2-Message Publicly Verifiable WI from (Subexponential) LWE

Alex Lombardi* Vinod Vaikuntanathan† Daniel Wichs‡

July 10, 2019

Abstract

We construct a 2-message publicly verifiable witness indistinguishable argument system for NP assuming that the Learning with Errors (LWE) problem is subexponentially hard. Moreover, the protocol is "delayed input"; that is, the verifier message in this protocol does not depend on the instance. This means that a single verifier message can be reused many times.

We construct two variants of this argument system: one variant is *adaptively sound*, while the other is *public-coin* (but only non-adaptively sound).

We obtain our result via a generic transformation showing that the correlation intractable hash families constructed by Canetti et al. (STOC 2019) and Peikert and Shiehian (CRYPTO 2019) suffice to construct such 2-message WI arguments when combined with an appropriately chosen "trapdoor Σ-protocol." Our construction can be seen as an adaptation of the Dwork-Naor "reverse randomization" paradigm (FOCS '00) for constructing ZAPs to the setting of *computational soundness* rather than statistical soundness. Our adaptation of the Dwork-Naor transformation crucially relies on complexity leveraging to prove that soundness is preserved.

*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.

†MIT. Email: vindov@mit.edu. 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

In this note, we consider the question of constructing 2-message witness indistinguishable (WI) arguments for NP that are *publicly verifiable*; that is, the argument system consists of a single verifier message followed by a single prover message, and anyone can verify a proof given only the transcript.

In a seminal work, Dwork and Naor [DN00] showed that such argument systems can be constructed given any non-interactive zero knowledge (NIZK) proof system in the common random string model; given the state-of-the-art on NIZK, this yields constructions assuming the hardness of factoring [FLS99] as well as under falsifiable assumptions on bilinear maps [CHK03, GOS06].

In recent work, Canetti et al. [CCH+19] and Peikert and Shiehian [PS19] gave constructions of NIZK argument systems from *lattice assumptions*[1]; however, the [DN00] transformation cannot be directly applied to these constructions in order to obtain 2-message WI arguments. The issue is that both of these works construct NIZKs that are either (1) statistically sound, but requiring a structured common reference string, or (2) using a uniformly random CRS, but only satisfying soundness against computationally bounded provers. On the other hand, the [DN00] transformation crucially assumes that the underlying NIZK satisfies statistical soundness and uses a uniformly random CRS.

In this work, we show that a slight modification of the [DN00] transformation can be applied to the [CCH+19, PS19] NIZKs in order to obtain 2-message publicly verifiable WI arguments for NP. Unlike the [DN00] construction, we rely on *complexity leveraging* in order to prove soundness of the 2-message argument system, so we must rely on the subexponential hardness of LWE in order to prove security. As a result, we obtain the following theorem.

Theorem 1.1. *Assuming the subexponential hardness of LWE, there exist two-message publicly verifiable WI arguments for NP.*

We construct two variants of such an argument system: in one variant, soundness is *adaptive* (that is, soundness holds even when the cheating prover is allowed to choose the false statement that he wants to prove), while in the other, the protocol is *public-coin* (that is, the verifier message is a uniformly random string). Both variants are "delayed-input" protocols – meaning that the verifier message does not depend on the instance x – so in either variant, the verifier message can be reused across many executions (even for different statements).

While our construction can be seen as a new variant of the [DN00] transformation from NIZKs to 2-message arguments, we choose to present the construction as a compiler from (sufficiently structured) "trapdoor Σ-protocols" [CCH+19] to 2-message arguments, combining a special-purpose instantiation of the Fiat-Shamir heuristic with a [DN00]-like transformation. More specifically, we give a construction combining dual Regev encryption with the correlation intractable hash families of [CCH+19, PS19].

