Subspace LWE

Krzysztof Pietrzak*

IST Austria

Abstract. The (decisional) learning with errors problem (LWE) asks to distinguish "noisy" inner products of a secret vector with random vectors from uniform. The learning parities with noise problem (LPN) is the special case where the elements of the vectors are bits. In recent years, the LWE and LPN problems have found many applications in cryptography.

In this paper we introduce a (seemingly) much stronger *adaptive* assumption, called "subspace LWE" (SLWE), where the adversary can learn the inner product of the secret and random vectors after they were projected into an adaptively and adversarially chosen subspace. We prove that, surprisingly, the SLWE problem mapping into subspaces of dimension d is almost as hard as LWE using secrets of length d (the other direction is trivial.)

This result immediately implies that several existing cryptosystems whose security is based on the hardness of the LWE/LPN problems are provably secure in a much stronger sense than anticipated. As an illustrative example we show that the standard way of using LPN for symmetric CPA secure encryption is even secure against a very powerful class of related key attacks.

1 Introduction

The (search version of the) learning with errors problem (LWE) is specified by parameters $\ell, q \in \mathbb{N}$ and an error distribution χ over $\mathbb{Z}_q$. It asks to find a secret vector $\mathbf{s} \in \mathbb{Z}_q^\ell$ given any number of "noisy" inner products of $\mathbf{s}$ with random vectors. Formally, these products are samples from a distribution $A_{\chi,\ell}(\mathbf{s})$ over $\mathbb{Z}_q^{\ell+1}$ which is defined by sampling a uniform $\mathbf{r} \xleftarrow{\$} \mathbb{Z}_q^\ell$ and an error $e \leftarrow \chi$, and outputting $(\mathbf{r}, \mathbf{r}^\mathsf{T}\mathbf{s} + e)$ (where multiplications and additions are all modulo q.)

An important special case of this problem is Regev's LWE problem [Reg05] where χ is a so called discrete Gaussian distribution and q is polynomial or exponential in a security parameter. Another important case is the learning parities with noise problem (LPN) where $q = 2$.

The decisional version of the LWE problem asks to *distinguish* samples of the form $A_{\chi,\ell}(\mathbf{s})$ from uniform (which might be easier than to actually output $\mathbf{s}$ as required by the computational version of the problem). The decisional LWE

* Supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC Starting Grant (259668-PSPC).

R. Cramer (Ed.): TCC 2012, LNCS 7194, pp. 548–563, 2012.

problem has been proven polynomially equivalent to the computational version if q is prime [Reg05], and in particular for LPN [BFKL94, KS06]. In this paper we will always consider the decisional version of the problem, and we also only prove the main result for the case where q is prime.

Regev's LWE. The LWE problem has proven to be extremely useful to construct cryptographic schemes. One reason is its versatility, pretty much any cryptographic primitive known to date can be based on LWE. Another reason is its hardness. The best known algorithms against Regev's LWE (where χ is a discrete Gaussian and $q = poly(\ell)$) need time and space $2^{\Theta(\ell)}$ [BKW00] to recover $\mathbf{s} \in \mathbb{Z}_q^\ell$,[1] and unlike for most other assumptions on which public-key crypto can be based, no faster quantum algorithms for the problem are known. But most strikingly, Regev's LWE is as hard as *worst-case* (standard) lattice assumptions [Reg05, Pei09].

An incomplete list of cryptosystems whose security can be reduced to LWE is public-key encryption secure against chosen plaintext [Reg05, KTX08, PVW08] and chosen ciphertext attacks [PW08, Pei09], circular-secure encryption [ACPS09], identity-based encryption [GPV08, CHKP10, ABB10a, ABB10b], oblivious transfer [PVW08], collision-resistant hash functions [PR06, LMPR08] and public-key identification schemes [Lyu08, Lyu09].

LPN. The learning parity with noise (LPN) problem [BFKL94, BKW00, Kea93] is the special case of the LWE problem where $q = 2$ (i.e. we work over bits) and the error distribution is the Bernoulli distribution for some constant parameter $\tau, 0 < \tau < 1/2$, denoted Ber_τ, and defined as $\Pr[x = 1 \,;\, x \leftarrow \mathsf{Ber}_\tau] = \tau$. The LPN problem is closely related to the problem of decoding random linear codes,[2] a well studied question in coding theory. The LPN problem seems less versatile than the general LWE problem, and so far only "minicrypt" primitives (i.e. primitives known to be equivalent to one-way functions) were constructed under the LPN assumption. Alekhnovich [Ale03] constructs a public-key encryption scheme from a relaxed LPN assumption where the error τ is not constant but upper bounded as a function of ℓ as $\tau = O(1/\sqrt{\ell})$.

The Appeal of the LPN problem comes from the fact that LPN based schemes can be extremely efficient, just requiring relatively few bit-level operations to compute an inner product of two bit-vectors. Constructions from LPN include PRGs [FS96] and encryption schemes [GRS08, ACPS09] and public-key authentication schemes [Ste94], but by far most work has been done on efficient LPN based authentication schemes which we'll discuss in more detail in Section 4.

Subspace LWE. The LWE problem has been shown to be very robust with respect to *leakage*. Distinguishing LWE samples remains hard even if we the adversary can learn a function $f(\mathbf{s})$ about the secret $\mathbf{s}$ as long as $f(.)$ is compressing

[1] This is slightly better than a trivial brute-force search which takes time $\approx 2^{\ell \log q} = 2^{\Theta(\ell \log \ell)}$ but only linear space.

[2] The only difference is that in the decoding problem one is given a fixed number of samples (typically a small multiple of the length of the secret), whereas in the LPN problem the adversary can ask for arbitrary many samples.

[AGV09] or hard to invert [DKL09, DGK+10, GKPV10]. In this paper we show that the LWE problem is also very robust to *tampering* with the secret vector $\mathbf{s}$ and the randomness vector $\mathbf{r}$ (albeit not with the noise e.)

We define a (seemingly) much stronger *adaptive* version of LWE which we call "Subspace LWE", or SLWE for short. In the SLWE problem the adversary is not restricted to just ask for samples $\mathbf{r}, \mathbf{r}^{\mathsf{T}}.\mathbf{s} + e$ from $A_{\chi,\ell}(\mathbf{s})$ as in LWE, but has access to a more powerful oracle which she can query adaptively. The oracle takes as input the description of two affine mappings $\phi_r, \phi_s : \mathbb{Z}_q^\ell \to \mathbb{Z}_q^\ell$ and outputs a sample

$$\mathbf{r}, \phi_r(\mathbf{r})^{\mathsf{T}}.\phi_s(\mathbf{s}) + e \quad \text{where} \quad \mathbf{r} \xleftarrow{\$} \mathbb{Z}_q^\ell \ , \ e \leftarrow \chi$$

An affine mapping $\phi_r : \mathbb{Z}_q^\ell \to \mathbb{Z}_q^\ell$ (similarly ϕ_s) is given by a matrix and a vector $\phi_r = [\mathbf{X}_r \in \mathbb{Z}_q^{\ell \times \ell}, \mathbf{x}_r \in \mathbb{Z}_q^\ell]$ and its evaluation is defined as

$$\phi_r(\mathbf{r}) = \mathbf{X}_r.\mathbf{r} + \mathbf{x}_r$$

Without additional restrictions, the SLWE problem as just defined is easy to break. By choosing the input to the oracle appropriately,[3] one can e.g. learn samples of the form $\mathbf{s}[i] + e$, $e \leftarrow \chi$ ($\mathbf{s}[i]$ denotes the ith element of $\mathbf{s}$.) For distributions χ as used in LPN or Regev's LWE one can efficiently learn $\mathbf{s}[i]$ (and thus the entire $\mathbf{s}$) from just a few such samples.

