On the statistical leak of the GGH13 multilinear map and its variants

Léo Ducas[1], Alice Pellet--Mary[2]

[1]Cryptology Group, CWI, Amsterdam

[2]LIP, ENS de Lyon.

25th April, 2017

European Research Council
Established by the European Commission

ENS DE LYON

Introduction

In this talk:

- Focus on the GGH13 multilinear map

Introduction

In this talk:

- Focus on the GGH13 multilinear map
- Classical attacks: zeroizing attacks
  $\Rightarrow$ main application of GGH today: obfuscators

Introduction

In this talk:

- Focus on the GGH13 multilinear map
- Classical attacks: zeroizing attacks
  $\Rightarrow$ main application of GGH today: obfuscators
- Contribution: analyze averaging attacks
  - In some case, we have a complete attack against GGH.
  - In some other cases, we get some leaked information.

Table of Contents

History of multilinear maps (until February 2015)

- 2000 Joux introduces bilinear maps (pairings) for cryptographic uses.
- 2003 Boneh and Silverberg introduce the concept of multilinear maps.
- $\geq$ 2003 Many applications.
- 2013 Garg, Gentry and Halevi publish the first candidate multilinear map (GGH13 map).
- 2013 Garg et al. publish the first candidate obfuscator, using the GGH13 map.
- 2013 Coron, Lepoint and Tibouchi propose another candidate multilinear map, relying on integers (CLT map).
- 2015 Gentry, Gorbunov and Halevi propose a graph-induced multilinear map (GGH15 map).

Cryptographic multilinear maps

Definition: κ-multilinear map

Different levels of encodings, from 0 to κ.
Denote by $C(a, i)$ a level-i encoding of the message a.
Level-0 encoding: a plaintext (message not encoded).

Addition: $\text{Add}(C(a_1, i), C(a_2, i)) = C(a_1 + a_2, i)$.

Multiplication: $\text{Mult}(C(a_1, i), C(a_2, j)) = C(a_1 \cdot a_2, i + j)$.

Zero-test: $\text{Zero-test}(C(a, \kappa)) = \text{True iff } a = 0$.

Cryptographic multilinear maps

Definition: κ-multilinear map

Different levels of encodings, from 0 to κ.
Denote by $C(a, i)$ a level-i encoding of the message a.
Level-0 encoding: a plaintext (message not encoded).

Addition: $\text{Add}(C(a_1, i), C(a_2, i)) = C(a_1 + a_2, i)$.

Multiplication: $\text{Mult}(C(a_1, i), C(a_2, j)) = C(a_1 \cdot a_2, i + j)$.

Zero-test: $\text{Zero-test}(C(a, \kappa)) = \text{True iff } a = 0$.

Security: What should be hard for a cryptographic multilinear map?

Application to multipartite key-exchange

Objective: $\kappa + 1$ users want to agree on a shared secret s.
Let D be a distribution over the message space.

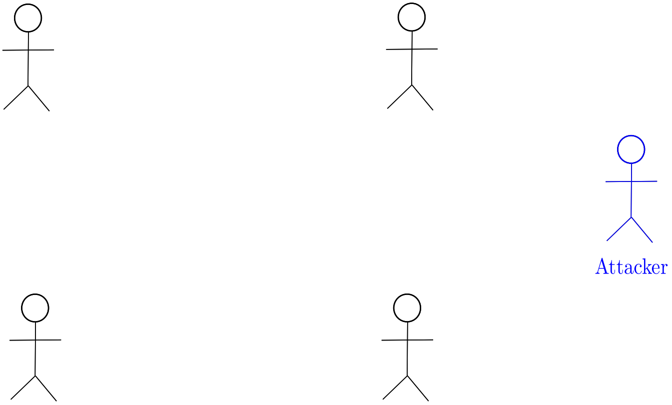

Application to multipartite key-exchange

Objective: $\kappa + 1$ users want to agree on a shared secret s.
Let D be a distribution over the message space.

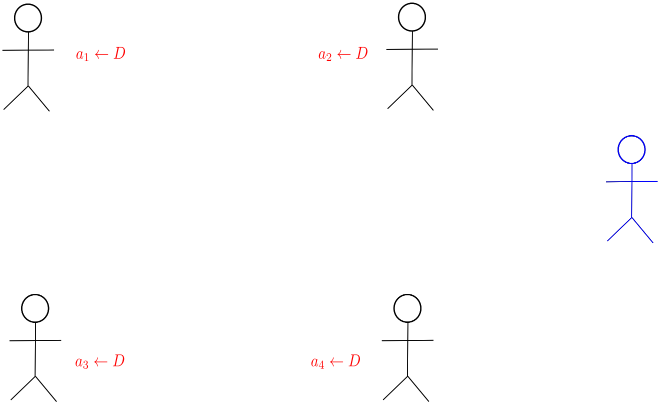

Application to multipartite key-exchange

Objective: $\kappa + 1$ users want to agree on a shared secret s.
Let D be a distribution over the message space.