1.1 Concurrent Work

In concurrent and independent work, Badrinarayan et al. [BFJ+19] note essentially the same construction of 2-message WI arguments from LWE. Moreover, they give an exciting extension of the

[1] [CCH+19] gave a construction from a circular-secure variant of the learning with errors (LWE) assumption, while [PS19] weakened the assumption to plain LWE.

result that yields a 2-message (publicly verifiable) WI argument system satisfying *statistical witness indistinguishability*. Such argument systems were not previously known under any standard cryptographic assumption, and we do not give such a construction in this note.

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 $\{\mathcal{A}_\lambda\}$,

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

2.1 Witness Indistinguishable Arguments

Definition 2.1. *A* witness indistinguishable arugment system Π *for an* NP *relation R consists of ppt interactive algorithms (P, V) with the following syntax.*

- *$P(x, w)$ is an interactive algorithm that takes as input an instance x and witness w that $(x, w) \in R$.*
- *$V(x)$ is an interactive algorithm that takes as input an instance x. At the end of an interaction, 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 accepts at the end of an interaction $\langle P(x, w), V(x)\rangle$.*
- **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 V accepts with probability $\nu(\lambda)$ at the end of an interaction $\langle P^*(x), V(x)\rangle$.*
- **Witness Indistinguishability.** *For every ppt (malicious) verifier V^* and every ensemble $\{(x_n, (w_{0,n}, w_{1,n}), z_n) : (x_n, w_{0,n}), (x_n, w_{1,n}) \in R_n\}$, the distribution ensembles*

$$\mathsf{view}_{V^*}\langle P(x, w_0), V^*(x, w_0, w_1, z)\rangle$$

and

$$\mathsf{view}_{V^*}\langle P(x, w_1), V^*(x, w_0, w_1, z)\rangle$$

are computationally indistinguishable.

In the work, we focus on obtaining two message WI arguments for NP. A (two message) WI argument system can also satisfy various stronger properties. We list some important variants below.

- **Publicly Verifiable**: A WI argument system is publicly verifiable if the verifier's accept/reject algorithm is an efficiently computable function of the transcript (independent of the verifier's internal state).
- **Public Coin**: A WI argument system is *public coin* if all (honest) verifier messages are uniformly random strings (sampled independently of the protocol so far). Note that any public coin protocol is publicly verifiable.
- **Delayed Input**: A *two-message* WI argument system is *delayed input* if the (honestly sampled) verifier message does not depend on the instance x.
- **Adaptive Soundness**: A *two-message, delayed-input* protocol Π is adaptively sound if for every polynomial size algorithm $P^* = \{P^*_\lambda\}$, there is a negligible function ν such that for all λ,

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

3 Correlation Intractable Hash Families

In this section, we recall the notion of correlation intractability [CGH04], specialization to "efficiently-searchable relations" [CCH+19], and LWE-based instantiation [PS19].

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 $\mathcal{H} = \{h_\lambda : \{0,1\}^{s(\lambda)} \times \{0,1\}^{n(\lambda)} \to \{0,1\}^{m(\lambda)}\}_{\lambda\in\mathbb{N}}$ of keyed hash functions, along with a pair of p.p.t. algorithms:*

- $\mathcal{H}.\mathsf{Gen}(1^\lambda)$ *outputs a hash key $k \in \{0,1\}^{s(\lambda)}$.*
- $\mathcal{H}.\mathsf{Hash}(k, x)$ *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 cay that $\mathcal{H}$ is public-coin[2] *if $\mathcal{H}.\mathsf{Gen}$ 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 $\mathcal{H} = \{h_\lambda : \{0,1\}^{s(\lambda)} \times \{0,1\}^{n(\lambda)} \to \{0,1\}^{m(\lambda)}\}$ is said to be R-*correlation intractable with security *(s, δ) if for every s-size $\mathcal{A} = \{\mathcal{A}_\lambda\}$,*

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

*We say that $\mathcal{H}$ is R-*correlation intractable with security *δ if it is (λ^c, δ)-correlation intractable for all $c > 1$. Finally, we say that $\mathcal{H}$ is R-*correlation intractable *if it is $(\lambda^c, \frac{1}{\lambda^c})$-correlation intractable for all $c > 1$.*

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 $\nu(\lambda) = \mathrm{negl}(\lambda)$ such that for every $R \in \mathcal{R}$,*

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

[2]Sometimes "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 $k \leftarrow \mathcal{H}.\mathsf{Gen}(1^\lambda)$. For our purposes (e.g. ignoring compactness), this definition is equivalent to ours.

