

Two-Message Statistically Sender-Private OT from LWE

Zvika Brakerski*

Nico Döttling†

Abstract

We construct a two-message oblivious transfer (OT) protocol without setup that guarantees statistical privacy for the sender even against malicious receivers. Receiver privacy is game based and relies on the hardness of learning with errors (LWE). This flavor of OT has been a central building block for minimizing the round complexity of witness indistinguishable and zero knowledge proof systems, non-malleable commitment schemes and multi-party computation protocols, as well as for achieving circuit privacy for homomorphic encryption in the malicious setting. Prior to this work, all candidates in the literature from standard assumptions relied on number theoretic assumptions and were thus insecure in the post-quantum setting. This work provides the first (presumed) post-quantum secure candidate and thus allows to instantiate the aforementioned applications in a post-quantum secure manner.

Technically, we rely on the transference principle: Either a lattice or its dual must have short vectors. Short vectors, in turn, can be translated to information loss in encryption. Thus encrypting one message with respect to the lattice and one with respect to its dual guarantees that at least one of them will be statistically hidden.

1 Introduction

Oblivious transfer (OT), introduced by Rabin [Rab81], is one of the most fundamental cryptographic tasks. A sender (S) holds two values μ_0, μ_1 and a receiver (R) holds a bit β . The functionality should allow the receiver to learn μ_β and nothing else, the sender should learn nothing. OT has been a fundamental building block for many cryptographic applications, in particular ones related to secure multi-party computation (MPC), starting with [Yao86, GMW87].

A central measure for the complexity of a protocol or a proof system is its round complexity. One could imagine a protocol implementing the OT functionality with only two messages: a first message from the receiver to the sender, and a second message from the sender to the receiver. Indeed, in the semi-honest setting, where parties are assumed to follow the protocol, this can be achieved based on a variety of concrete cryptographic assumptions (Decisional Diffie-Hellman, Quadratic Residuosity, Decisional Composite Residuosity, Learning with Errors, to name a few), as well as based on generic assumptions such as trapdoor permutations, additively homomorphic encryption and public key encryption with oblivious public key generation (e.g. [BM89, GKM⁺00]).

In the malicious setting, where an adversarial party might deviate from the designated protocol, the ultimate simulation based security notion cannot be achieved in a two message protocol (without

*Weizmann Institute of Science, zvika.brakerski@weizmann.ac.il. Supported by the Israel Science Foundation (Grant No. 468/14), Binational Science Foundation (Grants No. 2016726, 2014276), and by the European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482) and via Project PROMETHEUS (Grant 780701).

†CISPA Helmholtz Center, doettling@cispa.saarland.

assuming setup such as a common random string or a random oracle) [GO94]. The standard security notion in this setting, which originated from the works of Naor and Pinkas [NP01] and Aiello, Ishai and Reingold [AIR01], and was further studied in [Kal05, HK12, BGI⁺17], provides a meaningful relaxation of the standard (simulation-based) security notion. This definition requires that the receiver’s only message is computationally indistinguishable between the cases of $\beta = 0$ and $\beta = 1$ ¹, and that regardless of the receiver’s first message, the sender’s message statistically hides at least one of μ_0, μ_1 . Alternative equivalent formulations are simulation using a computationally unbounded (or exponential time) simulator, or the existence of a computationally unbounded (or exponential time) extractor, that can extract a β value from any receiver message.

With the aforementioned connection to secure MPC, it is not surprising that this notion of malicious statistical sender-private OT (SSP-OT) found numerous applications. In particular in recent years as the round complexity of MPC and related objects is taken to the necessary minimum. Badrinarayanan et al. [BGI⁺17], Jain et al. [JKKR17] and Kalai et al. [KKS18] used it to construct two-message witness indistinguishable proof systems, and even restricted forms of zero-knowledge proof systems.

Badrinarayanan et al. [BGJ⁺17a] used similar techniques to present malicious MPC with minimal round complexity (4-rounds). In particular, their building blocks are SSP-OT and a 3-round semi-malicious MPC protocol (a comparable result was achieved by Halevi et al. [HHPV17] using different techniques, in particular requiring NIZK/ZAP). Khurana and Sahai [KS17] used SSP-OT to construct two-message non-malleable commitment schemes (with respect to the commitment), and Khurana [Khu17] used it (together with ZAPs) to achieve 3-round non-malleable commitments from polynomial assumptions. Badrinarayanan et al. [BGJ⁺17b] relied on SSP-OT to construct 3-round concurrent MPC.

Ostrovsky, Paskin-Cherniavsky and Paskin-Cherniavsky [OPP14] used SSP-OT to show that any fully homomorphic encryption scheme (FHE) can be converted to one that is statistically circuit private even against maliciously generated public keys and ciphertexts.

Our Results and Applications. Prior to this work it was only known how to construct SSP-OT from number theoretic assumptions such as DDH [NP01, AIR01], QR and DCR [HK12]. If setup is allowed, specifically a common random string, then an LWE-based construction by Peikert, Vaikuntanathan and Waters [PVW08] achieves strong simulation security (even in the UC model). However, the aforementioned applications require a construction without setup and could therefore not be instantiated in a post-quantum secure manner. In this work, we construct SSP-OT from the learning with errors (LWE) assumption [Reg05], with polynomial noise-ratio, which translates to the hardness of polynomially approximating short-vector problems (such as SIVP or GapSVP) to within a polynomial factor. Currently, no polynomial time quantum algorithm is known for these problems, and thus they serve as a major candidate for constructing post-quantum secure cryptography.

Relying on our construction, it is possible for the first time, to instantiate the works of [BGI⁺17, JKKR17, KKS18, KS17, BGJ⁺17b] from LWE, i.e. in a post-quantum secure manner, and obtain proof systems with witness-indistinguishable or (limited) zero-knowledge properties, as well as non-malleable commitment schemes and concurrent MPC protocols. It is also possible to construct a round-optimal malicious MPC from LWE by applying the result of [BGJ⁺17a] using our SSP-OT

¹Notice that it is impossible to achieve statistical indistinguishability in this setting, at least against non-uniform malicious receivers.

and the LWE-based 3-round semi-malicious MPC of Brakerski, Halevi and Polychroniadou [BHP17]. Lastly, our result allows to achieve malicious circuit private FHE from LWE by instantiating the [OPP14] result with our LWE-based SSP-OT and relying on the numerous existing LWE-based FHE schemes. We stress that none of these applications had prior post-quantum secure candidates.

1.1 Technical Overview

Our construction relies on some fundamental properties of lattices. For our purposes we will only consider the so called q -ary lattices that can be described as follows. Given a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ for some modulus q and $m \geq n$, we can define $\Lambda_q(\mathbf{A}) = \{\mathbf{y} \in \mathbb{Z}^m : \mathbf{y} = \mathbf{sA} \pmod{q}\}$ which is the lattice defined by the row-span of \mathbf{A} , and $\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{Ax} = \mathbf{0} \pmod{q}\}$ which is the lattice defined by the kernel of \mathbf{A} . Note that both lattices have rank m over the integers, i.e. they contain a set of m linearly independent vectors over the integers (but not modulo q), since they contain $q \cdot \mathbb{Z}^m$. There is a duality relation between these two lattices, both induced by the matrix \mathbf{A} , and this relation will be instrumental for our methods.

An important fact about lattices is that a good basis implies decoding. Specifically, if $\Lambda_q^\perp(\mathbf{A})$ contains m linearly independent vectors (over the integers) of length at most ℓ , then it is possible to decode vectors of the form $\mathbf{sA} + \mathbf{e} \pmod{q}$, if $\|\mathbf{e}\|$ is sufficiently smaller than q/ℓ . Namely, to recover \mathbf{s}, \mathbf{e} . Such a short basis is sometimes called a trapdoor for \mathbf{A} .²

Consider sampling \mathbf{s} uniformly in \mathbb{Z}_q^n and \mathbf{e} from a Gaussian s.t. $\|\mathbf{e}\|$ is slightly below the decoding capability q/ℓ . Then if $\Lambda_q^\perp(\mathbf{A})$ indeed has an ℓ -basis then \mathbf{s}, \mathbf{e} can be recovered from $\mathbf{sA} + \mathbf{e} \pmod{q}$. However, a critical observation for us is that this encoding becomes *lossy* if the lattice $\Lambda_q(\mathbf{A})$ contains a vector of norm $\ll q/\ell$. That is, in this case it is information theoretically impossible to recover the original \mathbf{s} . This is because the component of \mathbf{sA} that is in the direction of the short vector is masked by the noise \mathbf{e} (which is Gaussian and thus has a component in every direction). This property was also used by Goldreich and Goldwasser [GG98] to show that some lattice problems are in **coAM**.

To utilize this structure for our purposes, we specify the OT receiver message to be a matrix \mathbf{A} . Then the OT sender generates $\mathbf{sA} + \mathbf{e} \pmod{q}$ and encodes one of its inputs, say μ_1 using entropy from the vector \mathbf{s} (e.g. using a randomness extractor). We get that this value is recoverable if \mathbf{A} has ℓ -basis and information-theoretically hidden if $\Lambda_q(\mathbf{A})$ has a short vector. If the receiver's choice bit is 1, all it needs to do is generate \mathbf{A} that has an ℓ -trapdoor, for which there are many well known methods to generate such \mathbf{A} 's that are statistically indistinguishable from uniform (starting from [Ajt99] with numerous followups). In order to complete the OT functionality we need to find a way to encode μ_0 in a way that is *lossy* if $\Lambda_q(\mathbf{A})$ has no short vector. This will guarantee that regardless of the (possibly malicious) choice of matrix \mathbf{A} , either μ_0 or μ_1 are information theoretically hidden.

Let us examine the case where all vectors in $\Lambda_q(\mathbf{A})$ are of length $\gg t$ for some parameter t . Then the duality relations expressed in Banaszczyk's transference theorems [Ban93] guarantees that $\Lambda_q^\perp(\mathbf{A})$ has a basis of length $\ll q/t$. In such case we can use the *smoothing* principle to conclude that if \mathbf{x} is a discrete Gaussian with parameter q/t then $\mathbf{Ax} \pmod{q}$ is statistically close to uniform. We can thus instruct the sender to compute $\mathbf{Ax} + \mathbf{d} \pmod{q}$ for some vector \mathbf{d} , and encode μ_1 using entropy extracted from \mathbf{d} . This guarantees lossiness if $\Lambda_q(\mathbf{A})$ has no short vectors as required.

²While the form $\mathbf{sA} + \mathbf{e} \pmod{q}$ bears resemblance to an instance of the LWE problem (to be discussed below), the matrix \mathbf{A} in our setting might be chosen by a malicious party and therefore cannot be assumed to be close to uniform.

Furthermore, it is possible to generate a pseudorandom \mathbf{A} (under the LWE assumption) and specify \mathbf{d} such that \mathbf{d} is recoverable (this \mathbf{A} corresponds to the public key in Regev’s original encryption scheme [Reg05]).

