Failure Probability Measured with Rare Event Simulation

Mathieu Ballandras, Jean-Baptiste Orfila, Samuel Tap

Irwin–Hall distribution

Let $(U_i)_{i=1,\dots,m}$ bet independent random variables uniformly distributed over $[-\frac{w}{2}, \frac{w}{2}]$. The distribution **Irwin-Hall**(m, w) is defined as the distribution of:

$$\sum_{i=1}^{m} U_i.$$

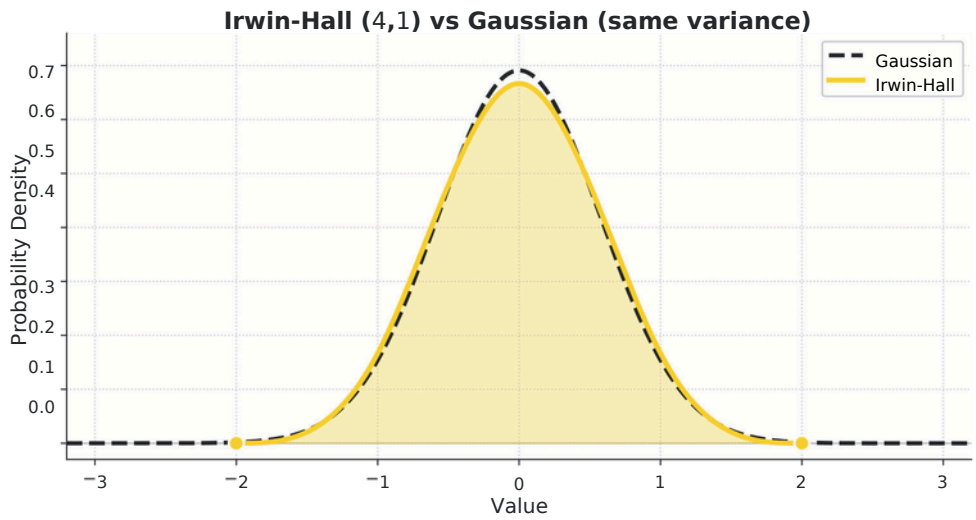

Modulus Switch noise: Irwin–Hall vs. Gaussian

Consider several variants of Modulus Switch (switching from modulus q' to q) for **LWE** ciphertexts with dimension n and binary secret key. The noise added by these operations is better modeled by an Irwin–Hall distribution than by a Gaussian distribution.

- Modulus Switch: **Irwin-Hall** $\left(\frac{n}{2}, \frac{1}{2q}\right)$
- Mean Compensated [RDV25] **MS**: **Irwin-Hall** $\left(n, \frac{1}{4q}\right)$
- Multi Bit [Zho+18] **MS**, grouping factor g: **Irwin-Hall** $\left(\left(1 - \frac{1}{2^g}\right)\frac{n}{g}, \frac{1}{2q}\right)$

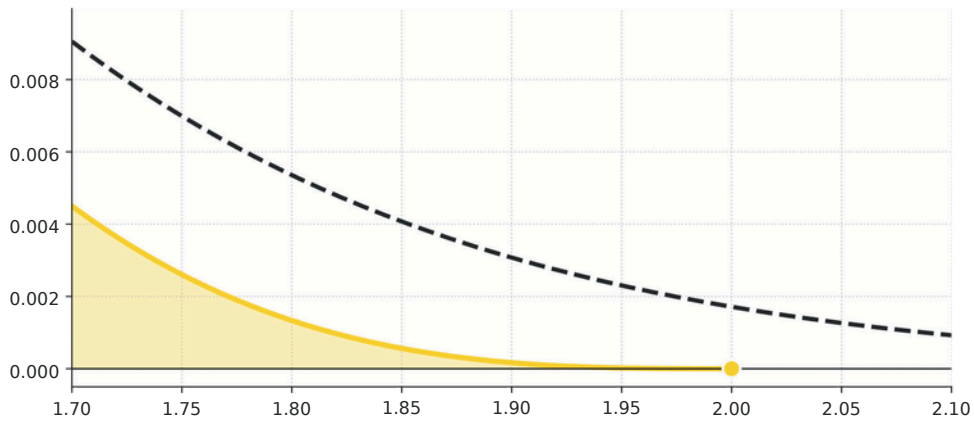

Zoom on the tail

	Modulus Switch	Mean Comp. MS	Multi Bit MS
Importance Splitting	9.33e−23	3.28e−40	2.81e−43
(Upper Bound 3σ)	1.20e−22	4.34e−40	5.18e−43
Gaussian Prediction	2.20e−22	9.82e−40	3.70e−41
Irwin–Hall Prediction	1.05e−22	3.27e−40	2.20e−43

Table: Failure Probability: Experimental measures (Importance Splitting) vs. Theoretical Predictions.

Importance Splitting: the algorithm

1.
   - Fix a sequence of thresholds $T_1 < T_2 < \cdots < T_l = T$
   - Initialize the population $X_1^{(0)}, \dots, X_n^{(0)}$
2. For $i = 1, \dots, l$ do
   The population is $X_1^{(i-1)}, \dots, X_n^{(i-1)}$ and $S(X_j^{(i-1)}) > T_{i-1}$ for all j
   - Keep only the samples with score larger than T_i and count them

$$\left\{Y_1^{(i)}, \dots, Y_{n_i}^{(i)}\right\} := \left\{X_j^{(i-1)} \text{ such that } S(X_j^{(i-1)}) > T_i\right\}$$

