Applications of Smoothing

Statistical Security in Lattice Schemes and Entropic Security

Nico Döttling (CISPA Helmholtz Center), July 16th, 2024

Learning with Errors and Gaussians

- Daniele's talk: LWE is *Interface* to Lattice Cryptography
- Do lattice crypto without fully understanding lattices!
- Gaussians are somewhat error distribution and easy to analyse
- **This talk:** Gaussians provide security features we don't know how to achieve with other error distributions

Overview

- **Basic Tools**
  - Drowning
  - Leftover Hashing
- **Refined Tools**
  - Singular Value Analysis
  - Smoothing

Applications

- FHE Circuit Privacy
- Entropic LWE

Learning with Errors [Reg05]

Decisional Version

$$\mathbb{Z}_q$$

$$\approx_c$$

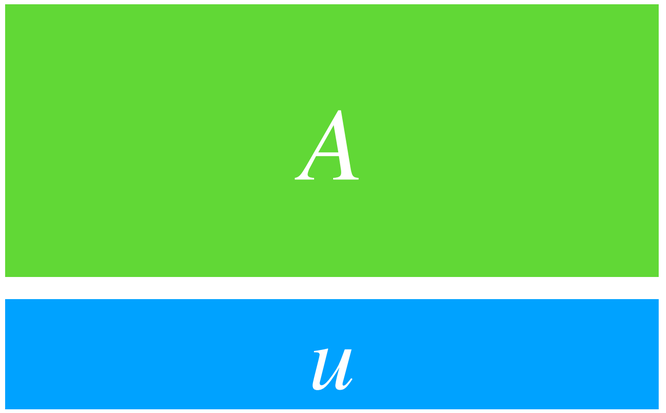

$$=$$

$$e \sim \chi$$

Worst-Case Hardness of LWE/Modulo-to-Noise Ratio

- For **gaussian** error distributions $\chi = D_{\sigma'}$ LWE enjoys worst-case hardness
- Quantum Reduction from (wc) SIVP to LWE [Reg05], classical reduction from (wc) GapSVP to LWE [Pei09,BLPRS13]
- Approxiation factor of worst-case problem relates to the modulus-to-noise ratio $\alpha = q/\sigma$

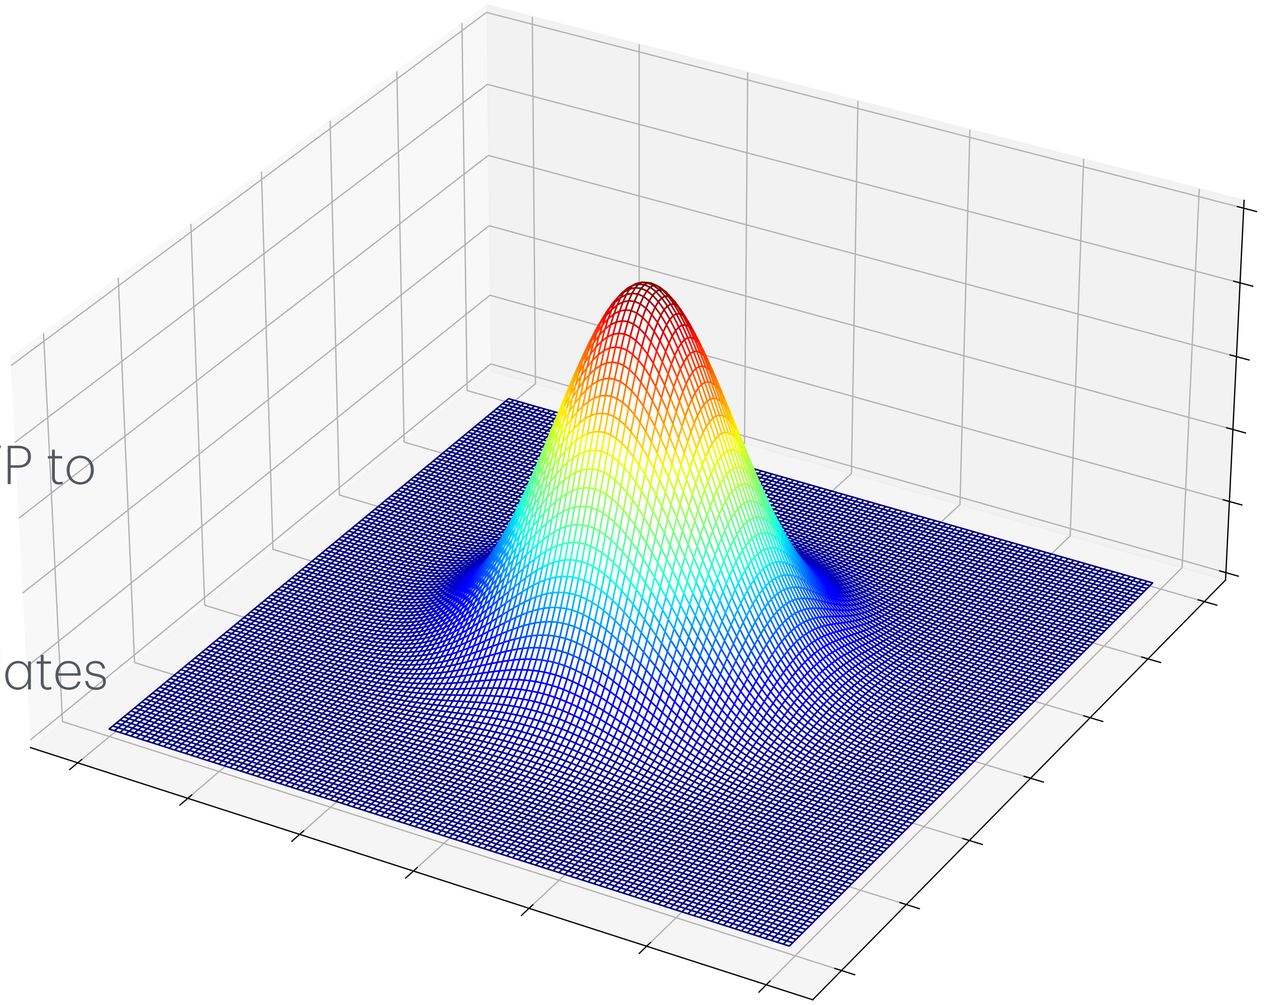

Basic Tools

Drowning/Flooding/Smudging

Lemma:

- χ **symmetric** and **monotonously decreasing** distribution
- $e \sim \chi, t \in \mathbb{R}$
- Then $\Delta(e + t, e) = \Pr[e \in [-t/2, t/2]]$
- **Bottom Line:** Good anti-concentration bound for χ $\Rightarrow e$ hides t

