Provable security against decryption failure attacks from LWE

Christian Majenz, Fabrizio Sisinni

Technical University of Denmark (DTU)

Funded by
the European Union

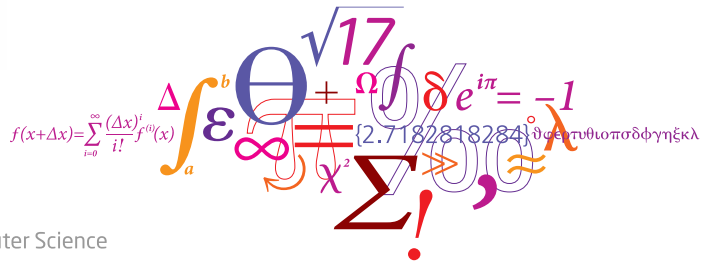

DTU Compute
Department of Applied Mathematics and Computer Science

Why are we interested in this topic?

- KEMs have received a lot of attentions due to the NIST standardization process.

Why are we interested in this topic?

- KEMs have received a lot of attentions due to the NIST standardization process.
- Most of the proposals submitted make use of the Fujisaki-Okamoto (FO) transformation to upgrade their security.

Why are we interested in this topic?

- KEMs have received a lot of attentions due to the NIST standardization process.
- Most of the proposals submitted make use of the Fujisaki-Okamoto (FO) transformation to upgrade their security.
- Most of the proposals are not perfectly correct. Thus, we need to handle with decryption failures.

A possible solution

- Introduce a security notion that handles with decryption failures in a more natural way.

A possible solution

- Introduce a security notion that handles with decryption failures in a more natural way.
- Test this notion to each PQ-secure PKE that is not perfectly correct.

Our contribution

Security reduction from Decision LWE to FFP-NG

Assuming LWE is hard, no adversary can use a given public key to find a message/randomness pair that triggers a decryption failure more likely with respect to the given key than with respect to an independently generated one.

Outline

- Find Failing Plaintext - Non Generic (FFP-NG)
  - The Find Failing Plaintext (FFP) family
  - Basic Properties of FFP-NG
- Discrete Gaussians
  - The smoothing parameter
  - GMPW framework
- Learning With Error
  - Improved version of Regev scheme
- Security reduction from LWE to FFP-NG
  - Main result
  - Before proving the reduction...
  - The LWE adversary
  - Further steps

FFP games

In a recent work, Hövelmanns, Hülsing and Majenz introduced a family of security games called Find Failing Plaintext. We are interested in the following member of this family

FFP-Non Generic	
$(sk_0, pk_0) \leftarrow \mathrm{KeyGen}$	$\underline{\textbf{FCO}_b(v; \mathbf{d}) :}$ # One query
$(sk_1, pk_1) \leftarrow \mathrm{KeyGen}$	$c := \mathrm{Enc}(pk_b, v; \mathbf{d})$
$b \leftarrow \{0, 1\}$	$v' := Dec(sk_b, c)$
$b' \leftarrow \mathcal{A}^{\mathrm{FCO}_b}(pk_0)$	**return** $[\![v \neq v']\!]$
return $[\![b = b']\!]$	

Table: FFP-NG games against a PKE $\Pi = (\mathrm{KeyGen}, \mathrm{Enc}, \mathrm{Dec})$.

K. Hövelmanns, A. Hülsing, C. Majenz, Failing gracefully: Decryption failures and the Fujisaki-Okamoto transform

FFP-NG

FFP-NG is trivially achievable

Any PKE that has perfect correctness is FFP-NG secure.

FFP-NG

FFP-NG is trivially achievable

Any PKE that has perfect correctness is FFP-NG secure.

It is also important to observe that this notion is independent of CPA security.

Independence from standard security notions

1. There exists a PKE Π that is OW-CPA insecure but FFP-NG secure.
2. Assuming the existence of IND-CPA secure PKE Π, there exists a PKE Π' that is IND-CPA secure but FFP-NG insecure.

Definition of Discrete Gaussian

Given an n-dimensional lattice Λ and a covariance matrix Σ, we define the discrete gaussian as follows

$$\mathcal{D}_{\Lambda,\sqrt{\Sigma}} = \frac{\rho_{\sqrt{\Sigma}}(\mathbf{x})}{\rho_{\sqrt{\Sigma}}(\Lambda)},$$

where $\rho_{\sqrt{\Sigma}}(\mathbf{x})$ is the multivariate gaussian function with covariance matrix Σ.

