Cryptographic Divergences: New Techniques and New Applications.

Marc Abboud
Université de Rennes 1

Thomas Prest
PQShield

SCN Conference 2020

Table of Contents

Section 1

Divergence in security proofs

General setting

Notations

- Distributions will be denoted cursively.
- Q will be the number of queries, λ the security level.
- $\mathcal{P}$ will be a "real" distribution and $\mathcal{Q}$ an "ideal" one.

Goal

Scheme $\mathcal{S}$, adversary $\mathcal{A}$.

$$\mathrm{adv}_{\mathcal{A}}(\mathcal{S}^{\mathcal{Q}}) \leq 2^{-\lambda-c} \Rightarrow \mathrm{adv}_{\mathcal{A}}(\mathcal{S}^{\mathcal{P}}) \leq 2^{-\lambda}$$

Subsection 1

Statistical distance and Renyi Divergence

Statistical distance

Definition 1 (Statistical distance)

Let $\mathcal{P}, \mathcal{Q}$ be two distributions over a space X.

$$\Delta_{\mathrm{SD}}(\mathcal{P}, \mathcal{Q}) := \frac{1}{2} \sum_{x \in X} |\mathcal{P}(x) - \mathcal{Q}(x)|$$

Statistical distance

Definition 1 (Statistical distance)

Let $\mathcal{P}, \mathcal{Q}$ be two distributions over a space X.

$$\Delta_{\mathrm{SD}}(\mathcal{P}, \mathcal{Q}) := \frac{1}{2} \sum_{x \in X} |\mathcal{P}(x) - \mathcal{Q}(x)|$$

Proposition 1

For all event $E \subset X$,

$$\mathcal{P}(E) \leq \mathcal{Q}(E) + \Delta_{\mathrm{SD}}(\mathcal{P}, \mathcal{Q})$$

Statistical distance

Definition 1 (Statistical distance)

Let $\mathcal{P}, \mathcal{Q}$ be two distributions over a space X.

$$\Delta_{\mathrm{SD}}(\mathcal{P}, \mathcal{Q}) := \frac{1}{2} \sum_{x \in X} |\mathcal{P}(x) - \mathcal{Q}(x)|$$

Proposition 1

For all event $E \subset X$,

$$\mathcal{P}(E) \leq \mathcal{Q}(E) + \Delta_{\mathrm{SD}}(\mathcal{P}, \mathcal{Q})$$

This requires $\Delta_{\mathrm{SD}}(\mathcal{P}, \mathcal{Q}) \leq 2^{-\lambda}$.

Renyi's Divergence

Definition 2 (Renyi's Divergence)

Suppose moreover than $\mathrm{Supp}\,\mathcal{P} \subset \mathrm{Supp}\,\mathcal{Q}$, *then for* $\alpha \in [1, \infty]$, *one defines*

- $\forall 1 < \alpha < \infty, \quad R_\alpha(\mathcal{P}; \mathcal{Q}) := \left(\sum_{x \in \mathrm{Supp}\,\mathcal{Q}} \frac{\mathcal{P}(x)^\alpha}{\mathcal{Q}(x)^{\alpha-1}} \right)^{1/\alpha - 1}$
- $R_\infty(\mathcal{P}; \mathcal{Q}) := \max_{x \in \mathrm{Supp}\,\mathcal{Q}} \frac{\mathcal{P}(x)}{\mathcal{Q}(x)}$
- $R_1(\mathcal{P}; \mathcal{Q}) = e^{D_{\mathrm{KL}}(\mathcal{P}; \mathcal{Q})}$ *where* $D_{\mathrm{KL}}(\mathcal{P}; \mathcal{Q}) = \sum_{x \in X} \mathcal{P}(x) \log(\frac{\mathcal{P}(x)}{\mathcal{Q}(x)})$

Properties

Proposition 2

- *For all event* $E \subset X$*, one has* $\mathcal{P}(E) \leq \mathcal{Q}(E)^{1-\frac{1}{\alpha}} R_\alpha(\mathcal{P};\mathcal{Q})^{1-1/\alpha}$.
- *If* $\mathcal{P}_1, \mathcal{Q}_1, ..., \mathcal{P}_n, \mathcal{Q}_n$ *are independent distributions, then* $R_\alpha(\mathcal{P}_1 \times \cdots \times \mathcal{P}_n, \mathcal{Q}_1 \times \cdots \times \mathcal{Q}_n) = R_\alpha(\mathcal{P}_1, \mathcal{Q}_1) \times \cdots \times R_\alpha(\mathcal{P}_n, \mathcal{Q}_n)$.

