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 low-norm $\boldsymbol{\pi} \in \mathbb{Z}_q^{m}$ *where* $\boldsymbol{A}\boldsymbol{\pi} = \boldsymbol{t}$

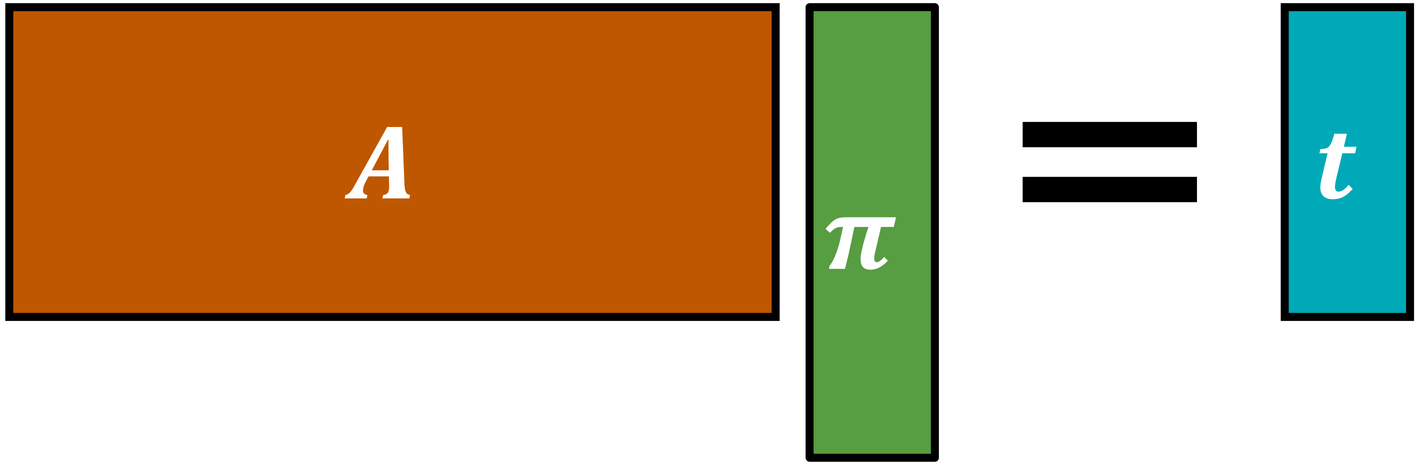

Problem is hard in general:

- 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

Multi-Preimage Sampling

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

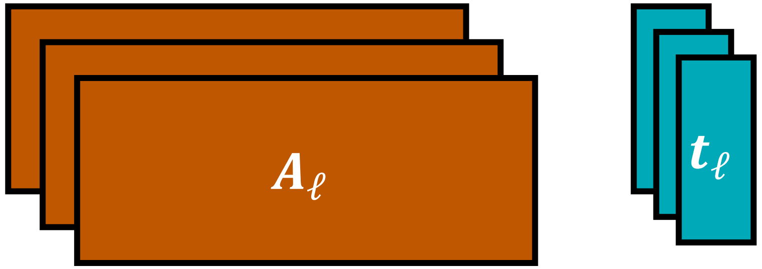

find low-norm $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^{m}$ *where* $\boldsymbol{A}_i\boldsymbol{\pi}_i = \boldsymbol{t}_i$ *for all* $i \in [\ell]$