We prove that the SLWE problem (using secrets in $\mathbb{Z}_q^\ell$ and error distribution χ) is almost as hard as the standard (q, χ, d)-LWE problem with secrets of length $d \leq \ell$ if the adversary is restricted in the sense that she is only allowed to query ϕ_r, ϕ_s which "overlap" in an $d + \delta$ (or more) dimensional subspace where $\delta \in \mathbb{N}$ is a statistical security parameter. Formally this means $\mathbf{X}_r^{\mathsf{T}}.\mathbf{X}_s$ must have rank at least $d + \delta$. We call this the $(q, \chi, \ell, d + \delta)$-SLWE problem. Let us mention that the other direction – showing that (q, χ, ℓ, d)-SLWE is at most as hard as (q, χ, d)-LWE – is trivial.

The precise statement of our result asserts that for any $\ell, d, \delta \in \mathbb{N}, d + \delta \leq \ell$, the $(q, \chi, \ell, d + \delta)$-SLWE problem is at most an additive term $2/q^{\delta+1}$ easier than the standard (q, χ, d)-LWE problem. For large fields, where q is superpolynomial, $2/q^{\delta+1}$ is negligible already for $\delta = 0$. For small fields, in particular the important case $q = 2$ as used in LPN, we must choose some δ to be a statistical security parameter.

The above formulation of SLWE is somewhat redundant, in the sense that an adversary who is restricted to always choose ϕ_s to be the identity function, is as powerful (i.e. can learn exactly the same distribution from the SLWE oracle) as the adversary described above. We chose to explicitly allow the adversary to choose affine mappings for the randomness and the secret separately, as for the applications it is sometimes more convenient to think of the adversary being able to apply a mapping to the secret key (like in the setting of related key attacks we'll discuss), or to the randomness (e.g. to show that LWE is hard, even if the randomness comes from a bit-fixing source.)

[3] Set $\mathbf{x}_r = \mathbf{x}_s = 0^\ell$ and $\mathbf{X}_r, \mathbf{X}_s$ the zero matrix with a single one in the ith diagonal element. The oracle will output $\mathbf{r}, \mathbf{r}[i]\mathbf{s}[i] + e$, the last element is $\mathbf{s}[i] + e$ if $\mathbf{r}[i] = 1$.

When q is not Prime. The reduction from SLWE to LWE assumes that q is prime, as we use the fact that $\mathbb{Z}_q$ is a field and $\mathbb{Z}_q^m$ is a vector space.[4] We believe that the proof of the reduction can be adapted to the case where q is composite.

The case where q is prime covers the cryptographically most interesting cases of LPN and Regev's LWE. Also the reduction from the search to decision version of LWE [Reg05] only works for prime q (of polynomial size.) But the case where q is not prime has found cryptographic applications too. In particular, the case where $q = p^e$ for a prime p and $e > 1$ has been used in the construction of an encryption scheme with circular security [ACPS09]. The case where q is a product of distinct, small primes has been used in [Pei09].

Applications of SLWE. In Section 4 we'll discuss some applications of the SLWE problem. In particular, the fact that SLWE is equivalent to LWE implies stronger security notions – like security against related-key attacks – that one can give for existing schemes whose security is reduced to the LWE problem. In subsequent work, the hardness of SLPN has been used to construct efficient authentication schemes and even MACs from LPN. These schemes differ significantly from previous schemes which all were extensions of the Hopper-Blum protocol.

Outline. In Section 2 we first define the LWE and the new subspace LWE (SLWE) problem. In Section 3 we state and prove our main technical result (Theorem 1) which bounds the hardness of the SLWE problem in terms of the hardness of he standard LWE problem. In Section 4 we describe in more detail some applications of this result which were already mentioned in the introduction.

2 Hard Learning Problems

2.1 Notation

We denote the set of integers modulo an integer $q \geq 1$ by $\mathbb{Z}_q$. We will use normal, bold and capital bold letters like x, $\mathbf{x}$, $\mathbf{X}$ to denote single elements, vectors and matrices over $\mathbb{Z}_q$, respectively. For $\mathbf{x} \in \mathbb{Z}_q^\ell$, $|\mathbf{x}| = \ell$ denotes the length and $\mathbf{wt}(\mathbf{x})$ denotes the Hamming weight of the vector $\mathbf{x}$, i.e. the number of indices $i \in \{1, \ldots, |\mathbf{x}|\}$ where $\mathbf{x}[i] \neq 0$. For $\mathbf{v} \in \mathbb{Z}_2^m$ we denote with $\overline{\mathbf{v}}$ its inverse, i.e. $\overline{\mathbf{v}}[i] = 1 - \mathbf{v}[i]$ for all i. For a distribution χ, $x \leftarrow \chi$ denotes sampling a value x with distribution χ. For a set $\mathcal{S}$, $x \xleftarrow{\$} \mathcal{S}$ denotes sampling a value x with the uniform distribution over $\mathcal{S}$.

$\mathbf{x}_{\downarrow\mathbf{v}}, \mathbf{X}_{\downarrow\mathbf{v}}$: For two vectors $\mathbf{v} \in \mathbb{Z}_2^\ell$ and $\mathbf{x} \in \mathbb{Z}_q^\ell$, we denote with $\mathbf{x}_{\downarrow\mathbf{v}}$ the vector (of length $\mathbf{wt}(\mathbf{v})$) which is derived from $\mathbf{x}$ by deleting all the bits $\mathbf{x}[i]$ where $\mathbf{v}[i] = 0$. If $\mathbf{X} \in \mathbb{Z}_q^{\ell \times m}$ is a matrix, then $\mathbf{X}_{\downarrow\mathbf{v}} \in \mathbb{Z}_q^{\mathbf{wt}(\mathbf{v}) \times m}$ denotes the submatrix we get by deleting the ith row if $\mathbf{v}[i] = 0$.

[4] The fact that $\mathbb{Z}_q^m$ is a vector space is e.g. used in the proof of Lemma 1.

$\mathbf{x_v}, \mathbf{X_v}$: For $\mathbf{x}, \mathbf{X}, \mathbf{v}$ as in the previous item, $\mathbf{x_v}$ denotes the vector where the ith entry is $\mathbf{x}[i] \wedge \mathbf{v}[i]$. Think of $\mathbf{x_v}$ as $\mathbf{x}$ where all entries of $\mathbf{x}$ where $\mathbf{v}$ is 0 are set to 0. $\mathbf{X_v}$ denotes the matrix $\mathbf{X}$ where the ith row is set to all 0 if $\mathbf{v}[i] = 0$.

2.2 The (Subspace) LWE Problem

The (search version of the) learning with errors (LWE) problem is specified by parameters $\ell, q \in \mathbb{N}$ and an error distribution χ over $\mathbb{Z}_q$. It asks to find a secret vector $\mathbf{s} \in \mathbb{Z}_q^\ell$ given any number of "noisy" inner products of $\mathbf{s}$ with random vectors.

Formally, let $\Lambda_{\chi,\ell}(\mathbf{s})$ be the distribution over $\mathbb{Z}_q^{\ell+1}$ where a sample is given by

$$(\mathbf{r}, \mathbf{r}^\mathsf{T}.\mathbf{s} + e) \leftarrow \Lambda_{\chi,\ell}(\mathbf{s}) \qquad \text{where} \qquad \mathbf{r} \xleftarrow{\$} \mathbb{Z}_q^\ell \ , \ e \leftarrow \chi$$

Let U_q^m denote the uniform distribution over $\mathbb{Z}_q^m$ and $U_q = U_q^1$. The decisional LWE problem asks to distinguish samples from $\Lambda_{\chi,\ell}(\mathbf{s})$ with a uniform $\mathbf{s}$ from a random oracle (outputting $U_q^{\ell+1}$ samples.) For any $\mathbf{s}$, $\Lambda_{U_q,\ell}(\mathbf{s})$ is the same as the uniform distribution $U_q^{\ell+1}$. It will be convenient for the proof to think of the random oracle as outputting samples from $\Lambda_{U_q,\ell}(\mathbf{s})$ for some random $\mathbf{s}$ instead of $U_q^{\ell+1}$.

