularity of f_{i^*} , we have that $y^{(i^*)}$ therefore hits every image point of f_{i^*} with inverse polynomial probability/probability negligibly close to $1/r$. This is preserved when moving to the full function f since all of the steps after f_{i^*} are injective (whp). \square

3.1.2 Instantiation for Bounded Adversary ELFs

Our construction of bounded adversary ELFs is based on the DDH-based lossy *trapdoor* functions of Peikert and Waters [PW08] and Freeman et al. [FGK⁺10]. We stress that we do not need the trapdoor property of their construction, only the lossy property. Security will be based on the exponential hardness of the decisional Diffie-Hellman problem, or its k -linear generalizations.

Definition 3.10. A *cryptographic group* consists of an algorithm GroupGen that takes as input a security parameter λ , and outputs the description of a cyclic group \mathbb{G} of prime order $p \in [2^\lambda, 2 \times 2^\lambda)$, and a generator g for \mathbb{G} such that:

- The group operation $\times : \mathbb{G}^2 \rightarrow \mathbb{G}$ can be computed in time polynomial in λ .
- Exponentiation by elements in \mathbb{Z}_p can be carried out in time polynomial in λ . This follows from the efficient group operation procedure by repeated doubling and the fact that $\log p \leq \lambda + 1$.
- The representation of a group element h has size polynomial in λ . This also follows implicitly from the assumption that the group operation is efficient.

We now introduce some notation. For a matrix $\mathbf{A} \in \mathbb{Z}_p^{m \times n}$, we write $g^{\mathbf{A}} \in \mathbb{G}^{m \times n}$ to be the $m \times n$ matrix of group elements $g^{A_{i,j}}$. Similarly define $g^{\mathbf{w}}$ for a vector $\mathbf{w} \in \mathbb{Z}_p^n$. Given a matrix $\hat{\mathbf{A}} \in \mathbb{G}^{m \times n}$ of group elements and a vector $\mathbf{v} \in \mathbb{Z}_p^n$, define $\hat{\mathbf{A}} \cdot \mathbf{v}$ to be $\hat{\mathbf{w}} \in \mathbb{G}^m$ where $\hat{w}_i = \prod_{j=1}^n \hat{A}_{i,j}^{v_j}$. Using this notation, $(g^{\mathbf{A}}) \cdot \mathbf{v} = g^{\mathbf{A} \cdot \mathbf{v}}$. Therefore, the map $g^{\mathbf{A}}, \mathbf{v} \mapsto g^{\mathbf{A} \cdot \mathbf{v}}$ is efficiently computable.

Definition 3.11. The *exponential decisional Diffie Hellman* (eDDH) assumption on a cryptographic group specified by `GroupGen` holds if there is a polynomial $q(\cdot, \cdot)$ such that the following is true. For any time bound t and probability ϵ , let $\lambda = \log q(t, 1/\epsilon)$. Then for any adversary \mathcal{A} running in time at most t , the following two distributions are indistinguishable, except with advantage at most ϵ :

$$\begin{aligned} (\mathbb{G}, g, g^a, g^b, g^c) &: (\mathbb{G}, g, p) \leftarrow \text{GroupGen}(\lambda), a, b, c \leftarrow \mathbb{Z}_p \text{ and} \\ (\mathbb{G}, g, g^a, g^b, g^{ab}) &: (\mathbb{G}, g, p) \leftarrow \text{GroupGen}(\lambda), a, b \leftarrow \mathbb{Z}_p \end{aligned}$$

More generally, we consider the exponential k -linear assumption:

Definition 3.12. The *exponential decisional k -linear assumption* (k -eLin) on a cryptographic group specified by `GroupGen` holds if there is a polynomial $q(\cdot)$ such that the following is true. For any time bound t and probability ϵ , let $\lambda = \log q(t, 1/\epsilon)$. Then for any adversary \mathcal{A} running in time at most t , the following two distributions are indistinguishable, except with advantage at most ϵ :

$$\begin{aligned} (\mathbb{G}, g, g^{a_1}, \dots, g^{a_k}, g^c, g^{a_1 b_1}, \dots, g^{a_k b_k}) &: (\mathbb{G}, g, p) \leftarrow \text{GroupGen}(\lambda), a_i, b_i, c \leftarrow \mathbb{Z}_p \text{ and} \\ (\mathbb{G}, g, g^{a_1}, \dots, g^{a_k}, g^{\sum_{i=1}^k b_i}, g^{a_1 b_1}, \dots, g^{a_k b_k}) &: (\mathbb{G}, g, p) \leftarrow \text{GroupGen}(\lambda), a_i, b_i \leftarrow \mathbb{Z}_p \end{aligned}$$

$k = 1$ corresponds to the eDDH assumption above.

The following will help us achieve a public coin ELF.

Definition 3.13. A cryptographic group is *public coin* if the following holds:

- The “description” of \mathbb{G}, g, p is just the random coins sampled by `GroupGen`.
- There is a (potentially redundant) efficiently computable representation of group elements in \mathbb{G} as strings in $\{0, 1\}^n$ such that (1) a random string in $\{0, 1\}^n$ corresponds to a random element in \mathbb{G} , and (2) a random representation of a random element in \mathbb{G} is a random string in $\{0, 1\}^n$.

A plausible candidate for a cryptographic group supporting the eDDH assumption are groups based on elliptic curves. Despite over a decade or research, the best attacks on many elliptic curves are generic attacks which require exponential time. Therefore, the eDDH assumption on these groups appears to be a very reasonable assumption. We note that groups based on elliptic curves can be made public coin.

Construction. Our construction is as follows, and will be parameterized by k . `ELF.Gen'_k(M, r, b)` does the following.

- Let λ be the largest integer such that $(2 \times 2^\lambda)^k < r$. Run $(\mathbb{G}, g, p) \leftarrow \text{GroupGen}(\lambda)$.
- Let m be the smallest integer such that $p^m \geq M$. Let R be a pairwise independent function from $[M]$ into \mathbb{Z}_p^m that is injective with overwhelming probability.
- Let $n \geq m$ be chosen such that a random matrix sampled from $\mathbb{Z}_p^{n \times m}$ has rank m with overwhelming probability. Note that a random *square* matrix will be full singular with probability $1/p$, and in our case, p is polynomial. Therefore, we need to choose an n somewhat larger than m . It suffices to set $n = 2m$.

- If $b = 0$, choose a random matrix of group elements $g^{\mathbf{A}}$. If $b = 1$, choose a random rank- k matrix $\mathbf{A} \in \mathbb{Z}_p^{n \times m}$ and compute $g^{\mathbf{A}}$.
- Output the function $f(x) = \mathbf{A} \cdot (R(x))$. The description of f will consist of $(\mathbb{G}, p, \mathbf{A}, R, m, n)$.

Theorem 3.14. *If GroupGen is a group where the k -eLin assumption holds for some constant k , then $\text{ELF.Gen}'_k$ is a strongly regular bounded adversary ELF. If GroupGen is public coin, then so is $\text{ELF.Gen}'$.*

Proof. If \mathbf{A} is full rank, then the map $\mathbf{y} \mapsto g^{\mathbf{A} \cdot \mathbf{y}}$ is injective. Therefore, injectivity of f follows from the fact that R is almost always injective and \mathbf{A} is almost always full rank.

For security, note that we just to show that the two distributions on $g^{\mathbf{A}}$ are indistinguishable. Note that the k -linear assumption implies that it is hard to distinguish $g^{\mathbf{B}}$ for a random full rank $\mathbf{B} \in \mathbb{Z}_p^{k+1, k+1}$ from $g^{\mathbf{B}}$ for a random rank k matrix \mathbf{B} with no loss in the security of the reduction [NS09]. From here, it is straightforward to show that it is hard to distinguish the full rank and rank k cases of $g^{\mathbf{A}}$, with a loss of a factor of $m - k$. Villar [Vil12] shows even more that the loss can even be made logarithmic in m :

Lemma 3.15 ([Vil12]). *Suppose \mathcal{A} is an adversary that, on input $g^{\mathbf{A}}$, distinguishes the cases where $\mathbf{A} \in \mathbb{Z}_p^{n \times m}$ is uniformly random from the case where \mathbf{A} has rank k with advantage δ . Then \mathcal{A} can be used to construct an adversary \mathcal{B} that distinguishes the two cases in the k -liner assumption with advantage at most $\delta / \log(m)$. If \mathcal{A} runs in time t , \mathcal{B} runs in time $t + u$ for some polynomial u in the bit-length of $g^{\mathbf{A}}$.*

Let q be the polynomial guaranteed by the k -eLin assumption. Let t be a polynomial and δ an inverse polynomial. Let $q' = 4q(t + u, m/\delta)^k$, where u is the computational overhead in the reduction in Lemma 3.15. Suppose an adversary runs in time t and distinguishes the two distributions on $g^{\mathbf{A}}$ with advantage δ . For any $r \geq q'$, we have that $\lambda \geq r^{1/k}/4 \geq q(t + u, \log(m)/\delta)$. This means no $(t + u)$ -time adversary can break the k -eLin assumption with advantage greater than $\delta / \log(m)$. By the reduction of Lemma 3.15, this means no t -time adversary can distinguish the two cases of $g^{\mathbf{A}}$, except with advantage at most δ , as desired.

Notice that if GroupGen is public coin, we can sample $g^{\mathbf{A}}$ directly in the injective mode since it is just a matrix of random group elements.

Finally, note that the function $\mathbf{y} \mapsto g^{\mathbf{A} \cdot \mathbf{y}}$ is *perfectly* regular. The strong regularity of $x \mapsto g^{\mathbf{A} \cdot R(x)}$ follows from the pairwise independence of R and a simple application of Lemma 2.3. \square

Corollary 3.16. *If there exists a constant k and a cryptographic group where the k -eLin assumption holds, then there exists an ELF with a strongly efficiently sampleable/enumerable image. Moreover, if the group is public coin, then so is the ELF.*

4 Point Function Obfuscation

A (expanding) random oracle H serves as a good point function obfuscator: to obfuscate the point

function $I_x(x') = \begin{cases} 1 & \text{if } x' = x \\ 0 & \text{if } x' \neq x \end{cases}$, simply output $y = H(x)$. Then to run the “program” on input

x' , simply check that $H(x') = y$. For any x that is drawn from an source with super-logarithmic min-entropy, an adversary making a polynomial number of queries to H will not be able to determine x from y . Thus, x is hidden to all efficient adversaries.

In this section, we show how to use ELF's to implement a concrete function H for which the strategy above still yields a secure point obfuscation (PO).

