Private Information Retrieval: Coding instead of Replication

Alexander Vardy

Information Theory & Applications Center University of California San Diego 9500 Gilman Drive, La Jolla, CA 92093, USA

Jont work with:



What is private information retrieval?



Private information retrieval (PIR)

Alice wishes to retrieve a data item x_i from the database $(x_1, x_2, ..., x_n)$ without revealing any information about *i* to the server.

Formal privacy condition: The distribution of randomized queries sent by the user to the server does not depend on *i*.

What is private information retrieval?



Private information retrieval (PIR)

Alice wishes to retrieve a data item x_i from the database $(x_1, x_2, ..., x_n)$ without revealing any information about *i* to the server.

Formal privacy condition: The distribution of randomized queries sent by the user to the server does not depend on *i*. **Naive Solution:** Ask the server to send the entire database!

What is private information retrieval?



Private information retrieval (PIR)

Alice wishes to retrieve a data item x_i from the database $(x_1, x_2, ..., x_n)$ without revealing any information about *i* to the server.

Formal privacy condition: The distribution of randomized queries sent by the user to the server does not depend on *i*.



Solution: Ask the server to send the entire database!

This is the only solution possible! Communication $cost = \Omega(n)$.

Two general classes of solutions

Computational PIR



The server is computationally bounded + **standard cryptographic assumptions** (one-way functions, quadratic residuosity).

E. Kushilevitz and R. Ostrovsky, Replication is not needed: Single database, computationally-private information retrieval, *Proc. 38-th IEEE Symp. Foundations Computer Science*, pp. 364–373, October 1997.

• Information-theoretic PIR

The **database is replicated** among $k \ge 2$ non-communicating servers, with guarantees of information-theoretic privacy.



Two general classes of solutions

Computational PIR



The server is computationally bounded + **standard cryptographic assumptions** (one-way functions, quadratic residuosity).

E. Kushilevitz and R. Ostrovsky, Replication is not needed: Single database, computationally-private information retrieval, *Proc. 38-th IEEE Symp. Foundations Computer Science*, pp. 364–373, October 1997.

• Information-theoretic PIR

The **database is replicated** among $k \ge 2$ non-communicating servers, with guarantees of information-theoretic privacy.



B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan, Private information retrieval, Proc. 36-th IEEE Symposium Foundations Computer Science, pp. 41–50, October 1995.

This talk: We consider only information-theoretic PIR!

Replication among k = 4 servers S_1 , S_2 , S_3 , S_4 with communication cost of $8\sqrt{n} + 4$ bits. The database is represented as a square of side \sqrt{n} .

Replication among k = 4 servers S_1 , S_2 , S_3 , S_4 with communication cost of $8\sqrt{n} + 4$ bits. The database is represented as a square of side \sqrt{n} .

B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan, Private information retrieval, Proceedings IEEE Symp. Foundations Computer Science, pp. 41–50, October 1995.

Query generation:

Alice wishes to retrieve $x_{s,t}$. She generates the vectors $y, z \in \{0, 1\}^{\sqrt{n}}$ uniformly at random, and sends

$$\mathcal{S}_1 \leftarrow (\boldsymbol{y}, \boldsymbol{z}), \ \mathcal{S}_2 \leftarrow (\boldsymbol{y} + \boldsymbol{e}_s, \boldsymbol{z}), \ \mathcal{S}_3 \leftarrow (\boldsymbol{y}, \boldsymbol{z} + \boldsymbol{e}_t), \ \mathcal{S}_4 \leftarrow (\boldsymbol{y} + \boldsymbol{e}_s, \boldsymbol{z} + \boldsymbol{e}_t)$$



Replication among k = 4 servers S_1 , S_2 , S_3 , S_4 with communication cost of $8\sqrt{n} + 4$ bits. The database is represented as a square of side \sqrt{n} .

B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan, Private information retrieval, Proceedings IEEE Symp. Foundations Computer Science, pp. 41–50, October 1995.

Query generation:

Alice wishes to retrieve $x_{s,t}$. She generates the vectors $y, z \in \{0, 1\}^{\sqrt{n}}$ uniformly at random, and sends

$$egin{aligned} \mathcal{S}_1 \leftarrow (m{y}, m{z}), \ \mathcal{S}_2 \leftarrow (m{y} + m{e}_s, m{z}), \ \mathcal{S}_3 \leftarrow (m{y}, m{z} + m{e}_t), \ \mathcal{S}_4 \leftarrow (m{y} + m{e}_s, m{z} + m{e}_t) \end{aligned}$$





Answer computation:

Given a query (u, v), each server S_i returns the following:

$$a = \sum_{i \in \text{supp}(u)} \sum_{j \in \text{supp}(v)} x_{i,j}$$

Query generation:

 $\mathcal{S}_1 \leftarrow (y, z), \ \mathcal{S}_2 \leftarrow (y + e_s, z), \ \mathcal{S}_3 \leftarrow (y, z + e_t), \ \mathcal{S}_4 \leftarrow (y + e_s, z + e_t)$

Answer computation:

$$a = \sum_{i \in \text{supp}(u)} \sum_{j \in \text{supp}(v)} x_{i,j}$$

Reconstruction:



Query generation:

 $\mathcal{S}_1 \leftarrow (y, z), \ \mathcal{S}_2 \leftarrow (y + e_s, z), \ \mathcal{S}_3 \leftarrow (y, z + e_t), \ \mathcal{S}_4 \leftarrow (y + e_s, z + e_t)$

Answer computation: $a = \sum_{i \in \text{supp}(u)} \sum_{j \in \text{supp}(v)} x_{ij}$ Reconstruction: t t s a_1 a_2 a_3 a_4

1 The bit $x_{s,t}$ contributes to **exactly one** of the answers a_1, a_2, a_3, a_4 .

Query generation:

 $\mathcal{S}_1 \leftarrow (y, z), \ \mathcal{S}_2 \leftarrow (y + e_s, z), \ \mathcal{S}_3 \leftarrow (y, z + e_t), \ \mathcal{S}_4 \leftarrow (y + e_s, z + e_t)$

Answer computation: $a = \sum_{i \in \text{supp}(u)} \sum_{j \in \text{supp}(v)} x_{ij}$ Reconstruction:



• The bit $x_{s,t}$ contributes to **exactly one** of the answers a_1, a_2, a_3, a_4 .

2 All other bits in the database contribute an **even number** of times.

Query generation:

 $\mathcal{S}_1 \leftarrow (y, z), \ \mathcal{S}_2 \leftarrow (y + e_s, z), \ \mathcal{S}_3 \leftarrow (y, z + e_t), \ \mathcal{S}_4 \leftarrow (y + e_s, z + e_t)$

Answer computation:

$$a = \sum_{i \in \text{supp}(u)} \sum_{j \in \text{supp}(v)} x_{i,j}$$

Reconstruction:



① The bit $x_{s,t}$ contributes to **exactly one** of the answers a_1, a_2, a_3, a_4 .

2 All other bits in the database contribute an **even number** of times.

It follows from 1 and 2 that:

$$a_1 + a_2 + a_3 + a_4 = x_{s,t}$$

# of Servers	Communication Complexity	Year	Reference
2	$O(n^{\frac{1}{3}})$	1995	B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan
k	$O\left(n^{\frac{1}{k}}\right)$	1995	B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan

# of Servers	Communication Complexity	Year	Reference
2	$O(n^{\frac{1}{3}})$	1995	B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan
k	$O(n^{\frac{1}{k}})$	1995	B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan
k	$O(n^{\frac{1}{2k-1}})$	1997	A. Ambainis

# of Servers	Communication Complexity	Year	Reference
2	$O(n^{\frac{1}{3}})$	1995	B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan
k	$O(n^{\frac{1}{k}})$	1995	B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan
k	$O(n^{\frac{1}{2k-1}})$	1997	A. Ambainis
k	$O(n^{\frac{\log\log k}{k\log k}})$	2002	A. Beimel, Y. Ishai, E. Kushilevitz, and J.F. Raymond

# of Servers	Communication Complexity	Year	Reference
2	$O(n^{\frac{1}{3}})$	1995	B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan
k	$O(n^{\frac{1}{k}})$	1995	B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan
k	$O(n^{\frac{1}{2k-1}})$	1997	A. Ambainis
k	$O(n^{\frac{\log\log k}{k\log k}})$	2002	A. Beimel, Y. Ishai, E. Kushilevitz, and J.F. Raymond
$k \ge 3$			

# of Servers	Communication Complexity	Year	Reference
2	$O(n^{\frac{1}{3}})$	1995	B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan
k	$O(n^{\frac{1}{k}})$	1995	B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan
k	$O(n^{\frac{1}{2k-1}})$	1997	A. Ambainis
k	$O(n^{\frac{\log \log k}{k \log k}})$	2002	A. Beimel, Y. Ishai, E. Kushilevitz, and J.F. Raymond
$k \ge 3$	$n^{O\left(\sqrt{\frac{\log\log n}{\log n}}\right)}$	2008	S. Yekhanin; K. Efremenko
2	$nO\left(\sqrt{\frac{\log\log n}{\log n}}\right)$	2014	Z. Dvir and S. Gopi

# of Servers	Communication Complexity	Year	Reference
2	$O(n^{\frac{1}{3}})$	1995	B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan
k	$O(n^{\frac{1}{k}})$	1995	B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan
k	$O(n^{\frac{1}{2k-1}})$	1997	A. Ambainis
k	$O(n^{\frac{\log\log k}{k\log k}})$	2002	A. Beimel, Y. Ishai, E. Kushilevitz, and J.F. Raymond
$k \ge 3$	$n^{O\left(\sqrt{\frac{\log\log n}{\log n}}\right)}$	2008	S. Yekhanin; K. Efremenko
2	$n^{O\left(\sqrt{\frac{\log\log n}{\log n}}\right)}$	2014	Z. Dvir and S. Gopi

During the past 20 years, the communication cost of information-theoretic PIR has been reduced dramatically by many researchers:

# of Servers	Communication Complexity	Year	Reference
2	$O(n^{\frac{1}{3}})$	1995	B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan
k	$O(n^{\frac{1}{k}})$	1995	B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan
k	$O(n^{\frac{1}{2k-1}})$	1997	A. Ambainis
k	$O(n^{\frac{\log\log k}{k\log k}})$	2002	A. Beimel, Y. Ishai, E. Kushilevitz, and J.F. Raymond
$k \ge 3$	$nO\left(\sqrt{\frac{\log\log n}{\log n}}\right)$	2008	S. Yekhanin; K. Efremenko
2	$n^{O\left(\sqrt{\frac{\log\log n}{\log n}}\right)}$	2014	Z. Dvir and S. Gopi

Note: Dvir-Gopi protocol gives an even better communication cost for large *k*.