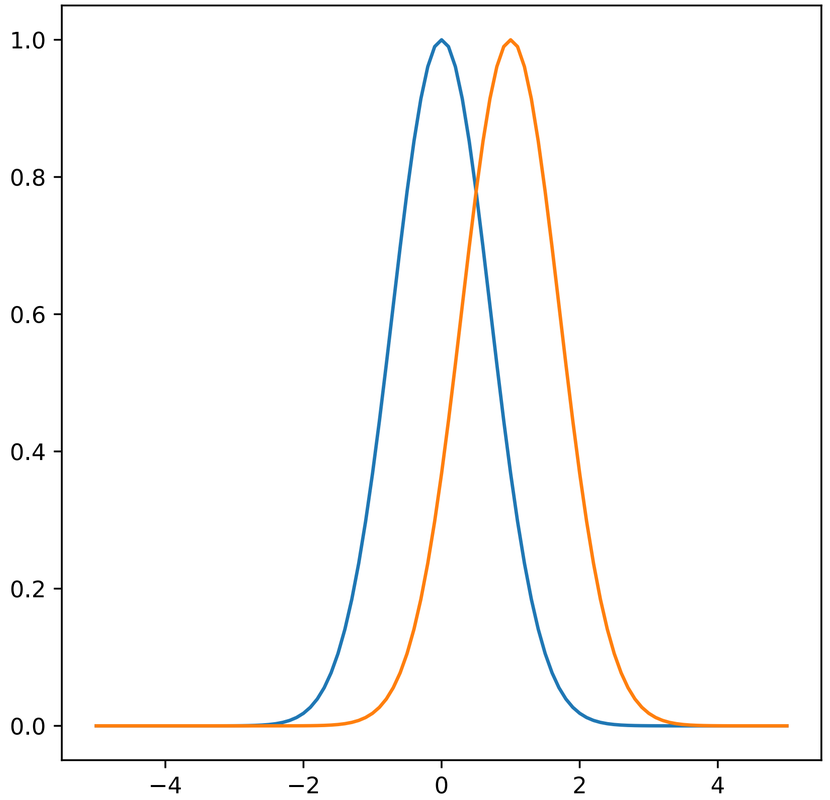

Proof

$$\Delta(X, X') = \sum_x |\Pr[X = x] - \Pr[X' = x]|$$

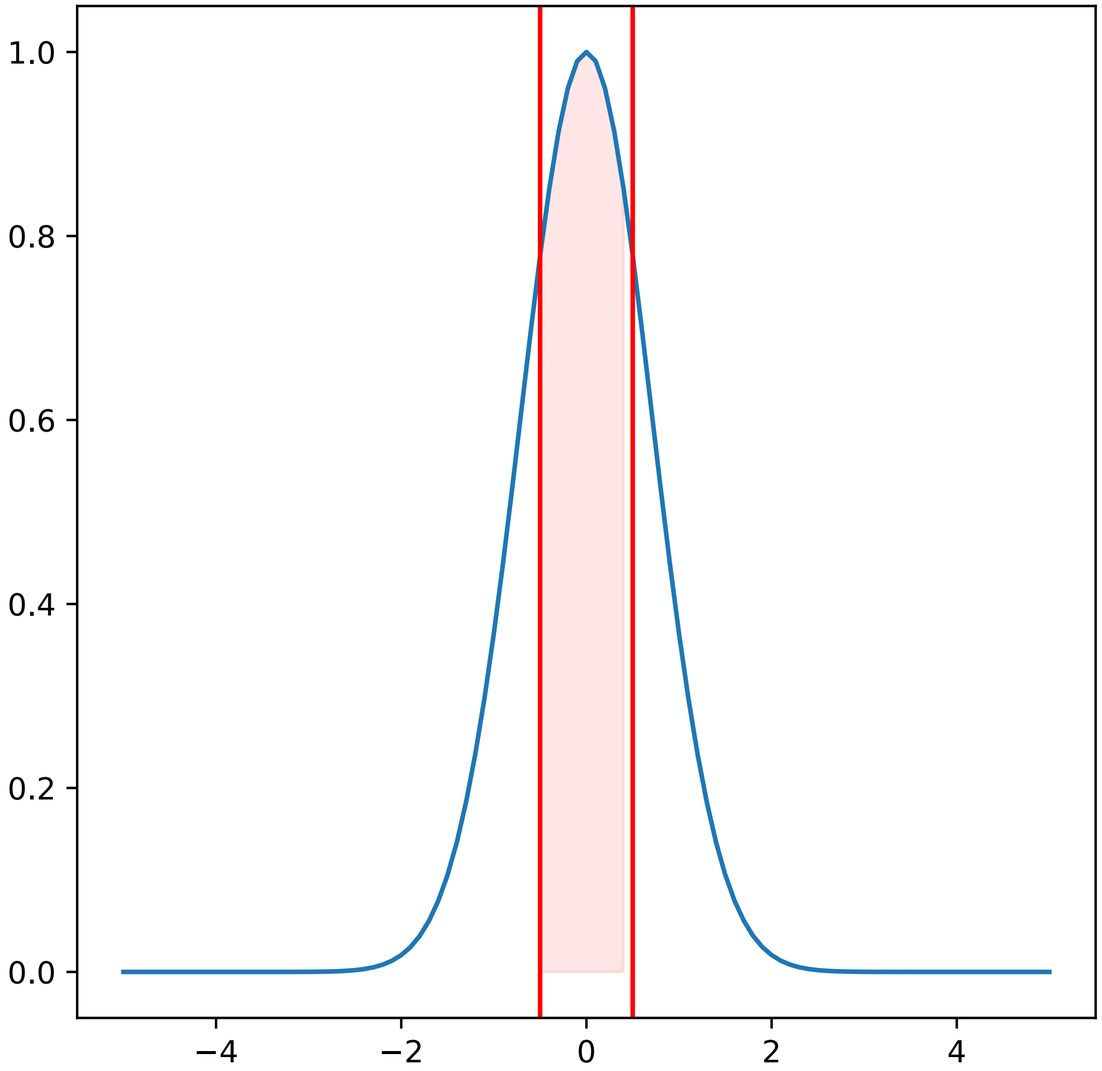

Drowning/Flooding/Smudging

- χ rectangular: $\Delta(e+t,e) = \Pr[e \in [-t/2,t/2]] = \frac{t}{r}$
- χ Gaussian: $\Delta(e+t,e) = \Pr[e \in [-t/2,t/2]] \leq \frac{t}{\sigma}$
- r , s must be superpoly larger than t for this expression to be negligible
- **Drawback:** Superpoly modulus-to-noise regime unavoidable

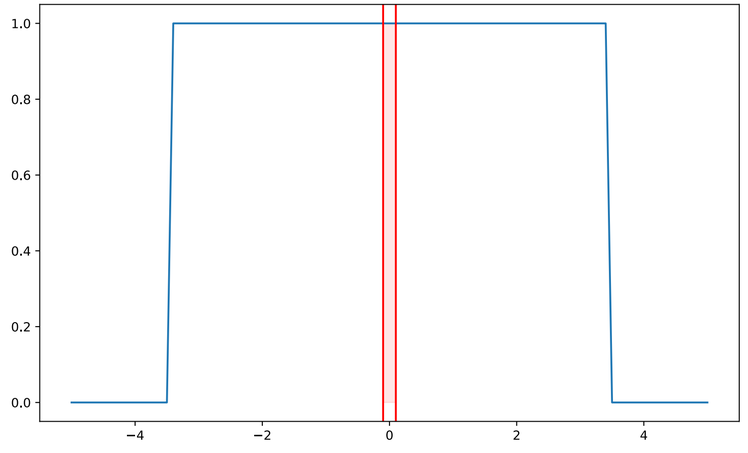

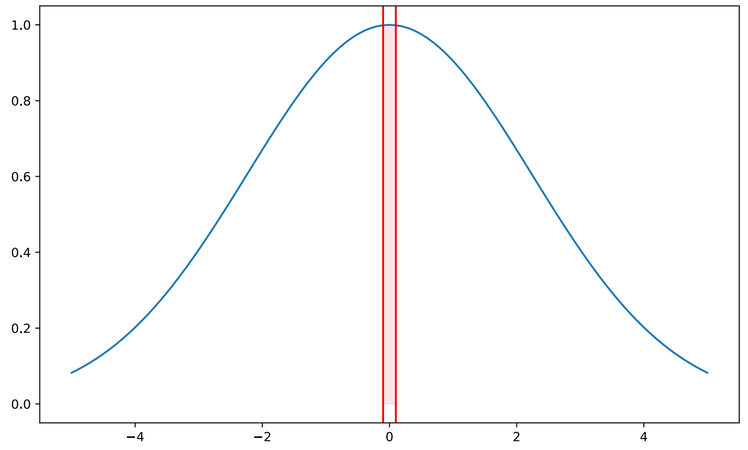

Leftover Hashing

- Min Entropy $\sim$ Log-Guessability: $H_\infty(x) = -\log(\max_\xi \Pr[x = \xi])$
- For simplicity q prime
- x supported on $\mathbb{Z}_q^m$ with $H_\infty(x) \geq n\log(q) + 2\log(1/\epsilon)$
- $A \in \mathbb{Z}_q^{n \times m}$ chosen uniformly random
- Then $(A, Ax) \approx_\epsilon (A, u)$ for uniformly random $u \in \mathbb{Z}_q^n$

GSW Encryption

- **Public Key:** Matrix A
- **Secret Key:** Vector s
- **Ciphertext:** Matrix $C = AR + mG$
  R random short, G gadget matrix

GSW Encryption

GSW Homomorphic Operations

- **Decryption:** $(-s^\top, 1)C = e^\top R + mG$
- $C_1 = AR_1 + m_1 G,\ C_2 = AR_2 + m_2 G$
- **Homomorphic Addition:** $C_1 + C_2 = A(R_1 + R_2) + (m_1 + m_2)G$
- **Homomorphic Multiplication:** $C_1 \cdot G^{-1}(C_2) = A(R_1 G^{-1}(C_2) + m_1 R_2)) + m_1 m_2 G$
- Ciphertext $C^* = AR^* + m^* G$ after homomorphic Operation: Norm of R grows moderately, but R' is far from random (comes from homomorphic evaluation)

Circuit Privacy

Ciphertext Sanitization

- Circuit Privacy: Homomorphically computed ciphertexts statistically look like fresh encryptions
- Choose fresh R and set $C' = C^* + AR$
- Recall that $A = \begin{pmatrix} A' \\ s^\top A' + e^\top \end{pmatrix}$
- First component of C': $A'R$ is statistically close to uniform and hence is $A'R^* + A'R$
- Second component of C': $s^\top(A'R^* + A'R) + e^\top(R^* + R)$
- $e^\top(R^* + R)$ leaks information about R^*

Circuit Privacy

Ciphertext Sanitization

- $e^{\mathsf{T}}(R^* + R)$ leaks information about R^*
- Drown this term out with a drowning term d, i.e. set $C' = C^* + AR + \begin{pmatrix} 0 \\ \tilde{e} \end{pmatrix}$
- Second component of C' now:
  $s^{\mathsf{T}}(A'R^* + A'R) + e^{\mathsf{T}}(R^* + R) + \tilde{e} \approx s^{\mathsf{T}}(A'R^* + A'R) + \tilde{e}$
- Statistically independent of R^*

