

Non-Interactive Zero Knowledge and Correlation Intractability from Circular-Secure FHE

Ran Canetti*

Alex Lombardi[†]

Daniel Wichs[‡]

December 31, 2018

Abstract

We construct non-interactive zero-knowledge (NIZK) arguments for NP from any circular-secure fully homomorphic encryption (FHE) scheme. In particular, we obtain such NIZKs under a circular-secure variant of the learning with errors (LWE) problem while only assuming a standard (poly/negligible) level of security. Our construction can be modified to obtain NIZKs which are either: (1) statistically zero-knowledge arguments in the common *random* string model or (2) statistically sound proofs in the common reference string model.

We obtain our result by constructing a new correlation-intractable hash family [Canetti, Goldreich, and Halevi, JACM '04] for a large class of relations, which suffices to apply the Fiat-Shamir heuristic to specific 3-message proof systems. In particular, assuming circular secure FHE, our hash function h ensures that for any function f of some a-priori bounded circuit size, it is hard to find an input x such that $h(x) = f(x)$. This continues a recent line of works [Holmgren and Lombardi, FOCS '18; Canetti et al., ePrint '18] focused on instantiating special forms of correlation intractability and Fiat-Shamir under weaker assumptions. Another consequence of our hash family construction is that, assuming circular-secure FHE, the classic quadratic residuosity protocol of [Goldwasser, Micali, and Rackoff, SICOMP '89] is *not* zero knowledge when repeated in parallel.

We also show that, under the plain LWE assumption (without circularity), our hash family is a *universal* correlation intractable family for general relations, in the following sense: If there exists *any* hash family of some description size that is correlation-intractable for general (even inefficient) relations, then our *specific* construction (with a comparable size) is correlation-intractable for general (efficiently verifiable) relations.

*Boston University and Tel Aviv University. Member of CPIIS. Supported by NSF awards 1413920 & 1801564, ISF award 1523/14. canetti@bu.edu.

[†]MIT. Email: alexjl@mit.edu. Research supported in part by an NDSEG fellowship. Research supported in part by NSF Grants CNS-1350619 and CNS-1414119, and by the Defense Advanced Research Projects Agency (DARPA) and the U.S. Army Research Office under contracts W911NF-15-C-0226 and W911NF-15-C-0236.

[‡]Northeastern. Email: wichs@ccs.neu.edu. Research supported by NSF grants CNS-1314722, CNS-1413964, CNS-1750795 and the Alfred P. Sloan Research Fellowship.

Contents

1	Introduction	1
1.1	Our Contributions	2
1.2	Prior Work on Correlation Intractability and Fiat-Shamir	5
1.3	Our Techniques	7
1.4	Organization	11
2	Preliminaries	11
2.1	(Lossy) Public Key Encryption	11
2.2	Fully Homomorphic Encryption and Circular Security	12
2.3	Non-Interactive Zero Knowledge Arguments (and Proofs)	13
3	Somewhere Statistically Correlation Intractable Hash Families	15
4	Correlation Intractability via Fully Homomorphic Encryption	17
4.1	Correlation Intractability for Efficiently Searchable Relations	17
4.2	Universal Correlation Intractability from LWE	19
4.3	Multi-Input Correlation Intractability	21
5	Non-Interactive Zero Knowledge Arguments	22
5.1	The [FLS99] Protocol	23
5.2	Our NIZK Protocol	23
5.3	Obtaining Theorem 1.3, Theorem 1.4, and LWE-based Instantiation	27
6	Fiat-Shamir for (Instance-Dependent) Trapdoor Σ-protocols	27
6.1	Instance-Dependent Trapdoor Σ -Protocols	28
6.2	Examples and Implications	29
	References	32

1 Introduction

Zero-knowledge (ZK) protocols, introduced by [GMR85], have been ubiquitous in cryptography for the last 30 years. At a high level, a zero-knowledge protocol Π is an interactive protocol between a prover P and a verifier V , in which the verifier V is *convinced* that some statement “ $x \in L$ ” is true but learns *nothing* beyond this fact. This “zero knowledge” property is formalized by the simulation paradigm: proving that an interaction between an honest prover P and any (potentially dishonest) verifier V^* can be *simulated* given only the verifier V^* and its input. We distinguish between *argument* systems with computational soundness and *proofs* where soundness is statistical.

An important and extremely useful variant of ZK protocols is *non-interactive zero-knowledge* (NIZK), where the protocol consists of a single message from the prover to the verifier. While it is known [GO94] that such NIZKs (or even 2 message zero-knowledge argument systems) for languages outside BPP do not exist in the plain model, we can construct NIZK proof systems in a setting where the prover and the verifier have access to a *common reference string* which is chosen from a predefined distribution, e.g. [BFM88, FLS99, CHK03, GOS12, SW14, BP15]. In this work, we make progress on two related open problems in the study of non-interactive zero-knowledge protocols.

Lattice-Based Non-Interactive Zero Knowledge. While it is known how to construct such NIZKs for NP under standard number-theoretic assumptions such as factoring and Bilinear Diffie-Hellman in prime-order elliptic-curve groups [BFM88, FLS99, CHK03], we do not know how to construct NIZK protocols based on *lattice assumptions* [Ajt96, AD97, Reg09] (except perhaps for extremely strong assumptions that suffice for indistinguishability obfuscation). In particular, we do not know how to construct NIZK protocols from any known variant of the learning with errors (LWE) problem [Reg09]. This stands in sharp contrast to the large body of work ([GPV08, PVW08, Gen09, BV11, BLMR13, GVW13, GKP⁺13, BV15, CC17, GKW17a, WZ17, GKW18], to name a few) that successfully constructed a variety of cryptographic applications from LWE and closely related lattice assumptions — including many applications where LWE-based realizations are the only known ones. In fact, NIZK has stood out as possibly *the* exceptional core cryptographic primitive that can be constructed from the above number-theoretic assumptions (which are notably all broken by polynomial-time quantum computers) but not lattice assumptions.

NIZK via the Fiat-Shamir Transform. A natural approach to constructing NIZK protocols is to use the Fiat-Shamir transform [FS86], which prescribes a general way to remove interaction from public-coin interactive proofs: To transform an interactive proof Π to a non-interactive argument, have the verifier first send a hash function h to the prover, and then have the prover compute the entire transcript of Π by itself, replacing the verifier’s challenges by the result of applying the hash function to the transcript so far (or portions thereof). The prover sends this entire transcript to the verifier in one message, and the verifier accepts if all checks verify.

Fiat and Shamir heuristically proposed to apply this methodology to a three-round identification protocol, using a fixed hash function such as DES (later instantiations use SHA and other cryptographic hash functions), with the goal of obtaining a signature scheme. Bellare and Rogaway [BR93] later showed that, when applied to a three-rounds honest-verifier Zero-Knowledge protocol with negligible soundness error, the Fiat-Shamir transform yields a NIZK protocol for the same language, *as long as the hash function is modeled as a random oracle*. Still, while this paradigm seems like a natural and attractive way to construct simple and efficient NIZK protocols,

finding explicit hash functions that suffice to make the approach work under well-defined hardness assumptions has proved to be elusive.

Furthermore, several works have demonstrated that such a hash function would have implications elsewhere, and might also be hard to come by. Specifically, [DNRS99] show that if there is a hash function \mathcal{H} instantiating the Fiat-Shamir transform for some three round public coin interactive proof Π , then Π is *not* (general verifier) zero knowledge. [Bar01, GK03] show that there exists some (artificially constructed) 3-round public-coin computationally sound argument Π for which the Fiat-Shamir heuristic fails to preserve soundness no matter what hash function is used to instantiate it. Furthermore, [BDG⁺13] show that no hash function family can be shown to suffice for the Fiat-Shamir via black-box reduction to a game-based assumption, even if one restrict attention to the case where the initial protocol is a three-round statistically sound proof. Nevertheless, it remains plausible that the Fiat-Shamir heuristic could be securely instantiated via some explicit hash family for specific classes of protocols. Showing that this is the case under standard assumptions is a long-standing open problem [BLV03].

Correlation Intractability and Recent Progress. Another hardness property for hash functions, which turns out to be easier to formalize, and closely related to “soundness for the Fiat-Shamir transform,” is *correlation intractability* (which was defined in [CGH04] for a different purpose). Roughly speaking, a hash function family \mathcal{H} is correlation intractable (CI) for relation $R(x, y)$ if it is computationally hard, given a random hash key k , to find any input x such that $(x, H_k(x)) \in R$. The most general class of relations typically considered is the set of all sparse relations: a relation R is sparse if for every x , the set of all y such that $(x, y) \in R$ is a negligible fraction of all possible values y . As observed in [HMR08], CI families (for this broadest possible class of relations) suffice for the soundness of the Fiat-Shamir transform, whenever the initial protocol is a statistically sound proof.

Initially this observation was taken as evidence for the hardness of constructing CI functions. Recently, however, a number of explicit hash function families were shown to be CI for certain classes of relations under well-defined assumptions [CCR16, KRR17, CCRR18, HL18, CCH⁺18]. Moreover, these hash functions were shown to be sound for the Fiat-Shamir transform for large classes of protocols. While these are significant advancements, the hardness assumptions used in these works are very strong and not well understood. See more discussion in Section 1.2.

1.1 Our Contributions

Our main result is a correlation-intractable hash family for a large class of relations, based on circular-secure fully homomorphic encryption (FHE). This is the first construction based on a “fully falsifiable” assumption: one defined via a game between an adversary and a polynomial-time challenger where we assume that every polynomial-time adversary has at most a negligible advantage in the game. Moreover, it is a cryptographic assumption that is widely used elsewhere; in particular, it is currently an essential step for obtaining fully (non-leveled) homomorphic encryption in the first place.

Our correlation-intractable hash family is powerful enough to instantiate the Fiat-Shamir transform for a certain class of public-coin proof systems. The class is quite broad; in particular, it suffices for obtaining NIZKs for all of NP. We provide two variants of this transformation: one variant results in NIZK protocol where the zero-knowledge property is *statistical*, and the CRS is “public coin” (in fact, it is uniformly distributed). The other variant results in a NIZK with *statistical*

soundness. The latter variant is especially surprising since, even in the random oracle model, the Fiat-Shamir transform only provides computational soundness and therefore our hash function has some advantages even over a random oracle. Furthermore, the two variants have reference strings that are indistinguishable from each other, so the resulting NIZK protocol has a “dual mode” property [DN01, GOS12].

In addition to the NIZK application, we show two other interesting applications of our hash family. One result – essentially following from [DNRS99] – is that assuming circular-secure FHE, a class of natural three-message public-coin protocols (which in particular includes the [GMR89] Quadratic Residuosity protocol) are *not* zero knowledge when repeated in parallel. This partially resolves open questions posed in [DNRS99, BLV03].

The other application (or extension) is that our hash family has the following interesting *universality* properties for correlation intractability, assuming only plain **LWE**: if any one of a class of hash functions is correlation intractable for all (even inefficiently verifiable) sparse relations, then our family is correlation intractable for all (efficiently verifiable and sufficiently sparse) relations. Remarkably, universality holds even for *multi-input* correlation intractability (namely, when the relation can depend on multiple inputs to the hash function and the corresponding outputs).

We now describe our contributions in more detail.

1.1.1 Correlation Intractability from Fully Homomorphic Encryption

We focus on obtaining correlation intractability for the following class of relations. First, we consider relations R where for every x there is a single y such that $R(x, y)$ holds. We let f denote the function that maps x to the corresponding y that makes $R(x, y)$ hold. That is, $R(x, y) = 1$ iff $y = f(x)$. We say that R is *searchable* in time T if f is computable in time T . We then construct, for each time bound T , a hash function family that is CI with respect to all relations that are searchable in time T . That is:

Theorem 1.1. *If there exists a circular secure fully homomorphic encryption scheme, then for every polynomial time bound $T = T(\lambda)$, every polynomial input size $n = n(\lambda)$ and every constant $\epsilon > 0$, there exists a hash family \mathcal{H} that is correlation intractable for all relations that are searchable in time T , with input size n and output size λ^ϵ .*

We emphasize that efficient searchability is quite different than the notions of efficient relations used in prior work [CCR16, HL18, CCH⁺18]. Still, we will see that it suffices for our needs. Moreover, our construction is very different from most prior work on correlation intractability: we show that a random key k is indistinguishable from a key k' for which there *do not exist* any x for which $h_k(x) = f(x)$. We call this property “somewhere statistical correlation intractability” in analogy to the notion of “somewhere statistically binding” hash functions of [HW15]. This statistical property is also what allows us to modify our NIZK argument system to obtain NIZK *proofs* rather than just arguments. See more details in Section 1.3

Universal CI. We also show that our *particular* hash function $h(k, x)$ with some fixed time bound T is correlation-intractable for general efficiently verifiable relations of sufficient sparsity, assuming that:

1. There exists *some* hash function $h'(\cdot, \cdot)$ of size T which is correlation-intractable for general (even inefficiently verifiable) relations of sufficiently smaller sparsity.

2. The FHE scheme is semantically secure (we do not rely on circular security for this result).

In addition, our universality argument even extends to the case of *multi-input* correlation intractability, about which very little is currently known.

We note that the flavor of universality demonstrated in this work is very different than other universality results which rely on “Levin’s trick” [Lev73]. Specifically, Levin’s trick involves guessing the description of a Turing Machine M that securely implements the primitive, and the resulting universal schemes incur a security loss which is exponential in the length of M . Although this is only a constant loss, it is likely to be quite large. In contrast, our universal scheme does not involve guessing a Turing Machine and does not incur the corresponding security loss. In fact, in contrast to Levin’s trick, our technique even works in the “non-uniform” setting: if we only start with the premise that there exists a non-uniform constructions of a secure correlation-intractable hash family, then our construction (which is uniform) is still secure (but the security reduction is non-uniform).

1.1.2 Applications to Fiat-Shamir and NIZK

By applying our hash family from Theorem 1.1 to a particular 3-round proof system for graph Hamiltonicity based on [FLS99], we obtain NIZK arguments in the common reference string model from any (circular-secure) fully homomorphic encryption scheme.

Theorem 1.2. *If there exists a circular-secure fully homomorphic encryption scheme, then there exist (adaptively sound) NIZK arguments for NP in the common reference string model.*

In fact, we prove two different strengthenings of Theorem 1.2: we construct (adaptively sound) non-interactive *statistical* zero-knowledge (NISZK) arguments for NP in the common *random* string model, and we construct statistically sound NIZK *proofs* for NP in the common reference string model.