Renyi divergence in security proofs

Lemma 1

Let $\mathcal{A}$ be an adversary and $\mathcal{S}$ be a search problem such that

$$\mathrm{adv}_{\mathcal{A}}(\mathcal{S}^{\mathcal{Q}}) \leq 2^{-\lambda}$$

Suppose $R_\lambda(\mathcal{P};\mathcal{Q}) \leq 1 + 1/Q$, then

$$\mathrm{adv}_{\mathcal{A}}(\mathcal{S}^{\mathcal{P}}) \leq 2^{-(\lambda-1)} \cdot e$$

i.e at most 3 bits of security are lost when going from $\mathcal{Q}$ to $\mathcal{P}$.

Renyi divergence in security proofs

Lemma 1

Let $\mathcal{A}$ be an adversary and $\mathcal{S}$ be a search problem such that

$$\mathrm{adv}_{\mathcal{A}}(\mathcal{S}^{\mathcal{Q}}) \leq 2^{-\lambda}$$

Suppose $R_\lambda(\mathcal{P};\mathcal{Q}) \leq 1 + 1/Q$, then

$$\mathrm{adv}_{\mathcal{A}}(\mathcal{S}^{\mathcal{P}}) \leq 2^{-(\lambda-1)} \cdot e$$

i.e at most 3 bits of security are lost when going from $\mathcal{Q}$ to $\mathcal{P}$.

$\Rightarrow$ tightness requires $R_\alpha = 1 + \frac{1}{Q}$ and typically $\alpha = \lambda$.

Renyi divergence in security proofs

Lemma 1

Let $\mathcal{A}$ be an adversary and $\mathcal{S}$ be a search problem such that

$$\mathrm{adv}_{\mathcal{A}}(\mathcal{S}^{\mathcal{Q}}) \leq 2^{-\lambda}$$

Suppose $R_\lambda(\mathcal{P};\mathcal{Q}) \leq 1 + 1/Q$, then

$$\mathrm{adv}_{\mathcal{A}}(\mathcal{S}^{\mathcal{P}}) \leq 2^{-(\lambda-1)} \cdot e$$

i.e at most 3 bits of security are lost when going from $\mathcal{Q}$ to $\mathcal{P}$.

$\Rightarrow$ tightness requires $R_\alpha = 1 + \frac{1}{Q}$ and typically $\alpha = \lambda$.

Not for decision problems

If E is the event of breaking a decision problem, then we want $\mathcal{P}(E) = \frac{1}{2} + \epsilon$ with $\epsilon << 1$.

Entropy

Definition 2 (Rényi entropy)

Let $\alpha \in [1, +\infty]$ and $\mathcal{X}$ be a discrete distribution. The α-entropy (or Rényi entropy) of $\mathcal{X}$ is:

$$H_\alpha(\mathcal{X}) = \begin{cases} -\sum_x \Pr[\mathcal{X} = x] \log_2 \Pr[\mathcal{X} = x] & \text{if } \alpha = 1 \\ \frac{1}{1-\alpha} \log_2 \left(\sum_x \Pr[\mathcal{X} = x]^\alpha \right) & \text{if } 1 < \alpha < \infty \\ -\max_x \log_2 \Pr[\mathcal{X} = x] & \text{if } \alpha = \infty \end{cases}$$

Entropy

Definition 2 (Rényi entropy)

Let $\alpha \in [1, +\infty]$ and $\mathcal{X}$ be a discrete distribution. The α-entropy (or Rényi entropy) of $\mathcal{X}$ is:

$$H_\alpha(\mathcal{X}) = \begin{cases} -\sum_x \Pr[\mathcal{X} = x] \log_2 \Pr[\mathcal{X} = x] & \text{if } \alpha = 1 \\ \frac{1}{1-\alpha} \log_2 \left(\sum_x \Pr[\mathcal{X} = x]^\alpha \right) & \text{if } 1 < \alpha < \infty \\ -\max_x \log_2 \Pr[\mathcal{X} = x] & \text{if } \alpha = \infty \end{cases}$$

Example

H_1 is the Shannon entropy, H_2 is the collision entropy and H_∞ the min-entropy.