Definition 1 (Decisional Learning with Errors Problem (LWE)). *The (decisional) (q, χ, ℓ)-LWE problem is (t, Q, ε) hard if for every distinguisher D running in time t and making Q oracle queries,*

$$\left| \Pr\left[\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^\ell \ : \ D^{\Lambda_{\chi,\ell}(\mathbf{s})} = 1\right] - \underbrace{\Pr\left[\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^\ell \ : \ D^{\Lambda_{U_q,\ell}(\mathbf{s})} = 1\right]}_{\Pr[D^{U_q^{\ell+1}}=1]} \right| \leq \varepsilon. \qquad (1)$$

Usually one defines the LWE problem by considering a distinguisher who gets a polynomial number of samples as input and not access to an oracle (which doesn't take inputs anyway.) We use this oracle based definition so it is more similar to the SLWE problem we define below, where the oracle does take adaptively chosen inputs.

An affine projection $\phi : \mathbb{Z}_q^\ell \to \mathbb{Z}_q^\ell$ is given by a matrix/vector tuple $\mathbf{X} \in \mathbb{Z}_q^{\ell\times\ell}, \mathbf{x} \in \mathbb{Z}_q^\ell$ and defined as $\phi(\mathbf{v}) \stackrel{\mathsf{def}}{=} \mathbf{X}^\mathsf{T}\mathbf{v} + \mathbf{x}$.

For $\mathbf{s} \in \mathbb{Z}_q^{\ell\times\ell}$ and affine projections $\phi_r = [\mathbf{X}_r, \mathbf{x}_r], \phi_s = [\mathbf{X}_s, \mathbf{x}_s]$ we define the distribution $\Gamma_{\chi,\ell,d}(\mathbf{s}, \phi_r, \phi_s)$ over $\mathbb{Z}_q^{\ell+1} \cup \perp$ as

$$\perp \quad \leftarrow \quad \Gamma_{\chi,\ell,d}(\mathbf{s}, \phi_r, \phi_s) \quad \text{if} \quad \mathsf{rank}(\mathbf{X}_r^\mathsf{T}\mathbf{X}_s) < d$$

and

$$[\ \mathbf{r}\ ,\ \phi_r(\mathbf{r})^\mathsf{T}.\phi_s(\mathbf{s}) + e\ ] \leftarrow \Gamma_{\chi,\ell,d}(\mathbf{s}, \phi_r, \phi_s) \quad \text{where} \quad \mathbf{r} \xleftarrow{\$} \mathbb{Z}_q^\ell \ , \ e \leftarrow \chi$$

otherwise. With $\Gamma_{\chi,\ell,d}(\mathbf{s}, .)$ we denote the oracle which on input ϕ_r, ϕ_s outputs a sample of $\Gamma_{\chi,\ell,d}(\mathbf{s}, \phi_r, \phi_s)$.

Definition 2 (Subspace Learning with Errors Problem (SLWE)). *The (decisional) (q, χ, ℓ, d)-SLWE problem is (t, Q, ε) hard if for every distinguisher D running in time t and making Q oracle queries,*

$$\left|\Pr\left[\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^\ell \ : \ D^{\Gamma_{\chi,\ell,d}(\mathbf{s},.)} = 1\right] - \Pr\left[\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^\ell \ : \ D^{\Gamma_{U_q,\ell,d}(\mathbf{s}.,)} = 1\right]\right| \leq \varepsilon. \quad (2)$$

Note that by definition the $\Gamma_{U_q,\ell,d}(\mathbf{s}.,)$ oracle outputs $\perp$ if the input satisfies $\mathsf{rank}(\mathbf{X}_r^\mathsf{T}\mathbf{X}_s) < d$ and a uniform sample $U_q^{\ell+1}$ otherwise. In particular, like $\Lambda_{U_q,\ell}(\mathbf{s})$, the output distribution of $\Gamma_{U_q,\ell,d}(\mathbf{s}.,)$ is independent of $\mathbf{s}$.

3 The Hardness of SLWE

Theorem 1 below is the main technical result stating that the SLWE problem mapping into subspaces of dimension d is almost as hard as the standard LWE problem with secrets of length d. But let's first look at the (easy) other direction as stated by Claim 1 below.

Claim 1 ((q, χ, ℓ, d)-SLWE at most as hard as (q, χ, d)-LWE). *If (q, χ, ℓ, d)-SLWE is (s, t, ϵ) hard then (q, χ, d)-LWE is (s', t, ϵ) hard where $s' = s - poly(t, \ell)$.*

Proof (of Claim). To prove this claim we will show how, for any error distribution χ', one can efficiently generate (q, χ', d)-LWE samples which have distribution $\Lambda_{\chi',d}(\mathbf{s}')$ (for some uniform $\mathbf{s}' \in \mathbb{Z}_q^d$) given access to a (q, χ', ℓ, d)-SLWE oracle $\Gamma_{\chi',\ell,d}(\mathbf{s}, .)$ (for some uniform $\mathbf{s} \in \mathbb{Z}_q^\ell$.) We do so without known knowing the distribution χ' or $\mathbf{s}$.

Given such a transformation, we then can use any distinguisher D who breaks the (q, χ, d)-LWE assumption with advantage ϵ as defined in eq.(1), to break the (q, χ, ℓ, d)-SLWE assumption as in eq.(2) with the same advantage by simply transforming the SLWE samples (where the oracle uses either the error distribution $\chi' = \chi$ or $\chi' = U_q^{\ell+1}$, but we don't know which) to LWE samples (with the same unknown error distribution χ') before forwarding them do D.

Let $\mathbf{v} \stackrel{\mathsf{def}}{=} 1^d \| 0^{\ell-d}$. To generate samples as described above, query $\Gamma_{\chi',\ell,d}(\mathbf{s}, .)$ so it outputs samples $\Lambda_{\chi',d}(\mathbf{s}')$ where $\mathbf{s}' \in \mathbb{Z}_q^d$ consists of, say, the first d elements of $\mathbf{s} \in \mathbb{Z}_q^\ell$, i.e. $\mathbf{s}' := \mathbf{s}_{\downarrow\mathbf{v}}$. This can be done by making q queries $\mathbf{X}_s, \mathbf{X}_r, \mathbf{x}_s, \mathbf{x}_r$ to $\Gamma_{\chi',\ell,d}(\mathbf{s}, .)$ where $\mathbf{x}_s = \mathbf{x}_r = 0^\ell$ and $\mathbf{X}_s = \mathbf{X}_r$ is 1 in the first d diagonal entries and 0 everywhere else. The output of the SLWE oracle on these queries are samples of the form

$$\mathbf{r}, \underbrace{(\mathbf{X}_r\mathbf{r} + \mathbf{x}_r)^\mathsf{T}}_{\mathbf{r}_{\downarrow\mathbf{v}}^\mathsf{T}} \underbrace{(\mathbf{X}_s\mathbf{s} + \mathbf{x}_s)}_{\mathbf{s}_{\downarrow\mathbf{v}}} + e \quad \text{where} \quad e \leftarrow \chi' \ , \ \mathbf{r} \xleftarrow{\$} \mathbb{Z}_q^\ell$$

from which we then get an $\Lambda_{\chi',d}(\mathbf{s}_{\downarrow\mathbf{v}})$ sample $\mathbf{r}_{\downarrow\mathbf{v}}, \mathbf{r}_{\downarrow\mathbf{v}}^\mathsf{T}\mathbf{s}_{\downarrow\mathbf{v}} + e$ by replacing $\mathbf{r}$ with $\mathbf{r}_{\downarrow\mathbf{v}}$. Note that these samples have the right distribution, which means $\mathbf{s}_{\downarrow\mathbf{v}}$ and the q $\mathbf{r}_{\downarrow\mathbf{v}}$'s are uniformly random as required. This is easy to see recalling that $\mathbf{s}$ and the q $\mathbf{r}$'s are uniform. □