Definition 4.1. A point obfuscator (PO) is an efficient probabilistic algorithm \mathcal{O} with the following properties:

- (Almost Perfect Correctness) On input a point function I_x , with overwhelming probability over the random coins of \mathcal{O} , \mathcal{O} outputs the description of a program P that is functionally equivalent to I_x . P must run in time polynomial in the length of x and the security parameter.
- (Secrecy) For any distribution \mathcal{D} over a set \mathcal{X} with super-logarithmic min-entropy, the distribution $\mathcal{O}(I_x)$ for $x \leftarrow \mathcal{D}$ is computationally indistinguishable from $\mathcal{O}(I_{x'})$ where $x' \leftarrow U_{\mathcal{X}}$.

Before giving our construction, we show that a point obfuscator implies a separation from NP with super-logarithmic non-determinism and P. Thus, any primitive used to build point obfuscation, such as ELF's, must necessarily imply such a separation. This is essentially the same statement as a theorem of Wee [Wee05].

Theorem 4.2. *If Point Obfuscators exist, then for any super-logarithmic function t , NP with t bits of non-determinism is not solvable in polynomial time.*

Proof. Fix a super-logarithmic function t , and let $T = 2^t$. Suppose that NP with t bits of non-determinism is solvable in polynomial time. Then the following problem is solvable in polynomial time: given a boolean circuit C on $t = t(|C|)$ input bits, determine if C has a satisfying input.

Now we use such a solution to break a point obfuscator. Given a point obfuscator for points $x \in \mathcal{X} = [2^m]$, let s be the size of the obfuscated circuit, which is polynomial in m . Let $t = t(s) = \omega(\log s) = \omega(\log m)$ as above. Let $T = 2^t$, and let \mathcal{D} be the uniform distribution on $[T]$. Then \mathcal{D} has super-logarithmic min-entropy. Our distinguisher works as follows. Given the description of a program P outputted by the obfuscator, let P^T be the program restricted to inputs in $[T]$ (equivalently, all but the last t bits are hard-coded to 0). Then run the circuit satisfiability algorithm above, and output the result.

Now notice that if $x \leftarrow \mathcal{D}$, then by (almost) perfect correctness, P^T has an accepting input, namely x . Moreover, if $x \leftarrow U_{\mathcal{X}}$, then with overwhelming probability $x \notin [T]$, and thus by (almost) perfect correctness, P^T will *not* have an accepting input. Therefore, by running the circuit satisfiability algorithm, we successfully distinguish the two cases. \square

4.1 The Construction

Construction 4.3. Let \mathcal{X} be the desired domain of H . To generate H , to the following:

- Let \mathcal{Z} be some set such that $|\mathcal{X}|/|\mathcal{Z}|$ is negligible. We will assume that $\mathcal{Z} = [T]$ for some integer T , or at least that \mathcal{Z} is in efficient bijection with $[T]$. This way, we can use \mathcal{Z} as the domain for an ELF.

Sample a hash function h from a uniform and pairwise independent function distribution from \mathcal{X} to \mathcal{Z} . We require that h is injective. There are two simple ways to ensure the injectivity of h . By setting \mathcal{Z} even larger so that $|\mathcal{X}|^2/|\mathcal{Z}|$ is negligible, h will be injective with overwhelming probability for any pairwise independent function distribution. Another option which does not require increasing the size of \mathcal{Z} is to let \mathcal{Z} be a large field and \mathcal{X} a subset of \mathcal{Z} . Then we can set $h(x) = ax + b$ for random a, b . With overwhelming probability, $a \neq 0$, in which case h is injective.

- Let $f \leftarrow \text{ELF.Gen}(|\mathcal{Z}|, |\mathcal{Z}|)$ to get an injective-mode f .
- Output $H = f \circ h$.




Figure 2: The function $H = f \circ h$.

Theorem 4.4. *Assuming ELF is a secure ELF, H in Construction 4.3 gives a secure point obfuscator. If ELF is public coin, then so is H .*

Proof. We will actually show something stronger: that the point function obfuscation of x is indistinguishable from an obfuscation of the all-zeros function. In particular, we will show that no efficient adversary can distinguish $y = f(h(x))$ from $y = f(z)$ for a uniformly random z . Notice that by injectivity of f , y has a pre-image under $H = f \circ h$ if and only if $z = f^{-1}(y)$ has a pre-image under h . Since we chose h to be expanding, when we sample z uniformly random, z will have no pre-image with overwhelming probability. Therefore, $y = f(z)$ has no pre-image with overwhelming probability.

The proof involves a sequence of hybrids. Suppose the adversary runs in time t and distinguishes $y = f(h(x))$ from $y = f(z)$ with non-negligible advantage ϵ . This means there is an inverse polynomial δ such that $\epsilon \geq \delta$ infinitely often.

Hybrid 0 This is the honestly generated $y = f(h(x))$ for f drawn in injective mode and x drawn from D .

Hybrid 1 Now, we change f to be lossy. That is, we generate $f \leftarrow \text{ELF.Gen}(|\mathcal{Z}|, r)$ where r is chosen so that no adversary running in time t can distinguish this lossy f from an injective f , except with advantage at most $\delta/3$. Thus by ELF security, the adversary cannot distinguish **Hybrid 0** from **Hybrid 1**, except with probability $\delta/3$.

Hybrid 2 Now we change y to be $y = f(z)$ for a random uniform $z \in \mathcal{Z}$. Fix f , and let E be the distribution of y . Then notice that by the pairwise independence and uniformity of h , the composition $H = f \circ h$ is pairwise independent and has output distribution E . Moreover, $\text{Supp}(E) \leq r$ is a polynomial. Therefore, by using our new leftover hash lemma (Lemma 2.3), we see that **Hybrid 1** and **Hybrid 2** are indistinguishable, except with probability $\frac{1}{2} \sqrt{CP(D)(|\text{Supp}(E)| - 1)}$. As long as the collision probability of \mathcal{X} is negligible (which in particular happens when \mathcal{X} has super-logarithmic min-entropy), this quantity will be negligible. In particular, the distinguishing advantage will be less than $\delta/3$.

Hybrid 3 Now we change f to be injective again. Similarly to before, the distinguishing advantage between **Hybrid 2** and **Hybrid 3** will be at most $\delta/3$. Notice that **Hybrid 3** is exactly our all-zeros obfuscation. Therefore, **Hybrid 0** and **Hybrid 3** are indistinguishable, except

with probability less than δ , meaning $\epsilon < \delta$. This contradicts our assumption about the adversary. \square

In Section 6, we will show how to strengthen our construction to get a point obfuscator that is secure even against auxiliary information about the point.

5 Output Intractability

Consider any $k + 1$ -ary relation R over $\mathcal{Y}^k \times \mathcal{W}$ that is *computationally intractable*: on a random input $\mathbf{y} \in \mathcal{Y}^k$, it is computationally infeasible to find a $w \in \mathcal{W}$ such that $R(\mathbf{y}, w)$ outputs 1. Let H be a hash function, and for a vector $\mathbf{x} = (x_1, \dots, x_k)$ of inputs, let $H(\mathbf{x}) = (H(x_1), \dots, H(x_k))$. If H is modeled as a random oracle, assuming k is a constant, it is computationally infeasible to find a set of distinct inputs \mathbf{x} , $x_i \neq x_j \forall i \neq j$, and a $w \in \mathcal{W}$, such that $R(H(\mathbf{x}), w) = 1$. We will now show how to build standard-model hash functions H that achieve the same property.

Definition 5.1. A family of hash functions $H : [M] \rightarrow \mathcal{Y}$ is *k -ary output intractable* if, for any computationally intractable $k + 1$ -ary relation $R : \mathcal{Y}^k \times \mathcal{W} \rightarrow \{0, 1\}$, no efficient adversary, given H , can find a set of distinct inputs $\mathbf{x} \in [M]^k$ and an element $w \in \mathcal{W}$, such that $R(H(\mathbf{x}), w) = 1$, except with negligible probability.

Note that binary output intractability implies as a special case collision resistance. In the unary case, and if \mathcal{W} is just a singleton set, then output intractability is a special case of *correlation intractability*, where the relation in addition depends on the *input*.

The unary case captures the following use case of hash functions: a given protocol may require a common reference string (crs), but some or all instances of the crs may admit a trapdoor that allows breaking the protocol. Of course, such a trapdoor should be difficult to find for a random crs. To “prove” that the crs is generated so that the generator of the crs does not know a trapdoor, the generator sets the crs to be the output of a public hash function on an arbitrary point. Since the potentially malicious generator does not control the hash function, he should be unable to find an output along with a corresponding trapdoor. Modeling the hash function as a random oracle, this methodology is sound. However, standard notions of security do not prevent the crs generator from choosing the input in such a way so that it knows a trapdoor. Unary output intractability precludes this case.

Moreover, consider, say, the case of Dual_EC_DRBG, where there are two crs components, namely two points on an elliptic curve. The discrete log between the two points can be used as a trapdoor that allows for breaking the generator [YY97]. One option to “prove” that nobody knows a trapdoor is the following: each point is set to be the output of a hash function on some input, and the pair of inputs is included in the description of the protocol as the proof. By using a *binary* output intractable hash function, finding a pair of inputs such that one knows the discrete log between the outputs is computationally intractable. More generally, if there are multiple crs components, k -ary output intractability for higher k can be used to prove that nobody knows a trapdoor between the outputs.

Of course, the hash function itself needs to be set up in a trusted manner; however, once the hash function is set up and trusted, it can be used to generate arbitrarily many different crs by even untrusted authorities.

We note, however, that the unary case on its own is not very interesting: the family of hash functions H parameterizes by a string $y \in \mathcal{Y}$ where $H(x) = y$ for all x is clearly unary intractable.

Depending on the application, one may want additional features such as collision resistance, which as noted above is implied by binary output intractability ($k = 2$). Moreover, collision-resistance ensures that different protocols will use different crs's, provided that different inputs are used. Therefore, $k = 2$ and above are likely to be the most interesting settings.

Trivial impossibility for arbitrary k . We note that no one family of hash functions H can satisfy k -ary output intractability for all k . That is, for different k , a different family will be required. Suppose to the contrary that a family H satisfied k -output intractability for all k . Let t be the size of the circuit computing H . Choose k so that $k \log |\mathcal{Y}| \geq t$. Then with overwhelming probability over the choice of random $\mathbf{y} \in \mathcal{Y}^k$, there is no circuit of size at most t that outputs y_i on input $i \in [k]$. Therefore, let \mathcal{W} be the set of circuits of size at most t , and let $R(\mathbf{y}, C)$ output 1 if and only if $C(i) = y_i$ for each $i \in [k]$. Then R is computationally (in fact statistically) intractable. However, it is trivial to find an \mathbf{x}, w that satisfy $R(H(\mathbf{x}), w) = 1$: set $\mathbf{x} = [k]$ and $w = H$. Therefore, output intractability is violated. We obtain the following:

Theorem 5.2. *For any family $H : [M] \rightarrow \mathcal{Y}$ of hash functions, let t be the size of circuits for H . Then H cannot be output intractable for any $k \geq t / \log |\mathcal{Y}|$.*

In the following, we show that it is nonetheless possible to obtain output intractability for any given constant k . Our functions will be described by strings of length $k(\log |\mathcal{Y}| + \text{poly}(\log M))$, which in the case $|\mathcal{Y}| \gg M$ gives a near-optimal relationship between k and t .