All that is left is to set the relation between ℓ, t, q so as to make sure that if one mode of the OT is decodable then the other is lossy. One may be suspicious whether there is a valid setting of parameters, but in fact there is quite some slackness in the choice of parameters. We can start by setting ℓ, t to be some fixed polynomial in n that is sufficient to guarantee correct recovery in the respective cases. This can be done regardless of the value of q . We will set the parameter q to ensure that if μ_1 is recoverable then μ_0 is not, which is sufficient to guarantee statistical sender privacy against malicious receiver. Specifically, if μ_1 is recoverable then $\Lambda_q(\mathbf{A})$ does not have vectors of length $q/(k\ell)$, where k is some polynomial in n (that does not depend on q), and thus $\Lambda_q^\perp(\mathbf{A})$ has a $k\ell$ basis. We therefore require that $q/t \gg k\ell$, or equivalently $q \gg k\ell t$, which guarantees that μ_0 is not recoverable in this case. Since k, ℓ, t are fixed polynomials in n , it is sufficient to choose q to be a sufficiently larger polynomial than the product $k\ell t$ to guarantee security. Receiver privacy is guaranteed since \mathbf{A} is either statistically indistinguishable from uniform if the choice bit β is 1, or computationally indistinguishable from uniform if $\beta = 0$.

Disadvantages of the Basic Solution, and Our Actual Improved Scheme. The proposal above can indeed be used to implement an SSP-OT. However, when actual parameters are assigned, it becomes apparent that the argument about the lossiness of \mathbf{s} given $\mathbf{s}\mathbf{A} + \mathbf{e} \pmod{q}$ when $\Lambda_q(\mathbf{A})$ has some short vector does not produce sufficient randomness to allow extraction. This can be resolved by repetition (many \mathbf{s} values with the same \mathbf{A}). However, the lossiness argument for \mathbf{d} guarantees much more and in fact allows to extract random bits from \mathbf{d} deterministically. The consequence is an unnecessarily inefficient scheme. In particular, the information rate is inverse polynomial in the security parameter of the scheme.

The scheme we actually introduce and analyze is therefore a balanced version of the above outline, where we “pay” in weakening the lossiness in \mathbf{d} in exchange for strengthening the lossiness for \mathbf{s} , which leads to a scheme with information rate $\tilde{\Omega}(1)$ (achieving constant information rate while preserving statistical security remains an intriguing question). Towards this end, we introduce refinements of known lattice tools that may be of independent interest.

The idea is to improve the lossiness in \mathbf{s} by considering the case where $\Lambda_q(\mathbf{A})$ has multiple short vectors, instead of just one. Intuitively, this will introduce entropy into additional components of \mathbf{s} , thus increasing the lossiness. We formalize this by considering the Gaussian measure of $\Lambda_q(\mathbf{A})$. A high Gaussian measure translates (at least intuitively) to the existence of a multitude of short vectors, formally it characterizes the potency of \mathbf{e} to hide information about \mathbf{s} . The formal argument goes through the optimal Voronoi cell decoder, see Section 3 for formal statement and additional details.

Of course the lossiness in \mathbf{s} needs to be complemented by lossiness in \mathbf{d} if the Gaussian measure of $\Lambda_q(\mathbf{A})$ is small, which translates to having few independent short vectors in $\Lambda_q(\mathbf{A})$. We show that in this case we can derive *partial smoothing* where for a Gaussian \mathbf{x} , the value $\mathbf{A}\mathbf{x} \pmod{q}$ is no longer uniform, but rather is uniform over some subspace modulo q . If the dimension of this subspace is large enough, we can get lossiness for the vector \mathbf{d} and complete the security proof. Partial smoothing and implications are discussed in Section 4.

To apply these principles we need to slightly modify the definition of the vector \mathbf{d} and the matrix \mathbf{A} in the case of $\beta = 0$. Now \mathbf{A} will no longer correspond to the public key of the Regev

scheme but rather, interestingly, to the public key of the batched scheme introduced in [PVW08] (which is also concerned with constructing OT, but allowing setup). The complete construction and analysis can be found in Section 5.

2 Preliminaries

2.1 Statistical Sender-Private Two-Message Oblivious Transfer

We now define the object of main interest in this work, namely SSP-OT. We only define the two-message perfect-correctness variant since this is what we achieve in this work. A two-message oblivious transfer protocol consists of a tuple PPT algorithms $(\text{OTR}, \text{OTS}, \text{OTD})$ with the following syntax.

- $\text{OTR}(1^\lambda, \beta)$ takes the security parameter λ and a selection bit β and outputs a message ot_1 and secret state st .
- $\text{OTS}(1^\lambda, (\mu_0, \mu_1), \text{ot}_1)$ takes the security parameter λ , two inputs $(\mu_0, \mu_1) \in \{0, 1\}^{\text{len}}$ (where len is a parameter of the scheme) and a message ot_1 . It outputs a message ot_2 .
- $\text{OTD}(1^\lambda, \beta, \text{st}, \text{ot}_2)$ takes the security parameter, the bit β , secret state st and message ot_2 and outputs $\mu' \in \{0, 1\}^{\text{len}}$.

Correctness and security are defined as follows.

Definition 2.1. *A tuple $(\text{OTR}, \text{OTS}, \text{OTD})$ is a SSP-OT scheme if the following hold.*

- **Correctness.** *For all $\lambda, \beta, \mu_0, \mu_1$, letting $(\text{ot}_1, \text{st}) = \text{OTR}(1^\lambda, \beta)$, $\text{ot}_2 = \text{OTS}(1^\lambda, (\mu_0, \mu_1), \text{ot}_1)$, $\mu' = \text{OTD}(1^\lambda, \beta, \text{st}, \text{ot}_2)$, it holds that $\mu' = \mu_\beta$ with probability 1.*
- **Receiver Privacy.** *Consider the distribution $\mathcal{D}_\beta(\lambda)$ defined by running $(\text{ot}_1, \text{st}) = \text{OTR}(1^\lambda, \beta)$ and outputting ot_1 . Then $\mathcal{D}_0, \mathcal{D}_1$ are computationally indistinguishable.*
- **Statistical Sender Privacy.** *There exists an extractor OTExt (possibly computationally unbounded) s.t. for any sequence of messages $\text{ot}_1 = \text{ot}_1(\lambda)$ and inputs $(\mu_0, \mu_1) = (\mu_0(\lambda), \mu_1(\lambda))$, the distribution ensembles $\text{OTS}(1^\lambda, (\mu_0, \mu_1), \text{ot}_1)$ and $\text{OTS}(1^\lambda, (\mu_{\beta'}, \mu_{\beta'}), \text{ot}_1)$, where $\beta' = \text{OTExt}(\text{ot}_1)$, are statistically indistinguishable.*

2.2 Linear Algebra, Min-Entropy and Extractors

Random Matrices: The probability that a uniformly random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_2^{n \times m}$ (with $m \geq n$) has full rank is given by

$$\Pr_{\mathbf{A}}[\text{rank}(\mathbf{A}) < n] = 1 - \prod_{i=0}^{n-1} (1 - 2^{i-m}) \leq \sum_{i=0}^{n-1} 2^{i-m} \leq 2^{n-m},$$

where the first inequality follows from the union-bound.

Average Conditional Min-Entropy Let X be a random-variable supported on a finite set \mathcal{X} and let Z be a (possibly correlated) random variable supported on a finite set \mathcal{Z} . The average-conditional min-entropy $\tilde{H}_\infty(X|Z)$ of X given Z is defined as

$$\tilde{H}_\infty(X|Z) = -\log \left(\mathbb{E}_z \left[\max_{x \in \mathcal{X}} \Pr[X = x|Z = z] \right] \right).$$

We will use the following easy-to-establish fact about uniform distributions on binary vector-spaces: If $U, V \subseteq \mathbb{Z}_2^n$ are sub-vectorspaces of \mathbb{Z}_2^n , and if $\mathbf{u} \stackrel{\$}{\leftarrow} U$ and $\mathbf{v} \stackrel{\$}{\leftarrow} V$, then it holds that

$$\tilde{H}_\infty(\mathbf{u}|\mathbf{u} + \mathbf{v}) = \dim(U \cap V).$$

Extractors A function $\text{Ext} : \{0, 1\}^d \times \mathcal{X} \rightarrow \{0, 1\}^\ell$ is called a *seeded strong average-case* (k, ϵ) -*extractor*, if it holds for all random variables X with support \mathcal{X} and Z defined on some finite support that if $\tilde{H}_\infty(X|Z) \geq k$, then it holds that

$$(\mathbf{s}, \text{Ext}(\mathbf{s}, X), Z) \approx_\epsilon (\mathbf{s}, U, Z),$$

where $\mathbf{s} \stackrel{\$}{\leftarrow} \{0, 1\}^d$ and $U \stackrel{\$}{\leftarrow} \{0, 1\}^\ell$. Such extractors can be constructed from universal hash functions [DRS04, DORS08]. In fact, any extractor is an average-case extractor for slightly worse parameters by the averaging principle³.

2.3 Lattices

We recall the standard facts about lattices. A lattice $\Lambda \subseteq \mathbb{R}^m$ is the set of all integer-linear combinations of a set of linearly independent basis-vectors, i.e. for every lattice Λ there exists a full-rank matrix $\mathbf{B} \in \mathbb{R}^{k \times m}$ such that $\Lambda = \Lambda(\mathbf{B}) = \{\mathbf{z} \cdot \mathbf{B} \mid \mathbf{z} \in \mathbb{Z}^k\}$. We call k the rank of Λ and \mathbf{B} a basis of Λ . More generally, for a set $S \subseteq \Lambda$ we denote by $\Lambda(S)$ the smallest sub-lattice of Λ which contains S . Moreover, we will write $\text{rank}(S)$ to denote $\text{rank}(\Lambda(S))$.

The dual-lattice $\Lambda^* = \Lambda^*(\Lambda)$ of a lattice Λ is defined by $\Lambda^*(\Lambda) = \{\mathbf{x} \in \mathbb{R}^n \mid \forall \mathbf{y} \in \Lambda : \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}\}$. Note that it holds that $(\Lambda^*)^* = \Lambda$. The determinant of a lattice Λ is defined by $\det \Lambda = \sqrt{\det(\mathbf{B} \cdot \mathbf{B}^\top)}$ where \mathbf{B} is any basis of Λ . It holds that $\det \Lambda^* = 1/\det \Lambda$. If $\Lambda = \Lambda(\mathbf{B})$ and the norm of each row of \mathbf{B} is at most ℓ , then an argument using Gram-Schmidt orthogonalization establishes $\det \mathbf{B} \leq \ell^k$.

For a basis $\mathbf{B} \in \mathbb{R}^{k \times m}$ of Λ , we define the parallel-epiped of \mathbf{B} by $\mathcal{P}(\mathbf{B}) = \{\mathbf{x} \cdot \mathbf{B} \mid \mathbf{x} \in [-1/2, 1/2)^k\}$. In abuse of notation we write $\mathcal{P}(\Lambda)$ to denote $\mathcal{P}(\mathbf{B})$ for some canonic basis \mathbf{B} of Λ (such as e.g. a Hermite basis). For lattices $\Lambda \subseteq \Lambda_0$, we will use $\mathcal{P}(\Lambda) \cap \Lambda_0$ as a system of (unique) representatives for the quotient group Λ_0/Λ .