Refined Tools

Lattices, Dual Lattices and Fourier Transforms

- Lattice $\Lambda = \{B \cdot x \mid x \in \mathbb{Z}^n\}$ for some full rank $B \in \mathbb{R}^{n \times n}$
- Dual Lattice $\Lambda^* = \{y \in \mathbb{R}^n \mid \forall z \in \Lambda : \langle y, z \rangle \in \mathbb{Z}\}$
- It holds $\Lambda^* = \{(B^{-1})^\top \cdot x \mid x \in \mathbb{Z}^n\}$

Singular Value Analysis

- $A \in \mathbb{R}^{n \times m}$ real-valued matrix with (say) $m \geq n$
- Singular values $\sigma_1(A) \geq \ldots \geq \sigma_n(A) \geq 0$ are square-roots of eigenvalues of $A \cdot A^\top$
- A can be written as $A = UDV^\top$ with $U \in O(\mathbb{R}^n)$, $V \in O(\mathbb{R}^m)$ and $D \in \mathbb{R}^{n \times m}$ a matrix with the $\sigma_i(A)$ on diagonal and 0 everywhere else

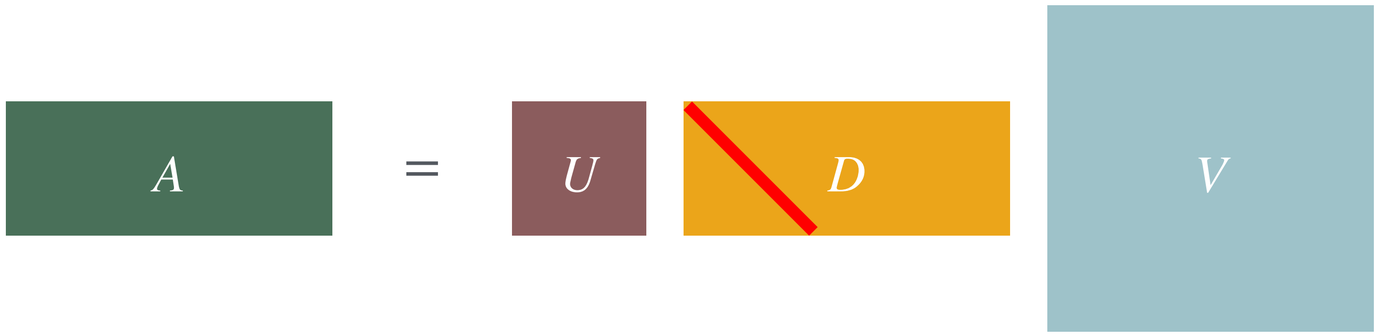

Singular Value Analysis

- Matrix $\boldsymbol{M} \in \mathbb{R}^{n \times n}$ positive definite $\Leftrightarrow \boldsymbol{M}$ symmetric and all singular values of $\boldsymbol{M}$ are positive
- Write $\boldsymbol{M} > 0$
- $\boldsymbol{M}$ is of the form $A^\top \cdot A$
- "Inner product matrices": For all $x \in \mathbb{R}^n \backslash \{0\}$ it holds $x^\top M x > 0$
- Additional Notation: Write $\boldsymbol{M} > \boldsymbol{M}'$ for $\boldsymbol{M} - \boldsymbol{M}' > 0$

Gaussians

$$D_{\mathbb{R}^n,\sqrt{\Sigma}}$$

$$\rho_{\sqrt{\Sigma}}(x) \propto e^{-\pi \cdot x^\top \Sigma^{-1} x}$$

$$x \sim D^{\mathbb{R}^n,\sqrt{\Sigma}} \Rightarrow \mathbb{E}[xx^\top] = \Sigma$$

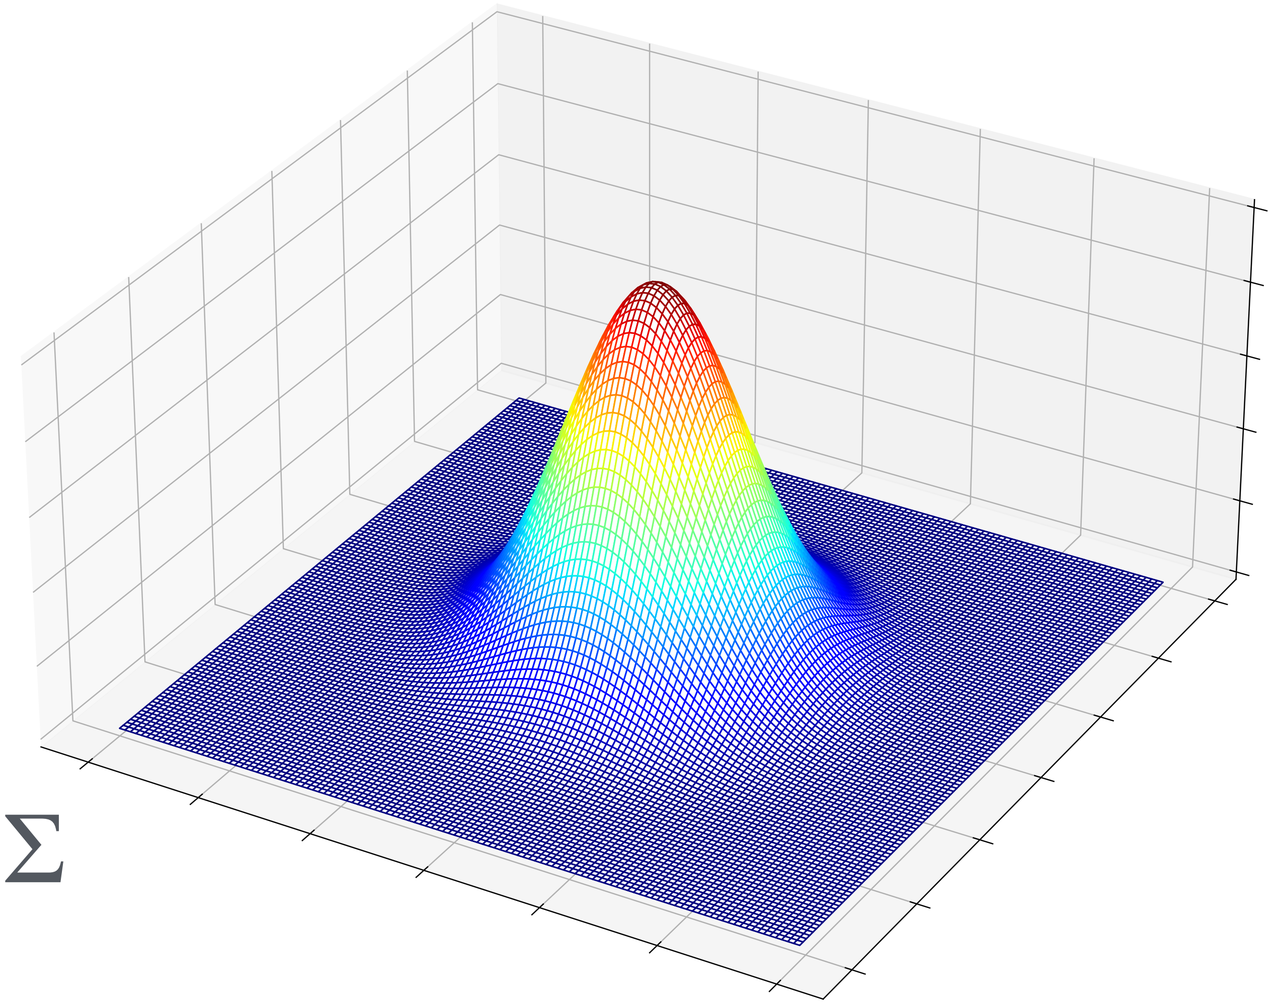

Discrete Gaussians

$$D_{\Lambda,\sqrt{\Sigma}}$$

$$\hat{\rho}_{\sqrt{\Sigma}}(z) = \rho_{\sqrt{\Sigma}}(z)/\rho_{\sqrt{\Sigma}}(\Lambda)$$