In the proof of Theorem 1, we'll need the following simple technical Lemma:

Lemma 1. *For $q, d, \delta \in \mathbb{N}$, let $\Delta(q, d, \delta)$ denote the probability that a random matrix in $\mathbb{Z}_q^{(d+\delta)\times d}$ has rank less than d, then*

$$\Delta(q, d, \delta) \leq \frac{2}{q^{\delta+1}} \ .$$

Proof. Assume we sample the d columns of a matrix $M \in \mathbb{Z}_q^{(d+\delta)\times d}$ one by one. For $i = 1, \ldots, d$ let E_i denote the event that the first i columns are linearly independent, then

$$\Pr[\neg E_i | E_{i-1}] = \frac{q^{i-1}}{q^{d+\delta}} = q^{i-1-d-\delta}$$

as $\neg E_i$ happens iff the ith column (sampled uniformly from a space of size $q^{d+\delta}$) falls into the space (of size q^{i-1}) spanned by the first $i-1$ columns. We get further

$$\Delta(q, d, \delta) = \Pr[\neg E_d] \leq \sum_{i=1}^{d} \Pr[\neg E_i | E_{i-1}] = \sum_{i=1}^{d} q^{i-1-d-\delta} \leq \frac{2}{q^{\delta+1}}$$

□

Theorem 1 ((q, χ, ℓ, d)-SLWE almost as hard as (q, χ, d)-LWE) *For $q, d, \delta, \ell \in \mathbb{N}$. If the (q, χ, d)-LWE problem is (s, t, ϵ) hard, then the $(q, \chi, \ell, d+\delta)$-SLWE problem is (s', t, ϵ') hard where*

$$s' = s - poly(\ell, t) \qquad \epsilon' = \epsilon + \frac{2t}{q^{\delta+1}}$$

Proof (of Theorem 1). To prove the theorem we will show how to sample outputs of an SLWE oracle $\Gamma_{\chi',\ell,d+\delta}(\hat{\mathbf{s}}, .)$ for some uniformly random $\hat{\mathbf{s}} \in \mathbb{Z}_q^\ell$ and adversarially chosen inputs, given only standard LWE samples $\Lambda_{\chi',d}(\mathbf{s})$ for some uniform $\mathbf{s} \in \mathbb{Z}_q^d$. This sampling is done without knowing $\mathbf{s}$ or the error distribution χ'.

Given such a transformation, we then can use any distinguisher D who breaks the $(q, \chi, \ell, d+\delta)$-SLWE assumption with advantage ϵ to break the standard (q, χ, d)-LWE assumption with the same advantage, minus the probability that the transformation will fail (which, unlike in the previous claim, is non-zero.)

Recall that an LWE sample $\Lambda_{\chi',d}(\mathbf{s} \in \mathbb{Z}_q^d)$ is of the form

$$\mathbf{r}, \mathbf{r}^\mathsf{T}.\mathbf{s} + e \quad \text{where} \quad e \leftarrow \chi' \quad \mathbf{r} \xleftarrow{\$} \mathbb{Z}_q^d \tag{3}$$

For $\mathbf{X}_r, \mathbf{X}_s \in \mathbb{Z}_q^{\ell\times\ell}$, $\mathbf{x}_s, \mathbf{x}_r \in \mathbb{Z}_q^\ell$, we'll show how to transform such a sample into a an SLWE sample $\Gamma_{\chi',\ell,d+\delta}(\hat{\mathbf{s}}, [\mathbf{X}_r, \mathbf{x}_r, \mathbf{X}_s, \mathbf{x}_s])$. If $\mathsf{rank}(\mathbf{X}_r^\mathsf{T}.\mathbf{X}_s) < d+\delta$ this sample is simply $\perp$, so from now one we assume that this rank is at least $d+\delta$, in this case the sample has the form

$$\hat{\mathbf{r}}, (\mathbf{X}_r.\hat{\mathbf{r}} + \mathbf{x}_r)^\mathsf{T}(\mathbf{X}_s.\hat{\mathbf{s}} + \mathbf{x}_s) + e \quad \text{where} \quad e \leftarrow \chi' \quad \hat{\mathbf{r}} \xleftarrow{\$} \mathbb{Z}_q^\ell \tag{4}$$

In our transformation, the SLWE secret $\hat{\mathbf{s}} \in \mathbb{Z}_q^\ell$ is defined as a function of the LWE secret $\mathbf{s} \in \mathbb{Z}_q^d$ as follows

$$\mathbf{R} \xleftarrow{\$} \mathbb{Z}_q^{\ell \times d} \qquad \mathbf{b} \xleftarrow{\$} \mathbb{Z}_q^\ell \qquad \hat{\mathbf{s}} = \mathbf{R}.\mathbf{s} + \mathbf{b} \tag{5}$$

Note that we only know $\mathbf{R}, \mathbf{b}$ (which we sampled), but will not get $\hat{\mathbf{s}}$ as we don't know $\mathbf{s}$. Also note that $\hat{\mathbf{s}}$ is uniformly random as it is blinded with a uniform $\mathbf{b}$. Define the set $\mathcal{L} \subseteq \mathbb{Z}_q^\ell$, which is the set of solutions to a system of linear equations, as

$$\mathcal{L} = \{\mathbf{y} \ : \ \mathbf{y}.\mathbf{X}_r^\mathsf{T}.\mathbf{X}_s.\mathbf{R} = \mathbf{r}^\mathsf{T} - \mathbf{x}_r^\mathsf{T}.\mathbf{X}_s.\mathbf{R}\}. \tag{6}$$

If $\mathbf{X}_r^\mathsf{T}.\mathbf{X}_s.\mathbf{R}$ has rank at least d, then $\mathcal{L}$ is not empty as the linear equation considered in eq.(6) is (over)defined (we will bound the probability that the rank is d later.) In this case the LWE sample is transformed into an SLWE sample as

$$\underbrace{\mathbf{r}, \mathbf{r}^\mathsf{T}.\mathbf{s} + e}_{\text{LWE Sample (3)}} \quad \rightarrow \quad \underbrace{\hat{\mathbf{r}}, \mathbf{r}^\mathsf{T}.\mathbf{s} + z + e}_{\text{SLWE Sample (4)}} \quad \text{where} \quad \hat{\mathbf{r}} \xleftarrow{\$} \mathcal{L} \tag{7}$$

and the z is computed from known values as

$$z \stackrel{\mathsf{def}}{=} (\hat{\mathbf{r}}^\mathsf{T}.\mathbf{X}_r^\mathsf{T} + \mathbf{x}_r^\mathsf{T}).\mathbf{X}_s.\mathbf{b} + (\mathbf{X}_r.\hat{\mathbf{r}} + \mathbf{x}_r)^\mathsf{T}.\mathbf{x}_s$$

It follows from the three claims below that this sampling gives the right distribution.