In addition to the communication cost, another important cost metric is the storage overhead, defined as follows:

storage overhead $\stackrel{\text{def}}{=} \frac{\text{total number of bits stored on all the servers}}{\text{number of bits in the database}}$

The storage overhead of replicating the database *k* times is trivially *k*. The Dvir and Gopi paper is considered a breakthrough in part because it reduces the storage overhead from *k* ≥ 3 to *k* = 2, for the same complexity.

In addition to the communication cost, another important cost metric is the storage overhead, defined as follows:

storage overhead $\stackrel{\text{def}}{=}$ <u>total number of bits stored on all the servers</u> number of bits in the database

The storage overhead of replicating the database *k* times is trivially *k*. The Dvir and Gopi paper is considered a breakthrough in part because it reduces the storage overhead from $k \ge 3$ to k = 2, for the same complexity.





In addition to the communication cost, another important cost metric is the storage overhead, defined as follows:

storage overhead $\stackrel{\text{def}}{=}$ <u>total number of bits stored on all the servers</u> number of bits in the database

• The storage overhead of replicating the database *k* times is trivially *k*. The Dvir and Gopi paper is considered a breakthrough in part because it reduces the storage overhead from $k \ge 3$ to k = 2, for the same complexity.





Should we be happy with k = 2? In coding theory, increasing the amount of stored data by a factor of two is often undesirable.

In addition to the communication cost, another important cost metric is the storage overhead, defined as follows:

storage overhead $\stackrel{\text{def}}{=}$ <u>total number of bits stored on all the servers</u> number of bits in the database

• The storage overhead of replicating the database *k* times is trivially *k*. The Dvir and Gopi paper is considered a breakthrough in part because it reduces the storage overhead from $k \ge 3$ to k = 2, for the same complexity.





Should we be happy with k = 2? In coding theory, increasing the amount of stored data by a factor of two is often undesirable.

But doing better than k = 2 **is impossible!** It was shown back in 1995 that the communication cost is $\Omega(n)$ unless the database is replicated on at least two non-communicating servers.

In addition to the communication cost, another important cost metric is the storage overhead, defined as follows:

storage overhead $\stackrel{\text{def}}{=}$ $\frac{\text{total number of bits stored on all the servers}}{1}$ number of bits in the database

• The storage overhead of replicating the database *k* times is trivially *k*. The Dvir and Gopi paper is considered a breakthrough in part because it reduces the storage overhead from $k \ge 3$ to k = 2, for the same complexity.





Should we be happy with k = 2? In coding theory, increasing the amount of stored data by a factor of two is often undesirable.

But doing better than k = 2 **is impossible!** It was shown back in 1995 that the communication cost is $\Omega(n)$ unless the database is replicated on at least two non-communicating servers. Or is it?

This talk: The main theme

Impossible





Open Problem: Can we achieve information-theoretic PIR with low communication cost but *without doubling* (or worse if $k \ge 3$) the number of bits we need to store?

This talk: The main theme

Impossible





Open Problem: Can we achieve information-theoretic PIR with low communication cost but *without doubling* (or worse if $k \ge 3$) the number of bits we need to store?

Taking cue from distributed storage: In practice, the database may need to be stored in a distributed manner (e.g., for security or reliability purposes).

This talk: The main theme

This is cryptography, people! We do the impossible for breakfast.



Open Problem: Can we achieve information-theoretic PIR with low communication cost but *without doubling* (or worse if $k \ge 3$) the number of bits we need to store?

Taking cue from distributed storage: In practice, the database may need to be stored in a distributed manner (e.g., for security or reliability purposes).

Key idea: Partitioning the database

Impossible

Partition the database string *x* into parts $x_1, x_2, ..., x_s$. We will use $m \ge k$ non-communicating servers. But **each server will store only part of the database,** so that the total number of bits stored is $(1 + \varepsilon)n$.

Conventional *k*-server **PIR**

Definition: *k***-server PIR scheme**

A *k*-server PIR scheme consists of the following: a binary string *x* of length *n*, called the **database**, *k* non-communicating servers $S_1, S_2, ..., S_k$ each storing a replica of *x*, a user Alice who wishes to retrieve x_i for some $i \in [n]$, without revealing *i* to any of the servers, and a *k*-server PIR protocol.

Definition: k-server PIR protocol [CKGS95]

The *k*-server PIR protocol \mathcal{P} involves a triple of algorithms \mathcal{Q} , \mathcal{A} , \mathcal{C} and consists of the following steps:

- Alice flips coins and uses the random outcome to invoke the query algorithm Q(k, n; i) that generates a k-tuple of queries q1, q2,..., qk.
- 2 For all $j \in [k]$, Alice sends the query q_i to the *j*-th server S_j .
- **③** For all *j* ∈ [*k*], the server S_j invokes the **answer algorithm** A to respond with the answer $a_i = A(k, j; x, q_i)$.
- Alice computes x_i using the **reconstruction algorithm** $C(k, n; i, a_1, \ldots, a_k)$.

