

Hintless Single-Server PIR

Baiyu Li ¹ Daniele Micciancio ² Mariana Raykova ¹
Mark Schultz-Wu ²³

¹Google

²University of California San Diego

19 August 2024

³Work performed at Google

Contents

- 1 PIR
- 2 Composable Preprocessing
- 3 Evaluation

Single-Server PIR

Private Information Retrieval

Client ($i \in [m]$)

Server ($db \in \Omega^m$)

$qu \leftarrow \text{Query}(i)$

qu



$rsp \leftarrow \text{Response}(qu, db)$

rsp



$db_* \leftarrow \text{Recover}(rsp, i)$

Single-Server PIR

Private Information Retrieval

Client ($i \in [m]$)

Server ($db \in \Omega^m$)

$qu \leftarrow \text{Query}(i)$

qu



$rsp \leftarrow \text{Response}(qu, db)$

rsp



$db_* \leftarrow \text{Recover}(rsp, i)$

- **Correct:** $db_* = db[i]$

Single-Server PIR

Private Information Retrieval

Client ($i \in [m]$)Server ($db \in \Omega^m$) $qu \leftarrow \text{Query}(i)$ qu  $rsp \leftarrow \text{Response}(qu, db)$ rsp  $db_* \leftarrow \text{Recover}(rsp, i)$

- **Correct:** $db_* = db[i]$
- **Secure:** $\forall i, j \in [m], \{\text{Query}(i)\} \approx_c \{\text{Query}(j)\}$

Single Server PIR

- **Trivial:** Server transmits db to the Client

Single Server PIR

- **Trivial:** Server transmits db to the Client
 - Info-theoretically secure + correct

Single Server PIR

- **Trivial:** Server transmits db to the Client
 - Info-theoretically secure + correct
 - Bandwidth $O(|db|)$, compute $O(|db|)$

Single Server PIR

- **Trivial:** Server transmits db to the Client
 - Info-theoretically secure + correct
 - Bandwidth $O(|db|)$, compute $O(|db|)$
- **HE-based:** Client sends $\text{Enc}(i)$ (and maybe e_k), server hom. computes $i \mapsto db[i]$
 - Secure+Correct if HE is (IND-CPA) Secure+Correct

Single Server PIR

- **Trivial:** Server transmits db to the Client
 - Info-theoretically secure + correct
 - Bandwidth $O(|db|)$, compute $O(|db|)$
- **HE-based:** Client sends $\text{Enc}(i)$ (and maybe e_k), server hom. computes $i \mapsto \text{db}[i]$
 - Secure+Correct if HE is (IND-CPA) Secure+Correct
- **Metrics to Optimize:**
 - Bandwidth

Single Server PIR

- **Trivial:** Server transmits db to the Client
 - Info-theoretically secure + correct
 - Bandwidth $O(|db|)$, compute $O(|db|)$
- **HE-based:** Client sends $\text{Enc}(i)$ (and maybe e_k), server hom. computes $i \mapsto db[i]$
 - Secure+Correct if HE is (IND-CPA) Secure+Correct
- **Metrics to Optimize:**
 - Bandwidth
 - Server Compute

Single Server PIR

- **Trivial:** Server transmits db to the Client
 - Info-theoretically secure + correct
 - Bandwidth $O(|db|)$, compute $O(|db|)$
- **HE-based:** Client sends $\text{Enc}(i)$ (and maybe e_k), server hom. computes $i \mapsto db[i]$
 - Secure+Correct if HE is (IND-CPA) Secure+Correct
- **Metrics to Optimize:**
 - Bandwidth
 - Server Compute
 - Preprocessing Costs

PIR from HE: SOTA

- **SimplePIR** and **DoublePIR** [HHCMV Usenix23]
 - **Pros**: Extremely fast (online) server processing
 - **Cons**: Large queries, require a database-dependent “hint” (which requires **expensive** server pre-processing to generate)

PIR from HE: SOTA

- **SimplePIR** and **DoublePIR** [HHCMV Usenix23]
 - **Pros**: Extremely fast (online) server processing
 - **Cons**: Large queries, require a database-dependent “hint” (which requires **expensive** server pre-processing to generate)
 - $2^{20} \times 256\text{B}$ (286MB) Database has hint $1/3$ DB size (**SimplePIR**), and $25 \times$ DB size (**DoublePIR**)

PIR from HE: SOTA

- **SimplePIR** and **DoublePIR** [HHCMV Usenix23]
 - **Pros**: Extremely fast (online) server processing
 - **Cons**: Large queries, require a database-dependent “hint” (which requires **expensive** server pre-processing to generate)
 - $2^{20} \times 256\text{B}$ (286MB) Database has hint $1/3$ DB size (**SimplePIR**), and $25 \times$ DB size (**DoublePIR**)