We say that a lattice is q -ary if $(q\mathbb{Z})^m \subseteq \Lambda \subseteq \mathbb{Z}^m$. In particular, for every q -ary lattice Λ there exists a matrix $\mathbf{A} \in \mathbb{Z}_q^{k \times m}$ such that $\Lambda = \Lambda_q(\mathbf{A}) = \{\mathbf{y} \in \mathbb{Z}^m \mid \exists \mathbf{x} \in \mathbb{Z}_q^k : \mathbf{y} = \mathbf{x} \cdot \mathbf{A} \pmod{q}\}$. We also define the lattice $\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{y} \in \mathbb{Z}_q^m \mid \mathbf{A} \cdot \mathbf{y} = 0 \pmod{q}\}$. It holds that $(\Lambda_q(\mathbf{A}))^* = \frac{1}{q} \Lambda_q^\perp(\mathbf{A})$.

Gaussians The Gaussian function $\rho_\sigma : \mathbb{R}^m \rightarrow \mathbb{R}$ is defined by

$$\rho_\sigma(\mathbf{x}) = e^{-\pi \cdot \frac{\|\mathbf{x}\|^2}{\sigma^2}}.$$

³i.e. a simple application of Markov's inequality

For a lattice $\Lambda \subseteq \mathbb{R}^m$ and a parameter $\sigma > 0$, we define the *discrete Gaussian distribution* $D_{\Lambda, \sigma}$ on Λ as the distribution with probability-mass function $\Pr[\mathbf{x} = \mathbf{x}'] = \rho_\sigma(\mathbf{x}') / \rho_\sigma(\Lambda)$ for all $\mathbf{x}' \in \Lambda$. Let in the following $\mathcal{B} = \{\mathbf{x} \in \mathbb{R}^m \mid \|\mathbf{x}\| \leq 1\}$ be the closed ball of radius 1 in \mathbb{R}^m . A standard concentration inequality for discrete gaussians on general lattices is provided by Banaszczyk's Theorem.

Theorem 2.2 ([Ban93]). *For any lattice $\Lambda \in \mathbb{R}^m$, parameter $\sigma > 0$ and $u \geq 1/\sqrt{2\pi}$ it holds that*

$$\rho_\sigma(\Lambda \setminus u\sigma\sqrt{m}\mathcal{B}) \leq 2^{-c_u \cdot m} \cdot \rho_\sigma(\Lambda),$$

where $c_u = -\log(\sqrt{2\pi}eu \cdot e^{-\pi u^2})$.

Setting $\Lambda = \mathbb{Z}^m$ and $u = 1$ in Theorem 2.2 we obtain the following corollary.

Corollary 2.3. *Let $\sigma > 0$ and $\mathbf{x} \stackrel{\$}{\leftarrow} D_{\mathbb{Z}^m, \sigma}$. Then it holds that $\|\mathbf{x}\| \leq \sigma \cdot \sqrt{m}$, except with probability 2^{-m} .*

Uniform Matrix Distributions with Decoding Trapdoor For our construction we will need an efficiently samplable ensemble of matrices which is statistically close to uniform and is equipped with an efficient bounded-distance-decoder. Such an ensemble was first constructed by Ajtai [Ajt99] for q -ary lattices with prime q . We use a more efficient ensemble due to Micciancio and Peikert [MP12] which works for arbitrary modulus.

Lemma 2.4 ([MP12]). *Let $\kappa(n) = \omega(\sqrt{\log(n)})$ be any function that grows faster than $\sqrt{\log(n)}$ and τ be a sufficiently large constant. There exists a pair of algorithms (SampleWithTrapdoor, Decode) such that if $(\mathbf{A}, \text{td}) \leftarrow \text{SampleWithTrapdoor}(q, n)$, then \mathbf{A} is of size $n \times m$ with $m = m(q, n) = O(n \cdot \log(q))$ and \mathbf{A} is 2^{-n} close to uniform. For any $\mathbf{s} \in \mathbb{Z}_q^m$ and $\boldsymbol{\eta} \in \mathbb{Z}_q^m$ with $\|\boldsymbol{\eta}\| < \frac{q}{\sqrt{m \cdot \kappa(n)}}$ the algorithm Decode on input td and $\mathbf{s} \cdot \mathbf{A} + \boldsymbol{\eta}$ will output \mathbf{s} .*

2.4 Learning with Errors

The learning with errors (LWE) problem was defined by Regev [Reg05]. In this work we exclusively use the decisional version. The $\text{LWE}_{n, m, q, \chi}$ problem, for $n, m, q \in \mathbb{N}$ and for a distribution χ supported over \mathbb{Z} is to distinguish between the distributions $(\mathbf{A}, \mathbf{s}\mathbf{A} + \mathbf{e} \pmod{q})$ and (\mathbf{A}, \mathbf{u}) , where \mathbf{A} is uniform in $\mathbb{Z}_q^{n \times m}$, \mathbf{s} is a uniform row vector in \mathbb{Z}_q^n , \mathbf{e} is a uniform row vector drawn from χ^m , and \mathbf{u} is a uniform vector in \mathbb{Z}_q^m . Often we consider the hardness of solving LWE for any $m = \text{poly}(n \log q)$. This problem is denoted $\text{LWE}_{n, q, \chi}$. The matrix version of this problem asks to distinguish $(\mathbf{A}, \mathbf{S} \cdot \mathbf{A} + \mathbf{E})$ from (\mathbf{A}, \mathbf{U}) , where $\mathbf{S} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{k \times n}$, $\mathbf{E} \stackrel{\$}{\leftarrow} \chi^{k \times m}$ and $\mathbf{U} \leftarrow \mathbb{Z}_q^{k \times m}$. The hardness of the matrix version for any $k = \text{poly}(n)$ can be established from $\text{LWE}_{n, m, q, \chi}$ via a routine hybrid-argument.

As shown in [Reg05, PRS17], the $\text{LWE}_{n, q, \chi}$ problem with χ being the discrete Gaussian distribution with parameter $\sigma = \alpha q \geq 2\sqrt{n}$ (i.e. the distribution over \mathbb{Z} where the probability of x is proportional to $e^{-\pi(|x|/\sigma)^2}$, see more details below), is at least as hard as approximating the shortest independent vector problem (SIVP) to within a factor of $\gamma = \tilde{O}(n/\alpha)$ in *worst case* dimension n lattices. This is proven using a quantum reduction. Classical reductions (to a slightly different problem) exist as well [Pei09, BLP⁺13] but with somewhat worse parameters. The best known (classical or quantum) algorithms for these problems run in time $2^{\tilde{O}(n/\log \gamma)}$, and in particular they are conjectured to be intractable for $\gamma = \text{poly}(n)$.

3 Lossy Modes for q -Ary Lattices

The following lemmata borrow techniques of the proofs of two lemmata by Chung et al. [CDLP13] (Lemma 3.3 and Lemma 3.4), but are not directly implied by these lemmata. In this section and Section 4, it will be instructive to think of Λ_0 as \mathbb{Z}^n , which will be the case in our application in Section 5.

Lemma 3.1. *Let $\Lambda \subseteq \Lambda_0 \subseteq \mathbb{R}^m$ be full rank lattices and let $T \subseteq \Lambda_0$ be a system of coset representatives of Λ_0/Λ , i.e. we can write every $\mathbf{x} \in \Lambda_0$ as $\mathbf{x} = \mathbf{t} + \mathbf{z}$ for unique $\mathbf{t} \in T$ and $\mathbf{z} \in \Lambda$. Then it holds for any parameter $\sigma > 0$ that*

$$\frac{\rho_\sigma(T)}{\rho_\sigma(\Lambda_0)} \leq \frac{1}{\rho_\sigma(\Lambda)}.$$

Proof. As the $T + \mathbf{y}$ cover Λ_0 it holds that

$$\begin{aligned} \rho_\sigma(\Lambda_0) &= \sum_{\mathbf{y} \in \Lambda} \frac{1}{2} (\rho_\sigma(T + \mathbf{y}) + \rho_\sigma(T - \mathbf{y})) \\ &= \sum_{\mathbf{y} \in \Lambda} \sum_{\mathbf{t} \in T} \frac{1}{2} (\rho_\sigma(\mathbf{t} + \mathbf{y}) + \rho_\sigma(\mathbf{t} - \mathbf{y})) \\ &= \sum_{\mathbf{y} \in \Lambda} \sum_{\mathbf{t} \in T} \rho_\sigma(\mathbf{y}) \cdot \rho_\sigma(\mathbf{t}) \cdot \underbrace{\frac{1}{2} (e^{-2\pi\langle \mathbf{t}, \mathbf{y} \rangle / \sigma^2} + e^{2\pi\langle \mathbf{t}, \mathbf{y} \rangle / \sigma^2})}_{\geq 1} \\ &\geq \sum_{\mathbf{y} \in \Lambda} \rho_\sigma(\mathbf{y}) \sum_{\mathbf{t} \in T} \rho_\sigma(\mathbf{t}) \\ &= \rho_\sigma(\Lambda) \cdot \rho_\sigma(T), \end{aligned}$$

where the first equality follows from the fact that $\sum_{\mathbf{y} \in \Lambda} \rho_\sigma(T + \mathbf{y}) = \sum_{\mathbf{y} \in \Lambda} \rho_\sigma(T - \mathbf{y}) = \rho_\sigma(\Lambda_0)$. The claim follows immediately. \square

Lemma 3.2. *Fix a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ with $m = O(n \log(q))$ and a parameter $0 < \sigma < \frac{q}{2\sqrt{m}}$. Let $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ and $\mathbf{e} \xleftarrow{\$} D_{\mathbb{Z}^m, \sigma}$. Then it holds that $\tilde{H}_\infty(\mathbf{s} | \mathbf{s}\mathbf{A} + \mathbf{e} \bmod q) \geq -\log\left(\frac{1}{\rho_\sigma(\Lambda_q(\mathbf{A}))} + 2^{-m}\right)$.*

Proof. Given arbitrary \mathbf{A} and \mathbf{y} , we would like to find an \mathbf{s}^* that maximizes the probability $\Pr[\mathbf{s} = \mathbf{s}^* | \mathbf{y} = \mathbf{s}\mathbf{A} + \mathbf{e}]$. By Bayes' rule, it holds that

$$\begin{aligned} \Pr[\mathbf{s} = \mathbf{s}^* | \mathbf{y} = \mathbf{s}\mathbf{A} + \mathbf{e}] &= \Pr[\mathbf{y} = \mathbf{s}\mathbf{A} + \mathbf{e} | \mathbf{s} = \mathbf{s}^*] \cdot \frac{\Pr[\mathbf{s} = \mathbf{s}^*]}{\Pr[\mathbf{y} = \mathbf{s}\mathbf{A} + \mathbf{e}]} \\ &= \Pr[\mathbf{e} = \mathbf{y} - \mathbf{s}^*\mathbf{A}] \cdot \frac{\Pr[\mathbf{s} = \mathbf{s}^*]}{\sum_{\mathbf{s}'} \Pr[\mathbf{y} = \mathbf{s}\mathbf{A} + \mathbf{e} | \mathbf{s} = \mathbf{s}'] \Pr[\mathbf{s} = \mathbf{s}']} \\ &= \Pr[\mathbf{e} = \mathbf{y} - \mathbf{s}^*\mathbf{A}] \cdot \frac{q^{-n}}{\sum_{\mathbf{s}'} \Pr[\mathbf{e} = \mathbf{y} - \mathbf{s}'\mathbf{A}] q^{-n}} \\ &= \frac{\Pr[\mathbf{e} = \mathbf{y} - \mathbf{s}^*\mathbf{A}]}{\sum_{\mathbf{s}'} \Pr[\mathbf{e} = \mathbf{y} - \mathbf{s}'\mathbf{A}]}. \end{aligned}$$

As the denominator $\sum_{\mathbf{s}'} \Pr[\mathbf{e} = \mathbf{y} - \mathbf{s}'\mathbf{A}]$ is independent of \mathbf{s}^* , it suffices to maximize the numerator $\Pr[\mathbf{e} = \mathbf{y} - \mathbf{s}^*\mathbf{A}]$ with respect to \mathbf{s}^* . As $\Pr[\mathbf{e} = \mathbf{y} - \mathbf{s}^*\mathbf{A}] = \frac{\rho_\sigma(\mathbf{y} - \mathbf{s}^*\mathbf{A})}{\rho_\sigma(\mathbb{Z}^m)}$ is monotonically decreasing in $\|\mathbf{y} - \mathbf{s}^*\mathbf{A}\|$, this probability is maximal for the \mathbf{s}^* that minimizes $\|\mathbf{y} - \mathbf{s}^*\mathbf{A}\|$.

