Olivier Bernard Marc Joye Nigel Smart Michael Walter

06.05.2025

Drifting Towards Better Error Probabilities in Fully Homomorphic Encryption Schemes

ZAMA

Motivation

Challenges in FHE Deployments

- Performance overhead due to **noise accumulation**
- **Decryption failure probability** must be minimized
- Trade-offs between **efficiency, security, and correctness**

Motivation

Challenges in FHE Deployments

- Performance overhead due to **noise accumulation**
- **Decryption failure probability** must be minimized
- Trade-offs between **efficiency, security, and correctness**

Questions

- How do **decryption failures** impact practical security?
- Can we develop **low-overhead solutions** to reduce failure probability?
- How do our techniques improve **theoretical and applied security**?

Ciphertext Drift

Definition and Impact

What is Ciphertext Drift?

- Ciphertext drift refers to the **accumulation of small errors** introduced during modulus switching in FHE
- It occurs due to **rounding effects** when converting ciphertexts between different moduli

Ciphertext Drift

Definition and Impact

What is Ciphertext Drift?

- Ciphertext drift refers to the **accumulation of small errors** introduced during modulus switching in FHE
- It occurs due to **rounding effects** when converting ciphertexts between different moduli

Why Does Ciphertext Drift Matter?

- Drift increases **failure probability**, which can be exploited in certain adversarial scenarios
- Reducing drift often requires **larger cryptographic parameters**, increasing computational cost

Security Models in FHE

IND-CPA Standard security against chosen-plaintext attacks

IND-CPA^D Attacker has [very restricted] access to a decryption oracle

- only decryption queries for which corresponding plaintext is known to the attacker are allowed

sIND-CPA^D Strengthened IND-CPA^D model

Security Models in FHE

IND-CPA Standard security against chosen-plaintext attacks

IND-CPA^D Attacker has [very restricted] access to a decryption oracle

- only decryption queries for which corresponding plaintext is known to the attacker are allowed

sIND-CPA^D Strengthened IND-CPA^D model

Why Standard IND-CPA May Be Insufficient

- **FHE decryption failures** may reveal information about secret keys
- Attacks leveraging failure probabilities can break **security assumptions**

Noise Growth and Decryption Failures

- **Noise accumulation in homomorphic operations**
  - Each operation increases ciphertext noise
  - Bootstrapping required for noise reduction but introduces additional drift

Noise Growth and Decryption Failures

- **Noise accumulation in homomorphic operations**
  - Each operation increases ciphertext noise
  - Bootstrapping required for noise reduction but introduces additional drift
- **Failure probability in FHE**
  - When noise exceeds a threshold, decryption fails
  - Even negligible failure probabilities impact security in adversarial settings

Noise Growth and Decryption Failures

- **Noise accumulation in homomorphic operations**
  - Each operation increases ciphertext noise
  - Bootstrapping required for noise reduction but introduces additional drift
- **Failure probability in FHE**
  - When noise exceeds a threshold, decryption fails
  - Even negligible failure probabilities impact security in adversarial settings
- **Practical impact on IND-CPA^{D} security**
  - Attacker can craft (honestly generated or evaluated) ciphertexts to probe decryption failures
  - Failure probabilities must be controlled at all computation steps

Ciphertext Drift & Modulus Switching

Ciphertext drift

- Accumulated error from rounding during modulus switching
- Leads to decryption failures and security vulnerabilities

Modulus switching

- Converts ciphertext modulus from q to q' through rescaling and rounding
- Essential step for bootstrapped FHE [FHEW, TFHE, FINAL, ...]
- Introduces drift, which must be controlled to ensure correctness

Ciphertext Drift & Modulus Switching

Ciphertext drift

- Accumulated error from rounding during modulus switching
- Leads to decryption failures and security vulnerabilities

Modulus switching

- Converts ciphertext modulus from q to q' through rescaling and rounding
- Essential step for bootstrapped FHE [FHEW, TFHE, FINAL, …]
- Introduces drift, which must be controlled to ensure correctness

- Increased drift raises failure probability
- Failure-aware adversaries can extract secret key information

Conventional IND-CPA parameters may not ensure IND-CPA^{D} security!