PIR from HE: SOTA

- **SimplePIR** and **DoublePIR** [HHCMV Usenix23]
 - **Pros:** Extremely fast (online) server processing
 - **Cons:** Large queries, require a database-dependent “hint” (which requires **expensive** server pre-processing to generate)
 - $2^{20} \times 256\text{B}$ (286MB) Database has hint $1/3$ DB size (**SimplePIR**), and $25 \times$ DB size (**DoublePIR**)
- **Spiral** [MW SSP22]
 - **Pros:** Small queries
 - **Cons:** $\approx 10 \times$ slower, moderately-sized evaluation keys

PIR from HE: SOTA

- **SimplePIR** and **DoublePIR** [HHCMV Usenix23]
 - **Pros:** Extremely fast (online) server processing
 - **Cons:** Large queries, require a database-dependent “hint” (which requires **expensive** server pre-processing to generate)
 - $2^{20} \times 256\text{B}$ (286MB) Database has hint $1/3$ DB size (**SimplePIR**), and $25 \times$ DB size (**DoublePIR**)
- **Spiral** [MW SSP22]
 - **Pros:** Small queries
 - **Cons:** $\approx 10 \times$ slower, moderately-sized evaluation keys
- **TiptoePIR** [HDCZ SOSP23] (Concurrent to our work)
 - **Pros:** Removes the hint from SimplePIR!
 - **Cons:** $+\Omega(n^2 \log q)$ ($\approx 23\text{MB}$) per-query bandwidth cost

LWE Encryption

- Ciphertexts: $[A, \vec{b}]$ s.t. $\vec{b} - A \cdot \vec{s} \approx \Delta \vec{m}$

LWE Encryption

- **Ciphertexts:** $[A, \vec{b}]$ s.t. $\vec{b} - A \cdot \vec{s} \approx \Delta \vec{m}$
 - $\|\vec{s}\|_\infty$ **small**,

LWE Encryption

- **Ciphertexts:** $[A, \vec{b}]$ s.t. $\vec{b} - A \cdot \vec{s} \approx \Delta \vec{m}$
 - $\|\vec{s}\|_\infty$ **small**,
 - $A \in \mathbb{Z}_q^{m \times n}$ can be **reused**

LWE Encryption

- **Ciphertexts:** $[A, \vec{b}]$ s.t. $\vec{b} - A \cdot \vec{s} \approx \Delta \vec{m}$
 - $\|\vec{s}\|_\infty$ **small**,
 - $A \in \mathbb{Z}_q^{m \times n}$ can be **reused**
 - **Easily** supports hom. computation of $\vec{x} \mapsto db \cdot \vec{x}$

LWE Encryption

- **Ciphertexts:** $[A, \vec{b}]$ s.t. $\vec{b} - A \cdot \vec{s} \approx \Delta \vec{m}$
 - $\|\vec{s}\|_\infty$ **small**,
 - $A \in \mathbb{Z}_q^{m \times n}$ can be **reused**
 - **Easily** supports hom. computation of $\vec{x} \mapsto db \cdot \vec{x}$
 - $db \cdot [A, \vec{b}]$ has **large** component $H := db \cdot A$

LWE Encryption

- **Ciphertexts:** $[A, \vec{b}]$ s.t. $\vec{b} - A \cdot \vec{s} \approx \Delta \vec{m}$
 - $\|\vec{s}\|_\infty$ **small**,
 - $A \in \mathbb{Z}_q^{m \times n}$ can be **reused**
 - **Easily** supports hom. computation of $\vec{x} \mapsto db \cdot \vec{x}$
 - $db \cdot [A, \vec{b}]$ has **large** component $H := db \cdot A$
 - **General** moduli q

LWE Encryption

- **Ciphertexts:** $[A, \vec{b}]$ s.t. $\vec{b} - A \cdot \vec{s} \approx \Delta \vec{m}$
 - $\|\vec{s}\|_\infty$ **small**,
 - $A \in \mathbb{Z}_q^{m \times n}$ can be **reused**
 - **Easily** supports hom. computation of $\vec{x} \mapsto db \cdot \vec{x}$
 - $db \cdot [A, \vec{b}]$ has **large** component $H := db \cdot A$
 - **General** moduli $q \in \{2^{32}, 2^{64}\}$

SimplePIR's Hint

- SimplePIR (and DoublePIR) requires clients to download a **hint** $H := db \cdot A$

SimplePIR's Hint

- SimplePIR (and DoublePIR) requires clients to download a **hint** $H := db \cdot A$
- Clients only use this to compute $\vec{s} \mapsto H \cdot \vec{s}$ for LWE key \vec{s}

SimplePIR's Hint

- SimplePIR (and DoublePIR) requires clients to download a **hint** $H := db \cdot A$
- Clients only use this to compute $\vec{s} \mapsto H \cdot \vec{s}$ for LWE key \vec{s}
 - **Almost** a PIR query, which is solved trivially