Let $V \subseteq \mathbb{Z}^m$ be the *discretized Voronoi-cell* of $\Lambda_q(\mathbf{A})$, that is V consists of all the points in \mathbb{Z}^m that are (strictly) closer to 0 than to any other point in Λ and, for any point $\mathbf{x} \in \mathbb{Z}^m$ that is equidistant to several lattice-points $\mathbf{z}_1, \dots, \mathbf{z}_\ell$ (where $\mathbf{z}_1 = 0$), assume that there is some tie-breaking rule $\mathbf{x} \mapsto i(\mathbf{x})$, such that $\mathbf{x} - \mathbf{z}_{i(\mathbf{x})} \in V$, but for all $j \in [\ell] \setminus \{i(\mathbf{x})\}$ it holds that $\mathbf{x} - \mathbf{z}_j \notin V$. By construction, V is a system of coset representatives of $\mathbb{Z}^m/\Lambda_q(\mathbf{A})$.

Moreover, for the maximum-likelihood \mathbf{s}^* it holds that $\Pr[\mathbf{s} = \mathbf{s}^* | \mathbf{y} = \mathbf{s}\mathbf{A} + \mathbf{e}] = \Pr[\mathbf{e} \bmod q \in V]$. By Corollary 2.3 it holds that $\|\mathbf{e}\| \leq \sigma \cdot \sqrt{m} < q/2$, except with probability 2^{-m} . Moreover, conditioned on $\|\mathbf{e}\| < q/2$ the events $\mathbf{e} \bmod q \in V$ and $\mathbf{e} \in V$ are equivalent. We can therefore bound $\Pr[\mathbf{e} \bmod q \in V] \leq \Pr[\mathbf{e} \in V] + 2^{-m}$. By Lemma 3.1 we obtain $\Pr[\mathbf{e} \in V] \leq \frac{\rho_\sigma(V)}{\rho_\sigma(\mathbb{Z}^m)} \leq \frac{1}{\rho_\sigma(\Lambda_q(\mathbf{A}))}$ and therefore $\Pr[\mathbf{e} \bmod q \in V] \leq \frac{1}{\rho_\sigma(\Lambda_q(\mathbf{A}))} + 2^{-m}$.

We conclude that $\max_{\mathbf{s}^* \in \mathbb{Z}_q^n} \Pr[\mathbf{s} = \mathbf{s}^* | \mathbf{y} = \mathbf{s}\mathbf{A} + \mathbf{e}] = \Pr[\mathbf{e} \bmod q \in V] \leq \frac{1}{\rho_\sigma(\Lambda_q(\mathbf{A}))} + 2^{-m}$. Thus, it holds that

$$\begin{aligned} \tilde{H}_\infty(\mathbf{s} | \mathbf{s}\mathbf{A} + \mathbf{e}) &= -\log(\mathbb{E}_{\mathbf{y}} \left[\max_{\mathbf{s}^*} \Pr[\mathbf{s} = \mathbf{s}^* | \mathbf{y} = \mathbf{s}\mathbf{A} + \mathbf{e}] \right]) \\ &= -\log(\mathbb{E}_{\mathbf{y}}[\Pr[\mathbf{e} \bmod q \in V]]) \\ &= -\log(\Pr[\mathbf{e} \bmod q \in V]) \\ &\geq -\log\left(\frac{1}{\rho_\sigma(\Lambda_q(\mathbf{A}))} + 2^{-m}\right) \end{aligned}$$

□

4 Partial Smoothing

In this section we will state a variant of the smoothing lemma of Micciancio and Regev [MR07]. Consider a discrete gaussian $D_{\Lambda_0, \sigma}$ on a lattice Λ_0 . As in the setting of the smoothing Lemma of [MR07], we want to analyze what happens to the distribution of this Gaussian when we reduce it modulo a sublattice $\Lambda \subseteq \Lambda_0$. The new lemma states that if the mass of the Fourier-transform of the probability-mass function of $D_{\Lambda_0, \sigma} \bmod \Lambda$ is concentrated on short vectors of the dual lattice Λ^* , then $D_{\Lambda_0, \sigma} \bmod \Lambda$ will be uniform on a certain sublattice Λ_1 with $\Lambda_0 \subseteq \Lambda_1 \subseteq \Lambda$.

Lemma 4.1. *Let $\sigma > 0$ and let $\Lambda \subseteq \Lambda_0 \subseteq \mathbb{R}^n$ be full-rank lattices where $\det(\Lambda_0) = 1$. Furthermore, let $\gamma > 0$. Define $\Lambda_1 = \{\mathbf{z} \in \Lambda_0 \mid \forall \mathbf{y} \in \Lambda^* \cap \gamma\mathcal{B} : \langle \mathbf{z}, \mathbf{y} \rangle \in \mathbb{Z}\}$. Given that $\rho_{1/\sigma}(\Lambda^* \setminus \gamma\mathcal{B}) \leq \epsilon$, it holds that*

$$\mathbf{x} \bmod \Lambda \approx_\epsilon (\mathbf{x} + \mathbf{u}) \bmod \Lambda,$$

where $\mathbf{x} \stackrel{\$}{\leftarrow} D_{\Lambda_0, \sigma}$ and $\mathbf{u} \stackrel{\$}{\leftarrow} \mathcal{P}(\Lambda) \cap \Lambda_1$.

Notice that for the case of $\Lambda^* \cap \gamma\mathcal{B} = \{0\}$ we recover the standard smoothing lemma of [MR07]. The proof of Lemma 4.1 uses standard Fourier-analytic techniques akin to [MR07] and is deferred to Appendix A. We will make use of the following consequence of Lemma 4.1.

Corollary 4.2. *Let $q > 0$ be an integer and let $\gamma > 0$. Let $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ and let $\sigma > 0$ and $\epsilon > 0$ be such that $\rho_{q/\sigma}(\Lambda_q(\mathbf{A}) \setminus \gamma\mathcal{B}) \leq \epsilon$. Let $\mathbf{D} \in \mathbb{Z}_q^{k \times m}$ be a full-rank (and therefore minimal) matrix with $\Lambda_q^\perp(\mathbf{D}) = \{\mathbf{x} \in \mathbb{Z}^m \mid \forall \mathbf{y} \in \Lambda_q(\mathbf{A}) \cap \gamma\mathcal{B} : \langle \mathbf{x}, \mathbf{y} \rangle = 0 \pmod{q}\}$. Let $\mathbf{x} \stackrel{\$}{\leftarrow} D_{\mathbb{Z}^m, \sigma}$ and $\mathbf{u} \stackrel{\$}{\leftarrow} \Lambda_q^\perp(\mathbf{D}) \bmod q$. Then it holds that*

$$\mathbf{A}\mathbf{x} \bmod q \approx_\epsilon \mathbf{A} \cdot (\mathbf{x} + \mathbf{u}) \bmod q.$$

Proof. Setting $\Lambda_0 = \mathbb{Z}^n$, $\Lambda = \Lambda_q^\perp(\mathbf{A})$ and $\gamma' = \gamma/q$, it holds that $\Lambda^* = \frac{1}{q}\Lambda_q(\mathbf{A})$ and

$$\epsilon \geq \rho_{q/\sigma}(\Lambda_q(\mathbf{A}) \setminus \gamma\mathcal{B}) = \rho_{1/\sigma} \left(\frac{1}{q}\Lambda_q(\mathbf{A}) \setminus \frac{\gamma}{q}\mathcal{B} \right) = \rho_{1/\sigma}(\Lambda^* \setminus \gamma'\mathcal{B}).$$

Therefore, we can set

$$\begin{aligned} \Lambda_1 &= \{\mathbf{x} \in \mathbb{Z}^m \mid \forall \mathbf{y} \in \Lambda^* \cap \gamma'\mathcal{B} : \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}\} \\ &= \{\mathbf{x} \in \mathbb{Z}^m \mid \forall \mathbf{y} \in \Lambda_q(\mathbf{A}) \cap \gamma\mathcal{B} : \langle \mathbf{x}, \mathbf{y} \rangle = 0 \pmod{q}\} \\ &= \Lambda_q^\perp(\mathbf{D}). \end{aligned}$$

Now it holds by Lemma 4.1 as $\mathbf{u} \stackrel{\S}{\leftarrow} \Lambda_q^\perp(\mathbf{D})$ that $\mathbf{x} \bmod \Lambda_q^\perp(\mathbf{A}) \approx_\epsilon (\mathbf{x} + \mathbf{u}) \bmod \Lambda_q^\perp(\mathbf{A})$. Write $\mathbf{y}_1 = \mathbf{x} \bmod \Lambda_q^\perp(\mathbf{A})$ as $\mathbf{y}_1 = \mathbf{x} + \mathbf{z}_1 \bmod q$ for a suitable $\mathbf{z}_1 \in \Lambda_q^\perp(\mathbf{A})$. Likewise, we can write $\mathbf{y}_2 = \mathbf{x} + \mathbf{u} \bmod \Lambda_q^\perp(\mathbf{A})$ as $\mathbf{y}_2 = \mathbf{x} + \mathbf{u} + \mathbf{z}_2 \bmod q$ for a suitable \mathbf{z}_2 . Thus it holds that

$$\mathbf{A}\mathbf{x} = \mathbf{A}(\mathbf{x} + \mathbf{z}_1) \approx_\epsilon \mathbf{A}(\mathbf{x} + \mathbf{u} + \mathbf{z}_2) = \mathbf{A}(\mathbf{x} + \mathbf{u}) \pmod{q}.$$

□

We will also use the following lower bound on the gaussian measure of lattices that have many short linearly independent vectors. The proof of Lemma 4.3 is technically similar to the proof of the transference theorem in [Ban93].