   - Compute $p_i = \frac{n_i}{n}$
   - Resample a population of size n from the population $Y_1^{(i)}, \dots, Y_{n_i}^{(i)}$

$$X_1^{(i)}, \dots, X_n^{(i)}$$

   such that $S(X_j^{(i)}) > T_i$ for all j
3. Return $p = p_1 p_2 \dots p_l$

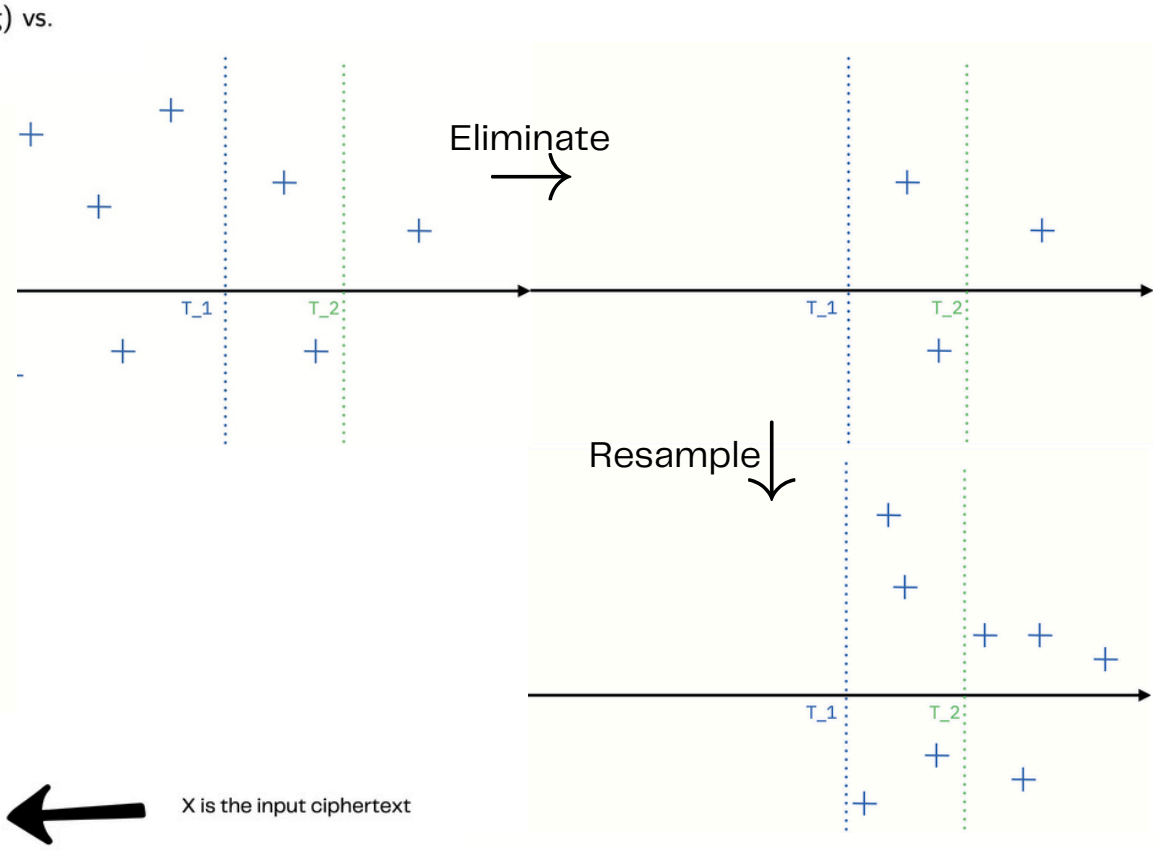

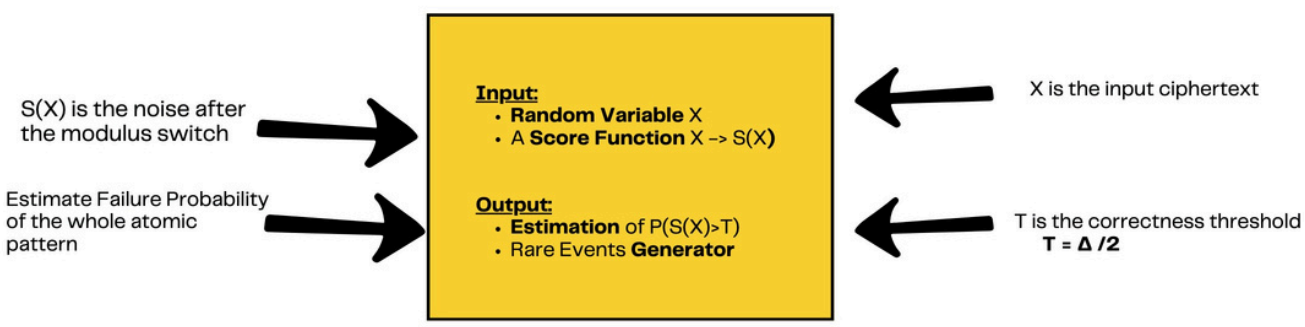

Bayes formula makes it work:

$$P(S > T_l) = P(S > T_1) P(S > T_2 \,|\, S > T_1) \dots P(S > T_l \,|\, S > T_{l-1})$$

[RDV25] Thomas de Ruijter, Jan-Pieter D'Anvers, and Ingrid Verbauwhede. *Don't be mean: Reducing Approximation Noise in TFHE through Mean Compensation*. Cryptology ePrint Archive, Paper 2025/809. 2025. URL: https://eprint.iacr.org/2025/809.

[Zho+18] Tanping Zhou, Xiaoyuan Yang, Longfei Liu, Wei Zhang, and Ningbo Li. "Faster Bootstrapping With Multiple Addends". In: *IEEE Access* 6 (2018), pp. 49868–49876. DOI: 10.1109/ACCESS.2018.2867655.

Resources
zama.org/blog
github.com/zama-ai