The three algorithms together satisfy the **correctness** ($C(k, n; i, a_1, ..., a_k) = x_i$) and the **privacy** (distibution of q_j independent of i) conditions defined earlier.

Conventional *k*-server PIR: Linearity

Our construction of distributed PIR schemes with low storage overhead uses **two main ingredients**:

- A binary linear code C with a certain special property, to be defined shortly.
- An existing *k*-server PIR protocol in which the answer algorithm is *linear in the database*.



Conventional *k*-server PIR: Linearity

Our construction of distributed PIR schemes with low storage overhead uses **two main ingredients**:

- A binary linear code C with a certain special property, to be defined shortly.
- An existing *k*-server PIR protocol in which the answer algorithm is *linear in the database*.



Definition: Linear *k*-server PIR protocol

A *k*-server PIR protocol $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ is **linear** if for all $x_1, x_2 \in \{0, 1\}^n$ and for all possible queries *q*, the following holds:

 $\mathcal{A}(k,j;\mathbf{x}_1+\mathbf{x}_2,q) = \mathcal{A}(k,j;\mathbf{x}_1,q) + \mathcal{A}(k,j;\mathbf{x}_2,q) \quad \text{for all } j \in [k]$

Conventional *k*-server PIR: Linearity

Our construction of distributed PIR schemes with low storage overhead uses **two main ingredients**:

- A binary linear code C with a certain special property, to be defined shortly.
- An existing *k*-server PIR protocol in which the answer algorithm is *linear in the database*.



Definition: Linear *k*-server PIR protocol

A *k*-server PIR protocol $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ is **linear** if for all $x_1, x_2 \in \{0, 1\}^n$ and for all possible queries *q*, the following holds:

 $\mathcal{A}(k,j;\mathbf{x}_1+\mathbf{x}_2,q) = \mathcal{A}(k,j;\mathbf{x}_1,q) + \mathcal{A}(k,j;\mathbf{x}_2,q) \quad \text{for all } j \in [k]$

Good news: All known PIR protocols are linear!

Note: We also assume that the answer algorithm A is public knowledge. This means that any server can simulate the answers of any other server.

Example: Coded 3-server PIR

Example: Reducing the storage overhead of 3-server PIR

Consider **any** existing 3-server PIR protocol $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$, and assume it is linear. We will reduce its storage overhead from k = 3 to m/s = 2.

Example: Coded 3-server PIR

Example: Reducing the storage overhead of 3-server PIR

Consider **any** existing 3-server PIR protocol $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$, and assume it is linear. We will reduce its storage overhead from k = 3 to m/s = 2.

We partition the database x of length n into 4 parts x_1, x_2, x_3, x_4 , each of length n/4. These parts are distributed among 8 servers as follows:

The result is a coded PIR scheme with s = 4 parts x_1, x_2, x_3, x_4 and m = 8 coded shares $c_1, c_2, c_3, c_4, c_5, c_6, c_7, c_8$.

Example: Coded 3-server PIR

Example: Reducing the storage overhead of 3-server PIR

Consider **any** existing 3-server PIR protocol $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$, and assume it is linear. We will reduce its storage overhead from k = 3 to m/s = 2.

We partition the database x of length n into 4 parts x_1, x_2, x_3, x_4 , each of length n/4. These parts are distributed among 8 servers as follows:

The result is a coded PIR scheme with s = 4 parts x_1, x_2, x_3, x_4 and m = 8 coded shares $c_1, c_2, c_3, c_4, c_5, c_6, c_7, c_8$.

staraga avarbaad —	n/s bits stored on <i>m</i> servers	_ <i>m</i>
storage overhead –	<i>n</i> bits in the database	$-\overline{s}$

Example: How to retrieve *x_i*?

Assume, for now, that Alice wishes to read the *i*-th bit from the first part x_1 . That is, she wants the bit $x_{1,i}$ for some $i \in [n/4]$. She proceeds as follows:

- Alice flips coins and invokes the **query algorithm of** $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ to generate three queries $q_1, q_2, q_3 := \mathcal{Q}(3, n/4; i)$.
- Is sends queries to the 8 servers as follows:

 $(S_1, S_2, S_3, S_4, S_5, S_6, S_7, S_8) \leftarrow (q_1, q_2, q_3, q_3, q_2, q_2, q_3, q_3)$

Alice ignores the answers from S₃, S₆, S₇ but collects the other five answers as follows:

• Since the *answer algorithm of* $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ *is linear* in the database, Alice can compute:

 $a_2' = a_2 + a_5 = \mathcal{A}(3,2;x_2,q_2) + \mathcal{A}(3,2;x_1+x_2,q_2) = \mathcal{A}(3,2;x_1,q_2)$
Assume, for now, that Alice wishes to read the *i*-th bit from the first part x_1 . That is, she wants the bit $x_{1,i}$ for some $i \in \lfloor n/4 \rfloor$. She proceeds as follows:

- Alice flips coins and invokes the **query algorithm of** $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ to generate three queries $q_1, q_2, q_3 := \mathcal{Q}(3, n/4; i)$.
- She sends queries to the 8 servers as follows:

 $(S_1, S_2, S_3, S_4, S_5, S_6, S_7, S_8) \leftarrow (q_1, q_2, q_3, q_3, q_2, q_2, q_3, q_3)$