Subsection 2

A Class of Divergences with Peculiar Properties

RE_α-divergence

Definition 3 (RE_α-divergence)

Let $\mathcal{P}, \mathcal{Q}$ be two distributions over a countable space X. Suppose moreover that $\mathrm{Supp}\,\mathcal{P} \subset \mathrm{Supp}\,\mathcal{Q}$. One defines the RE_α-divergence as

- $\forall \alpha \geq 1, \quad \mathrm{RE}_\alpha(\mathcal{P};\mathcal{Q}) = \left(\sum_{x\in X} \mathcal{Q}(x) \left|\frac{\mathcal{P}(x)}{\mathcal{Q}(x)} - 1\right|^\alpha\right)^{\frac{1}{\alpha}}$.
- $\mathrm{RE}_\infty(\mathcal{P};\mathcal{Q}) = \mathrm{RE}(\mathcal{P};\mathcal{Q}) = \max_{x\in \mathrm{Supp}\,\mathcal{Q}} \left|\frac{\mathcal{P}(x)}{\mathcal{Q}(x)} - 1\right|$.

Notice that $\mathrm{RE}_1 = 2\Delta_{\mathrm{SD}}$.

Properties

Proposition 3

- *For all event* $E \subset X, \mathcal{P}(E) \leq \mathcal{Q}(E)^{1-1/\alpha}(1 + \mathrm{RE}_\alpha(\mathcal{P}; \mathcal{Q}))$.
- $R_\alpha(\mathcal{P}; \mathcal{Q}) \leq (1 + \mathrm{RE}_\alpha(\mathcal{P}; \mathcal{Q}))^{\frac{\alpha-1}{\alpha}}$.

Properties

Proposition 3

- *For all event* $E \subset X, \mathcal{P}(E) \leq \mathcal{Q}(E)^{1-1/\alpha}(1 + \mathrm{RE}_\alpha(\mathcal{P}; \mathcal{Q}))$.
- $R_\alpha(\mathcal{P}; \mathcal{Q}) \leq (1 + \mathrm{RE}_\alpha(\mathcal{P}; \mathcal{Q}))^{\frac{\alpha-1}{\alpha}}$.

This requires also $\mathrm{RE}_\alpha(\mathcal{P}; \mathcal{Q}) = \frac{1}{Q}$ and typically $\alpha = \lambda$.

Properties

Proposition 3

- *For all event* $E \subset X, \mathcal{P}(E) \leq \mathcal{Q}(E)^{1-1/\alpha}(1 + \mathrm{RE}_\alpha(\mathcal{P}; \mathcal{Q}))$.
- $R_\alpha(\mathcal{P}; \mathcal{Q}) \leq (1 + \mathrm{RE}_\alpha(\mathcal{P}; \mathcal{Q}))^{\frac{\alpha-1}{\alpha}}$.

This requires also $\mathrm{RE}_\alpha(\mathcal{P}; \mathcal{Q}) = \frac{1}{Q}$ and typically $\alpha = \lambda$.

Proposition 4 ([Pre17])

Let $\alpha \geq 1$, $\mathcal{P}, \mathcal{Q}$ *be two distributions over a space* X. *Then,*

$$R_\alpha(\mathcal{P}; \mathcal{Q}) \lessapprox 1 + \frac{\alpha \, \mathrm{RE}(\mathcal{P}; \mathcal{Q})^2}{2}$$

What's new ?

What's new ? RE_α has also an amplification property . . .