Claim. If $\mathbf{T} \stackrel{\mathsf{def}}{=} \mathbf{X}_r^\mathsf{T}.\mathbf{X}_s.\mathbf{R}$ has rank $\geq d$ then $\hat{\mathbf{r}} \xleftarrow{\$} \mathcal{L}$ is uniformly random (given $\mathbf{x}_s, \mathbf{x}_r, \mathbf{X}_s^\mathsf{T}, \mathbf{X}_r^\mathsf{T}, \mathbf{R}, \mathbf{b}$.)

Proof (of Claim). Fix some $\mathbf{t} \in \mathbb{Z}_2^\ell$ of weight $\mathbf{wt}(\mathbf{t}) = d$ such that $\mathbf{T}_{\downarrow \mathbf{t}}$ has full rank. Such a $\mathbf{t}$ exists as $\mathbf{T}$ has rank d.

By eq.(6), $\hat{\mathbf{r}} \xleftarrow{\$} \mathcal{L}$ is a random solution to the equation

$$\hat{\mathbf{r}}.\mathbf{T} = \mathbf{r}^\mathsf{T} - \mathbf{x}_r^\mathsf{T}.\mathbf{X}_s.\mathbf{R}$$

or equivalently (using $\hat{\mathbf{r}}.\mathbf{T} = \hat{\mathbf{r}}_{\downarrow \mathbf{t}}.\mathbf{T}_{\downarrow \mathbf{t}} + \hat{\mathbf{r}}_{\downarrow \overline{\mathbf{t}}}.\mathbf{T}_{\downarrow \overline{\mathbf{t}}}$)

$$\hat{\mathbf{r}}_{\downarrow \mathbf{t}}.\mathbf{T}_{\downarrow \mathbf{t}} = \mathbf{r}^\mathsf{T} - \mathbf{x}_r^\mathsf{T}.\mathbf{X}_s.\mathbf{R} - \hat{\mathbf{r}}_{\downarrow \overline{\mathbf{t}}}.\mathbf{T}_{\downarrow \overline{\mathbf{t}}} \tag{8}$$

Now sampling a random $\hat{\mathbf{r}}$ can be done as follows. First sample $\hat{\mathbf{r}}_{\downarrow \overline{\mathbf{t}}} \xleftarrow{\$} \mathbb{Z}_q^{\ell - d}$ uniformly. The remaining d positions $\hat{\mathbf{r}}_{\downarrow \mathbf{t}} \in \mathbb{Z}_q^d$ are then uniquely determined by $\mathbf{r}$ and given by the solution to the equation (8).

As $\mathbf{T}_{\downarrow \mathbf{t}}$ is a full rank square matrix eq.(8) defines a bijection between $\hat{\mathbf{r}}_{\downarrow \mathbf{t}}$ and $\mathbf{r}$. As $\mathbf{r}$ is chosen uniformly at random, also $\hat{\mathbf{r}}_{\downarrow \mathbf{t}}$ is uniformly random. Thus the entire $\hat{\mathbf{r}}$ is uniform as claimed. □

Claim. The $\hat{\mathbf{r}}, \mathbf{r}^\mathsf{T}.\mathbf{s} + z + e$ as sampled in (7) is an SLWE sample for secret $\hat{\mathbf{s}}$, randomness $\hat{\mathbf{r}}$ and error e.

Proof (of Claim).

$$
\begin{aligned}
& \hat{\mathbf{r}}, (\mathbf{X}_r.\hat{\mathbf{r}} + \mathbf{x}_r)^\mathsf{T}.(\mathbf{X}_s.\hat{\mathbf{s}} + \mathbf{x}_s) + e && \text{(SLWE sample)} \\
&= \hat{\mathbf{r}}, (\mathbf{X}_r.\hat{\mathbf{r}} + \mathbf{x}_r)^\mathsf{T}.(\mathbf{X}_s.\hat{\mathbf{s}}) + (\mathbf{X}_r.\hat{\mathbf{r}} + \mathbf{x}_r)^\mathsf{T}.\mathbf{x}_s + e \\
&\overset{(5)}{=} \hat{\mathbf{r}}, (\mathbf{X}_r.\hat{\mathbf{r}} + \mathbf{x}_r)^\mathsf{T}.(\mathbf{X}_s.(\mathbf{R}.\mathbf{s} + \mathbf{b})) + (\mathbf{X}_r.\hat{\mathbf{r}} + \mathbf{x}_r)^\mathsf{T}.\mathbf{x}_s + e \\
&= \hat{\mathbf{r}}, (\hat{\mathbf{r}}^\mathsf{T}.\mathbf{X}_r^\mathsf{T} + \mathbf{x}_r^\mathsf{T}).(\mathbf{X}_s.\mathbf{R}.\mathbf{s}) + \underbrace{(\hat{\mathbf{r}}^\mathsf{T}.\mathbf{X}_r^\mathsf{T} + \mathbf{x}_r^\mathsf{T}).\mathbf{X}_s.\mathbf{b} + (\mathbf{X}_r.\hat{\mathbf{r}} + \mathbf{x}_r)^\mathsf{T}.\mathbf{x}_s}_{z} + e \\
&\overset{(6)}{=} \hat{\mathbf{r}}, \mathbf{r}^\mathsf{T}.\mathbf{s} + z + e
\end{aligned}
$$

□

We have shown how to simulate an SLWE oracle $\Gamma_{\chi',\ell,d+\delta}(\hat{\mathbf{s}}, .)$ from standard LWE samples $\Lambda_{\chi',d}(\mathbf{s})$. This simulation goes well as long as we never get a query containing $\mathbf{X}_r, \mathbf{X}_s$ where $\mathsf{rank}(\mathbf{X}_r^\mathsf{T}.\mathbf{X}_s) \geq d + \delta$ (so the sample is not just $\perp$) but where $\mathsf{rank}(\mathbf{X}_r^\mathsf{T}.\mathbf{X}_s.\mathbf{R}) < d$ (in this case $\mathcal{L}$ can be empty.) The following claims bounds the probability of this happening.

Claim. Consider any $\mathbf{X} \in \mathbb{Z}_q^{\ell \times \ell}$ with $\mathsf{rank}(\mathbf{X}) \geq d + \delta$, then (with Δ as defined in Lemma 1)

$$\Pr[\mathsf{rank}(\mathbf{X}.\mathbf{R}) < d \ : \ \mathbf{R} \xleftarrow{\$} \mathbb{Z}_q^{\ell \times d}] \leq \Delta(q, d, \delta)$$

Proof (of Claim). Since the matrix $\mathbf{X}$ has rank at least $d + \delta$, without loss of generality, we can assume that the first $d+\delta$ rows of $\mathbf{X}$ are linearly independent. Since $\mathbf{R}$ is a random matrix, the upper $(d+\delta) \times d$ submatrix of $\mathbf{X}.\mathbf{R}$ is a random matrix in $\mathbb{Z}_q^{(d+\delta)\times d}$ and (by definition) such a matrix has rank strictly less than d with probability at most $\Delta(q, d, \delta)$. Thus $\mathbf{X}.\mathbf{R}$ has rank strictly less than d with at most the same probability. □

Using the union bound, we can upper bound the probability that for any of the t queries the matrix $\mathbf{X} = \mathbf{X}_r^\mathsf{T}.\mathbf{X}_s$ chosen by the distinguisher D will satisfy $\mathsf{rank}(\mathbf{X}.\mathbf{R}) < d$ by

$$t \cdot \Delta(q, n, d) \leq \frac{2 \cdot t}{q^{\delta+1}}$$

This error probability is thus an upper bound on the gap of the success probability ϵ' of D (in breaking SLWE) and the success probability ϵ we get in breaking LWE using the transformation.