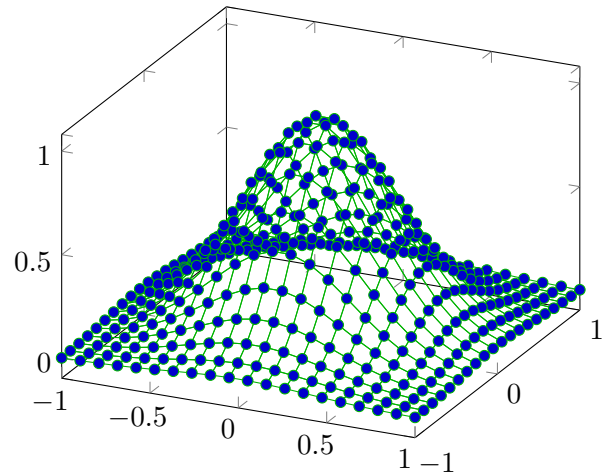

Smoothing parameter

Discrete Gaussians behave like standard Gaussians in so many ways, under some conditions on the so called **smoothing parameter**. Intuitively, beyond the smoothing parameter, the Gaussian measure no longer sees the discrete structure of Λ.

Beyond the smoothing parameter

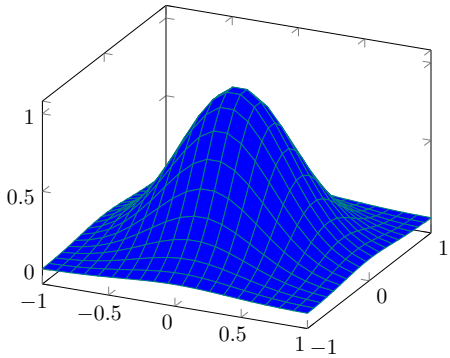

Smoothing parameter

Discrete Gaussians behave like standard Gaussians in so many ways, under some conditions on the so called **smoothing parameter**. Intuitively, beyond the smoothing parameter, the Gaussian measure no longer sees the discrete structure of Λ.

Beyond the smoothing parameter

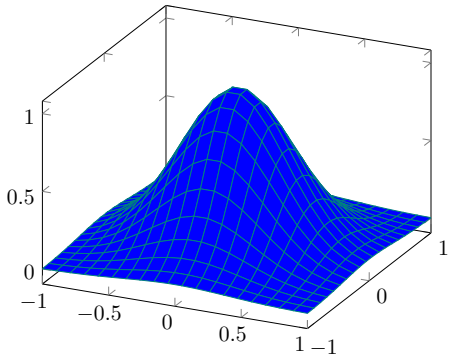

Within the smoothing parameter

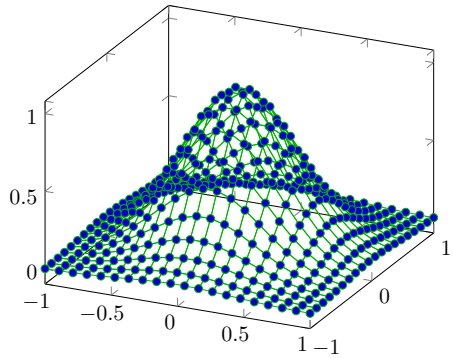

Discrete Gaussians
The GMPW framework

Genise, Micciancio, Peikert and Walter introduce a modular framework to describe several results about Discrete Gaussians. The two main features of this framework are

1. it is based on simple linear algebra and linear transformation;
2. it generalizes all previous frameworks.

Under some conditions on smoothing parameters of the lattices involved, we can show that

1. convolution of discrete gaussians is still a discrete gaussian;
2. linear combination of discrete gaussians is still a discrete gaussian.

N. Genise, D. Micciancio, C. Peikert and M. Walter, Improved Discrete Gaussian and Subgaussian Analysis for Lattice Cryptography

$PVW_{\mathcal{D}}$ PKE with Discrete Gaussian error distribution

Key Generation: choose $\mathbf{S} \in \mathbb{Z}_p^{n \times \ell}$ and $\mathbf{A} \in \mathbb{Z}_p^{n \times m}$ uniformly at random, and $\mathbf{E} \in \mathbb{Z}_p^{\ell \times m}$ by choosing each entry according to $\mathcal{D}_{\mathbb{Z}_p,\sigma}$. Define $\mathbf{B}$ as follows

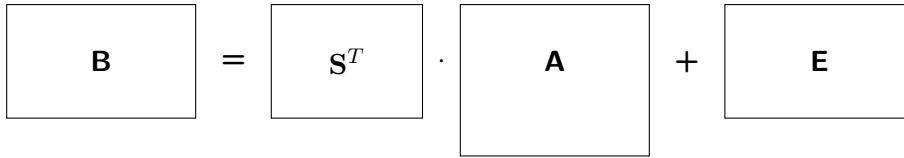

The private key is $\mathbf{S}$. The public key is $(\mathbf{A}, \mathbf{B}) \in \mathbb{Z}_p^{n \times m} \times \mathbb{Z}_p^{\ell \times m}$.