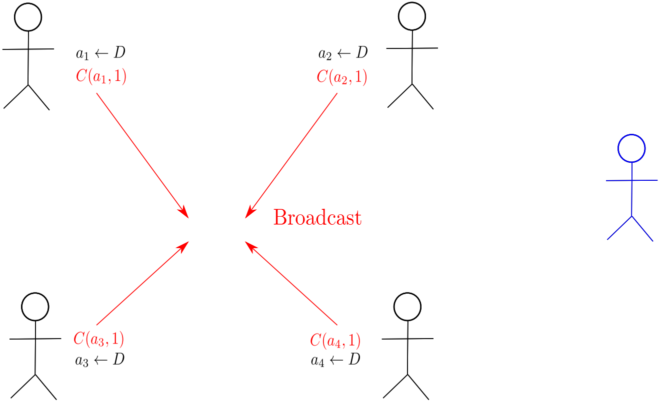

Application to multipartite key-exchange

Objective: $\kappa + 1$ users want to agree on a shared secret s.
Let D be a distribution over the message space.

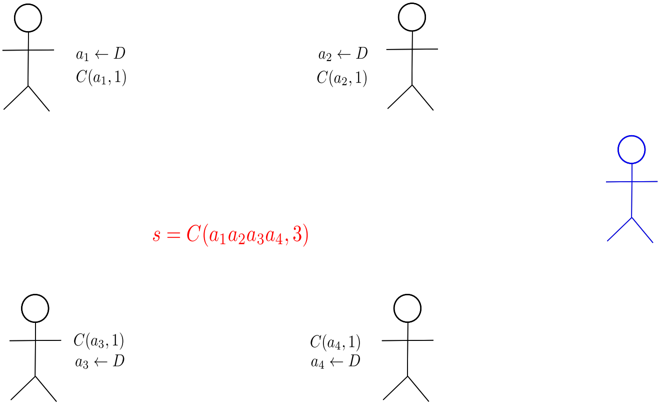

Application to multipartite key-exchange

Objective: $\kappa + 1$ users want to agree on a shared secret s.
Let D be a distribution over the message space.

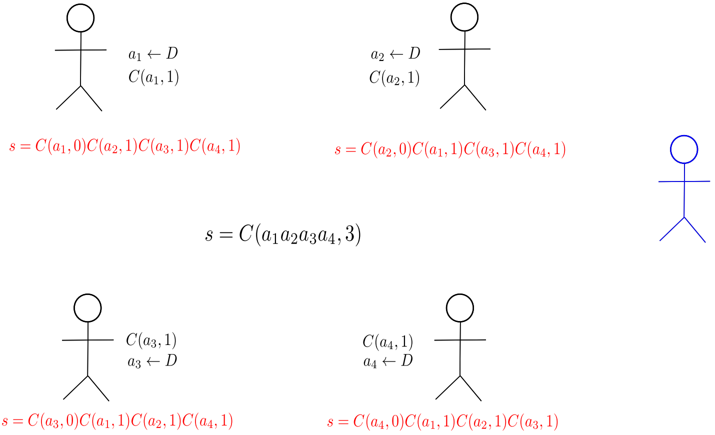

Application to multipartite key-exchange

Objective: $\kappa + 1$ users want to agree on a shared secret s.
Let D be a distribution over the message space.

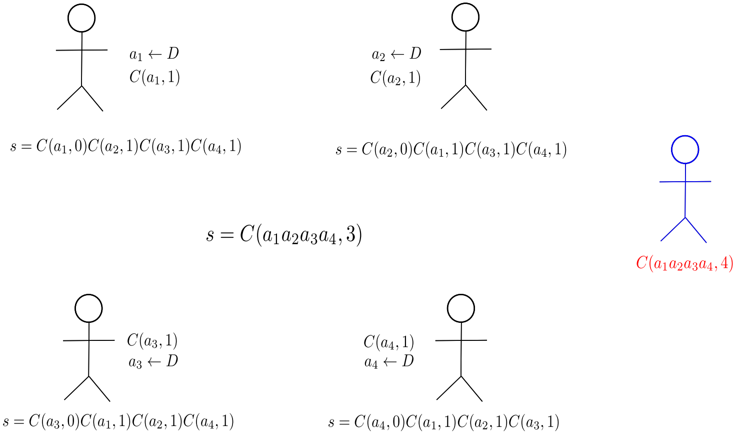

The GGH13 multilinear map

- Define $R = \mathbb{Z}[X]/(X^n + 1)$ with $n = 2^k$.