First, however, we show that, *even if the set \mathcal{W} is just a singleton set*, then output intractability implies that NP with super-logarithmic non-determinism is separated from P.

Theorem 5.3. *If binary output intractable hash functions exist, then for any super-logarithmic function t , NP with t bits of non-determinism is not solvable in polynomial time.*

Proof. The proof is analogous to the proof of Theorem 4.2. Suppose to the contrary that NP with t bits of non-determinism is solvable in polynomial time. This means that it is possible to solve circuit-SAT for circuits with t input bits in polynomial time. From such a polynomial-time algorithm, it is straightforward to construct a polynomial-time algorithm that actually finds a satisfying input.

Let $H : \{0, 1\}^{t/2} \rightarrow \{0, 1\}^{t/4}$ be unary output intractable. Let $R : (\{0, 1\}^{t/4})^2 \rightarrow \{0, 1\}$ be the relation such that $R(y_1, y_2) = 1$ if and only if $y_1 = y_2$ (here, the set \mathcal{W} is trivial). Then R is intractable. Let $C_H(x_1, x_2)$ be the circuit that checks that $x_1 \neq x_2$ and that $R(H(x_1), H(x_2)) = 1$. Clearly, since H is shrinking, satisfying inputs to C_H must exist. Moreover, C takes as input t -bits. Therefore, we can use the search algorithm above to efficiently find a satisfying input. This violates the output intractability of H . \square

5.1 The Construction

Construction 5.4. Let $[M]$ be the desired domain of H , and \mathcal{Y} the desired range. To generate H , to the following:

- Let $f \leftarrow \text{ELF.Gen}(M, M)$ to get an injective-mode f , with codomain \mathcal{Z} .
- Let g be a k -wise independent and uniform function from \mathcal{Z} to \mathcal{Y} .
- Output $H = g \circ f$.




Figure 3: The function $H = g \circ f$.

Theorem 5.5. *If ELF is a secure ELF with an efficiently enumerable image, then for any constant k the hash function H in Construction 5.4 is k -ary output intractable. If ELF is public coin, then so is H .*

Proof. Suppose toward contradiction that there is an intractable $k + 1$ -ary relation R and an adversary \mathcal{A} that on input H finds a set of distinct inputs \mathbf{x} and a value $w \in \mathcal{W}$ such that $R(H(\mathbf{x}), w) = 1$ with non-negligible probability ϵ . Let δ be an inverse polynomial such that $\epsilon \geq \delta$ infinitely often. We will switch to a lossy mode for f so that (1) f has polynomial image size, and (2) no adversary running in time t (for a t to be chosen later) can distinguish the lossy mode from injective, except with probability $\delta/3$. By choosing t to be larger than the running time of \mathcal{A} , we have that \mathcal{A} still outputs \mathbf{x} of distinct elements, and a string w , such that $R(H(\mathbf{x}), w) = 1$ with probability $\epsilon - \delta/3$.

We first argue that each of the elements of $f(\mathbf{x})$ are distinct except with probability $\delta/3$. Since this was true in the injective case (since \mathbf{x} is distinct), if this is not true in the lossy case, then the injective and lossy modes could be easily distinguished by an adversary taking slightly more time than \mathcal{A} . Let t be this time, so that this distinguisher is impossible. Thus, the adversary succeeds *and* the elements of $f(\mathbf{x})$ are distinct with probability at least $\epsilon - 2\delta/3$. This probability is larger than $\delta/3$ infinitely often, and is therefore non-negligible. Let S be the polynomial-sized set of image points of f . Then in other words, the adversary comes up with an ordered set \mathbf{z} of distinct elements in S , and a string w , such that $R(g(\mathbf{z}), w) = 1$.

Now, note that, for any ordered set \mathbf{z} of k distinct inputs, $g(\mathbf{z})$ is distributed uniformly at random, by the k -wise independence of g . Moreover, it is straightforward, given \mathbf{z} and a vector $\mathbf{y} \in \mathcal{Y}^k$, to sample a random g conditioned on $g(\mathbf{z}) = \mathbf{y}$. Sampling random \mathbf{y} , and then g in this way, gives a correctly distributed g .

We now describe an algorithm \mathcal{B} that breaks the intractability of R . \mathcal{B} , on input $\mathbf{y} \in \mathcal{Y}^k$, chooses lossy f as above, and then selects a random ordered set \mathbf{z} , $|\mathbf{z}| = k$ among the p outputs of f . Next, it chooses a random g such that $g(\mathbf{z}) = \mathbf{y}$. Finally, it runs \mathcal{A} on the hash function $H = g \circ f$. When \mathcal{A} outputs \mathbf{x}, w , if $f(\mathbf{x}) = \mathbf{z}$ (equivalently, $H(\mathbf{x}) = \mathbf{y}$), \mathcal{B} outputs w ; otherwise it aborts.

Since \mathbf{y} is hidden from \mathcal{A} 's view, g is distributed randomly according to the k -wise independent distribution. Therefore, \mathcal{A} will output a valid w with probability at least $\epsilon - 2\delta/3$. If \mathcal{B} 's guess for \mathbf{z} was correct, then w will break the correlation intractability of R on \mathbf{y} . Since \mathbf{z} is random and independent of \mathcal{A} 's view, the probability of a good guess is at least $1/p^k$ (since there are at most p^k such \mathbf{z}). Therefore, \mathcal{B} breaks the intractability of R with probability $(\epsilon - 2\delta/3)/p^k$, which is larger than $\delta/3p^k$ infinitely often, and is therefore non-negligible. \square

6 Leakage-resilient PRGs, AIPO and Poly-many Hardcore Bits

In this section, we use ELFs to give arbitrarily-many hardcore bits for any one-way function, or for constructing point function obfuscation secure in the presence of auxiliary information. Both of these can be seen as special cases of a very strong security requirement for pseudorandom generators.

Definition 6.1. A distribution \mathcal{D} on pairs $(x, z) \in \mathcal{X} \times \mathcal{Z}$ is *computationally unpredictable* if no efficient adversary can guess x given z .

Definition 6.2. A family of pseudorandom generators $H : \mathcal{X} \rightarrow \mathcal{Y}$ is secure for computationally unpredictable seeds if, for any computationally unpredictable distribution on $(\mathcal{X}, \mathcal{Z})$, no efficient adversary can distinguish $(H, z, H(x))$ from (H, z, S) where $(x, z) \leftarrow \mathcal{D}$ and $S \leftarrow U_{\mathcal{Y}}$.

Basically, this requirement states that H is a secure pseudorandom generator for arbitrary distributions on the seed, and even remains secure in the presence of arbitrary leakage about the seed, so long as the seed remains *computationally* unpredictable. The only restriction is that the distribution on the seed and the leakage must be chosen independently of H . However, in the absence of other restrictions, this independence between the source \mathcal{D} and function H can easily be seen to be necessary: if z contained a few bits of $H(x)$, then it is trivial to distinguish $H(x)$ from random.

6.1 The Construction

The intuition behind our construction is the following. The usual way of extracting pseudorandomness from computationally unpredictable source is to output a hardcore bit of the source, say using Goldreich-Levin [GL89]. While this can be used to generate a logarithmic number of pseudorandom bits, security is lost once a super-logarithmic number of hardcore bits have been generated in this way.

In order to get around this logarithmic barrier, we actually compute a *polynomial* number of Goldreich-Levin bits. Of course, we cannot output these in the clear or else the seed can be easily computed by linear algebra. Instead, we scramble the hardcore bits using a sequence of ELFs. We can argue that each of the (scrambled) hardcore bits really is “as good as” random, in the sense that we can replace each bit with a truly random bit before scrambling without detection. To do so, we use the lossiness of the ELFs to argue that, when the i th hardcore bit is incorporated into the scramble, enough information is lost about the previous bits that the i th bit actually still is hardcore. By iterating this for each bit, we replace each one with random. We now give the details.

Construction 6.3. Let q be the input length and m be the output length. Let λ be a security parameter. We will consider inputs x as q -dimensional vectors $\mathbf{x} \in \mathbb{F}_2^q$. Let **ELF** be an ELF. Let $M = 2^{m+\lambda+1}$, and let n be the bit-length of the ELF on input $m+1$. Set $N = 2^n$. Let ℓ be some polynomial in m, λ to be determined later. First, we will construct a function H' as follows.

Choose random $f_1, \dots, f_\ell \leftarrow \mathbf{ELF.Gen}(M, M)$ where $f_i : [M] \rightarrow [N]$, and let $h_1, \dots, h_{\ell-1} : [N] \rightarrow [M/2] = [2^{m+\lambda}]$ and $h_\ell : [N] \rightarrow [2^m]$ be sampled from pairwise independent and uniform function families. Define $\mathbf{f} = \{f_1, \dots, f_\ell\}$ and $\mathbf{h} = \{h_1, \dots, h_\ell\}$. Define $H'_i : \{0, 1\}^i \rightarrow [M/2]$ (and $H'_\ell : \{0, 1\}^\ell \rightarrow [2^m]$) as follows:

- $H'_0() = 1 \in [2^{m+\lambda}]$
- $H'_i(\mathbf{b}_{[1, i-1]}, b_i) : \text{compute } y_i = H'_{i-1}(\mathbf{b}_{[1, i-1]}), z_i \leftarrow f_i(y_i || b_i), \text{ and output } y_{i+1} \leftarrow h_i(z_i)$

Then we set $H' = H'_\ell$. Then to define H , choose a random matrix $\mathbf{R} \in \mathbb{F}_2^{\ell \times q}$. The description of H consists of $\mathbf{f}, \mathbf{h}, \mathbf{R}$. Then set $H(x) = H'(\mathbf{R} \cdot \mathbf{x})$. A diagram of H is given in Figure 4.




Figure 4: An example instantiation for $\ell = 3$. Notice that each iteration is identical, except for the final iteration, where h_ℓ has a smaller output.