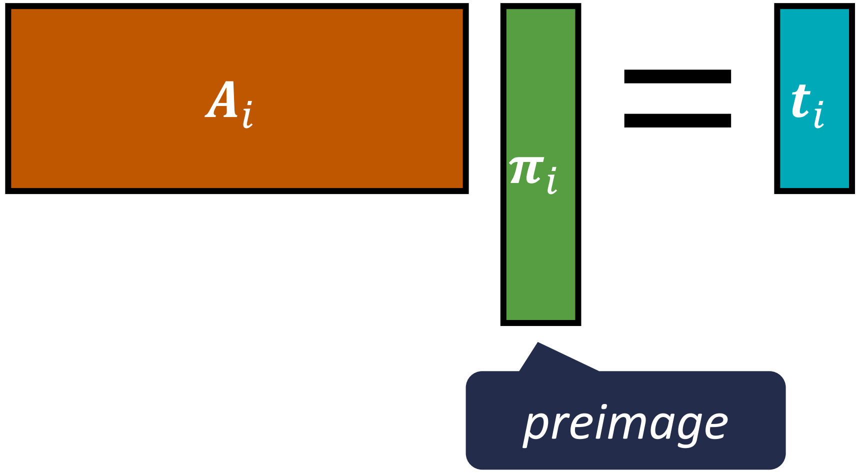

Shifted Multi-Preimage Sampling

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

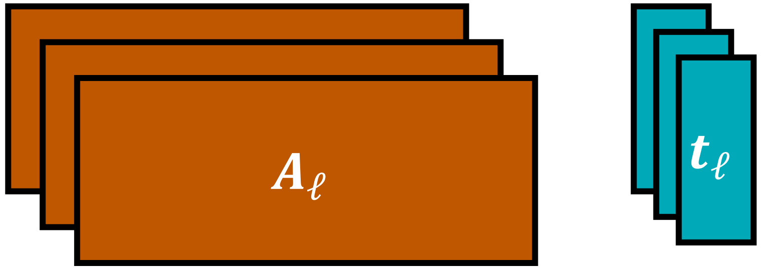

find $\boldsymbol{c} \in \mathbb{Z}_q^n$ *and low-norm* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

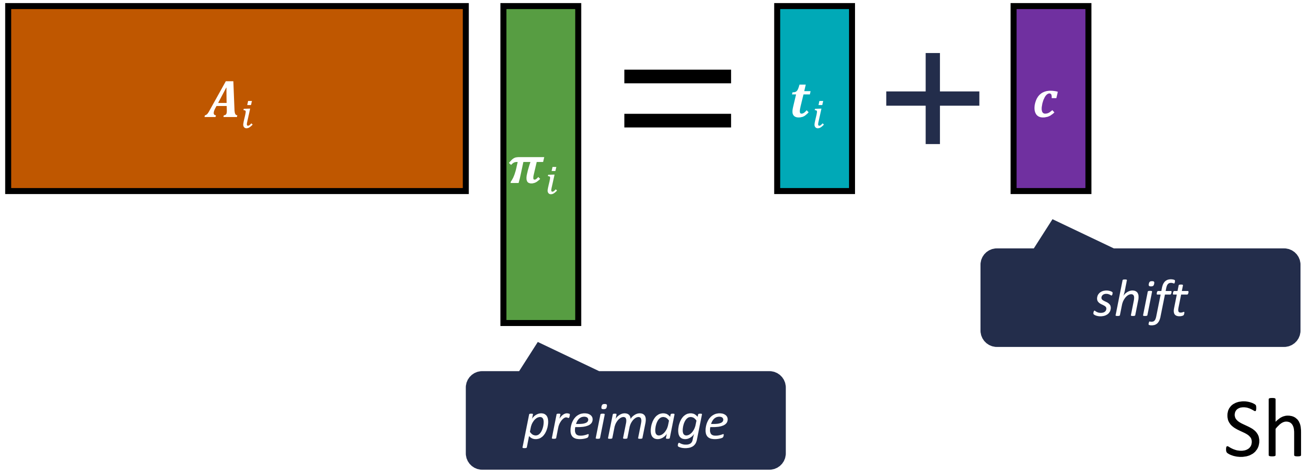

Shift gives **one** degree of freedom

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 low-norm $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ 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 $\boldsymbol{A}_i$

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 $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell \in \mathbb{Z}_q^{n \times m}$ and $\boldsymbol{t}_1, \ldots, \boldsymbol{t}_\ell \in \mathbb{Z}_q^n$,
find $\boldsymbol{c} \in \mathbb{Z}_q^n$ and low-norm $\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ 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 solve the shifted multi-preimage sampling problem

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 low-norm* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *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 solve the shifted multi-preimage sampling problem
- $(\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell, \mathrm{td})$ 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 $\boldsymbol{B}$

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 low-norm* $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *where* $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ *for all* $i \in [\ell]$

Implications:

- Statistically-hiding vector commitments from SIS with $\mathrm{poly}(\lambda, \log \ell)$-size public parameters, commitments, and openings (and transparent setup)
  - Previous lattice-based schemes had long *structured* CRS [WW23] or were computationally hiding [dCP23]

- Dual-mode NIZK from LWE via the hidden-bits model with polynomial modulus, CRS size linear in the length of the hidden-bits string, and transparent setup in statistical ZK mode
  - 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
  - Achieves properties as those obtained via the correlation-intractability framework [CCHLRRW19, PS19]
  - *Subsequent work [BLNWW24]:* statistical ZAP argument from LWE via the hidden-bits approach