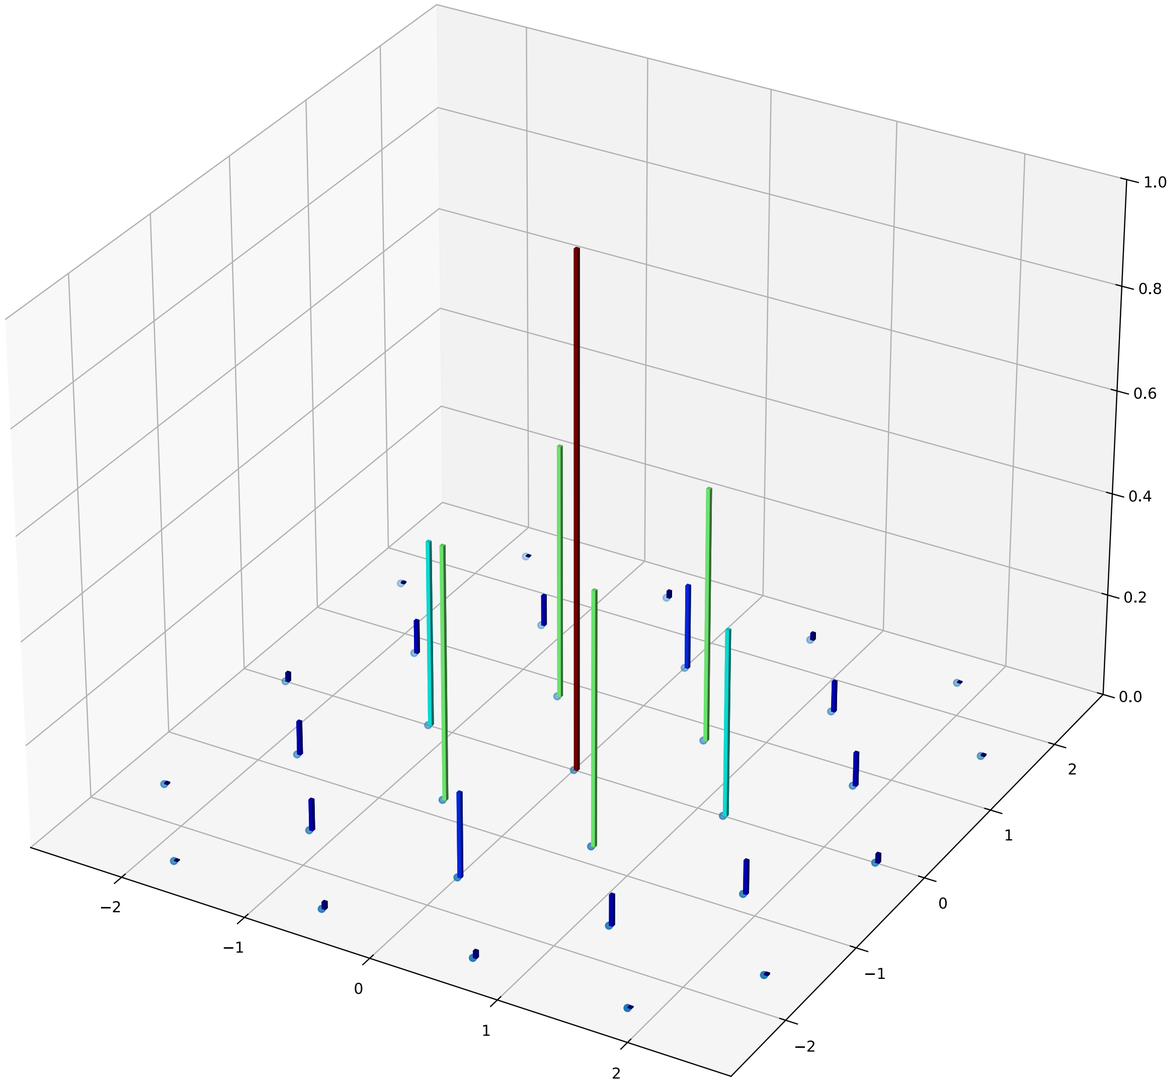

Smoothing

- Invented in [MR04] to “blur” a worst-case lattice (recall Daniele’s talk)
- Incredibly useful for more efficient schemes with statistical security (no drowning)
- **Smoothing parameter** $\eta_\epsilon(\Lambda)$

  Fix some $\epsilon > 0$

  We say $\sigma \geq \eta_\epsilon(\Lambda)$ if $\hat{\rho}_{1/\sigma}(\Lambda^*) \leq 1 + \epsilon$

Smoothing

Continuous Smoothing Lemma [Regev'05]

- Let $\sigma \geq \sqrt{2} \cdot \eta_\epsilon(\Lambda)$
- $x \sim D_{\Lambda,\sigma}$
- $e \sim D_{\mathbb{R}^n,\sigma}$
- $e^* \sim D_{\mathbb{R}^n,\sqrt{2}\sigma}$
- Then $x + e \approx_\epsilon e^*$

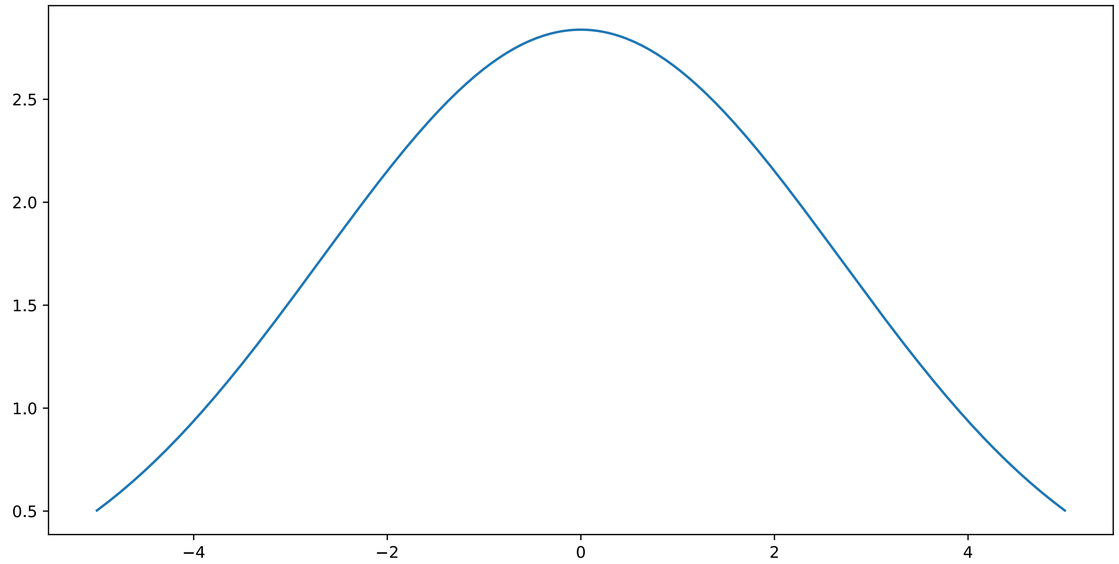

Gaussian Decompositions

Lemma

- Fix matrix $Z \in \mathbb{R}^{m \times n}$
- $e_1 \sim D_{\mathbb{R}^n, \sigma_1}$, $e_2 \sim D_{\mathbb{R}^m, \sqrt{\Sigma_2}}$
- Then $e_3 = Ze_1 + e_2 \sim D_{\mathbb{R}^m, \sqrt{\Sigma_3}}$ where $\Sigma_3 = \sigma_1^2 ZZ^\top + \Sigma_2$
- Conversely, if $\Sigma_2 \geq \Sigma_3 - \sigma_1^2 ZZ^\top \geq 0$ (positive definite), then $e_3 \sim D_{\mathbb{R}^m, \sqrt{\Sigma_3}}$ can be decomposed as $e_3 = Ze_1 + e_2$ for independent $e_1 \sim D_{\mathbb{R}^n, \sigma_1}$ and $e_2 \sim D_{\mathbb{R}^m, \sqrt{\Sigma_2}}$
- If $\Sigma_3 = \sigma^2 I$, then such a Σ_2 exists if $\sigma > \sigma_1 \cdot \sigma_1(Z)$

Proof

- The covariance matrix of $e_3 = Ze_1 + e_2$ is

$$\begin{aligned}\mathbb{E}[e_3 e_3^\top] &= \mathbb{E}[Ze_1 e_1^\top Z^\top] + \mathbb{E}[Ze_1 e_2^\top] + \mathbb{E}[e_2 e_1^\top Z^\top] + \mathbb{E}[e_2 e_2^\top] \\ &= Z\mathbb{E}[e_1 e_1^\top]Z^\top + \mathbb{E}[e_2 e_2^\top] \\ &= Z\sigma_1^2 Z^\top + \Sigma_2 \\ &= \sigma_1^2 ZZ^\top + \Sigma_2\end{aligned}$$

- If $\mathbb{E}[e_3 e_3^\top] = \sigma^2$, then $\Sigma_2 = \sigma^2 I - \sigma_1^2 ZZ^\top$ is positive definite, as for all $x \in \mathbb{R}^n \backslash \{0\}$

$$x^\top \Sigma_2 x = \sigma^2 x^\top x - \sigma_1^2 x^\top ZZ^\top x^\top = \sigma^2 \|x\|^2 - \sigma_1^2 \|Z^\top x\|^2 \geq \sigma^2 \|x\|^2 - \sigma_1^2 \sigma_1(Z)^2 \|x\|^2 > 0$$

Gaussian "Leftover Hash Lemma"

Goal

- Fix some short matrix $Z \in \mathbb{R}^{m \times n}$
- x discrete Gaussian
- e continuous Gaussian
- **Show**: $Zx + e$ Gaussian

Lemma

- $x \sim D_{\Lambda, \sigma_1}$
- $e \sim D_{\mathbb{R}^n, \sigma_2}$ with $\sigma_2^2 I \geq \sigma_1^2 Z Z^\top$
- Then $Zx + e \sim D_{\mathbb{R}^n, \sqrt{\Sigma}}$ with $\Sigma = \sigma_2^2 I + \sigma_1^2 Z Z^\top$