We now prove two important facts about H and H' :

Claim 6.4. *If $\ell \geq m + \lambda$, and if \mathbf{b} is drawn uniformly at random, then $(H', H'(\mathbf{b}))$ is statistically close to (H', R) where R is uniformly random in $[2^m]$.*

Proof. We will prove the case $\ell = m + \lambda$, the case of larger ℓ being similar though more tedious. We will consider f_1, \dots, f_ℓ as being fixed injective functions; since the f_i are injective with overwhelming probability, the claim follows from this case. This means that the composition $h_i \cdot f_i$ is pairwise independent, for all i .

Let $d_i(h_1, \dots, h_i)$ be the collision probability of y_{i+1} when $b_1 \dots, b_i$ are random bits, for fixed h_1, \dots, h_i . Let d_i be the expectation (over h_1, \dots, h_i) of this value. There are two possibilities for a collision at y_{i+1} :

- There is a collision at y_i and b_i . This happens with half the probability that there is a collision at y_i .
- There is not a collision at y_i and b_i , but $h_i \cdot f_i$ maps the two points to the same y_{i+1} . Conditioned on there being no collision at (y_i, b_i) , this occurs with probability $\frac{1}{2^{m+\lambda}}$ for $i < \ell$, and $\frac{1}{2^m}$ for $i = \ell$.

Then we have that $d_0 = 1$ and the following recurrence for $i < \ell$

$$\begin{aligned}
d_i &= \mathbb{E}_{h_1, \dots, h_i} d_i(h_1, \dots, h_i) = \mathbb{E}_{h_1, \dots, h_{i-1}} \sum_h \left(\mathbb{E}_{h_i} \Pr[y_i = y'_i : h_1, \dots, h_{i-1}, h_i] \right) \\
&= \mathbb{E}_{h_1, \dots, h_{i-1}} \left[\frac{\Pr[y_{i-1} = y'_{i-1} : h_1, \dots, h_{i-1}]}{2} + \left(1 - \frac{\Pr[y_{i-1} = y'_{i-1} : h_1, \dots, h_{i-1}]}{2} \right) \frac{1}{2^{m+\lambda}} \right] \\
&= \mathbb{E}_{h_1, \dots, h_{i-1}} \left[\Pr[y_{i-1} = y'_{i-1} : h_1, \dots, h_{i-1}] \left(\frac{1}{2} - \frac{1}{2 \times 2^{m+\lambda}} \right) + \frac{1}{2^{2+\lambda}} \right] \\
&= d_{i-1} \left(\frac{1}{2} - \frac{1}{2 \times 2^{m+\lambda}} \right) + \frac{1}{2^{m+\lambda}}
\end{aligned}$$

This recurrence has the form $d_i = a + bd_{i-1}$, which is solved by $d_i = b^i + a\frac{b^i-1}{b-1}$. Therefore, we have

$$d_i = \frac{2}{2^{m+\lambda} + 1} + \frac{(2^{m+\lambda} - 1)^{i+1}}{2^{(m+\lambda)i}(2^{m+\lambda} + 1)}$$

Now, for $i = m + \lambda - 1$, this becomes $d_{m+\lambda-1} = \frac{2(1+(1-2^{-m-\lambda})^{m+\lambda})}{2^{m+\lambda+1}} \leq \frac{4}{2^{m+\lambda}}$.

Next, we have the modified recurrence for $d_{m+\lambda}$: $d_{m+\lambda} = d_{m+\lambda-1} \left(\frac{1}{2} - \frac{1}{2 \times 2^m}\right) + \frac{1}{2^m}$. This gives us $d_{m+\lambda} \leq \frac{1}{2^m} + \frac{2}{2^{m+\lambda}}$. Now, the Rényi entropy of $y_{m+\lambda+1}$ is exactly the collision probability, scaled up by a factor of the 2^m . Therefore, the expected Rényi entropy of $y_{m+\lambda+1}$ is at most $1 + \frac{2}{2^\lambda}$: $\mathbb{E}_{h_1, \dots, h_\ell} RE(y_{\ell+1}) \leq 1 + \frac{2}{2^\lambda}$

Finally, we relate the Rényi entropy to the statistical distance from uniform using Lemma 2.1:

$$\mathbb{E}_{h_1, \dots, h_\ell} \Delta(y_{\ell+1}, R) \leq \mathbb{E}_{h_1, \dots, h_\ell} \frac{1}{2} \sqrt{RE(y_{\ell+1}) - 1} \leq \frac{1}{2} \sqrt{\mathbb{E}[RE(y_{\ell+1})] - 1} \leq 2^{-(\lambda+1)/2}$$

This is negligible in λ , completing the proof. \square

We will thus set $\ell = m + \lambda$ in our construction of H' . Claim 6.4 will be crucial for our proof that H meets our strong PRG security notion.

We also show our H is injective (whp) exactly when a truly random function with the same domain and co-domain is injective (whp).

Claim 6.5. *If $2^{-(m-2q)}$ is negligible (in q), and $\ell \geq m$, then with overwhelming probability H is injective.*

Proof. First, note that with overwhelming probability by our choice of $\ell \geq m \geq 2q$, \mathbf{R} is full rank. Next, let \mathcal{Y}_i be the set of possible y_i values as we vary \mathbf{x} , and \mathcal{Z}_i be the set of possible z_i values. By the injectivity of f_i , we have that $|\mathcal{Z}_i| \geq |\mathcal{Y}_i|$. Moreover, since h_i is pairwise independent and uniform, with overwhelming probability h_i is injective on \mathcal{Z}_i since $|\mathcal{Z}_i| \leq 2^q$ but the co-domain of h_i has size at least $2^m \gg (2^q)^2$. Therefore $|\mathcal{Y}_{i+1}| = |\mathcal{Z}_i| \geq |\mathcal{Y}_i|$. This means that as we increase i , the image size never decreases (with overwhelming probability)

Now pick q linearly independent rows of \mathbf{R} . We will assume that the q rows constitute the first q rows of \mathbf{R} ; the more general case is handled analogously. By performing an appropriate invertible transformation on the domain, we can assume that these q rows form the identity matrix. Therefore, we can take $b_i = x_i$ for $i \in [q]$. Next, observe that y_i for $i \in [q]$ only depends on the first $i - 1$ bits of \mathbf{x} . Thus the set of possible pairs $(y_i, b_i) = (y_i, x_i)$ is exactly $\mathcal{Y}_i \times \{0, 1\}$, which has size $2|\mathcal{Y}_i|$. By the injectivity of f_i , $|\mathcal{Z}_i| = 2|\mathcal{Y}_i|$. Since $|\mathcal{Y}_{i+1}| = |\mathcal{Z}_i| = 2|\mathcal{Y}_i|$, we have that the image size exactly doubles in each iteration for $i \in [q]$. Once we get to $i = q$, the image size is 2^q , and the remaining iterations do not introduce any collisions. Thus the image size of H is 2^q , meaning H is injective. \square

We now present our main theorem of the section:

Theorem 6.6. *If ELF is a secure ELF, then H in Construction 6.3 is a pseudorandom generator secure for computationally unpredictable seeds. If ELF is public coin, then so is H .*

Proof. Recall that $H(\mathbf{x}) = H'(\mathbf{R} \cdot \mathbf{x})$, and that $H'(\mathbf{b})$ is statistically close to random when \mathbf{b} is random. Therefore, it suffices to show that the following distributions are indistinguishable: $(\mathbf{f}, \mathbf{h}, \mathbf{R}, z, H'(\mathbf{R} \cdot \mathbf{x}))$ and $(\mathbf{f}, \mathbf{h}, \mathbf{R}, z, H'(\mathbf{b}))$ for a uniformly random \mathbf{b} .

Suppose an adversary \mathcal{A} has non-negligible advantage ϵ in distinguishing the two distributions. Define $\mathbf{b}^{(i)}$ so that the first i bits of $\mathbf{b}^{(i)}$ are equal to the first i bits of $\mathbf{R} \cdot \mathbf{x}$, and the remaining $\ell - i$ bits are chosen uniformly at random independently of \mathbf{x} . Define **Hybrid** i to be the case where \mathcal{A} is given the distribution $(\mathbf{f}, \mathbf{h}, \mathbf{R}, z, H'(\mathbf{b}^{(i)}))$.

Then \mathcal{A} distinguishes **Hybrid** 0 from **Hybrid** ℓ with probability ϵ . Choose an i at random from $[\ell]$. Then the adversary distinguishes **Hybrid** $i - 1$ from **Hybrid** i with expected advantage at least ϵ/ℓ . Next, observe that since bits $i + 1$ through ℓ are random in either case, they can be simulated independently of the challenge. Moreover, $H'(\mathbf{b})$ can be computed given $H'_{i-1}(\mathbf{b}_{[i-1]})$, b_i (be it random or equal to $\mathbf{R}_i \cdot \mathbf{x}$), and the random b_{i+1}, \dots, b_ℓ . Thus, we can construct an adversary \mathcal{A}' that distinguishes the following distributions:

$$(i, \mathbf{f}, \mathbf{h}, \mathbf{R}_{[i-1]}, z, H'_{i-1}(\mathbf{R}_{[i-1]} \cdot \mathbf{x}), \mathbf{R}_i, \mathbf{R}_i \cdot \mathbf{x}) \text{ and } (i, \mathbf{f}, \mathbf{h}, \mathbf{R}_{[i-1]}, z, H'_{i-1}(\mathbf{R}_{[i-1]} \cdot \mathbf{x}), \mathbf{R}_i, b_i)$$

with advantage ϵ/ℓ , where i is chosen randomly in $[\ell]$, $\mathbf{R}_{[i-1]}$ consists of the first $i - 1$ rows of \mathbf{R} , \mathbf{R}_i is the i th row of \mathbf{R} , and b_i is a random bit.

Next, since $\epsilon/3\ell$ is non-negligible, there is an inverse polynomial δ such that $\epsilon/3\ell \geq \delta$ infinitely often. Then, there is a polynomial r such \mathcal{A}' cannot distinguish f_i generated as $\text{ELF.Gen}(M, r)$ from the honest f_i generated from $\text{ELF.Gen}(M, M)$, except with probability at most δ . This means, if we generate $f_i \leftarrow \text{ELF.Gen}(M, r)$, we have that \mathcal{A}' still distinguishes the distributions

$$(i, \mathbf{f}, \mathbf{h}, \mathbf{R}_{[i-1]}, z, H'_{i-1}(\mathbf{R}_{[i-1]} \cdot \mathbf{x}), \mathbf{R}_i, \mathbf{R}_i \cdot \mathbf{x}) \text{ and } (i, \mathbf{f}, \mathbf{h}, \mathbf{R}_{[i-1]}, z, H'_{i-1}(\mathbf{R}_{[i-1]} \cdot \mathbf{x}), \mathbf{R}_i, b_i) \quad (1)$$

with advantage $\epsilon' = \epsilon/\ell - 2\delta$. Put another way, given $(\mathbf{f}, \mathbf{h}, \mathbf{R}_{[i-1]}, z, H'_{i-1}(\mathbf{R}_{[i-1]} \cdot \mathbf{x}), \mathbf{R}_i)$, \mathcal{A}' is able to compute $\mathbf{R}_i \cdot \mathbf{x}$ with probability $\frac{1}{2} + \epsilon'$. Note that $\epsilon' \geq \delta$ infinitely often, and is therefore non-negligible.