This talk

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

Concurrent work [BCDJMS25]: dual-mode NIZK in the hidden-bits model from LWE

- Polynomial modulus and transparent setup in statistical ZK mode
- CRS size is quadratic in the length of the hidden-bits string
- Multi-theorem zero-knowledge requires "or-proof" (need to apply NIZK to cryptographic language)
- Does not need lattice trapdoors

- Dual-mode NIZK from LWE via the hidden-bits model with polynomial modulus, CRS size linear in the length of the hidden-bits string, and transparent setup in statistical ZK mode
  - 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
  - Achieves properties as those obtained via the correlation-intractability framework [CCHLRRW19, PS19]
  - *Subsequent work [BLNWW24]:* statistical ZAP argument from LWE via the hidden-bits approach

This talk

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)

$\boldsymbol{c} \longrightarrow$

x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_9
0	1	0	0	1	0	1	1	1

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

π_1	π_2	π_3	π_4	π_5	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 (x_j, π_j) for $j \neq i$

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

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)

$\boldsymbol{c} \longrightarrow$

x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_9
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 (x_j, π_j) for $j \neq i$

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

Dual-Mode Hidden-Bits Generators

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

common reference string:

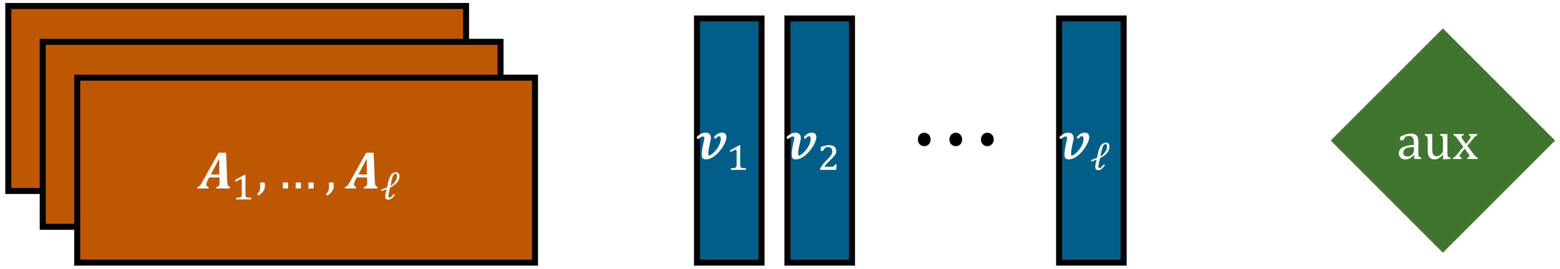

ℓ: length of hidden-bits string

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

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

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

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

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

Dual-Mode Hidden-Bits Generators

Our dual-mode hidden-bits generator from LWE:

common reference string:

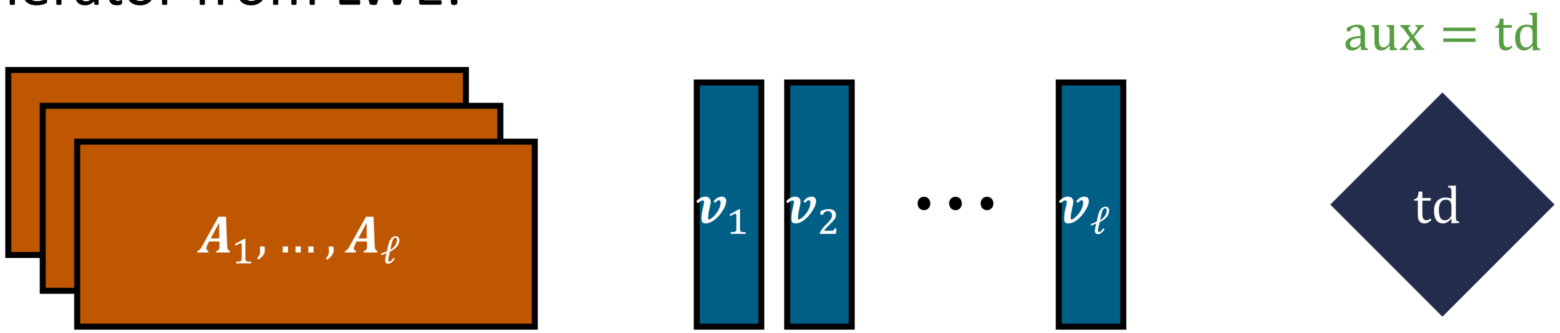

ℓ: length of hidden-bits string

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

openings are low-norm vectors $\boldsymbol{\pi}_i$ where $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{c}$ (via td for 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 $\boldsymbol{c}$:

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

value of $\boldsymbol{s}_i$ (from CRS) and $\boldsymbol{c}$ determine x_i

Dual-Mode Hidden-Bits Generators

Our dual-mode hidden-bits generator from LWE:

common reference string:

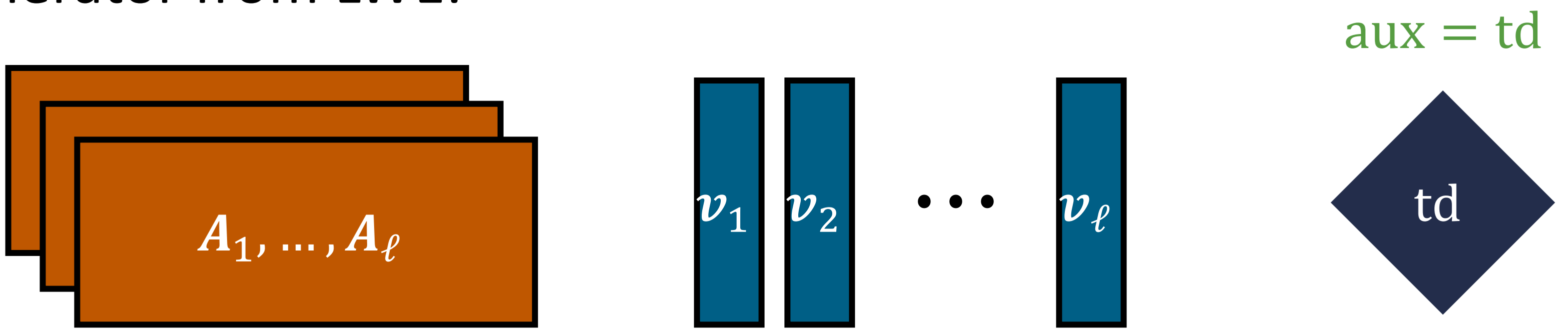

ℓ: length of hidden-bits string

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

openings are low-norm vectors $\boldsymbol{\pi}_i$ where A

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 (use $\boldsymbol{v}_i$ to extract entropy from $\boldsymbol{\pi}_i$), that $\boldsymbol{v}_i^{\mathrm{T}} \boldsymbol{\pi}_i$ is uniform

Dual-Mode Hidden-Bits Generators

Our dual-mode hidden-bits generator from LWE:

common reference string:

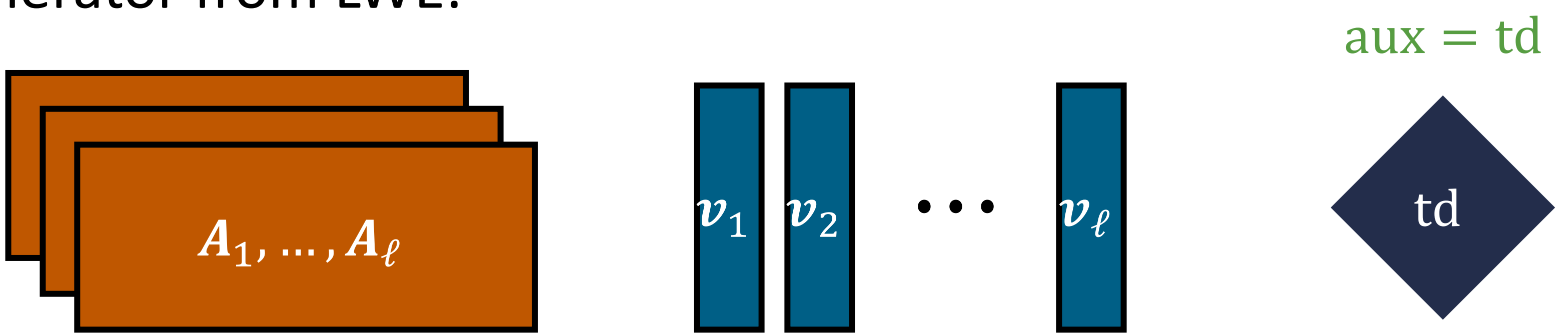

ℓ: length of hidden-bits string

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

openings are low-norm vectors $\boldsymbol{\pi}_i$ where $\boldsymbol{A}_i \boldsymbol{\pi}_i = \boldsymbol{c}$ (via td for 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 $\mathrm{td}, \boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$)

(by hybrid argument)