Theorem 1.3. *If there exists a circular secure fully homomorphic encryption scheme, then there exist statistically sound NIZK proofs for NP in the common reference string model.*

Theorem 1.4. *Suppose that there exists a circular secure fully homomorphic encryption scheme with pseudorandom ciphertexts and public keys. Furthermore, suppose that there exists a lossy public key encryption scheme [KN08, PVW08, BHY09] with uniformly random lossy public keys. Then, there exist adaptively sound NISZK arguments for NP in the common random string model.*

The additional hypotheses of Theorem 1.4 are satisfied under LWE. Interestingly, to the best of our knowledge, we did not previously have NISZK argument systems in the common random string model from any standard cryptographic assumption (the [GOS12] NISZK argument system requires a non-random common reference string). Also, we previously did not have any approach toward achieving statistically sound proofs via the Fiat-Shamir heuristic.

Fiat-Shamir for Trapdoor Σ -Protocols As explained above, our hash family \mathcal{H} can be used to soundly instantiate the Fiat-Shamir heuristic for a particular modification of the 3-round proof system of [FLS99]. More generally, we can apply Fiat-Shamir to “trapdoor Σ -protocols”: roughly speaking, these are 3-message protocols Π in the common reference/random string (CRS) model with the following two properties:

- If the statement x is false, then for every first message \mathbf{a} , there is a *unique* challenge \mathbf{e} for which there is an accepting third message \mathbf{z} that results in an accepting transcript $(\mathbf{a}, \mathbf{e}, \mathbf{z})$.
- There is a *trapdoor* τ associated with the CRS that allows us to efficiently compute this “bad challenge” \mathbf{e} from the first message \mathbf{a} .

In this language, we modify the [FLS99] 3-round proof system to make it a “trapdoor Σ -protocol” by choosing a commitment scheme that has a commitment public key (which we put in the CRS) for which there exists a trapdoor allowing for extraction. Moreover, we define a generalization called an “instance-dependent trapdoor Σ -protocol”, in which the trapdoor is allowed to depend on the instance x , that also captures the *unmodified* [GMR89] protocol. We prove that our hash family suffices to instantiate Fiat-Shamir for all such protocols. By [DNRS99], this implies that the (parallel repeated) [GMR89] protocol is not zero knowledge.

Corollary 1.5. *Suppose that there exists a circular secure fully homomorphic encryption scheme, and further assume the hardness of quadratic residuosity. Then for any $\epsilon > 0$, the [GMR89] protocol for quadratic residuosity, repeated λ^ϵ times in parallel, is not zero knowledge.*

1.1.3 About the Assumption: Circular Secure FHE

We know how to construct leveled FHE under the learning with errors (LWE) assumption [BV11, BGV12, Bra12, GSW13, BV14] which is in turn as hard as worst-case lattice problems [Reg09, BLP⁺13, PRSD17]. As far as we know, it is reasonable to assume that any of these FHE schemes is circular secure, meaning that if we encrypt the secret key (one bit at a time) then this is indistinguishable from encrypting a dummy message consisting of all 0s. This is a “fully falsifiable” assumption where we only need to assume a standard poly/negligible level of security. In fact, this assumption is needed to perform bootstrapping [Gen09] and is currently the only known approach to get *fully* (non-leveled) homomorphic encryption. Moreover, circular security appears to be a very mild assumption: as far as we know all *natural* encryption schemes that are semantically secure are also circular secure. In fact, it is highly non-trivial to come up with even contrived constructions of semantically secure encryption schemes which are *not* circular secure and we had no such examples until fairly recently with the works of [Rot13, GKW17b, GKW17a, WZ17]. Although we do not know how to prove the circular security of any FHE candidate under LWE directly, we consider circular secure FHE to be a mild, fully falsifiable, lattice-based assumption. Our work achieves the first constructions of correlation-intractable hash functions, instantiations of Fiat-Shamir and NIZKs under such assumptions.

1.2 Prior Work on Correlation Intractability and Fiat-Shamir

This work continues a recent line of works [CCR16, KRR17, CCRR18, HL18, CCH⁺18] focused on constructing correlation-intractable hash families and using them to instantiate the Fiat-Shamir transform in the standard model. Throughout, we consider and compare the following main aspects of the constructions and assumptions in these works (we consider game-based assumptions):

- (a) Our level of familiarity with the assumption
- (b) Formal characteristics of the assumption, namely:

- (b.1) The complexity of the algorithm conducting the security game and deciding whether a purported adversary won the game (is it exponential in the security parameter?)
- (b.2) The bound on the allowed success probability of the adversary (is it exponential in the security parameter?)
- (c) The class of relations for which correlation intractability is achieved; in particular, whether the hash family is *compact* (namely whether the a single family of functions can withstand relations of arbitrary polynomial size).

The works are described below.

- [CCR16] constructs hash functions that are correlation intractable for all efficiently verifiable relations assuming subexponentially secure indistinguishability obfuscation (IO) [BGI⁺01, GGH⁺13] for all circuits as well as input-hiding obfuscation [BCKP14] for all evasive circuits. Both of these assumptions are non-standard; indeed, IO has no constructions from standard assumptions, and input-hiding obfuscation is even less understood. The [CCR16] construction is non-compact: the description size of the hash function depends polynomially on the maximum description size of the relations covered. Using ideas from [HL18], the [CCR16] construction can be used to instantiate the Fiat-Shamir heuristic for specific 3-message protocols of interest in a similar way to what is done in later works.
- [KRR17] (independently of [CCR16]) constructs hash functions that are correlation intractable for *all sparse relations*; in particular, they instantiate the Fiat-Shamir transform for all constant-round interactive proofs, yielding a construction of NIZK arguments as well as showing that no constant-round public-coin zero knowledge proofs exist. They do so assuming subexponential indistinguishability obfuscation, and in addition a strong variant of point function obfuscation satisfying a form of “fully exponential KDM security.” Roughly speaking, this requires that it is fully exponentially hard for a polynomial-time adversary to recover a point x^* given an obfuscation of a program that outputs $y^* = f(x^*)$ on a particular (random) input x^* , for *any* (possibly inefficient) function f .
- [CCRR18] obtains results similar to [KRR17], with significantly simpler and more efficient constructions that avoid obfuscation. Furthermore, their assumptions pertain to security properties of known constructs such as Regev and El-Gamal encryption. Still, their assumptions have the same strong flavors as those of [KRR17]: in particular, they require the existence of an encryption scheme with the property that it is fully exponentially hard to recover the secret key sk given a KDM-encryption $\text{Enc}(\text{sk}, f(\text{sk}))$ for any (possibly inefficient) function f (where Enc is either Regev or El-Gamal encryption).
- [HL18] constructs a correlation-intractable hash family for all relations sampleable in a bounded polynomial time, assuming subexponential IO and exponentially secure one-way functions. This removes the KDM-style assumption (as compared to [KRR17]) but retains the reliance on indistinguishability obfuscation and fully exponential hardness. Their construction is also non-compact. Still, their hash family suffices to instantiate the Fiat-Shamir heuristic for a wide class of 3-message protocols – a strictly broader class than the protocols that we can handle in this work.

Reference (Functionality)	Assumes IO?	Exotic Assumption?	Exp-time Challenger?	Exponential Probability?
[CCR16] (Non-Compact, verifiable relations)	Yes	Yes++	No	No
[KRR17] (Compact, all relations)	Yes	Yes	Yes	Yes
[CCRR18] (Compact, all relations)	No	No	Yes	Yes
[HL18] (Non-Compact, sampleable relations)	Yes	No	No	Yes
[CCH ⁺ 18] (Compact, sampleable relations)	No	No	No	Yes
This work (Non-Compact, searchable relations)	No	No	No	No

Figure 1: Constructions of Correlation Intractible Hash Families

- [CCH⁺18] constructs two correlation-intractable hash families for efficiently sampleable relations: a compact one (i.e., a single poly-size family that covers all poly-size relations) and a non-compact one. Unlike the [CCRR18] constructions, both of these hash families are secure under lattice assumptions that are falsifiable in polynomial time (but with exponentially small success probability). The compact family requires a form of circular security, whereas the non-compact one relies on plain search-LWE (albeit with fully exponentially small success probability). Their non-compact family suffices to instantiate Fiat-Shamir for a broad class of protocols similarly to [HL18, CCR16]. Furthermore, their compact scheme suffices for applying the Fiat-Shamir paradigm to the GKR interactive proof [GKR08], thereby obtaining a publicly verifiable succinct non-interactive argument for logspace-uniform NC without assuming indistinguishability obfuscation or non-interactive knowledge extraction from general adversaries.

In Fig. 1, we compare key features of the above works. As evident from our description, [KRR17] and all subsequent works have a “fully exponential success probability” barrier: they can only prove security under an assumption that polynomial time adversaries cannot solve some problem with probability significantly better than random guessing. This seems somewhat inherent to the proof technique used in all of these results.

We note that this work is a direct follow-up to [CCH⁺18]. Indeed, the results and techniques of [CCH⁺18] were the initial inspiration for this work.

1.3 Our Techniques

We give a high-level overview of the proofs of Theorem 1.1, Theorem 1.2, and Corollary 1.5. We begin with our main contribution: a new correlation-intractable hash family.

CI for Efficiently Searchable Relations. We construct a hash-function family $h(k, x)$ with a public hash-key k and input x that satisfies correlation-intractability for all “efficiently searchable

relations” with some fixed polynomial time bound T , meaning the following. For any function f having circuit size T , if a polynomial time adversary is given a random k , he cannot find an input x such that $h(k, x) = f(x)$. The output length of the hash function can be as low as λ^ϵ for any $\epsilon > 0$ and the input length can be an arbitrary polynomial in λ . Note that our hash family h only depends on the bound T but not on f ; it is secure for *all* functions f of size T . This is what makes the problem non-trivial – if we only wanted to achieve security for a particular f and had allowed the hash function to depend on f , we would have a trivial solution by defining $h(k, x)$ to compute $f(x)$ and (e.g.) flip the last bit. Interestingly, while in our actual construction the hash function does not depend on f , in the proof we switch it for one that does and use a variant of the above trivial idea to argue security.

Our construction is simple:

$$h(k, x) = \text{FHE.Eval}_{\text{pk}}(U_x, \text{ct}) \quad \text{where } k = (\text{pk}, \text{ct}) \text{ and } U_x(g) = g(x). \quad (1)$$

In particular, the hash function interprets the hash-key $k = (\text{pk}, \text{ct})$ as an FHE public key pk along with a ciphertext ct encrypting some circuit g with 1-bit output. The hash function then homomorphically evaluates the universal circuit $U_x(g) = g(x)$ over the ciphertext ct . We can choose the hash-key k by setting ct to be an FHE encryption of (a sufficiently long) all 0 string, or, if the FHE has pseudorandom public-keys and ciphertexts, we can even choose k as a uniformly random string.

Our proof of security is also simple. Assume that an adversary gets k and is able to find x such that $h(k, x) = f(x)$ with non-negligible probability. We first switch the ciphertext ct in the key k to be an FHE encryption of the circuit $g(x) = \text{Dec}_{\text{sk}}(f(x)) \oplus 1$, where sk is the FHE secret key. In other words, g first computes $f(x)$, then interprets it as an FHE ciphertext of a 1-bit plaintext, decrypts it and flips the output.¹ We argue that this change is indistinguishable by the security of the FHE. This requires circular security of FHE since the circuit g depends on sk .² Since the adversary cannot distinguish this change, he still outputs x such that $h(k, x) = f(x)$ with non-negligible probability. So, we have:

$$f(x) = h(k, x) = \text{FHE.Eval}_{\text{pk}}(U_x, \text{ct}) = \text{Enc}_{\text{pk}}(g(x)) = \text{Enc}_{\text{pk}}(\text{Dec}_{\text{sk}}(f(x)) \oplus 1).$$

However, this is a contradiction since, by applying $\text{Dec}_{\text{sk}}(\cdot)$ to both sides, we get $\text{Dec}_{\text{sk}}(f(x)) = \text{Dec}_{\text{sk}}(f(x)) \oplus 1$. In other words, once we switched ct to be an encryption of g , we ensured that there is no x for which $h(k, x) = f(x)$. This is because we ensure that $h(k, x)$ outputs a ciphertext that is guaranteed to decrypt to a different value than the value v obtained by applying the decryption algorithm to $f(x)$. (We stress that we do not assume any “semantic” meaning to the value v . Indeed, $f(x)$ is in general not a valid ciphertext, so there are no guarantees as to what v might be. Still, it is a well-defined binary value.)

We note that the above proof actually demonstrates a property stronger than plain correlation intractability: For any function f of size T , we can switch the hash-key k to a computationally indistinguishable hash-key $k' = k'_f$ such that the hash function is statistically correlation intractable

¹Without loss of generality, we assume that the decryption algorithm always outputs some bit $b \in \{0, 1\}$ and never rejects; e.g., if a ciphertext is invalid decryption outputs 0.

²In slightly more detail, instead of encrypting g directly we separately encrypt f, sk and then homomorphically compute g . This allows us to just rely on circular security (encrypting the secret key itself) rather than key-dependent message security (encrypting functions of the secret key).

for f : there does not exist any x such that $h(k', x) = f(x)$. This is reminiscent of the somewhere statistically binding hashing of [HW15]; we call such hash functions *somewhere statistically correlation intractable*.

In terms of parameters, the output size of the hash function needs to be as large as a single FHE ciphertext encrypting a single bit, which can be set to be as low as λ^ϵ for any $\epsilon > 0$. On the other hand, the size of the key $k = (\text{pk}, \text{ct})$ and the time to evaluate $h(k, \cdot)$ depend on T ; this is because ct needs to be large enough to encrypt f and U_x needs to be large enough to evaluate f .³

NIZKs for NP via Fiat-Shamir. We show that CI for efficiently searchable relations is sufficient to instantiate the Fiat-Shamir heuristic for a particular Σ -protocol (i.e. 3 round public-coin protocol with “special soundness”) and get NIZK arguments for NP. This follows the general framework explored in [HL18] and [CCH⁺18].

We take the Σ -protocol of [FLS99] for showing that a graph G has a Hamiltonian cycle. The prover sends a commitment to a random cycle graph C . The verifier sends a challenge bit. If the bit is 0, the prover decommits to the entire graph and the verifier checks that it is indeed a cycle. If the bit is 1, the prover sends a random permutation of G which maps the Hamiltonian cycle in G to C along with the opening of all the non-edges of the permuted G . We amplify soundness via parallel repetition. Borrowing an idea from [HL18], we use a public-key encryption scheme to implement the commitment, where the public key pk is in the CRS.

The above Σ -protocol has the following property. If the statement is false, then given the prover’s first message a , there is a unique “bad” challenge e that for which a valid response z exists. Furthermore, this “bad” challenge would be efficiently computable if we had all the committed values. Since we use a public-key encryption scheme as a commitment, we can extract the committed values from the commitment a using the encryption secret key sk . Combining the above, given the encryption secret key sk , there is an efficiently computable function $f_{\text{sk}}(a)$ which maps the prover’s first message a to the unique “bad” challenge e that has a valid response. The size of the function f_{sk} is bounded by some bound T .