Amplification property

Proposition 5 (Amplification property for the statistical distance [DS16])

Let $f : X \to X$ be a random function over a space X such that there exists $\delta > 0$

$$\forall a, b \in X, \Delta_{\mathrm{SD}}(f(a), f(b)) \leq \delta.$$

Then, for all $k \geq 1$,

$$\Delta_{\mathrm{SD}}(f^k(a), f^k(b)) \leq \delta^k$$

Amplification property

Proposition 6 (Amplification property)

Let $f : X \to X$ be a random function over a space X. Suppose that $\forall a, b \in X, \mathrm{Supp}\, f(a) = \mathrm{Supp}\, f(b)$ and that there exists $\delta > 0$

$$\forall a, b \in X, \mathrm{RE}_\alpha(f(a), f(b)) \leq \delta.$$

Then, for all $k \geq 1$,

$$\mathrm{RE}_\alpha(f^k(a), f^k(b)) \leq \delta^k$$

Section 2

Security for prime number generators

Ideally, one would use the naive random generator:

Naive algorithm

Parameters: x.
Output: a prime number $p \leq x$.
repeat
  Sample $p \hookleftarrow \{1, ..., x\}$
until p is prime
Return p

But this is very costly in term of time and randomness.
We will denote by $\mathcal{U}$ the uniform distribution over primes.

Subsection 1

PrimeInc

PrimeInc generator: Naive version

PrimeInc generator: First version [BD93]

Parameters: x
Output: a prime number between $x/2$ and x
Sample m in $[x/2; x]$
repeat
  $m \leftarrow m + 2$
until m is prime

PrimeInc generator: Naive version

PrimeInc generator: First version [BD93]

Parameters: x
Output: a prime number between $x/2$ and x
Sample m in $[x/2; x]$
repeat
  $m \leftarrow m + 2$
until m is prime

Output distribution won't be uniform

then, $\Pr(q) << \Pr(p)$.

PrimeInc generator: Naive version

PrimeInc generator: First version [BD93]

Parameters: x
Output: a prime number between $x/2$ and x
Sample m in $[x/2; x]$
repeat
  $m \leftarrow m + 2$
until m is prime

Output distribution won't be uniform

then, $\Pr(q) << \Pr(p)$.

PrimeInc generator: Naive version

PrimeInc generator: First version [BD93]

Parameters: x
Output: a prime number between $x/2$ and x
Sample m in $[x/2; x]$
repeat
  $m \leftarrow m + 2$
until m is prime

Output distribution won't be uniform

then, $\Pr(q) << \Pr(p)$.

PrimeInc generator

PrimeInc generator [BD93]

Parameters: $x, s = c \log x$
Output: a prime number between $x/2$ and x
Sample p in $[x/2; x]$
for $i = 1$ to s **do**
  if p is prime **then**
    Return p
  else
    $p \leftarrow p + 2$
  end if
end for
If no prime has been found after s steps, return "failure".

Security Analysis

Proposition 7 (BD93)

Let $\mathcal{P}$ be the output distribution of the PrimeInc generator. Then, under the prime r-tuple conjecture, one has

$$\frac{H_1(\mathcal{P})}{H_1(\mathcal{U})} \xrightarrow[x\to\infty]{} 1.$$

Security Analysis

Proposition 8 (Asymptotical security of PrimeInc)

Denote by X the output distribution of PrimeInc and by U the uniform distribution over the prime numbers. Under the prime r-tuple conjecture, one has

$$R_\infty(X, U) \leq 2c(1 + o_{c,x}(1))$$

Equivalently for all $\alpha \geq 2$,

$$H_\alpha(\mathcal{P}) \geq H_\alpha(\mathcal{U}) - \log(2c) + o_{c,x}(1).$$

Practical Implications

- Security against collision attacks.
- With a constant number of calls, a scheme secure with a uniform generator remains secure when using PrimeInc instead.