The GGH13 multilinear map

- Define $R = \mathbb{Z}[X]/(X^n + 1)$ with $n = 2^k$.
- Sample g a "small" element in R.
  $\Rightarrow$ the plaintext space is $\mathcal{P} = R/\langle g \rangle$.

The GGH13 multilinear map

- Define $R = \mathbb{Z}[X]/(X^n + 1)$ with $n = 2^k$.
- Sample g a "small" element in R.
  $\Rightarrow$ the plaintext space is $\mathcal{P} = R/\langle g \rangle$.
- Sample q a "large" integer.
  $\Rightarrow$ the encoding space is $R_q = R/(qR) = \mathbb{Z}_q[X]/(X^n + 1)$.

Notation

We write $[r]_q$ or $[r]$ the elements in R_q, and r (without $[\cdot]$) the elements in R.

The GGH13 multilinear map: encodings

- Sample z uniformly in R_q.
- **Encoding:** An encoding of a at level i is

$$u = [(a + rg)z^{-i}]_q$$

where $a + rg$ is a small element in $a + \langle g \rangle$.

The GGH13 multilinear map: encodings

- Sample z uniformly in R_q.
- **Encoding:** An encoding of a at level i is

$$u = [(a + rg)z^{-i}]_q$$

where $a + rg$ is a small element in $a + \langle g \rangle$.

Addition and multiplication

Addition:

$$[(a_1 + r_1 g)z^{-i}]_q + [(a_2 + r_2 g)z^{-i}]_q = [(a_1 + a_2 + r' g)z^{-i}]_q.$$

Multiplication:

$$[(a_1 + r_1 g)z^{-i}]_q \cdot [(a_2 + r_2 g)z^{-j}]_q = [(a_1 \cdot a_2 + r' g)z^{-(i+j)}]_q.$$

The GGH13 multilinear map: zero-test

- Sample h in R of the order of $q^{1/2}$.
- Define

$$p_{zt} = [z^{\kappa} h g^{-1}]_q.$$

The GGH13 multilinear map: zero-test

- Sample h in R of the order of $q^{1/2}$.
- Define
$$p_{zt} = [z^{\kappa} h g^{-1}]_q.$$

Zero-test

To test if $u = [cz^{-\kappa}]$ is an encoding of zero (i.e. $c = 0 \mod g$), compute
$$[u \cdot p_{zt}]_q = [chg^{-1}]_q.$$
This is small iff c is a small multiple of g.

The GGH13 multilinear map: other public parameters

Question

How to compute an encoding of a at level 1 when we only have the public parameters R, q and p_{zt}?

The GGH13 multilinear map: other public parameters

Question

How to compute an encoding of a at level 1 when we only have the public parameters R, q and p_{zt}?

Solution. We add to the public parameters

- y an encoding of 1 at level 1
- x an encoding of 0 at level 1.

To compute $C(a, 1)$:

Sample r in R and output $u = [ay + rx]_q$.

Conclusion on the GGH13 map

- We have a mathematical object, that satisfies some properties (addition, multiplication, zero-test).
- What about its security ?

Table of contents: 2 - Zeroizing attacks and consequences

Zeroizing attacks

Idea

When $u = [cz^{-\kappa}]_q$ with $c = bg$ a small multiple of g, we have

$$[u \cdot p_{zt}]_q = [chg^{-1}]_q = bh$$

because bh is smaller than q so $[bh]_q = bh \in R$.

Example of attack (from GGH13)

Compute

$$[x^2 y^{\kappa-2} p_{zt}]_q = [g^2 \cdot r \cdot g^{-1}]_q = g \cdot r$$

$\Rightarrow$ recover multiples of g, and then $\langle g \rangle$.

Hu and Jia's attack

Hu and Jia, 2015[1]

An attacker can recover the shared secret s in the multipartite key exchange protocol, when using the GGH13 multilinear map.

For this attack, we need x, the level 1 encoding of zero.