Proof

- Decompose e as $e = Ze_1 + e_2$ with $e_1 \sim D_{\mathbb{R}^n,\sigma_1}$ and $e_2 \sim D_{\mathbb{R}^m,\sqrt{\sigma_2^2 I - \sigma_1^2 ZZ^\top}}$ as per last lemma
- Then $Zx + e = Zx + Ze_1 + e_2 = Z(x + e_1) + e_2$
- $x + e_1$ statistically close to $x' \sim D_{\mathbb{R}^n,\sqrt{2}\sigma_1}$ as per smoothing lemma
- Hence $Zx' + e_2$ follows $D_{R^m,\sqrt{\sigma_2^2 + \sigma_1^2 ZZ^\top}}$

Application: FHE Circuit Privacy [BDMW16]

Recall

- Gadget Matrix

$$G = \begin{bmatrix} 1,2,\ldots,2^{\log q} & & \\ & 1,2,\ldots,2^{\log q} & \\ & & \ddots \\ & & & 1,2,\ldots,2^{\log q} \end{bmatrix}$$

- Randomised $G_\sigma^{-1}(\cdot)$

  $x \sim G_\sigma^{-1}(c)$ follows $D_{Z^m,\sigma}$ conditioned on $Gx = c$

Application: FHE Circuit Privacy [BDMW16]

- Alternative homomorphic evaluation: $C = C_1 \cdot G_{rand}^{-1}(C_2) + \begin{pmatrix} 0 \\ \tilde{e} \end{pmatrix}$

- Show: C statistically close to a fresh encryption of $m_1 \cdot m_2$

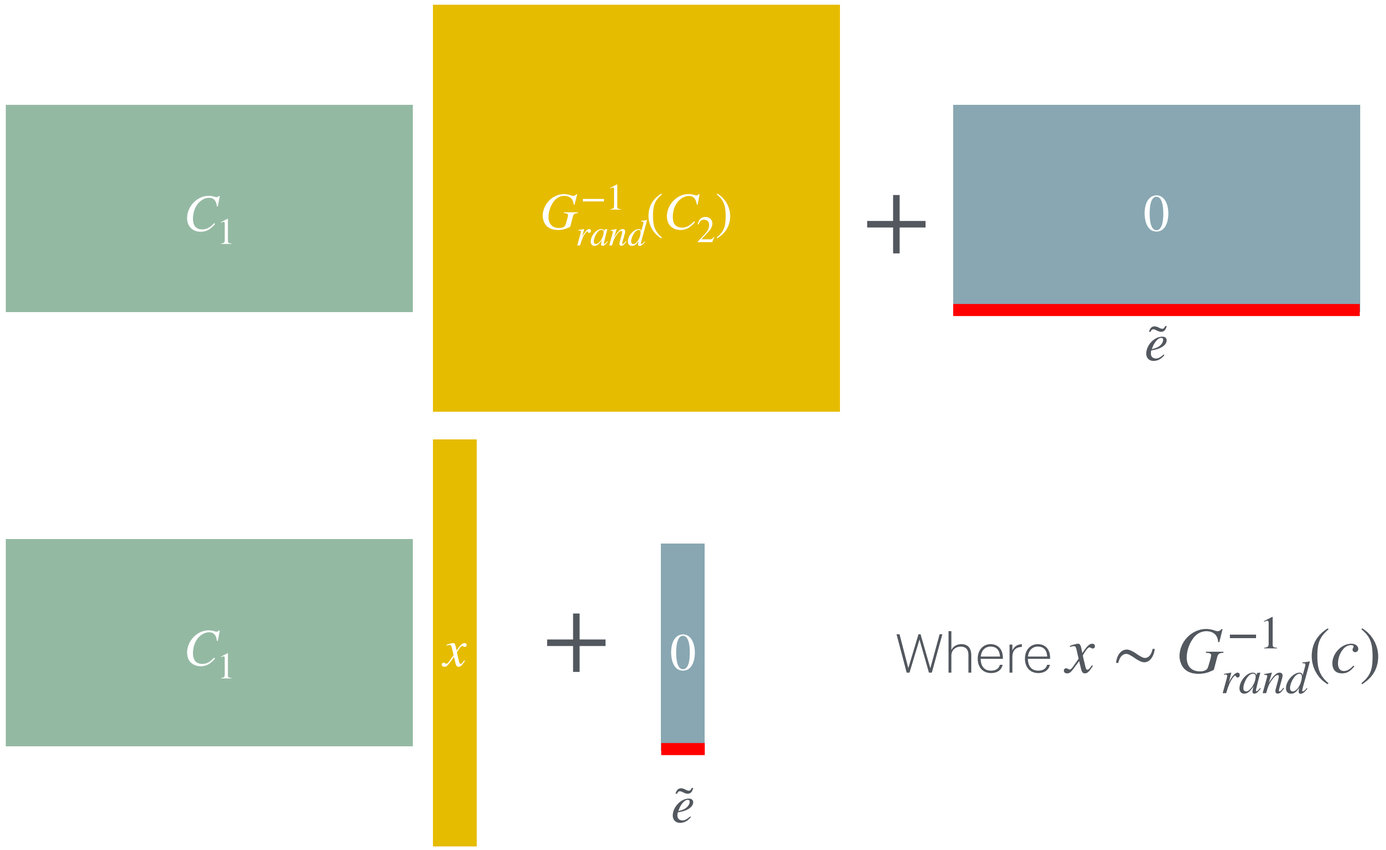

Application: FHE Circuit Privacy [BDMW16]

- Show: $\boldsymbol{C}$ statistically close to a fresh encryption of m

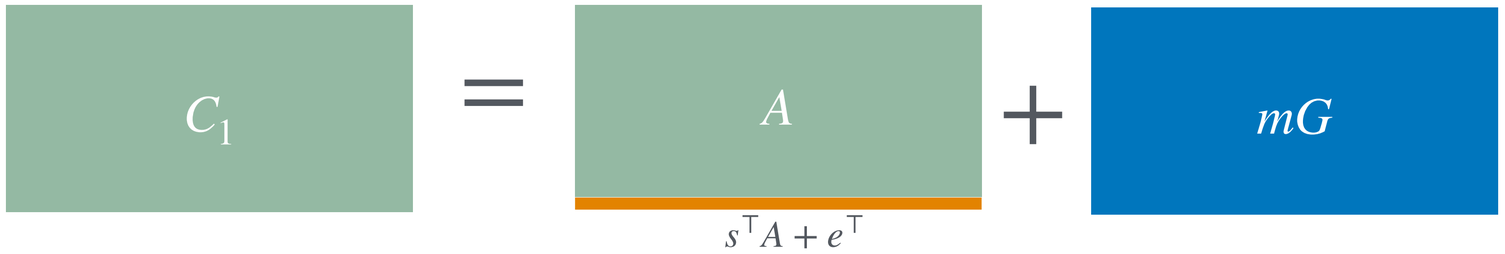

$x \sim G^{-1}_{rand}(c)$

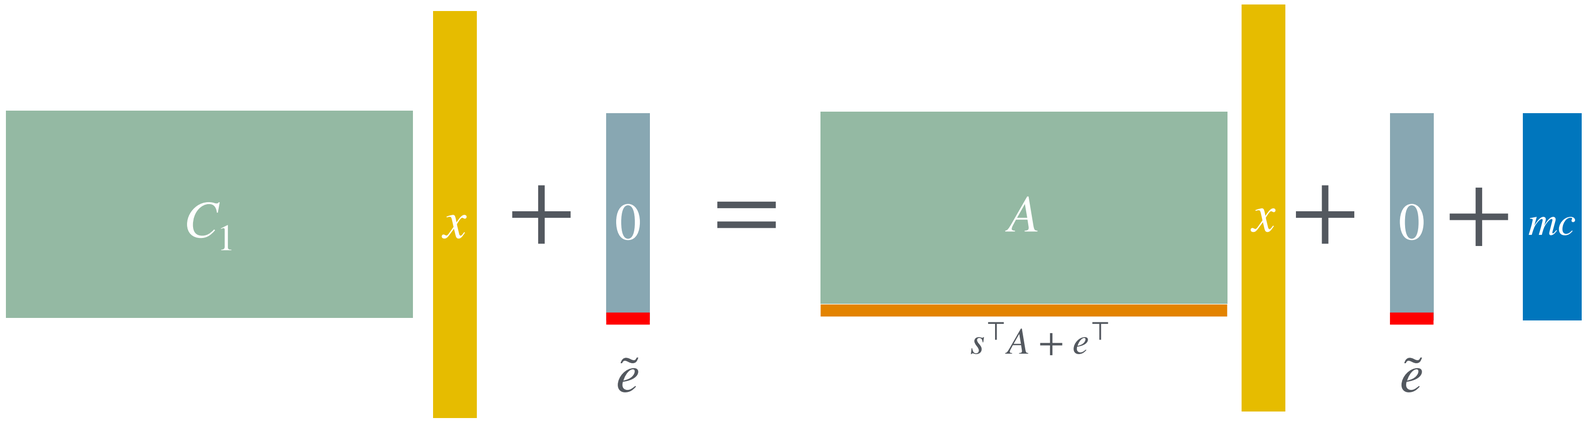

Application: FHE Circuit Privacy [BDMW16]