Lemma 4.3. *Let $\Lambda \in \mathbb{R}^m$, $\sigma > 0$ and $\gamma > 0$ be such that $\Lambda \cap \gamma\mathcal{B}$ contains at least k linearly independent vectors. Then it holds that $\rho_\sigma(\Lambda) \geq (\sigma/\gamma)^k$.*

Proof. Let $\Lambda' \subseteq \Lambda$ be the sublattice generated by the vectors in $\Lambda \cap \gamma\mathcal{B}$. Let k be the dimension of the span of Λ' . As $\Lambda' \subseteq \Lambda$, it holds that $\rho_\sigma(\Lambda) \geq \rho_\sigma(\Lambda')$. As Λ' has a basis of length at most γ , we have that $\det(\Lambda') \leq \gamma^k$ and conclude $\det((\Lambda')^*) = 1/\det(\Lambda') \geq \frac{1}{\gamma^k}$. By the Poisson-summation formula, we get that

$$\begin{aligned} \rho_\sigma(\Lambda') &= \sigma^k \cdot \det((\Lambda')^*) \cdot \rho_{1/\sigma}((\Lambda')^*) \\ &\geq (\sigma/\gamma)^k, \end{aligned}$$

as $\rho_{1/\sigma}((\Lambda')^*) \geq 1$. Thus we conclude that $\rho_\sigma(\Lambda) \geq (\sigma/\gamma)^k$. □

5 Our Oblivious Transfer Protocol

We are now ready to provide our statistically sender private oblivious transfer protocol. In the following, let $q, n, \ell = \text{poly}(\lambda)$ and assume that q is of the form $q = 2p$ for an odd p . Let $(\text{SampleWithTrapdoor}, \text{Decode})$ be the pair of algorithms provided in Lemma 2.4 and let $m = m(q, n)$ be such that the matrices \mathbf{A} generated by $\text{SampleWithTrapdoor}(q, 2n)$ are elements of $\mathbb{Z}^{2n \times m}$. Let $\text{Ext}_0 : \{0, 1\}^d \times \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ and $\text{Ext}_1 : \{0, 1\}^d \times \mathbb{Z}_q^{2n} \rightarrow \{0, 1\}^\ell$ be seeded extractors, both with seed-length d and ℓ bits of output. Finally, let $\sigma_0, \sigma_1 > 0$ be parameters for discrete Gaussians and χ be an LWE error-distribution.

The protocol $\text{OT} = (\text{OTR}, \text{OTS}, \text{OTD})$ is given as follows.

- $\text{OTR}(1^\lambda, \beta \in \{0, 1\})$:
 - If $\beta = 0$, choose a matrix $\mathbf{A}_1 \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$, a matrix $\mathbf{S} \leftarrow \mathbb{Z}_q^{n \times n}$, $\mathbf{E} \xleftarrow{\$} \chi^{n \times m}$. Set $\mathbf{A}_2 \leftarrow \mathbf{S} \cdot \mathbf{A}_1 + \mathbf{E}$ and $\mathbf{A} \leftarrow \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix}$. Repeat this step until $\mathbf{A} \bmod 2$ has full rank. Output $\text{ot}_1 \leftarrow \mathbf{A}$ and $\text{st} \leftarrow \mathbf{S}$.
 - If $\beta = 1$, sample $(\mathbf{A}, \text{td}) \xleftarrow{\$} \text{SampleWithTrapdoor}(q, 2n)$. Repeat this step until $\mathbf{A} \bmod 2$ has full rank. Output $\text{ot}_1 \leftarrow \mathbf{A}$ and $\text{st} \leftarrow \text{td}$.
- $\text{OTS}(1^\lambda, (\mu_0, \mu_1) \in (\{0, 1\}^\ell)^2, \text{ot}_1 = \mathbf{A})$:
 - Check if $\mathbf{A} \bmod 2$ has full rank, if not output \perp .
 - Parse $\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix}$. Sample and reject a discrete Gaussian $\mathbf{x} \xleftarrow{\$} D_{\mathbb{Z}^m, \sigma_0}$ until $\|\mathbf{x}\| < \sigma_0 \sqrt{m}$. Choose a uniformly random $\mathbf{r} \leftarrow \{0, 1\}^n$ and choose a random seed $\mathbf{s}_0 \xleftarrow{\$} \{0, 1\}^d$ for the extractor Ext_0 . Compute $\mathbf{y}_1 \leftarrow \mathbf{A}_1 \mathbf{x}$ and $\mathbf{y}_2 \leftarrow \mathbf{A}_2 \mathbf{x} + \frac{q}{2} \cdot \mathbf{r}$. Set $\mathbf{c}_0 \leftarrow (\mathbf{y}_1, \mathbf{y}_2, \mathbf{s}_0, \text{Ext}_0(\mathbf{s}_0, \mathbf{r}) \oplus \mu_0)$.
 - Sample and reject $\boldsymbol{\eta} \xleftarrow{\$} D_{\mathbb{Z}^m, \sigma_1}$ until $\|\boldsymbol{\eta}\| < \sigma_1 \sqrt{m}$. Choose a uniformly random $\mathbf{t} \xleftarrow{\$} \mathbb{Z}_q^{2n}$ and a seed $\mathbf{s}_1 \xleftarrow{\$} \{0, 1\}^d$ for the extractor Ext_1 . Compute $\mathbf{y} \leftarrow \mathbf{t} \cdot \mathbf{A} + \boldsymbol{\eta}$ set $\mathbf{c}_1 \leftarrow (\mathbf{y}, \mathbf{s}_1, \text{Ext}_1(\mathbf{s}_1, \mathbf{t}) \oplus \mu_1)$.
 - Output $\text{ot}_2 \leftarrow (\mathbf{c}_0, \mathbf{c}_1)$.
- $\text{OTD}(\beta, \text{st}, \text{ot}_2 = (\mathbf{c}_0, \mathbf{c}))$
 - If $\beta = 0$: Parse $\text{st} = \mathbf{S}$ and $\mathbf{c}_0 = (\mathbf{y}_1, \mathbf{y}_2, \mathbf{s}_0, \tau)$. Compute $\mathbf{r}' \leftarrow \lfloor \mathbf{y}_2 - \mathbf{S} \cdot \mathbf{y}_1 \rfloor_{q/2}$ and output $\mu'_0 \leftarrow \text{Ext}_0(\mathbf{s}_0, \mathbf{r}') \oplus \tau$.
 - If $\beta = 1$: Parse $\text{st} = \text{td}$ and $\mathbf{c}_1 = (\mathbf{y}, \mathbf{s}_1, \tau)$. Compute $\mathbf{t}' \leftarrow \text{Decode}(\text{td}, \mathbf{y})$ and output $\mu'_1 \leftarrow \text{Ext}_1(\mathbf{s}_1, \mathbf{t}') \oplus \tau$.

We will first show correctness of our protocol.

Lemma 5.1 (Correctness). *Assume that the distribution χ is a B -bounded. Provided that $\sigma_0 \leq \frac{q}{4B \cdot m}$ and $\sigma_1 \leq \frac{q}{m \cdot \kappa(n)}$ (where $\kappa(n) = \omega(\sqrt{\log(n)})$ as in Lemma 2.4), the protocol OT is perfectly correct.*

Proof. First note that as $m \geq n \cdot \log(q)$, it holds for a uniformly random $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{2n \times m}$ that $\mathbf{A} \bmod 2$ has full rank, except with negligible probability 2^{n-m} (as detailed in Section 2.2). Moreover for $\mathbf{x} \leftarrow D_{\mathbb{Z}^m, \sigma_0}$ and $\boldsymbol{\eta} \xleftarrow{\$} D_{\mathbb{Z}^m, \sigma_1}$ it holds by Corollary 2.3 that $\|\mathbf{x}\| < \sigma_0 \sqrt{m}$ and $\|\boldsymbol{\eta}\| < \sigma_1 \sqrt{m}$, except with negligible probability. Thus, rejection in OTR and OTS happens only with negligible probability.

In the case of $\beta = 0$, it holds that

$$\begin{aligned} \mathbf{y}_2 - \mathbf{S} \cdot \mathbf{y}_1 &= (\mathbf{S}\mathbf{A}_1 + \mathbf{E})\mathbf{x} + \frac{q}{2}\mathbf{r} - \mathbf{S}\mathbf{A}_1\mathbf{x} \\ &= \mathbf{E} \cdot \mathbf{x} + \frac{q}{2} \cdot \mathbf{r}. \end{aligned}$$

By the Cauchy-Schwarz inequality it holds for each row \mathbf{e}_i of \mathbf{E} that $|\langle \mathbf{e}_i, \mathbf{x} \rangle| \leq \|\mathbf{e}_i\| \cdot \|\mathbf{x}\|$. As the entries of \mathbf{e}_i are chosen according to χ , we can bound $\|\mathbf{e}_i\|$ by $\|\mathbf{e}_i\| \leq B \cdot \sqrt{m}$. As $\|\mathbf{x}\| < \sigma_0 \cdot \sqrt{m}$, we have that

$$|\langle \mathbf{e}_i, \mathbf{x} \rangle| \leq B \cdot \sigma_0 \cdot m < \frac{q}{4}$$

as $\sigma_0 \leq \frac{q}{4B \cdot m}$. We conclude that $\mathbf{r}' = \lfloor \mathbf{y}_2 - \mathbf{S} \cdot \mathbf{y}_1 \rfloor_{q/2}$ is identical to the vector \mathbf{r} used during encryption. Consequently, it holds that $\mu'_0 = \text{Ext}_0(\mathbf{s}_0, \mathbf{r}') \oplus \tau = \text{Ext}_0(\mathbf{s}_0, \mathbf{r}') \oplus \text{Ext}_0(\mathbf{s}_0, \mathbf{r}) \oplus \mu_0 = \mu_0$.

For the case of $\beta = 1$, as $\|\boldsymbol{\eta}\| < \sigma_1 \sqrt{m} \leq \frac{q}{\sqrt{m \cdot \kappa(n)}}$ it holds by Lemma 2.4 that $\text{Decode}(\text{td}, \mathbf{y}_1)$ outputs the correct $\mathbf{t}' = \mathbf{t}$. We conclude that $\mu'_1 = \text{Ext}_1(\mathbf{s}_1, \mathbf{t}') \oplus \tau = \text{Ext}_1(\mathbf{s}_1, \mathbf{t}') \oplus \text{Ext}_1(\mathbf{s}_1, \mathbf{t}) \oplus \mu = \mu_1$. \square

We now show that OT has computational receiver privacy under the decisional matrix LWE assumption.

Lemma 5.2 (Computational Receiver Security). *Given that the decisional $LWE_{n,q,\chi}$ -assumption holds, the protocol $\text{OT} = (\text{OTR}, \text{OTS}, \text{OTD})$ has receiver privacy.*

Proof. Let $(\mathbf{A}, \mathbf{st}_0) \leftarrow \text{OTR}(1^\lambda, 0)$ and $(\mathbf{A}', \mathbf{st}_1) \leftarrow \text{OTR}(1^\lambda, 1)$. Assume towards contradiction that there exists a PPT-distinguisher \mathcal{D} which distinguishes \mathbf{A} and \mathbf{A}' with non-negligible advantage ϵ . We can immediately use \mathcal{D} to distinguish decisional matrix LWE. Decomposing $\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix}$, it holds that \mathbf{A}_1 is uniformly random and $\mathbf{A}_2 = \mathbf{S} \cdot \mathbf{A}_1 + \mathbf{E}$, i.e. $(\mathbf{A}_1, \mathbf{A}_2)$ is a sample of the matrix LWE distribution. On the other hand, due to the uniformity property of $\text{SampleWithTrapdoor}$ (provided in Lemma 2.4) it holds that $\mathbf{A}' \approx_s \mathbf{A}^*$ for a uniformly random $\mathbf{A}^* \xleftarrow{\$} \mathbb{Z}_q^{2n \times m}$. Consequently

$$\begin{aligned} \text{Adv}_{LWE}(\mathcal{D}) &= |\Pr[\mathcal{D}(\mathbf{A}) = 1] - \Pr[\mathcal{D}(\mathbf{A}^*) = 1]| \\ &\geq |\Pr[\mathcal{D}(\mathbf{A}) = 1] - \Pr[\mathcal{D}(\mathbf{A}') = 1]| - |\Pr[\mathcal{D}(\mathbf{A}^*) = 1] - \Pr[\mathcal{D}(\mathbf{A}') = 1]| \\ &\geq \epsilon - \text{negl}, \end{aligned}$$

which contradicts the hardness of decisional matrix LWE. \square

We will now show that OT is statistically sender-private.