[1] Hu, Y., & Jia, H. (2016, May). "Cryptanalysis of GGH map".

Hu and Jia's attack

Hu and Jia, 2015[1]

An attacker can recover the shared secret s in the multipartite key exchange protocol, when using the GGH13 multilinear map.

For this attack, we need x, the level 1 encoding of zero.

Question

Maybe the GGH13 map is still safe if we do not have low level encodings of zero?

[1] Hu, Y., & Jia, H. (2016, May). "Cryptanalysis of GGH map".

Not all obfuscators are broken yet

Good news for obfuscators

We do not need the public parameters x and y in the GGH13 map when used for obfuscators.

$\Rightarrow$ the attack of Hu and Jia does not apply.

Not all obfuscators are broken yet

Good news for obfuscators

We do not need the public parameters x and y in the GGH13 map when used for obfuscators.

$\Rightarrow$ the attack of Hu and Jia does not apply.

Yes but...

Still, many obfuscators using the GGH13 map were proven insecure using zeroizing techniques.

Table of contents: 3 - Averaging attacks

Another approach: averaging

Idea

Instead of looking at the arithmetic properties of R, we use statistical properties.

This kind of attacks was already mentioned in the original article of GGH13.

Another approach: averaging

Idea

Instead of looking at the arithmetic properties of R, we use statistical properties.

This kind of attacks was already mentioned in the original article of GGH13.

Property: If D is a distribution over R and $x_1, \cdots, x_\ell$ are independent elements sampled from D, then

$$\frac{1}{\ell}\sum_{i=1}^{\ell} x_i \underset{\ell \to +\infty}{\longrightarrow} \mathbb{E}(x_1).$$

With ℓ samples, we expect to get $\log(\ell)$ bits of precision for $\mathbb{E}(x_1)$.

Notations and definitions (1)

Definitions

A distribution is said **centered** if its mean is zero.
A distribution is said **isotropic** if no direction is privileged.

Example

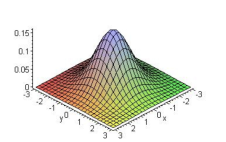

Notation: We write in red the centered isotropic variables.

Notations and definitions (1)

Definitions

A distribution is said **centered** if its mean is zero.
A distribution is said **isotropic** if no direction is privileged.

Example

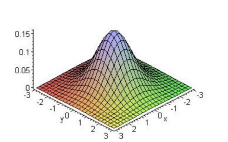

Notation: We write in red the centered isotropic variables.

Gaussian distribution

We denote by D_σ the (discrete) Gaussian distribution centered in 0 and of variance σ^2.

Remark. D_σ is a centered isotropic distribution (if σ is large enough).

Definitions and properties (2)

Definitions / Notation

- For $r \in R$, we denote $A(r) = r\bar{r}$ the **auto-correlation** of r, where $\bar{r}$ is the complex conjugate of r when seen in $\mathbb{C}$.
- The **variance** of a centered variable r is $\mathrm{Var}(r) := \mathbb{E}(r\bar{r})$.

Definitions and properties (2)

Definitions / Notation

- For $r \in R$, we denote $A(r) = r\bar{r}$ the **auto-correlation** of r, where $\bar{r}$ is the complex conjugate of r when seen in $\mathbb{C}$.
- The **variance** of a centered variable r is $\mathrm{Var}(r) := \mathbb{E}(r\bar{r})$.

Proposition: If r is sampled in R according to a centered isotropic distribution, then

$$\mathbb{E}(r) = 0$$
$$\mathrm{Var}(r) = \mu \in \mathbb{R}$$

Back to the attack: what do we know?

Reminder: We do not want to publicly give x and y anymore. So what is public?

Back to the attack: what do we know?

Reminder: We do not want to publicly give x and y anymore. So what is public?

Toy model inspired by obfuscators

- we are given R, q and p_{zt} as before.

Back to the attack: what do we know?

Reminder: We do not want to publicly give x and y anymore. So what is public?

Toy model inspired by obfuscators

- we are given R, q and p_{zt} as before.
- we are given $u_i = [c_i z^{-i}]$ for $1 \leq i < \kappa$ and $c_i \leftarrow D_\sigma$.
- such that $u_i u_{\kappa-i}$ is an encoding of 0 at level κ.

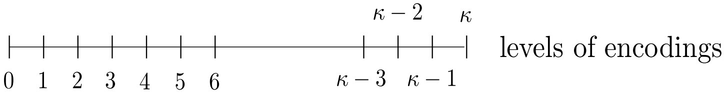

Back to the attack: what do we know?