Now fix $\mathbf{f}, \mathbf{h}, \mathbf{R}_{[i-1]}$, which fixes H'_{i-1} . Let $y_i = H'_{i-1}(\mathbf{R}_{[i-1]} \cdot \mathbf{x})$. Notice that since \mathbf{f}, \mathbf{h} are fixed, there are at most r possible values for y_i , and recall that r is a polynomial. We now make the following claim:

Claim 6.7. *Let \mathcal{D} be a computationally unpredictable distribution on $\mathcal{X} \times \mathcal{Z}$. Suppose $T : \mathcal{X} \rightarrow \mathcal{R}$ is drawn from a family \mathcal{T} of efficient functions where the size of the image of T is polynomial. Then the following distribution is also computationally unpredictable: $(x, (T, z, T(x)))$ where $T \leftarrow \mathcal{T}$, $(x, z) \leftarrow \mathcal{D}$.*

Proof. Suppose we have an efficient adversary \mathcal{B} that predicts x with non-negligible probability γ given $T, z, T(x)$, and suppose T has polynomial image size r . We then construct a new adversary \mathcal{C} that, given x , samples a random T , samples $(x', z') \leftarrow \mathcal{D}$, and sets $a = T(x')$. It then runs $\mathcal{B}(T, z, a)$ to get a string x'' , which it outputs. Notice that a is sampled from the same distribution as $T(x)$, so with probability at least $1/r$, $a = T(x)$. In this case, $x'' = x$ with probability γ . Therefore, \mathcal{C} outputs x with probability γ/r , which is non-negligible. \square

Using Claim 6.7 with $T = H'_{i-1}(\mathbf{R}_{[i-1]} \cdot \mathbf{x})$, we see that $(\mathbf{x}, (i, \mathbf{f}, \mathbf{h}, \mathbf{R}_{[i-1]}, z, H'_{i-1}(\mathbf{R}_{[i-1]} \cdot \mathbf{x})))$ is computationally unpredictable. Moreover, $\mathbf{R}_i \cdot \mathbf{x}$ is a Goldreich-Levin [GL89] hardcore bit. We therefore recall the following theorem, adapted from [GL89]:

Theorem 6.8 (Adapted from [GL89]). *Let \mathcal{D} be a computationally unpredictable distribution on $\mathcal{X} \times \mathcal{Z}$, where $\mathcal{X} = \{0, 1\}^q$. Then the following two distributions are computationally indistinguishable:*

$$(z, \mathbf{r}, \mathbf{r} \cdot \mathbf{x}) \text{ where } (\mathbf{x}, z) \leftarrow \mathcal{D} \text{ and } \mathbf{r} \leftarrow \{0, 1\}^q$$

$$(z, \mathbf{r}, b) \text{ where } (\mathbf{x}, z) \leftarrow \mathcal{D}, \mathbf{r} \leftarrow \{0, 1\}^q, \text{ and } b \leftarrow \{0, 1\}$$

Applying the Goldreich-Levin theorem to the computationally unpredictable distribution $(\mathbf{x}, (i, \mathbf{f}, \mathbf{h}, \mathbf{R}_{[i-1]}, z, H'_{i-1}(\mathbf{R}_{[i-1]} \cdot \mathbf{x})))$, we see that the two distributions from Line 1 above are indistinguishable. This contradicts the existence of \mathcal{A}' , proving Theorem 6.6. \square

6.2 Applications

Polynomially-many hardcore bits for any one-way function. We see that H immediately gives us a hardcore function of arbitrary stretch for any computationally unpredictable distribution. This includes any one-way function. To the best of our knowledge, this is the first hardcore function of arbitrary stretch for general computationally unpredictable sources. In the special case of one-way functions, the only prior constructions are due to Bellare, Stepanovs, and Tessaro [BST14] using differing inputs obfuscation (diO), and of Zhandry [Zha16] using extractable witness PRFs. Both diO and extractable witness PRFs constitute very strong knowledge assumptions, and the plausibility of both have been significantly challenged [GGHW14]. In the case of *injective* one-way functions, [BST14] can relax the assumption to indistinguishability obfuscation (iO), which still remains a very strong primitive based on very new and relatively untested assumptions on multilinear maps. Our construction offers an entirely different approach to constructing hardcore functions with arbitrary stretch, and is based on a very simple primitive. While the only instantiation of our primitive so far requires less-than-standard elliptic curve assumptions, elliptic curves and assumptions on them, including ours, are far better understood.

Strong injective one-way functions. Bitansky and Paneth [BP12] conjecture the existence of a very strong one-way permutation family. We demonstrate that our function H meets this notion of security. Unfortunately, however, it is only injective, not a permutation.

Definition 6.9. A [BP12] one-way function is a family of functions H such that for any computationally unpredictable distribution \mathcal{D} , the following two distributions are also unpredictable:

$$(x, (z, H, H(x))), \text{ and } (H(x), (z, H)) \quad (\text{where } (x, z) \leftarrow \mathcal{D})$$

The first property is a generalization of a strong uninvertability assumption of Wee [Wee05]. The second guarantees that if x is unpredictable, then so is $H(x)$. We now show that our construction H satisfies this definition:

Theorem 6.10. *H constructed above using a secure ELF, when set to be injective as in Claim 6.5, is a [BP12] injective one-way function.*

Proof. For the first property, suppose that an efficient adversary \mathcal{A} can predict x given $(z, H, H(x))$. We now build an adversary \mathcal{A}' that can distinguish $(z, H, H(x))$ from (z, H, S) for a random S . $\mathcal{A}'(z, H, S)$ runs $\mathcal{A}(z, H, S)$ to get a point x' . It then tests if $H(x') = S$. If so, it outputs 1 and otherwise outputs 0. In the case that $S = H(x)$, by assumption \mathcal{A} will output $x' = x$ with non-negligible probability, meaning \mathcal{A}' will output 1 with non-negligible probability. In the case

that S is random, since H is expanding, with overwhelming probability there is no pre-image of S under H . Therefore \mathcal{A} will output 1 with negligible probability. Thus \mathcal{A}' successfully distinguishes the two cases, violating Theorem 6.6.

For the second property, suppose an adversary \mathcal{A} can predict $H(x)$ from z, H . We construct an adversary \mathcal{A}' that can distinguish $(z, H, H(x))$ from (z, H, S) for a random S . $\mathcal{A}'(z, H, S)$ runs $\mathcal{A}(z, H)$ to get a string S' , and then outputs 1 if and only if $S' = S$. In the case where $S = H(x)$, by assumption \mathcal{A} will output S , and so \mathcal{A}' will output 1, with non-negligible probability. In contrast, when S is chosen at random it is independent of the view of \mathcal{A} . Therefore \mathcal{A} will output S , and so \mathcal{A}' will output 1, with negligible probability. Thus \mathcal{A}' successfully distinguishes the two cases, violating Theorem 6.6. \square

The main application of Bitansky and Paneths [BP12] assumption is to build auxiliary input point function obfuscation (AIPO). Unfortunately, they require the one-way function to be a permutation; since our H is not a permutation, it cannot be immediately plugged into their construction. Yet, next, we show that going through their construction is unnecessary in our case: we show that our function H gives an AIPO “out of the box” with no additional overhead.

Point function obfuscation with auxiliary input (AIPO). We now show how to achieve full AIPO using just the assumption of ELFs.

Definition 6.11. A *auxiliary input point obfuscator* (AIPO) is an efficient probabilistic algorithm \mathcal{O} that satisfies the *almost perfect correctness* requirement of Definition 4.1, as well as the following secrecy requirement: for any unpredictable distribution \mathcal{D} over pairs $(x, z) \in \mathcal{X} \times \mathcal{Z}$, the following distributions are computationally indistinguishable:

$$(\mathcal{O}(I_x), z) : (x, z) \leftarrow \mathcal{D} \text{ and } (\mathcal{O}(I_{x'}), z) : (x, z) \leftarrow \mathcal{D}; x' \leftarrow \mathcal{X}$$

As in Section 4, an expanding ideal hash function (random oracle) H gives a very natural AIPO: the obfuscation of a point function I_x is simply $S = H(x)$. Injectivity of H gives (almost perfect) correctness. Moreover, security is easily proved in the random oracle model.

We now show that that by choosing H to be as in the construction above, the same is true. In particular, by Claim 6.5, H is injective in the same regime of input/output sizes as a random oracle. For security, we have the following:

Theorem 6.12. *The obfuscation construction described above is a secure AIPO assuming H is constructed as in Construction 6.3 using a secure ELF.*

Proof. Note that since H is expanding, if we choose S at random from $[2^m]$, then with overwhelming probability there are no inputs \mathbf{x} that map to S . Therefore, the obfuscated program corresponding to S is just the all-zeros function.

Let \mathcal{D} be any computationally unpredictable source. We thus need to show that the following two distributions are indistinguishable: $(H, z, H(\mathbf{x}))$ and (H, z, S) (where $(\mathbf{x}, z) \leftarrow \mathcal{D}$). This follows immediately from Theorem 6.6. \square

Public key encryption from trapdoor permutations. We show how our hardcore function can be used in the hybrid encryption schemes of Bellare and Rogaway [BR93] for converting a trapdoor permutation into a public key encryption scheme. Recall the [BR93] public key encryption scheme:

- $\text{PKE.Gen}(\lambda)$: run $(P, P^{-1}) \leftarrow \text{TDP.Gen}(\lambda)$. Choose a random H . Then the secret key for the scheme is $\text{sk} = (P^{-1}, H)$ and the public key is $\text{pk} = (P, H)$.
- $\text{PKE.Enc}(\text{pk}, m)$: Choose a random r in the domain of P . Output $c = (P(r), H(r) \oplus m)$. That is, the ciphertext has two components, $P(r)$, and a one-time pad encryption of m using $H(r)$ as the secret key.
- $\text{PKE.Dec}(\text{sk}, c = (c_1, c_2))$: let $r = P^{-1}(c_1)$, and then output $m = H(r) \oplus c_2$.