We show that, if we apply the Fiat-Shamir heuristic to the above protocol and use a hash function which is correlation-intractable for all “efficiently searchable relations” with the bound T , we get a NIZK argument. Recall that the Fiat-Shamir heuristic adds a hash key k to the CRS and requires the prover to come up with a valid protocol transcript (a, e, z) where $e = h(k, a)$. Since the hash function is correlation-intractable for all “efficiently searchable relations” with the bound T , it is in particular correlation-intractable for f_{sk} . This means that an efficient prover cannot come up with a value a such that $h(k, a) = f_{\text{sk}}(a)$. But if the statement is false, then the only way a proof (a, e, z) can be valid is if $h(k, a) = e = f_{\text{sk}}(a)$. Therefore, the prover cannot come up with any valid proofs and we have soundness. The zero-knowledge (ZK) property of the NIZK follows from the honest-verifier zero-knowledge property of the Σ -protocol.⁴ We elaborate that in the above argument, we only need the hash function to be correlation-intractable for a particular function f_{sk} but since sk is secret (releasing it would break zero-knowledge) we rely on the fact that we can choose the hash key k in a way that does not reveal sk .

³If f has a succinct description as a Turing Machine that runs in time T , we can make the key k shorter than T by having ct encrypt the Turing Machine description of f and letting U_x be a universal circuit which takes as input a Turing Machine and runs it for T steps. Still, the time to evaluate $h(k, \cdot)$ depends on the run-time T .

⁴For this to work, we also need the hash function to be 1-universal, meaning that for any a the value $h(k, a)$ is uniformly random over the choice of k . We show that 1-universality can be generically added to any correlation-intractable hash function.

We obtain a statistically sound NIZK *proof* system by using a “statistically correlation intractable” hash key $k_{f_{\text{sk}}}$ instead of a plain hash key k . This makes use the fact that the function f_{sk} depends on a trapdoor to the 3-message CRS but not on the instance x . Now, when we argue zero-knowledge, we make use of the computational property that $k_{f_{\text{sk}}}$ is indistinguishable from a random k which does not depend on sk .

Finally, if we use a “lossy encryption” scheme to implement the commitments, we obtain statistical zero-knowledge. In this variant, we can also make the CRS truly random assuming that we have FHE where the public-keys and ciphertexts are pseudorandom (implied by circular LWE) and that we have “lossy encryption” with random public keys (implied by LWE).

Fiat-Shamir for the [GMR89] Protocol. By a very similar argument to our NIZK construction, we show that the [GMR89] Quadratic Residuosity protocol is not zero knowledge when repeated a large number of times in parallel, assuming the existence of correlation-intractable hash functions for efficiently searchable relations. This takes advantage of the aforementioned [DNRS99] result that if there exists a hash function that suffices for the Fiat-Shamir transform for a protocol Π (for a language $L \notin \text{BPP}$), then Π cannot be zero knowledge. In this overview, we use [GMR89] as an example for the general notion of an “instance-dependent trapdoor Σ -protocol” that we introduce.

Recall the [GMR89] protocol: in order to prove that a number y is a quadratic residue modulo a composite number $N = pq$, the prover sends to the verifier a random square $a = r^2$ modulo N ; the verifier sends a random bit e to the prover, at which point the prover reveals a square root of $a \cdot y^e$ (either r or rx for some square root x of y).

To show that our hash family suffices to instantiate Fiat-Shamir for the [GMR89] protocol (repeated in parallel), we note that the soundness of Fiat-Shamir for this protocol follows from correlation intractability for the function

$$f_N(\mathbf{a}) = \mathbf{e} := (e_i = QR(N, a_i))_i,$$

where $QR(N, a) = 1$ if and only if a_i is a square modulo N . This function f_N simply computes, for every first message \mathbf{a} in the (parallel repeated) [GMR89] protocol, the unique challenge $\mathbf{e} \in \{0, 1\}^t$ that a cheating prover has any hope of being able to win on (provided that the instance y is not a quadratic residue).

While f_N is not efficiently computable as a function of (N, \mathbf{a}) , it *is* efficiently computable given the factorization $N = pq$ as non-uniform advice, so we can show that our hash family is correlation intractable for every f_N and hence demonstrates our claimed result.

At a high level, the [GMR89] protocol is similar to our modified [FLS99] protocol in that no instances (N, y) have an associated “bad challenge function” f_N and there is a trapdoor (p, q) making f_N efficiently computable. However, in this case, the trapdoor depends on the instance (N, y) (as opposed to in the [FLS99] modification, where it only depends on the CRS). This motivates our definition of an “instance-dependent trapdoor Σ -protocol” in Section 6.

Universal CI. We also show that our *particular* hash function $h(k, x)$ described in (1) with some fixed time bound T is correlation-intractable for general efficiently decidable relations of sufficient sparsity, assuming that:

1. There exists *some* hash function $h'(\cdot, \cdot)$ of size T which is correlation-intractable for general (even inefficient) relations of sufficient (necessarily larger) sparsity.

2. The FHE scheme is semantically secure (we do not rely on circular security for this result).

To see why our construction is “universal”, assume for contradiction that there is some sufficiently sparse and efficiently computable relation R as well as an adversary that, given the key $k = (\text{pk}, \text{ct})$ of our hash function, computes x such that $(x, h(k, x)) \in R$ with non-negligible probability. We first switch ct to be an encryption of the correlation-intractable hash function $h'(k', \cdot)$ for a random k' . By the semantic security of the encryption, this is indistinguishable and therefore the adversary still produces x such that $(x, h(k, x)) \in R$ with non-negligible probability. Since $h'(k', \cdot)$ is correlation-intractable for all sufficiently sparse relations, it is in particular correlation-intractable for the (inefficient) relation⁵:

$$R_{\text{sk}}^* = \{(x, z) : \exists y \text{ such that } (x, y) \in R \text{ and } z = \text{Dec}(\text{sk}, y)\},$$

But if $(x, y = h(k, x)) \in R$ then $y = \text{Enc}_{\text{pk}}(h'(k', x))$ and therefore, for $z = \text{Dec}(\text{sk}, y)$, we have $(x, z = h'(k', x)) \in R_{\text{sk}}^*$. So the adversary also breaks the correlation intractability of $h'(k', \cdot)$ with respect to the relation R_{sk}^* , which should be impossible.

1.4 Organization

The remainder of the paper is organized as follows. In Section 2, we recall basic preliminaries, and in Section 3, we define correlation intractability [CGH04] and the specific variants focused on in this work. In Section 4, we present our main constructions of correlation intractable hash families from fully homomorphic encryption. Finally, we apply these hash families in Section 5 to obtain our main results (Theorem 1.2 and its extensions), and in Section 6 to obtain our most general Fiat-Shamir instantiation.

2 Preliminaries

We say that a function $\mu(\lambda)$ is *negligible* if $\mu(\lambda) = O(\lambda^{-c})$ for every constant c , and that two distribution ensembles $X = \{X_\lambda\}$ and $Y = \{Y_\lambda\}$ are computationally indistinguishable ($X \approx_c Y$) if for all polynomial-sized circuit ensembles $\{\mathcal{A}_\lambda\}$,

$$\left| \Pr[\mathcal{A}_\lambda(X_\lambda) = 1] - \Pr[\mathcal{A}_\lambda(Y_\lambda) = 1] \right| = \text{negl}(\lambda).$$

2.1 (Lossy) Public Key Encryption

Definition 2.1 (Public Key Encryption). *A public-key encryption scheme $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ consists of three p.p.t. algorithms:*

- $\text{Gen}(1^\lambda)$ takes as input the security parameter and outputs a public key pk and a secret key sk .
- $\text{Enc}(\text{pk}, m)$ takes as input the public key and a bit⁶ $m \in \{0, 1\}$; it outputs a ciphertext ct .
- $\text{Dec}(\text{sk}, \text{ct})$ takes as input the secret key and a ciphertext ct ; it outputs a message m' .

⁵For this argument to work, parameters must be set so that this relation is still sparse.

⁶As usual, this extends naturally to encrypting many-bit plaintexts.

PKE must furthermore satisfy the following properties.

- **Correctness:** For all λ , all $m \in \{0, 1\}$, and all (pk, sk) in the support of 1^λ , it holds with probability 1 that $\text{Dec}(sk, \text{Enc}(pk, m)) = m$.
- **Semantic Security:** The distribution ensembles $\{(pk, sk) \leftarrow \text{Gen}(1^\lambda) : (pk, \text{Enc}(pk, 0))\} \approx_c \{(pk, sk) \leftarrow \text{Gen}(1^\lambda) : (pk, \text{Enc}(pk, 1))\}$ are computationally indistinguishable.

We say that a public key encryption scheme PKE has pseudorandom ciphertexts if the distribution ensembles $\{(pk, sk) \leftarrow \text{Gen}(1^\lambda) : (pk, \text{Enc}(pk, 0))\} \approx_c \{(pk, sk) \leftarrow \text{Gen}(1^\lambda), u \leftarrow U_{|\text{Enc}(pk, 0)|} : (pk, u)\}$, where \approx_c denotes computational indistinguishability, and PKE has pseudorandom public keys if a public key pk sampled according to $\text{Gen}(1^\lambda)$ is computationally pseudorandom.

Definition 2.2 (Lossy PKE). A public-key encryption scheme PKE is said to be lossy [KN08, PVW08, BHY09] if there exists a fake key generation algorithm FakeGen such that:

- The (randomized) output \tilde{pk} of FakeGen(1^λ) is computationally indistinguishable from a public key pk sampled by Gen(1^λ).
- Encryption under fake keys is statistically hiding. That is,

$$\{(\tilde{pk}, \text{Enc}(\tilde{pk}, 0))\} \approx_s \{(\tilde{pk}, \text{Enc}(\tilde{pk}, 1))\},$$

where $\tilde{pk} \leftarrow \text{FakeGen}(1^\lambda)$ and \approx_s denotes statistical indistinguishability.

We say that PKE is lossy with uniformly random lossy public keys if (in addition) FakeGen outputs a uniformly random string.

In this work, we make use of the fact that public-key Regev encryption [Reg09] is lossy with uniformly random lossy public keys under the LWE assumption.

2.2 Fully Homomorphic Encryption and Circular Security

Definition 2.3. A fully homomorphic encryption scheme FHE = (Gen, Enc, Dec, Eval) consists of four p.p.t. algorithms such that (Gen, Enc, Dec) is a public key encryption scheme, and:

- Eval(pk, f, ct_1, \dots, ct_n) takes as input the public key, a function f (represented by a boolean circuit), and a vector of ciphertexts (ct_1, \dots, ct_n) ; it outputs another ciphertext ct' , which has size that is polynomial in λ (and, without loss of generality, linear in the output length of f).
- For any $(pk, sk) \leftarrow (\text{Gen}(1^\lambda))$, any $m_1, \dots, m_n \in \{0, 1\}$, and any circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}$, it holds with probability 1 that

$$\text{Dec}(sk, \text{Eval}(pk, C, \text{Enc}(pk, m_1), \dots, \text{Enc}(pk, m_n))) = C(m_1, \dots, m_n).$$

Definition 2.4. A leveled fully homomorphic encryption scheme FHE = (Gen, Enc, Dec, Eval) satisfies the same syntax, correctness, and security properties of a FHE scheme, except that

- Gen($1^\lambda, 1^d$) takes as additional input a circuit depth d .

- Homomorphic evaluation correctness is only guaranteed to hold for circuits of depth at most d .
- Ciphertexts output by $\text{Enc}(\text{pk}, m)$ and $\text{Eval}(\text{pk}, f, \text{ct})$ have size that are polynomial in λ (and the output length of f), independent of d .
- The decryption algorithm $\text{Dec}(\text{sk}, \text{ct})$ has a fixed $\text{poly}(\lambda)$ depth (independent of d).

Leveled fully homomorphic encryption schemes are known to exist from the learning with errors (LWE) assumption [BV11, BGV12, Bra12, GSW13, BV14]. Fully homomorphic encryption schemes are known to exist using Gentry’s bootstrapping technique [Gen09], which requires making a circular security assumption on an LWE-based encryption scheme. In this work, we consider the following variant of circular security.

Definition 2.5. A public key encryption scheme PKE is said to be circular secure if $\{(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) : (\text{pk}, \text{Enc}(\text{pk}, 0^{|\text{sk}|}))\} \approx_c \{(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) : (\text{pk}, \text{Enc}(\text{pk}, \text{sk}))\}$.

2.3 Non-Interactive Zero Knowledge Arguments (and Proofs)

The following preliminaries are taken (with edits) from [CCH⁺18].

Definition 2.6. A non-interactive zero knowledge (NIZK) argument system Π for an NP relation R consists of three ppt algorithms (Setup, P, V) with the following syntax.

- $\text{Setup}(1^n, 1^\lambda)$ takes as input a statement length n and a security parameter λ . It outputs a common reference string crs .
- $P(\text{crs}, x, w)$ takes as input the common reference string, as well as x and w such that $(x, w) \in R$. It outputs a proof π .
- $V(\text{crs}, x, \pi)$ takes as input the common reference string, a statement x , and a proof π . It outputs a bit b . If $b = 1$, we say that V accepts, and otherwise we say that V rejects.

The proof system Π must satisfy the following requirements for every polynomial function $n = n(\lambda)$. Recall that $\mathcal{L}(R)$ denotes the language $\{x : \exists w \text{ s.t. } (x, w) \in R\}$ and R_n denotes the set $R \cap (\{0, 1\}^n \times \{0, 1\}^*)$.

- **Completeness.** For every $(x, w) \in R$, it holds with probability 1 that $V(\text{crs}, x, \pi) = 1$ in the probability space defined by sampling $\text{crs} \leftarrow \text{Setup}(1^{|x|}, 1^\lambda)$ and $\pi \leftarrow P(\text{crs}, x, w)$.
- **Soundness.** For every $\{x_n \in \{0, 1\}^n \setminus \mathcal{L}(R)\}$ and every polynomial size $P^* = \{P_\lambda^*\}$, there is a negligible function ν such that

$$\Pr_{\substack{\text{crs} \leftarrow \text{Setup}(1^n, 1^\lambda) \\ \pi \leftarrow P_\lambda^*(\text{crs})}} [V(\text{crs}, x_n, \pi) = 1] \leq \nu(\lambda).$$

- **Zero Knowledge.** There is a ppt simulator Sim such that for every ensemble $\{(x_n, w_n) \in R_n\}$, the distribution ensembles

$$\left\{ (\text{crs}_\lambda, P(\text{crs}_\lambda, x_n, w_n)) \right\}_\lambda$$

and

$$\{\text{Sim}(x_n, 1^\lambda)\}_\lambda$$

are computationally indistinguishable in the probability space defined by sampling $\text{crs}_\lambda \leftarrow \text{Setup}(1^n, 1^\lambda)$ (and evaluating P and Sim with independent and uniform randomness).