Subsection 2

Fouque and Tibouchi's generator

Fouque and Tibouchi's generator

Fouque and Tibouchi's generator [FT19]

Parameters: $x, \epsilon, q \simeq x^{1-\epsilon}$
Output: a prime number $p \leq x$
Sample $a \xleftarrow{\$} (\mathbf{Z}/q\mathbf{Z})^*$
repeat
$t \xleftarrow{\$} \left\{0, ..., \lfloor \frac{x-a}{q} \rfloor \right\}$
until $p = a + tq$ is prime

Fouque and Tibouchi's generator

Fouque and Tibouchi's generator [FT19]

Parameters: $x, \epsilon, q \simeq x^{1-\epsilon}$
Output: a prime number $p \leq x$
Sample $a \xleftarrow{\$} (\mathbf{Z}/q\mathbf{Z})^*$
repeat
$t \xleftarrow{\$} \left\{0, ..., \lfloor \frac{x-a}{q} \rfloor \right\}$
until $p = a + tq$ is prime

Here we want ϵ to be as low as possible so that fewer randomness is needed to find a prime number. Indeed, the average entropy consumption of the generator is

$$(\epsilon + o(1)) \cdot \frac{\phi(q)}{q} \cdot \frac{(\log x)^2}{\log 2}.$$

Security analysis

Let $\mathcal{X}$ be the output distribution of Fouque and Tibouchi's generator.

Proposition 9 ([FT19])

Under the Friedlander-Granville-Montgomery conjecture, one has

$$\Delta_{\mathrm{SD}}(\mathcal{X},\mathcal{U}) \ll \frac{\log x}{x^{\epsilon/4}}$$

Security analysis

Let $\mathcal{X}$ be the output distribution of Fouque and Tibouchi's generator.

Proposition 9 ([FT19])

Under the Friedlander-Granville-Montgomery conjecture, one has

$$\Delta_{\mathrm{SD}}(\mathcal{X},\mathcal{U}) \ll \frac{\log x}{x^{\epsilon/4}}$$

We show that

Proposition 10

Under the Friedlander-Granville-Montgomery conjecture, one has

- $\mathrm{RE}(\mathcal{X},\mathcal{U}_{|\operatorname{Supp}\mathcal{X}}) \ll \frac{\log x}{x^{\epsilon/4}}$ *and* $\Pr[(\operatorname{Supp}\mathcal{X})^c] \leq \frac{\log(x)^2}{x}$.
- $\mathrm{RE}_\alpha(\mathcal{X},\mathcal{U}) \ll \frac{\log x}{x^{\epsilon/4}} + \left(\frac{\log x}{x^2}\right)^\alpha$.

Practical Implications

- A statistical distance argument requires that $\frac{\log x}{x^{\epsilon/4}} < 2^{-\lambda}$ which yields

$$\epsilon \geq \frac{2\lambda \log\log x}{\log x} \simeq 0.36$$

- A Renyi-based argument requires $\lambda Q \frac{(\log x)^2}{2x^{\epsilon/2}} \leq 1$ which yields

$$\epsilon \geq \frac{2\log(\lambda Q)\log\log x}{\log x} \simeq 0.076$$

So we gain a factor roughly 5.

Section 3

Circuit private FHE

Subsection 1

Putting Washing Machine in Eco mode.

Context

Setup for FHE

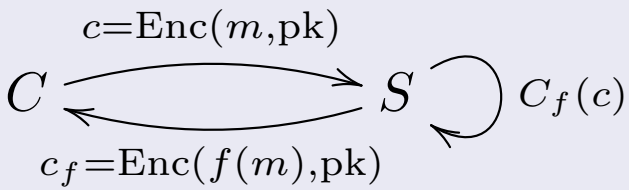

This is secure if the server doesn't learn anything about m or $f(m)$.

Context

Setup for FHE

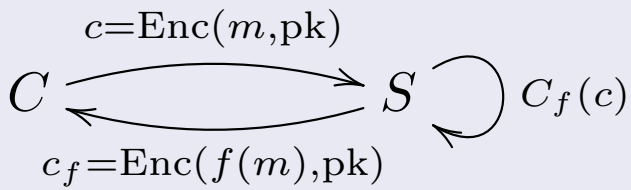

This is secure if the server doesn't learn anything about m or $f(m)$.

We want the additional property that the client does not learn anything from the circuit used by the server.

Context

Setup for FHE

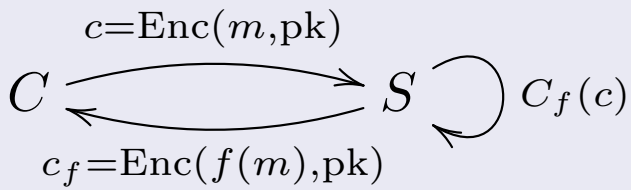

This is secure if the server doesn't learn anything about m or $f(m)$.