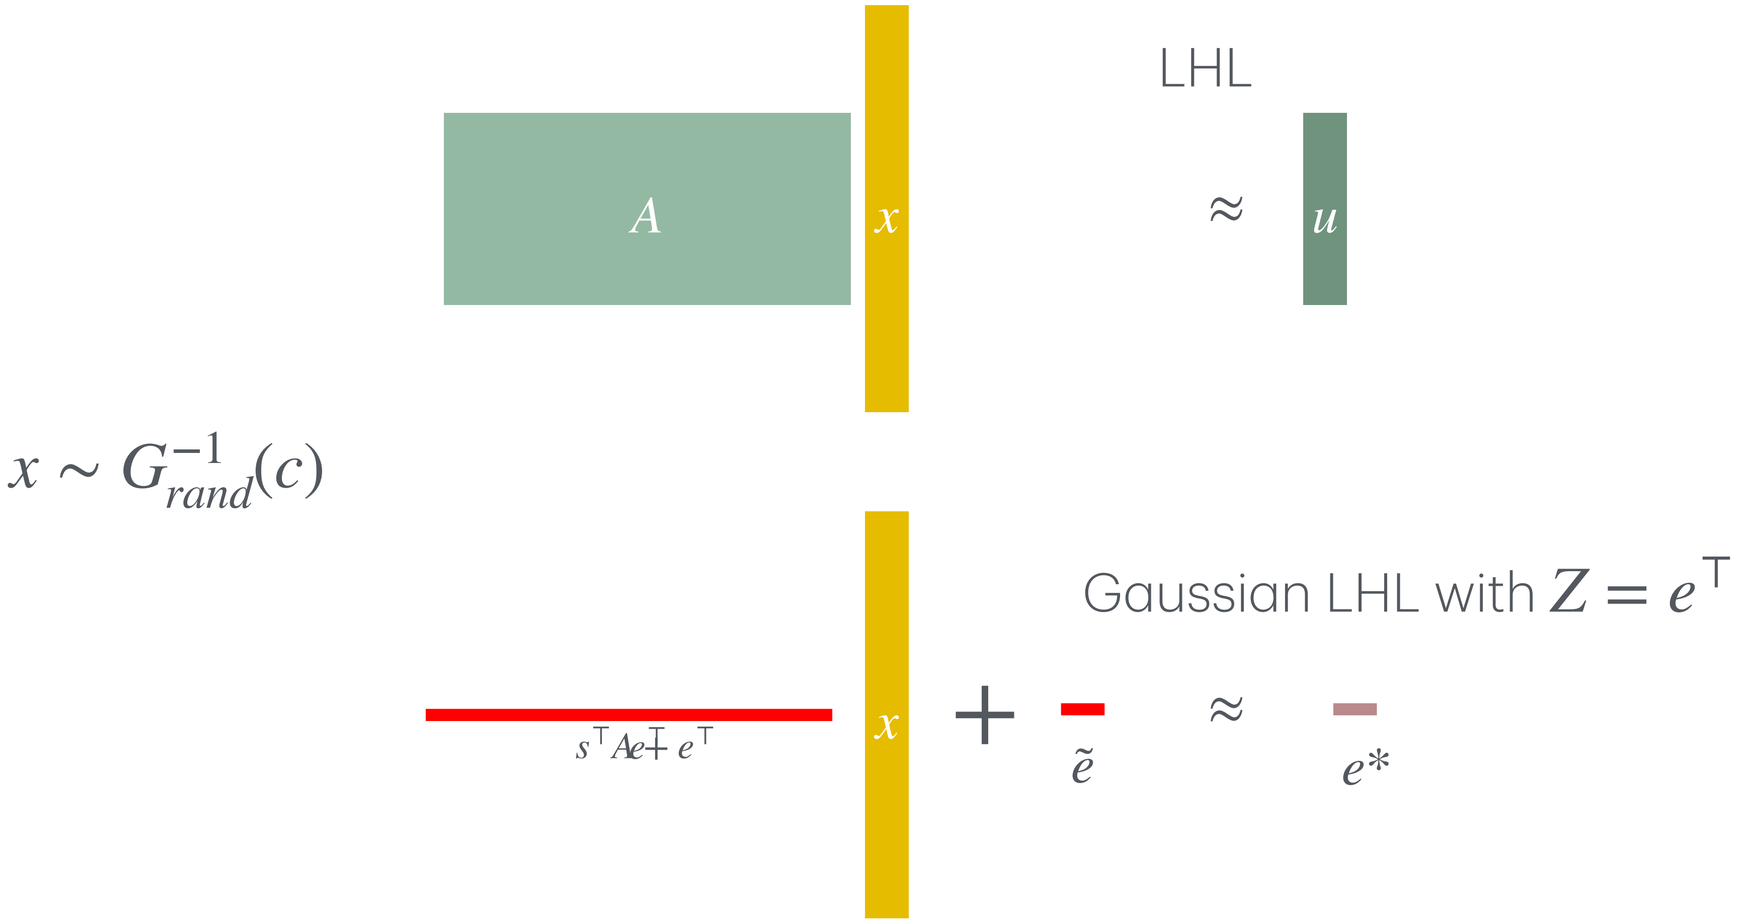

Application: Entropic LWE

- LWE: Secret $\boldsymbol{s}$ is uniform
- Entropic LWE: $\boldsymbol{s}$ only comes from a min-entropy distribution
- Think: LWE with leaky secret

Learning with Errors [Reg05]

Decisional Version

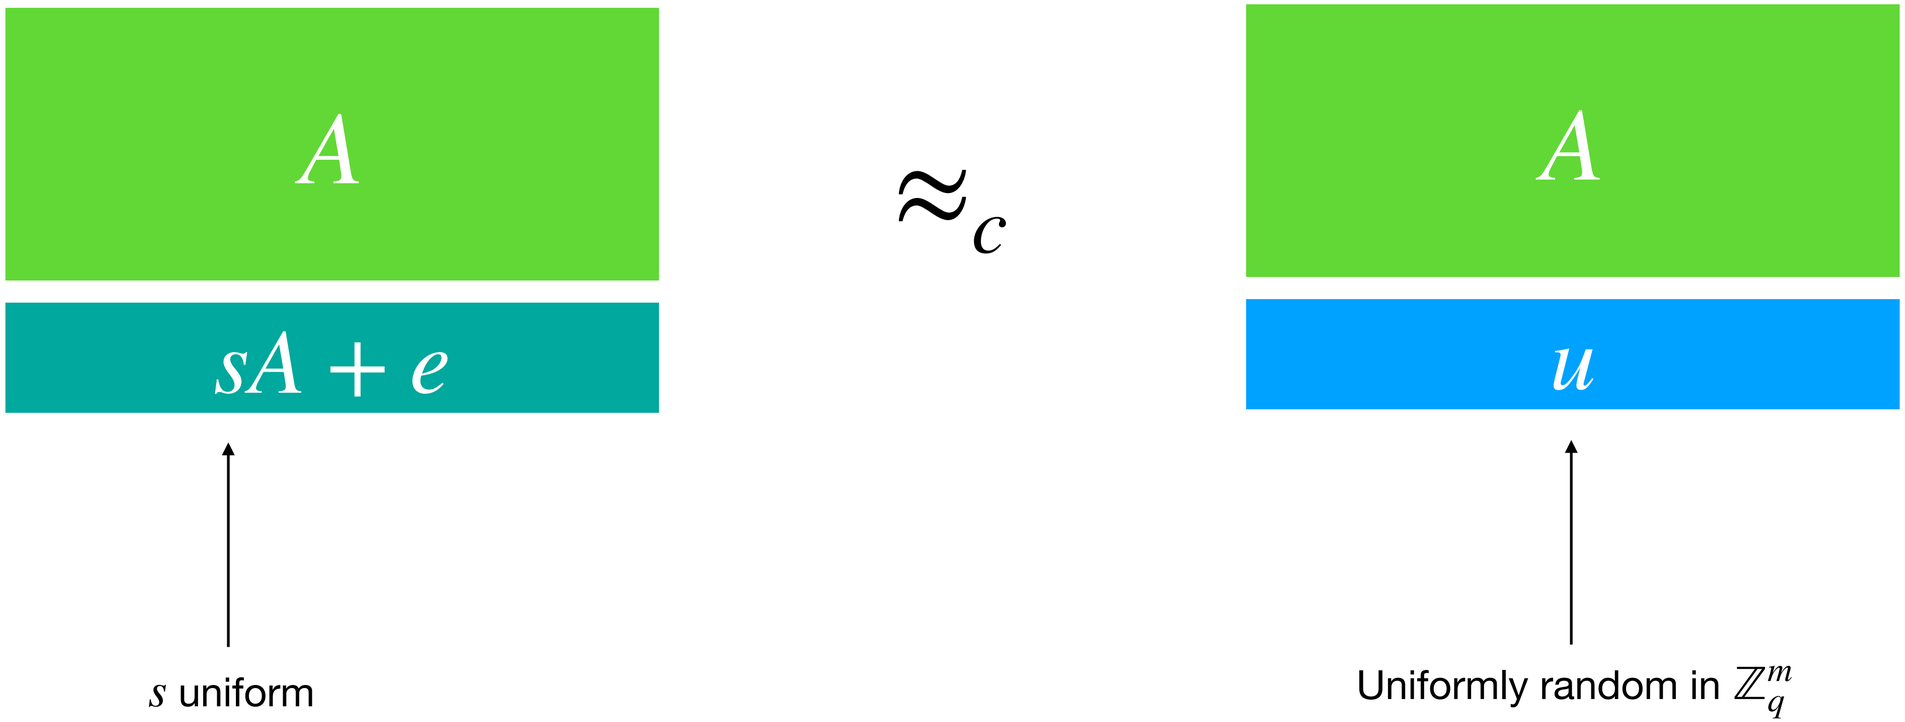

Entropic LWE

Decisional Version

The Lossiness Technique [GKPV10]

