New Techniques for Preimage Sampling:
NIZKs and More from LWE

Brent Waters, Hoeteck Wee, and David Wu

The Preimage Sampling Problem

Given $\boldsymbol{A} \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t} \in \mathbb{Z}_q^n$

find short $\boldsymbol{\pi} \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A\pi} = \boldsymbol{t}$

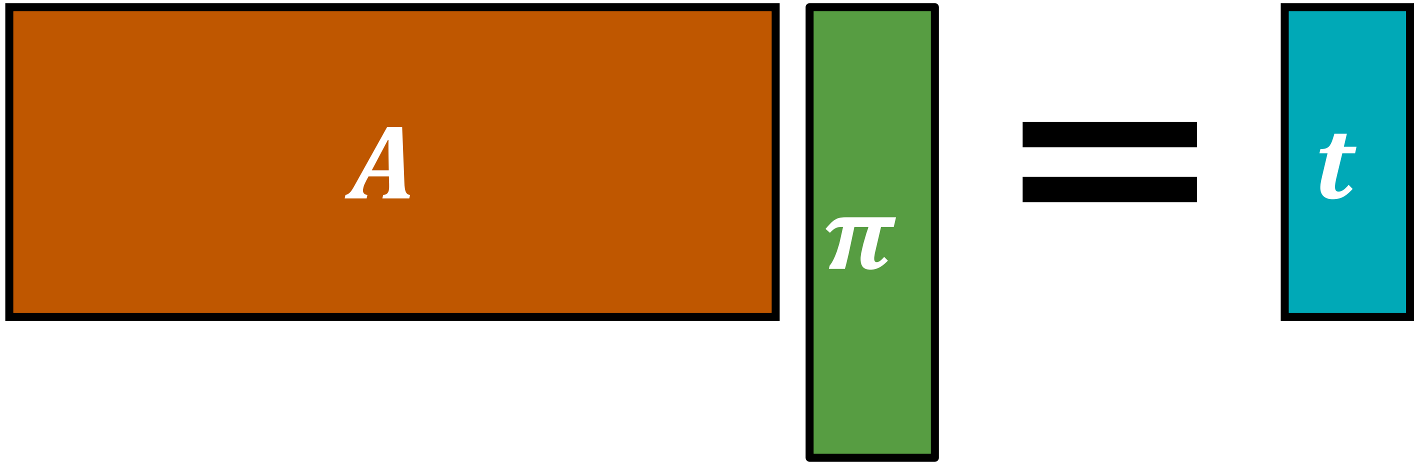

Problem is hard in general:

- Short integer solutions (SIS)
- Inhomogeneous SIS

But easy given a trapdoor for $\boldsymbol{A}$

[Ajt96, GPV08, MP12]

Many applications!

digital signatures, IBE, ABE, SNARGs, NIZKs

Shifted Multi-Preimage Sampling

Given $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_1, \ldots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$

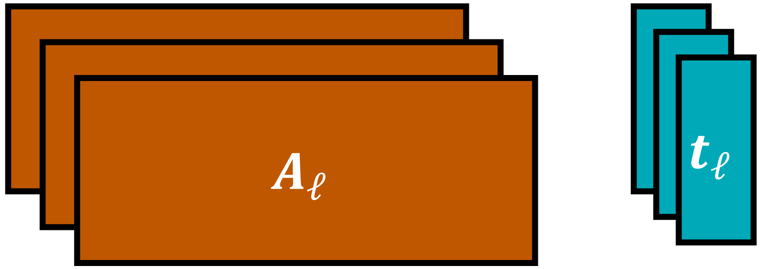

find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

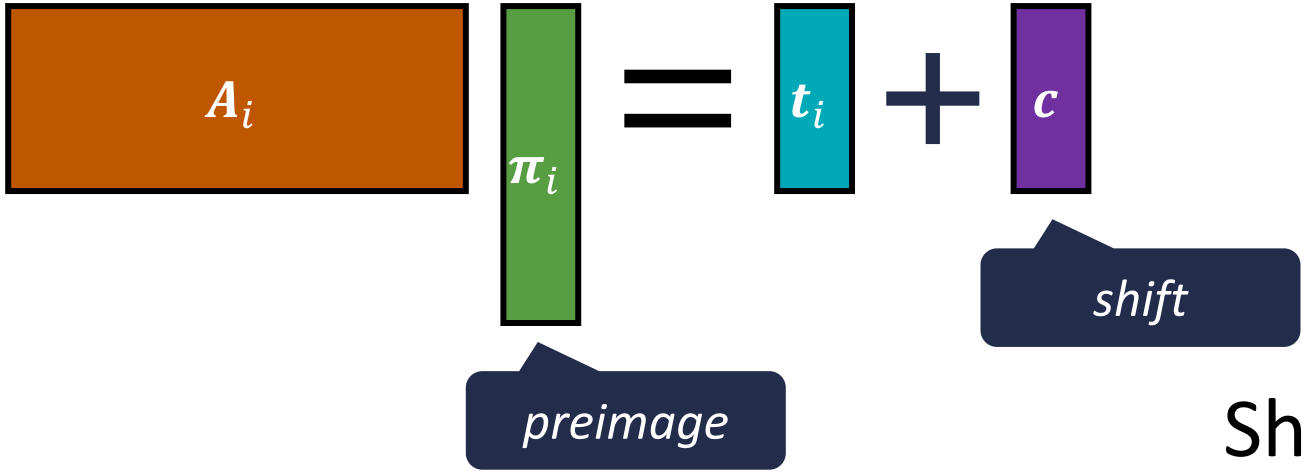

Shift gives **one** degree of freedom

Shifted Multi-Preimage Sampling

Given $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$,
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

Problem is implicitly considered in several recent lattice-based constructions:

- Vector commitments [PPS21, WW23]
- Dual-mode NIZKs via the hidden-bits model [Wat24]

Solving this problem typically requires a hint (i.e., trapdoor information) related to $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$

Trivial solution: $\mathrm{hint} = (\mathrm{td}_1, \dots, \mathrm{td}_\ell)$ where td_i is trapdoor for $\boldsymbol{A}_i$

When $\ell = 1$, problem is easy (*without* hints):
sample (arbitrary) $\boldsymbol{\pi}_1$ and set $\boldsymbol{c} = \boldsymbol{A}_1 \boldsymbol{\pi}_1 - \boldsymbol{t}_1$

Problem is also easy for some special choices of $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$ (e.g., $\boldsymbol{A}_1 = \boldsymbol{A}_2 = \cdots = \boldsymbol{A}_\ell = \boldsymbol{G}$)

Shifted Multi-Preimage Sampling

Given $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$,
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

Problem is implicitly considered in several recent lattice-based constructions:

- Vector commitments [PPS21, WW23]
- Dual-mode NIZKs via the hidden-bits model [Wat24]

Solving this problem typically requires a hint (i.e., trapdoor information) related to $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$

Trivial solution: $\text{hint} = (\text{td}_1, \dots, \text{td}_\ell)$ where td_i is trapdoor for $\boldsymbol{A}_i$

Above applications require that SIS/LWE remains hard with respect to any $\boldsymbol{A}_i$ even given the hint (rules out trivial solution)

Feasible only if we allow for the shift

This Work

Given $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n\times m}$ *and* $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$*,*
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i\boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

New approach to sample $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$ together with a trapdoor td where:

- td can be used to sample (Gaussian-distributed) solutions the shifted multi-preimage sampling problem with respect to $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$ and arbitrary targets $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell$

In fact, td can be used to sample solutions that are statistically close to the following distribution:

- $\boldsymbol{c} \leftarrow \mathbb{Z}_q^n$
- $\boldsymbol{\pi}_i \leftarrow \boldsymbol{A}_i^{-1}(\boldsymbol{t}_i + \boldsymbol{c})$; $\boldsymbol{\pi}_i$ is a discrete Gaussian vector satisfying $\boldsymbol{A}_i\boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$

This Work

Given $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$*,*
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

New approach to sample $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$ together with a trapdoor td where:

- td can be used to sample (Gaussian-distributed) solutions the shifted multi-preimage sampling problem with respect to $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$ and arbitrary targets $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell$
- $(\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell, \mathrm{td})$ can be *publicly* derived from a uniform random matrix $\boldsymbol{B} \leftarrow \mathbb{Z}_q^{n \times m \lceil \log \ell \rceil}$
- SIS/LWE problems are hard with respect to any $\boldsymbol{A}_i$ given $\boldsymbol{B}$

Note: $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$ are also elements of $\mathbb{Z}_q^{n \times m \lceil \log \ell \rceil}$ (slightly wider by a $\log \ell$ factor)

This Work

Given $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$,
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

New approach to sample $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$ together with a trapdoor td where:

- td can be used to sample (Gaussian-distributed) solutions the shifted multi-preimage sampling problem with respect to $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$ and arbitrary targets $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell$
- $(\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell, \mathrm{td})$ c
- SIS/LWE problem

Previously lattice-based schemes: either has long *structured* CRS [WW23] or not statistically hiding [dCP23]

Applications:

- Statistically-hiding vector commitments from SIS with $\mathrm{poly}(\lambda, \log \ell)$-size public parameters, commitments, and openings (and transparent setup)

This Work

Given $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$,
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

New approach to sample $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$ together with a trapdoor td where:

- td can be used to sample (Gaussian-distributed) solutions the shifted multi-preimage sampling problem with respect to $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$ and arbitrary targets $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell$
- $(\boldsymbol{A}_1$
- SIS,

Applicati

- Sta
  commitments, and opening nsparent setup)
- Dual-mode NIZK from LWE with polynomial modulus and a transparent setup in statistical ZK mode (and CRS size linear in the length of the hidden-bits string)

Previous construction [Wat24]: structured CRS in both modes, required sub-exponential modulus, and CRS size is quadratic in the length of the hidden-bit string

Our NIZK essentially achieves the same set of properties as those obtained via the correlation-intractability framework

This Work

Given $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_1, \ldots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$*,*
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

New approach to sample $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell$ together with a trapdoor td where:

- td can be used to sample (Gaussian-distributed) solutions the shifted multi-preimage sampling problem with respect to $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell$ and arbitrary targets $\boldsymbol{t}_1, \ldots, \boldsymbol{t}_\ell$
- $(\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell, \mathrm{td})$ can be *publicly* derived from a uniform random matrix $\boldsymbol{B} \leftarrow \mathbb{Z}_q^{n \times m \lceil \log \ell \rceil}$
- SIS/LWE problems are hard with respect to any $\boldsymbol{A}_i$ given $\boldsymbol{B}$

Applications:

- Statistically-hiding vector commitments from SIS with $\mathrm{poly}(\lambda, \log \ell)$-size public parameters, commitments, and openings (and transparent setup)
- Dual-mode NIZK from LWE with polynomial modulus and a transparent setup in statistical ZK mode (and CRS size linear in the length of the hidden-bits string)
- *Subsequent work [BLNWW24]:* statistical ZAP argument from LWE via the hidden-bits approach

Application to Vector Commitments

Given $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_1, \ldots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$,
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

The Wee-Wu blueprint [WW23] (in the language of shifted multi-preimage sampling):

common reference string: $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell$ td

commitment to vector $\boldsymbol{x} \in \mathbb{Z}_q^\ell$

$$\boldsymbol{A}_i \boldsymbol{\pi}_i = x_i \boldsymbol{e}_1 + \boldsymbol{c}$$

opening for bit i

commitment

To commit to $\boldsymbol{x}$: use td to sample $(\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_\ell, \boldsymbol{c})$

Verification checks $\boldsymbol{\pi}_i$ is small and

$$\boldsymbol{A}_i \boldsymbol{\pi}_i = x_i \boldsymbol{e}_1 + \boldsymbol{c}$$

$\boldsymbol{e}_1$: first basis vector

Application to Vector Commitments

> *Given* $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$,
> *find* $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

The Wee-Wu blueprint [WW23] (in the language of shifted multi-preimage sampling):

common reference string:

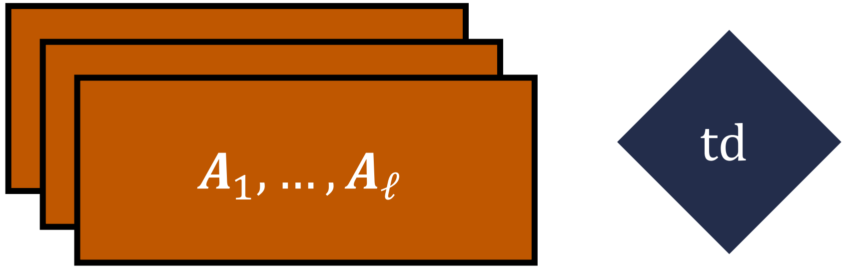

commitment to vector $\boldsymbol{x} \in \mathbb{Z}_q^\ell$

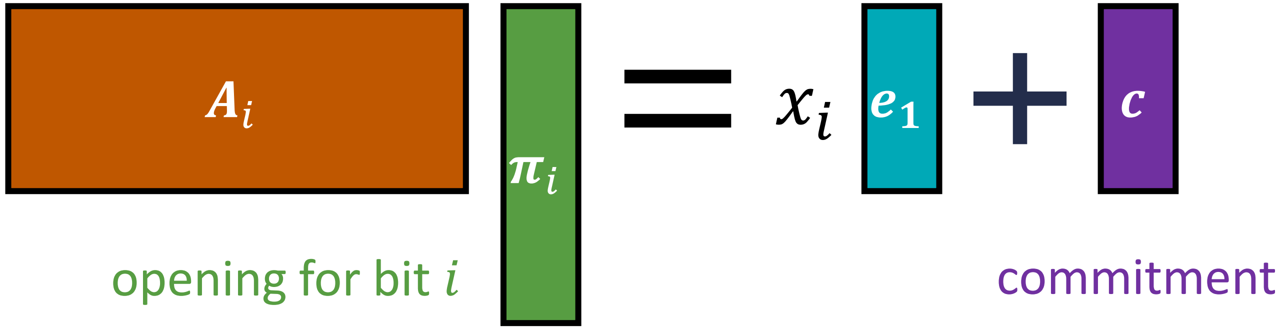

Binding proof:

- Suppose adversary comes up with $\boldsymbol{c}$ and two openings $(x_i, \boldsymbol{\pi}_i), (x_i', \boldsymbol{\pi}_i')$
- Then $\boldsymbol{A}_i(\boldsymbol{\pi}_i - \boldsymbol{\pi}_i') = (x_i - x_i')\, \boldsymbol{e}_1$

Application to Vector Commitments

Given $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$,
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

The Wee-Wu blueprint [WW23] (in the language of shifted multi-preimage sampling):

common reference string: $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$

td

commitment to vector $\boldsymbol{x} \in \mathbb{Z}_q^\ell$

$$\boldsymbol{A}_i \boldsymbol{\pi}_i = x_i \boldsymbol{e}_1 + \boldsymbol{c}$$

opening for bit i

commitment

Binding

- Supp…
  two openings …
- Then $\boldsymbol{A}_i(\boldsymbol{\pi}_i - \boldsymbol{\pi}_i') = (x_i - x_i')\, \boldsymbol{e}_1$
  non-zero

$\boldsymbol{e}_1$ is zero in all but the first row

$\boldsymbol{\pi}_i - \boldsymbol{\pi}_i'$ is a SIS solution
to $\boldsymbol{A}_i$ without the first row

Application to Vector Commitments

> *Given* $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$,
> *find* $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

The Wee-Wu blueprint [WW23] (in the language of shifted multi-preimage sampling):

common reference string:

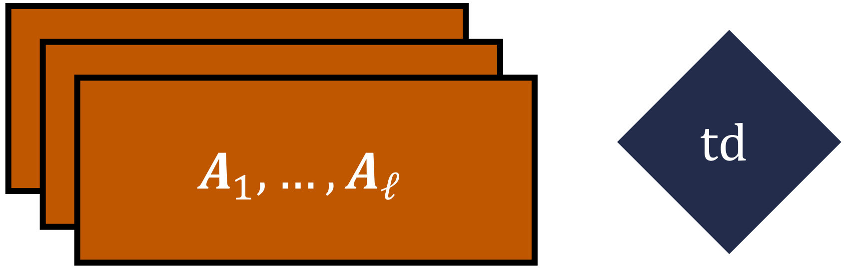

commitment to vector $\boldsymbol{x} \in \mathbb{Z}_q^\ell$

$$\boldsymbol{A}_i \boldsymbol{\pi}_i = x_i \boldsymbol{e}_1 + \boldsymbol{c}$$

opening for bit i

commitment

Hiding proof:

- Distribution of $(\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell, \boldsymbol{c})$ is statistically close to sampling
  $$\boldsymbol{c} \leftarrow \mathbb{Z}_q^n \text{ and } \boldsymbol{\pi}_i \leftarrow \boldsymbol{A}_i^{-1}(x_i \boldsymbol{e}_1 + \boldsymbol{c})$$
- Commitment and openings independent of the values of unopened inputs!