Theorem 5.3 (Statistical Sender Security). *Let $q = 2p$ for an odd p . Given that $\sigma_0 \cdot \sigma_1 \geq 4\sqrt{m} \cdot q$, $\sigma_1 < \frac{q}{2\sqrt{m}}$ and both Ext_0 and Ext_1 are strong average-case $(n/2, \text{negl})$ -extractors, then the above scheme enjoys statistical sender security.*

Proof. Fix a maliciously generated ot_1 -message $\text{ot}_1 = \mathbf{A}$. Let in the following $\gamma := \sqrt{m} \cdot \frac{q}{\sigma_0}$. Consider the following two cases.

1. $\rho_{q/\sigma_0}(\Lambda_q(\mathbf{A})) > 2^{n/2+1}$ or $\text{rank}(\Lambda_q(\mathbf{A}) \cap \gamma\mathcal{B}) > n/2$.
2. $\rho_{q/\sigma_0}(\Lambda_q(\mathbf{A})) \leq 2^{n/2+2}$ and $\text{rank}(\Lambda_q(\mathbf{A}) \cap \gamma\mathcal{B}) \leq n/2$.

First notices that the two cases are slightly overlapping, but for any choice of \mathbf{A} one of the two cases must be true.

The unbounded message extractor OTExt takes input \mathbf{A} and decides if item 1 or item 2 holds. If item 1 holds it outputs 0, otherwise 1. Note that $\text{rank}(\Lambda_q(\mathbf{A}) \cap \gamma\mathcal{B})$ can be computed exactly. On the other hand, it is sufficient approximate $\rho_{q/\sigma_0}(\Lambda_q(\mathbf{A}))$ to a certain precision to determine which case holds.

We will now show that in case 1 the sender-message μ_1 is statistically hidden, whereas in case 2 the sender-message μ_0 is statistically hidden.

Case 1. We will start with the (easier) first case. We will show that either statement implies $\rho_{\sigma_1}(\Lambda_q(\mathbf{A})) \geq 2^{n/2+1}$. If it holds that $\rho_{q/\sigma_0}(\Lambda_q(\mathbf{A})) > 2^{n/2+1}$, we can directly conclude that

$$\rho_{\sigma_1}(\Lambda_q(\mathbf{A})) \geq \rho_{4\sqrt{m} \cdot \frac{q}{\sigma_0}}(\Lambda_q(\mathbf{A})) \geq \rho_{\frac{q}{\sigma_0}}(\Lambda_q(\mathbf{A})) > 2^{n/2+1}.$$

If the second statement $\text{rank}(\Lambda_q(\mathbf{A}) \cap \gamma\mathcal{B}) > n/2$ holds, Lemma 4.3 implies

$$\rho_{\sigma_1}(\Lambda_q(\mathbf{A})) \geq (\sigma_1/\gamma)^{n/2+1} \geq 2^{n+2} \geq 2^{n/2+1},$$

as $\sigma_1 \geq 4\gamma$.

Now let $\mathbf{c}_1 = (\mathbf{y}, \mathbf{s}_1, \tau)$, where $\mathbf{y} \leftarrow \mathbf{t} \cdot \mathbf{A} + \boldsymbol{\eta}$. Note that we can switch to a hybrid in which the distribution of $\boldsymbol{\eta}$ is $D_{\mathbb{Z}^m, \sigma_1}$ instead of the truncated version while only incurring a negligible statistical error.

As $\rho_{\sigma_1}(\Lambda_q(\mathbf{A})) \geq 2^{n/2+1}$ and $\sigma_1 < \frac{q}{2\sqrt{m}}$, Lemma 3.2 implies that

$$\tilde{H}_\infty(\mathbf{t}|\mathbf{y}) \geq -\log(1/\rho_{\sigma_1}(\Lambda_q(\mathbf{A})) + 2^{-m}) \geq -\log(2^{-n/2-1} + 2^{-m}) \geq n/2$$

Thus, as Ext_1 is a strong $(n/2, \text{negl})$ -extractor, we get that $\text{Ext}_1(\mathbf{s}_1, \mathbf{t})$ is statistically close to uniform given \mathbf{y} . Consequently, $\tau = \text{Ext}_1(\mathbf{s}_1, \mathbf{t}) \oplus \mu_1$ is statistically close to uniform given \mathbf{s}_1 and \mathbf{y} , which concludes the first case.

Case 2. We will now turn to the second case, i.e. it holds that $\rho_{q/\sigma_0}(\Lambda_q(\mathbf{A})) \leq 2^{n/2+2}$ and $\text{rank}(\Lambda_q(\mathbf{A}) \cap \gamma\mathcal{B}) \leq n/2$. Theorem 2.2 yields that

$$\begin{aligned} \rho_{q/\sigma_0}(\Lambda_q(\mathbf{A}) \setminus \gamma\mathcal{B}) &= \rho_{q/\sigma_0} \left(\Lambda_q(\mathbf{A}) \setminus \frac{\sqrt{m} \cdot q}{\sigma_0} \mathcal{B} \right) \\ &\leq 2^{-C \cdot m} \cdot \rho_{q/\sigma_0}(\Lambda_q(\mathbf{A})) \\ &\leq 2^{-C \cdot m} \cdot 2^{n/2+2} = 2^{n/2+2-C \cdot m} \end{aligned}$$

where $C > 0$ is a constant. This expression is negligible as $m \geq n \cdot \log(q)$. Consequently, the precondition $\rho_{q/\sigma_0}(\Lambda_q(\mathbf{A}) \setminus \gamma\mathcal{B}) \leq \text{negl}$ of Corollary 4.2 is fulfilled.

Now let $\mathbf{D} \in \mathbb{Z}_q^{k \times m}$ be a full-rank matrix with $\Lambda_q^\perp(\mathbf{D}) = \{\mathbf{z} \in \mathbb{Z}^m \mid \forall \mathbf{v} \in \Lambda_q(\mathbf{A}) \cap \gamma\mathcal{B} : \langle \mathbf{z}, \mathbf{v} \rangle = 0 \pmod{q}\}$. Thus it holds that $\Lambda_q(\mathbf{A}) \cap \gamma\mathcal{B} \subset \Lambda_q(\mathbf{D})$ and there is no matrix with fewer than k rows with this property. As $\text{rank}(\Lambda_q(\mathbf{A}) \cap \gamma\mathcal{B}) \leq n/2$, it holds that $k \leq n/2$.

Decompose the matrix \mathbf{A} into $\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix}$ with $\mathbf{A}_1 \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{A}_2 \in \mathbb{Z}_q^{n \times m}$. Let $\mathbf{c}_0 = (\mathbf{y}_1, \mathbf{y}_2, \mathbf{s}_0, \tau)$, where $\mathbf{y}_1 = \mathbf{A}_1 \mathbf{x}$ and $\mathbf{y}_2 = \mathbf{A}_2 \mathbf{x} + \frac{q}{2} \mathbf{r}$ with $\mathbf{x} \stackrel{\$}{\leftarrow} D_{\mathbb{Z}^m, \sigma_0}$ and $\mathbf{r} \stackrel{\$}{\leftarrow} \{0, 1\}^n$. As $\rho_{q/\sigma_0}(\Lambda_q(\mathbf{A}) \setminus \gamma\mathcal{B}) \leq \epsilon$, Corollary 4.2 implies that

$$(\mathbf{y}_1, \mathbf{y}_2) = (\mathbf{A}_1 \mathbf{x}, \mathbf{A}_2 \mathbf{x} + \frac{q}{2} \mathbf{r}) \approx_\epsilon (\mathbf{A}_1(\mathbf{x} + \mathbf{u}), \mathbf{A}_2(\mathbf{x} + \mathbf{u}) + \frac{q}{2} \mathbf{r}) =: (\mathbf{y}'_1, \mathbf{y}'_2)$$

where $\mathbf{u} \stackrel{\$}{\leftarrow} \Lambda_q^\perp(\mathbf{D})$. We can therefore switch to a hybrid experiment in which we replace \mathbf{x} with $\mathbf{x} + \mathbf{u}$ while only incurring negligible statistical distance. We will now show that $\tilde{H}_\infty(\mathbf{r}|\mathbf{y}'_1, \mathbf{y}'_2) \geq n/2$.

As $q = 2p$ and p is odd, it holds by the Chinese remainder theorem that

$$\begin{aligned} \mathbf{y}'_1 &\equiv (\mathbf{A}_1(\mathbf{x} + \mathbf{u}) \pmod{2}, \mathbf{A}_1(\mathbf{x} + \mathbf{u}) \pmod{p}) \\ \mathbf{y}'_2 &\equiv (\mathbf{A}_2(\mathbf{x} + \mathbf{u}) + \mathbf{r} \pmod{2}, \mathbf{A}_1(\mathbf{x} + \mathbf{u}) \pmod{p}) \end{aligned}$$

Note that $\mathbf{u} \bmod 2$ and $\mathbf{u} \bmod p$ are independent. As the $\bmod p$ part does not depend on \mathbf{r} , we only need to consider the $\bmod 2$ part. Let in the following variables with a hat denote this variable is reduced modulo 2, e.g. $\hat{\mathbf{x}} = \mathbf{x} \bmod 2$. It holds that $\hat{\mathbf{u}}$ is chosen uniformly from $\ker(\hat{\mathbf{D}}) = \{\mathbf{w} \in \mathbb{Z}_2^m \mid \hat{\mathbf{D}} \cdot \mathbf{w} = 0\}$. The dimension of $\ker(\hat{\mathbf{D}})$ is at least $m - k \geq m - n/2$. Let $\hat{\mathbf{B}} \in \mathbb{Z}_2^{m \times m}$ be a basis of $\ker(\hat{\mathbf{D}})$. As $\hat{\mathbf{A}}$ has full rank and therefore $\text{rank}(\ker(\hat{\mathbf{A}})) = m - 2n$, it holds that $\text{rank}(\hat{\mathbf{A}} \cdot \hat{\mathbf{B}}) \geq \frac{3}{2}n$. Therefore $\hat{\mathbf{A}} \cdot \hat{\mathbf{u}}$ is uniformly random in an $\frac{3}{2}n$ dimensional subspace. But this means that $(\hat{\mathbf{y}}'_1, \hat{\mathbf{y}}'_2) = (\hat{\mathbf{A}}_1 \hat{\mathbf{x}} + \hat{\mathbf{A}}_1 \hat{\mathbf{u}}, \hat{\mathbf{A}}_2 \hat{\mathbf{x}} + \hat{\mathbf{A}}_2 \hat{\mathbf{u}} + \mathbf{r})$ loses at least $n/2$ bits of information about \mathbf{r} (c.f. Section 2.2). Consequently, it holds that $\tilde{H}_\infty(\mathbf{r} | \mathbf{y}'_1, \mathbf{y}'_2) \geq n/2$. Therefore, as Ext_0 is a strong $(n/2, \text{negl})$ -extractor, we get that $\text{Ext}_0(\mathbf{s}_0, \mathbf{r})$ is statistically close to uniform given $\mathbf{y}'_1, \mathbf{y}'_2$. Finally, $\tau = \text{Ext}_0(\mathbf{s}_0, \mathbf{r}) \oplus \mu_0$ is statistically close to uniform given \mathbf{s}_0 and $\mathbf{y}'_1, \mathbf{y}'_2$, which concludes the second case. \square

5.1 Setting the Parameters

We will now show that the parameters of the scheme can be chosen such that correctness, statistical sender privacy and computational receiver privacy hold.