If the distributions are statistically indistinguishable, then Π is said to be statistically zero knowledge.

A NIZK argument system can also satisfy various stronger properties. We list some important variants below.

- **“Common Random String”**: A NIZK argument system in the common *random* string model is a NIZK argument system Π such that $\text{Setup}(1^n, 1^\lambda)$ simply samples and outputs a uniformly random string.
- **Adaptive Soundness**: Π is adaptively sound if for every polynomial size algorithm $P^* = \{P_\lambda^*\}$, there is a negligible function ν such that for all λ ,

$$\Pr_{\substack{\text{crs} \leftarrow \text{Setup}(1^n, 1^\lambda) \\ (x, \pi) := P_\lambda^*(\text{crs})}} [x \notin \mathcal{L}(R) \wedge V(\text{crs}, x, \pi) = 1] \leq \nu(\lambda).$$

- **Statistical Soundness**: Π is statistically sound if there is a negligible function ν such that for all λ ,

$$\Pr_{\text{crs} \leftarrow \text{Setup}(1^n, 1^\lambda)} [\exists(x, \pi) \text{ such that } x \notin \mathcal{L}(R) \wedge V(\text{crs}, x, \pi) = 1] \leq \nu(\lambda).$$

A NIZK argument system satisfying statistical soundness is called a **NIZK proof system**.

- **Multi-Theorem Zero Knowledge**: Π is multi-theorem zero knowledge if for every polynomial function $p(\lambda)$, there is a p.p.t. simulator Sim such that for every ensemble $\{(x_i^{(n(\lambda))}, w_i^{(n(\lambda))})_{i=1}^{p(\lambda)}\}$, the distribution ensembles

$$\left\{ (\text{crs}_\lambda, P(\text{crs}_\lambda, x_i^{(n)}, w_i^{(n)}))_{i=1}^{p(\lambda)} \right\}_\lambda$$

and

$$\{\text{Sim}(x_1, \dots, x_{p(\lambda)})\}_\lambda$$

are computationally indistinguishable. [FLS99] showed a generic transformation from a NIZK proof or argument system to one satisfying multi-theorem zero knowledge. This transformation preserves computational zero knowledge in the common random string model and statistical zero knowledge in the common reference string model.

- **Adaptive Zero Knowledge**: Π is adaptive zero knowledge if for every p.p.t. verifier V^* , there is a p.p.t. simulator Sim such that the following distribution ensembles are computationally indistinguishable:

$$\left\{ \text{crs} \leftarrow \text{Setup}(1^n, 1^\lambda), (x, w, \text{aux}) \leftarrow V^*(\text{crs}) : (\text{crs}, P(\text{crs}, x, w), \text{aux}) \right\}$$

and

$$\{\text{Sim}(1^n, 1^\lambda)\}.$$

This can (analogously to above) be extended to a definition of **adaptive multi-theorem zero knowledge**, which can be obtained generically from adaptive zero-knowledge by [FLS99].

3 Somewhere Statistically Correlation Intractable Hash Families

In this section, we recall the notion of correlation intractability [CGH04], which is a particular security property associated to a hash family \mathcal{H} . We then introduce a new strengthening of this definition, which we call “somewhere statistical correlation intractability” by analogy to the “somewhere statistically binding” hash functions of [HW15].

We also define new classes of relations – “efficiently searchable” and “efficiently enumerable” relations – for which we later (1) achieve correlation intractability (“somewhere statistical,” in the case of efficiently searchable relations), and (2) obtain applications of interest.

Definition 3.1. For a pair of efficiently computable functions $(n(\cdot), m(\cdot))$, a hash family with input length n and output length m is a collection $\mathcal{H} = \{h_\lambda : \{0, 1\}^{s(\lambda)} \times \{0, 1\}^{n(\lambda)} \rightarrow \{0, 1\}^{m(\lambda)}\}_{\lambda \in \mathbb{N}}$ of keyed hash functions, along with a pair of p.p.t. algorithms:

- $\mathcal{H}.\text{Gen}(1^\lambda)$ outputs a hash key $k \in \{0, 1\}^{s(\lambda)}$.
- $\mathcal{H}.\text{Hash}(k, x)$ computes the function $h_\lambda(k, x)$. We may use the notation $h(k, x)$ to denote hash evaluation when the hash family is clear from context.

We say that \mathcal{H} is public-coin⁷ if $\mathcal{H}.\text{Gen}$ outputs a uniformly random string $k \leftarrow \{0, 1\}^{s(\lambda)}$.

Definition 3.2 (Correlation Intractability). For a given relation ensemble $R = \{R_\lambda \subseteq \{0, 1\}^{n(\lambda)} \times \{0, 1\}^{m(\lambda)}\}$, a hash family $\{h_\lambda : \{0, 1\}^{s(\lambda)} \times \{0, 1\}^{n(\lambda)} \rightarrow \{0, 1\}^{m(\lambda)}\}$ is said to be R -correlation intractable if for every polynomial-size $\mathcal{A} = \{\mathcal{A}_\lambda\}$,

$$\Pr_{\substack{k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda) \\ x \leftarrow \mathcal{A}(k)}}} [(x, h(k, x)) \in R] = \text{negl}(\lambda)$$

If \mathcal{R} is a collection of relation ensembles, then \mathcal{H} is said to be uniformly \mathcal{R} -correlation intractable if for every polynomial-size \mathcal{A} , there exists a function $\nu(\lambda) = \text{negl}(\lambda)$ such that for every $R \in \mathcal{R}$,

$$\Pr_{\substack{k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda) \\ x \leftarrow \mathcal{A}(k)}}} [(x, h(k, x)) \in R] \leq \nu(\lambda).$$

As noted in [CGH04], a random oracle (typically thought of as an “ideal hash function” [BR93]) behaves like an R -correlation intractable for all sparse relations R .

Definition 3.3 (Sparsity). For any relation ensemble $R = \{R_\lambda \subseteq \{0, 1\}^{n(\lambda)} \times \{0, 1\}^{m(\lambda)}\}$, we say that R is $\rho(\cdot)$ -sparse if for $\lambda \in \mathbb{N}$ and any $x \in \{0, 1\}^n$,

$$\Pr_{y \leftarrow \{0, 1\}^m} [(x, y) \in R] \leq \rho(\lambda).$$

When ρ is a negligible function, we say that R is sparse.

We now introduce our new notion of “somewhere statistical correlation intractability.”

⁷Sometimes “public-coin” hash families are defined to be hash families whose security properties hold even when the adversary is given the random coins used to sample $k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda)$. For our purposes (e.g. ignoring compactness), this definition is equivalent to ours.

Definition 3.4 (Somewhere Statistical Correlation Intractability). *Given a collection \mathcal{R} of relation ensembles, we say that a hash family \mathcal{H} is somewhere statistically correlation intractable with respect to \mathcal{R} if there is an additional key generation algorithm StatGen with the following properties.*

- **Syntax:** $\text{StatGen}(1^\lambda, \text{aux}_\lambda)$ takes as input the security parameter λ as well as an auxiliary input aux_λ . It outputs a hash key k .
- **Security:** For any relation ensemble $R \in \mathcal{R}$, there exists an auxiliary input ensemble aux such that the following two properties hold.
 - **Key Indistinguishability:** An honestly generated key $k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda)$ is computationally indistinguishable from a fake key $k \leftarrow \mathcal{H}.\text{StatGen}(1^\lambda, \text{aux}_\lambda)$.
 - **Statistical Correlation Intractability:**

$$\Pr_{k \leftarrow \mathcal{H}.\text{StatGen}(1^\lambda, \text{aux}_\lambda)} \left[\exists x \in \{0, 1\}^{n(\lambda)} : (x, h(k, x)) \in R_\lambda \right] = \text{negl}(\lambda).$$

That is, with high probability over the choice of $k \leftarrow \text{StatGen}(1^\lambda, \text{aux}_\lambda)$, input-output pairs satisfying R_λ do not exist.

Remark 3.1. By a simple hybrid argument, somewhere statistical correlation intractability with respect to a family of relations \mathcal{R} implies (ordinary) correlation intractability for \mathcal{R} .

In this work, we focus further on achieving (somewhere statistical) correlation intractability for relations R with a *unique* output $y = f(x)$ associated to each input x , and such that $y = f(x)$ is an efficiently computable function of x .

Definition 3.5 (Unique Output Relation). *We say that a relation R is a unique output relation if for every input x , there exists at most one output y such that $(x, y) \in R$.*

Remark 3.2. When restricted to the case of unique output relations, obtaining correlation intractable hash functions for output length m becomes strictly harder as m becomes smaller (as a function of n).

Definition 3.6 (Efficiently Searchable Relation). *We say that a (necessarily unique-output) relation ensemble R is searchable in (non-uniform) time T if there exists a function $f = f_R : \{0, 1\}^* \rightarrow \{0, 1\}^*$ computable in (non-uniform) time T such that for any input x , if $(x, y) \in R$ then $y = f(x)$; that is, $f(x)$ is the unique y such that $(x, y) \in R$, provided that such a y exists. We say that R is efficiently searchable if it is searchable in time $\text{poly}(n)$.*

We further consider relations R that are not necessarily unique-output, but such that the set of all outputs y satisfying $(x, y) \in R$ is efficiently computable given x :

Definition 3.7 (Efficiently Enumerable Relation). *We say that a relation R is enumerable in (non-uniform) time T if there exists a (non-uniform) time T algorithm \mathcal{A} that on any input x outputs the set S_x of all y such that $(x, y) \in R$, provided that S_x is nonempty. We say that R is efficiently enumerable if it is enumerable in time $\text{poly}(n)$.*

Remark 3.3. If a hash family \mathcal{H} is correlation intractable for all relations *searchable* in time T , then it is also correlation intractable for all relations *enumerable* in time $T - O(\log(T))$.

4 Correlation Intractability via Fully Homomorphic Encryption

In this section, we describe a new candidate correlation intractable hash family \mathcal{H} that can be based on any fully homomorphic encryption scheme. We then prove that \mathcal{H} satisfies various notions of correlation intractability under different assumptions. Namely, we show:

- One variant of our hash family is (somewhere statistically) correlation intractable for efficiently searchable relations assuming that the FHE scheme is circular secure.
- Another variant of our hash family is “universal” for correlation intractable hash families (in a specific sense defined in Section 4.2), assuming that the FHE scheme is semantically secure. This holds both for single-input and multi-input correlation intractability.

4.1 Correlation Intractability for Efficiently Searchable Relations

Construction 4.1. Let $\text{FHE} = (\text{Gen}, \text{Enc}, \text{Dec}, \text{Eval})$ be any circular secure fully homomorphic encryption scheme. We define the following hash family $\mathcal{H} = \mathcal{H}_{\text{FHE}}^{\text{circ}}$ associated to FHE along with some circuit size bound $L(\lambda)$, input length $n(\lambda)$, and constant $\epsilon > 0$:

- $\mathcal{H}.\text{Gen}(1^\lambda)$ calls $\text{Gen}(1^\lambda)$, obtaining a pair (pk, sk) . It then computes $\text{ct} = \text{Enc}(\text{pk}, 0^{|\text{sk}|+L})$ and outputs $k = (\text{pk}, \text{ct})$.
- $\mathcal{H}.\text{Hash}(k, x)$ interprets $k = (\text{pk}, \text{ct})$ and outputs

$$y = \text{Eval}(\text{pk}, \tilde{U}_x, \text{ct}),$$

where \tilde{U}_x denotes the universal circuit for evaluation of circuits of size $L(\lambda)$, single bit output, and partially hardcoded input $x||s$ with (non-hardcoded) $s \in \{0, 1\}^{|\text{sk}|}$; that is,

$$\tilde{U}_x(s, C) = C(x, s).$$

We note that the output length of this hash function is some fixed polynomial $\text{poly}(\lambda)$. By setting the security parameter λ appropriately (in relation to n), this can result in a hash function with arbitrary polynomial relationship between input and output length.

Remark 4.1. We use the modified universal circuit \tilde{U} as above in order to rely on the *circular-security* (rather than more general KDM-security) of FHE in our security proof.

Theorem 4.2. Suppose that FHE is a circular secure fully-homomorphic encryption scheme, let $n(\lambda) = \lambda^{\Theta(1)}$, and let $T = \text{poly}(n, \lambda)$ be given. Then, $\mathcal{H} = \mathcal{H}^{\text{circ}}$ with input length $n(\lambda)$ and size parameter $L = T + \text{poly}(\lambda)$ is correlation intractable for all relations R that are searchable in time T (and appropriate input/output lengths).

In fact, \mathcal{H} is somewhere statistically correlation intractable for this class of relations, such that the function $\text{StatGen}(1^\lambda, \text{aux})$ can use any circuit computing the search function f_R as its auxiliary input.

Proof. Let FHE and R be fixed; recall by the definition of T -searchability, there exists a function $f = f_R : \{0, 1\}^* \rightarrow \{0, 1\}^*$ computable in (non-uniform) time T such that if $(x, y) \in R$ then $f(x) = y$.

To show that \mathcal{H} is somewhere statistically correlation intractable for all T -searchable relations R , we define the auxiliary algorithm $\mathcal{H}.\text{StatGen}(1^\lambda, \text{aux})$, which operates as follows:

- Interpret aux as a circuit C of size T .
- Call $\text{Gen}(1^\lambda)$, obtaining a pair (pk, sk) . Then, compute $\text{ct} = \text{Enc}(\text{pk}, \text{sk} || C')$, where

$$C'(x, \text{sk}) = 1 \oplus \text{Dec}(\text{sk}, C(x)).$$

- Output $k = (\text{pk}, \text{ct})$.

It now suffices to prove that our augmented hash family \mathcal{H} satisfies key indistinguishability and statistical correlation intractability (for keys generated by StatGen).

- **Key indistinguishability** follows immediately from the circular security of FHE, as the only difference between $\mathcal{H}.\text{Gen}$ and $\mathcal{H}.\text{StatGen}$ is that the former samples $\text{ct} = \text{Enc}(\text{pk}, 0^{|\text{sk}|+L})$ while the latter samples $\text{ct} = \text{Enc}(\text{pk}, \text{sk} || C')$.
- **Statistical correlation intractability** holds by the following argument. Let R be any T -searchable relation, let C be any circuit computing the search function f_R , and let $k = (\text{pk}, \text{ct}) \leftarrow \text{StatGen}(1^\lambda, C)$. Then, for any $x \in \{0, 1\}^n$ and $y = h(k, x)$, we see that by the correctness of FHE-evaluation,

$$\text{Dec}(\text{sk}, y) = \text{Dec}(\text{sk}, \text{Eval}(\tilde{U}_x, \text{ct})) = 1 \oplus \text{Dec}(\text{sk}, C(x)).$$

Therefore, if $(x, y) \in R$, then $y = f_R(x) = C(x)$ and we obtain the equation $\text{Dec}(\text{sk}, y) = 1 \oplus \text{Dec}(\text{sk}, y)$, a contradiction. Thus, input-output pairs satisfying R (unconditionally) do not exist.