Our Techniques and Results

New Theoretical Insights

- Separation between **IND-CPA^D and sIND-CPA^D security**
- **Characterization** of failure probability in practical FHE schemes

Our Techniques and Results

New Theoretical Insights

- Separation between **IND-CPA^D and sIND-CPA^D security**
- **Characterization** of failure probability in practical FHE schemes

New Modulus Switching Methods

- **Controlled noise management** reducing failure probability
- Requires **no significant parameter inflation**

Illustration: The Case of LWE

Modulus Switching

Input: LWE-type ciphertext modulo q

$$\boldsymbol{c} \leftarrow (a_1, \ldots, a_n, b) \in (\mathbb{Z}/q\mathbb{Z})^{n+1}$$

with $b = \sum_{i=1}^{n} a_i s_i + \Delta m + e_{in}$ and where e_{in} is some input noise error

Output: LWE-type ciphertext modulo $2N$

$$\tilde{\boldsymbol{c}} \leftarrow \left(\tilde{a}_1, \ldots, \tilde{a}_n, \tilde{b}\right) \in (\mathbb{Z}/2N\mathbb{Z})^{n+1}$$

with

$$\begin{cases} \tilde{a}_i = \left\lfloor \frac{a_i}{q} 2N \right\rceil & \text{for } i \in \{1, \ldots, n\} \\ \tilde{b} = \left\lfloor \frac{b}{q} 2N \right\rceil \end{cases}$$

Illustration: The Case of LWE

Modulus Switching

Input: LWE-type ciphertext modulo q

$$\boldsymbol{c} \leftarrow (a_1, \ldots, a_n, b) \in (\mathbb{Z}/q\mathbb{Z})^{n+1}$$

with $b = \sum_{i=1}^{n} a_i s_i + \Delta m + e_{in}$ and where e_{in} is some input noise error

Output: LWE-type ciphertext modulo $2N$

$$\tilde{\boldsymbol{c}} \leftarrow (\tilde{a}_1, \ldots, \tilde{a}_n, \tilde{b}) \in (\mathbb{Z}/2N\mathbb{Z})^{n+1}$$

with

$$\begin{cases} \tilde{a}_i = \left\lfloor \frac{a_i}{q} 2N \right\rceil & \text{for } i \in \{1, \ldots, n\} \\ \tilde{b} = \left\lfloor \frac{b}{q} 2N \right\rceil \end{cases}$$

Ciphertext Drift

Write

$$\begin{cases} a_i = \tilde{a}_i \frac{q}{2N} - \alpha_i \\ b = \tilde{b} \frac{q}{2N} - \beta \end{cases}$$

for some $\alpha_i, \beta \in \left[\!\left[-\frac{q}{4N}, \frac{q}{4N} \right]\!\right]$

Then

$$\left(\tilde{b} - \sum_{i=1}^{n} \tilde{a}_i s_i - \frac{2N}{p} m\right) \bmod 2N = \frac{\overbrace{\left(e_{in} + \left(\beta - \sum_{i=1}^{n} \alpha_i s_i\right)\right)}^{e_{drift}} \bmod q}{q/2N}$$

Two Important Observations

1. An LWE ciphertext $(a_1, ..., a_n, b)$ can be **publicly** re-randomized
   - simply add an encryption 0
2. The drift vector $(\alpha_1, ..., \alpha_n, \beta)$ can be **publicly** computed
   [Drift error cannot]

Two Important Observations

1. An LWE ciphertext $(a_1, ..., a_n, b)$ can be **publicly** re-randomized
   - simply add an encryption 0
2. The drift vector $(\alpha_1, ..., \alpha_n, \beta)$ can be **publicly** computed
   [Drift error cannot]

Introducing Drift-Aware Modulus Switching (Public by Design)

- **Select a ciphertext representative**
  - by adding an encryption of 0 to the input ciphertext
- Use the **drift vector** to test the **'quality'** of the representative

(Repeat until a 'good' ciphertext is found)

Controlling the Drift

Probabilistic Approach

- For a fixed ciphertext $\boldsymbol{C} = (a_1, \ldots, a_n, b)$ with drift vector $(\alpha_1, \ldots, \alpha_n, \beta)$, corresponding drift error is $e_{drift} = \beta - \sum_{i=1}^{n} \alpha_i s_i$

Controlling the Drift

Probabilistic Approach

- For a fixed ciphertext $\boldsymbol{C} = (a_1, \ldots, a_n, b)$ with drift vector $(\alpha_1, \ldots, \alpha_n, \beta)$, corresponding drift error is $e_{drift} = \beta - \sum_{i=1}^{n} \alpha_i s_i$

- We have

$$\begin{aligned} \mu := \mathbb{E}[e_{drift}] &= \beta - \sum_{i=1}^{n} \alpha_i \mathbb{E}[s_i] \\ &= \beta - \frac{1}{2} \sum_{i=1}^{n} \alpha_i \\ \sigma^2 := \mathrm{Var}(e_{drift}) &= \sum_{i=1}^{n} \alpha_i^2 \, \mathrm{Var}(s_i) \\ &= \frac{1}{4} \sum_{i=1}^{n} \alpha_i^2 \end{aligned}$$