3.1 Efficiently Searchable Relations

As in [CCH+19,PS19] we make use of hash functions that are correlation intractable 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 3.3 (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 $(x, y) \in R$.*

Definition 3.4 (Efficiently Searchable Relation, [CLW18])**.** *We say that a (necessarily unique-output) relation ensemble R is* searchable in (non-uniform) time T *if there exists a function $f = f_R : \{0,1\}^* \to \{0,1\}^*$ 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* $\mathrm{poly}(n)$.

In this work, we make use of the hash functions of [PS19], which are correlation-intractable for efficiently searchable relations under the LWE assumption (with polynomial modulus). Moreover, we use the fact that under subexponential LWE, the [PS19] hash family is in fact 2^{-m^δ}-correlation intractable for some $\delta > 0$.

Theorem 3.5 ([PS19])**.** *Assume the subexponential hardness of LWE. Then, there exists some $\delta > 0$ such that for all polynomial functions $(n(\cdot), m(\cdot), T(\cdot))$, there is a hash family $\mathcal{H} = \{h_\lambda : \{0,1\}^s \times \{0,1\}^n \to \{0,1\}^m\}$ that is $2^{-m(\lambda)^\delta}$-correlation intractable for all relations searchable in time T.*

4 Reverse Randomization-Compatible Trapdoor Σ-Protocols

In this section, we present a variant of "trapdoor Σ-protocols" [CCH+19] that suffice for our transformation. The key differences as compared to the trapdoor Σ-protocols of [CCH+19] are as follows.

- We require that the honestly generated CRS is uniformly random and that the "fake CRS" distribution is statistically close to uniform.
- We require malicious-verifier witness indistinguishability rather than just honest-verifier zero knowledge (these two properties are equivalent for protocols with polynomial-size challenge spaces and their parallel repetitions).

As we will explain, this can be achieved by instantiating the generic commitment scheme used in the [Blu86, FLS99] Σ-protocols using dual Regev encryption.

Definition 4.1 (Reverse Randomization-Compatible Trapdoor Σ-Protocol)**.** *We say that a 3-message protocol $\Pi = (\mathsf{Gen}, P, V)$ in the CRS model is a* reverse randomization-compatible trapdoor Σ-protocol *if there are p.p.t. algorithms* $\mathsf{TrapGen}, \mathsf{BadChallenge}$ *with the following syntax.*

- $\mathsf{TrapGen}(1^\lambda)$ *takes as input the security parameter. It outputs a common reference string* $\mathsf{crs} \in \{0,1\}^\ell$ *along with a trapdoor* td.
- $\mathsf{BadChallenge}(\mathsf{td}, \mathsf{crs}, x, \mathbf{a})$ *takes as input a trapdoor* td*, common reference string* crs*, instance x, and first message $\mathbf{a}$. It outputs a challenge $\mathbf{e}$.*

We additionally require the following properties.

- ***Witness Indistinguishability with Uniform CRS.***
- ***CRS Indistinguishability:*** *The* crs *distribution output by* $\mathsf{TrapGen}(1^\lambda)$ *is statistically indistinguishable from the uniform distribution* U_ℓ*.*
- ***Efficient Special Soundness:*** *for every instance* $x \notin L$ *and for all* $(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{TrapGen}(1^\lambda)$*, if* $(\mathsf{crs}, x, \mathbf{a}, \mathbf{e}, \mathbf{z})$ *is a valid transcript for* Π*, then* $\mathbf{e} = \mathsf{BadChallenge}(\mathsf{td}, \mathsf{crs}, x, \mathbf{a})$*.*

Remark 4.1. Assuming the (polynomial) hardness of LWE, there is a reverse randomization-compatible trapdoor Σ-protocol for all of NP.

