Hardness of LWE on General Entropic Distributions

Nico Döttling
Helmholtz Center for Information Security (CISPA)

Joint work with Zvika Brakerski

Leakage Resilient Cryptography

- **General Question:** What if the secret key of a scheme was accidentally chosen from a not fully random distribution or additional side-information about the secret key was later leaked?

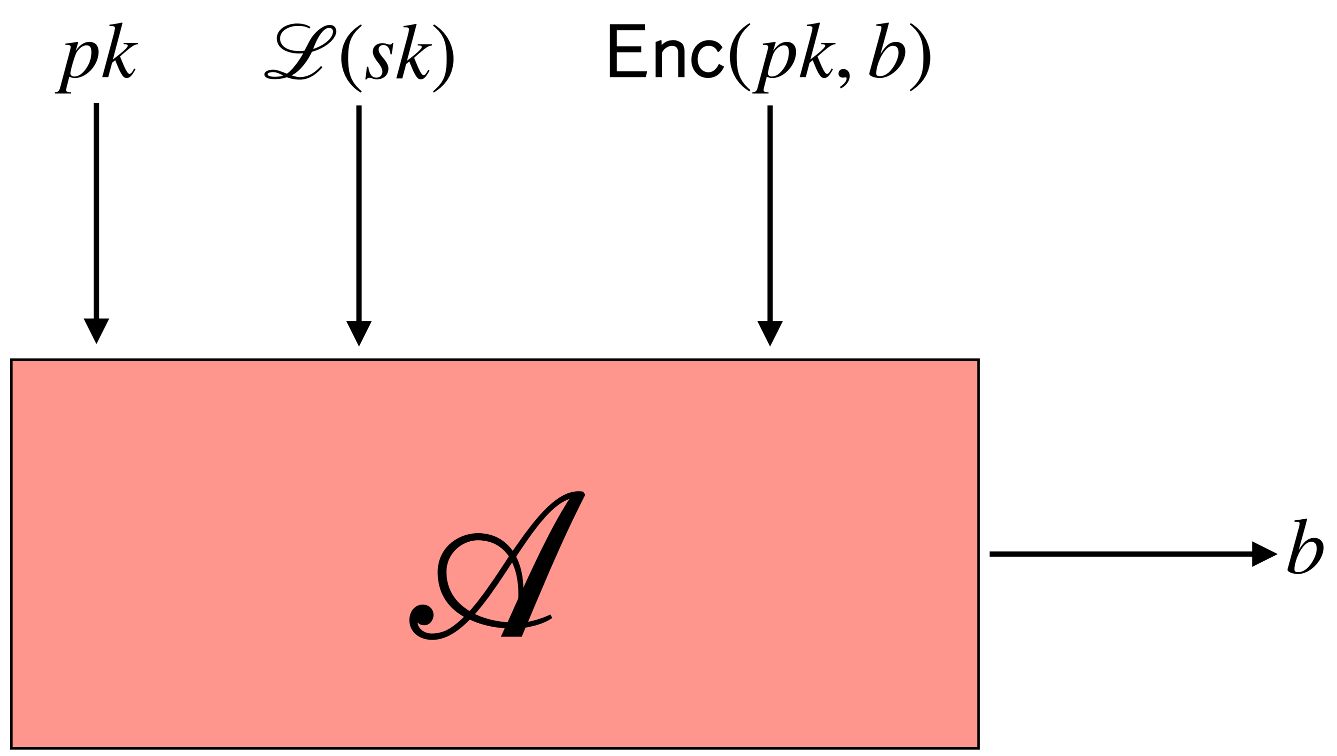

Overview

- Entropic LWE: LWE with weak secrets
- What was known
- Our Approach
- Lower Bounds

Learning with Errors [Reg05]

Search Version:

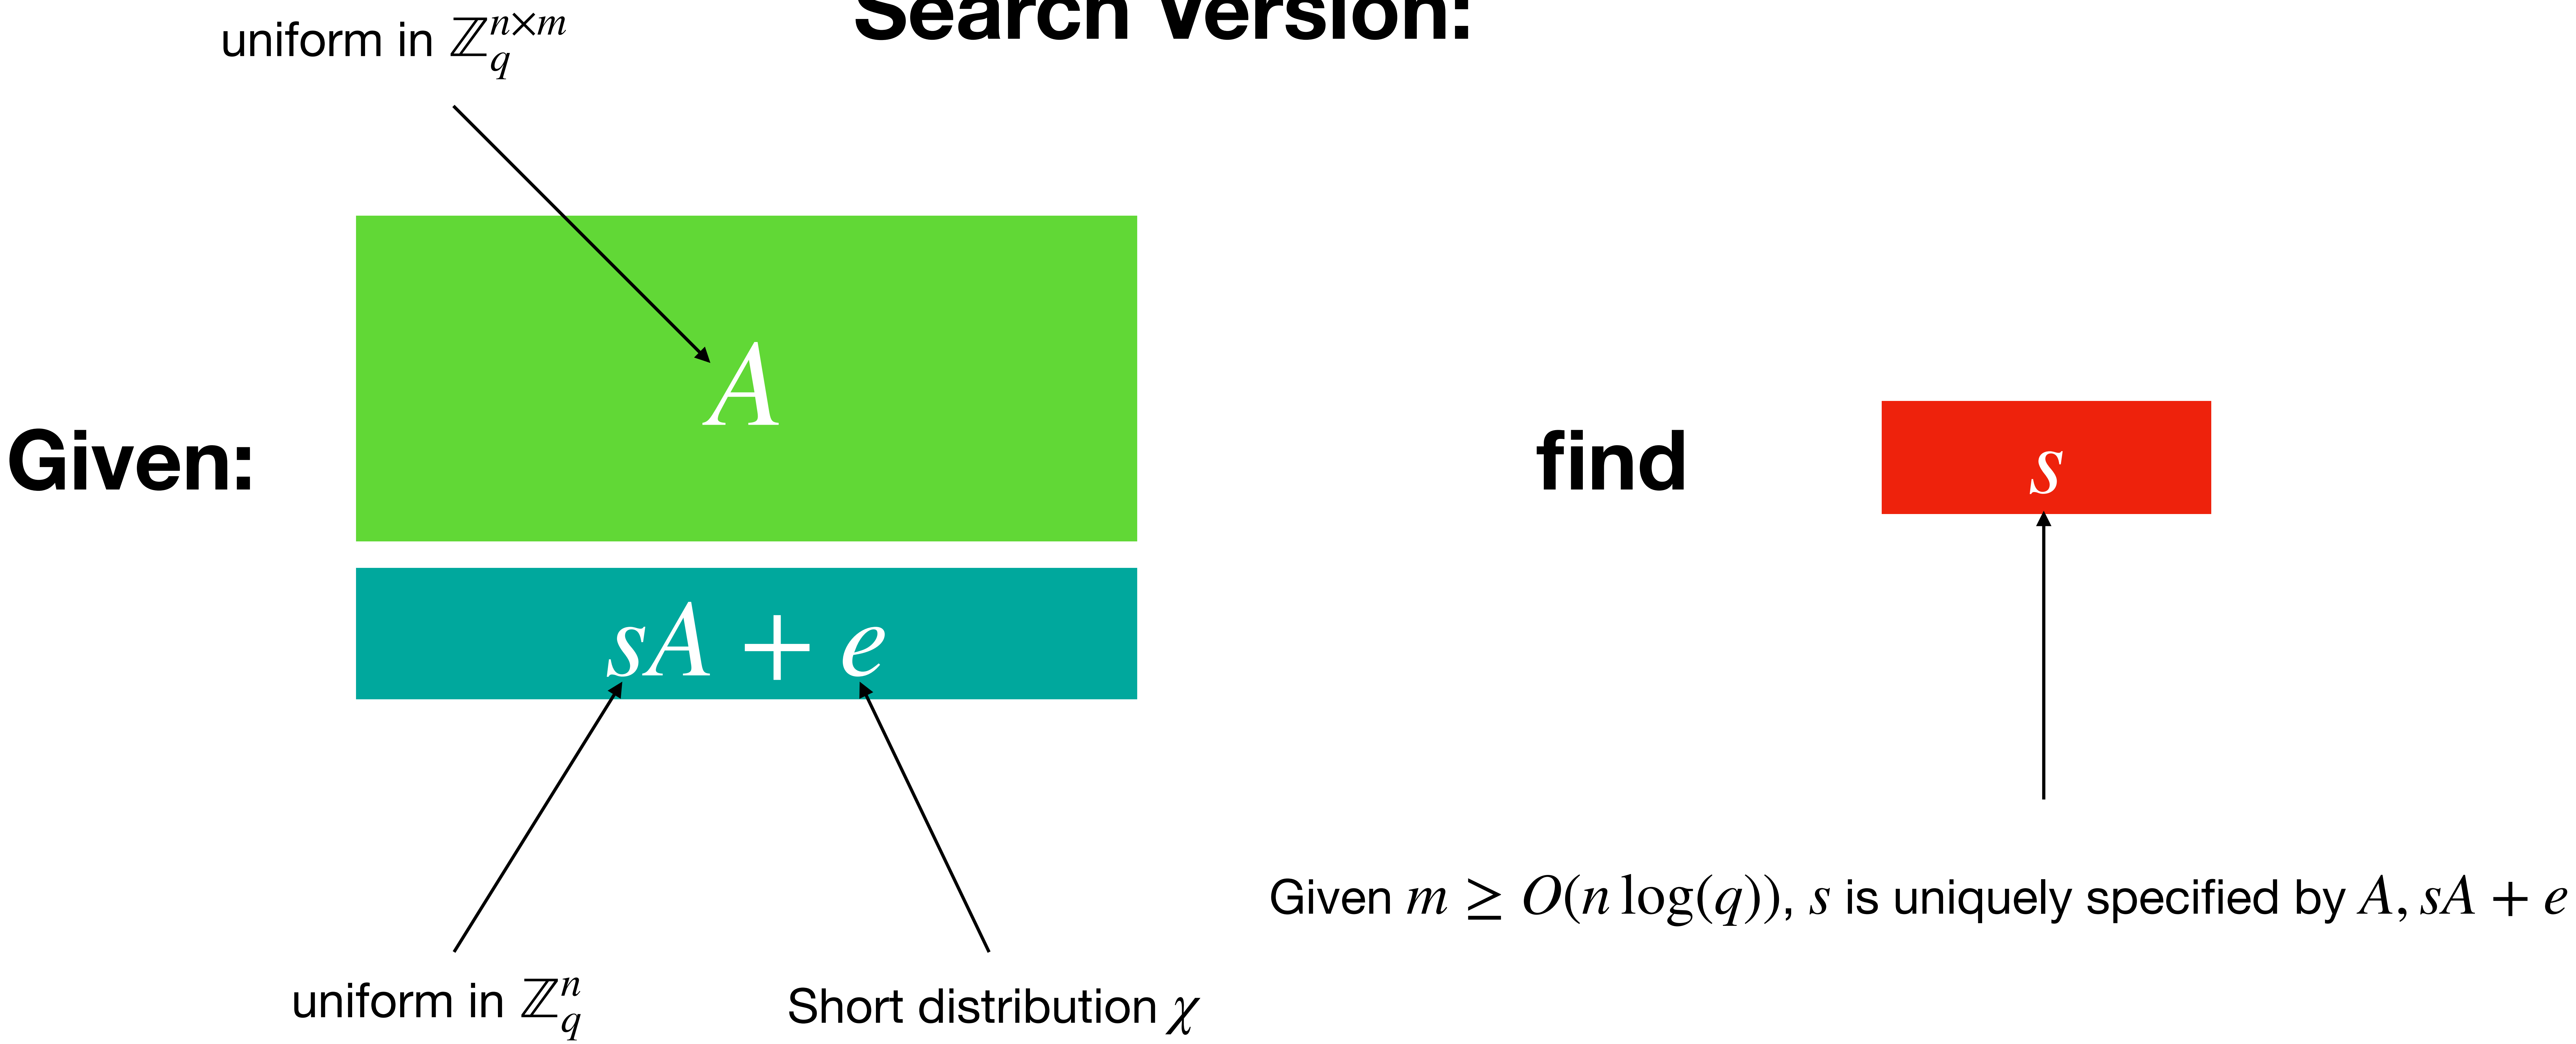

Learning with Errors [Reg05]

Decisional Version:

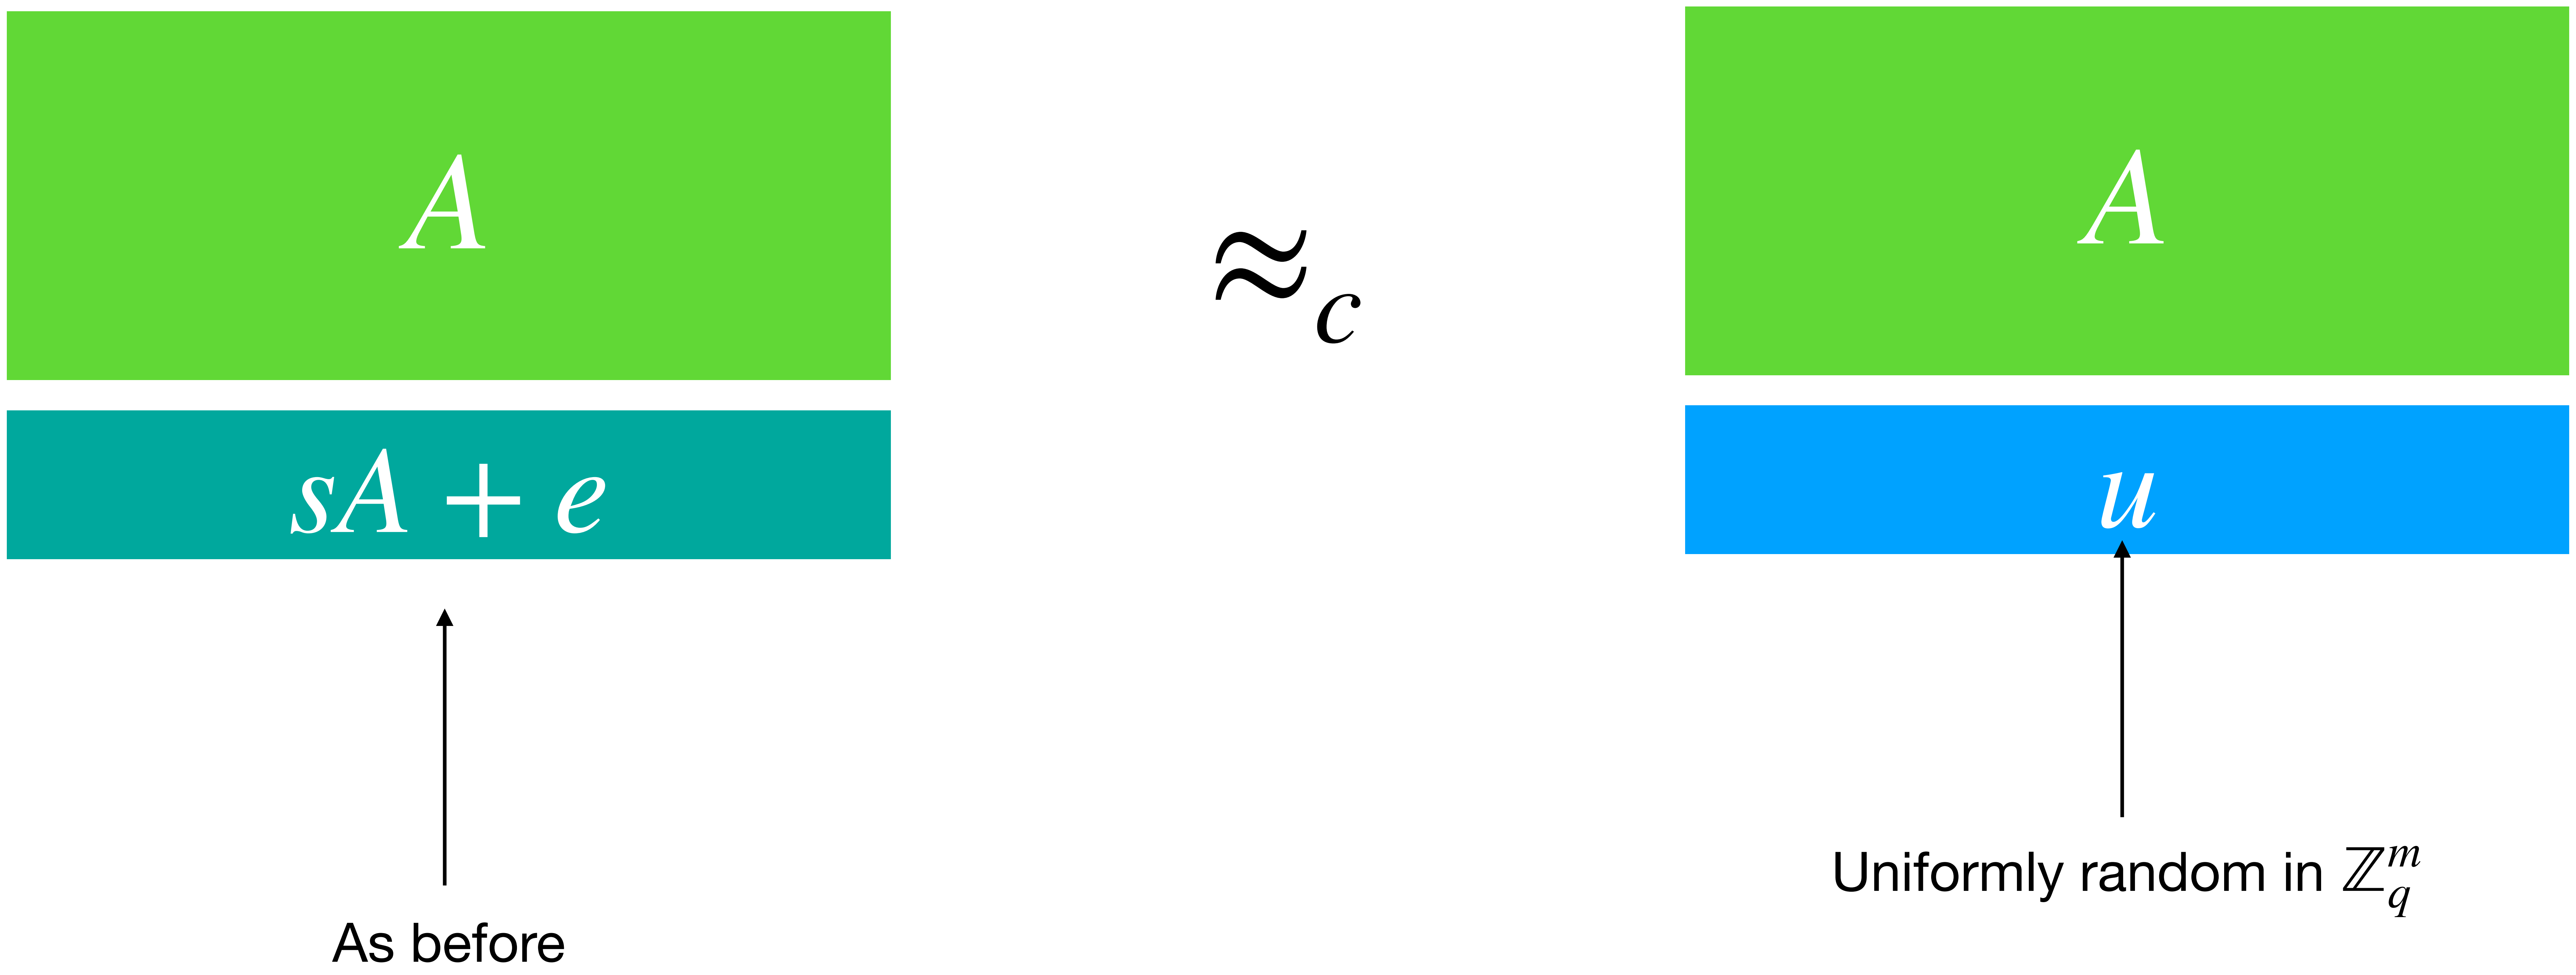

Worst-Case Hardness of LWE

- For **gaussian** error distributions D_σ, 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 $\alpha = q/\sigma$

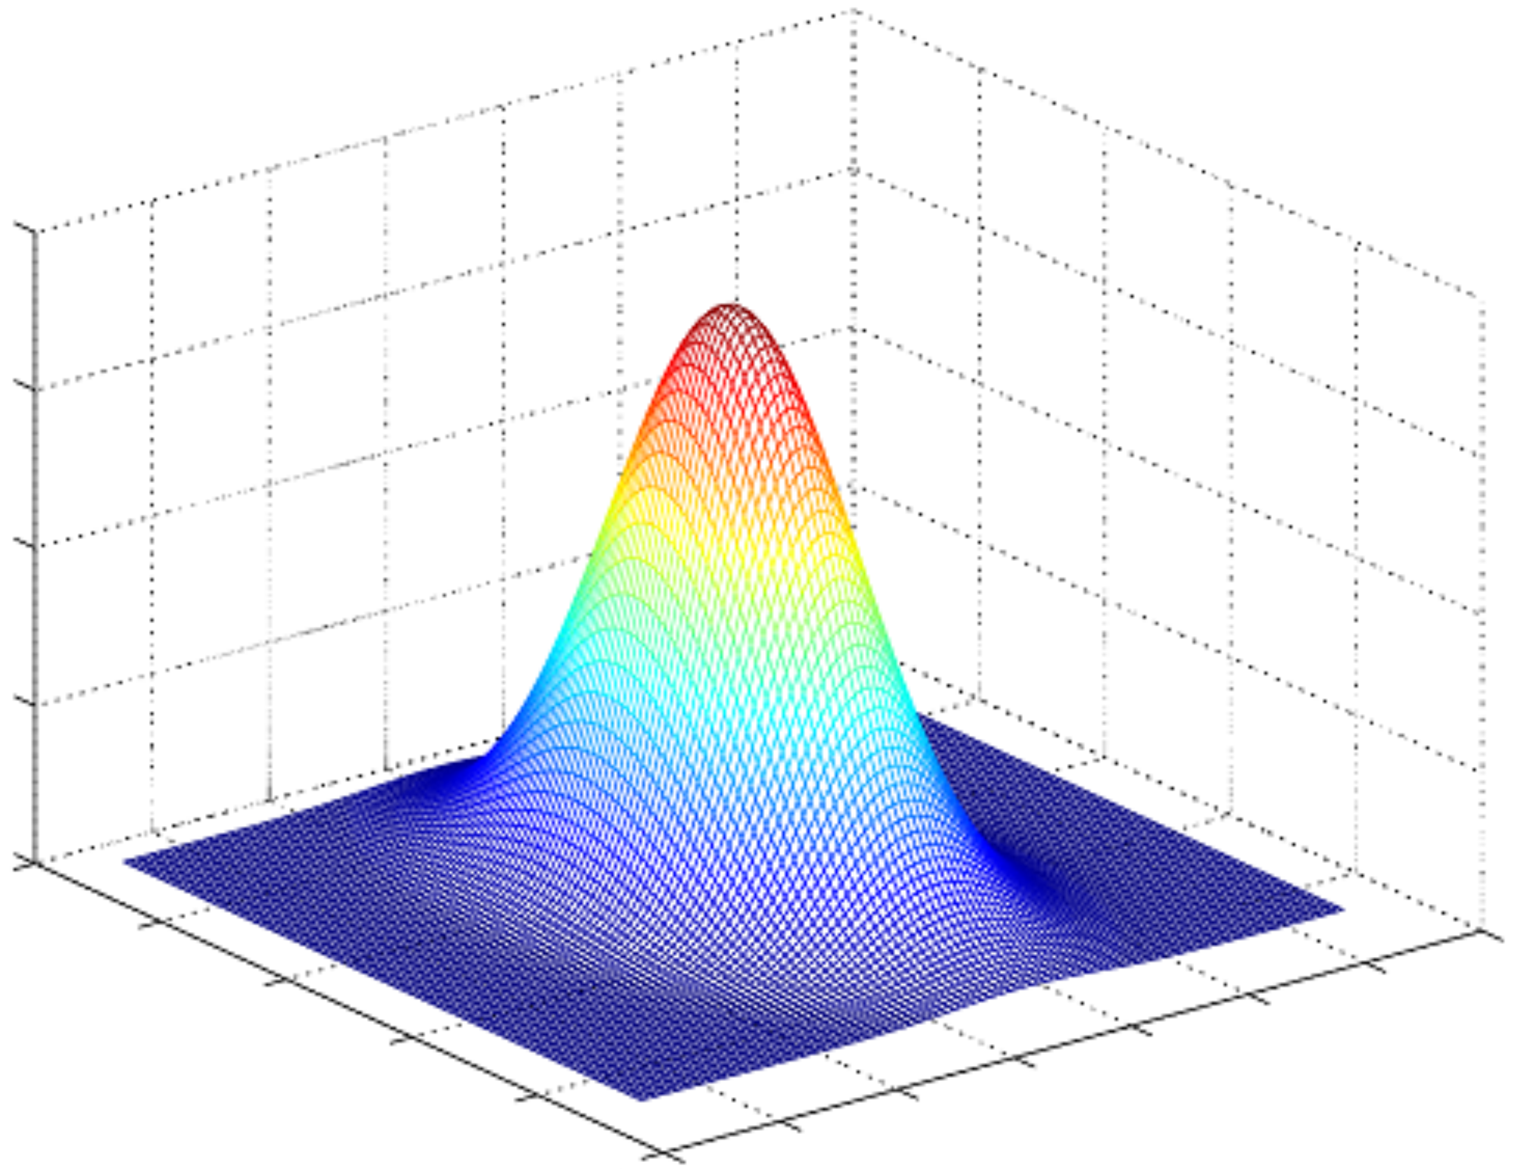

LWE-based Crypto

- Public Key Encryption
- Oblivious Transfer/Mutliparty Computation
- Fully Homomorphic Encryption (only under LWE)
- Attribute-based Encryption for all Circuits (only under LWE)
- Non-Interactive Zero-Knowledge

Leakage Resilience on Assumption Level

- For many schemes the LWE secret s constitutes the secret key
- A leakage resilient version of LWE we can generically add leakage resilience to many of these schemes, e.g. Regev encryption
- Tuesday Session: Version of LWE with (very strong) leakage can be used to build iO
- Given the importance of LWE, this can even be considered a self-supporting goal

Entropic LWE

Search Version:

uniform in $\mathbb{Z}_q^{n \times m}$

Given: A

$sA + e$

find s

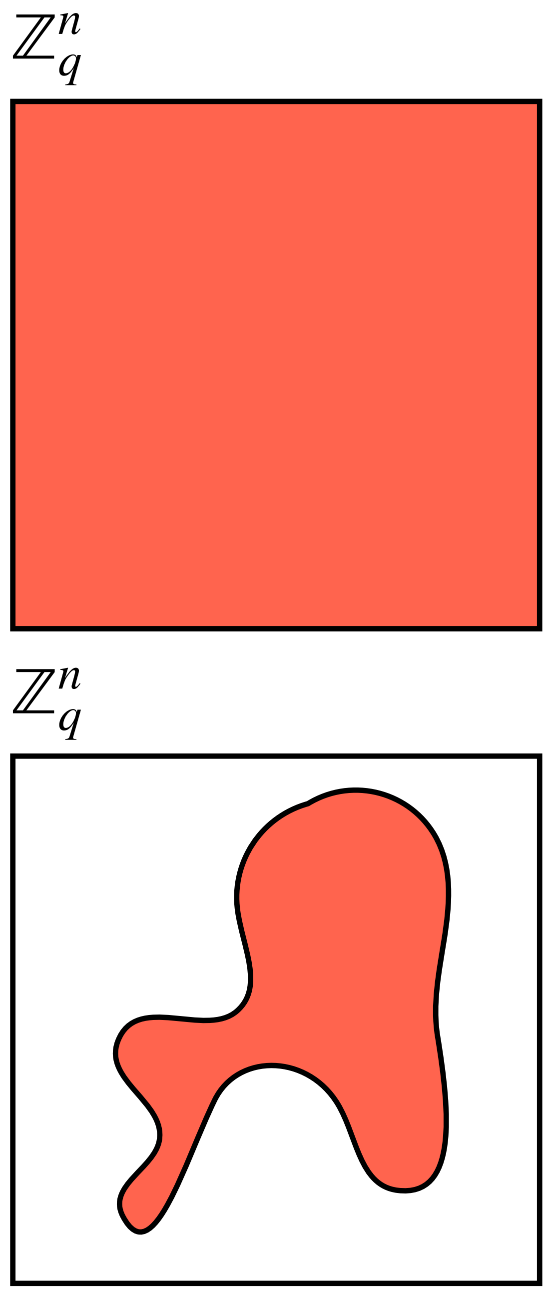

~~uniform in $\mathbb{Z}_q^n$~~

gaussian with parameter σ

chosen from a min-entropy distribution $\mathcal{S}$

<u>Distribution $\mathcal{S}$ is adversarially chosen from a class of distributions</u>

Entropic LWE

Decisional Version:

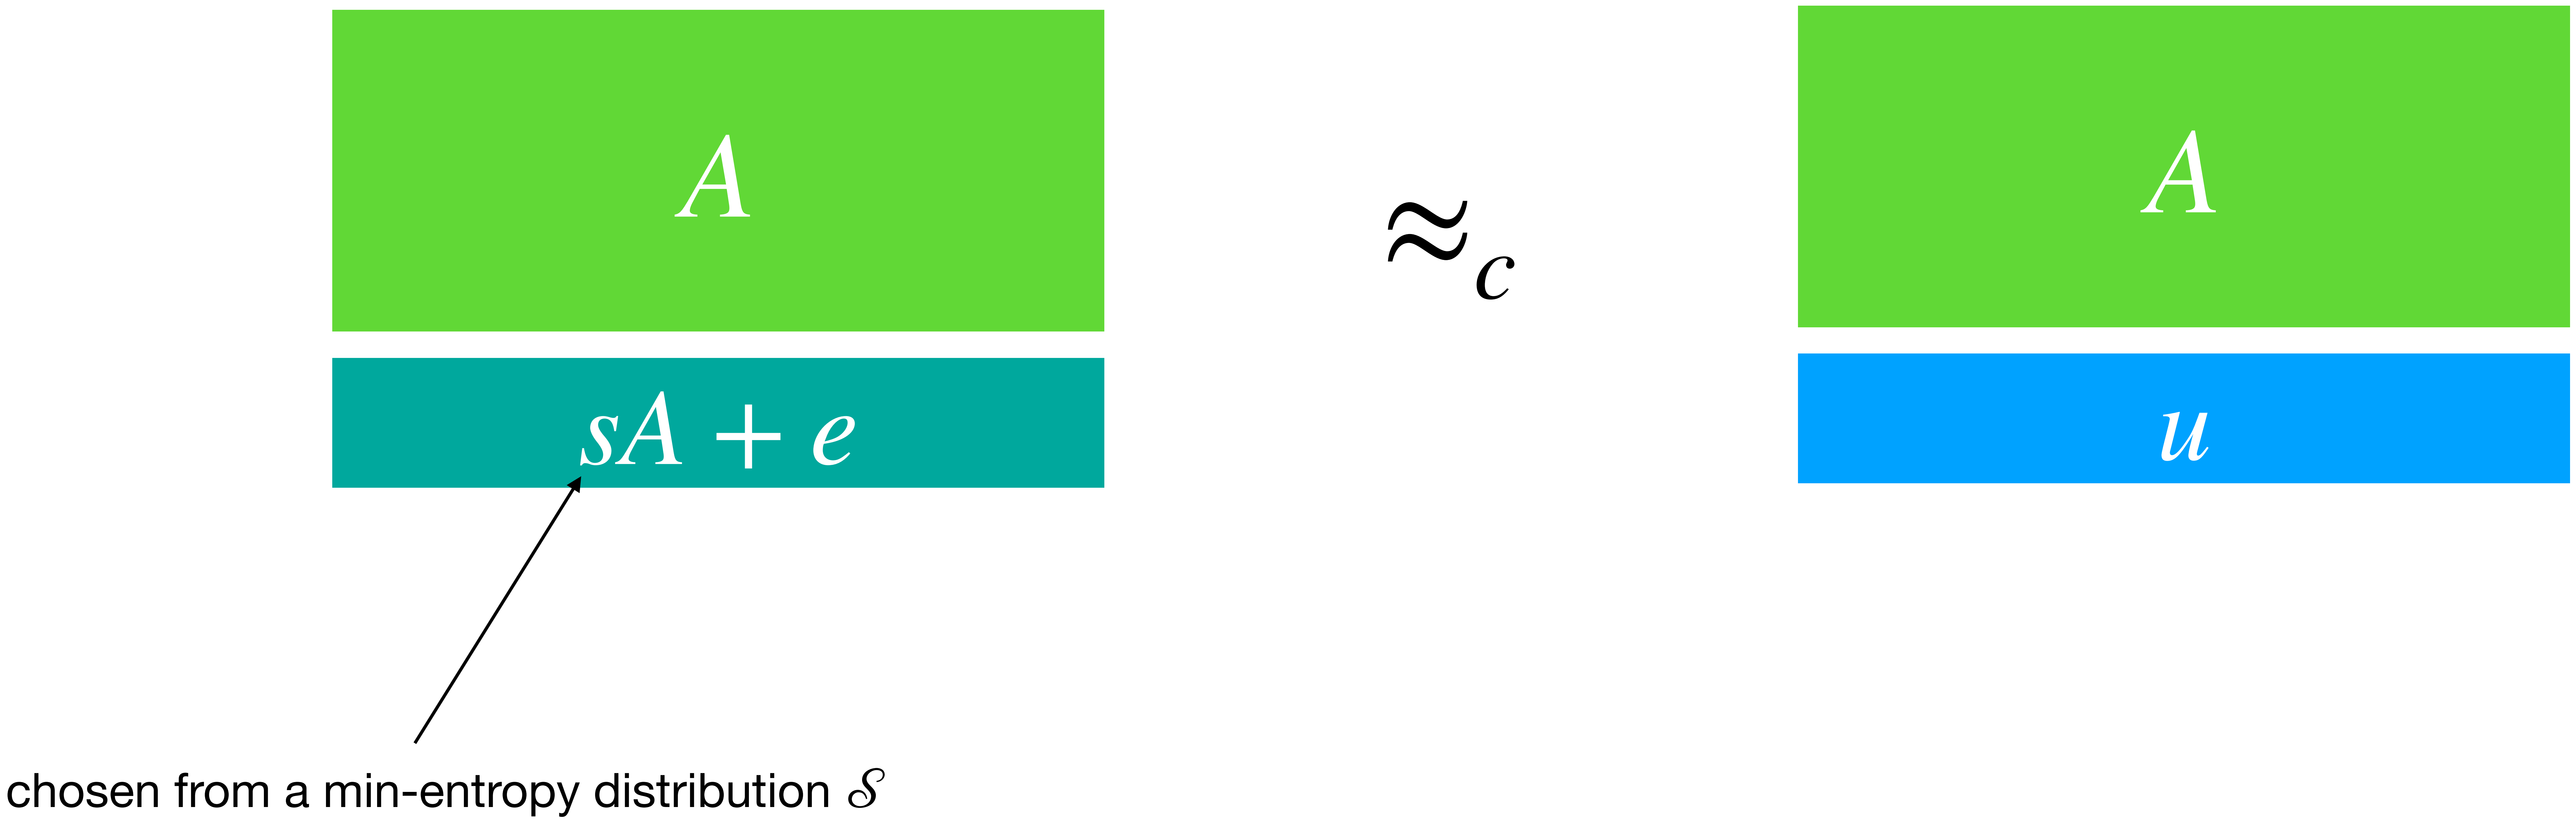

Hardness LWE with Entropic Secrets

- [GKPV10]: For super-polynomial α, reduction from LWE to eLWE for entropic secrets supported on short vectors
- [BLPRS13]: Hardness of LWE with binary secrets which preserves α exactly
- [AKPW13]: More refined version of the [GKPV10] argument, α degrades polynomially in the number of samples q, but also limited to short secrets

Recap: The Lossiness Technique [GKPV10]

The Lossiness Technique

- 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

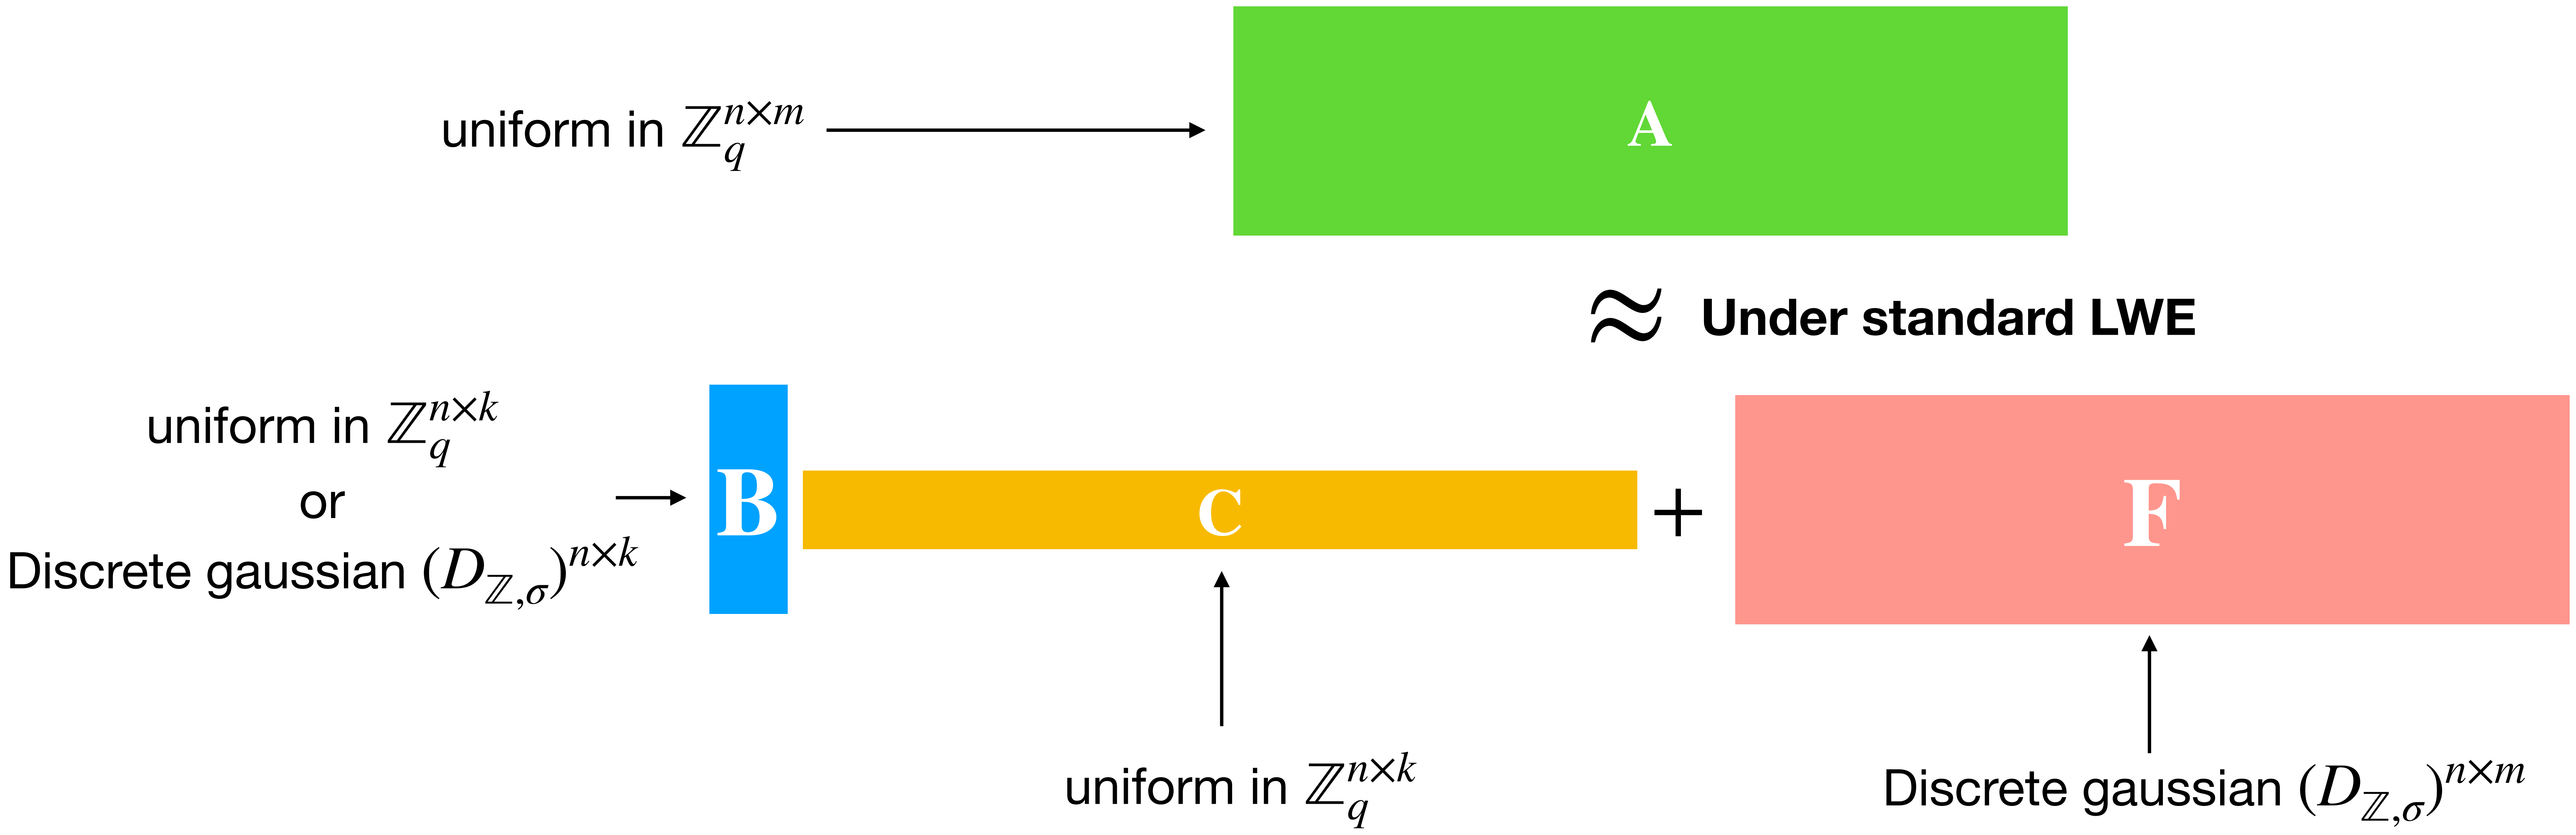

The Lossiness Technique [GKPV10]

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

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

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

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

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

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

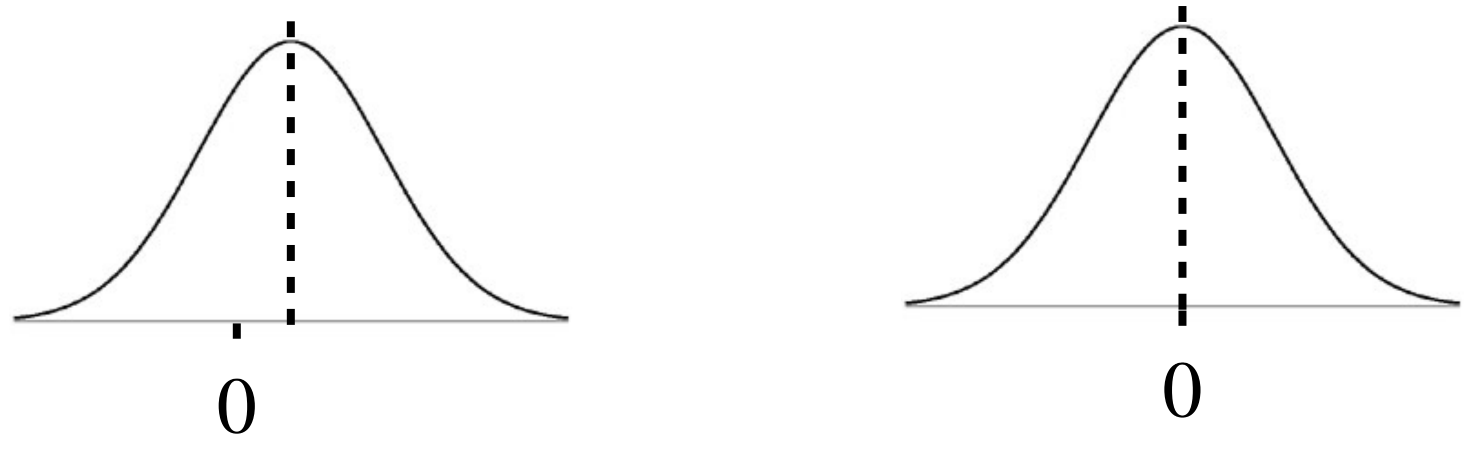

The Lossiness Technique

- This proof fundamentally relies on the fact that s is short
- Otherwise the term sF cannot be “drowned” by e
- Furthermore: modulus-to-noise ratio deteriorates drastically (overcome by [AKPW13])
- Natural Question: Is the requirement of s being short fundamental or rather a limitation of the proof technique?