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 low-norm $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ 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 $\boldsymbol{c} = -\boldsymbol{G}\hat{\boldsymbol{c}}$

Limitation: trapdoor for $\boldsymbol{D}_\ell$ is a **structured** matrix (and size ℓ^2)

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 low-norm $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ 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

$\boldsymbol{B} =$	$\boldsymbol{B}_1$	$\boldsymbol{B}_2$	$\cdots$	$\boldsymbol{B}_\ell$
$\boldsymbol{A}_1 =$	$\boldsymbol{B}_1 - 0 \cdot \boldsymbol{G}$	$\boldsymbol{B}_2 - 0 \cdot \boldsymbol{G}$	$\cdots$	$\boldsymbol{B}_\ell - 1 \cdot \boldsymbol{G}$

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 low-norm $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ 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

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

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

Claim: The matrix $\boldsymbol{D}_\ell$ has a **public** trapdoor

Constructing a Shifted Multi-Preimage Sampler

Sample $\boldsymbol{B} \leftarrow \mathbb{Z}_q^{n \times m\lceil \log \ell \rceil}$ and set

$\boldsymbol{u}_i$: binary representation of i

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

Claim: The matrix $\boldsymbol{D}_\ell$ has a **public** trapdoor

Idea: Define low-norm matrices $\boldsymbol{H}_{i,j}$ and $\boldsymbol{V}_i$ where

$$(\boldsymbol{B} - \boldsymbol{u}_i^{\mathrm{T}} \otimes \boldsymbol{G})\boldsymbol{H}_{i,j} = \begin{cases} -\boldsymbol{V}_j + \boldsymbol{G} & i = j \\ -\boldsymbol{V}_j & i \neq j \end{cases}$$

Matrices $\boldsymbol{H}_{i,j}$ and $\boldsymbol{V}_i$ can be **publicly** derived using homomorphic evaluation techniques from [GSW13, BGGHNSVV14]

$(i,j)^{\text{th}}$ block of product:

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

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

Constructing a Shifted Multi-Preimage Sampler

$$\overbrace{\begin{bmatrix} \boldsymbol{B} - \boldsymbol{u}_1^{\mathrm{T}} \otimes \boldsymbol{G} & & & \boldsymbol{G} \\ & \ddots & & \vdots \\ & & \boldsymbol{B} - \boldsymbol{u}_\ell^{\mathrm{T}} \otimes \boldsymbol{G} & \boldsymbol{G} \end{bmatrix}}^{\boldsymbol{D}_\ell} \times \overbrace{\begin{bmatrix} \boldsymbol{H}_{1,1} & \cdots & \boldsymbol{H}_{1,\ell} \\ \vdots & \ddots & \vdots \\ \boldsymbol{H}_{\ell,1} & \cdots & \boldsymbol{H}_{\ell,\ell} \\ \boldsymbol{G}^{-1}(\boldsymbol{V}_1) & \cdots & \boldsymbol{G}^{-1}(\boldsymbol{V}_\ell) \end{bmatrix}}^{\text{trapdoor for } \boldsymbol{D}_\ell} = \begin{bmatrix} \boldsymbol{G} & & \\ & \ddots & \\ & & \boldsymbol{G} \end{bmatrix}$$

Letting $\boldsymbol{A}_i = \boldsymbol{B} - \boldsymbol{u}_i \otimes \boldsymbol{G}$, the trapdoor for $\boldsymbol{D}_\ell$ gives a solution to the shifted multi-preimage sampling problem for matrices $\boldsymbol{A}_1, \dots, \boldsymbol{A}_\ell$:

Observe: Matrix $\boldsymbol{D}_\ell$ and its trapdoor is completely specified by matrix $\boldsymbol{B}$ (recall that $\boldsymbol{u}_i$ is vector corresponding to binary representation of i)

Yields vector commitments and dual-mode NIZKs with transparent setup

Summary

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 low-norm* $\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ *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 solve the shifted multi-preimage sampling
- $(\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell, \mathrm{td})$ 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 via the hidden-bits model with polynomial modulus, CRS size linear in the length of the hidden-bits string, and transparent setup in statistical ZK mode
- *Subsequent work [BLNWW24]:* statistical ZAP argument from LWE via the hidden-bits approach

Summary

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 low-norm $\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_\ell \in \mathbb{Z}_q^m$ where $\boldsymbol{A}_i\boldsymbol{\pi}_i = \boldsymbol{t}_i + \boldsymbol{c}$ for all $i \in [\ell]$

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

Thank you!

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