The proof of the following trivially follows from Theorem 6.6 and the fact that P is a one-way function so that $H(r)$ is pseudorandom given $P(r)$:

Theorem 6.13. *If P is a secure trapdoor permutation and H is constructed as in Construction 6.3 using a secure ELF , then PKE described above is a secure public key encryption scheme.*

6.3 Difficulty of Realizing Applications

Since AIPO implies PO, AIPO implies that NP with a super-logarithmic amount of non-determinism cannot be solved in polynomial time. Hence, this separation is inherent to the AIPO application. As an immediate corollary, we also have that our pseudorandom generator definition also implies such a separation. Since our pseudorandom generator definition is essentially equivalent to obtaining hardcore functions of arbitrary span for any unpredictable source, we also see that such a separation is inherent to such hardcore functions.

Here, we show that, if we are only interested in hardcore functions for the special case of one-way functions, some form of exponential hardness seems necessary without a major breakthrough in complexity theory.

Definition 6.14. A family of one-way functions $F_N : [N] \rightarrow \mathcal{Y}_N$ is *almost exponentially secure* if, for any super-logarithmic N , any polynomial-time algorithm can invert F_N on a random input with negligible probability.

Theorem 6.15. *Let H be a family of hardcore functions for any one-way function, where the hardcore function doubles the length of its input. Then if almost exponentially secure one-way functions exist, H must also be an almost exponentially secure one-way function.*

Therefore, proving that any H is a family of hardcore functions seems to require either (1) proving that no almost exponentially secure one way functions exist (which would be a major complexity theoretic breakthrough), or (2) that H itself is an almost exponentially secure one-way function⁹. (2) seems like the most likely scenario, which would mean that whatever underlying primitive is used to build H must also have some sort of (almost) exponential hardness.

Proof. Let F be almost exponentially secure. Let $N = N(\lambda)$ be any super-logarithmic function, and consider the function $H : \{0, 1\}^n \rightarrow \{0, 1\}^{2n}$. Then $F'_\lambda = F_N$ is a one-way function. Since H is hardcore, we have that no polynomial time algorithm can distinguish $(H, F(x), H(x))$ from $(H, F(x), S)$ where x is random in $\{0, 1\}^n$ and S is random in $\{0, 1\}^{2n}$. In particular, no polynomial time algorithm can distinguish $(H, H(x))$ from (H, S) . Since H is expanding, this means it is impossible in polynomial time to compute x from $H, H(x)$. Thus H is an almost exponentially secure one-way function. \square

⁹A third option is that H is a “universal” almost exponentially secure one-way functions, meaning that H is an exponentially secure one-way function if and only if such one-way functions exist. [Gol01].

7 Full Domain Hash Signatures

The Full Domain Hash signature scheme [BR93] makes use of a trapdoor permutation P with trapdoor P^{-1} , and a random oracle O whose range is the same as the domain of P . The signature on a message m is the pre-image under P of $H(m)$: $P^{-1}(H(m))$. Verifying the signature is straightforward. The proof of security in the random oracle model works roughly as follows. We are given a challenge $y = P(x)$ for some unknown x , and our goal is to use a forger to find x . We need to embed our challenge y into the forger’s view. We do this by programming the random oracle seen by the adversary: at a randomly chosen oracle query (say on message m^*), we set the output of H to be y (the rest of the outputs are sampled at random). If the adversary makes q queries, with probability $1/q$, the adversary’s forgery will be on m^* . In this case, the forgery is the pre-image of y , namely x . Thus, with probability $1/q$ the forger enables us to invert the trapdoor permutation.

Unfortunately, this proof technique seems to inherently require programming the random oracle H , and the programming needs to be *adaptive* (since we do not know a priori which points the adversary will query on). There is no standard model analog for such programming: for any concrete implementation of H , once we publish H (which must be a part of the public key), we have fixed the value of H at all messages.

One alternative is to settle for *selective* security for the signature scheme, where the adversary commits to the message m he will sign before receiving the public key. Here, H only needs to be programmed at m , which will be known in the security proof when generating H . Therefore, a standard-model instantiation may be possible. Indeed, Hohenberger, Sahai, and Waters [HSW14] show that this is possible assuming indistinguishability obfuscation. Adaptive security can be obtained at the cost of an exponential loss in the security reduction, therefore requiring the exponential hardness of *all* primitives involved, including the trapdoor permutation and the indistinguishability obfuscator. [HSW14] additionally show how to obtain adaptive security in the particular case where the trapdoor permutation is instantiated with the RSA permutation.

Beyond achieving only static security for general TDPs, another limitation of [HSW14] is that the hash function code depends not only on the TDP, but also on the specific public key for the TDP that the hash function will be used with. Therefore, H cannot be a global reusable hash function, but must instead be re-instantiated for every separate public key. This is in contrast to typical heuristic instantiations of the random oracle used in practice, which use public standardized hash functions such as SHA256.

We note that again, (sub)exponential hardness provides a solution to this limitation as well. The idea is to start with a universal H , and then in the reduction convert it to a permutation-dependent H as in [HSW14], and proceed using their proof to argue security. The reduction involves stepping through every input to H one at a time to change the output on that particular input to match [HSW14]; hence an exponential loss.

However, using complexity leveraging in this setting, besides having to assume that the various primitives involved are *very* secure, has a practical implication. If some trusted authority were to generate a global hash function as above, it likely will not be useable with existing trapdoor permutation keys, at least not in a sound manner. The issue is that, due to the exponential loss, all the building blocks, including the trapdoor permutation, need to have key sizes set appropriately large to as to handle this loss. Most currently deployed cryptographic keys are *not* set with this loss in mind. Therefore, in order to use this new hash function, all users would have to republish new, longer keys, which is a significant administrative hurdle.

In contrast, if the permutation was only required to be polynomially secure, existing deployed

keys for trapdoor permutations could be used with the new hash function.

To avoid the exponential loss in the reduction, we now describe another proof strategy, derived from Zhandry [Zha12], that does not require adaptive programming, and is independent of the TDP and public key (except for depending on a size bound for the TDP evaluation circuit). Suppose the adversary makes q_H oracle queries to H . Consider the following non-uniform distribution on H : first sample $r = O(q_H^{1/2})$ random values y_1, \dots, y_r . Then for each input x , set the output $H(x)$ to be chosen randomly from $\{y_i\}_{i \in [r]}$. Put another way, define H to be the composition $H = H_2 \circ H_1$, where $H_1 : \mathcal{X} \rightarrow [r]$ and $H_2 : [r] \rightarrow \mathcal{Y}$. Following Zhandry, we call functions from this distribution *small-range* functions.

From the adversary’s perspective, as long as none of the q queries form a collision in H , such small range functions are indistinguishable from the truly random case. Moreover by our choice of r , with probability $1/2$, there will be no collisions in the small-range case. Thus, if our adversary forges with probability 1 in the truly random case, he will still forge with probability $1/2$ in the small-range case.

Now, rather than choosing one of the adversary’s queries at random to insert y , we instead set one of the y_i to be y . That is, we choose a random $i \in [r]$, and set $y_i = y$. For all other i , we set $y_i = P(x_i)$ for randomly chosen x_i . As long as the adversary never makes a signature query on a message m such that $H(m) = y$, the reduction will be able to produce the signature as one of the x_i . Let q_S be the number of signing queries. We can assume that for every signing query on message m , the adversary also makes a oracle query to H on m , so that $q_S \leq q_H$. Then the probability that we can answer all signing queries is at least $1 - (1 - 1/q_H)^{q_S} \geq 1 - (1 - 1/q_S)^{q_S} \geq 1 - 1/e$. Moreover, with probability $1/r$, a forgery produced by the adversary will invert y . Therefore, with non-negligible probability, we can successfully use the adversary to invert the permutation.

Boneh and Zhandry [BZ13] uses this flavor of proof strategies to prove the security of several signature schemes against certain models of quantum attacks. Even though we are not concerned with quantum attacks, the proof strategy is still useful in our context because now the random oracle can be programmed *statically* — all of the outputs of O can be chosen up front. Now we may hope for a standard-model instantiation of O that allows for this proof technique to work. Indeed, by setting H_1 above to be an ELF, we get the necessary polynomial image size for the proof. We can then set H_2 to be a PRF, giving us (pseudo)random outputs.

Finally, since the hash function is made public, some difficulties arise. In the proof, we need to switch from the honest hash function where the y_i are random to one where $y_i = P(x_i)$ for known x_i . This is needed so we can sign messages other than the challenge. The two distributions on y_i are the same, but the code for the hash function will be different in the two cases. To make the proof work, we obfuscate the composition using an indistinguishably obfuscator, which allows us to hide the change.

7.1 Definitions

Puncturable PRFs. An $(\mathcal{X}, \mathcal{Y})$ -pseudorandom function PRF with domain \mathcal{X} and co-domain \mathcal{Y} consists of a polynomial-time algorithm $\text{PRF.Gen}()$ that outputs the description of an efficiently computable function $g : \mathcal{X} \rightarrow \mathcal{Y}$. The usual security requirement for a PRF is that oracle access to a random $g \leftarrow \text{PRF.Gen}()$ is indistinguishable from oracle access to a truly random function from \mathcal{X} to \mathcal{Y} .

A puncturable PRF has an additional algorithm PRF.Punct that takes as input g and a domain

point $x \in \mathcal{X}$, and outputs a “punctured key”, namely the description of a function $g^x : \mathcal{X} \rightarrow \mathcal{Y} \cup \{\perp\}$ where $g^x(y) = \begin{cases} g(y) & \text{if } y \neq x \\ \perp & \text{if } y = x \end{cases}$. Punctured PRF security is the following: the adversary chooses a point $x \in \mathcal{X}$, and receives g^x where $g \leftarrow \text{PRF.Gen}()$. Then the adversary, given g^x , is tasked with distinguishing $g(x)$ from a random string in \mathcal{Y} .

Indistinguishability Obfuscation. An *indistinguishability obfuscator* IO is a PPT uniform algorithm satisfying the following conditions:

- $\text{IO}(C)$ preserves the functionality of C . That is, for any polynomial-sized circuit C , if we compute $C' = \text{IO}(C)$, then $C'(x) = C(x)$ for all inputs x .
- For any two circuits C_0, C_1 with the same functionality and size, the circuits $\text{IO}(C_0)$ and $\text{IO}(C_1)$ are computationally indistinguishable.

The first candidate construction of such obfuscators is due to Garg et al. [GGH⁺13].

7.2 The Construction

Construction 7.1. Let \mathcal{M} be the domain of H (the message space), and \mathcal{Y} be the co-domain. To generate H , to the following:

- Let ELF be an ELF. Let $f \leftarrow \text{ELF.Gen}(|\mathcal{M}|, |\mathcal{M}|)$ to get an injective-mode f .
- Let PRF be a puncturable PRF with domain \mathcal{Z} and co-domain \mathcal{X} . Let $g \leftarrow \text{PRF.Gen}()$.
- Let IO be an indistinguishability obfuscator.
- Output $H = \text{IO}(g \circ f)$, where $g \circ f$ is appropriately padded before obfuscating.