Application to Dual-Mode Hidden-Bits Generators

Hidden-bits generator [FLS90, QRW19]

Used to compile (information-theoretic) NIZK in the hidden-bits model to NIZK in CRS model

common reference string (CRS)

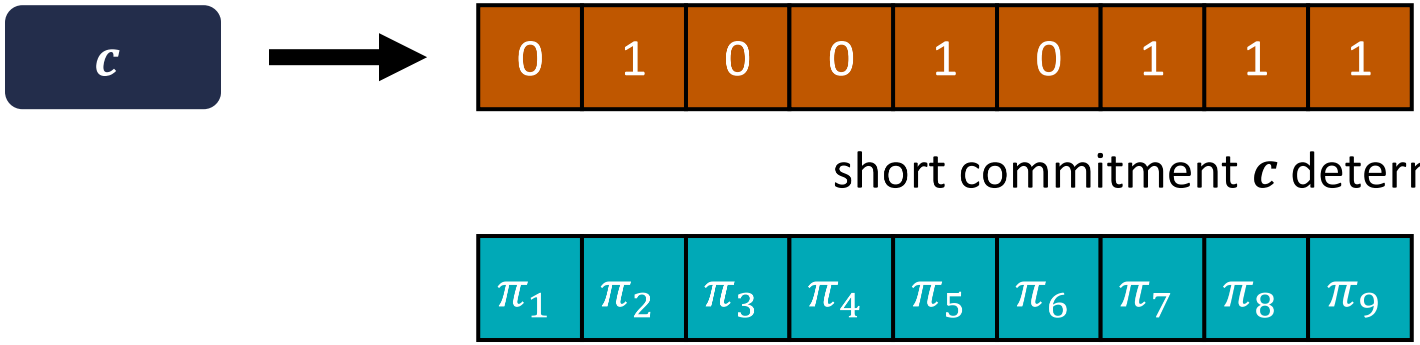

short commitment $\boldsymbol{c}$ determines a long pseudorandom string (length ℓ)

local openings for each bit x_i with respect to $\boldsymbol{c}$ and CRS

Binding: Can only open $\boldsymbol{c}$ to single bit $x_i \in \{0,1\}$ at each index $i \in [\ell]$

Hiding: x_i is pseudorandom given $\boldsymbol{c}$ and (x_j, π_j) for $j \neq i$

Succinctness: $|\boldsymbol{c}| = \text{poly}(\lambda, \log \ell)$

Application to Dual-Mode Hidden-Bits Generators

Hidden-bits generator [FLS90, QRW19]

Used to compile (information-theoretic) NIZK in the hidden-bits model to NIZK in CRS model

common reference string (CRS)

$\boldsymbol{c}$ → 0 1 0 0 1 0 1 1 1

short commitment $\boldsymbol{c}$ determines a long pseudorandom string (length ℓ)

Dual mode if CRS can be sampled to be either statistically binding or statistically hiding

π_6 π_7 π_8 π_9

local openings for each bit x_i with respect to $\boldsymbol{c}$ and CRS

Binding: Can only open $\boldsymbol{c}$ to single bit $x_i \in \{0,1\}$ at each index $i \in [\ell]$

Hiding: x_i is pseudorandom given $\boldsymbol{c}$ and $\left(x_j, \pi_j\right)$ for $j \neq i$

Succinctness: $|\boldsymbol{c}| = \mathrm{poly}(\lambda, \log \ell)$

Application to Dual-Mode Hidden-Bits Generators

The Waters [Wat24] (dual-mode) hidden-bits generator from LWE:

common reference string:

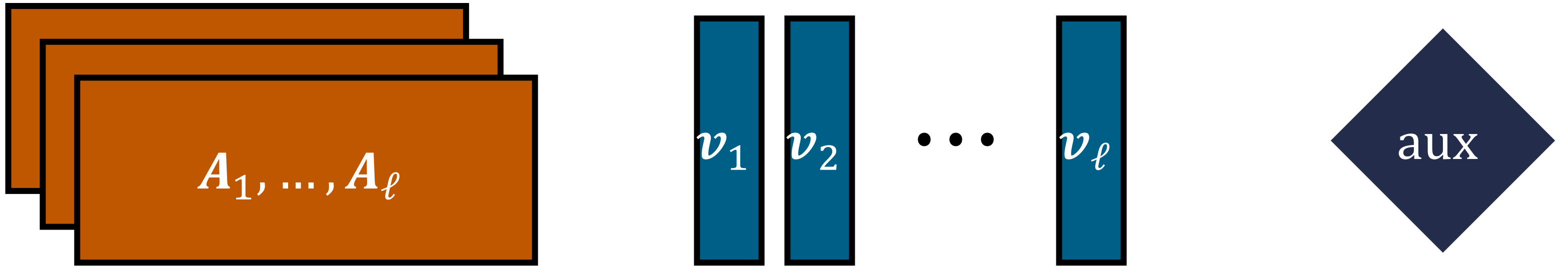

commitment is a vector $\boldsymbol{c} \in \mathbb{Z}_q^n$

openings are short vectors $\boldsymbol{\pi}_i$ where $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{c}$ (sampled using aux)

Application to Dual-Mode Hidden-Bits Generators

The Waters [Wat24] (dual-mode) hidden-bits generator from LWE:

common reference string:

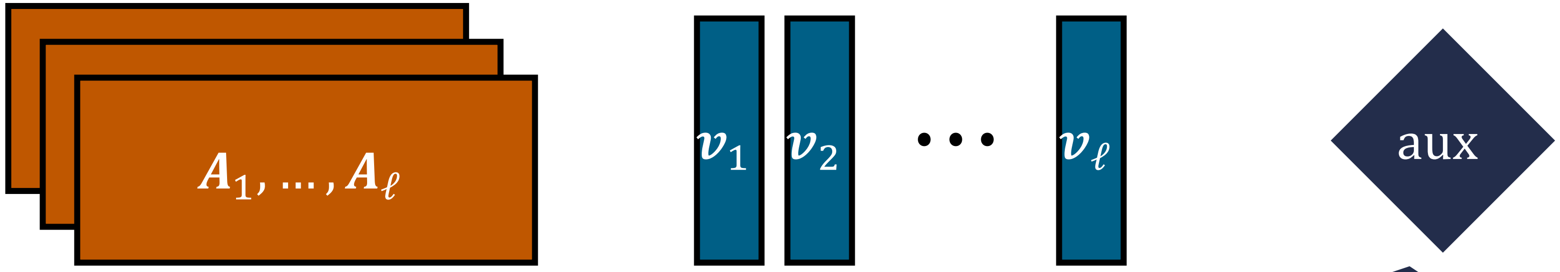

aux contains short preimages $\boldsymbol{A}_i \boldsymbol{W}_i = \boldsymbol{U}$

can derive commitment as $\boldsymbol{c} = \boldsymbol{U}\boldsymbol{t}$ and openings as $\boldsymbol{W}_i \boldsymbol{t}$ $(\boldsymbol{A}_i \boldsymbol{W}_i \boldsymbol{t} = \boldsymbol{U}\boldsymbol{t} = \boldsymbol{c})$

commitment is a vector $\boldsymbol{c} \in \mathbb{Z}_q^n$

openings are short vectors $\boldsymbol{\pi}_i$ where $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{c}$ (sa