This completes the proof of Theorem 4.2. □

Theorem 4.2 gives us a kind of correlation intractability that suffices to prove the soundness of the Fiat-Shamir transform for certain proof systems, as seen in Section 5 and Section 6. Since we hope to use our correlation intractable hash families to build NIZK arguments (which in particular must also be *zero knowledge*), we would like to have correlation intractable hash families satisfying a weak notion of *programmability*.

Definition 4.3. *We say that a hash family \mathcal{H} is 1-universal if for any λ , input $x \in \{0, 1\}^{n(\lambda)}$, and output $y \in \{0, 1\}^{m(\lambda)}$, we have that*

$$\Pr_{k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda)} [h(k, x) = y] = 2^{-m}.$$

We say that a hash family \mathcal{H} is programmable if it is 1-universal, and if there exists an efficient sampling algorithm $\text{Samp}(1^\lambda, x, y)$ that samples from the conditional distribution $k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda) \mid h(k, x) = y$.

We conclude our discussion of correlation intractability for efficiently searchable relations by describing a simple transformation showing that programmability holds without loss of generality.

Construction 4.4. *Let \mathcal{H} be any hash family. We define the programmable variant $\mathcal{H}' = \mathcal{H}^{\text{prog}}$ of \mathcal{H} as follows:*

- $\mathcal{H}'.\text{Gen}(1^\lambda)$ calls $k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda)$, samples a uniformly random $\alpha \leftarrow \{0,1\}^m$, and outputs (k, α) .
- $\mathcal{H}'.\text{Hash}((k, \alpha), x)$ outputs $\mathcal{H}.\text{Hash}(k, x) \oplus \alpha$.

We first remark that \mathcal{H}' is evidently programmable: 1-universality follows from the randomness of α , and the algorithm $\text{Samp}(1^\lambda, x, y)$ calls $k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda)$ and outputs $(k, \mathcal{H}.\text{Hash}(k, x) \oplus y)$.

Moreover, we note that if \mathcal{H}' directly inherits correlation intractability properties from \mathcal{H} .

Remark 4.2. For any relation class \mathcal{R} , if \mathcal{H} is (somewhere statistically) correlation intractable for the class of relations

$$\{R_\alpha^* = \{(x, y) : (x, y \oplus \alpha) \in R\}\}_{R \in \mathcal{R}},$$

then \mathcal{H}' is (somewhere statistically) correlation intractable for \mathcal{R} .

In particular, we obtain the following corollary.

Corollary 4.5. Fix functions $m(\lambda) = n(\lambda)^{\Theta(1)}$. If circular secure FHE exists, then for every polynomial function T , there exists a programmable hash family \mathcal{H} that is somewhere statistically correlation intractable for the class of relation ensembles $R = \{R_\lambda \subset \{0,1\}^{n(\lambda)} \times \{0,1\}^{m(\lambda)}\}$ that are searchable in (non-uniform) time T .

4.2 Universal Correlation Intractability from LWE

We now show that a simplified version of Construction 4.1 yields a hash family satisfying interesting notions of *universality* for correlation intractable hash families. We obtain results based on the LWE assumption, either with polynomial or subexponential (that is, 2^{λ^δ} -) security.

Construction 4.6. Let $\text{FHE} = (\text{Gen}, \text{Enc}, \text{Dec}, \text{Eval})$ be any (leveled) fully homomorphic encryption scheme. We define the following hash family $\mathcal{H} = \mathcal{H}_{\text{FHE}}^{\text{univ}}$ associated to FHE along with some circuit size bound $L(\lambda)$, input length $n(\lambda)$, plaintext length $m'(\lambda)$, and constant $\epsilon > 0$:

- $\mathcal{H}.\text{Gen}(1^\lambda)$ calls $\text{Gen}(1^{\lambda^\epsilon})$. It then computes $\text{ct} = \text{Enc}(\text{pk}, 0^L)$ and outputs $k = (\text{pk}, \text{ct})$.
- $\mathcal{H}.\text{Hash}(k, x)$ interprets $k = (\text{pk}, \text{ct})$ and outputs

$$y = \text{Eval}(\text{pk}, U_x, \text{ct}),$$

where U_x denotes the universal circuit for evaluation of any given circuit C , whose size is $L(\lambda)$ and whose output length is $m'(\lambda)$, on input x . That is, $U_x(C) = C(x)$.

We note that the output length of this hash function is not m' but $m(\lambda) = |\text{Enc}(\text{pk}, 0^{m'(\lambda)})| = m' \cdot \text{poly}(\lambda^\epsilon)$.

Theorem 4.7. Suppose that FHE is a (leveled) fully-homomorphic encryption scheme, and let $R = \{R_\lambda \subset \{0,1\}^{n(\lambda)} \times \{0,1\}^{m(\lambda)=m'(\lambda)\text{poly}(\lambda^\epsilon)}\}$ be a relation ensemble that is decidable in polynomial time. Then, if there exists a hash family \mathcal{H}_R , computable by circuits of size at most L , that is correlation intractable for all relations on $\{0,1\}^n \times \{0,1\}^{m'}$ of the form

$$R_{\text{sk}}^* = \{(x, z) : \exists y \text{ such that } (x, y) \in R \text{ and } z = \text{Dec}(\text{sk}, y)\},$$

then $\mathcal{H} = \mathcal{H}^{\text{univ}}$ with parameters (L, n, m', ϵ) is correlation intractable for R . Moreover, if FHE (with security parameter λ^ϵ) is secure against 2^{λ^δ} -time adversaries, then the same statement holds for all relations R decidable in (non-uniform) time $2^{\lambda^\delta} - \lambda^{\omega(1)}$.

Proof. Let FHE and R be fixed. Suppose that some p.p.t. adversary \mathcal{A} , given $k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda)$, outputs some $x \in \{0, 1\}^n$ such that $(x, h(k, x)) \in R$ with non-negligible probability. Let \mathcal{H}_R denote the correlation intractable hash family hypothesized to exist in the theorem statement.

We first claim that the adversary \mathcal{A} still succeeds with non-negligible probability when given a key $\tilde{k} = (\text{pk}, \text{Enc}(\text{pk}, H_{R, k'}))$, where $H_{R, k'}$ is a circuit computing H_R with randomly sampled key $k' \leftarrow \mathcal{H}_R.\text{Gen}(1^\lambda)$. This follows immediately from the semantic security of FHE, as \mathcal{A} 's win condition is decidable in the time required to decide R .

We now describe a p.p.t. adversary \mathcal{A}' that breaks the correlation intractability of \mathcal{H}_R :

1. \mathcal{A}' samples FHE parameters $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^{\lambda^\epsilon})$ and declares the relation R_{sk}^* as its challenge.
2. \mathcal{A}' is given a hash key $k' \leftarrow \mathcal{H}_R.\text{Gen}(1^\lambda)$.
3. \mathcal{A}' computes $\text{ct} \leftarrow \text{Enc}(\text{pk}, H_{R, k'})$, and runs $\mathcal{A}'(\text{pk}|\text{ct})$.
4. \mathcal{A}' obtains an input $x \in \{0, 1\}^n$ and returns x .

By construction, whenever $\mathcal{A}(\text{pk}|\text{ct})$ breaks the R -correlation intractability of \mathcal{H} , we have that \mathcal{A}' breaks the R_{sk}^* -correlation intractability of \mathcal{H}_R . To see this, note that if $y = h(\text{pk}|\text{ct}, x)$ in the above experiment, then

$$\text{Dec}(\text{sk}, y) = \text{Dec}(\text{sk}, \text{Eval}(\text{pk}, U_x, \text{ct})) = h_R(k', x).$$

Thus, if $(x, y) \in R$, then $(x, h_R(k', x)) \in R_{\text{sk}}^*$. This completes the proof of Theorem 4.7. \square

As an immediate corollary of Theorem 4.7, we conclude that $\mathcal{H}^{\text{univ}}$ is weakly *universal* for correlation intractable hash families (for sufficiently sparse relations), in the following sense.

Corollary 4.8. *Let $\gamma < 1$ and $\delta < 1$ be arbitrary, and let $m'(\lambda) = \lambda^{\Omega(1)}$ grow at least polynomially with λ . Set*

$$\epsilon = \Omega\left(\frac{\delta}{1 - \delta} \log_\lambda(m')\right)$$

and

$$\beta = \frac{\gamma}{1 - \delta}.$$

Finally, suppose that there exists a hash family (computable by a size L circuit) \mathcal{H}_β that is correlation intractable for all relations on $\{0, 1\}^n \times \{0, 1\}^{m'}$ with sparsity $\frac{2^{m'\beta}}{2^{m'}}$. Then, assuming the security of FHE, $\mathcal{H}^{\text{univ}}$ implemented with parameters (L, n, m', ϵ) is correlation intractable for all efficiently decidable relations on $\{0, 1\}^n \times \{0, 1\}^m$ of sparsity $\frac{2^{m'\beta}}{2^m} = \frac{2^{m\gamma}}{2^m}$. The same is true for all (sufficiently sparse) subexponentially decidable relations on $\{0, 1\}^n \times \{0, 1\}^m$ if FHE is assumed to be subexponentially secure.

Finally, we note that Theorem 4.2 – our construction of CI for efficiently searchable relations from circular secure FHE – can be thought of as a twist on the proof of Theorem 4.7. The difference between the two proofs is that in Theorem 4.2, we use the circular security of FHE to reduce from the security of $\mathcal{H}^{\text{circ}}$ for f to the existence of a hash family $\mathcal{H}_{f, \text{sk}}$ that is correlation intractable for *single* relation that depends on f and sk , the FHE secret key. Moreover, again due to the circular security assumption, we can use sk in the construction of $\mathcal{H}_{f, \text{sk}}$. This allows for an unconditional construction of the “inner hash function” that we have to assume exists in Theorem 4.7.

4.3 Multi-Input Correlation Intractability

Our universality results even extend to the notion of *multi-input* correlation intractability, about which very little is known.⁸

Definition 4.9 (Multi-Input Correlation Intractability). *Let ℓ (possibly depending on λ) denote an arity. For a given relation ensemble $R = \{R_\lambda \subseteq (\{0, 1\}^{n(\lambda)})^{\ell(\lambda)} \times (\{0, 1\}^{m(\lambda)})^{\ell(\lambda)}\}$, a hash family $\{h_\lambda : \{0, 1\}^{s(\lambda)} \times \{0, 1\}^{n(\lambda)} \rightarrow \{0, 1\}^{m(\lambda)}\}$ is said to be R -correlation intractable if for every polynomial-size $\mathcal{A} = \{\mathcal{A}_\lambda\}$,*

$$\Pr_{\substack{k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda) \\ \mathbf{x} = (x_1, \dots, x_\ell) \leftarrow \mathcal{A}(k)}}} \left[(x_1, \dots, x_\ell, h(k, x_1), \dots, h(k, x_\ell)) \in R \right] = \text{negl}(\lambda)$$

If \mathcal{R} is a collection of relation ensembles, then \mathcal{H} is said to be uniformly \mathcal{R} -correlation intractable if for every polynomial-size \mathcal{A} , there exists a function $\nu(\lambda) = \text{negl}(\lambda)$ such that for every $R \in \mathcal{R}$,

$$\Pr_{\substack{k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda) \\ \mathbf{x} = (x_1, \dots, x_\ell) \leftarrow \mathcal{A}(k)}}} (x_1, \dots, x_\ell, h(k, x_1), \dots, h(k, x_\ell)) \leq \nu(\lambda).$$

Analogously to Theorem 4.7, we observe that Construction 4.6 satisfies interesting notions of universality for *multi-input* correlation intractable hash families.

Theorem 4.10. *Suppose that FHE is a (leveled) fully-homomorphic encryption scheme, and let $R = \{R_\lambda \subset (\{0, 1\}^{n(\lambda)})^{\ell(\lambda)} \times (\{0, 1\}^{m(\lambda)=m'(\lambda)\text{poly}(\lambda^\epsilon)})^{\ell(\lambda)}\}$ be a relation ensemble that is decidable in polynomial time. Then, if there exists a hash family \mathcal{H}_R , computable by circuits of size at most L , that is correlation intractable for all relations on $(\{0, 1\}^n)^\ell \times (\{0, 1\}^{m'})^\ell$ of the form*

$$R_{\text{sk}}^* = \{(\mathbf{x}, \mathbf{z}) : \exists \mathbf{y} \text{ such that } (\mathbf{x}, \mathbf{y}) \in R \text{ and } \mathbf{z} = \text{Dec}(\text{sk}, \mathbf{y})\},$$

then $\mathcal{H} = \mathcal{H}^{\text{univ}}$ with parameters (L, n, m', ϵ) is correlation intractable for R . Moreover, if FHE (with security parameter λ^ϵ) is secure against 2^{λ^δ} -time adversaries, then the same statement holds for all relations R decidable in (non-uniform) time $2^{\lambda^\delta} - \lambda^{\omega(1)}$.

Proof. This is largely identical to the proof of Theorem 4.7, but we include a proof for completeness.

Let FHE and R be fixed. Suppose that some p.p.t. adversary \mathcal{A} , given $k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda)$, outputs some $\mathbf{x} \in (\{0, 1\}^n)^\ell$ such that $(\mathbf{x}, h(k, x_1), \dots, h(k, x_\ell)) \in R$ with non-negligible probability. Let \mathcal{H}_R denote the correlation intractable hash family hypothesized to exist in the theorem statement.

We first claim that the adversary \mathcal{A} still succeeds with non-negligible probability when given a key $\tilde{k} = (\text{pk}, \text{Enc}(\text{pk}, H_{R, k'}))$, where $H_{R, k'}$ is a circuit computing H_R with randomly sampled key $k' \leftarrow \mathcal{H}_R.\text{Gen}(1^\lambda)$. This follows immediately from the semantic security of FHE, as \mathcal{A} 's win condition is decidable in the time required to decide R .

We now describe a p.p.t. adversary \mathcal{A}' that breaks the correlation intractability of \mathcal{H}_R :

1. \mathcal{A}' samples FHE parameters $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^{\lambda^\epsilon})$ and declares the relation R_{sk}^* as its challenge.
2. \mathcal{A}' is given a hash key $k' \leftarrow \mathcal{H}_R.\text{Gen}(1^\lambda)$.

⁸The only known instantiations of multi-input correlation intractable hash families focus on the special case where the relation depends only on the output [Zha16, HL18] or rely on indistinguishability obfuscation [HL18].