- Common proof strategy: Replace uniformly chosen matrix A with a pseudorandom matrix which has unusually many short vectors in its (row-)span
- Now use that $A, sA + e$ loses information about s

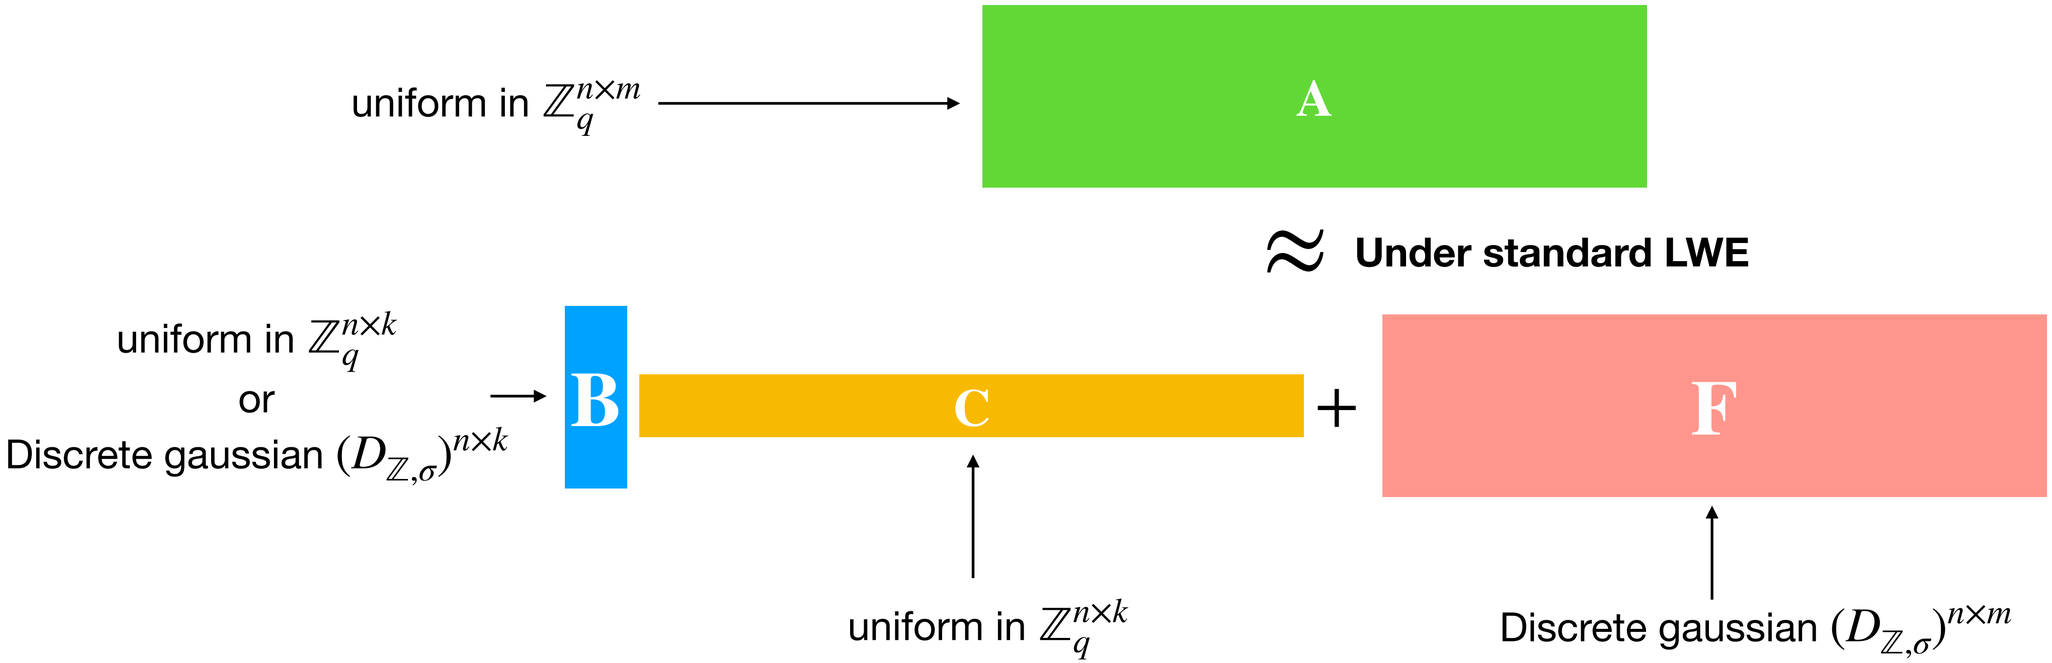

The Lossiness Technique [GKPV10]

Warmup: with drowning

Chosen from a min-entropy distribution $\mathcal{S}$ supported on $\{0,1\}^n$

$$A, sA + e \qquad\qquad A, u$$

$$\approx_{LWE} \qquad\qquad \approx_{LWE}$$

$$BC + F, s(BC + F) + e \qquad\qquad BC + F, u$$

$$= \qquad\qquad \approx_{LWE}$$

$$BC + F, sBC + sF + e \approx_s BC + F, sBC + e' \approx_{LHL} BC + F, tC + e'$$

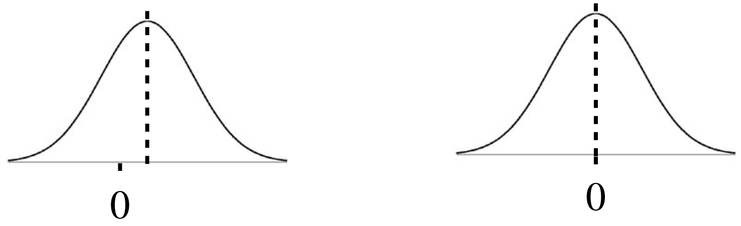

Noise-Lossiness [BD'20]

- Fix distribution $\mathcal{S}$ supported on $\mathbb{Z}_q^n$
- $s \leftarrow \mathcal{S}$, e is a gaussian with parameter σ
- Measures the information lost about s after passing it through a gaussian channel
- Different Perspective: How bad is $\mathcal{S}$ as an error correcting code?