SimplePIR's Hint

- SimplePIR (and DoublePIR) requires clients to download a **hint** $H := db \cdot A$
- Clients only use this to compute $\vec{s} \mapsto H \cdot \vec{s}$ for LWE key \vec{s}
 - **Almost** a PIR query, which is solved trivially
- **Our Work:** Use non-trivial (albeit technically straightforward) RLWE-based PIR instead

SimplePIR's Hint

- SimplePIR (and DoublePIR) requires clients to download a **hint** $H := db \cdot A$
- Clients only use this to compute $\vec{s} \mapsto H \cdot \vec{s}$ for LWE key \vec{s}
 - **Almost** a PIR query, which is solved trivially
- **Our Work**: Use non-trivial (albeit technically straightforward) RLWE-based PIR instead
 - Speed up with novel **precomputation** argument

SimplePIR's Hint

- SimplePIR (and DoublePIR) requires clients to download a **hint** $H := db \cdot A$
- Clients only use this to compute $\vec{s} \mapsto H \cdot \vec{s}$ for LWE key \vec{s}
 - **Almost** a PIR query, which is solved trivially
- **Our Work**: Use non-trivial (albeit technically straightforward) RLWE-based PIR instead
 - Speed up with novel **precomputation** argument
 - Also a PIR with $O(\sqrt[3]{|db|})$ bandwidth scaling (TensorPIR)

Why Keep SimplePIR?

- $T_{\text{SimplePIR}}(|\text{db}|) = 2|\text{db}|$ operations in \mathbb{Z}_q

Why Keep SimplePIR?

- $T_{\text{SimplePIR}}(|\text{db}|) = 2|\text{db}|$ operations in \mathbb{Z}_q
- $T_{\text{NTTlessPIR}}(|\text{db}|) = O(|\text{db}|)$ operations in \mathbb{Z}_{q^*}

Why Keep SimplePIR?

- $T_{\text{SimplePIR}}(|\text{db}|) = 2|\text{db}|$ operations in \mathbb{Z}_q
- $T_{\text{NTTlessPIR}}(|\text{db}|) = O(|\text{db}|)$ operations in \mathbb{Z}_{q^*}
 - for $q^* \notin \{2^{32}, 2^{64}\}$

Why Keep SimplePIR?

- $T_{\text{SimplePIR}}(|\text{db}|) = 2|\text{db}|$ operations in \mathbb{Z}_q
- $T_{\text{NTTlessPIR}}(|\text{db}|) = O(|\text{db}|)$ operations in \mathbb{Z}_{q^*}
 - for $q^* \notin \{2^{32}, 2^{64}\}$
 - $\approx 10\times$ slower than $q \in \{2^{32}, 2^{64}\}$

Why Keep SimplePIR?

- $T_{\text{SimplePIR}}(|\text{db}|) = 2|\text{db}|$ operations in \mathbb{Z}_q
- $T_{\text{NTTlessPIR}}(|\text{db}|) = O(|\text{db}|)$ operations in \mathbb{Z}_{q^*}
 - for $q^* \notin \{2^{32}, 2^{64}\}$
 - $\approx 10\times$ slower than $q \in \{2^{32}, 2^{64}\}$
- $T_{\text{HintlessPIR}}(|\text{db}|) = T_{\text{SimplePIR}}(|\text{db}|) + T_{\text{NTTlessPIR}}(n\sqrt{|\text{db}|})$

Why Keep SimplePIR?

- $T_{\text{SimplePIR}}(|\text{db}|) = 2|\text{db}|$ operations in \mathbb{Z}_q
- $T_{\text{NTTlessPIR}}(|\text{db}|) = O(|\text{db}|)$ operations in \mathbb{Z}_{q^*}
 - for $q^* \notin \{2^{32}, 2^{64}\}$
 - $\approx 10\times$ slower than $q \in \{2^{32}, 2^{64}\}$
- $T_{\text{HintlessPIR}}(|\text{db}|) = T_{\text{SimplePIR}}(|\text{db}|) + T_{\text{NTTlessPIR}}(n\sqrt{|\text{db}|})$
 - $2|\text{db}|$ operations in \mathbb{Z}_q , and $O(n\sqrt{|\text{db}|})$ operations in \mathbb{Z}_{q^*} .