Entropic LWE on General Min-Entropy Distributions via Gentle Flooding at the Source

Our Approach

- We also pursue lossiness approach, but with a twist
- Change of Perspective: Instead of analyzing the interference of the secret with the noise term, we analyze what effect the noise has on the secret directly
- We relate this to a new quantity we call *noise-lossiness* of the secret s

Noise-Lossiness

- Fix a distribution of secrets $\mathcal{S}$ 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?

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

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

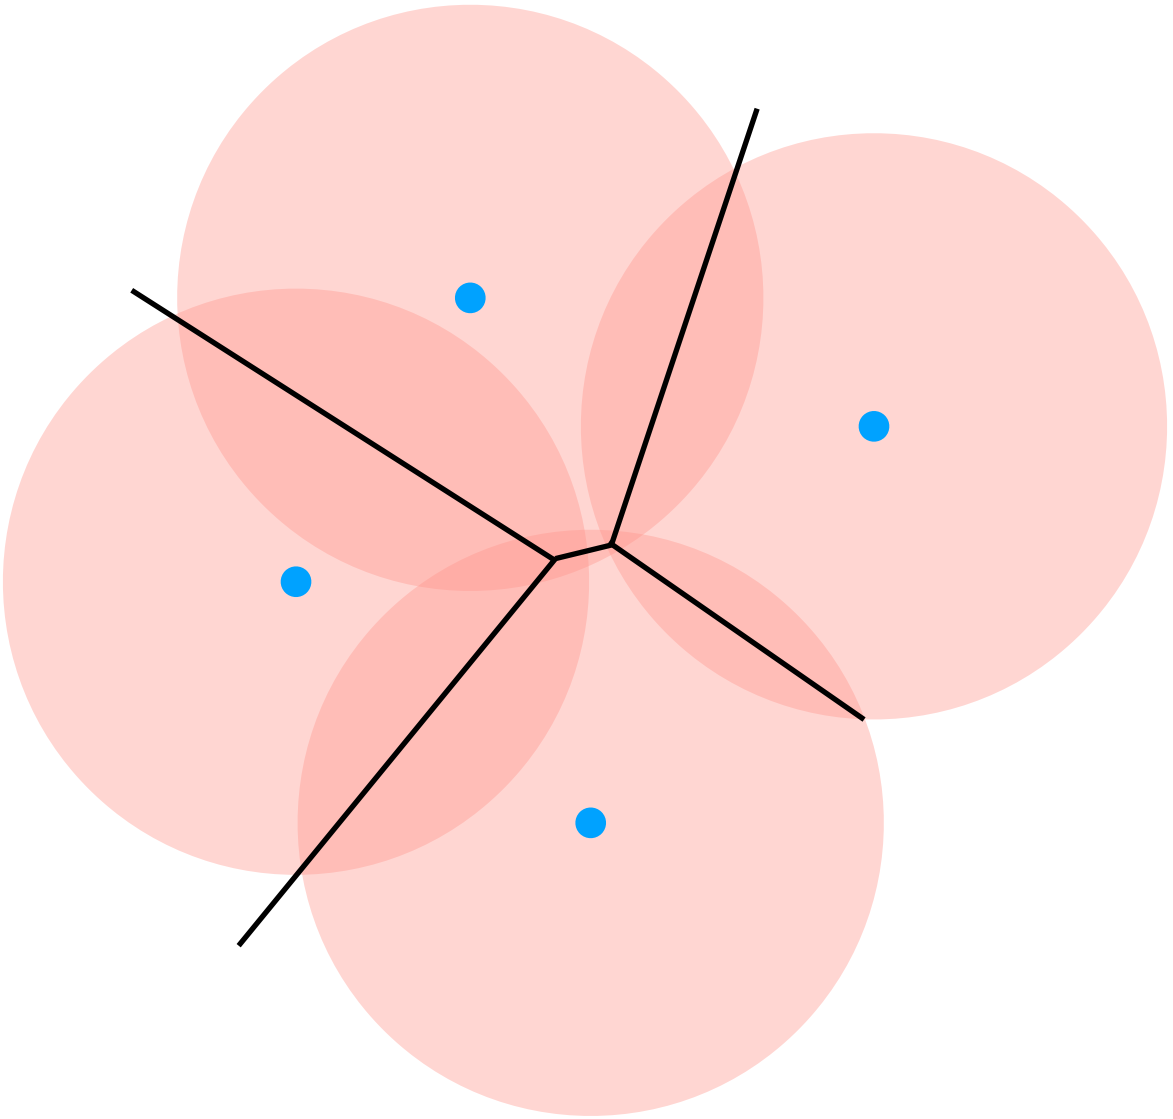

Decomposing Gaussians

- Well known: Sum of two continuous and independent gaussians is again a gaussian
- Reverse Perspective: Express a given gaussian as the sum of two independent gaussians
- For a given matrix F we want to decompose a spherical gaussian e with parameter σ into $e = e_1 F + e_2$
- e_1 is a spherical gaussian with parameter σ_1
- Such a decomposition exists if $\sigma \geq \|F\| \cdot \sigma_1$
- For a discrete gaussian $F \in \mathbb{Z}^{n \times m}$ with parameter γ, we can bound $\|F\| \leq O(\gamma\sqrt{m})$

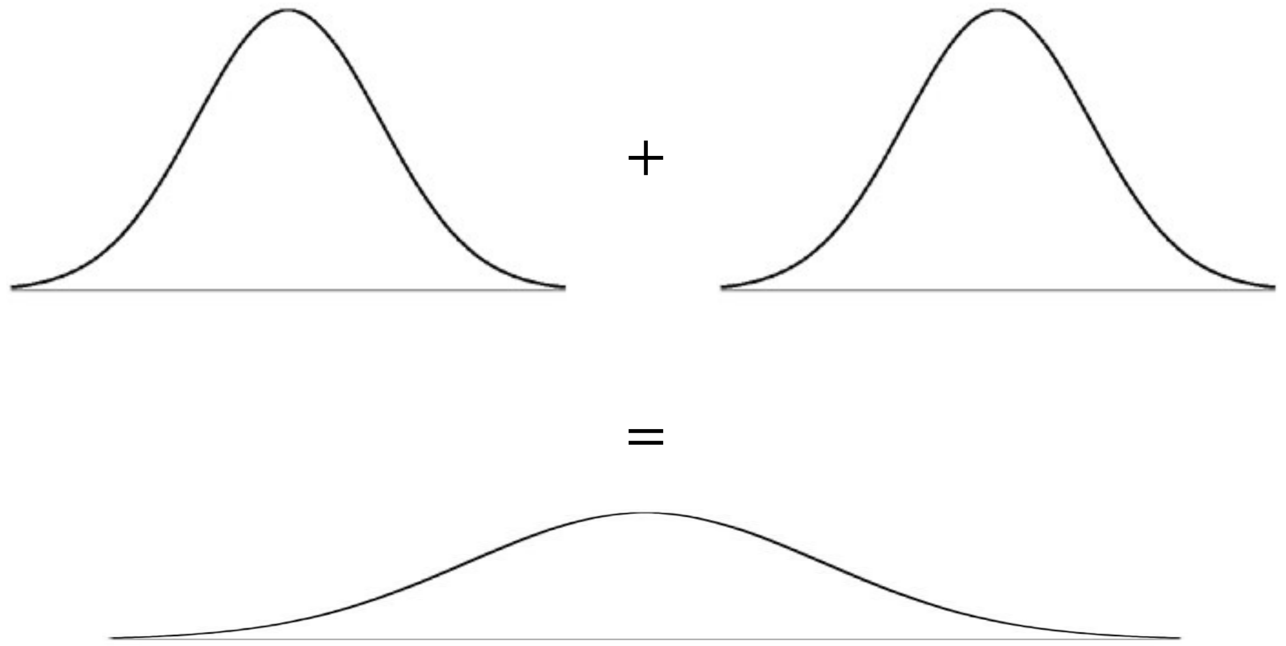

From Noise-Lossiness to Hardness of Entropic LWE

$$A, sA + e$$

$$\approx_{LWE}$$

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

$$=$$

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

$$=$$ ← $\|F\|$ **small**

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

$$=$$

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

From Noise-Lossiness to Hardness of Entropic LWE

$$A, sA + e$$

$$\approx_{LWE}$$

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

$$=$$

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

$$=$$

$$BC + F, sBC + sF + e_1 F + e_2$$

$$=$$

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

Search Version:

$$\tilde{H}_\infty(s \,|\, BC + F, sBC + (s + e_1)F + e_2)$$

$$= \tilde{H}_\infty(s \,|\, sB, s + e_1)$$

$$= \tilde{H}_\infty(s \,|\, s + e_1) - k \log(q)$$

$$= \nu_{\sigma_1}(\mathcal{S}) - k \log(q)$$

Can be improved if both s and B are short

Hard if $\nu_{\sigma_1}(\mathcal{S}) \geq k \log(q) + \omega(\log(\lambda))$

From Noise-Lossiness to Hardness of Entropic LWE

Decisional Version: Need that $\mathcal{S}$ extractable via LHL

$$A, sA + e$$

$$\approx$$

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

$$=$$

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

$$=$$

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

$$=$$

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

$$A, u$$

$$\approx_{LWE}$$

$$BC + F, u$$

$$\approx_{LWE}$$

Parameters

- We need to assume LWE with parameter σ
- We get hardness of entropic LWE with parameter $\sigma_1 \cdot \sigma \cdot \sqrt{m}$
- I.e. Modulus-to-noise ratio deteriorates by a factor $\sigma_1 \cdot \sqrt{m}$