O. Regev, On lattices, learning with errors, random linear codes, and cryptography,
C. Peikert, V. Vaikuntanathan, B. Waters, A Framework for Efficient and Composable Oblivious Transfer

$PVW_{\mathcal{D}}$ PKE with Discrete Gaussian error distribution

Encryption: given an element of the message space $\mathbf{v} \in \mathbb{Z}_t^\ell$ and a public key $(\mathbf{A}, \mathbf{B})$, choose a vector $\mathbf{d} \in \{-r, \ldots, r\}^m$ uniformly at random. Compute

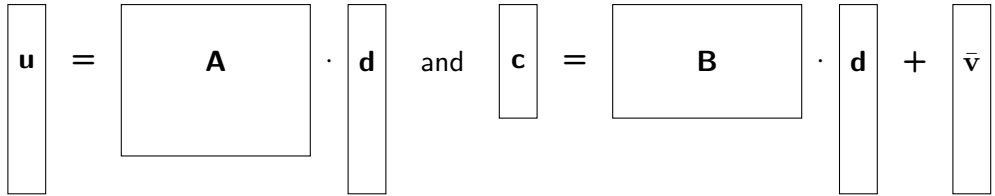

where $\bar{\mathbf{v}} = Encode(\mathbf{v}) \in \mathbb{Z}_p^\ell$. The ciphertext is $(\mathbf{u}, \mathbf{c}) \in \mathbb{Z}_p^n \times \mathbb{Z}_p^\ell$.

Decryption: given a ciphertext $(\mathbf{u}, \mathbf{c}) \in \mathbb{Z}_p^n \times \mathbb{Z}_p^\ell$ and a private key $\mathbf{S} \in \mathbb{Z}_p^{n \times \ell}$, compute $Decode(\mathbf{c} - \mathbf{S}^T\mathbf{u})$.

Main result

From LWE to FFP-NG

Let $\mathcal{A}$ be an FFP-NG adversary against the $PVW_{\mathcal{D}}$ encryption scheme, with error distribution $\mathcal{D}_{\mathbb{Z}_p,\sigma}$, denoted by Π. If the FFP-NG advantage $Adv_{\Pi}^{FFP-NG}(\mathcal{A}) \geq \delta$ is non-negligible in n, there exists an adversary $\mathcal{B}$ that solves the Decision LWE problem with error distribution $\mathcal{D}_{\mathbb{Z}_p,\hat{\sigma}}$, for $\hat{\sigma} = \sigma/\varphi$ such that

$$\mathrm{Adv}_{\Pi}^{\text{FFP-NG}}(\mathcal{A}) \leq \mathrm{Adv}_{\Pi}^{\text{LWE}}(\mathcal{B}) + \bar{\Gamma}(n,\varphi),$$

where $\bar{\Gamma}(n,\varphi) \in negl(n)$.

Security reduction from LWE to FFP-NG
Structure of the proof

1. Describe a decryption failure condition.
2. Get rid of the reduction modulo p.
3. Narrow down the possible strategies of the FFP-NG adversary.
4. Define an LWE adversary that exploits the FFP-NG adversary.
5. Analyze the winning probability of the adversary:

1. Decryption failure condition

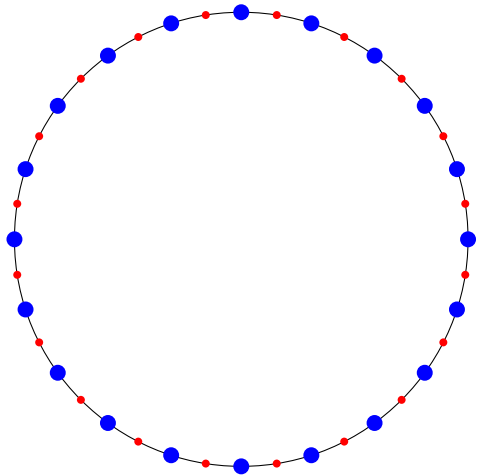

Blue dots are elements of $\mathbb{Z}_p$.
Lines between two red dots represent the range of a correct decryption.
If the value

$$|\langle \mathbf{E}, \mathbf{d} \rangle|$$

is too large, we get a decryption failure.
The decryption failure condition is independent of the message.