Column-Major Matrix-Vector Multiplication

$$\begin{pmatrix} \vdots & \vdots & \dots & \vdots \\ \vec{a}_1 & \vec{a}_2 & \dots & \vec{a}_n \\ \vdots & \vdots & \dots & \vdots \end{pmatrix} \cdot \begin{pmatrix} \vec{s}_1 \\ \vec{s}_2 \\ \vdots \\ \vdots \\ \vec{s}_n \end{pmatrix} = \sum_{i \in [n]} \vec{s}_i \cdot \begin{pmatrix} \vdots \\ \vec{a}_i \\ \vdots \end{pmatrix}$$

Column-Major Matrix-Vector Multiplication

$$\begin{pmatrix} \vdots & \vdots & \dots & \vdots \\ \vec{a}_1 & \vec{a}_2 & \dots & \vec{a}_n \\ \vdots & \vdots & \dots & \vdots \end{pmatrix} \cdot \begin{pmatrix} \vec{s}_1 \\ \vec{s}_2 \\ \vdots \\ \vdots \\ \vec{s}_n \end{pmatrix} = \sum_{i \in [n]} \vec{s}_i \cdot \begin{pmatrix} \vdots \\ \vdots \\ \vdots \\ \vdots \\ \vdots \end{pmatrix}$$

- Can **encrypt** $\{\vec{s}_i\}_i$, compute above with linearly homomorphic encryption (e.g. RLWE)

Column-Major Matrix-Vector Multiplication

$$\begin{pmatrix} \vdots & \vdots & \dots & \vdots \\ \vec{a}_1 & \vec{a}_2 & \dots & \vec{a}_n \\ \vdots & \vdots & \dots & \vdots \end{pmatrix} \cdot \begin{pmatrix} \vec{s}_1 \\ \vec{s}_2 \\ \vdots \\ \vdots \\ \vec{s}_n \end{pmatrix} = \sum_{i \in [n]} \vec{s}_i \cdot \begin{pmatrix} \vdots \\ \vdots \\ \vdots \\ \vdots \\ \vdots \end{pmatrix}$$

- Can **encrypt** $\{\vec{s}_i\}_i$, compute above with linearly homomorphic encryption (e.g. RLWE)
 - What **TiptoePIR** does

Column-Major Matrix-Vector Multiplication

$$\begin{pmatrix} \vdots & \vdots & \dots & \vdots \\ \vec{a}_1 & \vec{a}_2 & \dots & \vec{a}_n \\ \vdots & \vdots & \dots & \vdots \end{pmatrix} \cdot \begin{pmatrix} \vec{s}_1 \\ \vec{s}_2 \\ \vdots \\ \vdots \\ \vec{s}_n \end{pmatrix} = \sum_{i \in [n]} \vec{s}_i \cdot \begin{pmatrix} \vdots \\ \vec{a}_i \\ \vdots \end{pmatrix}$$

- Can **encrypt** $\{\vec{s}_i\}_i$, compute above with linearly homomorphic encryption (e.g. RLWE)
 - What **TiptoePIR** does
- **Con**: Requires n RLWE ciphertexts, $\Theta(n^2 \log q)$ bandwidth

Diagonally-Dominant Matrix-Vector Multiplication

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix} \cdot \begin{pmatrix} \vec{s}_1 \\ \vec{s}_2 \\ \vec{s}_3 \end{pmatrix}$$

Diagonally-Dominant Matrix-Vector Multiplication

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix} \cdot \begin{pmatrix} \vec{s}_1 \\ \vec{s}_2 \\ \vec{s}_3 \end{pmatrix} = \begin{pmatrix} a_{11}\vec{s}_1 + a_{12}\vec{s}_2 + a_{13}\vec{s}_3 \\ a_{21}\vec{s}_1 + a_{22}\vec{s}_2 + a_{23}\vec{s}_3 \\ a_{31}\vec{s}_1 + a_{32}\vec{s}_2 + a_{33}\vec{s}_3 \end{pmatrix}$$

Diagonally-Dominant Matrix-Vector Multiplication

$$\begin{aligned}
 \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix} \cdot \begin{pmatrix} \vec{s}_1 \\ \vec{s}_2 \\ \vec{s}_3 \end{pmatrix} &= \begin{pmatrix} a_{11}\vec{s}_1 + a_{12}\vec{s}_2 + a_{13}\vec{s}_3 \\ a_{21}\vec{s}_1 + a_{22}\vec{s}_2 + a_{23}\vec{s}_3 \\ a_{31}\vec{s}_1 + a_{32}\vec{s}_2 + a_{33}\vec{s}_3 \end{pmatrix} \\
 &= \begin{pmatrix} a_{11}\vec{s}_1 + a_{12}\vec{s}_2 + a_{13}\vec{s}_3 \\ a_{22}\vec{s}_2 + a_{23}\vec{s}_3 + a_{21}\vec{s}_1 \\ a_{33}\vec{s}_3 + a_{31}\vec{s}_1 + a_{32}\vec{s}_2 \end{pmatrix}
 \end{aligned}$$