- By Lemma 5.1, OT is correct if $\sigma_0 \leq \frac{q}{4B \cdot m}$ and $\sigma_1 \leq \frac{q}{m \cdot \kappa(n)}$ (where $\kappa(n) = \omega(\sqrt{\log(n)})$).
- By Theorem 5.3, OT is statistically sender private if $\sigma_0 \cdot \sigma_1 \geq 4\sqrt{m} \cdot q$ and $\sigma_1 < \frac{q}{2\sqrt{m}}$.

These requirements can be met if

$$\frac{q^2}{4\kappa(n)Bm^2} \geq 4\sqrt{m} \cdot q,$$

which is equivalent to

$$q \geq 16\kappa(n) \cdot B \cdot m^{2.5}. \tag{1}$$

If χ is a discrete Gaussian on \mathbb{Z} with parameter αq , i.e. $\chi = D_{\mathbb{Z}, \alpha q}$, then, given that $\alpha q \geq \eta_\epsilon(\mathbb{Z}) = \omega(\sqrt{\log(n)})$ it holds that χ is αq bounded, i.e. $B \leq \alpha q$ (with overwhelming probability). This means that

$$\alpha \leq \frac{1}{16 \cdot \kappa(n) m^{2.5}} = \tilde{O}(n^{-2.5})$$

implies inequality (1). Thus, we get a worst-case approximation factor $\tilde{O}(n/\alpha) = \tilde{O}(n^{3.5})$ for SIVP (compared to $\tilde{O}(n^{1.5})$ for primal Regev encryption). With this choice of α , we can choose $q = \tilde{O}(n^3)$, $\sigma_0 = \tilde{O}(n^{2.5})$ and $\sigma_1 = \tilde{O}(n)$.

Acknowledgments

We would like to thank the anonymous TCC 2018 reviewers for insightful comments that helped to improve the presentation of the paper.

References

- [AIR01] William Aiello, Yuval Ishai, and Omer Reingold. Priced oblivious transfer: How to sell digital goods. In Birgit Pfitzmann, editor, *Advances in Cryptology - EUROCRYPT 2001, International Conference on the Theory and Application of Cryptographic Techniques, Innsbruck, Austria, May 6-10, 2001, Proceeding*, volume 2045 of *Lecture Notes in Computer Science*, pages 119–135. Springer, 2001.
- [Ajt99] Miklós Ajtai. Generating hard instances of the short basis problem. In *ICALP*, volume 1644 of *Lecture Notes in Computer Science*, pages 1–9. Springer, 1999.
- [Ban93] Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(1):625–635, 1993.
- [BGI⁺17] Saikrishna Badrinarayanan, Sanjam Garg, Yuval Ishai, Amit Sahai, and Akshay Wadia. Two-message witness indistinguishability and secure computation in the plain model from new assumptions. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part III*, volume 10626 of *Lecture Notes in Computer Science*, pages 275–303. Springer, 2017.
- [BGJ⁺17a] Saikrishna Badrinarayanan, Vipul Goyal, Abhishek Jain, Yael Tauman Kalai, Dakshita Khurana, and Amit Sahai. Promise zero knowledge and its applications to round optimal MPC. *IACR Cryptology ePrint Archive*, 2017:1088, 2017.
- [BGJ⁺17b] Saikrishna Badrinarayanan, Vipul Goyal, Abhishek Jain, Dakshita Khurana, and Amit Sahai. Round optimal concurrent MPC via strong simulation. In Yael Kalai and Leonid Reyzin, editors, *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part I*, volume 10677 of *Lecture Notes in Computer Science*, pages 743–775. Springer, 2017.
- [BHP17] Zvika Brakerski, Shai Halevi, and Antigoni Polychroniadou. Four round secure computation without setup. In Yael Kalai and Leonid Reyzin, editors, *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part I*, volume 10677 of *Lecture Notes in Computer Science*, pages 645–677. Springer, 2017.
- [BLP⁺13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 575–584. ACM, 2013.
- [BM89] Mihir Bellare and Silvio Micali. Non-interactive oblivious transfer and applications. In *Conference on the Theory and Application of Cryptology*, pages 547–557. Springer, 1989.
- [CDLP13] Kai-Min Chung, Daniel Dadush, Feng-Hao Liu, and Chris Peikert. On the lattice smoothing parameter problem. In *IEEE Conference on Computational Complexity*, pages 230–241. IEEE Computer Society, 2013.

- [DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam D. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008.
- [DRS04] Yevgeniy Dodis, Leonid Reyzin, and Adam D. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In *EUROCRYPT*, volume 3027 of *Lecture Notes in Computer Science*, pages 523–540. Springer, 2004.
- [GG98] Oded Goldreich and Shafi Goldwasser. On the limits of non-approximability of lattice problems. In Jeffrey Scott Vitter, editor, *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 1–9. ACM, 1998.
- [GKM⁺00] Yael Gertner, Sampath Kannan, Tal Malkin, Omer Reingold, and Mahesh Viswanathan. The relationship between public key encryption and oblivious transfer. In *FOCS*, pages 325–335. IEEE Computer Society, 2000.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred V. Aho, editor, *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 218–229. ACM, 1987.
- [GO94] Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, 1994.
- [HHPV17] Shai Halevi, Carmit Hazay, Antigoni Polychroniadou, and Muthuramakrishnan Venkatasubramanian. Round-optimal secure multi-party computation. *IACR Cryptology ePrint Archive*, 2017:1056, 2017.
- [HK12] Shai Halevi and Yael Tauman Kalai. Smooth projective hashing and two-message oblivious transfer. *J. Cryptology*, 25(1):158–193, 2012.
- [JKKR17] Abhishek Jain, Yael Tauman Kalai, Dakshita Khurana, and Ron Rothblum. Distinguisher-dependent simulation in two rounds and its applications. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part II*, volume 10402 of *Lecture Notes in Computer Science*, pages 158–189. Springer, 2017.
- [Kal05] Yael Tauman Kalai. Smooth projective hashing and two-message oblivious transfer. In Ronald Cramer, editor, *Advances in Cryptology - EUROCRYPT 2005, 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22-26, 2005, Proceedings*, volume 3494 of *Lecture Notes in Computer Science*, pages 78–95. Springer, 2005.
- [Khu17] Dakshita Khurana. Round optimal concurrent non-malleability from polynomial hardness. In Yael Kalai and Leonid Reyzin, editors, *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part II*, volume 10678 of *Lecture Notes in Computer Science*, pages 139–171. Springer, 2017.

- [KKS18] Yael Tauman Kalai, Dakshita Khurana, and Amit Sahai. Statistical witness indistinguishability (and more) in two messages. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part III*, volume 10822 of *Lecture Notes in Computer Science*, pages 34–65. Springer, 2018.
- [KS17] Dakshita Khurana and Amit Sahai. How to achieve non-malleability in one or two rounds. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 564–575. IEEE Computer Society, 2017.
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, volume 7237 of *Lecture Notes in Computer Science*, pages 700–718. Springer, 2012.
- [MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007.
- [NP01] Moni Naor and Benny Pinkas. Efficient oblivious transfer protocols. In S. Rao Kosaraju, editor, *Proceedings of the Twelfth Annual Symposium on Discrete Algorithms, January 7-9, 2001, Washington, DC, USA.*, pages 448–457. ACM/SIAM, 2001.
- [OPP14] Rafail Ostrovsky, Anat Paskin-Cherniavsky, and Beni Paskin-Cherniavsky. Maliciously circuit-private FHE. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 536–553. Springer, 2014.
- [Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *STOC*, pages 333–342. ACM, 2009.
- [PRS17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-lwe for any ring and modulus. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 461–473. ACM, 2017.
- [PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In David A. Wagner, editor, *Advances in Cryptology - CRYPTO 2008, 28th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2008. Proceedings*, volume 5157 of *Lecture Notes in Computer Science*, pages 554–571. Springer, 2008.
- [Rab81] Michael O. Rabin. How to exchange secrets with oblivious transfer. *Harvard University Technical Report*, 1981.

- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *STOC*, pages 84–93. ACM, 2005. Full version in [Reg09].
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.
- [Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *FOCS*, pages 162–167, 1986.

A Appendix

In this Section we will provide the proof for Lemma 4.1. We will first provide some additional preliminaries.

A.1 Additional Preliminaries

Fourier Transforms We now recall a few basic facts about Fourier-transforms on lattices. Let $f : \mathbb{R}^m \rightarrow \mathbb{C}$ and Λ be a lattice, if it exists, we will write $f(\Lambda) := \sum_{\mathbf{x} \in \Lambda} f(\mathbf{x})$. For a *nice enough*⁴ function $f : \mathbb{R}^m \rightarrow \mathbb{C}$, we define the continuous Fourier-transform $\hat{f} : \mathbb{R}^m \rightarrow \mathbb{C}$ by $\hat{f}(\boldsymbol{\omega}) = \int_{\mathbf{x} \in \mathbb{R}^m} f(\mathbf{x}) \cdot e^{-2\pi i \langle \boldsymbol{\omega}, \mathbf{x} \rangle} d\mathbf{x}$. The Poisson summation formula states that $f(\Lambda) = \det(\Lambda^*) \cdot \hat{f}(\Lambda^*)$. The Fourier-transform of the Gaussian function $\rho_\sigma(\mathbf{x})$ is $\sigma^m \cdot \rho_{1/\sigma}(\boldsymbol{\omega})$. Consequently, we get by the Poisson summation formula that

$$\rho_\sigma(\Lambda) = \sigma^m \cdot \det(\Lambda^*) \cdot \rho_{1/\sigma}(\Lambda^*).$$

Fix a full-rank lattice $\Lambda_0 \subseteq \mathbb{R}^m$ and assume henceforth that $\Lambda \subseteq \Lambda_0$. We say a function $f : \Lambda_0 \rightarrow \mathbb{C}$ is Λ -periodic if it holds for all $\mathbf{x} \in \Lambda_0$ and all $\mathbf{z} \in \Lambda$ that $f(\mathbf{x} + \mathbf{z}) = f(\mathbf{x})$. Now let $f : \Lambda_0 \rightarrow \mathbb{C}$ be a Λ -periodic function. We define the discrete Fourier transform $\hat{f} : \Lambda^* \rightarrow \mathbb{C}$ of f by

$$\hat{f}(\boldsymbol{\omega}) = \sum_{\mathbf{x} \in \mathcal{P}(\Lambda) \cap \Lambda_0} f(\mathbf{x}) \cdot e^{-2\pi i \langle \boldsymbol{\omega}, \mathbf{x} \rangle}.$$

Here, $\mathcal{P}(\Lambda) \cap \Lambda_0$ can be replaced by any system of representatives for the quotient group Λ_0/Λ . Using Fourier-inversion, we can express f as

$$f(\mathbf{x}) = \frac{\det \Lambda_0}{\det \Lambda} \cdot \sum_{\boldsymbol{\omega} \in \mathcal{P}(\Lambda_0^*) \cap \Lambda^*} \hat{f}(\boldsymbol{\omega}) \cdot e^{2\pi i \langle \boldsymbol{\omega}, \mathbf{x} \rangle}.$$

Note that \hat{f} is Λ_0^* periodic.