Controlling the Drift

Probabilistic Approach

- For a fixed ciphertext $\boldsymbol{C} = (a_1, \ldots, a_n, b)$ with drift vector $(\alpha_1, \ldots, \alpha_n, \beta)$, corresponding drift error is $e_{drift} = \beta - \sum_{i=1}^{n} \alpha_i s_i$

- We have

$$\begin{aligned}\mu := \mathbb{E}[e_{drift}] &= \beta - \textstyle\sum_{i=1}^{n} \alpha_i \mathbb{E}[s_i] \\ &= \beta - \tfrac{1}{2} \textstyle\sum_{i=1}^{n} \alpha_i \\ \sigma^2 := \mathrm{Var}(e_{drift}) &= \textstyle\sum_{i=1}^{n} \alpha_i^2 \, \mathrm{Var}(s_i) \\ &= \tfrac{1}{4} \textstyle\sum_{i=1}^{n} \alpha_i^2\end{aligned}$$

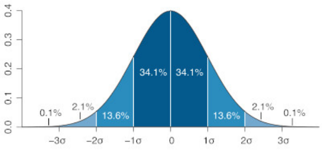

$$p = \Pr\Big[X \in [\![\mu - r\sigma, \mu + r\sigma]\!]\Big]$$

r	$1-p$
7.15	2^{-40}
9.16	2^{-64}
10.29	2^{-80}
13.11	2^{-128}

Controlling the Drift

Probabilistic Approach

- For a fixed ciphertext $\boldsymbol{C} = (a_1, \ldots, a_n, b)$ with drift vector $(\alpha_1, \ldots, \alpha_n, \beta)$, corresponding drift error is $e_{drift} = \beta - \sum_{i=1}^{n} \alpha_i s_i$
- We have

$$\begin{aligned} \mu &:= \mathbb{E}[e_{drift}] = \beta - \textstyle\sum_{i=1}^{n} \alpha_i \mathbb{E}[s_i] \\ &= \beta - \tfrac{1}{2} \textstyle\sum_{i=1}^{n} \alpha_i \\ \sigma^2 &:= \mathrm{Var}(e_{drift}) = \textstyle\sum_{i=1}^{n} \alpha_i^2 \, \mathrm{Var}(s_i) \\ &= \tfrac{1}{4} \textstyle\sum_{i=1}^{n} \alpha_i^2 \end{aligned}$$

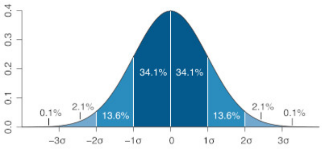

$$p = \Pr\left[X \in [\![\mu - r\sigma, \mu + r\sigma]\!]\right]$$

r	$1-p$
7.15	2^{-40}
9.16	2^{-64}
10.29	2^{-80}
13.11	2^{-128}

Proposed 'Quality' Test

Check that

$$|\mu| + r\sigma \overset{?}{\leq} T$$

where T is a bound to the maximum allowed drift error

Application

Parameter set	q	n	N	k	p_{err}	T	# trials	$p_{\mathrm{err}}^{(\mathrm{new})}$
$\mathfrak{Z}^{64}_{\mathrm{b}}$	2^{64}	739	512	3	2^{-64}	$2^{59.67}$	50	$\mathbf{2^{-128.83}}$
							100	$2^{-130.41}$
							1000	$2^{-134.75}$
$\mathfrak{Z}^{64}_{\mathrm{4b}}$	2^{64}	834	2048	1	2^{-64}	$2^{57.76}$	50	$\mathbf{2^{-128.44}}$
							100	$2^{-129.94}$
							1000	$2^{-134.02}$

Validation & Results

- Without our techniques
  - Failure probability is of $2^{-\kappa}$
- **With our techniques**
  - Failure probability is reduced to roughly $2^{-2\kappa}$

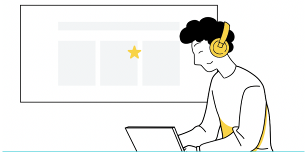

Validation & Results

- Without our techniques
  - Failure probability is of $2^{-\kappa}$
- **With our techniques**
  - Failure probability is reduced to roughly $2^{-2\kappa}$

Implications

- Exponential improvement in failure rate
- Unnoticeable performance penalty
- Enables stronger correctness guarantees for existing FHE schemes

Conclusion

Summary of Key Contributions

- Studied **ciphertext drift in FHE**
- Developed **novel modulus switching methods** reducing **failure probability**
- Strengthened FHE security models with **sIND-CPA^{D} refinements**

🕶 Read the full paper at ePrint 2024/1718

Contact and Links

michael.walter@zama.ai

zama.ai

GitHub

Community

ZAMA