Reminder: We do not want to publicly give x and y anymore.
So what is public?

Toy model inspired by obfuscators

- we are given R, q and p_{zt} as before.
- we are given $u_i = [c_i z^{-i}]$ for $1 \leq i < \kappa$ and $c_i \leftarrow D_\sigma$.
- such that $u_i u_{\kappa-i}$ is an encoding of 0 at level κ.

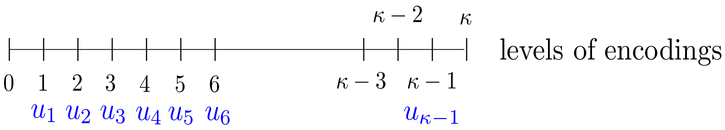

Back to the attack: what do we know?

Reminder: We do not want to publicly give x and y anymore. So what is public?

Toy model inspired by obfuscators

- we are given R, q and p_{zt} as before.
- we are given $u_i = [c_i z^{-i}]$ for $1 \leq i < \kappa$ and $c_i \leftarrow D_\sigma$.
- such that $u_i u_{\kappa-i}$ is an encoding of 0 at level κ.

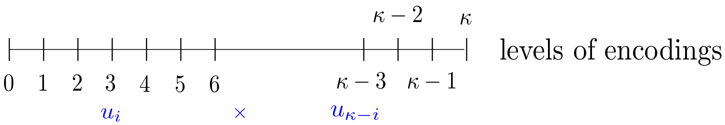

Idea of the attack

Recall our model

- we are given $u_i = [c_i z^{-i}]$ for $1 \leq i \leq \kappa - 1$ and $c_i \leftarrow D_\sigma$.
- such that $u_i u_{\kappa-i}$ is an encoding of 0 at level κ.

Observation:

$$\begin{aligned} [u_i u_{\kappa-i} \cdot p_{zt}] &= [c_i c_{\kappa-i} \cdot h/g] \\ &= c_i c_{\kappa-i} \cdot h/g \\ &= c_i^* \cdot h/g \end{aligned}$$

Idea of the attack (2)

Recall

We know

$$c_i^* \cdot h/g$$

for $1 \leq i \leq \kappa$, with c_i^* centered and isotropic.

Idea of the attack (2)

Recall

We know

$$c_i^* \cdot h/g$$

for $1 \leq i \leq \kappa$, with c_i^* centered and isotropic.

- $\mathbb{E}(c_i^*) = 0 \Rightarrow$ we do not learn anything with $\mathbb{E}(c_i^* \cdot h/g)$.

Idea of the attack (2)

Recall

We know

$$c_i^* \cdot h/g$$

for $1 \leq i \leq \kappa$, with c_i^* centered and isotropic.

- $\mathbb{E}(c_i^*) = 0 \Rightarrow$ we do not learn anything with $\mathbb{E}(c_i^* \cdot h/g)$.
- $\mathrm{Var}(c_i^*) = \mathbb{E}(A(c_i^*)) = \mu \in \mathbb{R}$ is some scalar $\Rightarrow$ we obtain

$$\frac{1}{\kappa}\sum_{i=1}^{\kappa} A(c_i^* \cdot h/g) \underset{\kappa \to +\infty}{\to} \mu A(h/g).$$

Idea of the attack (2)

Recall

We know

$$c_i^* \cdot h/g$$

for $1 \leq i \leq \kappa$, with c_i^* centered and isotropic.

- $\mathbb{E}(c_i^*) = 0 \Rightarrow$ we do not learn anything with $\mathbb{E}(c_i^* \cdot h/g)$.
- $\mathrm{Var}(c_i^*) = \mathbb{E}(A(c_i^*)) = \mu \in \mathbb{R}$ is some scalar $\Rightarrow$ we obtain

$$\frac{1}{\kappa}\sum_{i=1}^{\kappa} A(c_i^* \cdot h/g) \underset{\kappa \to +\infty}{\to} \mu A(h/g).$$

We get an approximation of $A(h/g)$ with $\log(\kappa)$ bits of precision.

GGH13 counter-measure

GGH13's authors noticed that their scheme was subject to averaging attacks $\Rightarrow$ they proposed a countermeasure.

Definition

Let z_i be the representative of $[z^i]$ in R with coefficients in $[-q/2, q/2]$.

Idea: choose c_i such that c_i/z_i is isotropic.