Computing the Noise Lossiness

Noise Lossiness: General Distributions

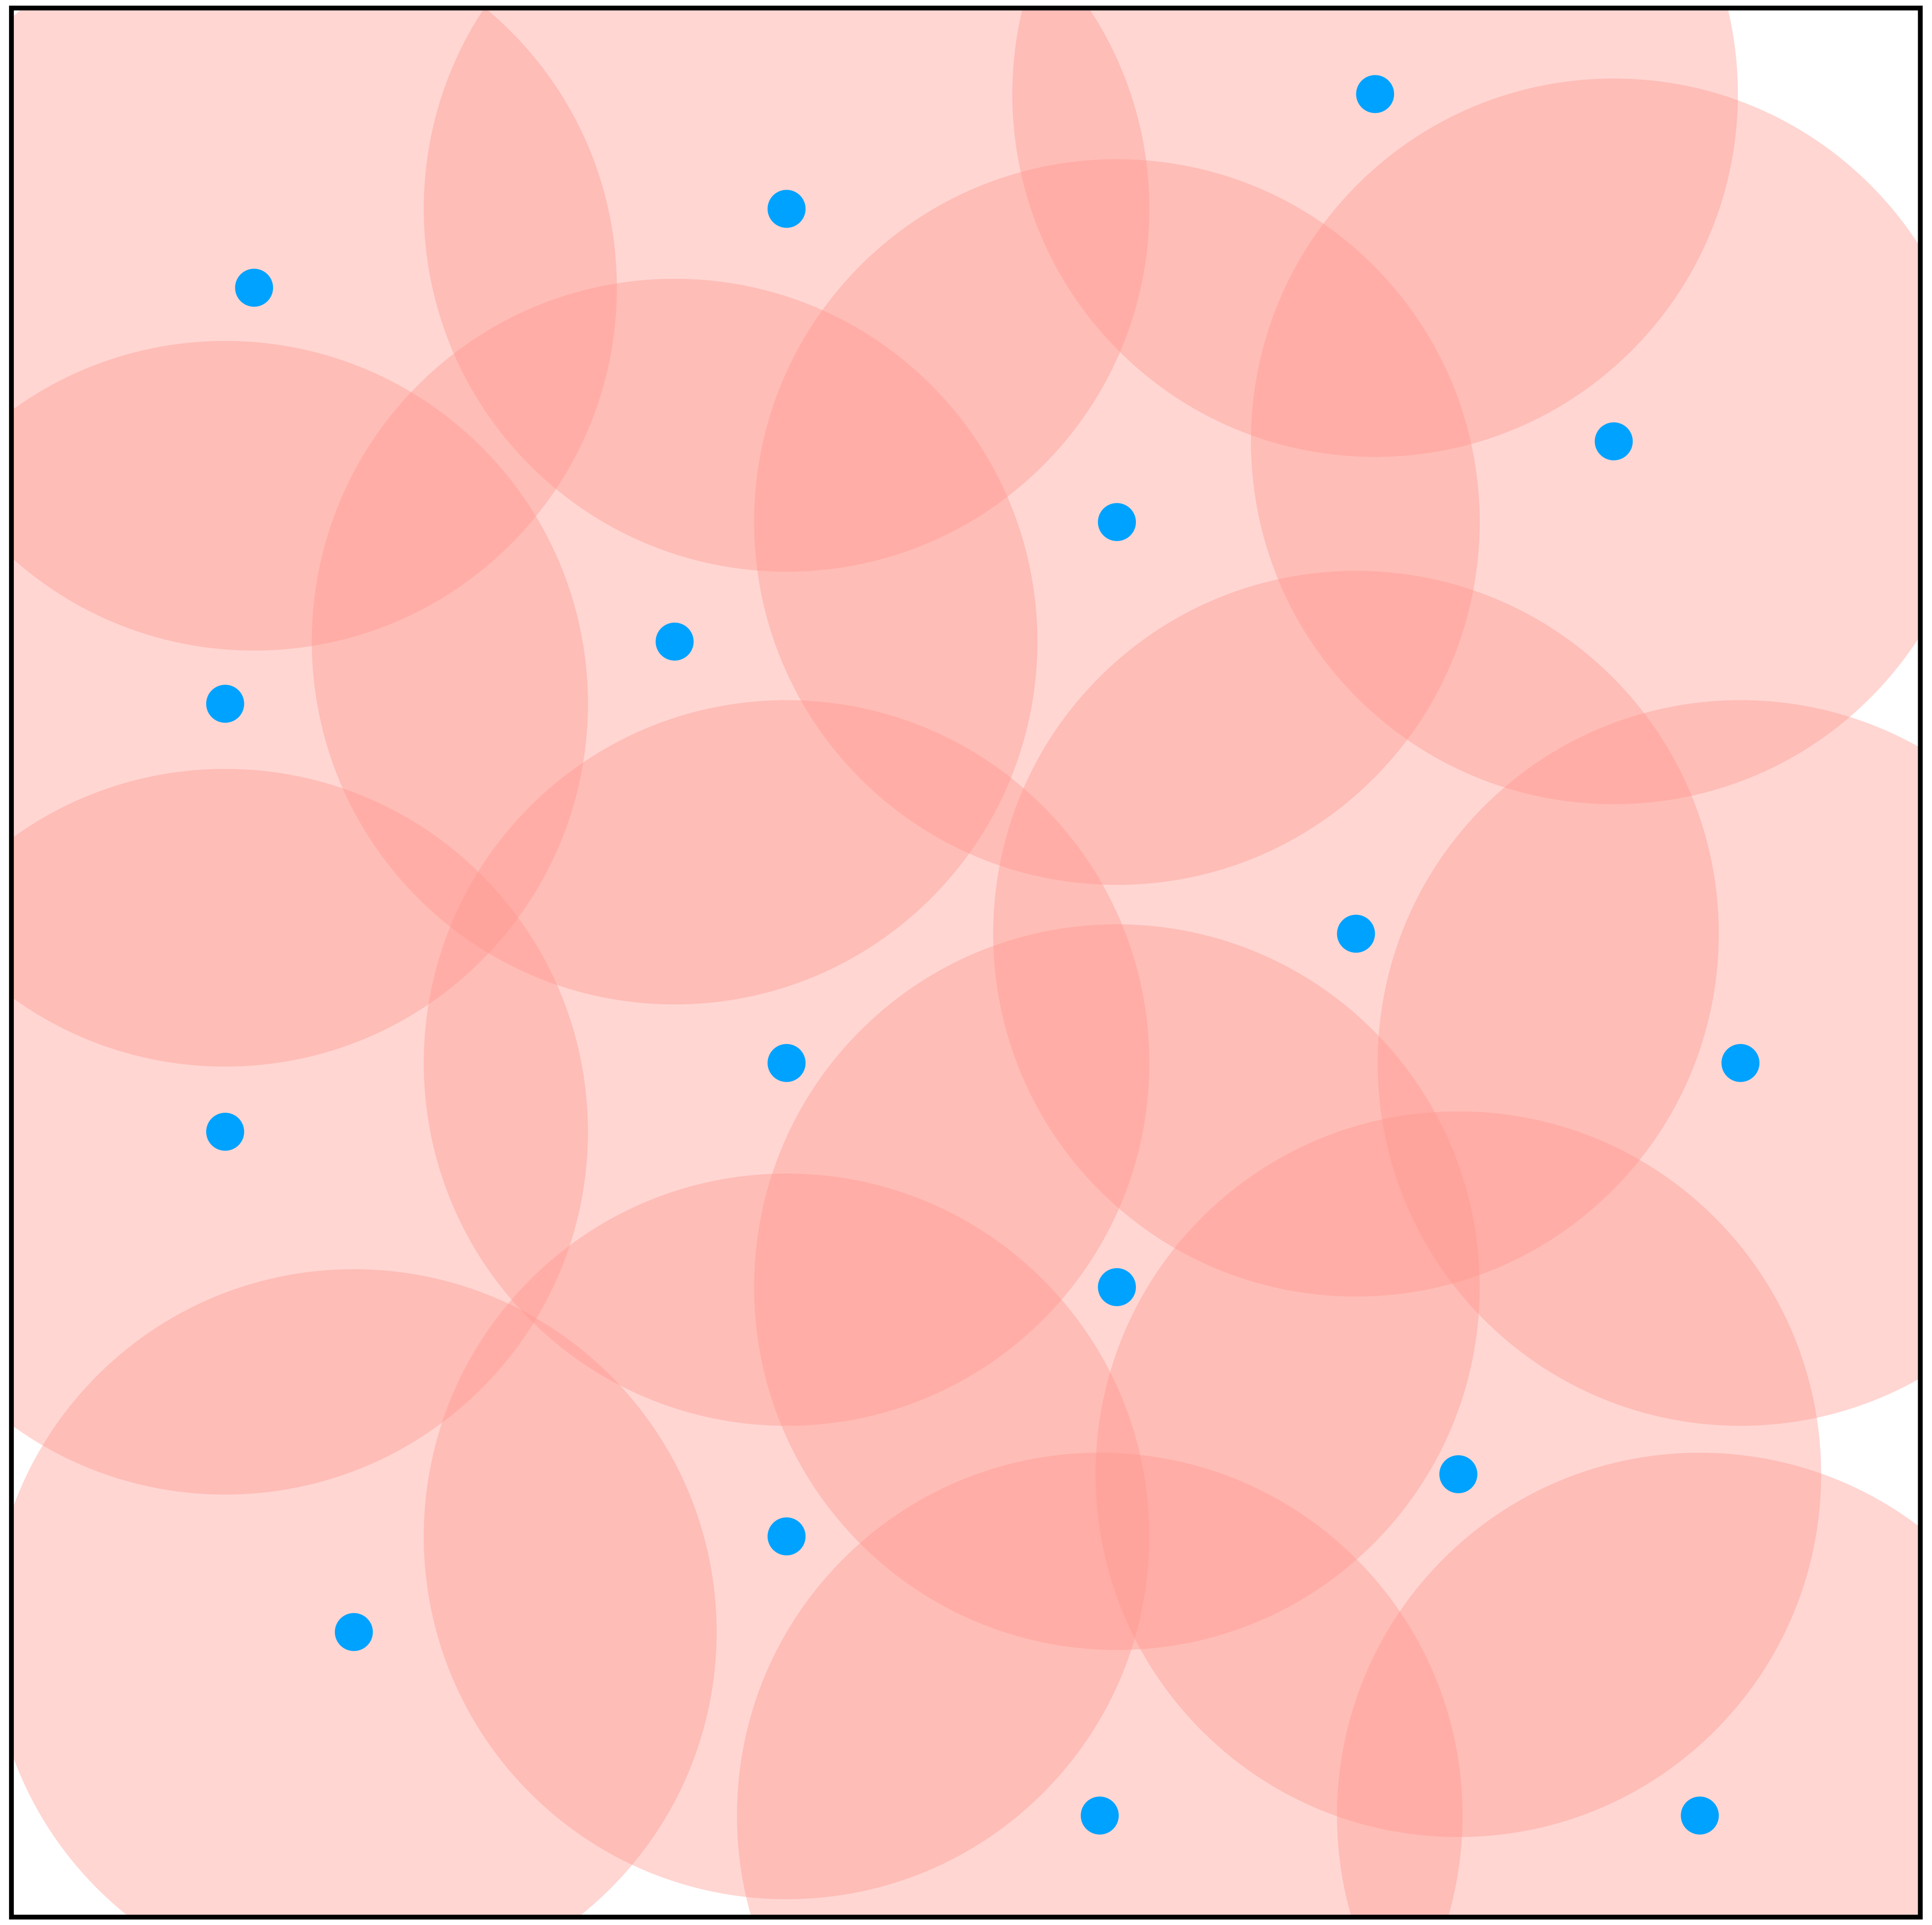

$\mathbb{Z}_q^n$

$$\nu_\sigma(\mathcal{S}) \geq H_\infty(s) - n \cdot \log(q/\sigma) - 1$$

Noise Lossiness: Short Distributions

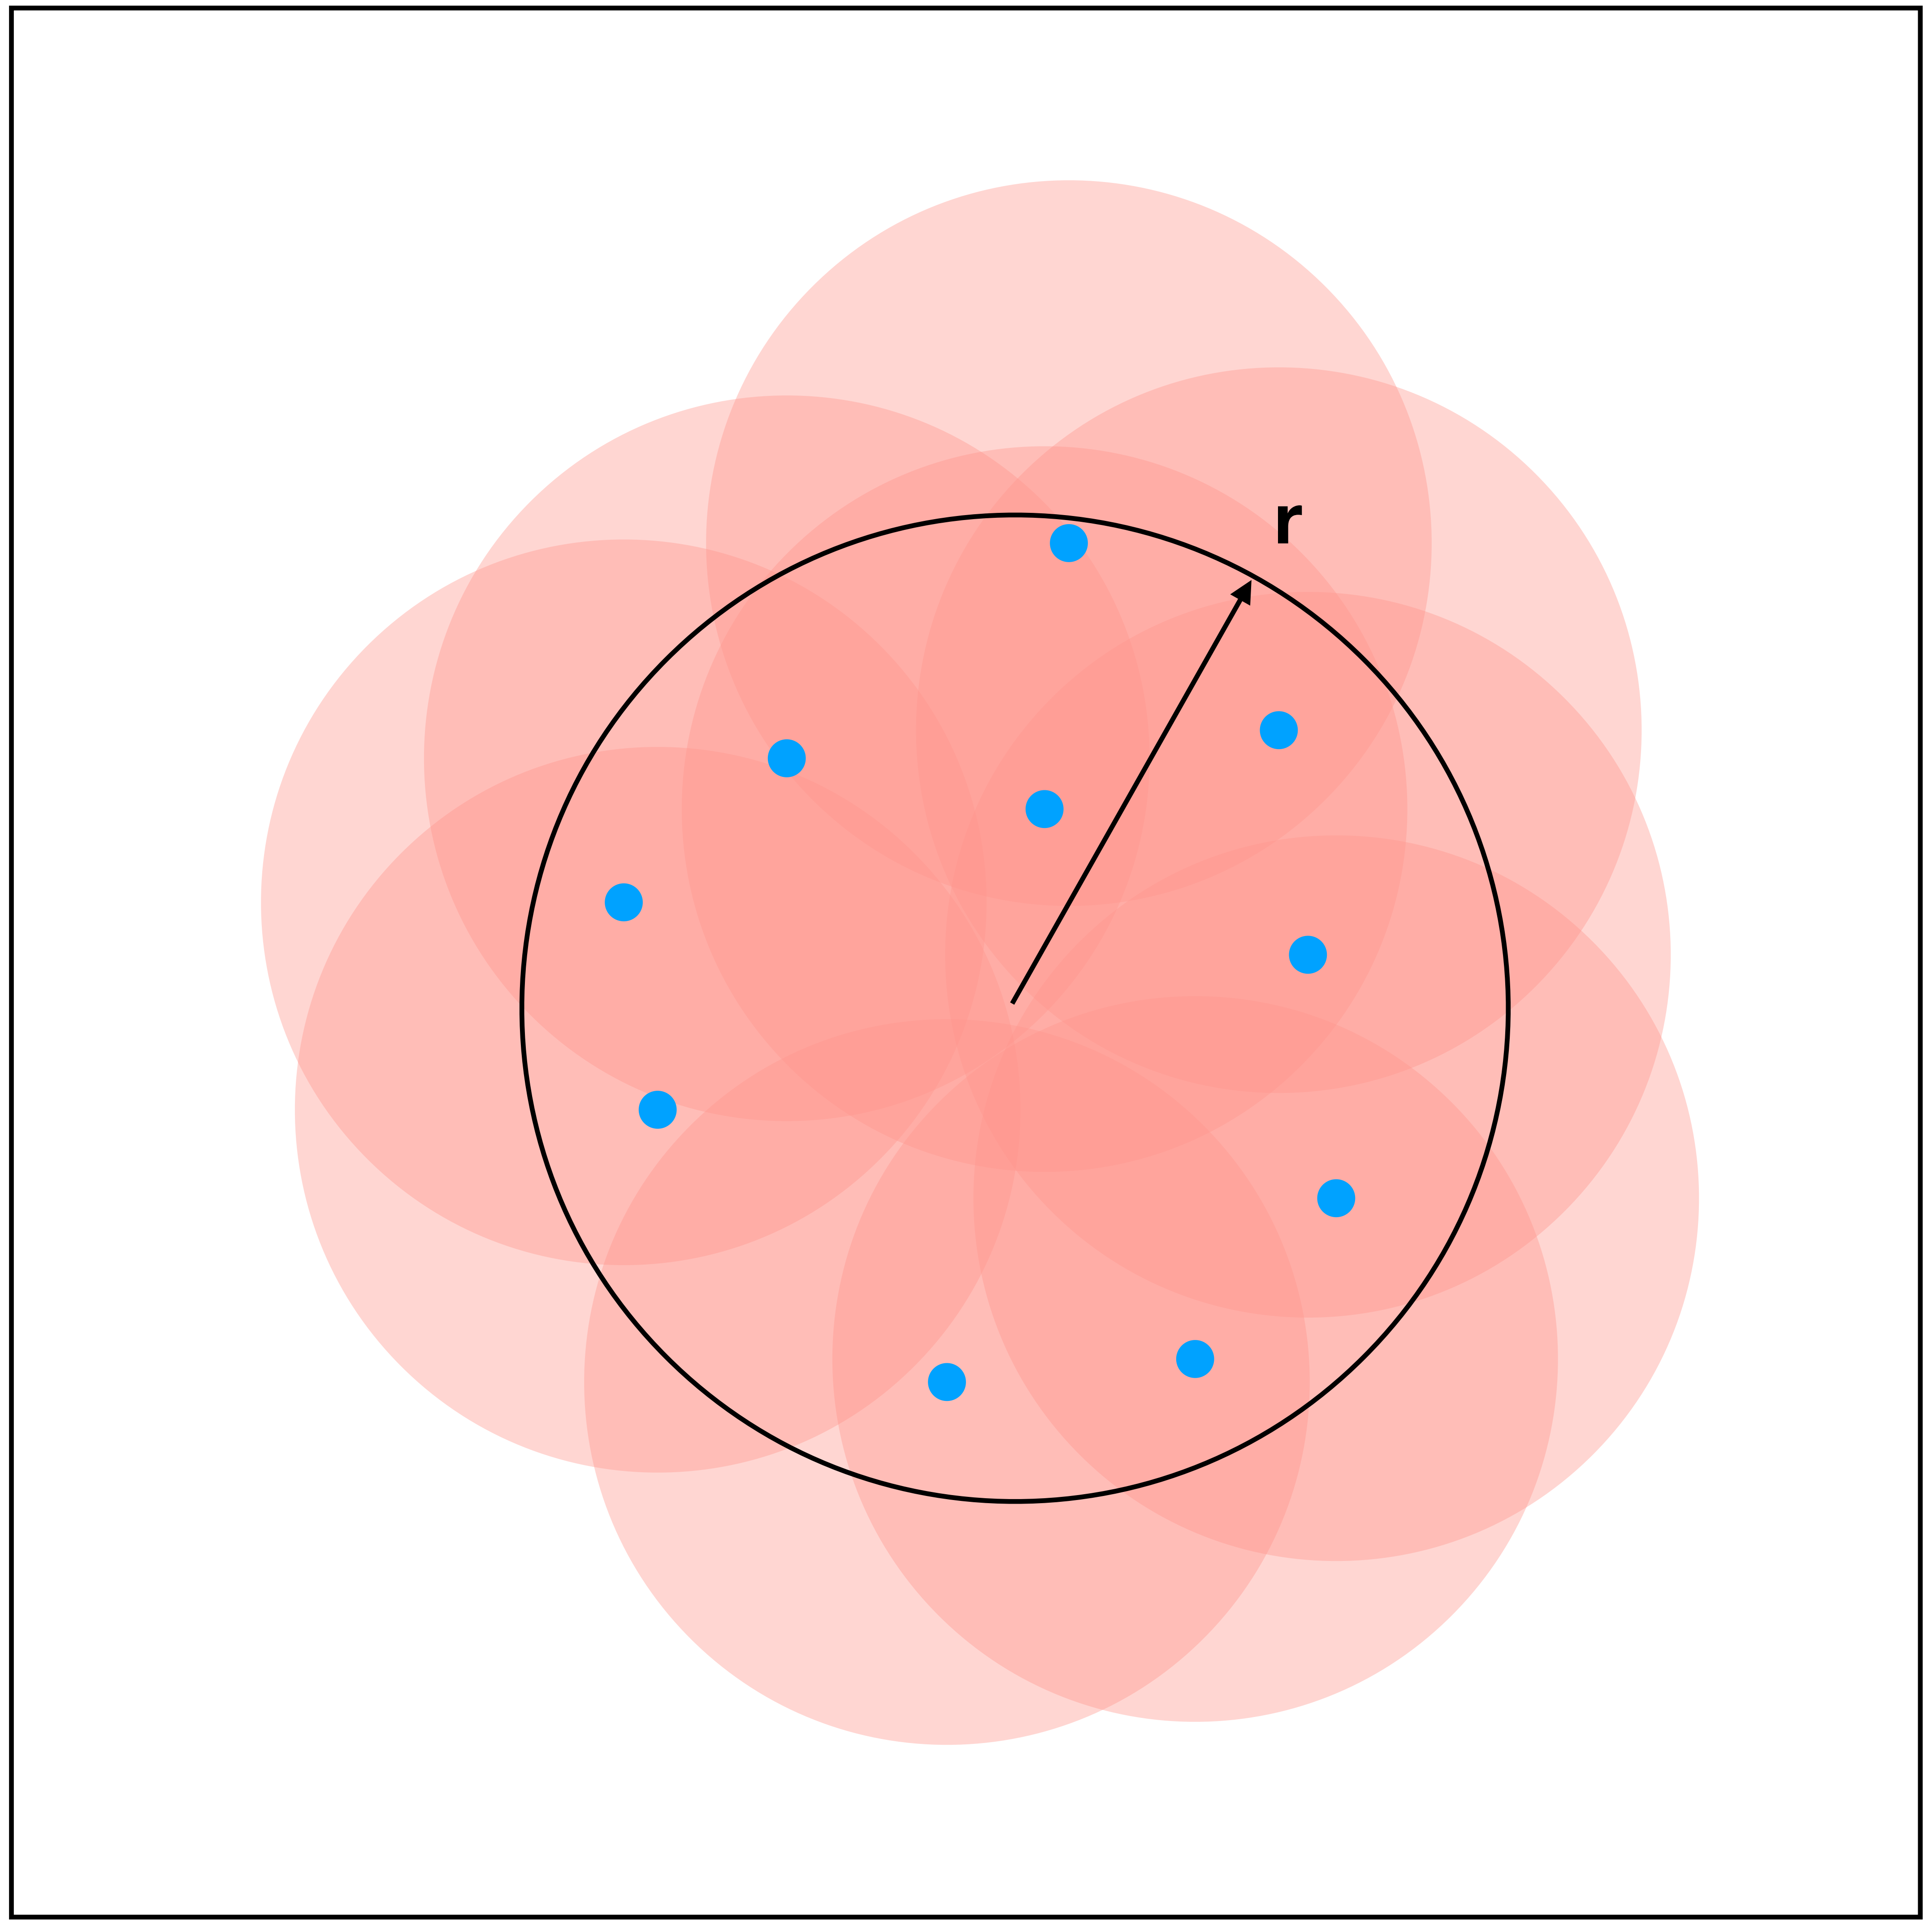

$$\mathbb{Z}_q^n$$

$$\nu_\sigma(\mathcal{S}) \geq H_\infty(s) - 2r\sqrt{n}/\sigma$$

Main Result

- Putting everything together, assuming $LWE(k, q, \gamma)$ is hard:
- For general (non-short) min-entropy distributions $\mathcal{S}$ we get that $eLWE(\mathcal{S}, n, q, m, \sigma)$ is hard given that $H_\infty(s) \gtrsim k \cdot \log(q) + n \cdot \log(q\gamma\sqrt{m}/\sigma)$
- For r-bounded distributions $\mathcal{S}$ we need $H_\infty(s) \gtrsim k \log(\gamma r) + 2r\sqrt{nm}\gamma/\sigma$

Lower Bounds