Diagonally-Dominant Matrix-Vector Multiplication

$$\begin{aligned}
 \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix} \cdot \begin{pmatrix} \vec{s}_1 \\ \vec{s}_2 \\ \vec{s}_3 \end{pmatrix} &= \begin{pmatrix} a_{11}\vec{s}_1 + a_{12}\vec{s}_2 + a_{13}\vec{s}_3 \\ a_{21}\vec{s}_1 + a_{22}\vec{s}_2 + a_{23}\vec{s}_3 \\ a_{31}\vec{s}_1 + a_{32}\vec{s}_2 + a_{33}\vec{s}_3 \end{pmatrix} \\
 &= \begin{pmatrix} a_{11}\vec{s}_1 + a_{12}\vec{s}_2 + a_{13}\vec{s}_3 \\ a_{22}\vec{s}_2 + a_{23}\vec{s}_3 + a_{21}\vec{s}_1 \\ a_{33}\vec{s}_3 + a_{31}\vec{s}_1 + a_{32}\vec{s}_2 \end{pmatrix} \\
 &= \sum_i \text{diag}_i(A) \circ \text{rot}^i(\vec{s})
 \end{aligned}$$

Diagonally-Dominant Matrix-Vector Multiplication

$$\begin{aligned}
 \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix} \cdot \begin{pmatrix} \vec{s}_1 \\ \vec{s}_2 \\ \vec{s}_3 \end{pmatrix} &= \begin{pmatrix} a_{11}\vec{s}_1 + a_{12}\vec{s}_2 + a_{13}\vec{s}_3 \\ a_{21}\vec{s}_1 + a_{22}\vec{s}_2 + a_{23}\vec{s}_3 \\ a_{31}\vec{s}_1 + a_{32}\vec{s}_2 + a_{33}\vec{s}_3 \end{pmatrix} \\
 &= \begin{pmatrix} a_{11}\vec{s}_1 + a_{12}\vec{s}_2 + a_{13}\vec{s}_3 \\ a_{22}\vec{s}_2 + a_{23}\vec{s}_3 + a_{21}\vec{s}_1 \\ a_{33}\vec{s}_3 + a_{31}\vec{s}_1 + a_{32}\vec{s}_2 \end{pmatrix} \\
 &= \sum_i \text{diag}_i(A) \circ \text{rot}^i(\vec{s})
 \end{aligned}$$

- **Small** overhead if one can hom. evaluate $\vec{s} \mapsto \text{rot}(\vec{s})$

Diagonally-Dominant Matrix-Vector Multiplication

$$\begin{aligned}
 \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix} \cdot \begin{pmatrix} \vec{s}_1 \\ \vec{s}_2 \\ \vec{s}_3 \end{pmatrix} &= \begin{pmatrix} a_{11}\vec{s}_1 + a_{12}\vec{s}_2 + a_{13}\vec{s}_3 \\ a_{21}\vec{s}_1 + a_{22}\vec{s}_2 + a_{23}\vec{s}_3 \\ a_{31}\vec{s}_1 + a_{32}\vec{s}_2 + a_{33}\vec{s}_3 \end{pmatrix} \\
 &= \begin{pmatrix} a_{11}\vec{s}_1 + a_{12}\vec{s}_2 + a_{13}\vec{s}_3 \\ a_{22}\vec{s}_2 + a_{23}\vec{s}_3 + a_{21}\vec{s}_1 \\ a_{33}\vec{s}_3 + a_{31}\vec{s}_1 + a_{32}\vec{s}_2 \end{pmatrix} \\
 &= \sum_i \text{diag}_i(A) \circ \text{rot}^i(\vec{s})
 \end{aligned}$$

- **Small** overhead if one can hom. evaluate $\vec{s} \mapsto \text{rot}(\vec{s})$
 - **Our Work**: Use **precomputation** to make this **fast**

Contents

- 1 PIR
- 2 Composable Preprocessing
- 3 Evaluation

RLWE Notation

- $R_{n,q} := \mathbb{Z}_q[x]/(x^n + 1)$ for $n = 2^k$

RLWE Notation

- $R_{n,q} := \mathbb{Z}_q[x]/(x^n + 1)$ for $n = 2^k$
- For $q \equiv 1 \pmod{2n}$ (NTT friendly q), exists a ring isomorphism

$$\text{NTT} : (R_{n,q}, +, *) \rightarrow (\mathbb{Z}_q^n, +, \circ)$$

computable in $O(n \log n)$ \mathbb{Z}_q -ops

RLWE Notation

- $R_{n,q} := \mathbb{Z}_q[x]/(x^n + 1)$ for $n = 2^k$
- For $q \equiv 1 \pmod{2n}$ (NTT friendly q), exists a ring isomorphism

$$\text{NTT} : (R_{n,q}, +, *) \rightarrow (\mathbb{Z}_q^n, +, \circ)$$

computable in $O(n \log n)$ \mathbb{Z}_q -ops

- Use standard RLWE (and gadget RLWE) encryption with plaintext packing

◇-Product: Standard Domain

◇-Product

Let $\vec{ct} \in (R_{n,q} \times R_{n,q})^\ell$ be a collection of ℓ RLWE ciphertexts. For a polynomial $a \in R_{n,q}$, define