Security Proof For FDH Signatures. We recall the full-domain hash (FDH) signature scheme using a hash function H . The starting point for the scheme is trapdoor permutation TDP, which is used to build a signature scheme SIG .

- $\text{SIG.Gen}()$: Run $(P, P^{-1}) \leftarrow \text{TDP.Gen}()$. Output $\text{sk} = P^{-1}$ and $\text{pk} = P$.
- $\text{SIG.Sign}(\text{sk}, m)$: run $\tilde{m} \leftarrow H(m)$ and output $\sigma \leftarrow P^{-1}(\tilde{m})$
- $\text{SIG.Ver}(\text{pk}, m, \sigma)$: run $\tilde{m} \leftarrow H(m)$, and verifies that $P(\sigma) = \tilde{m}$.

Theorem 7.2. *Assuming that ELF is a secure ELF, PRF is a secure PRF, and IO is a secure indistinguishability obfuscation, then if H is instantiated as in Construction 7.1 (and the obfuscation is appropriately padded), SIG is existentially unforgeable under a chosen message attack*

Proof. We prove security through a sequence of hybrids.

Hybrid 0 This is the standard chosen message attack. Suppose the adversary runs in time s and produces an existential forgery (m^*, σ^*) with non-negligible probability ϵ . Let δ be an inverse polynomial such that $\epsilon \geq \delta$ infinitely often.

Let $\{m_1, \dots, m_q\}$ be the messages for which the adversary requested signatures, and let σ_i be the corresponding signature. Let $\widehat{m}_i = f(m_i)$ and $\widehat{m}^* = f(m^*)$.

Hybrid 1 Here, we add a check that $\widehat{m}^* \notin \{\widehat{m}_1, \dots, \widehat{m}_q\}$. If the check fails, then we abort and say that the adversary failed. Notice that since f is injective, this is equivalent to $m^* \notin \{m_1, \dots, m_q\}$, which is already required. Thus this check is redundant and the adversary still succeeds with probability ϵ in **Hybrid 1**.

Hybrid 2 Here, we change f to be lossy. That is, we generate $f \leftarrow \text{ELF.Gen}(p(s, 2/\delta))$. By the security of the ELF, **Hybrid 2** is indistinguishable from **Hybrid 1**, except with probability $\delta/2$. Thus, the adversary produces an existential forgery with probability at least $\epsilon - \delta/2$.

Hybrid 3 Let $r \leq p(s, 2/\delta)$ be the image size of f . We modify **Hybrid 2** as follows. Choose a random \widehat{m}^* in the image of f using the efficient enumeration procedure. Now accept the adversary's forgery only if $f(m^*) = \widehat{m}^*$. Notice that \widehat{m}^* is independent of the adversary's view, and therefore $f(m^*) = \widehat{m}^*$ happens with probability at least $1/r$. Thus, the adversary succeeds in **Hybrid 3** with probability at least $(\epsilon - \delta/2)/r$. Notice that this quantity is at least $\delta/2r$ infinitely often, and is therefore non-negligible.

Hybrid 4 Let y^* be a random point in \mathcal{X} (equivalently, let $y^* = P(x^*)$ for a random $x^* \in \mathcal{X}$). We now replace H with an obfuscation of the program $H_{P,f,g,y^*,\widehat{m}^*}$ given in Figure 5.

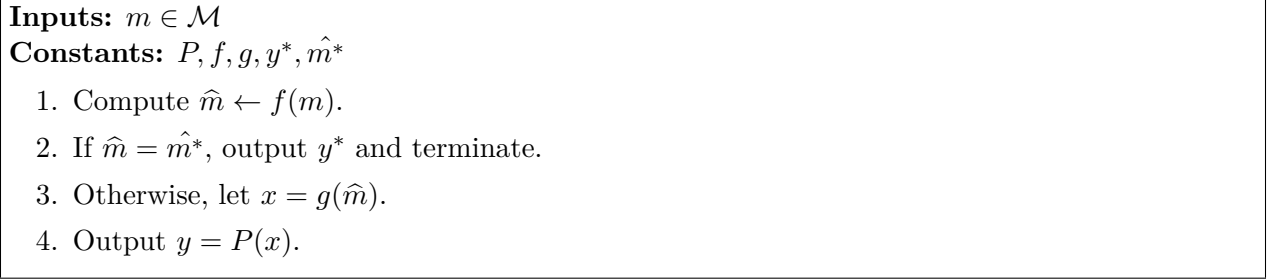


Figure 5: The program $H_{P,f,g,y^*,\widehat{m}^*}$.

Suppose our adversary succeeds in forging in **Hybrid 4** with non-negligible advantage. Then we can use such an adversary to invert P . Given a point $y \in \mathcal{X}$, we simulate **Hybrid 4** using $y^* = y$. If the adversary ever queries on a message m such that $\widehat{m} = f(m)$ equals \widehat{m}^* , abort the game (since **Hybrid 4** would have rejected anyway). To answer a signing query on message m , simply output $g(f(m))$. It is straightforward to see that, since $\widehat{m} \neq \widehat{m}^*$, this is the correct signature on m . Now the adversary's forgery is a message m^* and signature σ^* such that $P(\sigma^*) = H(m^*)$ and $f(m^*) = \widehat{m}^*$. Since $f(m^*) = \widehat{m}^*$, we have that $H(m^*) = H_{P,f,g,y^*,\widehat{m}^*}(m^*) = y^* = y$. Thus, σ^* is an inverse of y under P . Therefore we simply output σ^* . Our success probability is the same as the success probability of the adversary in **Hybrid 4**. It remains to prove that **Hybrid 4** is indistinguishable from **Hybrid 3**, which will complete the security proof.

To that end, endow \mathcal{Z} with a total order. Fix f and let $L = \{z_1, \dots, z_r\}$ be the sorted list of all r outputs of f (which can be efficiently enumerated by assumption). Let z_0 be smaller than the smallest element in \mathcal{Z} . Let i^* be the index such that $z_{i^*} = \widehat{m}^*$. Consider the following hybrids:

Hybrid 3.i for $i = 0, \dots, r$ This is the same as **Hybrid 3**, except that we generate two PRFs g_1, g_2 , and H is an obfuscation of the program $H_{P,f,g_1,g_2,y^*,\widehat{m}^*,z_i}$ given in in Figure 6

Inputs: $m \in \mathcal{M}$

Constants: $P, f, g_1, g_2, y^*, \hat{m}^*, z$

1. Compute $\hat{m} \leftarrow f(m)$.
2. If $\hat{m} \leq z$ and $\hat{m} \neq \hat{m}^*$, output $y = P(g_1(\hat{m}))$.
3. Otherwise, if $\hat{m} \leq z$ and $\hat{m} = \hat{m}^*$, output y^* and terminate.
4. Lastly, if $\hat{m} > z$, output $y = g_2(\hat{m})$.

Figure 6: The program $H_{P,f,g_1,g_2,y^*,\hat{m}^*,z}$.

First notice that $H_{P,f,g_1,g_2,y^*,\hat{m}^*,z_0}$ is functionally equivalent to $g_2 \circ f$. Thus the obfuscations are indistinguishable. This means **Hybrid 3** is indistinguishable from **Hybrid 3.0**. Similarly, $H_{P,f,g_1,g_2,y^*,\hat{m}^*,z_r}$ is functionally equivalent to $H_{P,f,g_2,y^*,\hat{m}^*}$ since all the image points of f are at most z_r . Therefore **Hybrid 3.r** is indistinguishable from **Hybrid 4**.

It remains to show that **Hybrid 3.i** – 1 is indistinguishable from **Hybrid 3.i**. This follows from a straightforward application of the punctured programming approach of Sahai and Waters [SW14] and the fact that $P(x)$ for a uniform x is identically distributed to a uniform y . We omit the details.

Piecing together, we see that **Hybrid 3** is indistinguishable from **Hybrid 4**, as desired.

Therefore **Hybrid 0** is indistinguishable from **Hybrid 4**, in which the adversary cannot forge. Thus no efficient adversary can forge in **Hybrid 0**. This completes the proof. \square

References

- [ABG⁺13] Prabhanjan Ananth, Dan Boneh, Sanjam Garg, Amit Sahai, and Mark Zhandry. Differing-inputs obfuscation and applications. Cryptology ePrint Archive, Report 2013/689, 2013. <http://eprint.iacr.org/2013/689>.
- [AGV09] Adi Akavia, Shafi Goldwasser, and Vinod Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In Omer Reingold, editor, *TCC 2009: 6th Theory of Cryptography Conference*, volume 5444 of *Lecture Notes in Computer Science*, pages 474–495. Springer, Heidelberg, Germany, March 15–17, 2009.
- [AHKM14] Daniel Apon, Yan Huang, Jonathan Katz, and Alex J. Malozemoff. Implementing cryptographic program obfuscation. Cryptology ePrint Archive, Report 2014/779, 2014. <http://eprint.iacr.org/2014/779>.
- [BCP14] Elette Boyle, Kai-Min Chung, and Rafael Pass. On extractability obfuscation. In Yehuda Lindell, editor, *TCC 2014: 11th Theory of Cryptography Conference*, volume 8349 of *Lecture Notes in Computer Science*, pages 52–73, San Diego, CA, USA, February 24–26, 2014. Springer, Heidelberg, Germany.
- [BCPR14] Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. On the existence of extractable one-way functions. In David B. Shmoys, editor, *46th Annual ACM Symposium*

on *Theory of Computing*, pages 505–514, New York, NY, USA, May 31 – June 3, 2014. ACM Press.