Let \mathbf{x} and \mathbf{y} be random variables defined on Λ_0/Λ . Let the probability-mass function of the distribution of \mathbf{x} be given by a Λ -periodic function $X : \Lambda_0 \rightarrow \mathbb{R}$, and let the probability-mass function of \mathbf{y} be given by a Λ -periodic function $Y : \Lambda_0 \rightarrow \mathbb{R}$. Finally, let $Z : \Lambda_0 \rightarrow \mathbb{R}$ be the probability mass function of $\mathbf{x} + \mathbf{y}$. The convolution theorem states that it holds that

$$\hat{Z}(\boldsymbol{\omega}) = \hat{X}(\boldsymbol{\omega}) \cdot \hat{Y}(\boldsymbol{\omega}),$$

⁴where *nice enough* means that $\int_{\mathbf{x} \in \mathbb{R}^m} |f(\mathbf{x})| d\mathbf{x}$ is finite

for all $\boldsymbol{\omega} \in \Lambda^*$.

If \mathbf{x} is distributed according to a discrete Gaussian $D_{\Lambda_0, \sigma}$ and $\Lambda \subseteq \Lambda_0$, then $\mathbf{x} \bmod \Lambda$ has the probability-mass function of a periodic gaussian given by

$$\psi_\sigma(\mathbf{x}') = \Pr[\mathbf{x} = \mathbf{x}'] = \frac{1}{\rho_\sigma(\Lambda_0)} \cdot \sum_{\mathbf{z} \in \Lambda} \rho_\sigma(\mathbf{x}' + \mathbf{z})$$

and it holds that

$$\widehat{\psi}_\sigma(\boldsymbol{\omega}) = \frac{1}{\det(\Lambda_0^*) \cdot \rho_{1/\sigma}(\Lambda_0^*)} \cdot \sum_{\boldsymbol{\xi} \in \Lambda_0^*} \rho_{1/\sigma}(\boldsymbol{\omega} + \boldsymbol{\xi})$$

for $\boldsymbol{\omega} \in \Lambda^*$.

We define the Dirac-function $\delta : \Lambda_0 \rightarrow \mathbb{R}$ as $\delta(0) = 1$ and $\delta(\mathbf{x}) = 0$ for $\mathbf{x} \neq 0$. If $\Lambda \subseteq \Lambda_1 \subset \Lambda_0$ and \mathbf{u} is distributed uniformly random on $\mathcal{P}(\Lambda) \cap \Lambda_1$, then \mathbf{u} has the probability-mass function

$$U(\mathbf{x}) = \frac{\det \Lambda_1}{\det \Lambda} \sum_{\mathbf{y} \in \Lambda_1} \delta(\mathbf{x} + \mathbf{y})$$

and the Fourier-transform

$$\widehat{U}(\boldsymbol{\omega}) = \sum_{\boldsymbol{\xi} \in \Lambda_1^*} \delta(\boldsymbol{\omega} + \boldsymbol{\xi}).$$

A.2 Proof of the Partial Smoothing Lemma

Lemma A.1. *Let $\sigma > 0$ and let $\Lambda \subseteq \Lambda_0 \subseteq \mathbb{R}^n$ be full-rank lattices where $\det(\Lambda_0) = 1$. Furthermore, let $\gamma > 0$. Define $\Lambda_1 = \{\mathbf{z} \in \Lambda_0 \mid \forall \mathbf{y} \in \Lambda^* \cap \gamma\mathcal{B} : \langle \mathbf{y}, \mathbf{z} \rangle \in \mathbb{Z}\}$. Given that $\rho_{1/\sigma}(\Lambda^* \setminus \gamma\mathcal{B}) \leq \epsilon$, it holds that*

$$\mathbf{y} \bmod \Lambda \approx_\epsilon (\mathbf{y} + \mathbf{u}) \bmod \Lambda,$$

where $\mathbf{y} \stackrel{\$}{\leftarrow} D_{\Lambda_0, \sigma}$ and $\mathbf{u} \stackrel{\$}{\leftarrow} \mathcal{P}(\Lambda) \cap \Lambda_1$.

Proof. First notice that $\Lambda \subseteq \Lambda_1 \subseteq \Lambda_0$ and $\Lambda^* \cap \gamma\mathcal{B} \subseteq \Lambda_1^*$. The probability-mass function of \mathbf{y} is given by

$$Y(\mathbf{x}) = \frac{1}{\rho_\sigma(\Lambda_0)} \sum_{\mathbf{z} \in \Lambda} \rho_\sigma(\mathbf{x} + \mathbf{z})$$

for $\mathbf{x} \in \mathcal{P}(\Lambda) \cap \Lambda_0$. The Fourier-transform of Y is

$$\widehat{Y}(\boldsymbol{\omega}) = \frac{1}{\det(\Lambda_0^*) \cdot \rho_{1/\sigma}(\Lambda_0^*)} \cdot \sum_{\boldsymbol{\xi} \in \Lambda_0^*} \rho_{1/\sigma}(\boldsymbol{\omega} + \boldsymbol{\xi}) = \frac{1}{\rho_{1/\sigma}(\Lambda_0^*)} \cdot \sum_{\boldsymbol{\xi} \in \Lambda_0^*} \rho_{1/\sigma}(\boldsymbol{\omega} + \boldsymbol{\xi})$$

for $\boldsymbol{\omega} \in \mathcal{P}(\Lambda_0^*) \cap \Lambda^*$.

The probability-mass function of \mathbf{u} is

$$U(\mathbf{x}) = \frac{\det \Lambda_1}{\det \Lambda} \sum_{\mathbf{y} \in \Lambda_1} \delta(\mathbf{x} + \mathbf{y})$$

for $\mathbf{x} \in \mathcal{P}(\Lambda) \cap \Lambda_0$.

Note that $U(\mathbf{x})$ is Λ -periodic as $\Lambda \subseteq \Lambda_1$. We can therefore compute the Fourier-transform of U and obtain

$$\hat{U}(\boldsymbol{\omega}) = \sum_{\boldsymbol{\xi} \in \Lambda_1^*} \delta(\boldsymbol{\omega} + \boldsymbol{\xi}),$$

i.e. $\hat{U}(\boldsymbol{\omega})$ is constant 1 on $\mathcal{P}(\Lambda_0^*) \cap \Lambda_1^\top$ and 0 everywhere else.

By the convolution theorem, the Fourier-transform of the probability mass function of $\mathbf{r} = \mathbf{y} + \mathbf{u} \bmod \Lambda$ is

$$R(\boldsymbol{\omega}) = \hat{Y}(\boldsymbol{\omega}) \cdot \hat{U}(\boldsymbol{\omega})$$

for $\boldsymbol{\omega} \in \mathcal{P}(\Lambda_0^*) \cap \Lambda^*$.

Consequently, we can bound the statistical distance between \mathbf{y} and \mathbf{r} by

$$\begin{aligned} 2 \cdot \Delta(\mathbf{y}, \mathbf{r}) &= \sum_{\mathbf{x} \in \mathcal{P}(\Lambda) \cap \Lambda_0} |Y(\mathbf{x}) - R(\mathbf{x})| \\ &\leq \frac{\det \Lambda}{\det \Lambda_0} \cdot \max_{\mathbf{x} \in \mathcal{P}(\Lambda) \cap \Lambda_0} |Y(\mathbf{x}) - R(\mathbf{x})| \\ &= \frac{\det \Lambda}{\det \Lambda_0} \cdot \max_{\mathbf{x} \in \mathcal{P}(\Lambda) \cap \Lambda_0} \left| \frac{\det \Lambda_0}{\det \Lambda} \cdot \sum_{\boldsymbol{\omega} \in \mathcal{P}(\Lambda_0^*) \cap \Lambda^*} \hat{Y}(\boldsymbol{\omega})(1 - \hat{U}(\boldsymbol{\omega}))e^{2\pi i \langle \boldsymbol{\omega}, \mathbf{x} \rangle} \right| \\ &= \max_{\mathbf{x} \in \mathcal{P}(\Lambda) \cap \Lambda_0} \left| \sum_{\boldsymbol{\omega} \in \mathcal{P}(\Lambda_0^*) \cap \Lambda^* \setminus \Lambda_1^*} \hat{Y}(\boldsymbol{\omega})e^{2\pi i \langle \boldsymbol{\omega}, \mathbf{x} \rangle} \right| \\ &= \max_{\mathbf{x} \in \mathcal{P}(\Lambda) \cap \Lambda_0} \left| \sum_{\boldsymbol{\omega} \in \mathcal{P}(\Lambda_0^*) \cap \Lambda^* \setminus \Lambda_1^*} \frac{1}{\rho_{1/\sigma}(\Lambda_0^*)} \sum_{\boldsymbol{\xi} \in \Lambda_0^*} \rho_{1/\sigma}(\boldsymbol{\omega} + \boldsymbol{\xi})e^{2\pi i \langle \boldsymbol{\omega}, \mathbf{x} \rangle} \right| \\ &\leq \max_{\mathbf{x} \in \mathcal{P}(\Lambda) \cap \Lambda_0} \left| \sum_{\boldsymbol{\omega} \in \Lambda^* \setminus \Lambda_1^*} \rho_{1/\sigma}(\boldsymbol{\omega})e^{2\pi i \langle \boldsymbol{\omega}, \mathbf{x} \rangle} \right| \\ &\leq \max_{\mathbf{x} \in \mathcal{P}(\Lambda) \cap \Lambda_0} \sum_{\boldsymbol{\omega} \in \Lambda^* \setminus \Lambda_1^*} |\rho_{1/\sigma}(\boldsymbol{\omega})e^{2\pi i \langle \boldsymbol{\omega}, \mathbf{x} \rangle}| \\ &= \sum_{\boldsymbol{\omega} \in \Lambda^* \setminus \Lambda_1^*} \rho_{1/\sigma}(\boldsymbol{\omega}) = \rho_{1/\sigma}(\Lambda^* \setminus \Lambda_1^*) \leq \rho_{1/\sigma}(\Lambda^* \setminus \gamma \mathcal{B}) \leq \epsilon \end{aligned}$$

The second inequality holds as $\frac{1}{\rho_{1/\sigma}(\Lambda_0^*)} \leq 1$ and the third inequality is an application of the triangle inequality. \square