We want the additional property that the client does not learn anything from the circuit used by the server.

Definition 4 (Circuit-privacy)

A FHE scheme is circuit private if there exists an algorithm M such that

- *Given $f(m)$ and* pk, *M outputs a ciphertext c' such that $c' = \mathrm{Enc}(f(m), pk)$.*
- $c' \simeq c_f$.

Sanitization of ciphertext

Question: How to create such an algorithm M ?

Sanitization of ciphertext

Question: How to create such an algorithm M ?

Solution: The washing machine [DS16]

The server has at its disposal a random function (Washing Machine)

$$\mathrm{Wash} : \mathcal{C} \to \mathcal{C}$$

such that

$$\Delta_{\mathrm{SD}}(\mathrm{Wash}(c_1, \mathrm{pk}), \mathrm{Wash}(c_2, \mathrm{pk})) \leq \delta.$$

and we will do k round of washing machine to use the amplification property.

Sanitization of ciphertext

Question: How to create such an algorithm M ?

Solution: The washing machine [DS16]

The server has at its disposal a random function (Washing Machine)

$$\mathrm{Wash} : \mathcal{C} \to \mathcal{C}$$

such that

$$\Delta_{\mathrm{SD}}(\mathrm{Wash}(c_1, \mathrm{pk}), \mathrm{Wash}(c_2, \mathrm{pk})) \leq \delta.$$

and we will do k round of washing machine to use the amplification property.

$\Rightarrow$ The security analysis can also be done with the RE_α-divergence leading to fewer calls to the washing machine.

Rerandomizing LWE-ciphertexts: Encryption of one bit

Notation

- $\mathrm{LWE}_s^q(\mu, \eta) = \left\{(\mathbf{a}, \langle \mathbf{a}, s\rangle + \mu \cdot \lfloor q/2 \rfloor + e) \in \mathbf{Z}_q^{n+1} \;\middle|\; |e| < \eta q\right\}$.

Rerandomizing LWE-ciphertexts: Encryption of one bit

Notation

- $\mathrm{LWE}_s^q(\mu, \eta) = \left\{(\mathbf{a}, \langle \mathbf{a}, s\rangle + \mu \cdot \lfloor q/2 \rfloor + e) \in \mathbf{Z}_q^{n+1} \;\middle|\; |e| < \eta q\right\}$.
- Correctness of decryption is ensured only if $\eta < 1/4$.

Rerandomizing LWE-ciphertexts: Encryption of one bit

Notation

- $\mathrm{LWE}_s^q(\mu, \eta) = \left\{(\mathbf{a}, \langle \mathbf{a}, s\rangle + \mu \cdot \lfloor q/2 \rfloor + e) \in \mathbf{Z}_q^{n+1} \;\middle|\; |e| < \eta q\right\}$.
- Correctness of decryption is ensured only if $\eta < 1/4$.
- Wash : $\mathrm{LWE}_s^q(\mu, \Upsilon) \underset{\text{Rerand}}{\overset{\text{Refresh}}{\rightleftarrows}} \mathrm{LWE}_s^q(\mu, \eta) \quad (\eta \ll \Upsilon)$.

Rerandomizing LWE-ciphertexts: Encryption of one bit

Notation

- $\mathrm{LWE}_s^q(\mu,\eta) = \left\{(\mathbf{a}, \langle \mathbf{a}, s\rangle + \mu \cdot \lfloor q/2 \rfloor + e) \in \mathbf{Z}_q^{n+1} \;\middle|\; |e| < \eta q\right\}$.
- Correctness of decryption is ensured only if $\eta < 1/4$.
- Wash : $\mathrm{LWE}_s^q(\mu,\Upsilon) \underset{\text{Rerand}}{\overset{\text{Refresh}}{\rightleftarrows}} \mathrm{LWE}_s^q(\mu,\eta) \quad (\eta \ll \Upsilon)$.

Wash function

- Refresh is the bootstrapping process to reduce the noise of the ciphertext. This step is very time-consuming.
- We focus on Rerand.