- [BFM14] Christina Brzuska, Pooya Farshim, and Arno Mittelbach. Indistinguishability obfuscation and UCEs: The case of computationally unpredictable sources. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 188–205, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
- [BH12] Itay Berman and Iftach Haitner. From non-adaptive to adaptive pseudorandom functions. In Ronald Cramer, editor, *TCC 2012: 9th Theory of Cryptography Conference*, volume 7194 of *Lecture Notes in Computer Science*, pages 357–368, Taormina, Sicily, Italy, March 19–21, 2012. Springer, Heidelberg, Germany.
- [BHJ⁺13] Florian Böhl, Dennis Hofheinz, Tibor Jäger, Jessica Koch, Jae Hong Seo, and Christoph Striecks. Practical signatures from standard assumptions. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, volume 7881 of *Lecture Notes in Computer Science*, pages 461–485, Athens, Greece, May 26–30, 2013. Springer, Heidelberg, Germany.
- [BHK13] Mihir Bellare, Viet Tung Hoang, and Sriram Keelveedhi. Instantiating random oracles via UCEs. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 398–415, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.
- [BP12] Nir Bitansky and Omer Paneth. Point obfuscation and 3-round zero-knowledge. In Ronald Cramer, editor, *TCC 2012: 9th Theory of Cryptography Conference*, volume 7194 of *Lecture Notes in Computer Science*, pages 190–208, Taormina, Sicily, Italy, March 19–21, 2012. Springer, Heidelberg, Germany.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In V. Ashby, editor, *ACM CCS 93: 1st Conference on Computer and Communications Security*, pages 62–73, Fairfax, Virginia, USA, November 3–5, 1993. ACM Press.
- [BST14] Mihir Bellare, Igors Stepanovs, and Stefano Tessaro. Poly-many hardcore bits for any one-way function and a framework for differing-inputs obfuscation. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology – ASIACRYPT 2014, Part II*, volume 8874 of *Lecture Notes in Computer Science*, pages 102–121, Kaoshiung, Taiwan, R.O.C., December 7–11, 2014. Springer, Heidelberg, Germany.
- [BST16] Mihir Bellare, Igors Stepanovs, and Stefano Tessaro. Contention in cryptoland: Obfuscation, leakage and UCE. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A: 13th Theory of Cryptography Conference, Part II*, volume 9563 of *Lecture Notes in Computer Science*, pages 542–564, Tel Aviv, Israel, January 10–13, 2016. Springer, Heidelberg, Germany.
- [BZ13] Dan Boneh and Mark Zhandry. Secure signatures and chosen ciphertext security in a quantum computing world. In Ran Canetti and Juan A. Garay, editors, *Advances in*

- Cryptology – CRYPTO 2013, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 361–379, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.
- [Can97] Ran Canetti. Towards realizing random oracles: Hash functions that hide all partial information. In Burton S. Kaliski Jr., editor, *Advances in Cryptology – CRYPTO’97*, volume 1294 of *Lecture Notes in Computer Science*, pages 455–469, Santa Barbara, CA, USA, August 17–21, 1997. Springer, Heidelberg, Germany.
- [CCR16] Ran Canetti, Yilei Chen, and Leonid Reyzin. On the correlation intractability of obfuscated pseudorandom functions. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A: 13th Theory of Cryptography Conference, Part I*, volume 9562 of *Lecture Notes in Computer Science*, pages 389–415, Tel Aviv, Israel, January 10–13, 2016. Springer, Heidelberg, Germany.
- [CGH98] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited (preliminary version). In *30th Annual ACM Symposium on Theory of Computing*, pages 209–218, Dallas, Texas, USA, May 23–26, 1998. ACM Press.
- [CMR98] Ran Canetti, Daniele Micciancio, and Omer Reingold. Perfectly one-way probabilistic hash functions (preliminary version). In *30th Annual ACM Symposium on Theory of Computing*, pages 131–140, Dallas, Texas, USA, May 23–26, 1998. ACM Press.
- [DP08] Stefan Dziembowski and Krzysztof Pietrzak. Leakage-resilient cryptography. In *49th Annual Symposium on Foundations of Computer Science*, pages 293–302, Philadelphia, Pennsylvania, USA, October 25–28, 2008. IEEE Computer Society Press.
- [DS05] Yevgeniy Dodis and Adam Smith. Correcting errors without leaking partial information. In Harold N. Gabow and Ronald Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 654–663, Baltimore, Maryland, USA, May 22–24, 2005. ACM Press.
- [DS15] Nico Döttling and Dominique Schröder. Efficient pseudorandom functions via on-the-fly adaptation. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *Advances in Cryptology – CRYPTO 2015, Part I*, volume 9215 of *Lecture Notes in Computer Science*, pages 329–350, Santa Barbara, CA, USA, August 16–20, 2015. Springer, Heidelberg, Germany.
- [FGK⁺10] David Mandell Freeman, Oded Goldreich, Eike Kiltz, Alon Rosen, and Gil Segev. More constructions of lossy and correlation-secure trapdoor functions. In Phong Q. Nguyen and David Pointcheval, editors, *PKC 2010: 13th International Conference on Theory and Practice of Public Key Cryptography*, volume 6056 of *Lecture Notes in Computer Science*, pages 279–295, Paris, France, May 26–28, 2010. Springer, Heidelberg, Germany.
- [GGH⁺13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual Symposium on Foundations of Computer Science*, pages 40–49, Berkeley, CA, USA, October 26–29, 2013. IEEE Computer Society Press.

- [GGHW14] Sanjam Garg, Craig Gentry, Shai Halevi, and Daniel Wichs. On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 518–535, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
- [GGSW13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th Annual ACM Symposium on Theory of Computing*, pages 467–476, Palo Alto, CA, USA, June 1–4, 2013. ACM Press.
- [GK16] Shafi Goldwasser and Yael Tauman Kalai. Cryptographic assumptions: A position paper. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A: 13th Theory of Cryptography Conference, Part I*, volume 9562 of *Lecture Notes in Computer Science*, pages 505–522, Tel Aviv, Israel, January 10–13, 2016. Springer, Heidelberg, Germany.
- [GL89] Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *21st Annual ACM Symposium on Theory of Computing*, pages 25–32, Seattle, Washington, USA, May 15–17, 1989. ACM Press.
- [Gol01] Oded Goldreich. *Foundations of Cryptography: Basic Tools*, volume 1. Cambridge University Press, Cambridge, UK, 2001.
- [GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd Annual ACM Symposium on Theory of Computing*, pages 99–108, San Jose, California, USA, June 6–8, 2011. ACM Press.
- [HJK⁺14] Dennis Hofheinz, Tibor Jager, Dakshita Khurana, Amit Sahai, Brent Waters, and Mark Zhandry. How to generate and use universal samplers. Cryptology ePrint Archive, Report 2014/507, 2014. <http://eprint.iacr.org/2014/507>.
- [HSW14] Susan Hohenberger, Amit Sahai, and Brent Waters. Replacing a random oracle: Full domain hash from indistinguishability obfuscation. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology – EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 201–220, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany.
- [ILL89] Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions (extended abstracts). In *21st Annual ACM Symposium on Theory of Computing*, pages 12–24, Seattle, Washington, USA, May 15–17, 1989. ACM Press.
- [KR00] Hugo Krawczyk and Tal Rabin. Chameleon signatures. In *ISOC Network and Distributed System Security Symposium – NDSS 2000*, San Diego, California, USA, February 2–4, 2000. The Internet Society.
- [Nao03] Moni Naor. On cryptographic assumptions and challenges (invited talk). In Dan Boneh, editor, *Advances in Cryptology – CRYPTO 2003*, volume 2729 of *Lecture Notes*

- in Computer Science*, pages 96–109, Santa Barbara, CA, USA, August 17–21, 2003. Springer, Heidelberg, Germany.
- [NS09] Moni Naor and Gil Segev. Public-key cryptosystems resilient to key leakage. In Shai Halevi, editor, *Advances in Cryptology – CRYPTO 2009*, volume 5677 of *Lecture Notes in Computer Science*, pages 18–35, Santa Barbara, CA, USA, August 16–20, 2009. Springer, Heidelberg, Germany.
- [PRS12] Krzysztof Pietrzak, Alon Rosen, and Gil Segev. Lossy functions do not amplify well. In Ronald Cramer, editor, *TCC 2012: 9th Theory of Cryptography Conference*, volume 7194 of *Lecture Notes in Computer Science*, pages 458–475, Taormina, Sicily, Italy, March 19–21, 2012. Springer, Heidelberg, Germany.
- [PS98] Sarvar Patel and Ganapathy S. Sundaram. An efficient discrete log pseudo random generator. In Hugo Krawczyk, editor, *Advances in Cryptology – CRYPTO’98*, volume 1462 of *Lecture Notes in Computer Science*, pages 304–317, Santa Barbara, CA, USA, August 23–27, 1998. Springer, Heidelberg, Germany.
- [PW08] Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In Richard E. Ladner and Cynthia Dwork, editors, *40th Annual ACM Symposium on Theory of Computing*, pages 187–196, Victoria, British Columbia, Canada, May 17–20, 2008. ACM Press.
- [Rao14] Vanishree Rao. Adaptive multiparty non-interactive key exchange without setup in the standard model. Cryptology ePrint Archive, Report 2014/910, 2014. <http://eprint.iacr.org/2014/910>.
- [RS09] Alon Rosen and Gil Segev. Chosen-ciphertext security via correlated products. In Omer Reingold, editor, *TCC 2009: 6th Theory of Cryptography Conference*, volume 5444 of *Lecture Notes in Computer Science*, pages 419–436. Springer, Heidelberg, Germany, March 15–17, 2009.
- [Sim98] 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*, volume 1403 of *Lecture Notes in Computer Science*, pages 334–345, Espoo, Finland, May 31 – June 4, 1998. Springer, Heidelberg, Germany.
- [SS90] A. W. Schift and Adi Shamir. The discrete log is very discreet. In *22nd Annual ACM Symposium on Theory of Computing*, pages 405–415, Baltimore, Maryland, USA, May 14–16, 1990. ACM Press.
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th Annual ACM Symposium on Theory of Computing*, pages 475–484, New York, NY, USA, May 31 – June 3, 2014. ACM Press.
- [Vil12] Jorge Luis Villar. Optimal reductions of some decisional problems to the rank problem. In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology – ASIACRYPT 2012*, volume 7658 of *Lecture Notes in Computer Science*, pages 80–97, Beijing, China, December 2–6, 2012. Springer, Heidelberg, Germany.

- [Wee05] Hoeteck Wee. On obfuscating point functions. In Harold N. Gabow and Ronald Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 523–532, Baltimore, Maryland, USA, May 22–24, 2005. ACM Press.
- [YY97] Adam Young and Moti Yung. Kleptography: Using cryptography against cryptography. In Walter Fumy, editor, *Advances in Cryptology – EUROCRYPT’97*, volume 1233 of *Lecture Notes in Computer Science*, pages 62–74, Konstanz, Germany, May 11–15, 1997. Springer, Heidelberg, Germany.
- [Zha12] Mark Zhandry. How to construct quantum random functions. In *53rd Annual Symposium on Foundations of Computer Science*, pages 679–687, New Brunswick, NJ, USA, October 20–23, 2012. IEEE Computer Society Press.
- [Zha16] Mark Zhandry. How to avoid obfuscation using witness PRFs. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A: 13th Theory of Cryptography Conference, Part II*, volume 9563 of *Lecture Notes in Computer Science*, pages 421–448, Tel Aviv, Israel, January 10–13, 2016. Springer, Heidelberg, Germany.