- For the general case, min-entropy of $\mathcal{S}$ must close to $n\log(q)$ or σ of the same order as q
- Can we do better for general entropic distributions?
- Specific Moduli: **No**!

Counterexample

$q = p \cdot q'$

Let $\mathcal{S}$ be the uniform distribution on $p \cdot \mathbb{Z}_q^n$

sA is supported on $p \cdot \mathbb{Z}_q^m$

$\|e\|_\infty < p/2$

$\Rightarrow sA + e \mod p = e$

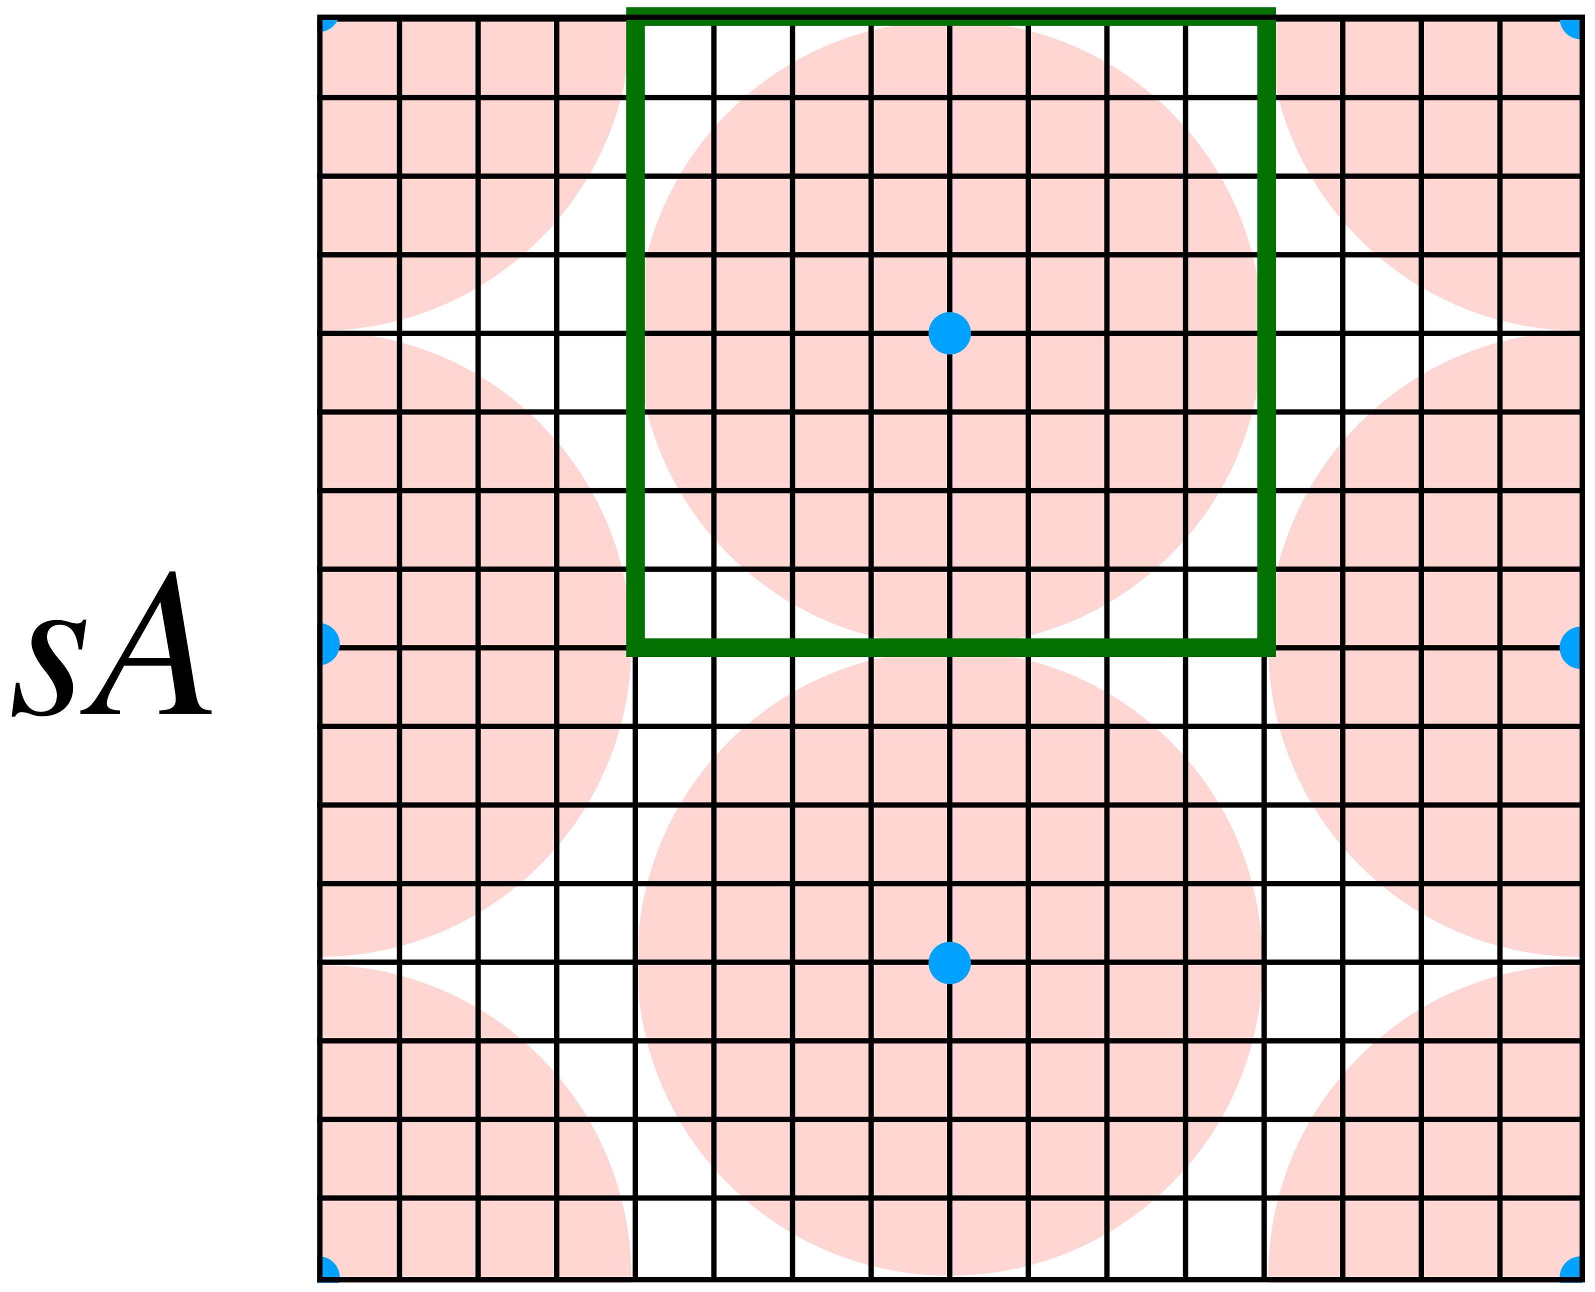

Lower Bounds

- What if $\mathbb{Z}_q$ does not have a sub-structure?
- Meta-Reduction Framework: Show that BB-reduction can be used to break the underlying assumption without using an adversary
- Simulatable Adversaries [Wichs13]: From the view of a BB-reduction, an unbounded adversary can be simulated efficiently
- **Main Idea:** Simulator knows all the samples that were given to the adversary