Application to Dual-Mode Hidden-Bits Generators

The Waters [Wat24] (dual-mode) hidden-bits generator from LWE:

common reference string:

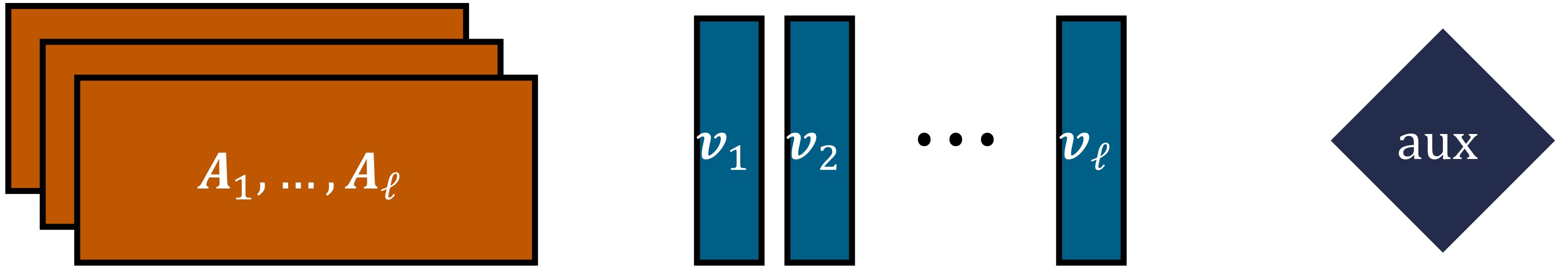

commitment is a vector $\boldsymbol{c} \in \mathbb{Z}_q^n$

openings are short vectors $\boldsymbol{\pi}_i$ where $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{c}$ (sampled using aux)

hidden bits are $x_1, \ldots, x_\ell \in \{0,1\}$ where $x_i = \lfloor \boldsymbol{v}_i^{\mathrm{T}} \boldsymbol{\pi}_i \rceil$

Application to Dual-Mode Hidden-Bits Generators

The Waters [Wat24] (dual-mode) hidden-bits generator from LWE:

common reference string:

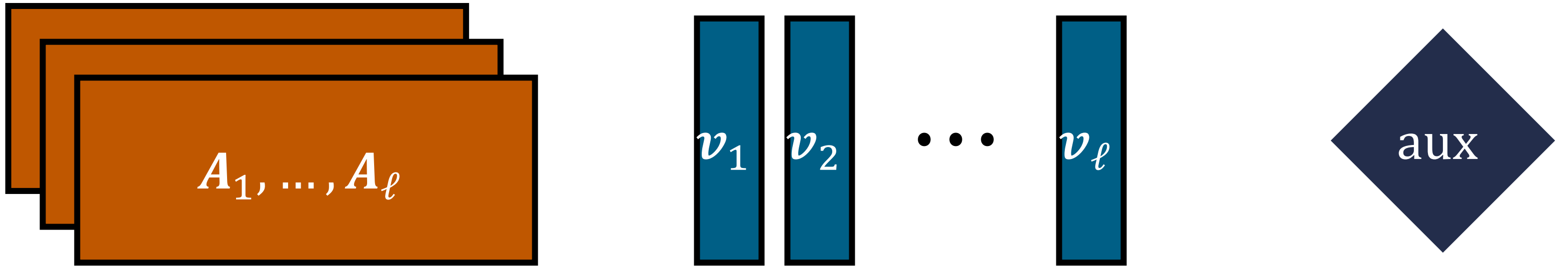

commitment is a vector $\boldsymbol{c} \in \mathbb{Z}_q^n$

openings are short vectors $\boldsymbol{\pi}_i$ where $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{c}$ (sampled using aux)

hidden bits are $x_1, \ldots, x_\ell \in \{0,1\}$ where $x_i = \lfloor \boldsymbol{v}_i^{\mathrm{T}} \boldsymbol{\pi}_i \rceil$

Observe: aux is basically used to solve the shifted multi-preimage sampling problem with respect to $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell$ and targets $\boldsymbol{t}_1, \ldots, \boldsymbol{t}_\ell = \boldsymbol{0}$

Solution is $(\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_\ell, \boldsymbol{c})$ where $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c} = \boldsymbol{c}$

Application to Dual-Mode Hidden-Bits Generators

Our dual-mode hidden-bits generator from LWE:

common reference string:

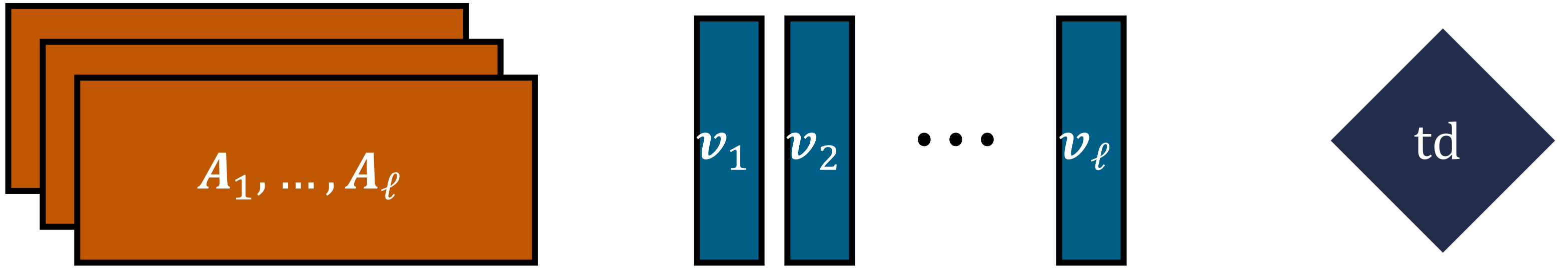

commitment is a vector $\boldsymbol{c} \in \mathbb{Z}_q^n$

openings are short vectors $\boldsymbol{\pi}_i$ where $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{c}$ (from shifted multi-preimage sampler)

hidden bits are $x_1, \dots, x_\ell \in \{0,1\}$ where $x_i = \lfloor \boldsymbol{v}_i^{\mathrm{T}} \boldsymbol{\pi}_i \rceil$

binding mode: $\boldsymbol{v}_i^{\mathrm{T}} = \boldsymbol{s}_i^{\mathrm{T}} \boldsymbol{A}_i + \boldsymbol{e}_i^{\mathrm{T}}$

essentially the same argument as in [Wat24]

value x_i is essentially determined by CRS and $\boldsymbol{c}$:

$$\boldsymbol{v}_i^{\mathrm{T}} \boldsymbol{\pi}_i = \boldsymbol{s}_i^{\mathrm{T}} \boldsymbol{A}_i \boldsymbol{\pi}_i + \boldsymbol{e}_i^{\mathrm{T}} \boldsymbol{\pi}_i \approx \boldsymbol{s}_i^{\mathrm{T}} \boldsymbol{c} \quad \text{(since } \boldsymbol{e}_i^{\mathrm{T}} \boldsymbol{\pi}_i \text{ is small)}$$

Application to Dual-Mode Hidden-Bits Generators

Our dual-mode hidden-bits generator from LWE:

common reference string:

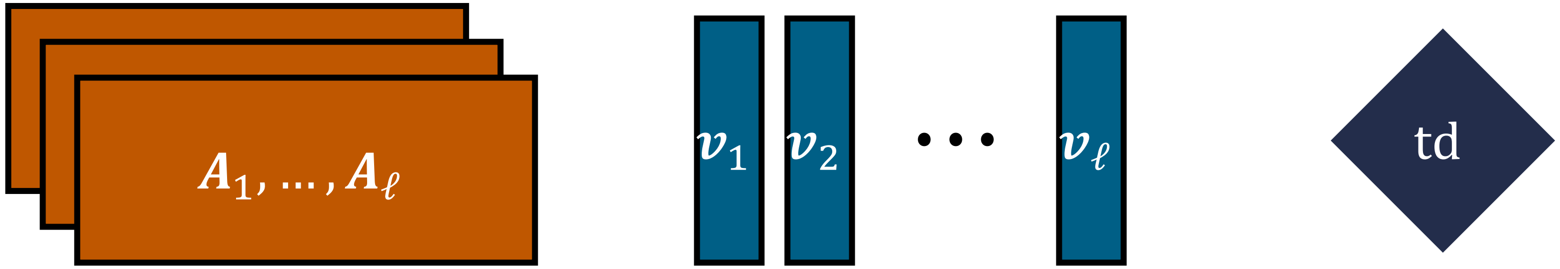

commitment is a vector $\boldsymbol{c} \in \mathbb{Z}_q^n$

openings are short vectors $\boldsymbol{\pi}_i$ where $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{c}$ (from shifted multi-preimage sampler)

hidden bits are $x_1, \dots, x_\ell \in \{0,1\}$ where $x_i = \lfloor \boldsymbol{v}_i^{\mathrm{T}} \boldsymbol{\pi}_i \rceil$

hiding mode: $\boldsymbol{v}_i \leftarrow \mathbb{Z}_q^m$

different argument from [Wat24]

distribution of $(\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell, \boldsymbol{c})$ is statistically close to sampling $\boldsymbol{c} \leftarrow \mathbb{Z}_q^n$ and $\boldsymbol{\pi}_i \leftarrow \boldsymbol{A}_i^{-1}(\boldsymbol{c})$

by leftover hash lemma (with seed $\boldsymbol{v}_i$, source $\boldsymbol{\pi}_i$, we conclude that $\boldsymbol{v}_i^{\mathrm{T}} \boldsymbol{\pi}_i$ is uniform)

Application to Dual-Mode Hidden-Bits Generators

Our dual-mode hidden-bits generator from LWE:

common reference string:

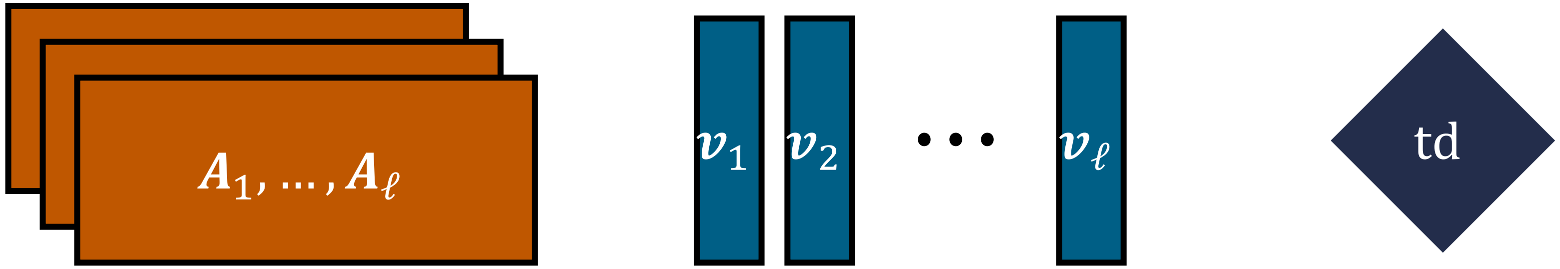

commitment is a vector $\boldsymbol{c} \in \mathbb{Z}_q^n$

openings are short vectors $\boldsymbol{\pi}_i$ where $\boldsymbol{A}_i \boldsymbol{\pi}_i$

hidden bits are $x_1, \dots, x_\ell \in \{0,1\}$ where x_i

Argument in [Wat24] relied on noise smudging (and thus, super-polynomial modulus q)

hiding mode: $\boldsymbol{v}_i \leftarrow \mathbb{Z}_q^m$

different argument from [Wat24]

distribution of $(\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell, \boldsymbol{c})$ is statistically close to sampling $\boldsymbol{c} \leftarrow \mathbb{Z}_q^n$ and $\boldsymbol{\pi}_i \leftarrow \boldsymbol{A}_i^{-1}(\boldsymbol{c})$

by leftover hash lemma (with seed $\boldsymbol{v}_i$, source $\boldsymbol{\pi}_i$, we conclude that $\boldsymbol{v}_i^{\mathrm{T}} \boldsymbol{\pi}_i$ is uniform)

Application to Dual-Mode Hidden-Bits Generators

Our dual-mode hidden-bits generator from LWE:

common reference string:

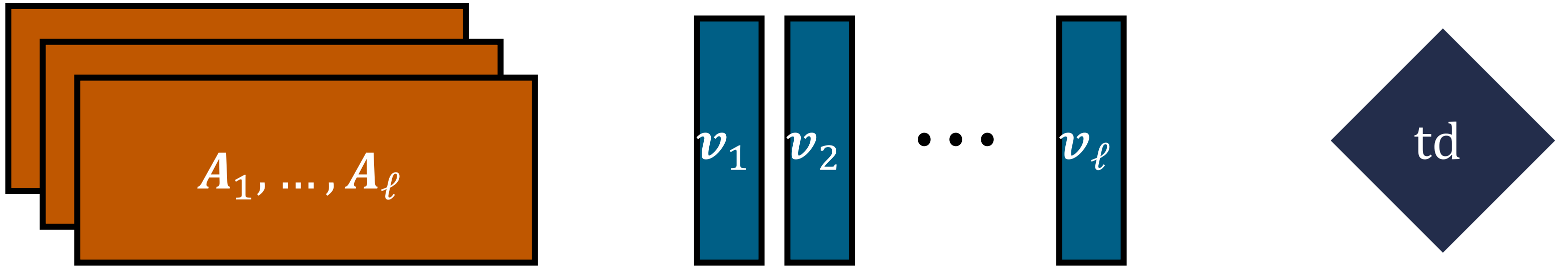

commitment is a vector $\boldsymbol{c} \in \mathbb{Z}_q^n$

openings are short vectors $\boldsymbol{\pi}_i$ where $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{c}$ (from shifted multi-preimage sampler)

hidden bits are $x_1, \dots, x_\ell \in \{0,1\}$ where $x_i = \lfloor \boldsymbol{v}_i^{\mathrm{T}} \boldsymbol{\pi}_i \rceil$

binding mode: $\boldsymbol{v}_i^{\mathrm{T}} = \boldsymbol{s}_i^{\mathrm{T}} \boldsymbol{A}_i + \boldsymbol{e}_i^{\mathrm{T}}$

hiding mode: $\boldsymbol{v}_i \leftarrow \mathbb{Z}_q^m$

modes are indistinguishable if LWE holds with respect to $\boldsymbol{A}_i$ (given $\mathrm{td}, \boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$)

Techniques also give a statistical ZAP argument from quasi-polynomial-hard LWE [BLNWW24]

Constructing a Shifted Multi-Preimage Sampler

Given $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$,
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

The Wee-Wu approach [WW23] for shifted multi-preimage sampling:

Sample $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \leftarrow \mathbb{Z}_q^{n \times m}$ and give out a **trapdoor** for the matrix

$$\boldsymbol{D}_\ell = \left[\begin{array}{ccc|c} \boldsymbol{A}_1 & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{A}_\ell & \boldsymbol{G} \end{array}\right] \qquad \boldsymbol{G} = \begin{bmatrix} 1 & 2 & \cdots & 2^t & & & & & \\ & & & & \ddots & & & & \\ & & & & & 1 & 2 & \cdots & 2^t \end{bmatrix}$$

$$t = \lceil \log q \rceil - 1$$

Using trapdoor for $\boldsymbol{D}_\ell$, can sample (Gaussian) solutions to the linear system

$$\left[\begin{array}{ccc|c} \boldsymbol{A}_1 & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{A}_\ell & \boldsymbol{G} \end{array}\right] \cdot \begin{bmatrix} \boldsymbol{\pi}_1 \\ \vdots \\ \boldsymbol{\pi}_\ell \\ \hat{\boldsymbol{c}} \end{bmatrix} = \begin{bmatrix} \boldsymbol{t}_1 \\ \vdots \\ \boldsymbol{t}_\ell \end{bmatrix} \quad \Longrightarrow$$

for all $i \in [\ell]$, $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i - \boldsymbol{G}\hat{\boldsymbol{c}}$
set $\boldsymbol{c} = -\boldsymbol{G}\hat{\boldsymbol{c}}$