2. Ignoring the reduction modulo p

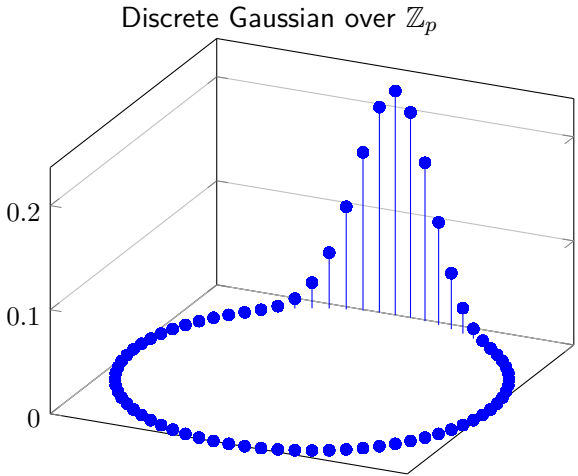

2. Ignoring the reduction modulo p

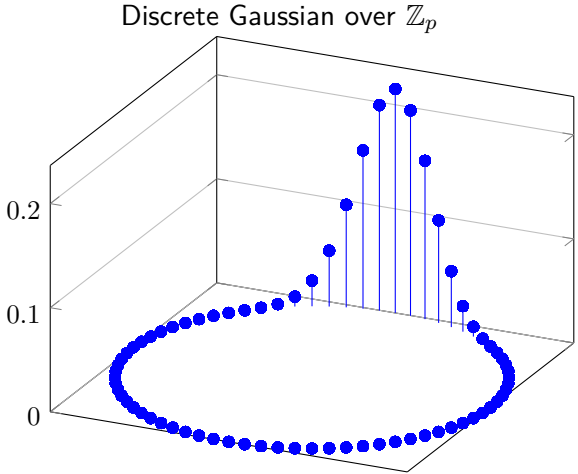

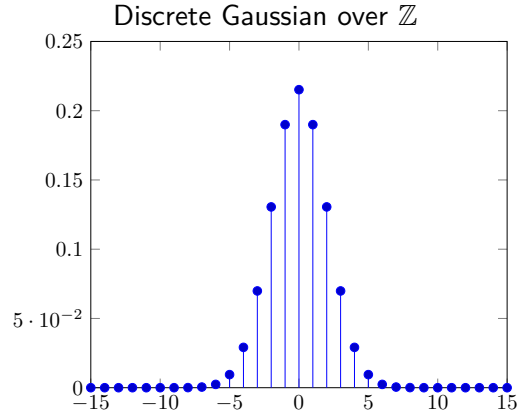

3. Possible strategies of the FFP-NG adversary

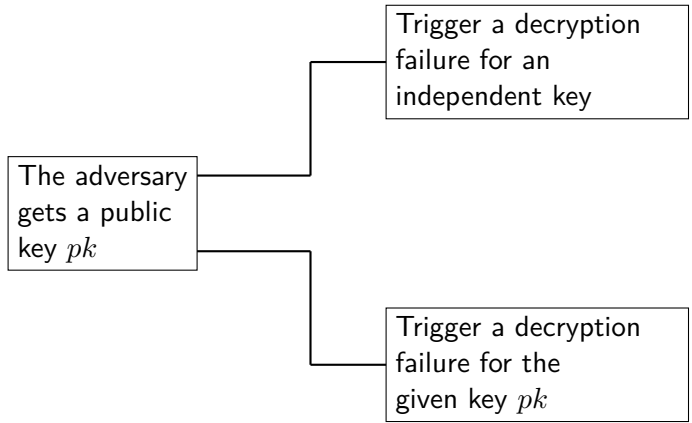

3. Possible strategies of the FFP-NG adversary

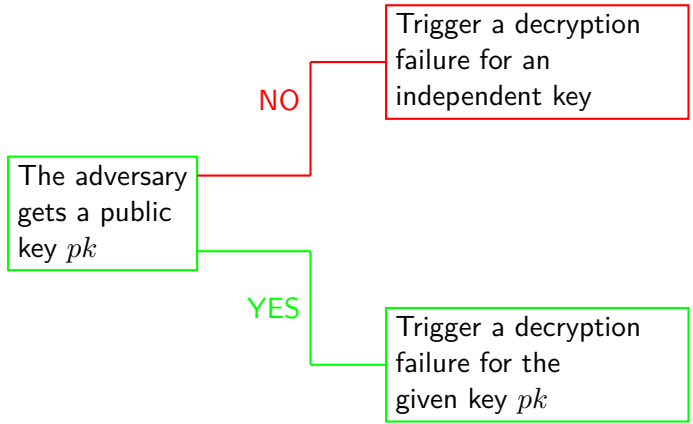

4. The LWE adversary

Given an FFP-NG adversary $\mathcal{A}$ with non negligible winning probability, we build the following LWE adversary $\mathcal{B}$. It takes as input a public key $pk = (\mathbf{A}, \mathbf{B})$ and does the following