3 Alice ignores the answers from S_3 , S_6 , S_7 but collects the other five answers as follows:

Since the answer algorithm of P(Q, A, C) is linear in the database, Alice can compute:

 $a'_{2} = a_{2} + a_{5} = \mathcal{A}(3,2;x_{2},q_{2}) + \mathcal{A}(3,2;x_{1} + x_{2},q_{2}) = \mathcal{A}(3,2;x_{1},q_{2})$

Assume, for now, that Alice wishes to read the *i*-th bit from the first part x_1 . That is, she wants the bit $x_{1,i}$ for some $i \in \lfloor n/4 \rfloor$. She proceeds as follows:

- Alice flips coins and invokes the **query algorithm of** $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ to generate three queries $q_1, q_2, q_3 := \mathcal{Q}(3, n/4; i)$.
- 2 She sends queries to the 8 servers as follows:

 $(\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_4, \mathcal{S}_5, \mathcal{S}_6, \mathcal{S}_7, \mathcal{S}_8) \leftarrow (q_1, q_2, q_3, q_3, q_2, q_2, q_3, q_3)$

3 Alice ignores the answers from S_3 , S_6 , S_7 but collects the other five answers as follows:

Server	Query	Response
S_1		$a_1 = \mathcal{A}(3, 1; c_1, q_1) = \mathcal{A}(3, 1; x_1, q_1)$
S_2		$a_2 = \mathcal{A}(3, 2; c_2, q_2) = \mathcal{A}(3, 2; x_2, q_2)$
\mathcal{S}_4		$a_4 = \mathcal{A}(3,3;c_4,q_3) = \mathcal{A}(3,3;x_4,q_3)$
S_5		$a_5 = \mathcal{A}(3, 2; c_5, q_2) = \mathcal{A}(3, 2; x_1 + x_2, q_2)$
S_8		$a_8 = \mathcal{A}(3,3;c_5,q_3) = \mathcal{A}(3,3;x_4+x_1,q_3)$

Since the **answer algorithm of** $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ **is linear** in the database, Alice can compute:

 $a'_{2} = a_{2} + a_{5} = \mathcal{A}(3,2;x_{2},q_{2}) + \mathcal{A}(3,2;x_{1} + x_{2},q_{2}) = \mathcal{A}(3,2;x_{1},q_{2})$

Assume, for now, that Alice wishes to read the *i*-th bit from the first part x_1 . That is, she wants the bit $x_{1,i}$ for some $i \in \lfloor n/4 \rfloor$. She proceeds as follows:

- Alice flips coins and invokes the **query algorithm of** $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ to generate three queries $q_1, q_2, q_3 := \mathcal{Q}(3, n/4; i)$.
- 2 She sends queries to the 8 servers as follows:

 $(S_1, S_2, S_3, S_4, S_5, S_6, S_7, S_8) \leftarrow (q_1, q_2, q_3, q_3, q_2, q_2, q_3, q_3)$

Alice ignores the answers from S₃, S₆, S₇ but collects the other five answers as follows:

Server	Query	Response
\mathcal{S}_1	q_1	$a_1 = \mathcal{A}(3, 1; c_1, q_1) = \mathcal{A}(3, 1; x_1, q_1)$
S_2	q_2	$a_2 = \mathcal{A}(3, 2; c_2, q_2) = \mathcal{A}(3, 2; x_2, q_2)$
\mathcal{S}_4	<i>q</i> ₃	$a_4 = \mathcal{A}(3,3;c_4,q_3) = \mathcal{A}(3,3;x_4,q_3)$
S_5	q_2	$a_5 = \mathcal{A}(3,2;c_5,q_2) = \mathcal{A}(3,2;x_1+x_2,q_2)$
\mathcal{S}_8	<i>q</i> ₃	$a_8 = \mathcal{A}(3,3;c_5,q_3) = \mathcal{A}(3,3;x_4+x_1,q_3)$

Since the *answer algorithm of* P(Q, A, C) *is linear* in the database, Alice can compute:

 $a'_{2} = a_{2} + a_{5} = \mathcal{A}(3,2;x_{2},q_{2}) + \mathcal{A}(3,2;x_{1}+x_{2},q_{2}) = \mathcal{A}(3,2;x_{1},q_{2})$

Assume, for now, that Alice wishes to read the *i*-th bit from the first part x_1 . That is, she wants the bit $x_{1,i}$ for some $i \in \lfloor n/4 \rfloor$. She proceeds as follows:

2 She sends queries to the 8 servers as follows:

 $(\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_4, \mathcal{S}_5, \mathcal{S}_6, \mathcal{S}_7, \mathcal{S}_8) \leftarrow (q_1, q_2, q_3, q_3, q_2, q_2, q_3, q_3)$

Alice ignores the answers from S₃, S₆, S₇ but collects the other five answers as follows:

Server	Query	Response
\mathcal{S}_1	q_1	$a_1 = \mathcal{A}(3, 1; c_1, q_1) = \mathcal{A}(3, 1; x_1, q_1)$
S_2	<i>q</i> ₂	$a_2 = \mathcal{A}(3, 2; c_2, q_2) = \mathcal{A}(3, 2; x_2, q_2)$
\mathcal{S}_4	<i>q</i> 3	$a_4 = \mathcal{A}(3,3;c_4,q_3) = \mathcal{A}(3,3;x_4,q_3)$
S_5	q_2	$a_5 = \mathcal{A}(3, 2; c_5, q_2) = \mathcal{A}(3, 2; x_1 + x_2, q_2)$
\mathcal{S}_8	<i>q</i> ₃	$a_8 = \mathcal{A}(3,3;c_5,q_3) = \mathcal{A}(3,3;x_4+x_1,q_3)$