Proof. We instantiate Blum's Hamiltonicity protocol [Blu86] (or the [FLS99] Hamiltonicity protocol) in the CRS model using dual Regev encryption [GPV08]. The fact that these schemes satisfy efficient special soundness was already argued in [CCH+19]. Since dual Regev public keys are statistically indistinguishable from uniformly random, we are done. □

5 Constructing 2-Message WI

In this section, we show that correlation intractable hash functions for efficiently searchable relations (Section 3) can be combined with reverse randomization-compatible trapdoor Σ-protocols (Section 4) to obtain 2-message publicly verifiable WI arguments.

As we described in the introduction, this can be seen as an extension of the Dwork-Naor "reverse randomization" paradigm to the setting of comptuational soundness.

Construction 5.1 (2-Message WI Protocol)**.** *Let* Π *be a reverse randomization-compatible trapdoor* Σ*-protocol with the following three efficiency properties:*

- *Common reference strings have length* $\ell(\lambda)$*.*
- *Challenges have length* $m(\lambda)$ *for some polynomial function* $m(\cdot)$*.*
- *The algorithm* $\mathsf{BadChallenge}(\tau, \mathsf{crs}, x, \mathbf{a})$ *is computable by a size* T *circuit for some polynomial function* $T(\lambda, n(\lambda))$*.*

Moreover, let $\mathcal{H}$ *denote a hash family that is* $2^{-\ell}\mathrm{negl}(\lambda)$*-correlation intractable for relations searchable in time* T*. We then define the following* 2*-message protocol* $\widetilde{\Pi}$*, which is a combination of the Fiat-Shamir transform (using* $\mathcal{H}$*) and [DN00]-style "reverse randomization."*

- *Verifier message: the verifier samples* λ *common random strings* $\mathsf{crs}_1, \ldots, \mathsf{crs}_t \xleftarrow{\$} \{0,1\}^\ell$ *(for* $t = 2\ell$*) along with a hash key* $k \leftarrow \mathcal{H}.\mathsf{Gen}(1^\lambda)$*.*
- *Prover message: given an instance* x*, witness* w*, and verifier message* $(\mathsf{crs}_1, \ldots, \mathsf{crs}_t, k)$*, the prover does the following.*
  - *Sample a random string* $\mathsf{crs}_P \xleftarrow{\$} \{0,1\}^\ell$ *and set* $\widetilde{\mathsf{crs}}_i = \mathsf{crs}_P \oplus \mathsf{crs}_i$*.*
  - *For* $1 \le i \le t$*, compute* $\mathbf{a}_i \leftarrow \Pi.P(\widetilde{\mathsf{crs}}_i, x, w)$*,* $\mathbf{e}_i = h(k, x||\mathbf{a}_i)$*,* $\mathbf{z} = \Pi.P(\widetilde{\mathsf{crs}}_i, x, w, \mathbf{a}_i, \mathbf{e}_i)$*.*

– *Output* $(\mathbf{a}_i, \mathbf{e}_i, \mathbf{z}_i)_{i=1}^t$.

- *The verifier accepts a transcript* $\left((\mathsf{crs}_i)_{i\le t}, k, x, \mathsf{crs}_P, (\mathbf{a}_i, \mathbf{e}_i, \mathbf{z}_i)_{i\le t}\right)$ *if for all* i, $\mathbf{e}_i = h(k, x||\mathbf{a}_i)$ *and* $\Pi.V(\widetilde{\mathsf{crs}}_i, x, \mathbf{a}_i, \mathbf{e}_i, \mathbf{z}_i) = 1$.