1. Sets a value $\varphi > 1$
2. Samples an extra error $\mathbf{E}' \leftarrow \mathcal{D}_{\mathbb{Z}_p^m, \sigma'}$, where $\sigma' = \sigma'(\sigma, \varphi)$
3. Modifies the public $pk' \leftarrow (\mathbf{A}, \mathbf{B} + \mathbf{E}')$
4. $(v, \mathbf{d}) \leftarrow \mathcal{A}(pk')$
5. Sets the value $k = k(\mathbf{d})$
6. Defines the inner product threshold $= \frac{p}{2t}(1 + \frac{1}{2^w}) - k$
7. If $|\langle \mathbf{E}', \mathbf{d} \rangle| \geq$ the inner product threshold output LWE distribution
8. If $|\langle \mathbf{E}', \mathbf{d} \rangle| <$ the inner product threshold output Uniform distribution

5. Analyzing the winning probability of $\mathcal{B}$

Uniform distribution

Given a pair error/randomness $(\mathbf{E}', \mathbf{d})$. If $(\mathbf{A}, \mathbf{B})$ is sampled uniformly at random, then

$$\Pr_{(\mathbf{A},\mathbf{B})\leftarrow U}\left[\left|\langle \mathbf{E}', \mathbf{d}\rangle\right| \geq \text{inner product threshold}\right]$$

is negligible in n.

5. Analyzing the winning probability of $\mathcal{B}$

Uniform distribution

Given a pair error/randomness $(\mathbf{E}', \mathbf{d})$. If $(\mathbf{A}, \mathbf{B})$ is sampled uniformly at random, then

$$\Pr_{(\mathbf{A},\mathbf{B})\leftarrow U}\left[\left|\langle \mathbf{E}', \mathbf{d}\rangle\right| \geq \text{inner product threshold}\right]$$

is negligible in n.

LWE distribution

Given a pair error/randomness $(\mathbf{E}', \mathbf{d})$. If $(\mathbf{A}, \mathbf{B})$ is sampled according to the LWE distribution, then

$$\Pr_{(\mathbf{A},\mathbf{B})\leftarrow LWE}\left[\left|\langle \mathbf{E}', \mathbf{d}\rangle\right| \geq \text{inner product threshold}\right]$$

is NOT negligible in n.

5. Analyzing the winning probability of $\mathcal{B}$

From LWE to FFP-NG

Let $\mathcal{A}$ be an FFP-NG adversary against the $PVW_{\mathcal{D}}$ encryption scheme, with error distribution $\mathcal{D}_{\mathbb{Z}_p,\sigma}$, denoted by Π. If the FFP-NG advantage $Adv_{\Pi}^{FFP-NG}(\mathcal{A}) \geq \delta$ is non-negligible in n, there exists an adversary $\mathcal{B}$ that solves the Decision LWE problem with error distribution $\mathcal{D}_{\mathbb{Z}_p,\hat{\sigma}}$, for $\hat{\sigma} = \sigma/\varphi$ such that

$$\mathrm{Adv}_{\Pi}^{\text{FFP-NG}}(\mathcal{A}) \leq \mathrm{Adv}_{\Pi}^{\text{LWE}}(\mathcal{B}) + \bar{\Gamma}(n,\varphi),$$

where $\bar{\Gamma}(n,\varphi) \in negl(n)$.

Further steps

We are now looking at the following questions:

1. can we also prove the reduction by using as underlying scheme ones introduced by Lindner and Peikert that uses the LWE assumption for both key generation and encryption?
2. can we prove a similar reduction from Ring-LWE?

R. Lindner, C. Peikert, Better Key Sizes (and Attacks) for LWE-Based Encryption

Further steps

We are now looking at the following questions:

1. can we also prove the reduction by using as underlying scheme ones introduced by Lindner and Peikert that uses the LWE assumption for both key generation and encryption?
2. can we prove a similar reduction from Ring-LWE?

KYBER!!!

R. Lindner, C. Peikert, Better Key Sizes (and Attacks) for LWE-Based Encryption

THANK YOU FOR YOUR ATTENTION!
ANY QUESTIONS?