Above we ignored the fact that D can choose its queries, and thus the matrix $\mathbf{X} = \mathbf{X}_r^\mathsf{T}.\mathbf{X}_s$, *adaptively*. To show that adaptivity does not help in picking an $\mathbf{X}$ where $\mathbf{X}.\mathbf{R}$ has rank $< d$ we must show that the view of D is *independent* of $\mathbf{R}$ (except for the fact that so far no query was made where $\mathsf{rank}(\mathbf{X}.\mathbf{R}) < d$.) To see this first note that $\hat{\mathbf{s}} = \mathbf{s}.\mathbf{R} + \mathbf{b}$ is independent of $\mathbf{R}$ as it is blinded with a uniform $\mathbf{b}$. In fact, the only reason we use this blinding is to enforce this independence. The $\hat{\mathbf{r}}$ are independent as they are uniform given $\mathbf{R}$ as shown in the first Claim in the proof of this theorem. □

4 Applications

In this section we discuss some consequences and applications which use the fact that the new subspace LWE problem is as hard as the classical LWE problem.

4.1 Security against Related Key Attacks

Theorem 1 implies that many existing schemes whose security is based on the standard LWE/LPN assumption are secure against attacks not anticipated by the designers of the schemes.

As an illustrative example below we discuss the simple construction of symmetric CPA secure secure encryption from LPN [GRS08]. We show that this simple scheme is not only CPA secure, but it's even secure against powerful related key-attacks. The scheme from [GRS08] is defined as follows

Public Parameters
- Constants $0 < \tau < 0.5$, $\delta > 0$, $\ell \in \mathbb{N}$.
- An error correcting code $\mathsf{E} : \mathbb{Z}_2^m \to \mathbb{Z}_2^n$, $\mathsf{D} : \mathbb{Z}_2^n \to \mathbb{Z}_2^m$, where D can correct up to $(\tau + \delta)\ell$ errors.

Key Generation: $\mathsf{KG}(\ell)$ samples and outputs $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_2^\ell$.

Encryption: $\mathsf{Enc}(K, \mathbf{m})$ samples $\mathbf{R} \xleftarrow{\$} \mathbb{Z}_2^{\ell \times n}$, $\mathbf{e} \xleftarrow{\$} \mathsf{Ber}_\tau^n$ and outputs the ciphertext $(\mathbf{R}, \mathbf{R}^\mathsf{T}.\mathbf{s} \oplus \mathbf{e} \oplus \mathsf{E}(\mathbf{m}))$.

Decryption: $\mathsf{Dec}(K, (\mathbf{R}, \mathbf{z}))$ outputs $\mathsf{D}(\mathbf{z} \oplus \mathbf{R}^\mathsf{T}.\mathbf{s})$.

Correctness. To see that this scheme is correct, note that on input a correctly generated ciphertext $(\mathbf{R}, \mathbf{R}^T.\mathbf{s} \oplus \mathbf{e} \oplus \mathsf{E}(\mathbf{m}))$, the decryption algorithm outputs $\mathsf{D}(\mathbf{e} \oplus \mathsf{E}(\mathbf{m}))$, which is equal to $\mathbf{m}$ unless the error vector $\mathbf{e}$ has weight more than $(\tau + \delta)\ell$. As the bits of $\mathbf{e}$ are i.i.d. with each bit being one with probability τ, the probability of $\mathbf{e}$ having such high weight can be upper bounded (using the Chernoff bound) by an exponentially small probability $2^{-\gamma \cdot \ell}$ (for some $\gamma > 0$ which depends on τ, δ).

CPA Security. Recall that an encryption scheme is IND-CPA secure if no efficient adversary $\mathcal{A}$ can win the following game with probability noticeably better than 1/2:

1. We sample a key $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_2^\ell$ and a bit $b \xleftarrow{\$} \{0, 1\}$.
2. $\mathcal{A}$ gets access to an oracle $\mathsf{Enc}_b(\mathbf{s}, .)$ where
   - $\mathsf{Enc}_0(\mathbf{s}, \mathbf{m}) = \mathsf{Enc}(\mathbf{s}, \mathbf{m})$ (encrypt $\mathbf{m}$)
   - $\mathsf{Enc}_1(\mathbf{s}, \mathbf{m}) = \mathsf{Enc}(\mathbf{s}, 0^{|\mathbf{m}|})$ (encrypt dummy message)
3. $\mathcal{A}$ outputs b' and wins if $b = b'$.

The IND-CPA security of the [GRS08] encryption scheme follows quite easily from the LPN assumption, i.e. the fact that samples $(\mathbf{R}, \mathbf{R}^\mathsf{T}.\mathbf{s} + \mathbf{e})$ are pseudorandom.

RKA Security. Classical security notions, like IND-CPA security, model the encryption scheme as a "black-box", where an adversary can only observe the legitimate input-output behavior of the scheme. Unfortunately, in the last decade it became evident that such idealized models fail to capture many real-world attacks where an adversary can attack an actual physical implementation of the scheme. An important example is direct leakage from the secret state, typically by side-channel attacks or malware. To deal with this issue, in the last years many "intrusion-resilient" and "leakage-resilient" schemes have been proposed [ISW03, MR04, Dzi06, DP07, DP08, ADW09, Pie09, CDD+07, KV09, DW09, DKL09].

But the key can also leak indirectly, for example due to key-dependent messages [BRS03, HK07, BHHO08, HU08, ACPS09, BHHI10, BG10, ABBC10]. Here, as the name suggest, one considers a setting where the encrypted message can depend on the secret key. Another important setting are related-key attacks (RKA). In an RKA attack on a encryption scheme the adversary can not only ask for encryptions under the secret key $\mathbf{s}$, but also under "related" keys. RKA attacks were first considered by Biham [Bih94] and Knudsen [Knu92], and were extensively studied in the last decade [Luc04, BDK06, BDK08, FKL+00, JD04, ZZWF07, BC10]. Bellare and Kohno [BK03] initiated a formal study of RKA attacks. All this works consider RKA security of deterministic primitives, usually block-ciphers.

Very recently [AHI11] initiated a formal study of RKA security for probabilistic encryption [GM84]. As in [BK03], they define RKA with respect to related-key-deriving functions (RKD) Φ. Φ-RKA-IND-CPA security of an encryption scheme is then defined almost like standard IND-CPA security, but where the adversary can additionally apply any function $\phi \in \Phi$ to the secret key $\mathbf{s}$, i.e.

1. We sample a key $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_2^\ell$ and a bit $b \xleftarrow{\$} \{0,1\}$.
2. $\mathcal{A}$ gets access to an oracle $\mathsf{Enc}_b^{\Phi}(\mathbf{s},.,.)$ where
   - $\mathsf{Enc}_0^{\Phi}(\mathbf{s},\mathbf{m},\phi \in \Phi) = \mathsf{Enc}(\phi(\mathbf{s}),\mathbf{m})$ (encrypt $\mathbf{m}$)
   - $\mathsf{Enc}_1^{\Phi}(\mathbf{s},\mathbf{m},\phi \in \Phi) = \mathsf{Enc}(\phi(\mathbf{s}),0^{|\mathbf{m}|})$ (encrypt dummy message)
3. $\mathcal{A}$ outputs b' and wins if $b = b'$.

In [AHI11] it is shown that [GRS08] is $\Phi^{\oplus}$-RKA-IND-CPA secure where $\Phi^{\oplus}$ is the class of XOR relations. This class contains, for every $\Delta \in \mathbb{Z}_2^\ell$, the function $\phi_\Delta(\mathbf{s}) \stackrel{\mathsf{def}}{=} \mathbf{s} \oplus \Delta$.

This is an interesting class of relations as (1) it captures realistic RKA and (2) many existing schemes (mostly block-ciphers) have actually been shown to be insecure against $\Phi^{\oplus}$-RKA. Unfortunately $\Phi^{\oplus}$-RKA security does not imply any security in the realistic scenario where an adversary can not only flip, but set some of the bit of the secret key. Neither does it cover the case where the adversary can exchange the *position* of the key bits.