Constructing a Shifted Multi-Preimage Sampler

Given $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n\times m}$ *and* $\boldsymbol{t}_1, \ldots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$,
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i\boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

The Wee-Wu approach [WW23] for shifted multi-preimage sampling:

Sample $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell \leftarrow \mathbb{Z}_q^{n\times m}$ and give out a **trapdoor** for the matrix

$$\boldsymbol{D}_\ell = \left[\begin{array}{ccc|c} \boldsymbol{A}_1 & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{A}_\ell & \boldsymbol{G} \end{array}\right] \qquad \boldsymbol{G} = \begin{bmatrix} 1 & 2 & \cdots & 2^t & & & & & \\ & & & & \ddots & & & & \\ & & & & & 1 & 2 & \cdots & 2^t \end{bmatrix}$$

- Trapdoor for $\boldsymbol{D}_\ell$ is sufficient to solve the shifted multi-preimage sampling problem with respect to $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell$
- When $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell$ are uniform, trapdoor for $\boldsymbol{D}_\ell$ can be obtained given **all but one** trapdoors of $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell$ (because $\boldsymbol{G}$ has a public trapdoor)
- LWE/SIS hold with respect to any $\boldsymbol{A}_i$ given $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell$ and trapdoor for $\boldsymbol{D}_\ell$
- **Limitation:** trapdoor for $\boldsymbol{D}_\ell$ is a **structured** matrix (and size ℓ^2)

Constructing a Shifted Multi-Preimage Sampler

Given $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n\times m}$ *and* $\boldsymbol{t}_1, \ldots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$,
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i\boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

The Wee-Wu approach [WW23] for shifted multi-preimage sampling:

Sample $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell \leftarrow \mathbb{Z}_q^{n\times m}$ and give out a **trapdoor** for the matrix

$$\boldsymbol{D}_\ell = \left[\begin{array}{ccc|c} \boldsymbol{A}_1 & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{A}_\ell & \boldsymbol{G} \end{array}\right] \qquad \boldsymbol{G} = \begin{bmatrix} 1 & 2 & \cdots & 2^t & & & & & \\ & & & & \ddots & & & & \\ & & & & & 1 & 2 & \cdots & 2^t \end{bmatrix}$$

This work: set $\boldsymbol{A}_i = \boldsymbol{B} - \boldsymbol{u}_i^{\mathrm{T}} \otimes \boldsymbol{G}$ where $\boldsymbol{u}_i$ is binary representation of i

$$\boldsymbol{A}_1 = \begin{array}{|c|c|c|c|} \hline \boldsymbol{B}_1 - 0\cdot\boldsymbol{G} & \boldsymbol{B}_2 - 0\cdot\boldsymbol{G} & \cdots & \boldsymbol{B}_\ell - 1\cdot\boldsymbol{G} \\ \hline \end{array}$$

Constructing a Shifted Multi-Preimage Sampler

Given $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n\times m}$ *and* $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$*,*
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i\boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

The Wee-Wu approach [WW23] for shifted multi-preimage sampling:

Sample $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \leftarrow \mathbb{Z}_q^{n\times m}$ and give out a **trapdoor** for the matrix

$$\boldsymbol{D}_\ell = \left[\begin{array}{ccc|c} \boldsymbol{A}_1 & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{A}_\ell & \boldsymbol{G} \end{array}\right] \qquad \boldsymbol{G} = \begin{bmatrix} 1 & 2 & \cdots & 2^t & & & & & \\ & & & & \ddots & & & & \\ & & & & & 1 & 2 & \cdots & 2^t \end{bmatrix}$$

This work: set $\boldsymbol{A}_i = \boldsymbol{B} - \boldsymbol{u}_i^{\mathrm{T}} \otimes \boldsymbol{G}$ where $\boldsymbol{u}_i$ is binary representation of i

- $\boldsymbol{B} \in \mathbb{Z}_q^{n\times m\lceil\log \ell\rceil}$ (slightly wider)
- The matrix $\boldsymbol{D}_\ell$ has a **public** trapdoor

Constructing a Shifted Multi-Preimage Sampler

Homomorphic computation using lattices [GSW13, BGGHNSVV14]

Encodes a vector $\boldsymbol{x} \in \{0,1\}^\ell$ with respect to matrix $\boldsymbol{B} = [\boldsymbol{B}_1 | \cdots | \boldsymbol{B}_\ell] \in \mathbb{Z}_q^{n \times \ell m}$

$\boldsymbol{B}_1 - x_1 \boldsymbol{G}$	$\boldsymbol{B}_2 - x_2 \boldsymbol{G}$	$\cdots$	$\boldsymbol{B}_\ell - x_\ell \boldsymbol{G}$

$\boldsymbol{B} - \boldsymbol{x}^{\mathrm{T}} \otimes \boldsymbol{G}$

Given any function $f: \{0,1\}^\ell \to \{0,1\}$, there exists a short matrix $\boldsymbol{H}_{\boldsymbol{B},f,\boldsymbol{x}}$ where

$$\left(\boldsymbol{B} - \boldsymbol{x}^{\mathrm{T}} \otimes \boldsymbol{G}\right) \cdot \boldsymbol{H}_{\boldsymbol{B},f,\boldsymbol{x}} = \boldsymbol{B}_f - f(x) \cdot \boldsymbol{G}$$

encoding of $\boldsymbol{x}$ with respect to $\boldsymbol{B}$

encoding of $f(\boldsymbol{x})$ with respect to $\boldsymbol{B}_f$

Given $\boldsymbol{B}$ and f, can efficiently compute the matrix $\boldsymbol{B}_f$

Constructing a Shifted Multi-Preimage Sampler

Define the indicator function

$$\delta_{\boldsymbol{u}}(\boldsymbol{x}) = \begin{cases} 1, & \boldsymbol{x} = \boldsymbol{u} \\ 0, & \boldsymbol{x} \neq \boldsymbol{u} \end{cases}$$

For simplicity, we will write
- $B_{\boldsymbol{u}} \coloneqq B_{\delta_{\boldsymbol{u}}}$
- $H_{\boldsymbol{B},\boldsymbol{u},\boldsymbol{x}} \coloneqq H_{\boldsymbol{B},\delta_{\boldsymbol{u}},\boldsymbol{x}}$

$$\left(\boldsymbol{B} - \boldsymbol{x}^{\mathrm{T}} \otimes \boldsymbol{G}\right) \cdot \boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u},\boldsymbol{x}} = \boldsymbol{B}_{\boldsymbol{u}} - \delta_{\boldsymbol{u}}(\boldsymbol{x}) \cdot \boldsymbol{G} = \begin{cases} \boldsymbol{B}_{\boldsymbol{u}} - \boldsymbol{G} & \boldsymbol{x} = \boldsymbol{u} \\ \boldsymbol{B}_{\boldsymbol{u}} & \boldsymbol{x} \neq \boldsymbol{u} \end{cases}$$

Let $\boldsymbol{u}_1, \dots, \boldsymbol{u}_\ell \in \{0,1\}^{\lceil \log \ell \rceil}$ be distinct vectors (e.g., $\boldsymbol{u}_i$ is bit representation of i)

Consider the matrix $\boldsymbol{D}_\ell = \left[\begin{array}{ccc|c} \boldsymbol{A}_1 & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{A}_\ell & \boldsymbol{G} \end{array}\right]$

$$\boldsymbol{A}_i \coloneqq \boldsymbol{B} - \boldsymbol{u}_i \otimes \boldsymbol{G}$$