BB-Lower Bound

Unbounded Adversary

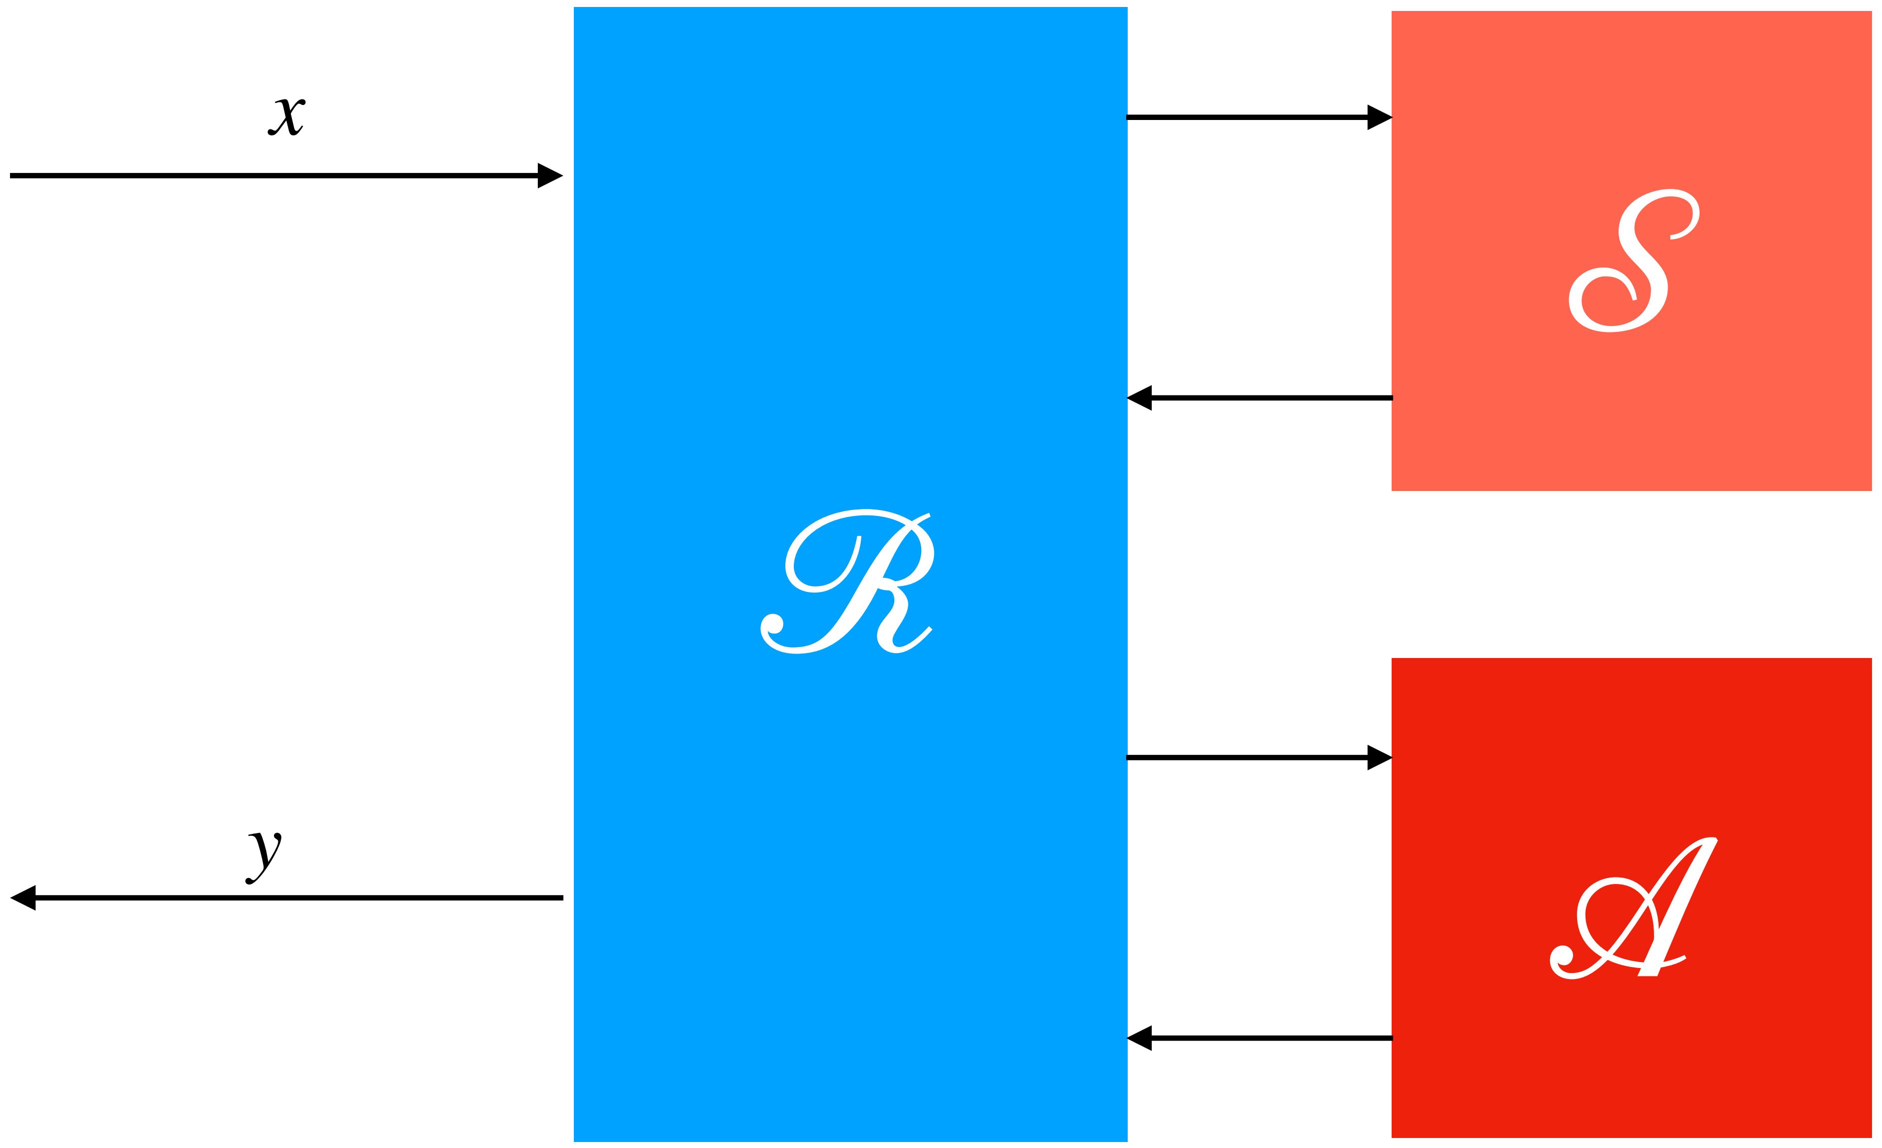

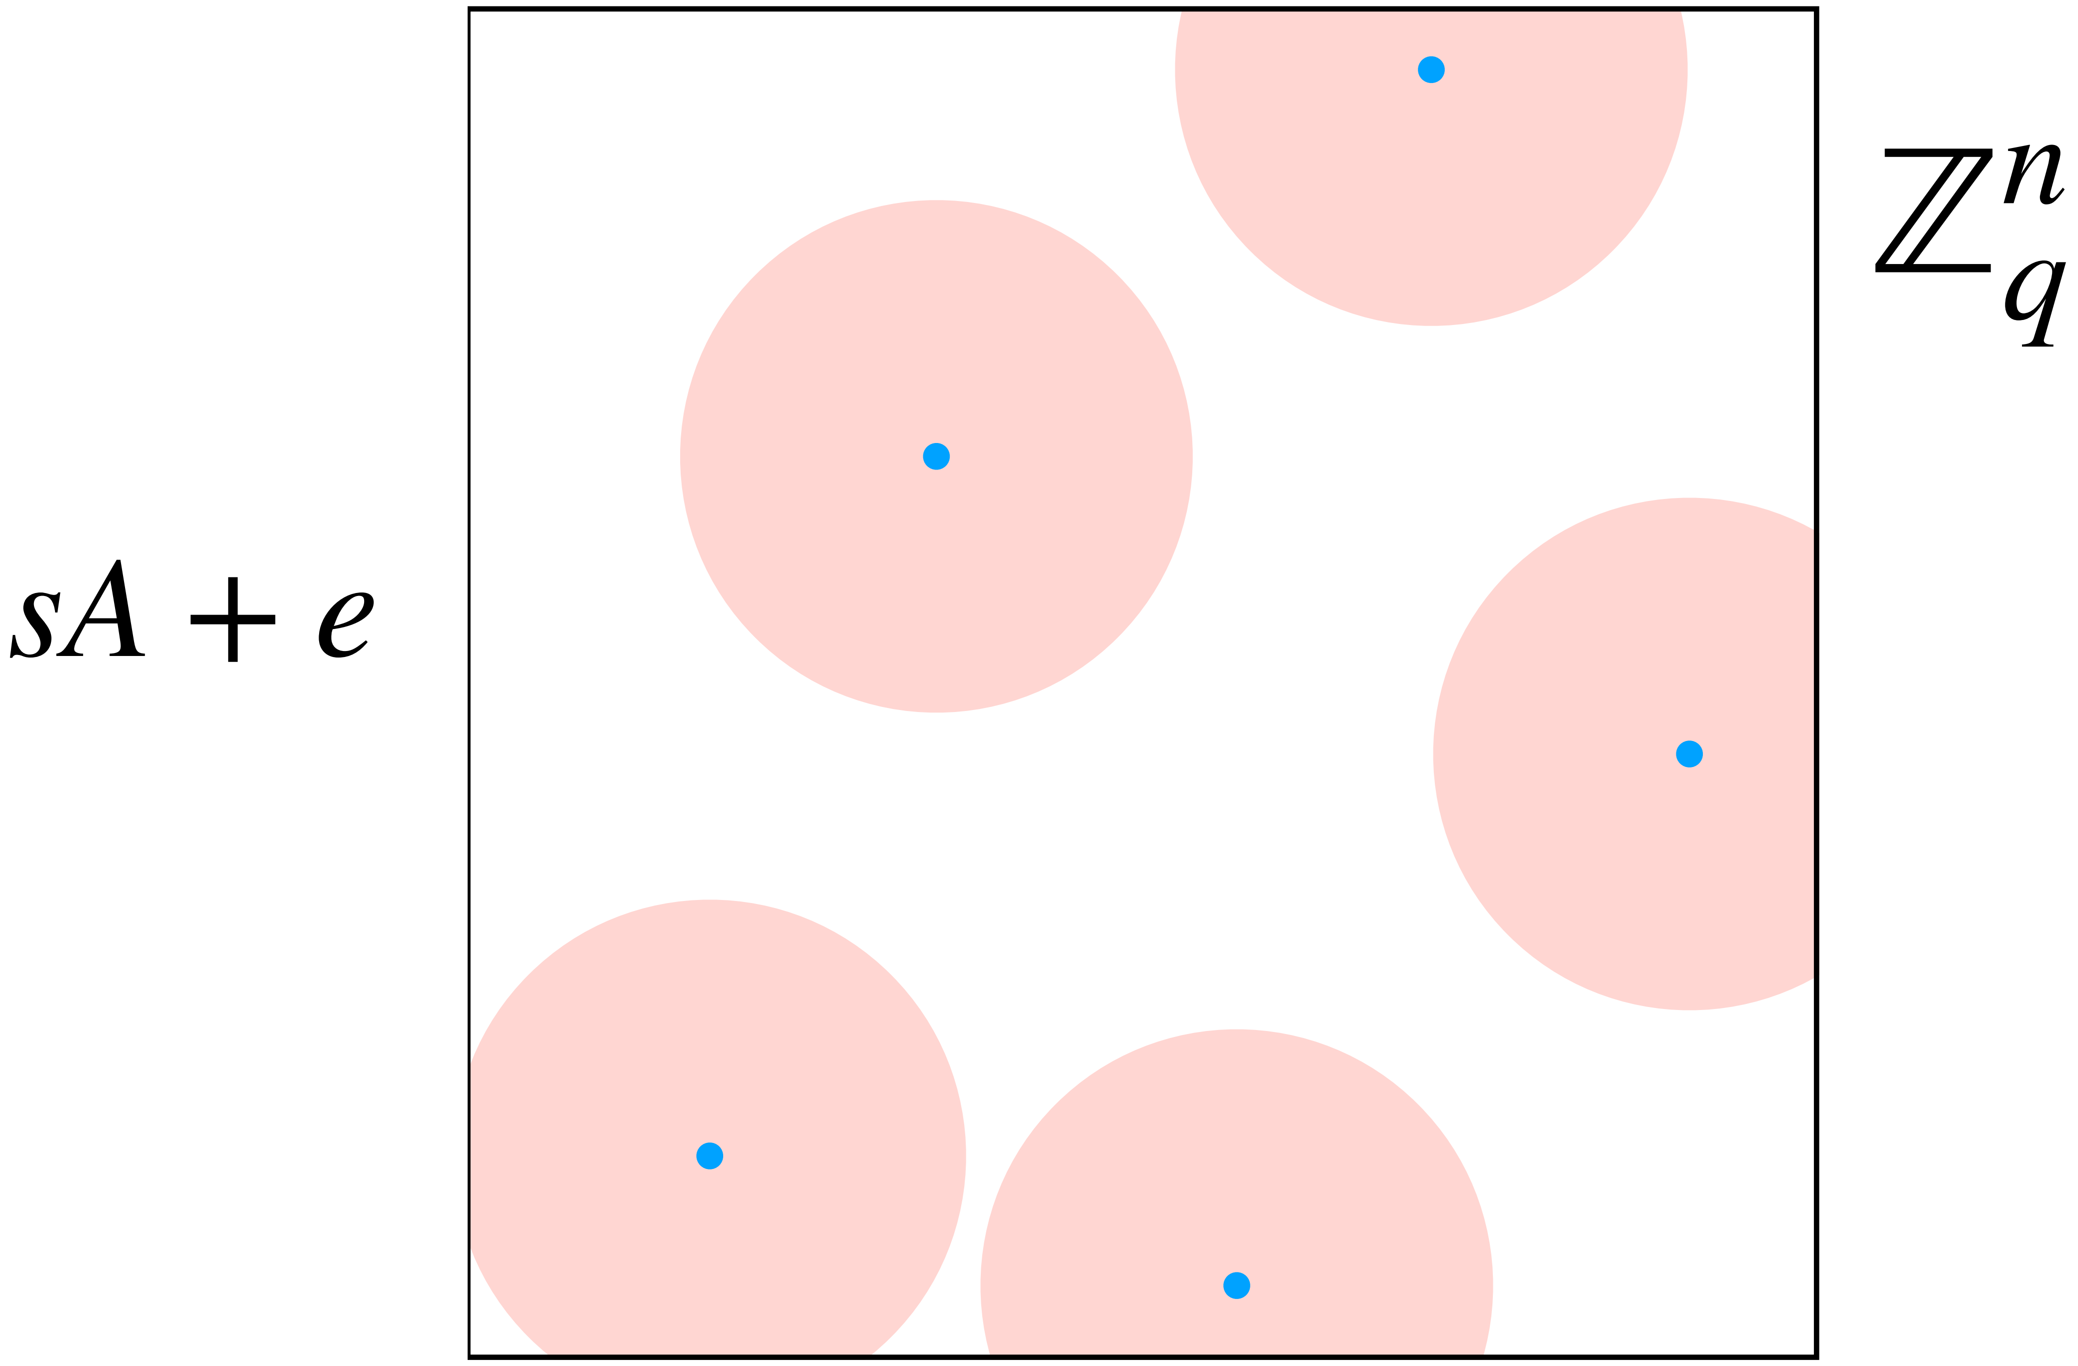

Support of $\mathscr{S}$ is chosen uniformly random of size 2^k where $k \lesssim n \log(q/B)$