Rerand

Suppose that the public key pk contains $\ell = O(n \log q)$ encryptions of 0 that we call rerandomizers $r_i = \mathrm{Enc}(0, \mathrm{pk})$.

Rerand

Suppose that the public key pk contains $\ell = O(n \log q)$ encryptions of 0 that we call rerandomizers $r_i = \mathrm{Enc}(0, \mathrm{pk})$.

Take $c = (\mathbf{a}, \langle \mathbf{a}, s \rangle + \mu \lfloor \frac{q}{2} \rfloor + e)$, then

Rerand

Suppose that the public key pk contains $\ell = O(n \log q)$ encryptions of 0 that we call rerandomizers $r_i = \mathrm{Enc}(0, \mathrm{pk})$.

Take $c = (\mathbf{a}, \langle \mathbf{a}, s \rangle + \mu \lfloor \frac{q}{2} \rfloor + e)$, then

$$\mathrm{Rerand}(c, \mathrm{pk}) = \underbrace{c + \sum_{i=1}^{\ell} \varepsilon_i r_i}_{=:c'} + (\mathbf{0}, f)$$

where $\varepsilon_i \hookleftarrow \{0, \pm 1\}$ and $f \in \mathbf{Z}_q$ is sampled from a chosen distribution.

Rerand

Suppose that the public key pk contains $\ell = O(n \log q)$ encryptions of 0 that we call rerandomizers $r_i = \mathrm{Enc}(0, \mathrm{pk})$.

Take $c = (\mathbf{a}, \langle \mathbf{a}, s \rangle + \mu \lfloor \frac{q}{2} \rfloor + e)$, then

$$\mathrm{Rerand}(c, \mathrm{pk}) = \underbrace{c + \sum_{i=1}^{\ell} \varepsilon_i r_i}_{=:c'} + (\mathbf{0}, f)$$

where $\varepsilon_i \hookleftarrow \{0, \pm 1\}$ and $f \in \mathbf{Z}_q$ is sampled from a chosen distribution.

Write $c' = (\mathbf{a}', \langle \mathbf{a}', s \rangle + \mu \lfloor \frac{q}{2} \rfloor + e')$. The only leakage of Rerand comes from the noise $e' + f$. Notice that $|e'| \leq (\ell + 1)\eta q$.

How to sample f to minimize the number of iterations k ?

How to sample f to minimize the number of iterations k ?

Uniform Distribution using statistical distance

In [DS16], $f \hookleftarrow \mathcal{U}([-B, B])$.

- To ensure correctness, one needs $B + (\ell + 1)\eta q < q/4$.
- To ensure security one needs $\Delta_{\mathrm{SD}}(\mathrm{Rerand}(c, \mathrm{pk}), \mathrm{Rerand}(c', \mathrm{pk})) \leq 2^{-\lambda/k}$.

This leads to $k = O(\lambda)$.

Gaussian distribution and relative error

In the paper, $f \hookleftarrow G_\sigma$.
Pick T such that $\Pr(|t| > T\sigma; t \hookleftarrow G_\sigma) \leq 2^{-\lambda}$.

- To ensure correctness, one needs $(\ell+1)\eta q + T\sigma < q/4$.
- To ensure security, one needs

$$\mathrm{RE}(\mathrm{Rerand}(c, \mathrm{pk}), \mathrm{Rerand}(c', \mathrm{pk})) \leq Q^{-1/2k}$$

  in order to apply Proposition 4.

This leads to $k = O(\log Q)$.

Conclusion

This procedure divides the number of iterations needed by $\frac{\lambda}{\log Q}$.

Thank you !

Jørgen Brandt and Ivan Damgård.

On generation of probable primes by incremental search.

In Ernest F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 358–370. Springer, Heidelberg, August 1993.

Léo Ducas and Damien Stehlé.

Sanitization of FHE ciphertexts.

In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 294–310. Springer, Heidelberg, May 2016.

Pierre-Alain Fouque and Mehdi Tibouchi.

Close to uniform prime number generation with fewer random bits.

IEEE Trans. Information Theory, 65(2):1307–1317, 2019.

Thomas Prest.

Sharper bounds in lattice-based cryptography using the Rényi divergence.

In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 347–374. Springer, Heidelberg, December 2017.