Using Theorem 1 we can show that the scheme is in fact secure against a much more powerful class of "affine relations", which as special cases contains the relations just mentioned. Let Φ_d^{aff} be the class which contains the functions

$$\phi_{\mathbf{X},\mathbf{x}}(\mathbf{s}) = \mathbf{X}^T.\mathbf{s} \oplus \mathbf{x}$$

for every $\mathbf{X} \in \mathbb{Z}_2^{\ell \times \ell}, \mathbf{x} \in \mathbb{Z}_2^\ell$ where $\mathsf{rank}(\mathbf{X}) \geq d$.

Proposition 1. *Under the (decisional) (τ, ℓ, d)-SLPN assumption*[5] *(which by Theorem 1 is equivalent to the standard LPN assumption), the encryption scheme from [GRS08] is Φ_d^{aff}-RKA-IND-CPA secure.*

Proof. For any $\phi \in \Phi_d^{\mathsf{aff}}$, samples of the from $\mathbf{R}, \mathbf{R}^{\mathsf{T}}.\phi(\mathbf{s}) + \mathbf{e}$ are pseudorandom by assumption. So the outputs of both $\mathsf{Enc}_0^{\Phi_d^{\mathsf{aff}}}(\mathbf{s}, ., .)$ and $\mathsf{Enc}_1^{\Phi_d^{\mathsf{aff}}}(\mathbf{s}, ., .)$ are pseudorandom and thus indistinguishable. □

Φ_d^{aff} is a very powerful class of relations, and captures many realistic settings. It contains $\Phi^{\oplus}$, but also the class of relations $\Phi_d^{\mathsf{set}} \subset \Phi_d^{\mathsf{aff}}$ which allows to overwrite all but d bits of the input, and the class $\Phi^{\mathsf{perm}} \subset \Phi_0^{\mathsf{aff}}$ which allows to permute the key bits.[6] Previous to our work no scheme was known to be provably secure against Φ_d^{aff}, or even just for one of the special cases Φ_d^{set} (for $d > 0$) or Φ^{perm}. In fact, no *deterministic* encryption scheme can be secure against Φ^{perm}, and no "natural"[7] deterministic scheme can be secure against Φ_1^{set}.

4.2 Weak Randomness and New Constructions

The RKA security example from the previous section used the fact that an adversary can apply any affine function to the LWE secret. There are also natural implications from the fact that she can apply a mapping to the randomness $\mathbf{r}$. For example, it implies that LWE is hard, even if the randomness $\mathbf{r}$ used to generate the samples $\mathbf{r}^{\mathsf{T}}\mathbf{s}+e$ is not uniform, but comes from a bit-fixing source [CGH+85]. Let us stress that the (comparably small) amount of randomness necessary to sample the error e must be uniform.

Theorem 1 not only has implications for existing constructions, but in subsequent work has inspired completely new constructions, most notably the authentication schemes and message authentications codes proposed in [KPC+11].

References

[ABB10a] Agrawal, S., Boneh, D., Boyen, X.: Efficient Lattice (H)IBE in the Standard Model. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 553–572. Springer, Heidelberg (2010)

[ABB10b] Agrawal, S., Boneh, D., Boyen, X.: Lattice Basis Delegation in Fixed Dimension and Shorter-Ciphertext Hierarchical IBE. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 98–115. Springer, Heidelberg (2010)

[ABBC10] Acar, T., Belenkiy, M., Bellare, M., Cash, D.: Cryptographic Agility and Its Relation to Circular Encryption. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 403–422. Springer, Heidelberg (2010)

[5] This is the $(2, \mathsf{Ber}_\tau, \ell, d)$-SLWE problem as given in Definition 2.

[6] $\Phi^{\mathsf{perm}} \subset \Phi_0^{\mathsf{aff}}$ as it only contains $\phi_{\mathbf{X},\mathbf{x}}$ where $\mathbf{x} = 0^\ell$ and $\mathbf{X}$ is a (full rank) permutation matrix.

[7] We need that every bit of the secret key is relevant, i.e. $\mathsf{Enc}(\mathbf{s}, \mathbf{m}) \neq \mathsf{Enc}(\mathbf{s}', \mathbf{m})$ with good probability for $\mathbf{s} \neq \mathbf{s}'$.

[ACPS09] Applebaum, B., Cash, D., Peikert, C., Sahai, A.: Fast Cryptographic Primitives and Circular-Secure Encryption Based on Hard Learning Problems. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 595–618. Springer, Heidelberg (2009)

[ADW09] Alwen, J., Dodis, Y., Wichs, D.: Leakage-Resilient Public-Key Cryptography in the Bounded-Retrieval Model. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 36–54. Springer, Heidelberg (2009)

[AGV09] Akavia, A., Goldwasser, S., Vaikuntanathan, V.: Simultaneous Hardcore Bits and Cryptography against Memory Attacks. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 474–495. Springer, Heidelberg (2009)

[AHI11] Applebaum, B., Harnik, D., Ishai, Y.: Semantic security under related-key attacks and applications. In: 2nd Innovations in Computer Science (ICS), pp. 45–60 (2011)

[Ale03] Alekhnovich, M.: More on average case vs approximation complexity. In: 44th FOCS, pp. 298–307. IEEE Computer Society Press (October 2003)

[BC10] Bellare, M., Cash, D.: Pseudorandom Functions and Permutations Provably Secure against Related-Key Attacks. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 666–684. Springer, Heidelberg (2010)

[BDK06] Biham, E., Dunkelman, O., Keller, N.: New Cryptanalytic Results on IDEA. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 412–427. Springer, Heidelberg (2006)

[BDK08] Biham, E., Dunkelman, O., Keller, N.: A Unified Approach to Related-Key Attacks. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 73–96. Springer, Heidelberg (2008)

[BFKL94] Blum, A., Furst, M.L., Kearns, M., Lipton, R.J.: Cryptographic Primitives Based on Hard Learning Problems. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 278–291. Springer, Heidelberg (1994)

[BG10] Brakerski, Z., Goldwasser, S.: Circular and Leakage Resilient Public-Key Encryption under Subgroup Indistinguishability (or: Quadratic Residuosity Strikes Back). In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 1–20. Springer, Heidelberg (2010)

[BHHI10] Barak, B., Haitner, I., Hofheinz, D., Ishai, Y.: Bounded Key-Dependent Message Security. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 423–444. Springer, Heidelberg (2010)

[BHHO08] Boneh, D., Halevi, S., Hamburg, M., Ostrovsky, R.: Circular-Secure Encryption from Decision Diffie-Hellman. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 108–125. Springer, Heidelberg (2008)

[Bih94] Biham, E.: New types of cryptanalytic attacks using related keys. Journal of Cryptology 7(4), 229–246 (1994)

[BK03] Bellare, M., Kohno, T.: A Theoretical Treatment of Related-key Attacks: RKA-PRPs, RKA-PRFs, and Applications. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 491–506. Springer, Heidelberg (2003)

[BKW00] Blum, A., Kalai, A., Wasserman, H.: Noise-tolerant learning, the parity problem, and the statistical query model. In: 32nd ACM STOC, pp. 435–440. ACM Press (May 2000)

[BRS03] Black, J., Rogaway, P., Shrimpton, T.: Encryption-Scheme Security in the Presence of Key-Dependent Messages. In: Nyberg, K., Heys, H.M. (eds.) SAC 2002. LNCS, vol. 2595, pp. 62–75. Springer, Heidelberg (2003)