BB-Lower Bound

Efficient Simulator

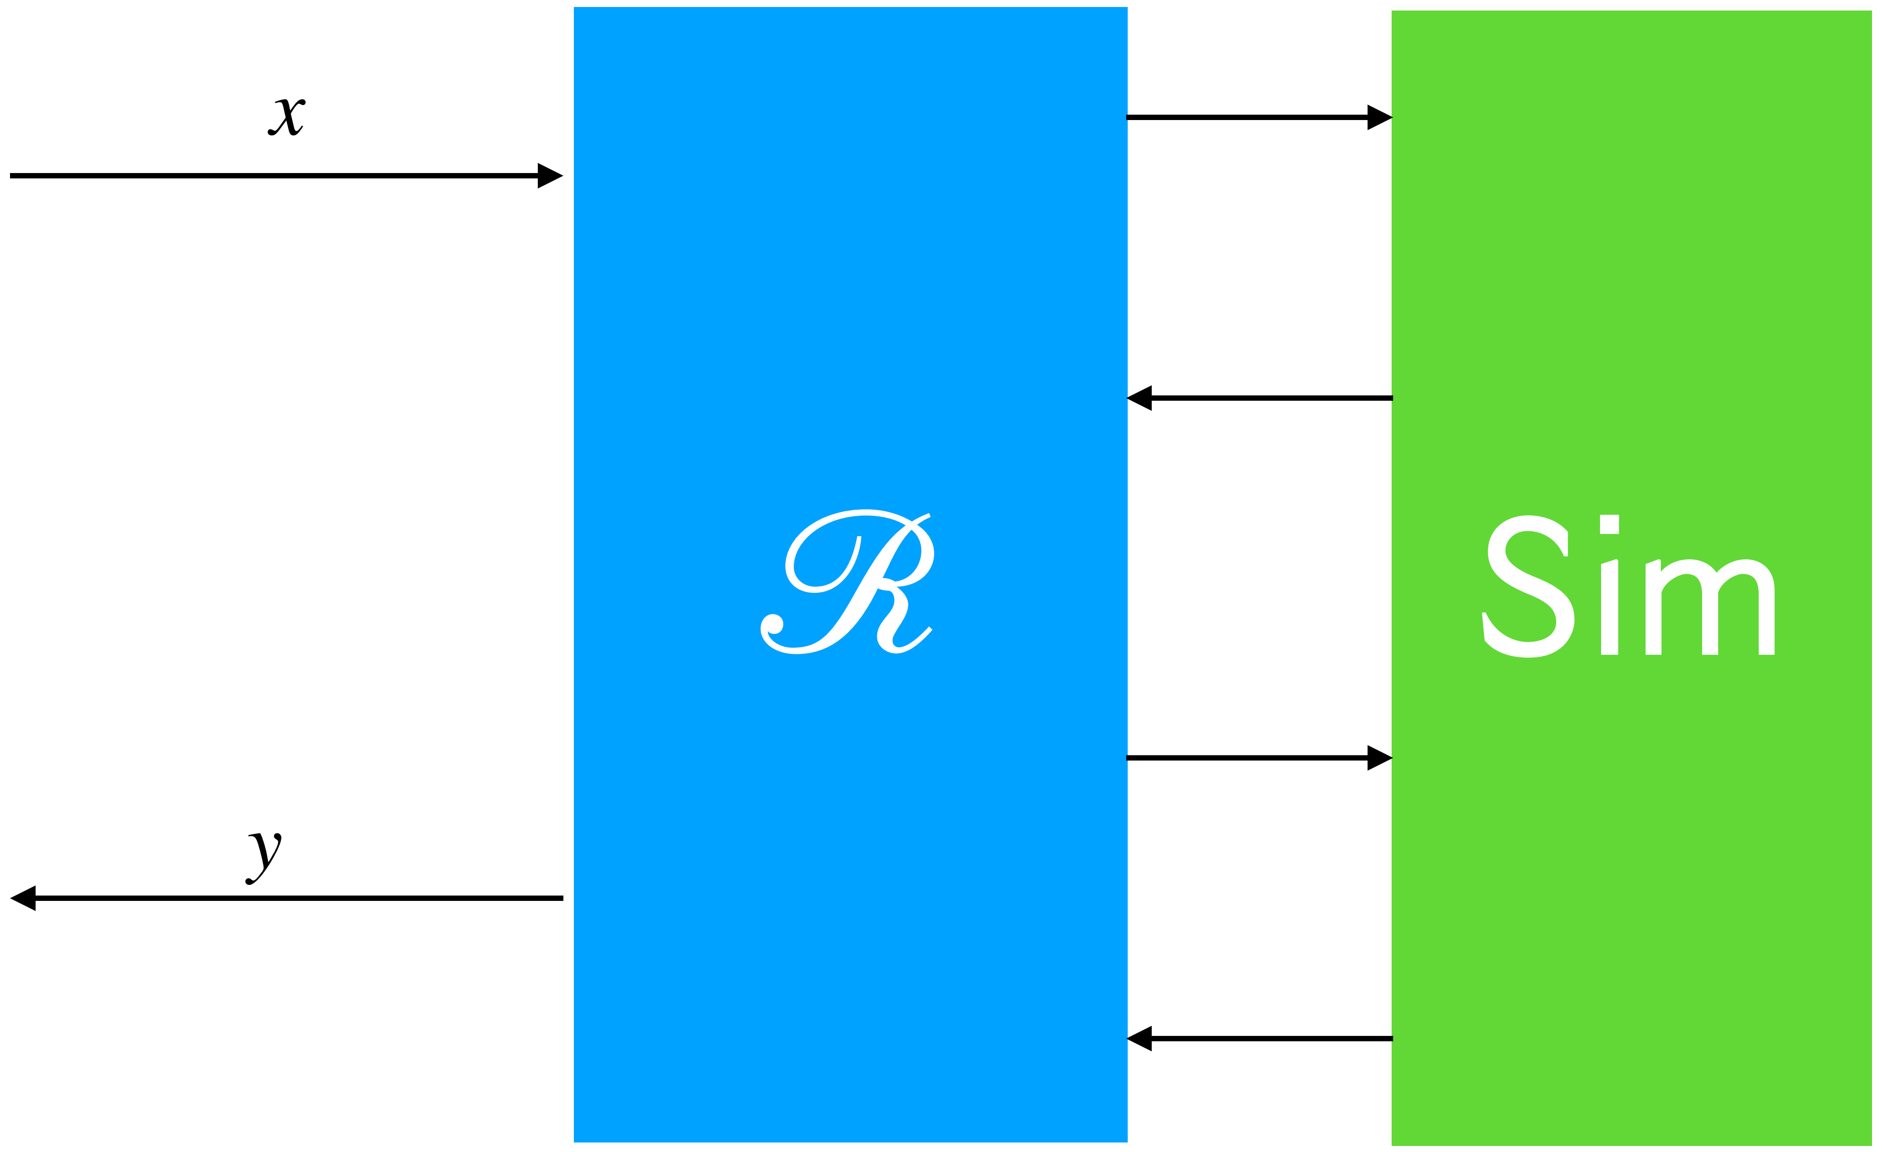

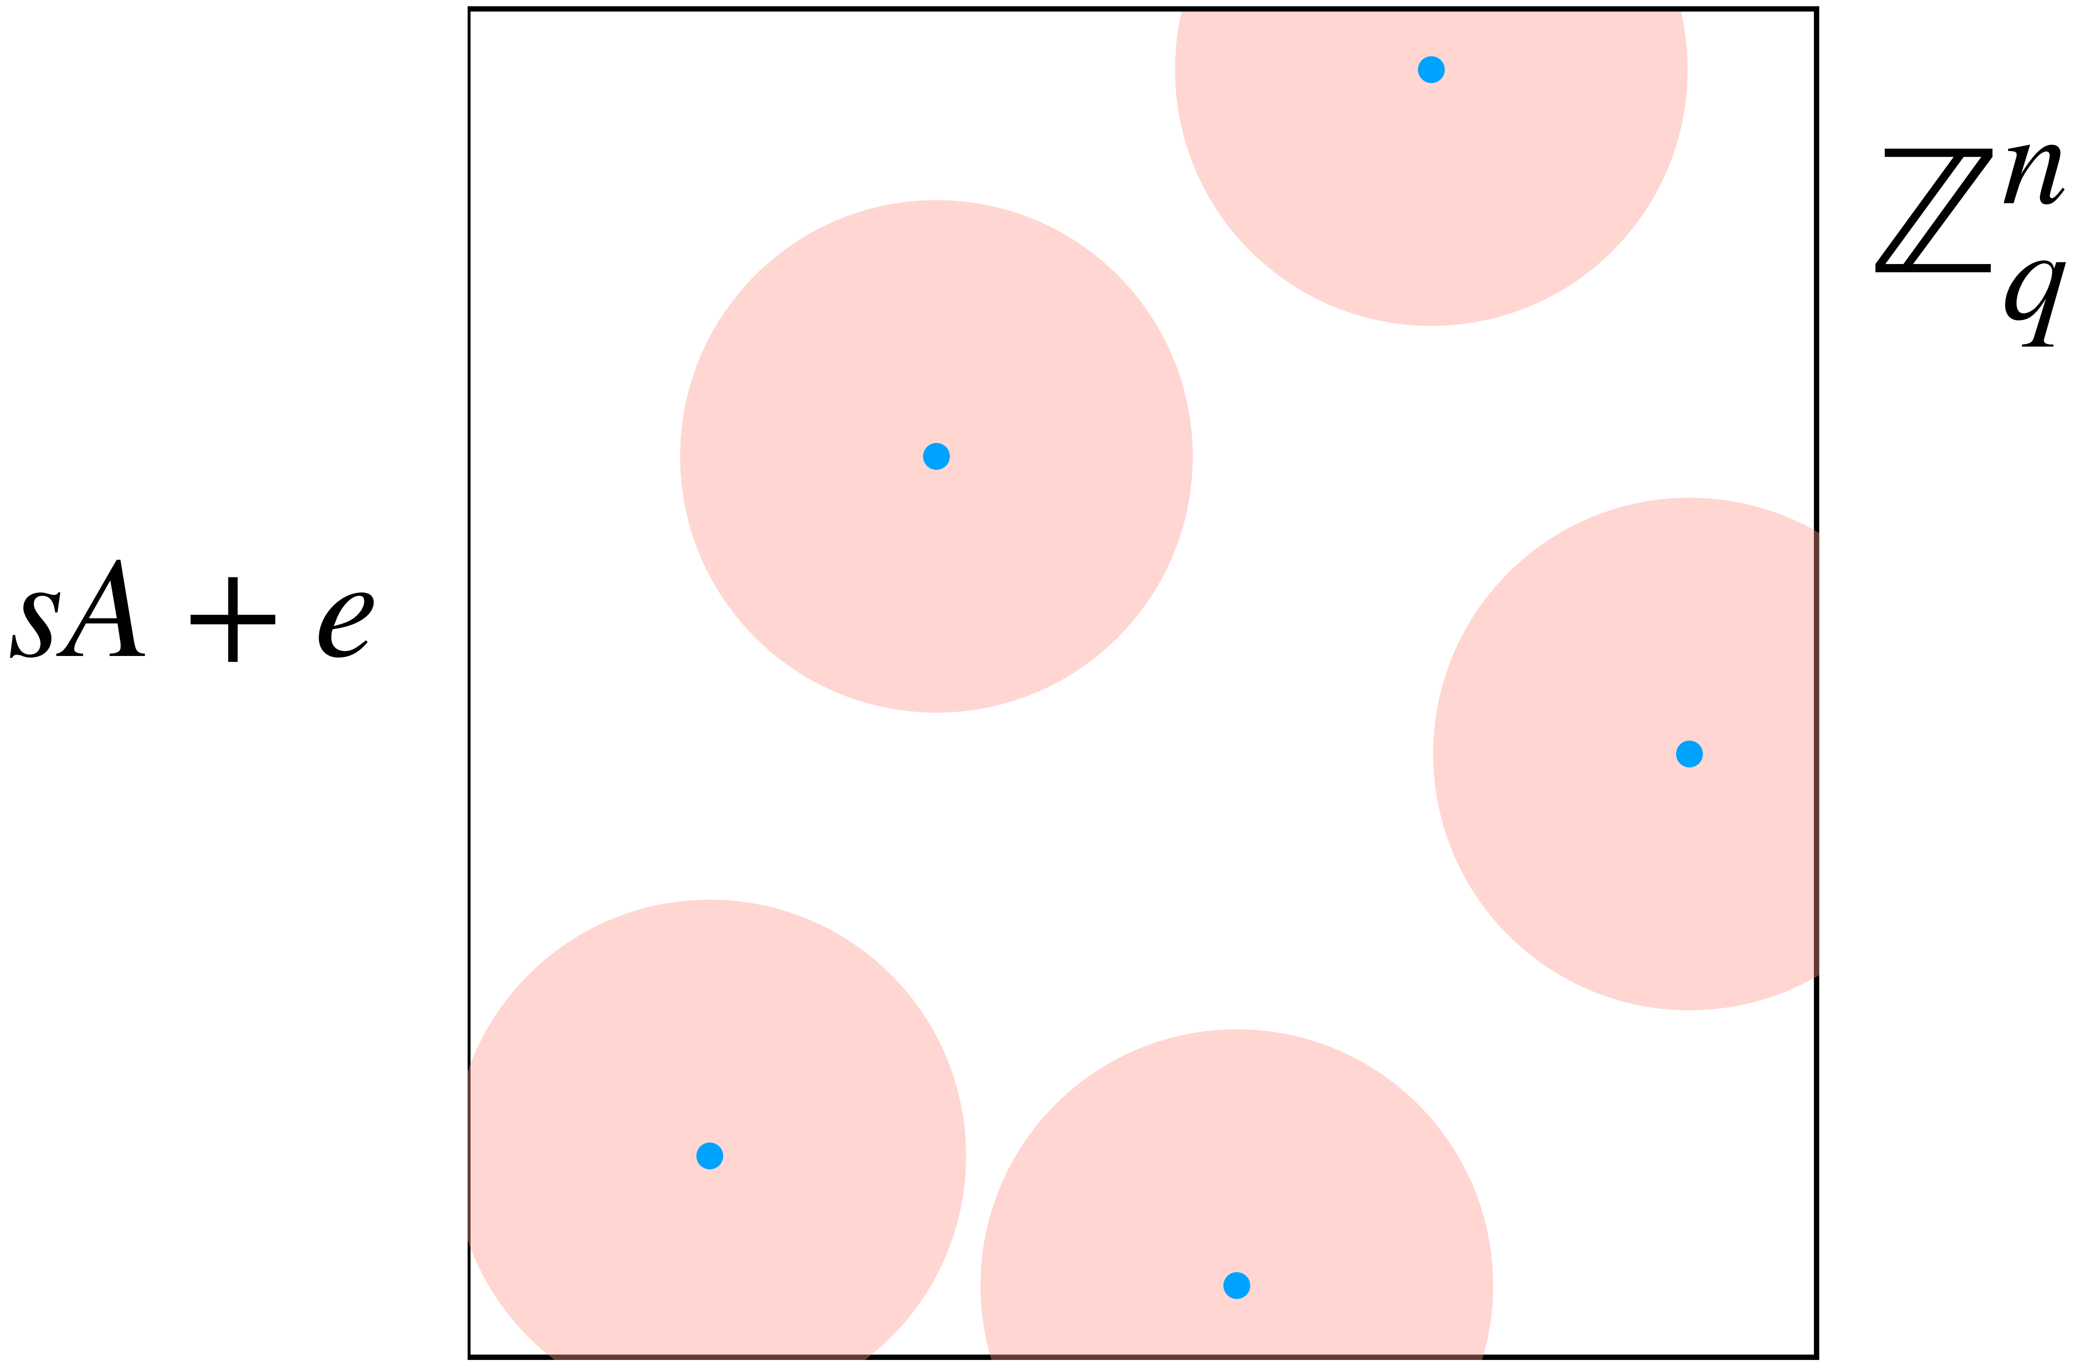

Support of $\mathscr{S}$ is chosen uniformly random of size 2^k where $k \lesssim n \log(q/B)$

Take Away and Open Problems

Conclusions

- Standard LWE (non-short secrets) can tolerate a small amount of leakage,
- This has inherent reasons, either attacks or BB-impossibility
- LWE with short/binary secret tolerates a much higher leakage rate, but in general this comes at the cost of large public keys (factor $\approx \log(q)$)

Open Problems

- What about more specific classes of distributions/leakage functions?
- Leakage that includes the noise?
- Techniques do translate to Learning-with-Rounding, but not “nicely”
- Does the BB-impossibility extend e.g. to quantum reductions?
- Structured LWE, e.g. Ring-LWE?

Thanks!