$$\tilde{H}_\infty(s \mid s+e) = -\log(\Pr_{s,e}[\mathcal{A}^*(s+e) = s])$$

$\mathcal{A}^*$ is maximum likelihood decoder for $\mathcal{S}$

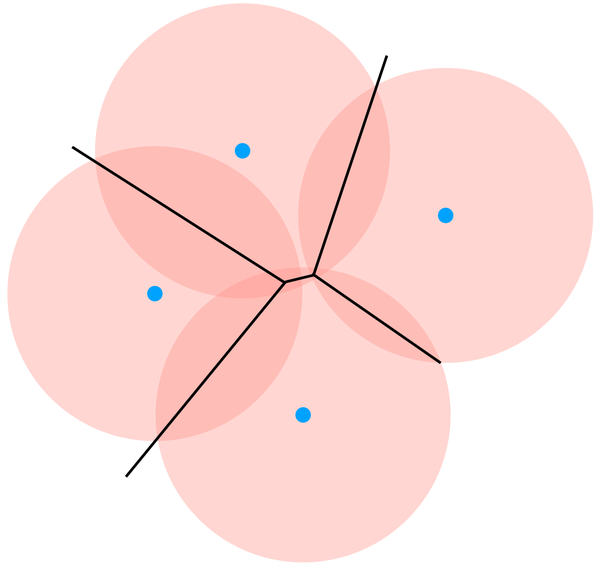

Noise Lossiness: General Distributions

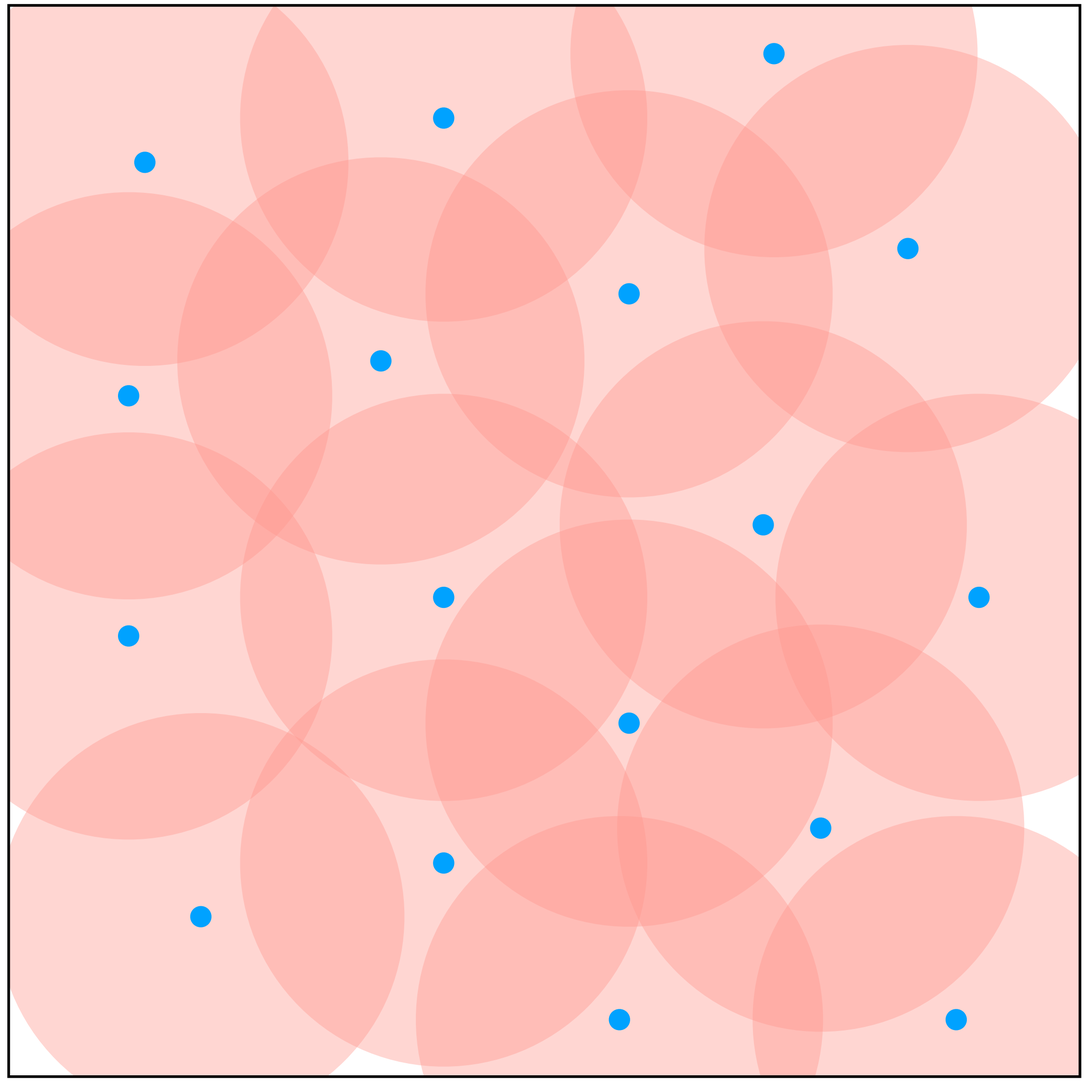

$$\mathbb{Z}_q^n$$

$$\tilde{H}_\infty(s \mid s + e) \geq H_\infty(s) - n \cdot \log(q/\sigma) - 1$$

Noise Lossiness: Short Distributions

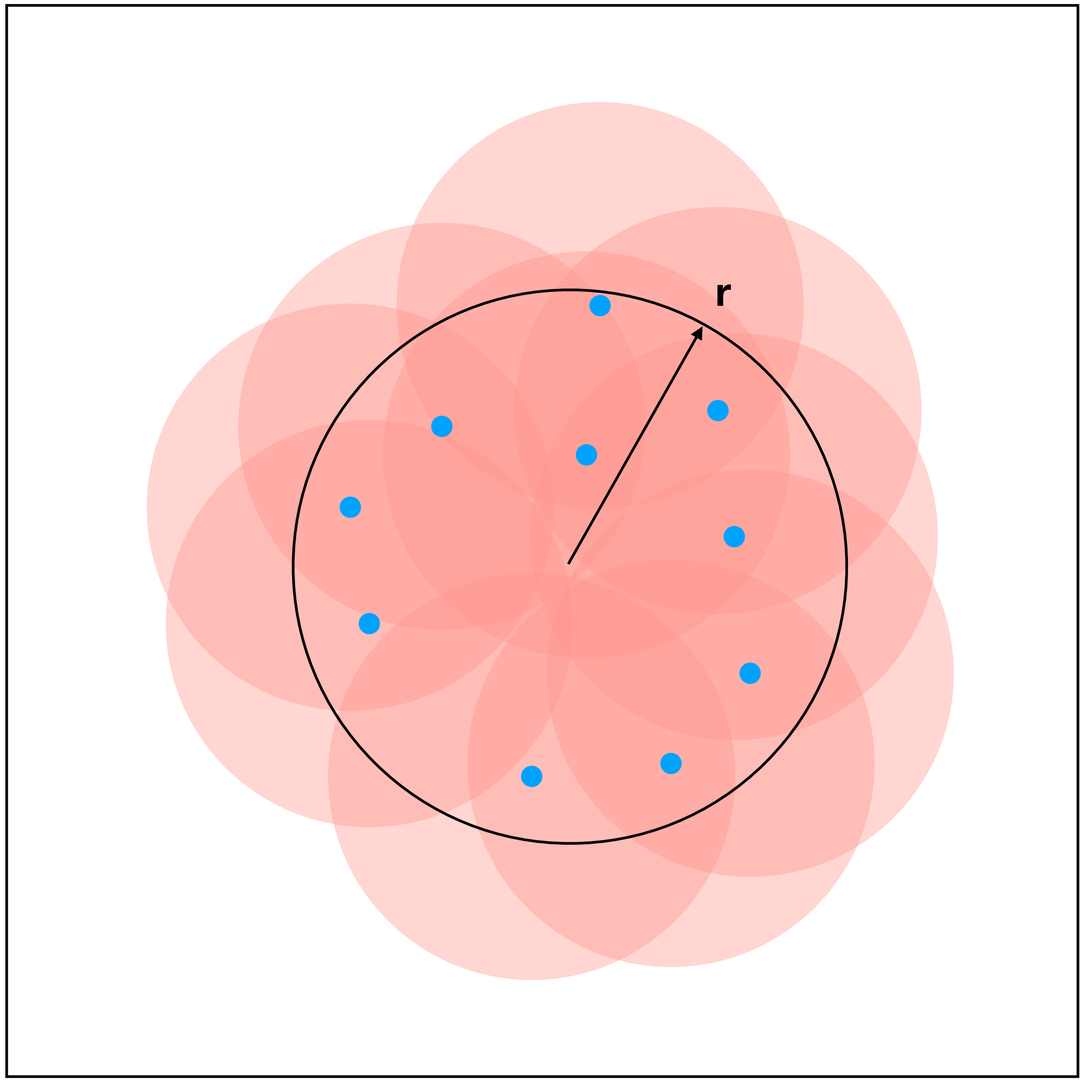

$$\mathbb{Z}_q^n$$

$$\tilde{H}_\infty(s \mid s + e) \geq H_\infty(s) - 2r\sqrt{n}/\sigma$$

From Noise-Lossiness to Hardness of Entropic LWE [BD'20]

$$A, sA + e$$

$$\approx_{LWE}$$

$$BC + F, s(BC + F) + e$$

$$=$$

$$BC + F, sBC + sF + e$$

$$=$$

$$BC + F, sBC + sF + e_1F + e_2$$

$$=$$

$$BC + F, sBC + (s + e_1)F + e_2$$

From Noise-Lossiness to Hardness of Entropic LWE [BD'20]

$$
\begin{array}{l}
A, sA + e \\
\quad \approx \\
BC + F, s(BC + F) + e \\
\quad = \\
BC + F, sBC + sF + e \\
\quad = \\
BC + F, sBC + sF + e_1F + e_2 \\
\quad = \\
BC + F, sBC + (s + e_1)F + e_2
\end{array}
\qquad
\begin{array}{l}
A, u \\
\quad \approx_{LWE} \\
BC + F, u \\
\quad \approx_{LWE}
\end{array}
$$

$$
\approx_{LHL} BC + F, tC + (s + e_1)F + e_2 = BC + F, tC + sF + e
$$

Conclusions and Further Applications

- Drowning is a general technique to remove *noise artefacts*
- *Requires a stronger LWE assumption and leads to (practically) undesirable parameters*
- *Gaussian Smoothing is (sometimes) an alternative to drowning*
- *Uses specific features of Gaussians (Decomposability), and works with poly modulus-to-noise LWE*
- *Other Recent Applications:*
- Ring LWE with entropic hardness [BD'20b]
- Laconic Encryption (Simple type of Laconic Function Evalution) with polynomial modulus-to-noise ratio [DKFLMR'23]

Thanks!