[CDD+07] Cash, D., Ding, Y.Z., Dodis, Y., Lee, W., Lipton, R.J., Walfish, S.: Intrusion-Resilient Key Exchange in the Bounded Retrieval Model. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 479–498. Springer, Heidelberg (2007)

[CGH+85] Chor, B., Goldreich, O., Hastad, J., Friedman, J., Rudich, S., Smolensky, R.: The bit extraction problem of t-resilient functions (preliminary version). In: FOCS, pp. 396–407 (1985)

[CHKP10] Cash, D., Hofheinz, D., Kiltz, E., Peikert, C.: Bonsai Trees, or How to Delegate a Lattice Basis. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 523–552. Springer, Heidelberg (2010)

[DGK+10] Dodis, Y., Goldwasser, S., Tauman Kalai, Y., Peikert, C., Vaikuntanathan, V.: Public-Key Encryption Schemes with Auxiliary Inputs. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 361–381. Springer, Heidelberg (2010)

[DKL09] Dodis, Y., Kalai, Y.T., Lovett, S.: On cryptography with auxiliary input. In: Mitzenmacher, M. (ed.) 41st ACM STOC, pp. 621–630. ACM Press (May/June 2009)

[DP07] Dziembowski, S., Pietrzak, K.: Intrusion-resilient secret sharing. In: 48th FOCS, pp. 227–237. IEEE Computer Society Press (October 2007)

[DP08] Dziembowski, S., Pietrzak, K.: Leakage-resilient cryptography. In: 49th FOCS, pp. 293–302. IEEE Computer Society Press (October 2008)

[DW09] Dodis, Y., Wichs, D.: Non-malleable extractors and symmetric key cryptography from weak secrets. In: Mitzenmacher, M. (ed.) 41st ACM STOC, pp. 601–610. ACM Press (May/June 2009)

[Dzi06] Dziembowski, S.: Intrusion-Resilience via the Bounded-Storage Model. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 207–224. Springer, Heidelberg (2006)

[FKL+00] Ferguson, N., Kelsey, J., Lucks, S., Schneier, B., Stay, M., Wagner, D., Whiting, D.L.: Improved Cryptanalysis of Rijndael. In: Schneier, B. (ed.) FSE 2000. LNCS, vol. 1978, pp. 213–230. Springer, Heidelberg (2001)

[FS96] Fischer, J.-B., Stern, J.: An Efficient Pseudo-random Generator Provably as Secure as Syndrome Decoding. In: Maurer, U.M. (ed.) EUROCRYPT 1996. LNCS, vol. 1070, pp. 245–255. Springer, Heidelberg (1996)

[GKPV10] Goldwasser, S., Kalai, Y., Peikert, C., Vaikuntanathan, V.: Robustness of the learning with errors assumption. In: 1st Innovations in Computer Science (ICS) (2010)

[GM84] Goldwasser, S., Micali, S.: Probabilistic encryption. Journal of Computer and System Sciences 28(2), 270–299 (1984)

[GPV08] Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC, pp. 197–206. ACM Press (May 2008)

[GRS08] Gilbert, H., Robshaw, M., Seurin, Y.: How to Encrypt with the LPN Problem. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfsdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part II. LNCS, vol. 5126, pp. 679–690. Springer, Heidelberg (2008)

[HK07] Halevi, S., Krawczyk, H.: Security under key-dependent inputs. In: Ning, P., De Capitani di Vimercati, S., Syverson, P.F. (eds.) ACM CCS 2007, pp. 466–475. ACM Press (October 2007)

[HU08] Hofheinz, D., Unruh, D.: Towards Key-Dependent Message Security in the Standard Model. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 108–126. Springer, Heidelberg (2008)

[ISW03] Ishai, Y., Sahai, A., Wagner, D.: Private Circuits: Securing Hardware against Probing Attacks. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 463–481. Springer, Heidelberg (2003)

[JD04] Jakimoski, G., Desmedt, Y.: Related-key Differential Cryptanalysis of 192-Bit Key AES Variants. In: Matsui, M., Zuccherato, R.J. (eds.) SAC 2003. LNCS, vol. 3006, pp. 208–221. Springer, Heidelberg (2004)

[Kea93] Kearns, M.J.: Efficient noise-tolerant learning from statistical queries. In: 25th ACM STOC, pp. 392–401. ACM Press (May 1993)

[Knu92] Knudsen, L.R.: Cryptanalysis of LOKI91. In: Zheng, Y., Seberry, J. (eds.) AUSCRYPT 1992. LNCS, vol. 718, pp. 196–208. Springer, Heidelberg (1993)

[KPC+11] Kiltz, E., Pietrzak, K., Cash, D., Jain, A., Venturi, D.: Efficient Authentication from Hard Learning Problems. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 7–26. Springer, Heidelberg (2011)

[KS06] Katz, J., Shin, J.S.: Parallel and Concurrent Security of the HB and HB+ Protocols. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 73–87. Springer, Heidelberg (2006)

[KTX08] Kawachi, A., Tanaka, K., Xagawa, K.: Concurrently Secure Identification Schemes Based on the Worst-Case Hardness of Lattice Problems. In: Pieprzyk, J. (ed.) ASIACRYPT 2008. LNCS, vol. 5350, pp. 372–389. Springer, Heidelberg (2008)

[KV09] Katz, J., Vaikuntanathan, V.: Signature Schemes with Bounded Leakage Resilience. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 703–720. Springer, Heidelberg (2009)

[LMPR08] Lyubashevsky, V., Micciancio, D., Peikert, C., Rosen, A.: SWIFFT: A Modest Proposal for FFT Hashing. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 54–72. Springer, Heidelberg (2008)

[Luc04] Lucks, S.: Ciphers Secure against Related-Key Attacks. In: Roy, B., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 359–370. Springer, Heidelberg (2004)

[Lyu08] Lyubashevsky, V.: Lattice-Based Identification Schemes Secure Under Active Attacks. In: Cramer, R. (ed.) PKC 2008. LNCS, vol. 4939, pp. 162–179. Springer, Heidelberg (2008)

[Lyu09] Lyubashevsky, V.: Fiat-Shamir with Aborts: Applications to Lattice and Factoring-Based Signatures. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 598–616. Springer, Heidelberg (2009)

[MR04] Micali, S., Reyzin, L.: Physically Observable Cryptography (Extended Abstract). In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 278–296. Springer, Heidelberg (2004)

[Pei09] Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In: Mitzenmacher, M. (ed.) 41st ACM STOC, pp. 333–342. ACM Press (May/June 2009)

[Pie09] Pietrzak, K.: A Leakage-Resilient Mode of Operation. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 462–482. Springer, Heidelberg (2009)

[PR06] Peikert, C., Rosen, A.: Efficient Collision-Resistant Hashing from Worst-Case Assumptions on Cyclic Lattices. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 145–166. Springer, Heidelberg (2006)

[PVW08] Peikert, C., Vaikuntanathan, V., Waters, B.: A Framework for Efficient and Composable Oblivious Transfer. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 554–571. Springer, Heidelberg (2008)

[PW08] Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC, pp. 187–196. ACM Press (May 2008)

[Reg05] Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Gabow, H.N., Fagin, R. (eds.) 37th ACM STOC, pp. 84–93. ACM Press (May 2005)

[Ste94] Stern, J.: A New Identification Scheme Based on Syndrome Decoding. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 13–21. Springer, Heidelberg (1994)

[ZZWF07] Zhang, W., Zhang, L., Wu, W., Feng, D.: Related-Key Differential-Linear Attacks on Reduced AES-192. In: Srinathan, K., Pandu Rangan, C., Yung, M. (eds.) INDOCRYPT 2007. LNCS, vol. 4859, pp. 73–85. Springer, Heidelberg (2007)