$$a \diamond \vec{ct} = \sum_{i \in [\ell]} \text{GadgetDecompose}(a)_i * ct_i$$

◇-Product: Standard Domain

◇-Product

Let $\vec{ct} \in (R_{n,q} \times R_{n,q})^\ell$ be a collection of ℓ RLWE ciphertexts. For a polynomial $a \in R_{n,q}$, define

$$a \diamond \vec{ct} = \sum_{i \in [\ell]} \text{GadgetDecompose}(a)_i * ct_i$$

- **Fundamental** Operation for RLWE-based HE

◇-Product: Standard Domain

◇-Product

Let $\vec{ct} \in (R_{n,q} \times R_{n,q})^\ell$ be a collection of ℓ RLWE ciphertexts. For a polynomial $a \in R_{n,q}$, define

$$a \diamond \vec{ct} = \sum_{i \in [\ell]} \text{GadgetDecompose}(a)_i * ct_i$$

- **Fundamental** Operation for RLWE-based HE
 - Main part of gadget-based key-switching

◇-Product: Standard Domain

◇-Product

Let $\vec{ct} \in (R_{n,q} \times R_{n,q})^\ell$ be a collection of ℓ RLWE ciphertexts. For a polynomial $a \in R_{n,q}$, define

$$a \diamond \vec{ct} = \sum_{i \in [\ell]} \text{GadgetDecompose}(a)_i * ct_i$$

- **Fundamental** Operation for RLWE-based HE
 - Main part of gadget-based key-switching
 - And GSW/FHEW/TFHE/etc.

◇-Product: Standard Domain

◇-Product

Let $\vec{ct} \in (R_{n,q} \times R_{n,q})^\ell$ be a collection of ℓ RLWE ciphertexts. For a polynomial $a \in R_{n,q}$, define

$$a \diamond \vec{ct} = \sum_{i \in [\ell]} \text{GadgetDecompose}(a)_i * ct_i$$

- **Fundamental** Operation for RLWE-based HE
 - Main part of gadget-based key-switching
 - And GSW/FHEW/TFHE/etc.
 - Implies hom. rotation $\vec{s} \mapsto \text{rot}(\vec{s})$

◇-Product: Standard Domain

◇-Product

Let $\vec{ct} \in (R_{n,q} \times R_{n,q})^\ell$ be a collection of ℓ RLWE ciphertexts. For a polynomial $a \in R_{n,q}$, define

$$a \diamond \vec{ct} = \sum_{i \in [\ell]} \text{GadgetDecompose}(a)_i * \vec{ct}_i$$

- **Fundamental** Operation for RLWE-based HE
 - Main part of gadget-based key-switching
 - And GSW/FHEW/TFHE/etc.
 - Implies hom. rotation $\vec{s} \mapsto \text{rot}(\vec{s})$
- As written, is $O(n^2 \log q)$ time due to $*$

◇-Product: NTT Domain

- If we assume a, \vec{ct} start in NTT domain $(\widehat{a}, \widehat{\vec{ct}})$, and want result in NTT domain, instead must compute

$$\widehat{a \diamond \vec{ct}} = \sum_{i \in [\ell]} \text{NTT}(\text{GadgetDecompose}(\text{NTT}^{-1}(\widehat{a}))_i) \circ \widehat{\vec{ct}}_i$$

◇-Product: NTT Domain

- If we assume a, \vec{ct} start in NTT domain $(\widehat{a}, \widehat{\vec{ct}})$, and want result in NTT domain, instead must compute

$$\widehat{a \diamond \vec{ct}} = \sum_{i \in [\ell]} \text{NTT}(\text{GadgetDecompose}(\text{NTT}^{-1}(\widehat{a}))_i) \circ \widehat{\vec{ct}}_i$$

- Have to switch back to coefficient domain to compute the coefficients

$$\{\text{NTT}(\text{GadgetDecompose}(\text{NTT}^{-1}(\widehat{a}))_i)\}_{i \in [\ell]}$$

◇-Product: NTT Domain

- If we assume a, \vec{ct} start in NTT domain $(\widehat{a}, \widehat{\vec{ct}})$, and want result in NTT domain, instead must compute

$$\widehat{a \diamond \vec{ct}} = \sum_{i \in [\ell]} \text{NTT}(\text{GadgetDecompose}(\text{NTT}^{-1}(\widehat{a}))_i) \circ \widehat{\vec{ct}}_i$$

- Have to switch back to coefficient domain to compute the coefficients

$$\{\text{NTT}(\text{GadgetDecompose}(\text{NTT}^{-1}(\widehat{a}))_i)\}_{i \in [\ell]}$$

- Precompute?