Counter-measure

- Sample $\widetilde{c_i} \leftarrow D_\sigma$.
- Define $c_i = \widetilde{c_i} \cdot z_i$.
- And $u_i = [c_i z^{-i}]$ as before.

Adapting the attack to the counter-measure

Recall

- $c_i = \widetilde{c_i} \cdot z_i$.
- $u_i = [c_i z^{-i}]$.
- $u_i u_{\kappa-i}$ is an encoding of 0 at level κ.

Observation:

$$\begin{aligned}[u_i u_{\kappa-i} \cdot p_{zt}] &= \widetilde{c_i}\widetilde{c_{\kappa-i}} \cdot z_i z_{\kappa-i} \cdot h/g \\ &= c_i^* \cdot z_i z_{\kappa-i} \cdot h/g\end{aligned}$$

Adapting the attack to the counter-measure

Recall

- $c_i = \widetilde{c_i} \cdot z_i$.
- $u_i = [c_i z^{-i}]$.
- $u_i u_{\kappa-i}$ is an encoding of 0 at level κ.

Observation:

$$[u_i u_{\kappa-i} \cdot p_{zt}] = \widetilde{c_i}\widetilde{c_{\kappa-i}} \cdot z_i z_{\kappa-i} \cdot h/g$$
$$= c_i^* \cdot z_i z_{\kappa-i} \cdot h/g$$

But: the z_i are isotropic and independent.

Adapting the attack to the counter-measure

Recall

- $c_i = \widetilde{c_i} \cdot z_i$.
- $u_i = [c_i z^{-i}]$.
- $u_i u_{\kappa-i}$ is an encoding of 0 at level κ.

Observation:

$$[u_i u_{\kappa-i} \cdot p_{zt}] = \widetilde{c_i}\,\widetilde{c_{\kappa-i}} \cdot z_i z_{\kappa-i} \cdot h/g$$
$$= c_i^* \cdot z_i z_{\kappa-i} \cdot h/g$$

But: the z_i are isotropic and independent.

Adapting the attack to the counter-measure

Recall

- $c_i = \widetilde{c_i} \cdot z_i$.
- $u_i = [c_i z^{-i}]$.
- $u_i u_{\kappa-i}$ is an encoding of 0 at level κ.

Observation:

$$[u_i u_{\kappa-i} \cdot p_{zt}] = \widetilde{c_i}\widetilde{c_{\kappa-i}} \cdot z_i z_{\kappa-i} \cdot h/g$$
$$= c_i^* \cdot z_i z_{\kappa-i} \cdot h/g$$

But: the z_i are isotropic and independent.

Averaging: we get an approx of $\mu A(h/g)$, for some constant μ.

Conclude the attack

Lemma

If we have

- an approximation of $A(h/g)$ with $\log \ell$ bits of precision,
- a guarantee that for any encoding $[cz^{-i}]$, the coefficients of c are less than $\ell/2$.

Then, we can recover $A(h/g)$ exactly and attack the GGH13 map.

Conclude the attack

Lemma

If we have

- an approximation of $A(h/g)$ with $\log \ell$ bits of precision,
- a guarantee that for any encoding $[cz^{-i}]$, the coefficients of c are less than $\ell/2$.

Then, we can recover $A(h/g)$ exactly and attack the GGH13 map.

Do we get enough samples for recovering $A(h/g)$ exactly?

- Without the counter-measure $\Rightarrow$ yes.
- With the counter-measure $\Rightarrow$ no, but this is because of constraints in the sampling procedure.

Conclusion

In the case where q is polynomial:

- complete attack without the counter-measure (if κ is large enough).
- leaked information with the counter-measure.
- other variants (adapted from [DGG+16][2]): leaked information but no complete attack.

[2] Döttling, N. et al. "Obfuscation from Low Noise Multilinear Maps".

Conclusion

In the case where q is polynomial:

- complete attack without the counter-measure (if κ is large enough).
- leaked information with the counter-measure.
- other variants (adapted from [DGG+16][2]): leaked information but no complete attack.

$\Rightarrow$ Not clear what could be a hard problem for the GGH map.

[2] Döttling, N. et al. "Obfuscation from Low Noise Multilinear Maps".

Conclusion

In the case where q is polynomial:

- complete attack without the counter-measure (if κ is large enough).
- leaked information with the counter-measure.
- other variants (adapted from [DGG+16][2]): leaked information but no complete attack.

$\Rightarrow$ Not clear what could be a hard problem for the GGH map.

Thank you for your attention.

[2] Döttling, N. et al. "Obfuscation from Low Noise Multilinear Maps".