We claim that this construction yields a 2-message (publicly verifiable) WI argument system for NP. Completeness and public verifiability are clear by construction, so we proceed to prove that this protocol is both WI and sound.

Lemma 5.2. *Assuming that* Π *is WI,* $\widetilde{\Pi}$ *is also WI.*

Proof. This is identical to the [DN00] proof of witness indistinguishability, which we sketch here. Fix a malicious verifier V^* along with a statement, pair of witnesses, and auxiliary information (x, w_1, w_2, z). Then, consider the following views $\mathsf{view}^{(j)}$ for $0 \le j \le t$: for every j, let

$$\tau^{(j)} = \left((\mathsf{crs}_i)_{i\le t}, k, x, \mathsf{crs}_P, (\mathbf{a}_i, \mathbf{e}_i, \mathbf{z}_i)_{i\le t}\right)$$

and $\text{view}^{(j)} = (\tau^{(j)}, r)$, where:

- r is the internal randomness of V^*, and $\left((\mathsf{crs}_i)_{i\le t}, k\right) = V^*(x, w_1, w_2, z; r)$.
- For every i, $(\mathbf{a}_i, \mathbf{e}_i, \mathbf{z}_i)$ is computed using $\widetilde{\mathsf{crs}}_i := \mathsf{crs}_i \oplus \mathsf{crs}_P$. Moreover, it is computed using witness w_1 if and only if $j \ge i$ (and witness w_2 otherwise).

By construction, $\text{view}^{(0)}$ is the view of V^* in an interaction with an honest prover using w_1, and $\text{view}^{(t)}$ is the interaction between V^* and an honest prover using w_2. The computational indistinguishability of $\text{view}^{(j)}$ and $\text{view}^{(j+1)}$ for every j follows from the (malicious verifier) witness indistinguishability of Π. □

Lemma 5.3. *Assuming that* $\mathcal{H}$ *is* $2^{-\ell}\text{negl}(\lambda)$*-correlation intractable for all relations searchable in time* $T(\lambda, n(\lambda))$, $\widetilde{\Pi}$ *is adaptively sound.*

Proof. Suppose that P^* is an efficient cheating prover that breaks the adaptive soundness of $\widetilde{\Pi}$ with non-negligible probability, meaning that

$$\Pr_{\substack{(\mathsf{crs}_1,\ldots,\mathsf{crs}_t),k \\ (x,\mathsf{crs}_P,\tilde{\pi})\leftarrow P(\mathsf{crs}_1,\ldots,\mathsf{crs}_t,k)}} [x \notin L \wedge V \text{ accepts } (x, \mathsf{crs}_1, \ldots, \mathsf{crs}_t, k, \mathsf{crs}_P, \tilde{\pi})] = \epsilon(\lambda)$$

for some non-negligible function $\epsilon(\cdot)$. We proceed to define a sequence of hybrid experiments where we change the underlying distributions and win conditions. Let $\mathsf{crs}^* \leftarrow \{0,1\}^\ell$ denote a uniformly random string of length ℓ sampled independently of the above random variables. Then, we have that

$$\Pr_{\substack{\mathsf{crs}^*,(\mathsf{crs}_1,\ldots,\mathsf{crs}_t),k \\ (x,\mathsf{crs}_P,\tilde{\pi})\leftarrow P(\mathsf{crs}_1,\ldots,\mathsf{crs}_t,k)}} [x \notin L \wedge V \text{ accepts } (x, \mathsf{crs}_1, \ldots, \mathsf{crs}_t, k, \mathsf{crs}_P, \tilde{\pi}) \wedge \mathsf{crs}_P = \mathsf{crs}^*] = \epsilon(\lambda)2^{-\ell}.$$