Composable Preprocessing

- For RLWE ctxt $ct = [a, b]$, define

$$\alpha(ct) = a, \quad \beta(ct) = b$$

Composable Preprocessing

- For RLWE ctxt $ct = [a, b]$, define

$$\alpha(ct) = a, \quad \beta(ct) = b$$

- These are **homomorphic** with respect $\{+, *\}$, e.g.

$$\alpha(ct_0 + ct_1) = \alpha(ct_0) + \alpha(ct_1)$$

Composable Preprocessing

- For RLWE ctxt $ct = [a, b]$, define

$$\alpha(ct) = a, \quad \beta(ct) = b$$

- These are **homomorphic** with respect $\{+, *\}$, e.g.

$$\alpha(ct_0 + ct_1) = \alpha(ct_0) + \alpha(ct_1)$$

- More interestingly, $\alpha(c \diamond \vec{ct}) = c \diamond \alpha(\vec{ct})$ is also homomorphic

Composable Preprocessing

- For RLWE ctxt $ct = [a, b]$, define

$$\alpha(ct) = a, \quad \beta(ct) = b$$

- These are **homomorphic** with respect $\{+, *\}$, e.g.

$$\alpha(ct_0 + ct_1) = \alpha(ct_0) + \alpha(ct_1)$$

- More interestingly, $\alpha(c \diamond \vec{ct}) = c \diamond \alpha(\vec{ct})$ is also homomorphic
 - **Precompute** $\alpha(\text{Eval}_{\text{ek}}(C, \{ct_i\}_i))$ for C built from $\{+, *, \diamond\}$
($2\times$ speedup)

Composable Preprocessing

- For RLWE ctxt $ct = [a, b]$, define

$$\alpha(ct) = a, \quad \beta(ct) = b$$

- These are **homomorphic** with respect $\{+, *\}$, e.g.

$$\alpha(ct_0 + ct_1) = \alpha(ct_0) + \alpha(ct_1)$$

- More interestingly, $\alpha(c \diamond \vec{ct}) = c \diamond \alpha(\vec{ct})$ is also homomorphic
 - **Precompute** $\alpha(\text{Eval}_{\text{ek}}(C, \{ct_i\}_i))$ for C built from $\{+, *, \diamond\}$
($2\times$ speedup)
- Hom. rotations only call GadgetDecompose on $\alpha(ct)$

Composable Preprocessing

- For RLWE ctxt $ct = [a, b]$, define

$$\alpha(ct) = a, \quad \beta(ct) = b$$

- These are **homomorphic** with respect $\{+, *\}$, e.g.

$$\alpha(ct_0 + ct_1) = \alpha(ct_0) + \alpha(ct_1)$$

- More interestingly, $\alpha(c \diamond \vec{ct}) = c \diamond \alpha(\vec{ct})$ is also homomorphic
 - **Precompute** $\alpha(\text{Eval}_{\text{ek}}(C, \{ct_i\}_i))$ for C built from $\{+, *, \diamond\}$
($2\times$ speedup)
- Hom. rotations only call GadgetDecompose on $\alpha(ct)$
 - Precompute all NTTs ($O(\log n)\times$ speedup)

NTTlessPIR

- Use (standard) Diagonally-Dominant matrix-vector multiplication algorithm

NTTlessPIR

- Use (standard) Diagonally-Dominant matrix-vector multiplication algorithm
 - With client-side CRT interpolation to handle NTT-unfriendly $q \in \{2^{32}, 2^{64}\}$

NTTlessPIR

- Use (standard) Diagonally-Dominant matrix-vector multiplication algorithm
 - With client-side CRT interpolation to handle NTT-unfriendly $q \in \{2^{32}, 2^{64}\}$
- Apply composable preprocessing to preprocess all NTTs that occur

NTTlessPIR

- Use (standard) Diagonally-Dominant matrix-vector multiplication algorithm
 - With client-side CRT interpolation to handle NTT-unfriendly $q \in \{2^{32}, 2^{64}\}$
- Apply composable preprocessing to preprocess all NTTs that occur
- **Net Result:** Can hom. compute $(A, \text{Enc}_{\vec{s}}(\vec{x})) \mapsto \text{Enc}_{\vec{s}}(A \cdot \vec{x})$ within a multiplicative constant of the cost in plaintext

Composable Preprocessing Limitations

- Doesn't seem to allow hom. multiplication

Composable Preprocessing Limitations

- Doesn't seem to allow hom. multiplication
 - Can non-composably preprocess algorithms still

Composable Preprocessing Limitations

- Doesn't seem to allow hom. multiplication
 - Can non-composably preprocess algorithms still (**TensorPIR**)