3. \mathcal{A}' computes $\text{ct} \leftarrow \text{Enc}(\text{pk}, H_{R,k'})$, and runs $\mathcal{A}'(\text{pk}|\text{ct})$.
4. \mathcal{A}' obtains an input $\mathbf{x} \in (\{0, 1\}^n)^\ell$ and returns \mathbf{x} .

By construction, whenever $\mathcal{A}(\text{pk}|\text{ct})$ breaks the R -correlation intractability of \mathcal{H} , we have that \mathcal{A}' breaks the R_{sk}^* -correlation intractability of \mathcal{H}_R . To see this, note that if $y_i = h(\text{pk}|\text{ct}, x_i)$ in the above experiment, then

$$\text{Dec}(\text{sk}, y_i) = \text{Dec}(\text{sk}, \text{Eval}(\text{pk}, U_{x_i}, \text{ct})) = h_R(k', x_i).$$

Thus, if $(\mathbf{x}, \mathbf{y}) \in R$, then $(\mathbf{x}, h_R(k', x_1), \dots, h_R(k', x_\ell)) \in R_{\text{sk}}^*$. This completes the proof of Theorem 4.10. \square

It follows, as a corollary of Theorem 4.10, that if there *exists* a hash family that is correlation intractable for all (sufficiently sparse) multi-input relations, then for an appropriate parameter setting, $\mathcal{H}^{\text{univ}}$ is correlation intractable for all (efficiently decidable, sufficiently sparse) multi-input relations.

Although little is known about the existence of *general* multi-input correlation intractable hash families, the security reduction in Theorem 4.10 maps output relations (i.e. relations $R(\mathbf{x}, \mathbf{y})$ that depend only on \mathbf{y}) to output relations, so we obtain a concrete hash function that combines a family of candidates from [HL18].⁹ As an example, we obtain the following corollary.

Corollary 4.11. *Assume the hardness of LWE. In addition, assume that there exists family of (symmetric, injective) k -one way product functions (OWPFs) with security $2^{-kn^\beta} \cdot \text{negl}(n)$ for some $\beta < 1$.¹⁰*

Then for arbitrary $\gamma < \beta$, the hash family $\mathcal{H}^{\text{univ}}$ (for appropriate parameter settings) using an LWE-based (leveled) FHE scheme is correlation intractable for efficiently decidable output relations of sparsity $2^{kn^\gamma - kn}$.

Note that we only need to assume that the [HL18] OWPFs *exist*; we do not need an explicit description of one in the construction of $\mathcal{H}^{\text{univ}}$. This is similar to the obfuscation-based result of [HL18], but Corollary 4.11 replaces indistinguishability obfuscation with LWE. However, our result only applies in the regime of fairly low ($2^{kn^\gamma - kn}$) sparsity.

5 Non-Interactive Zero Knowledge Arguments

In this section, we apply Theorem 4.2 to obtain Theorem 1.2, that is, NIZK arguments for NP assuming circular secure FHE. This closely follows the framework of [HL18, CCH⁺18] for obtaining NIZK arguments from weak forms of correlation intractability, but we apply the framework to the 3-message [FLS99] protocol for graph Hamiltonicity. This allows us to rely on correlation intractable hash functions for efficiently searchable relations (as constructed in Theorem 4.2) as opposed to efficiently samplable relations (as defined in [HL18, CCH⁺18]).

We augment our basic NIZK argument system in two different ways:

⁹The [Zha16] construction does not give a hash family that is correlation intractable for all output relations – only efficiently decidable relations – so we cannot use it as is.

¹⁰We refer the reader to [HL18] for a discussion of OWPFs.

- By using our *somewhere statistically correlation intractable* hash functions (Definition 3.4), we show that our NIZK argument system has a *statistically sound mode*, yielding NIZK *proofs* in the common reference string model.
- By using a lossy public key encryption scheme with uniformly lossy public keys (Definition 2.2), we show that our NIZK argument system has a *statistical zero knowledge mode* in which the CRS is uniformly random, yielding NISZK arguments in the common *random* string model.

We begin by recalling the 3-message [FLS99] protocol.

5.1 The [FLS99] Protocol

We construct NIZK arguments by applying the Fiat-Shamir transform to a variant of the 3-message [FLS99] proof system for graph Hamiltonicity. Recall that the Hamiltonicity language L_{Ham} consists of all graphs G with a Hamiltonian cycle, and the standard NP-relation for L_{Ham} uses a permutation σ as a witness exhibiting a Hamiltonian cycle $\sigma^{-1}(C_n)$ in G . We describe the 3-message protocol $\Pi = \Pi_{\text{FLS}}$ in Fig. 2, and we will make use of the following fact.




Figure 2: The Zero Knowledge Proof System Π_{FLS} for Graph Hamiltonicity.

Fact 5.1. *For any computationally hiding commitment scheme Com (potentially in the CRS model), Π is honest-verifier zero knowledge. If Com is a statistically hiding commitment scheme, then Π is honest-verifier statistical zero knowledge.*

Finally, Π is “honest-verifier adaptive zero knowledge”, meaning that Π remains honest-verifier zero knowledge when the adversary is allowed to choose (x, w) as an (arbitrary) efficient function of the transcript $(\text{crs}, \mathbf{a}, \mathbf{e})$ up to the second message.

We emphasize that our variant of the [FLS99] protocol explicitly allows for the commitment scheme Com to rely on a public commitment key pk .

5.2 Our NIZK Protocol

We start by defining three different modes for our protocol and use them to help prove security. We use the following tools in our construction.

- A hash family \mathcal{H} satisfying two properties:

- Correlation intractability for all (subexponentially sparse) relations that are (non-uniformly) searchable in a fixed polynomial time T .
- Programmability, as in Definition 4.3.

In Construction 5.3, we will assume that \mathcal{H} is somewhere statistically correlation intractable for the above class of relations, and that $\mathcal{H}.\text{StatGen}(1^\lambda, C)$ can use any circuit C computing the search function f_R as auxiliary input. In Construction 5.4, we will assume that \mathcal{H} has pseudorandom keys, and that the modified hash family \mathcal{H}' using a uniformly random key is also programmable.¹¹

- A public key encryption scheme $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$. In Construction 5.4, we will assume that PKE is lossy (Definition 2.2) with uniformly random lossy public keys.

Construction 5.2 (Basic NIZK). *Let $\Pi = \Pi_{\text{FLS}}$ be the [FLS99] protocol from Section 5.1, in which we instantiate the commitment scheme Com in the CRS model using PKE in the natural way: $\text{Com}(\text{pk}, b; \rho) = \text{Enc}(\text{pk}, b; \rho)$. We apply the Fiat-Shamir transform, using \mathcal{H} , to the protocol Π^λ ; that is, the protocol Π repeated λ times in parallel. We call the resulting protocol $\tilde{\Pi}$, which is formally defined as follows.*

- *Common reference string: a PKE -public key pk along with a hash key $k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda)$.*
- *Prover message: given an instance x , witness w , and common reference string $\text{crs} = (\text{pk}, k)$, the prover computes $\mathbf{a} \leftarrow \Pi^\lambda.P(\text{crs}, x, w)$, $\mathbf{e} = h(k, \mathbf{a})$, $\mathbf{z} = \Pi^\lambda.P(\text{crs}, x, w, \mathbf{a}, \mathbf{e})$, and outputs $(\mathbf{a}, \mathbf{e}, \mathbf{z})$.*
- *The verifier accepts a transcript $(\text{crs}, x, \mathbf{a}, \mathbf{e}, \mathbf{z})$ if $\mathbf{e} = h(k, \mathbf{a})$ and $\Pi^\lambda.V(\text{crs}, x, \mathbf{a}, \mathbf{e}, \mathbf{z}) = 1$.*

Our two modified constructions change only the common reference string distribution.

Construction 5.3 (Statistically Sound Mode). *Let $\Pi = \Pi_{\text{FLS}}$ be the [FLS99] protocol from Section 5.1, in which we instantiate the commitment scheme Com in the CRS model using PKE . The statistically sound mode of our protocol $\tilde{\Pi}_{\text{Sound}}$ is then defined as follows.*

- *Common reference string: a PKE -public key pk along with a **fake hash key** $k \leftarrow \mathcal{H}.\text{StatGen}(1^\lambda, C_{\text{sk}})$. Here, sk is the PKE -secret key associated to pk , and C_{sk} is a (poly-size) circuit computing the function $f_{\text{sk}}(\mathbf{a}) = \mathbf{e}$ such that for every $i \in [\lambda]$, $e_i = 0$ if and only if $\text{Dec}(\text{sk}, a_i)$ is a cycle.*
- *Prover message: given an instance x , witness w , and common reference string $\text{crs} = (\text{pk}, k)$, the prover computes $\mathbf{a} \leftarrow \Pi^\lambda.P(\text{crs}, x, w)$, $\mathbf{e} = h(k, \mathbf{a})$, $\mathbf{z} = \Pi^\lambda.P(\text{crs}, x, w, \mathbf{a}, \mathbf{e})$, and outputs $(\mathbf{a}, \mathbf{e}, \mathbf{z})$.*
- *The verifier accepts a transcript $(\text{crs}, x, \mathbf{a}, \mathbf{e}, \mathbf{z})$ if $\mathbf{e} = h(k, \mathbf{a})$ and $\Pi^\lambda.V(\text{crs}, x, \mathbf{a}, \mathbf{e}, \mathbf{z}) = 1$.*

Construction 5.4 (Statistical Zero Knowledge Mode). *Let $\Pi = \Pi_{\text{FLS}}$ be the [FLS99] protocol from Section 5.1, in which we instantiate the commitment scheme Com in the CRS model using PKE . The statistical zero knowledge mode of our protocol $\tilde{\Pi}_{\text{ZK}}$ is then defined as follows.*

¹¹The generic transformation (Construction 4.4) from correlation intractable hash families to programmable correlation intractable hash families guarantees this property.

- *Common reference string: a (uniformly random) PKE-lossy public key $\text{pk} \leftarrow \text{FakeGen}(1^\lambda)$ along with a uniformly random hash key k .*
- *Prover message: given an instance x , witness w , and common reference string $\text{crs} = (\text{pk}, k)$, the prover computes $\mathbf{a} \leftarrow \Pi^\lambda.P(\text{crs}, x, w)$, $\mathbf{e} = h(k, \mathbf{a})$, $\mathbf{z} = \Pi^\lambda.P(\text{crs}, x, w, \mathbf{a}, \mathbf{e})$, and outputs $(\mathbf{a}, \mathbf{e}, \mathbf{z})$.*
- *The verifier accepts a transcript $(\text{crs}, x, \mathbf{a}, \mathbf{e}, \mathbf{z})$ if $\mathbf{e} = h(k, \mathbf{a})$ and $\Pi^\lambda.V(\text{crs}, x, \mathbf{a}, \mathbf{e}, \mathbf{z}) = 1$.*

Theorem 5.5. $\tilde{\Pi}$ is a NIZK argument system for NP in the common reference string model satisfying both adaptive soundness and adaptive zero knowledge. Moreover, $\tilde{\Pi}_{\text{Sound}}$ is a NIZK proof system for NP in the common reference string model satisfying adaptive zero knowledge, and $\tilde{\Pi}_{\text{ZK}}$ is an adaptively sound NISZK argument system for NP in the common random string model.

We now proceed to prove Theorem 5.5 by proving a sequence of lemmas about our construction.

Lemma 5.6. $\tilde{\Pi}$ is (adaptively) sound.

Proof. To see this, we argue that $\tilde{\Pi}$ is adaptively sound if \mathcal{H} is correlation intractable for all relations of the form

$$R_{\text{sk}} = \{(\mathbf{a}, \mathbf{e}) : \mathbf{e} = f_{\text{sk}}(\mathbf{a})\}.$$

This holds because if $\text{Dec}(\text{sk}, a_i)$ is not a cycle and $e_i = 0$, then there is no input graph x and third message z_i such that (x, a_i, e_i, z_i) is an accepting transcript in the original protocol Π . Similarly, if $\text{Dec}(\text{sk}, a_i)$ is a cycle and $e_i = 1$, then there is no input graph x that is not Hamiltonian (i.e. there is no false statement x) and third message z_i such that (x, a_i, e_i, z_i) is an accepting transcript in Π . These two facts make use of the perfect decryption correctness of PKE and the standard (adaptive) soundness analysis of Π .

From this analysis, we conclude that any adversary \mathcal{A} breaking the (adaptive) soundness of $\tilde{\Pi}$ breaks the correlation intractability of \mathcal{H} with respect to some sk (indeed, a random sk sampled according to PKE.Gen suffices).

Finally, we note that for every secret key sk , R_{sk} is an efficiently searchable relation: indeed, the function f_{sk} is efficiently computable given sk . Thus, since we assumed \mathcal{H} is correlation intractable for all efficiently searchable relations, we conclude that $\tilde{\Pi}$ is adaptively sound. \square

Lemma 5.7. $\tilde{\Pi}$ is (adaptive) zero knowledge.

Proof. The proof that $\tilde{\Pi}$ is zero knowledge is almost identical to the proof of zero knowledge in ([CCH⁺18] Theorem 7.7). For completeness, we describe a simulator for $\tilde{\Pi}$:

- **Input:** a graph x .
- Sample $(\text{pk}, \text{sk}) \leftarrow \text{PKE.Gen}(1^\lambda)$.
- Call the honest verifier simulator $\Pi^t.\text{HVSIM}(x, \text{pk})$ associated to the parallel repeated protocol Π^t , producing $(\mathbf{a}, \mathbf{e}, \mathbf{z})$.
- Sample a hash key \tilde{k} from the conditional distribution $k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda) \mid h(k, \mathbf{a}) = \mathbf{e}$.
- Output $(\text{pk}, \tilde{k}, \mathbf{a}, \mathbf{e}, \mathbf{z})$.

Zero knowledge then follows from Fact 5.1 and the programmability of \mathcal{H} (by a standard hybrid argument).

To see that $\tilde{\Pi}$ is *adaptive* zero knowledge, we use the fact that Π is honest-verifier adaptive zero knowledge (and use this two-part simulator in place of HVSIm above). We refer the reader to [CCRR18] (Proposition 7.6) for more details on obtaining adaptive zero knowledge using Fiat-Shamir. \square

This completes the proof that $\tilde{\Pi}$ is an adaptively sound NIZK argument system in the common reference string model.

Lemma 5.8. $\tilde{\Pi}_{\text{Sound}}$ is statistically sound.

Proof. Let $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$ and $k \leftarrow \mathcal{H}.\text{StatGen}(1^\lambda, C_{\text{sk}})$ as in Construction 5.3. Then, for any first message \mathbf{a} for the protocol Π^λ , if $\text{Dec}(\text{sk}, a_i)$ is not a cycle and $e_i = 0$, then there is no input graph G and third message z_i such that (G, a_i, e_i, z_i) is an accepting transcript in the original protocol Π . Similarly, if $\text{Dec}(\text{sk}, a_i)$ is a cycle and $e_i = 1$, then there is no input graph G that is not *Hamiltonian* (i.e. there is no false statement G) and third message z_i such that (G, a_i, e_i, z_i) is an accepting transcript in Π . These two facts make use of the perfect decryption correctness of PKE and the standard (adaptive) soundness analysis of Π .