Constructing a Shifted Multi-Preimage Sampler

$$\left(\boldsymbol{B} - \boldsymbol{x}^{\mathrm{T}} \otimes \boldsymbol{G}\right) \cdot \boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u},\boldsymbol{x}} = \boldsymbol{B}_{\boldsymbol{u}} - \delta_{\boldsymbol{u}}(\boldsymbol{x}) \cdot \boldsymbol{G} \quad = \begin{cases} \boldsymbol{B}_{\boldsymbol{u}} - \boldsymbol{G} & \boldsymbol{x} = \boldsymbol{u} \\ \boldsymbol{B}_{\boldsymbol{u}} & \boldsymbol{x} \neq \boldsymbol{u} \end{cases}$$

$$\boldsymbol{D}_{\ell} = \left[\begin{array}{ccc|c} \boldsymbol{A}_1 & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{A}_{\ell} & \boldsymbol{G} \end{array}\right] = \left[\begin{array}{ccc|c} \boldsymbol{B} - \boldsymbol{u}_1 \otimes \boldsymbol{G} & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{B} - \boldsymbol{u}_{\ell} \otimes \boldsymbol{G} & \boldsymbol{G} \end{array}\right]$$

$$\left[\begin{array}{ccc|c} \boldsymbol{B} - \boldsymbol{u}_1 \otimes \boldsymbol{G} & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{B} - \boldsymbol{u}_{\ell} \otimes \boldsymbol{G} & \boldsymbol{G} \end{array}\right] \times \begin{bmatrix} -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_1,\boldsymbol{u}_1} & \cdots & -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_\ell,\boldsymbol{u}_1} \\ \vdots & \ddots & \vdots \\ -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_1,\boldsymbol{u}_\ell} & \cdots & -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_\ell,\boldsymbol{u}_\ell} \\ \boldsymbol{G}^{-1}\left(\boldsymbol{B}_{\boldsymbol{u}_1}\right) & \cdots & \boldsymbol{G}^{-1}\left(\boldsymbol{B}_{\boldsymbol{u}_\ell}\right) \end{bmatrix}$$

Constructing a Shifted Multi-Preimage Sampler

$$\left(\boldsymbol{B} - \boldsymbol{x}^{\mathrm{T}} \otimes \boldsymbol{G}\right) \cdot \boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u},\boldsymbol{x}} = \boldsymbol{B}_{\boldsymbol{u}} - \delta_{\boldsymbol{u}}(\boldsymbol{x}) \cdot \boldsymbol{G} \quad = \begin{cases} \boldsymbol{B}_{\boldsymbol{u}} - \boldsymbol{G} & \boldsymbol{x} = \boldsymbol{u} \\ \boldsymbol{B}_{\boldsymbol{u}} & \boldsymbol{x} \neq \boldsymbol{u} \end{cases}$$

Block in row i and column j:

$$(\boldsymbol{B} - \boldsymbol{u}_i \otimes \boldsymbol{G}) \cdot \left(-\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_j,\boldsymbol{u}_i}\right) + \boldsymbol{G} \cdot \boldsymbol{G}^{-1}\left(\boldsymbol{B}_{\boldsymbol{u}_j}\right)$$

$$\left[\begin{array}{ccc|c} \boldsymbol{B} - \boldsymbol{u}_1 \otimes \boldsymbol{G} & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{B} - \boldsymbol{u}_\ell \otimes \boldsymbol{G} & \boldsymbol{G} \end{array}\right] \times \begin{bmatrix} -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_1,\boldsymbol{u}_1} & \cdots & -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_\ell,\boldsymbol{u}_1} \\ \vdots & \ddots & \vdots \\ -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_1,\boldsymbol{u}_\ell} & \cdots & -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_\ell,\boldsymbol{u}_\ell} \\ \boldsymbol{G}^{-1}\left(\boldsymbol{B}_{\boldsymbol{u}_1}\right) & \cdots & \boldsymbol{G}^{-1}\left(\boldsymbol{B}_{\boldsymbol{u}_\ell}\right) \end{bmatrix}$$

Constructing a Shifted Multi-Preimage Sampler

$$(\boldsymbol{B} - \boldsymbol{x}^{\mathrm{T}} \otimes \boldsymbol{G}) \cdot \boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u},\boldsymbol{x}} = \boldsymbol{B}_{\boldsymbol{u}} - \delta_{\boldsymbol{u}}(\boldsymbol{x}) \cdot \boldsymbol{G} \quad = \begin{cases} \boldsymbol{B}_{\boldsymbol{u}} - \boldsymbol{G} & \boldsymbol{x} = \boldsymbol{u} \\ \boldsymbol{B}_{\boldsymbol{u}} & \boldsymbol{x} \neq \boldsymbol{u} \end{cases}$$

Block in row i and column j:

$$(\boldsymbol{B} - \boldsymbol{u}_i \otimes \boldsymbol{G}) \cdot \left(-\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_j,\boldsymbol{u}_i}\right) + \boldsymbol{G} \cdot \boldsymbol{G}^{-1}\left(\boldsymbol{B}_{\boldsymbol{u}_j}\right) = -\boldsymbol{B}_{\boldsymbol{u}_j} + \delta_{\boldsymbol{u}_i}(\boldsymbol{u}_j) \cdot \boldsymbol{G}$$

$$\left[\begin{array}{ccc|c} \boldsymbol{B} - \boldsymbol{u}_1 \otimes \boldsymbol{G} & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{B} - \boldsymbol{u}_\ell \otimes \boldsymbol{G} & \boldsymbol{G} \end{array}\right] \times \begin{bmatrix} -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_1,\boldsymbol{u}_1} & \cdots & -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_\ell,\boldsymbol{u}_1} \\ \vdots & \ddots & \vdots \\ -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_1,\boldsymbol{u}_\ell} & \cdots & -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_\ell,\boldsymbol{u}_\ell} \\ \boldsymbol{G}^{-1}\left(\boldsymbol{B}_{\boldsymbol{u}_1}\right) & \cdots & \boldsymbol{G}^{-1}\left(\boldsymbol{B}_{\boldsymbol{u}_\ell}\right) \end{bmatrix}$$

Constructing a Shifted Multi-Preimage Sampler

$$\left(\boldsymbol{B} - \boldsymbol{x}^{\mathrm{T}} \otimes \boldsymbol{G}\right) \cdot \boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u},\boldsymbol{x}} = \boldsymbol{B}_{\boldsymbol{u}} - \delta_{\boldsymbol{u}}(\boldsymbol{x}) \cdot \boldsymbol{G} \quad = \begin{cases} \boldsymbol{B}_{\boldsymbol{u}} - \boldsymbol{G} & \boldsymbol{x} = \boldsymbol{u} \\ \boldsymbol{B}_{\boldsymbol{u}} & \boldsymbol{x} \neq \boldsymbol{u} \end{cases}$$

Block in row i and column j:

$$\left(\boldsymbol{B} - \boldsymbol{u}_i \otimes \boldsymbol{G}\right) \cdot \left(-\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_j,\boldsymbol{u}_i}\right) + \boldsymbol{G} \cdot \boldsymbol{G}^{-1}\left(\boldsymbol{B}_{\boldsymbol{u}_j}\right) = -\boldsymbol{B}_{\boldsymbol{u}_j} + \delta_{\boldsymbol{u}_i}\left(\boldsymbol{u}_j\right) \cdot \boldsymbol{G} + \boldsymbol{B}_{\boldsymbol{u}_j} = \begin{cases} \boldsymbol{G}, & i = j \\ \boldsymbol{0}, & i \neq j \end{cases}$$

$$\left[\begin{array}{ccc|c} \boldsymbol{B} - \boldsymbol{u}_1 \otimes \boldsymbol{G} & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{B} - \boldsymbol{u}_\ell \otimes \boldsymbol{G} & \boldsymbol{G} \end{array}\right] \times \begin{bmatrix} -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_1,\boldsymbol{u}_1} & \cdots & -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_\ell,\boldsymbol{u}_1} \\ \vdots & \ddots & \vdots \\ -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_1,\boldsymbol{u}_\ell} & \cdots & -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_\ell,\boldsymbol{u}_\ell} \\ \boldsymbol{G}^{-1}\left(\boldsymbol{B}_{\boldsymbol{u}_1}\right) & \cdots & \boldsymbol{G}^{-1}\left(\boldsymbol{B}_{\boldsymbol{u}_\ell}\right) \end{bmatrix} = \begin{bmatrix} \boldsymbol{G} & & \\ & \ddots & \\ & & \boldsymbol{G} \end{bmatrix}$$