Next, in order to invoke correlation intractability, we need to argue that P^* must win while some $\widetilde{\mathsf{crs}}_i$ has a valid trapdoor. In order to have a uniform security reduction, we argue as follows. Since

the CRS distribution output by $\mathsf{TrapGen}(1^\lambda)$ is statistically close to uniform, we know that there exists a set $\mathcal{S} \subset \{0,1\}^\ell$ of size $\frac{1}{2}2^\ell$ such that for every $\mathsf{crs} \in \mathcal{S}$, $\mathsf{TrapGen}(1^\lambda)$ outputs crs with probability at least $\frac{1}{2}2^{-\ell}$. By independence, we conclude that for every fixed string crs^*,

$$\Pr_{\mathsf{crs}_1,\ldots,\mathsf{crs}_t}[\mathsf{crs}^* \oplus \mathsf{crs}_i \notin \mathcal{S} \text{ for all i}] = 2^{-t} = 2^{-2\ell},$$

so we have that

$$\Pr_{\substack{\mathsf{crs}^*,(\mathsf{crs}_1,\ldots,\mathsf{crs}_t),k \\ (x,\mathsf{crs}_P,\tilde{\pi})\leftarrow P(\mathsf{crs}_1,\ldots,\mathsf{crs}_t,k)}}[x \notin L \wedge V \text{ accepts } \wedge \mathsf{crs}_P = \mathsf{crs}^* \wedge \widetilde{\mathsf{crs}}_i \in \mathcal{S} \text{ for some } i] \geq \epsilon 2^{-\ell} - 2^{-2\ell}.$$

Picking a uniformly random $i^* \xleftarrow{\$} [t]$, we further see that

$$\Pr_{\substack{i^*,\mathsf{crs}^*,(\mathsf{crs}_1,\ldots,\mathsf{crs}_t),k \\ (x,\mathsf{crs}_P,\tilde{\pi})\leftarrow P(\mathsf{crs}_1,\ldots,\mathsf{crs}_t,k)}}[x \notin L \wedge V \text{ accepts } \wedge \mathsf{crs}_P = \mathsf{crs}^* \wedge \widetilde{\mathsf{crs}}_{i^*} \in \mathcal{S}] \geq \frac{1}{4\ell}\epsilon 2^{-\ell}.$$

We next consider an alternate experiment in which the uniformly random crs_{i^*} is replaced by the string $\mathsf{crs}^* \oplus \overline{\mathsf{crs}}_{i^*}$ for $(\overline{\mathsf{crs}}_{i^*}, \mathsf{td}_{i^*}) \leftarrow \mathsf{TrapGen}(1^\lambda)$. Since every string in $\mathcal{S}$ has weight at least $\frac{1}{2}2^{-\ell}$ in the $\mathsf{TrapGen}$ crs distribution, we see that

$$\Pr_{\substack{i^*,\mathsf{crs}^*,\overline{\mathsf{crs}}_{i^*},(\mathsf{crs}_1,\ldots,\mathsf{crs}_t),k \\ (x,\mathsf{crs}_P,\tilde{\pi})\leftarrow P(\mathsf{crs}_1,\ldots,\mathsf{crs}^*\oplus\overline{\mathsf{crs}}_{i^*},\ldots,\mathsf{crs}_t,k)}}[x \notin L \wedge V \text{ accepts } \wedge \mathsf{crs}_P = \mathsf{crs}^* \wedge \widetilde{\mathsf{crs}}_{i^*} \in \mathcal{S}] \geq \frac{1}{8\ell}\epsilon 2^{-\ell}.$$

Finally, we claim that this violates the $2^{-\ell}\mathrm{negl}(\lambda)$-correlation intractability of $\mathcal{H}$. Formally, an adversary $\mathcal{A}'$ can sample $i^*, (\overline{\mathsf{crs}}_{i^*}, \mathsf{td}_{i^*})$ and declare the relation

$$R_{\overline{\mathsf{crs}}_{i^*},\mathsf{td}_{i^*}} = \{(x||\mathbf{a},\mathbf{e}) : \mathbf{e} = \mathsf{BadChallenge}(\mathsf{td}_{i^*}, \overline{\mathsf{crs}}_{i^*}, x, \mathbf{a}).\}$$