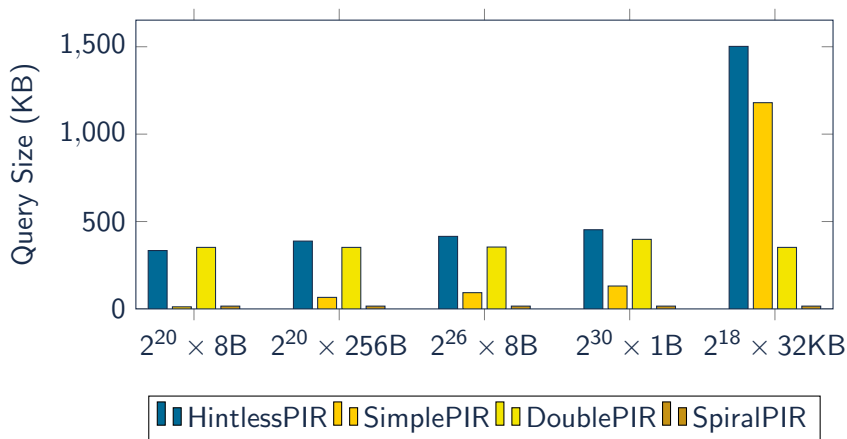
Contents

- 1 PIR
- 2 Composable Preprocessing
- 3 Evaluation

Evaluation

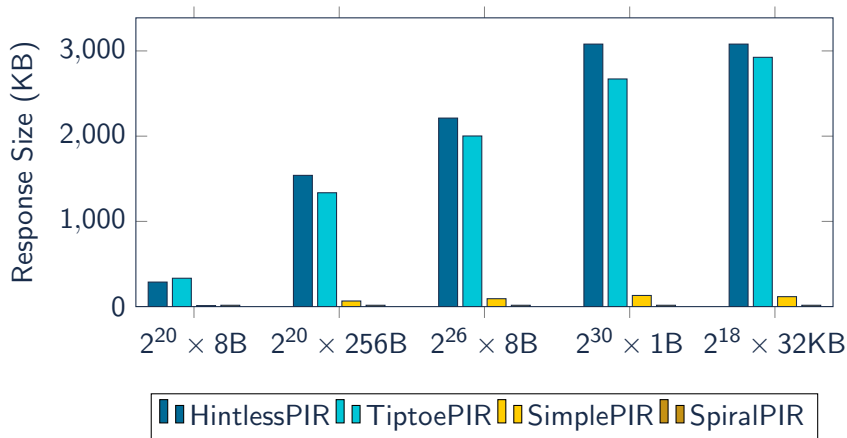
- Implemented our protocols
 - https://github.com/google/hintless_pir
- Compared HintlessPIR to other HE-based PIR
 - SimplePIR, DoublePIR, TiptoePIR, SpiralPIR

Bandwidth: Query Size



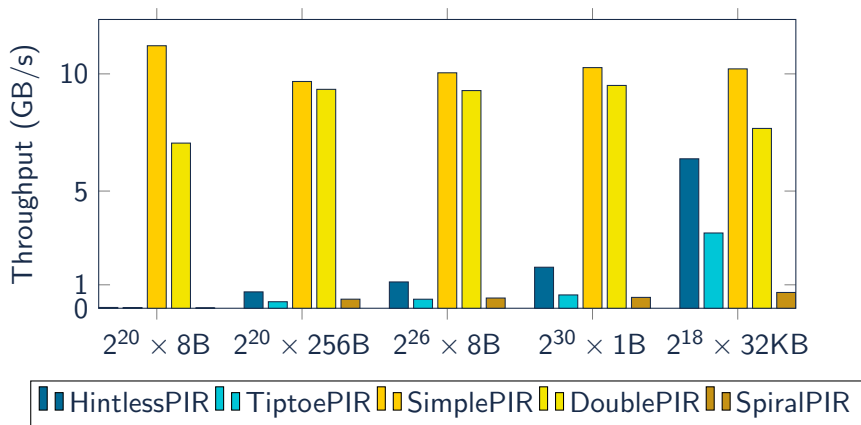
- TiptoePIR: $\approx 23MB$ for all database sizes

Bandwidth: Response Size



- DoublePIR's response depends only on record size. Ranges from 44KB (1B records) to 1273MB (32KB records)

Throughput



Hint Size

Database	SimplePIR	DoublePIR
$2^{20} \times 8\text{B}$ (8MB)	200% (16MB)	3025% (242MB)
$2^{20} \times 256\text{B}$ (286MB)	34% (92MB)	2573% (6897MB)
$2^{26} \times 8\text{B}$ (537MB)	24% (131MB)	45% (242MB)
$2^{30} \times 1\text{B}$ (1074MB)	17% (185MB)	2.8% (31MB)
$2^{18} \times 32\text{KB}$ (8389MB)	2% (185MB)	10418% (874000MB)