This tells us that any accepting transcript $(\text{pk}, G, \mathbf{a}, \mathbf{e}, \mathbf{z})$ for $\tilde{\Pi}_{\text{Sound}}$ must satisfy $\mathbf{e} = f_{\text{sk}}(\mathbf{a})$. However, by the statistical correlation intractability of \mathcal{H} , we know that there does not exist an input \mathbf{a} such that $h(k, \mathbf{a}) = f_{\text{sk}}(\mathbf{a})$, so we conclude that $\tilde{\Pi}_{\text{Sound}}$ is statistically sound. \square

Lemma 5.9. For any p.p.t. verifier $V^*(\text{crs})$ outputting a triple (x, w, aux) , honestly generated transcripts in $\tilde{\Pi}$ and $\tilde{\Pi}_{\text{Sound}}$ are computationally indistinguishable (even given (x, w, aux)).

Proof. This follows immediately from the key indistinguishability of \mathcal{H} . \square

Combining Lemma 5.8, Lemma 5.9, and Lemma 5.7, we conclude that Construction 5.3 is a NIZK proof system for NP in the common reference string model satisfying adaptive zero knowledge.

Lemma 5.10. $\tilde{\Pi}_{\text{ZK}}$ is statistical zero knowledge.

Proof. We define a simulator for $\tilde{\Pi}_{\text{ZK}}$ as follows.

- **Input:** a graph x .
- Sample a uniformly random public key $\tilde{\text{pk}}$.
- Call the statistical honest verifier simulator $\Pi^\lambda.\text{HVSIm}(x, \text{pk})$ associated to the parallel repeated protocol Π^λ , producing $(\mathbf{a}, \mathbf{e}, \mathbf{z})$.
- Sample a hash key \tilde{k} from the conditional distribution $k \leftarrow \mathcal{H}'.\text{Gen}(1^\lambda) \mid h(k, \mathbf{a}) = \mathbf{e}$, where \mathcal{H}' is the modified hash family using uniformly random hash keys.
- Output $(\tilde{\text{pk}}, \tilde{k}, \mathbf{a}, \mathbf{e}, \mathbf{z})$.

Zero knowledge then follows directly from the lossiness of PKE (which implies that the resulting commitment scheme is statistically hiding), Fact 5.1, and the programmability of \mathcal{H}' (by a standard hybrid argument). \square

Lemma 5.11. *The common reference strings in $\tilde{\Pi}$ and $\tilde{\Pi}_{\text{ZK}}$ are computationally indistinguishable.*

Proof. This follows immediately from the key indistinguishability of PKE (between real and lossy keys) and the pseudorandomness of the \mathcal{H} -keys. \square

Combining Lemma 5.10, Lemma 5.11 and Lemma 5.6, we conclude that $\tilde{\Pi}_{\text{ZK}}$ is an adaptively sound NISZK argument system for NP in the common random string model. This completes the proof of Theorem 5.5.

5.3 Obtaining Theorem 1.3, Theorem 1.4, and LWE-based Instantiation

Recall the statements of Theorem 1.3 and Theorem 1.4, our main results on obtaining NIZK arguments.

Theorem 5.12. *Suppose that circular-secure fully homomorphic encryption exists. Then, there exist NIZK proofs for NP in the common reference string model.*

Theorem 5.13. *Suppose that there exists a circular-secure fully homomorphic encryption scheme with pseudorandom ciphertexts and public keys. Furthermore, suppose that there exists a lossy public key encryption scheme [KN08, PVW08, BHY09] with uniformly random lossy public keys. Then, there exist adaptively sound NISZK arguments for NP in the common random string model.*

We obtain these results by a direct combination of Theorem 4.2 and Theorem 5.5. Theorem 5.5 states that (1) the desired NIZK proofs follow from the existence of somewhere statistically correlation intractable hash functions for (subexponentially sparse) efficiently searchable relations, and (2) under the lossy PKE assumption, the desired NIZK arguments follow from the existence of correlation intractable hash functions for (subexponentially sparse) efficiently searchable relations (with pseudorandom hash keys). Theorem 4.2 states that assuming circular secure FHE (with pseudorandom ciphertexts and public keys), such hash families exist.

Moreover, we note that both of the generic primitives in Theorem 1.4 can be instantiated from (circular secure) LWE. Namely, under plain LWE, Regev encryption [Reg09] is a lossy PKE scheme and has uniformly random lossy public keys, and under various circular security assumptions on LWE [BV11, BGV12, Bra12, GSW13, BV14], there exist circular secure FHE schemes.

6 Fiat-Shamir for (Instance-Dependent) Trapdoor Σ -protocols

In this section, we formalize our notion of an “instance-dependent trapdoor Σ -protocol” and prove that our hash family from Theorem 4.2 suffices to instantiate the Fiat-Shamir heuristic for such protocols. Examples of instance-dependent trapdoor Σ -protocols include variants of the [Blu86] and [FLS99] protocols for graph Hamiltonicity (in which the commitment scheme is instantiated using public-key encryption as in Section 5) as well as the *unmodified* [GMR89] protocol for quadratic residuosity. By the connection between Fiat-Shamir and (malicious verifier) zero knowledge [DNRS99], we conclude that these protocols cannot be malicious verifier zero knowledge, assuming the existence of circular-secure FHE and the hardness of deciding the underlying languages. This partially resolves open questions due to [DNRS99, BLV03].

6.1 Instance-Dependent Trapdoor Σ -Protocols

We provide the following definition of a Σ -protocol, which suffices for our purposes. We do not require any extractability (“proof of knowledge”) property.

Definition 6.1 (Σ -Protocol). *We say that a three-message proof system $\Pi = (\text{Gen}, P, V)$ in the common reference string model is a Σ -protocol if for every common reference string crs , every instance $x \notin L$, and every first message \mathbf{a} , there is at most one challenge $\mathbf{e} := f(\text{crs}, x, \mathbf{a})$ such that $(\text{crs}, x, \mathbf{a}, \mathbf{e}, \mathbf{z})$ is an accepting transcript for any choice of third message \mathbf{z} .*

We informally call f the “bad-challenge function” associated to Π , and note that f may not be efficiently computable.

We now define an instance-dependent trapdoor Σ -protocol to be, roughly speaking, a Σ -protocol that has a trapdoor making the bad-challenge function f efficiently computable.

Definition 6.2 (Instance-Dependent Trapdoor Σ -Protocol). *We say that a Σ -protocol $\Pi = (\text{Gen}, P, V)$ with bad-challenge function f is an instance-dependent trapdoor Σ -protocol if there are p.p.t. algorithms $\text{TrapGen}, \text{BadChallenge}$ with the following syntax.*

- $\text{TrapGen}(1^\lambda, x, \text{aux})$ takes as input the security parameter, an instance x , and an auxiliary input aux . It outputs a common reference string crs along with a trapdoor τ .
- $\text{BadChallenge}(\tau, \text{crs}, x, \mathbf{a})$ takes as input a trapdoor τ , common reference string crs , instance x , and first message \mathbf{a} . It outputs a challenge \mathbf{e} .

We additionally require the following properties.

- **CRS Indistinguishability:** *For any (x, aux) , an honestly generated common reference string crs is computationally indistinguishable from a common reference string output by $\text{TrapGen}(1^\lambda, x, \text{aux})$.*
- **Correctness:** *for every instance $x \notin L$, there exists an auxiliary input aux such that for all $(\text{crs}, \tau) \leftarrow \text{TrapGen}(1^\lambda, x, \text{aux})$, we have that $\text{BadChallenge}(\tau, \text{crs}, x, \mathbf{a}) = f(\text{crs}, x, \mathbf{a})$.*

Given this definition, we can now state our result on Fiat-Shamir for instance-dependent trapdoor Σ -protocols.

Theorem 6.3. *Suppose that \mathcal{H} is a hash family that is correlation-intractable for all subexponentially sparse relations that are searchable in time T . Moreover, suppose that $\Pi = (\text{Gen}, P, V, \text{TrapGen}, \text{BadChallenge})$ is an instance-dependent trapdoor Σ -protocol with $2^{-\lambda^\epsilon}$ soundness for some $\epsilon > 0$, such that $\text{BadChallenge}(\tau, \text{crs}, x, \mathbf{a})$ is computable in time T . Then, \mathcal{H} soundly instantiates the Fiat-Shamir heuristic for Π .*

Proof. Let $\tilde{\Pi}$ denote the one-message protocol resulting from applying the Fiat-Shamir transform, using \mathcal{H} , to Π . Explicitly, $\tilde{\Pi}$ is defined as follows.

- The **common reference string** consists of a common reference string crs_Π associated to Π , along with a hash key $k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda)$.
- **Prover message:** a triple $(\mathbf{a}, \mathbf{e}, \mathbf{z})$, where \mathbf{a} is computed by $\Pi.P(\text{crs}_\Pi, x, w)$, $\mathbf{e} = h(k, \mathbf{a})$, and \mathbf{z} is computed by $\Pi.P(\text{crs}_\Pi, x, w, \mathbf{a}, \mathbf{e})$.

- The verifier accepts a transcript $(\text{crs}_\Pi, k, x, \mathbf{a}, \mathbf{e}, \mathbf{z})$ if $\Pi.V(\text{crs}_\Pi, \mathbf{a}, \mathbf{e}, \mathbf{z}) = 1$ and $\mathbf{e} = h(k, \mathbf{a})$.

By construction, and by the definition of a Σ -protocol, we know that for every $x \notin L$ and every crs_Π , an accepting $\tilde{\Pi}$ -transcript $(\text{crs}_\Pi, k, x, \mathbf{a}, \mathbf{e}, \mathbf{z})$ must satisfy the condition that $h(k, \mathbf{a}) = \mathbf{e} = f(\text{crs}_\Pi, x, \mathbf{a})$.

Suppose that some efficient prover P^* , given $x \notin L$ and a random $\text{crs} = (\text{crs}_\Pi, k)$, could find $(\mathbf{a}, \mathbf{e}, \mathbf{z})$ making the transcript $(\text{crs}_\Pi, k, x, \mathbf{a}, \mathbf{e}, \mathbf{z})$ accepting with non-negligible probability. Then, by CRS indistinguishability, the same would be true for crs_Π sampled by the algorithm $\text{TrapGen}(1^\lambda, x, \text{aux})$ for an auxiliary input aux satisfying the correctness property of Definition 6.2. In other words, for $(\text{crs}_\Pi, \tau_\Pi) \leftarrow \text{TrapGen}(1^\lambda, x, \text{aux})$ and $\text{crs} = (\text{crs}_\Pi, k)$, $P^*(x, \text{crs})$ would output (with non-negligible probability) some \mathbf{a} such that $h(k, \mathbf{a}) = f(\text{crs}_\Pi, x, \mathbf{a}) = \text{BadChallenge}(\tau_\Pi, \text{crs}_\Pi, x, \mathbf{a})$.

This directly contradicts the correlation intractability of \mathcal{H} for the relation $R_{\tau_\Pi, \text{crs}_\Pi, x} = \{(\mathbf{a}, \mathbf{e}) : \mathbf{e} = \text{BadChallenge}(\tau_\Pi, \text{crs}_\Pi, x, \mathbf{a})\}$. In more detail, a correlation-intractability adversary \mathcal{A} could break the correlation intractability of \mathcal{H} by sampling $(\text{crs}_\Pi, \tau_\Pi)$ itself, declaring the relation $R_{\tau_\Pi, \text{crs}_\Pi, x}$ to be broken, and then running $P^*(x, \text{crs})$ after being given $k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda)$. Since Π originally had $2^{-\lambda^\epsilon}$ soundness, the relation $R_{\tau_\Pi, \text{crs}_\Pi, x}$ indeed has subexponential sparsity, so this contradicts our assumption on \mathcal{H} . Thus, we conclude that \mathcal{H} soundly instantiates the Fiat-Shamir heuristic for Π , as desired. \square

6.2 Examples and Implications

It is easy to see that the variant of the [FLS99] Hamiltonicity protocol described in Section 5 satisfies Definition 6.2; the trapdoor generation algorithm ignores the instance x (and auxiliary input aux) and simply samples $(\text{pk}, \text{sk}) \leftarrow \text{PKE.Gen}(1^\lambda)$ and outputs pk as the common reference string and sk as the trapdoor. As already described, the bad-challenge function associated to Π_{FLS} is indeed efficiently computable given sk . Similarly, variants of the [Blu86] Hamiltonicity protocol in which the commitment scheme is instantiated using public-key encryption also satisfy Definition 6.2.¹²

We now describe an interesting example of an instance-dependent trapdoor Σ -protocol such that trapdoor generation actually depends on the instance: the [GMR89] protocol for quadratic residuosity. Recall that an input $x = (N, y)$ to this protocol consists of an integer $N = pq$ that is a product of two primes along with an element $y \in \mathbb{Z}_N^\times$. An instance x is in the language QR if y is a quadratic residue modulo N . A witness w for this fact is a square root of y modulo N . The [GMR89] protocol $\Pi = \Pi_{\text{GMR}}$ is described in Fig. 3.




Figure 3: The Zero Knowledge Proof System Π_{GMR} for Quadratic Residuosity.

¹²This requires one further modification: the prover must additionally commit to the hidden permutation π and reveal it when asked to reveal the entire graph. We require this so that the bad-challenge function is computable given the PKE secret key – naively, the bad-challenge function would require solving a graph isomorphism problem.

We additionally consider the protocol Π_{GMR}^t repeated t times in parallel for $t = \Omega(\lambda^\epsilon)$. Note that this is indeed a Σ -protocol (with an empty common reference string) with bad-challenge function $f(x, \mathbf{a}) = \mathbf{e}$ such that $e_i = QR(N, a_i)$ for all i , and $QR(N, a)$ is defined to be 1 if and only if a is a square mod N . This holds because for any $x = (N, y)$ such that y is not a quadratic residue modulo N , if $a \in \mathbb{Z}_N^\times$ and $QR(N, a) = 1$, then $QR(N, ay) = 0$ and hence then “1” challenge associated to a cannot be answered by any third message z ; similarly, if $QR(N, a) = 0$ then the “0” challenge associated to a cannot be answered by any third message z .