Constructing a Shifted Multi-Preimage Sampler

Key observations:

- Matrix $\boldsymbol{D}_\ell$ can be described entirely by matrix $\boldsymbol{B}$
- Vectors $\boldsymbol{u}_1, \dots, \boldsymbol{u}_\ell$ just need to be distinct (e.g., $\boldsymbol{u}_i$ is binary representation of i)
- $\boldsymbol{D}_\ell$ has a public trapdoor (determined by $\boldsymbol{B}, \boldsymbol{u}_1, \dots, \boldsymbol{u}_\ell$)
- Since we are considering indicator functions, $\left\| \boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_i,\boldsymbol{u}_j} \right\| = 1$

$$\underbrace{\left[\begin{array}{ccc|c} \boldsymbol{B} - \boldsymbol{u}_1 \otimes \boldsymbol{G} & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{B} - \boldsymbol{u}_\ell \otimes \boldsymbol{G} & \boldsymbol{G} \end{array}\right]}_{\boldsymbol{D}_\ell} \times \underbrace{\begin{bmatrix} -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_1,\boldsymbol{u}_1} & \cdots & -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_\ell,\boldsymbol{u}_1} \\ \vdots & \ddots & \vdots \\ -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_1,\boldsymbol{u}_\ell} & \cdots & -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_\ell,\boldsymbol{u}_\ell} \\ \boldsymbol{G}^{-1}(\boldsymbol{B}_{\boldsymbol{u}_1}) & \cdots & \boldsymbol{G}^{-1}(\boldsymbol{B}_{\boldsymbol{u}_\ell}) \end{bmatrix}}_{\text{trapdoor for } \boldsymbol{D}_\ell} = \begin{bmatrix} \boldsymbol{G} & & \\ & \ddots & \\ & & \boldsymbol{G} \end{bmatrix}$$

Constructing a Shifted Multi-Preimage Sampler

$$\underbrace{\left[\begin{array}{ccc|c} \boldsymbol{B}-\boldsymbol{u}_1\otimes\boldsymbol{G} & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{B}-\boldsymbol{u}_\ell\otimes\boldsymbol{G} & \boldsymbol{G} \end{array}\right]}_{\boldsymbol{D}_\ell} \underbrace{\begin{bmatrix} -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_1,\boldsymbol{u}_1} & \cdots & -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_\ell,\boldsymbol{u}_1} \\ \vdots & \ddots & \vdots \\ -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_1,\boldsymbol{u}_\ell} & \cdots & -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_\ell,\boldsymbol{u}_\ell} \\ \boldsymbol{G}^{-1}(\boldsymbol{B}_{\boldsymbol{u}_1}) & \cdots & \boldsymbol{G}^{-1}(\boldsymbol{B}_{\boldsymbol{u}_\ell}) \end{bmatrix}}_{\text{trapdoor for } \boldsymbol{D}_\ell} = \begin{bmatrix} \boldsymbol{G} & & \\ & \ddots & \\ & & \boldsymbol{G} \end{bmatrix}$$

For *any* matrix $\boldsymbol{B} \in \mathbb{Z}_q^{n\times m\lceil\log \ell\rceil}$, the matrix $\boldsymbol{D}_\ell$ has a public trapdoor which can be used to solve the shifted multi-preimage sampling problem with respect to $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$ where $\boldsymbol{A}_i = \boldsymbol{B} - \boldsymbol{u}_i \otimes \boldsymbol{G}$

$$\begin{bmatrix} \boldsymbol{A}_1 & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{A}_\ell & \boldsymbol{G} \end{bmatrix} \cdot \begin{bmatrix} \boldsymbol{\pi}_1 \\ \vdots \\ \boldsymbol{\pi}_\ell \\ \hat{\boldsymbol{c}} \end{bmatrix} = \begin{bmatrix} \boldsymbol{t}_1 \\ \vdots \\ \boldsymbol{t}_\ell \end{bmatrix} \quad \Longrightarrow \quad \begin{array}{l} \text{for all } i \in [\ell],\ \boldsymbol{A}_i\boldsymbol{\pi}_i = \boldsymbol{t}_i - \boldsymbol{G}\hat{\boldsymbol{c}} \\ \text{set } \boldsymbol{c} = -\boldsymbol{G}\hat{\boldsymbol{c}} \end{array}$$

Summary

Given $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_{\boldsymbol{1}}, \dots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$*,*
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

New approach to sample $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$ together with a trapdoor td where:

- td can be used to sample (Gaussian-distributed) solutions the shifted multi-preimage sampling problem with respect to $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$ and arbitrary targets $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell$
- $(\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell, \mathrm{td})$ can be *publicly* derived from a uniform random matrix $\boldsymbol{B} \leftarrow \mathbb{Z}_q^{n \times m \lceil \log \ell \rceil}$
- SIS/LWE problems are hard with respect to any $\boldsymbol{A}_i$ given $\boldsymbol{B}$

Applications:

- Statistically-hiding vector commitments from SIS with $\mathrm{poly}(\lambda, \log \ell)$-size public parameters, commitments, and openings (and transparent setup)
- Dual-mode NIZK from LWE with polynomial modulus and a transparent setup in statistical ZK mode (and CRS size linear in the length of the hidden-bits string)
- *Subsequent work [BLNWW24]:* statistical ZAP argument from LWE via the hidden-bits approach

Summary

Given $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$,
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

More broadly: ability to sample *structured preimages* is very useful for many applications

Can enable this by either publishing hints or a trapdoor in the public parameters

Trapdoor for $\boldsymbol{D}_\ell = \left[\begin{array}{ccc|c} \boldsymbol{A}_1 & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{A}_\ell & \boldsymbol{G} \end{array}\right]$ $\Rightarrow$

Vector commitments
Dual-mode NIZK
Statistical ZAP arguments [BLNWW24]

security based on standard SIS/LWE

More power available with other types of trapdoors!

security based on succinct LWE [Wee24]

Trapdoor for $\boldsymbol{D}_\ell = \left[\begin{array}{ccc|c} \boldsymbol{A} & & & \boldsymbol{W}_1 \\ & \ddots & & \vdots \\ & & \boldsymbol{A} & \boldsymbol{W}_\ell \end{array}\right]$ $\Rightarrow$

ABE with succinct ciphertexts [Wee24]
Functional commitments [WW23b]
Distributed broadcast encryption [CW24]
(Succinct) registered ABE [CHW24]

Very useful for *compression*

and more!

Summary

Given $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ *and* $\boldsymbol{t}_1, \dots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$,
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and short* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i\boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

$$\underbrace{\left[\begin{array}{ccc|c} \boldsymbol{B} - \boldsymbol{u}_1 \otimes \boldsymbol{G} & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{B} - \boldsymbol{u}_\ell \otimes \boldsymbol{G} & \boldsymbol{G} \end{array}\right]}_{\boldsymbol{D}_\ell} \times \underbrace{\begin{bmatrix} -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_1,\boldsymbol{u}_1} & \cdots & -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_\ell,\boldsymbol{u}_1} \\ \vdots & \ddots & \vdots \\ -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_1,\boldsymbol{u}_\ell} & \cdots & -\boldsymbol{H}_{\boldsymbol{B},\boldsymbol{u}_\ell,\boldsymbol{u}_\ell} \\ \boldsymbol{G}^{-1}(\boldsymbol{B}_{\boldsymbol{u}_1}) & \cdots & \boldsymbol{G}^{-1}(\boldsymbol{B}_{\boldsymbol{u}_\ell}) \end{bmatrix}}_{\text{trapdoor for } \boldsymbol{D}_\ell} = \begin{bmatrix} \boldsymbol{G} & & \\ & \ddots & \\ & & \boldsymbol{G} \end{bmatrix}$$

Thank you!

`https://eprint.iacr.org/2024/1401`