Then, upon receiving a hash key k, $\mathcal{A}'$ can sample crs^* and $(\mathsf{crs}_1,\ldots,\mathsf{crs}_t)$ itself and call $(x,\mathsf{crs}_P,\tilde{\pi}) \leftarrow P^*(\mathsf{crs}_1,\ldots,\mathsf{crs}^*\oplus\overline{\mathsf{crs}}_i,\ldots,\mathsf{crs}_t)$. Finally, $\mathcal{A}'$ outputs the pair $(x,\mathbf{a}_{i^*})$. Whenever $x \notin L$, $\mathsf{crs}_P = \mathsf{crs}^*$, and V accepts the output of P^* in the above experiment, by the efficient special soundness of Π, we will have that $(x,\mathbf{a}_{i^*}) \in R_{\overline{\mathsf{crs}}_{i^*},\mathsf{td}_{i^*}}$, completing the reduction. □

5.1 Parameter Settings and Instantiation

Combining Section 5 with Theorem 3.5 and Remark 4.1, we obtain the following LWE-based instantiation of 2-message publicly verifiable WI. Assume that LWE is $2^{-\lambda^\delta} \cdot \mathrm{negl}(\lambda)$-hard for some fixed $\delta > 0$.

- Using dual Regev encryption and the [Blu86] proof system for Hamiltonicity (repeated $\lambda^{\frac{2}{\delta}}$ times in parallel), there is a reverse randomization-compatible trapdoor Σ-protocol Π with a crs of size λ and challenges of length $\lambda^{\frac{2}{\delta}}$.
- Using Theorem 3.5, there is a hash family that is $2^{-\lambda^2} \cdot \mathrm{negl}(\lambda)$-correlation intractable for all relations that are searchable in time $T(\lambda)$ sufficient to compute the $\mathsf{BadChallenge}$ function associated to Π.

- Applying Section 5, we conclude that the protocol $\widetilde{\Pi}$ in Construction 5.1 (using these building blocks) is a 2-message publicly verifiaible WI argument system for NP. Moreover, it satisfies adaptive soundness (again by Section 5). Finally, since hash keys in the hash family $\mathcal{H}$ are pseudorandom, we conclude that another variant of $\widetilde{\Pi}$ (in which the verifier message is uniformly random) is a non-adaptively sound publicly-verifiable WI argument.

References

[BFJ+19] Saikrishna Badrinarayan, Rex Fernando, Aayush Jain, Dakshita Khurana, and Amit Sahai, *Statistical zap arguments*, Cryptology ePrint Archive, Report 2019/780, 2019, https://eprint.iacr.org/2019/780.

[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.

[CCH+19] Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, Ron D. Rothblum, and Daniel Wichs, *Fiat-shamir: From practice to theory*, Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, ACM, 2019.

[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.

[CLW18] Ran Canetti, Alex Lombardi, and Daniel Wichs, *Fiat-shamir: From practice to theory, part ii (non-interactive zero knowledge and correlation intractability from circular-secure fhe)*, IACR Cryptology ePrint Archive **2018** (2018).

[DN00] Cynthia Dwork and Moni Naor, *Zaps and their applications*, Proceedings 41st Annual Symposium on Foundations of Computer Science, IEEE, 2000, pp. 283–293.

[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.

[GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai, *Non-interactive zaps and new techniques for nizk*, Annual International Cryptology Conference, Springer, 2006, pp. 97–111.

[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.

[PS19] Chris Peikert and Sina Shiehian, *Noninteractive zero knowledge for np from (plain) learning with errors*, Tech. report, IACR Cryptology ePrint Archive, 2019.