Finally, we note that the function $f(x, \mathbf{a})$ is efficiently computable given the factorization of $N = p \cdot q$, so we conclude that Π_{GMR}^t is an instance-dependent trapdoor Σ -protocol with auxiliary information $\text{aux} = (p, q)$ and trapdoor $\tau = \text{aux}$ (satisfying subexponential soundness if $t \geq \lambda^\epsilon$). Thus, we conclude

Corollary 6.4. *Assuming the existence of circular-secure FHE, for any $t \geq \lambda^\epsilon$, there exists a hash family \mathcal{H} soundly instantiating the Fiat-Shamir heuristic for Π_{GMR}^t .*

We obtain Corollary 1.5 as a consequence of Corollary 6.4 along with one of the main results from [DNRS99] (additionally assuming that $\text{QR} \notin \text{BPP}$), which generalizes to any protocol (not just Π_{GMR}^t):

Theorem 6.5 ([DNRS99]). *If there exists a hash family \mathcal{H} that soundly instantiates the Fiat-Shamir transform for a 3-message protocol Π for a language $L \notin \text{BPP}$, then Π is not zero knowledge.*

For clarity, we include a proof of Theorem 6.5; we only consider the standard definition of (auxiliary input) zero knowledge as opposed to the weakenings introduced in [DNRS99], but the argument extends to such weakenings.

Proof. (sketch) Let \mathcal{H} soundly instantiate the Fiat-Shamir transform for Π , and suppose for the sake of contradiction that Π is zero knowledge. We then define the following cheating verifier V^* for Π :

- Auxiliary input: a hash key k (sampled according to $\mathcal{H}.\text{Gen}$).
- Second message: upon receiving the first message \mathbf{a} , V^* computes $\mathbf{e} = h(k, \mathbf{a})$ and sends \mathbf{e} to the prover.
- Output: upon receiving a third message \mathbf{z} , V^* outputs the transcript $(\mathbf{a}, \mathbf{e}, \mathbf{z})$.

If Π is (auxiliary input) zero knowledge, then there is a PPT simulator $S^*(x, k)$ that produces a computationally indistinguishable transcript $(\mathbf{a}, \mathbf{e}, \mathbf{z})$ (for all $x \in L$ and for $k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda)$). In particular, this guarantees that transcripts output by S^* satisfy (with all but negligible probability)

1. $\mathbf{e} = H_k(\mathbf{a})$, and
2. $V(x, \mathbf{a}, \mathbf{e}, \mathbf{z}) = 1$,

because these are both efficiently checkable conditions given (x, k) .

Now, since we assumed that $L \notin \text{BPP}$, there must *also* exist some $x^* \notin L$ such that $S^*(x^*, k)$ produces a transcript $(\mathbf{a}, \mathbf{e}, \mathbf{z})$ satisfying conditions (1) and (2) with all but negligible probability;

otherwise, we would have a BPP algorithm deciding the language L (sample k , run $S^*(x, k)$, and check if conditions (1) and (2) are satisfied).

This allows us to contradict the soundness of the Fiat-Shamir protocol $\tilde{\Pi}$ defined by Π and \mathcal{H} : the simulator S^* , given x^* and $k \leftarrow \mathcal{H}.\text{Gen}(1^\lambda)$, exactly breaks the soundness of $\tilde{\Pi}$ by the previous paragraph. This completes the proof. \square

Theorem 6.5 and Corollary 6.4 directly imply that the (parallel repeated) [GMR89] protocol is not zero-knowledge (assuming the existence of circular-secure FHE and the hardness of quadratic residuosity), as desired.

Acknowledgements

We thank Adam Sealton for comments on an earlier version of this work, and in particular for pointing out that our NIZK argument system has a statistically sound mode.

References

- [AD97] Miklós Ajtai and Cynthia Dwork, *A public-key cryptosystem with worst-case/average-case equivalence*, Proceedings of the twenty-ninth annual ACM symposium on Theory of computing, ACM, 1997, pp. 284–293.
- [Ajt96] Miklós Ajtai, *Generating hard instances of lattice problems*, Proceedings of the twenty-eighth annual ACM symposium on Theory of computing, ACM, 1996, pp. 99–108.
- [Bar01] Boaz Barak, *How to go beyond the black-box simulation barrier*, FOCS, IEEE, 2001, p. 106.
- [BCKP14] Nir Bitansky, Ran Canetti, Yael Tauman Kalai, and Omer Paneth, *On virtual grey box obfuscation for general circuits*, International Cryptology Conference, Springer, 2014, pp. 108–125.
- [BDG⁺13] Nir Bitansky, Dana Dachman-Soled, Sanjam Garg, Abhishek Jain, Yael Tauman Kalai, Adriana López-Alt, and Daniel Wichs, *Why "fiat-shamir for proofs" lacks a proof*, Theory of Cryptography - 10th Theory of Cryptography Conference, TCC 2013, Tokyo, Japan, March 3-6, 2013. Proceedings, 2013, pp. 182–201.
- [BFM88] M Blum, P Feldman, and S Micali, *Non-interactive zero-knowledge proof systems and applications*, Proceedings of the 20th Annual ACM Symposium on Theory of Computing, 1988, pp. 103–112.
- [BGI⁺01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil Vadhan, and Ke Yang, *On the (im) possibility of obfuscating programs*, Annual International Cryptology Conference – CRYPTO 2001, Springer, 2001, Journal version appears in JACM 2012, pp. 1–18.
- [BGJ⁺16] Nir Bitansky, Shafi Goldwasser, Abhishek Jain, Omer Paneth, Vinod Vaikuntanathan, and Brent Waters, *Time-lock puzzles from randomized encodings*, Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016 (Madhu Sudan, ed.), ACM, 2016, pp. 345–356.
- [BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan, *(leveled) fully homomorphic encryption without bootstrapping*, Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ACM, 2012, pp. 309–325.
- [BHY09] Mihir Bellare, Dennis Hofheinz, and Scott Yilek, *Possibility and impossibility results for encryption and commitment secure under selective opening*, Annual International Conference on the Theory and Applications of Cryptographic Techniques, Springer, 2009, pp. 1–35.
- [BLMR13] Dan Boneh, Kevin Lewi, Hart Montgomery, and Ananth Raghunathan, *Key homomorphic prfs and their applications*, Advances in Cryptology–CRYPTO 2013, Springer, 2013, pp. 410–428.

- [BLP⁺13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé, *Classical hardness of learning with errors*, Proceedings of the forty-fifth annual ACM symposium on Theory of computing, ACM, 2013, pp. 575–584.
- [Blu86] Manuel Blum, *How to prove a theorem so no one else can claim it*, Proceedings of the International Congress of Mathematicians, vol. 1, 1986, p. 2.
- [BLV03] B Barak, Y Lindell, and S Vadhan, *Lower bounds for non-black-box zero knowledge*, Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on, IEEE, 2003, pp. 384–393.
- [BP15] Nir Bitansky and Omer Paneth, *ZAPs and Non-Interactive Witness Indistinguishability from Indistinguishability Obfuscation*, Theory of Cryptography - TCC 2015, 2015.
- [BR93] Mihir Bellare and Phillip Rogaway, *Random oracles are practical: A paradigm for designing efficient protocols*, Proceedings of the 1st ACM conference on Computer and communications security, ACM, 1993, pp. 62–73.
- [Bra12] Zvika Brakerski, *Fully homomorphic encryption without modulus switching from classical GapSVP*, Advances in Cryptology–CRYPTO 2012, Springer, 2012, pp. 868–886.
- [BV11] Zvika Brakerski and Vinod Vaikuntanathan, *Efficient fully homomorphic encryption from (standard) LWE*, Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science, IEEE Computer Society, 2011, pp. 97–106.
- [BV14] ———, *Lattice-based FHE as secure as PKE*, Proceedings of the 5th conference on Innovations in theoretical computer science, ACM, 2014, pp. 1–12.
- [BV15] ———, *Constrained key-homomorphic prfs from standard lattice assumptions*, Theory of Cryptography Conference, Springer, 2015, pp. 1–30.
- [CC17] Ran Canetti and Yilei Chen, *Constraint-hiding constrained prfs for nc1 from lwe*, Annual International Conference on the Theory and Applications of Cryptographic Techniques, Springer, 2017, pp. 446–476.
- [CCH⁺18] Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, and Ron D. Rothblum, *Fiat-shamir from simpler assumptions*, IACR Cryptology ePrint Archive **2018** (2018).
- [CCR16] Ran Canetti, Yilei Chen, and Leonid Reyzin, *On the correlation intractability of obfuscated pseudorandom functions*, Theory of Cryptography Conference, Springer, 2016, pp. 389–415.
- [CCRR18] Ran Canetti, Yilei Chen, Leonid Reyzin, and Ron D Rothblum, *Fiat-Shamir and correlation intractability from strong KDM-secure encryption*, Annual International Conference on the Theory and Applications of Cryptographic Techniques – EUROCRYPT 2018, Springer, 2018, pp. 91–122.
- [CGH04] Ran Canetti, Oded Goldreich, and Shai Halevi, *The random oracle methodology, revisited*, Journal of the ACM (JACM) **51** (2004), no. 4, 557–594.

- [CHK03] Ran Canetti, Shai Halevi, and Jonathan Katz, *A forward-secure public-key encryption scheme*, IACR Cryptology ePrint Archive **2003** (2003), 83.
- [DN01] Ivan Damgård and Jesper Buus Nielsen, *Perfect hiding and perfect binding universally composable commitment schemes with constant expansion factor*, IACR Cryptology ePrint Archive **2001** (2001), 91.
- [DNRS99] Cynthia Dwork, Moni Naor, Omer Reingold, and Larry J. Stockmeyer, *Magic functions*, FOCS, IEEE Computer Society, 1999, pp. 523–534.
- [FLS99] Uriel Feige, Dror Lapidot, and Adi Shamir, *Multiple noninteractive zero knowledge proofs under general assumptions*, SIAM Journal on Computing **29** (1999), no. 1, 1–28.
- [FS86] Amos Fiat and Adi Shamir, *How to prove yourself: Practical solutions to identification and signature problems*, Conference on the Theory and Application of Cryptographic Techniques, Springer, 1986, pp. 186–194.
- [Gen09] Craig Gentry, *Fully homomorphic encryption using ideal lattices*, STOC, ACM, 2009, pp. 169–178.
- [GGH⁺13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters, *Candidate Indistinguishability Obfuscation and Functional Encryption for all Circuits*, Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on, 2013, pp. 40–49.
- [GK03] Shafi Goldwasser and Yael Tauman Kalai, *On the (in) security of the Fiat-Shamir paradigm*, Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on, IEEE, 2003, pp. 102–113.
- [GKP⁺13] Shafi Goldwasser, Yael Kalai, Raluca Ada Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich, *Reusable garbled circuits and succinct functional encryption*, Proceedings of the forty-fifth annual ACM symposium on Theory of computing, ACM, 2013, pp. 555–564.
- [GKR08] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum, *Delegating computation: interactive proofs for muggles*, Proceedings of the fortieth annual ACM symposium on Theory of computing, ACM, 2008, pp. 113–122.
- [GKW17a] Rishab Goyal, Venkata Koppula, and Brent Waters, *Lockable obfuscation*, Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on, IEEE, 2017, pp. 612–621.
- [GKW17b] ———, *Separating semantic and circular security for symmetric-key bit encryption from the learning with errors assumption*, Annual International Conference on the Theory and Applications of Cryptographic Techniques, Springer, 2017, pp. 528–557.
- [GKW18] ———, *Collusion resistant traitor tracing from learning with errors*, Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, ACM, 2018, pp. 660–670.

- [GMR85] S Goldwasser, S Micali, and C Rackoff, *The knowledge complexity of interactive proof systems*, Proceedings of the seventeenth annual ACM symposium on Theory of computing, ACM, 1985, pp. 291–304.
- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff, *The knowledge complexity of interactive proof systems*, SIAM Journal on computing **18** (1989), no. 1, 186–208.
- [GO94] Oded Goldreich and Yair Oren, *Definitions and properties of zero-knowledge proof systems*, Journal of Cryptology **7** (1994), no. 1, 1–32.
- [GOS12] Jens Groth, Rafail Ostrovsky, and Amit Sahai, *New techniques for noninteractive zero-knowledge*, J. ACM **59** (2012), no. 3, 11:1–11:35.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan, *Trapdoors for hard lattices and new cryptographic constructions*, Proceedings of the fortieth annual ACM symposium on Theory of computing, ACM, 2008, pp. 197–206.
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters, *Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based*, Advances in Cryptology–CRYPTO 2013, Springer, 2013, pp. 75–92.
- [GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee, *Attribute-based encryption for circuits*, Proceedings of the forty-fifth annual ACM symposium on Theory of computing, ACM, 2013, pp. 545–554.
- [HL18] Justin Holmgren and Alex Lombardi, *Cryptographic hashing from strong one-way functions*, Proceedings - Annual IEEE Symposium on Foundations of Computer Science, FOCS, 2018, to appear.
- [HMR08] Shai Halevi, Steven Myers, and Charles Rackoff, *On seed-incompressible functions*, Theory of Cryptography Conference, Springer, 2008, pp. 19–36.
- [HW15] Pavel Hubacek and Daniel Wichs, *On the communication complexity of secure function evaluation with long output*, Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ACM, 2015, pp. 163–172.
- [KN08] Gillat Kol and Moni Naor, *Cryptography and game theory: Designing protocols for exchanging information*, Theory of Cryptography Conference, Springer, 2008, pp. 320–339.
- [KRR17] Yael Tauman Kalai, Guy N. Rothblum, and Ron D. Rothblum, *From obfuscation to the security of fiat-shamir for proofs*, CRYPTO (2), Lecture Notes in Computer Science, vol. 10402, Springer, 2017, pp. 224–251.
- [Lev73] Leonid Anatolevich Levin, *Universal sequential search problems*, Problemy Peredachi Informatsii **9** (1973), no. 3, 115–116.
- [PRSD17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz, *Pseudorandomness of ring-lwe for any ring and modulus*, Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, ACM, 2017, pp. 461–473.

- [PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters, *A framework for efficient and composable oblivious transfer*, Annual international cryptology conference, Springer, 2008, pp. 554–571.
- [Reg09] Oded Regev, *On lattices, learning with errors, random linear codes, and cryptography*, Journal of the ACM (JACM) **56** (2009), no. 6, 34.
- [Rot13] Ron D Rothblum, *On the circular security of bit-encryption*, Theory of Cryptography, Springer, 2013, pp. 579–598.
- [SW14] Amit Sahai and Brent Waters, *How to use indistinguishability obfuscation: deniable encryption, and more*, Proceedings of the forty-sixth annual ACM symposium on Theory of computing, ACM, 2014, pp. 475–484.
- [WZ17] Daniel Wichs and Giorgos Zirdelis, *Obfuscating compute-and-compare programs under lwe*, 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2017, pp. 600–611.
- [Zha16] Mark Zhandry, *The magic of elfs*, Proceedings, Part I, of the 36th Annual International Cryptology Conference on Advances in Cryptology—CRYPTO 2016-Volume 9814, Springer-Verlag New York, Inc., 2016, pp. 479–508.