Since the *answer algorithm of* P(Q, A, C) *is linear* in the database, Alice can compute:

 $\begin{aligned} a_2' &= a_2 + a_5 = \mathcal{A}(3,2;x_2,q_2) + \mathcal{A}(3,2;x_1 + x_2,q_2) = \mathcal{A}(3,2;x_1,q_2) \\ a_3' &= a_4 + a_8 = \mathcal{A}(3,3;x_4,q_3) + \mathcal{A}(3,3;x_4 + x_1,q_3) = \mathcal{A}(3,3;x_1,q_3) \end{aligned}$

Assume, for now, that Alice wishes to read the *i*-th bit from the first part x_1 . That is, she wants the bit $x_{1,i}$ for some $i \in \lfloor n/4 \rfloor$. She proceeds as follows:

She sends queries to the 8 servers as follows:

 $(\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_4, \mathcal{S}_5, \mathcal{S}_6, \mathcal{S}_7, \mathcal{S}_8) \leftarrow (q_1, q_2, q_3, q_3, q_2, q_2, q_3, q_3)$

Alice ignores the answers from S₃, S₆, S₇ but collects the other five answers as follows:

Server	Query	Response
\mathcal{S}_1	q_1	$a_1 = \mathcal{A}(3, 1; c_1, q_1) = \mathcal{A}(3, 1; x_1, q_1)$
S_2	<i>q</i> ₂	$a_2 = \mathcal{A}(3, 2; c_2, q_2) = \mathcal{A}(3, 2; x_2, q_2)$
\mathcal{S}_4	<i>q</i> 3	$a_4 = \mathcal{A}(3,3;c_4,q_3) = \mathcal{A}(3,3;x_4,q_3)$
S_5	q_2	$a_5 = \mathcal{A}(3,2;c_5,q_2) = \mathcal{A}(3,2;x_1+x_2,q_2)$
\mathcal{S}_8	<i>q</i> ₃	$a_8 = \mathcal{A}(3,3;c_5,q_3) = \mathcal{A}(3,3;x_4+x_1,q_3)$

Since the *answer algorithm of* P(Q, A, C) *is linear* in the database, Alice can compute:

 $\begin{aligned} a_2' &= a_2 + a_5 = \mathcal{A}(3,2;x_2,q_2) + \mathcal{A}(3,2;x_1+x_2,q_2) = \mathcal{A}(3,2;x_1,q_2) \\ a_3' &= a_4 + a_8 = \mathcal{A}(3,3;x_4,q_3) + \mathcal{A}(3,3;x_4+x_1,q_3) = \mathcal{A}(3,3;x_1,q_3) \end{aligned}$

Assume, for now, that Alice wishes to read the *i*-th bit from the first part x_1 . That is, she wants the bit $x_{1,i}$ for some $i \in [n/4]$. She proceeds as follows:

Alice ignores the answers from S₃, S₆, S₇ but collects the other five answers as follows:

Server	Query	Response
\mathcal{S}_1	q_1	$a_1 = \mathcal{A}(3, 1; c_1, q_1) = \mathcal{A}(3, 1; x_1, q_1)$
S_2	q_2	$a_2 = \mathcal{A}(3, 2; c_2, q_2) = \mathcal{A}(3, 2; x_2, q_2)$
\mathcal{S}_4	<i>q</i> 3	$a_4 = \mathcal{A}(3,3;c_4,q_3) = \mathcal{A}(3,3;x_4,q_3)$
S_5	q_2	$a_5 = \mathcal{A}(3,2;c_5,q_2) = \mathcal{A}(3,2;x_1+x_2,q_2)$
S_8	<i>q</i> ₃	$a_8 = \mathcal{A}(3,3;c_5,q_3) = \mathcal{A}(3,3;x_4+x_1,q_3)$

Since the *answer algorithm of* P(Q, A, C) *is linear* in the database, Alice can compute:

$$\begin{aligned} a_2' &= a_2 + a_5 = \mathcal{A}(3,2;x_2,q_2) + \mathcal{A}(3,2;x_1 + x_2,q_2) = \mathcal{A}(3,2;x_1,q_2) \\ a_3' &= a_4 + a_8 = \mathcal{A}(3,3;x_4,q_3) + \mathcal{A}(3,3;x_4 + x_1,q_3) = \mathcal{A}(3,3;x_1,q_3) \end{aligned}$$

Substitution Using the reconstruction algorithm of $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$, Alice now computes $\mathcal{C}(3, n/4; i, a_1, a'_2, a'_3)$, which is given by: $\mathcal{C}(3, n/4; i, \mathcal{A}(3, 1; x_1, q_1), \mathcal{A}(3, 2; x_1, q_2), \mathcal{A}(3, 3; x_1, q_3)) = x_{1,i}$

Now assume that Alice wishes to read the *i*-th bit from the second part x_2 . That is, she wants the bit $x_{2,i}$ for some $i \in [n/4]$. She proceeds as follows:

- Alice flips coins and invokes the **query algorithm of** $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ to generate three queries $q_1, q_2, q_3 := \mathcal{Q}(3, n/4; i)$, exactly as before.
- She sends queries to the 8 servers as follows:

 $(S_1, S_2, S_3, S_4, S_5, S_6, S_7, S_8) \leftarrow (q_2, q_1, q_3, q_3, q_2, q_3, q_3, q_3)$

Alice ignores the answers from S₄, S₇, S₈ but collects the other five answers as follows:

Since the *answer algorithm of* P(Q, A, C) *is linear* in the database, Alice can compute:

 $a_2' = a_1 + a_5 = \mathcal{A}(3,2;x_1,q_2) + \mathcal{A}(3,2;x_1+x_2,q_2) = \mathcal{A}(3,2;x_2,q_2)$

Now assume that Alice wishes to read the *i*-th bit from the second part x_2 . That is, she wants the bit $x_{2,i}$ for some $i \in [n/4]$. She proceeds as follows:

• Alice flips coins and invokes the **query algorithm of** $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ to generate three queries $q_1, q_2, q_3 := \mathcal{Q}(3, n/4; i)$, exactly as before.

She sends queries to the 8 servers as follows:

 $(S_1, S_2, S_3, S_4, S_5, S_6, S_7, S_8) \leftarrow (q_2, q_1, q_3, q_3, q_2, q_3, q_3, q_3)$

3 Alice ignores the answers from S_4 , S_7 , S_8 but collects the other five answers as follows:

Since the answer algorithm of P(Q, A, C) is linear in the database, Alice can compute:

 $a'_{2} = a_{1} + a_{5} = \mathcal{A}(3,2;x_{1},q_{2}) + \mathcal{A}(3,2;x_{1}+x_{2},q_{2}) = \mathcal{A}(3,2;x_{2},q_{2})$

Now assume that Alice wishes to read the *i*-th bit from the second part x_2 . That is, she wants the bit $x_{2,i}$ for some $i \in [n/4]$. She proceeds as follows:

- Alice flips coins and invokes the **query algorithm of** $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ to generate three queries $q_1, q_2, q_3 := \mathcal{Q}(3, n/4; i)$, exactly as before.
- 2 She sends queries to the 8 servers as follows:

 $(S_1, S_2, S_3, S_4, S_5, S_6, S_7, S_8) \leftarrow (q_2, q_1, q_3, q_3, q_2, q_3, q_3, q_3)$

3 Alice ignores the answers from S_4 , S_7 , S_8 but collects the other five answers as follows:

Server	Query	Response
S_1		$a_1 = \mathcal{A}(3, 2; c_1, q_2) = \mathcal{A}(3, 2; x_1, q_2)$
S_2		$a_2 = \mathcal{A}(3, 1; c_2, q_1) = \mathcal{A}(3, 1; x_2, q_1)$
S_3		$a_3 = \mathcal{A}(3,3;c_3,q_3) = \mathcal{A}(3,3;x_3,q_3)$
S_5		$a_5 = \mathcal{A}(3, 2; c_5, q_2) = \mathcal{A}(3, 2; x_1 + x_2, q_2)$
S_6		$a_6 = \mathcal{A}(3,3;c_6,q_3) = \mathcal{A}(3,3;x_2+x_3,q_3)$

Since the **answer algorithm of** $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ **is linear** in the database, Alice can compute:

 $a'_{2} = a_{1} + a_{5} = \mathcal{A}(3,2;x_{1},q_{2}) + \mathcal{A}(3,2;x_{1}+x_{2},q_{2}) = \mathcal{A}(3,2;x_{2},q_{2})$

Now assume that Alice wishes to read the *i*-th bit from the second part x_2 . That is, she wants the bit $x_{2,i}$ for some $i \in [n/4]$. She proceeds as follows:

- Alice flips coins and invokes the **query algorithm of** $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ to generate three queries $q_1, q_2, q_3 := \mathcal{Q}(3, n/4; i)$, exactly as before.
- 2 She sends queries to the 8 servers as follows:

 $(\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_4, \mathcal{S}_5, \mathcal{S}_6, \mathcal{S}_7, \mathcal{S}_8) \leftarrow (q_2, q_1, q_3, q_3, q_2, q_3, q_3, q_3)$

Alice ignores the answers from S₄, S₇, S₈ but collects the other five answers as follows:

Server	Query	Response
\mathcal{S}_1	q_2	$a_1 = \mathcal{A}(3, 2; c_1, q_2) = \mathcal{A}(3, 2; x_1, q_2)$
S_2	q_1	$a_2 = \mathcal{A}(3, 1; c_2, q_1) = \mathcal{A}(3, 1; x_2, q_1)$
S_3	<i>q</i> 3	$a_3 = \mathcal{A}(3,3;c_3,q_3) = \mathcal{A}(3,3;x_3,q_3)$
S_5	<i>q</i> ₂	$a_5 = \mathcal{A}(3,2;c_5,q_2) = \mathcal{A}(3,2;x_1+x_2,q_2)$
S_6	<i>q</i> ₃	$a_6 = \mathcal{A}(3,3;c_6,q_3) = \mathcal{A}(3,3;x_2+x_3,q_3)$

Since the *answer algorithm of* P(Q, A, C) *is linear* in the database, Alice can compute:

 $a'_{2} = a_{1} + a_{5} = \mathcal{A}(3,2;x_{1},q_{2}) + \mathcal{A}(3,2;x_{1}+x_{2},q_{2}) = \mathcal{A}(3,2;x_{2},q_{2})$

Now assume that Alice wishes to read the *i*-th bit from the second part x_2 . That is, she wants the bit $x_{2,i}$ for some $i \in [n/4]$. She proceeds as follows:

She sends queries to the 8 servers as follows:

 $(\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_4, \mathcal{S}_5, \mathcal{S}_6, \mathcal{S}_7, \mathcal{S}_8) \leftarrow (q_2, q_1, q_3, q_3, q_2, q_3, q_3, q_3)$

3 Alice ignores the answers from S₄, S₇, S₈ but collects the other five answers as follows:

Server	Query	Response
\mathcal{S}_1	q_2	$a_1 = \mathcal{A}(3, 2; c_1, q_2) = \mathcal{A}(3, 2; x_1, q_2)$
S_2	q_1	$a_2 = \mathcal{A}(3, 1; c_2, q_1) = \mathcal{A}(3, 1; x_2, q_1)$
S_3	<i>q</i> 3	$a_3 = \mathcal{A}(3,3;c_3,q_3) = \mathcal{A}(3,3;x_3,q_3)$
S_5	q_2	$a_5 = \mathcal{A}(3,2;c_5,q_2) = \mathcal{A}(3,2;x_1+x_2,q_2)$
S_6	<i>q</i> ₃	$a_6 = \mathcal{A}(3,3;c_6,q_3) = \mathcal{A}(3,3;x_2+x_3,q_3)$

Since the *answer algorithm of* P(Q, A, C) *is linear* in the database, Alice can compute:

$$\begin{aligned} a_2' &= a_1 + a_5 = \mathcal{A}(3,2; \mathbf{x}_1, q_2) + \mathcal{A}(3,2; \mathbf{x}_1 + \mathbf{x}_2, q_2) = \mathcal{A}(3,2; \mathbf{x}_2, q_2) \\ a_3' &= a_3 + a_6 = \mathcal{A}(3,3; \mathbf{x}_3, q_3) + \mathcal{A}(3,3; \mathbf{x}_2 + \mathbf{x}_3, q_3) = \mathcal{A}(3,3; \mathbf{x}_2, q_3) \end{aligned}$$

Now assume that Alice wishes to read the *i*-th bit from the second part x_2 . That is, she wants the bit $x_{2,i}$ for some $i \in [n/4]$. She proceeds as follows:

Alice ignores the answers from S₄, S₇, S₈ but collects the other five answers as follows:

Server	Query	Response
\mathcal{S}_1	q_2	$a_1 = \mathcal{A}(3, 2; c_1, q_2) = \mathcal{A}(3, 2; x_1, q_2)$
S_2	q_1	$a_2 = \mathcal{A}(3, 1; c_2, q_1) = \mathcal{A}(3, 1; x_2, q_1)$
\mathcal{S}_3	<i>q</i> 3	$a_3 = \mathcal{A}(3,3;c_3,q_3) = \mathcal{A}(3,3;x_3,q_3)$
S_5	q_2	$a_5 = \mathcal{A}(3,2;c_5,q_2) = \mathcal{A}(3,2;x_1+x_2,q_2)$
S_6	<i>q</i> ₃	$a_6 = \mathcal{A}(3,3;c_6,q_3) = \mathcal{A}(3,3;x_2+x_3,q_3)$

Since the *answer algorithm of* P(Q, A, C) *is linear* in the database, Alice can compute:

$$\begin{aligned} a_2' &= a_1 + a_5 = \mathcal{A}(3,2; \mathbf{x}_1, q_2) + \mathcal{A}(3,2; \mathbf{x}_1 + \mathbf{x}_2, q_2) = \mathcal{A}(3,2; \mathbf{x}_2, q_2) \\ a_3' &= a_3 + a_6 = \mathcal{A}(3,3; \mathbf{x}_3, q_3) + \mathcal{A}(3,3; \mathbf{x}_2 + \mathbf{x}_3, q_3) = \mathcal{A}(3,3; \mathbf{x}_2, q_3) \end{aligned}$$

Substitution Using the reconstruction algorithm of $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$, Alice now computes $\mathcal{C}(3, n/4; i, a_2, a'_2, a'_3)$, which is given by: $\mathcal{C}(3, n/4; i, \mathcal{A}(3, 1; x_2, q_1), \mathcal{A}(3, 2; x_2, q_2), \mathcal{A}(3, 3; x_2, q_3)) = x_{2,i}$

Coded *k***-server PIR: Definition**

Definition: Coded *k*-server PIR scheme

A **coded** *k***-server PIR scheme with** *s* **parts and** *m* **shares** consists of the following ingredients:

- A binary string *x* of length *n*, called the **database**, that is partitioned into *s* **parts** *x*₁, *x*₂, ..., *x*_s, each of length *n*/s.
- Coded shares c₁, c₂,..., c_m of length n/s, where c_j is a linear function of x₁, x₂,..., x_s for all j ∈ [m], stored in m non-communicating servers S₁, S₂,..., S_m.
- A user Alice who wishes to retrieve x_i for some i ∈ [n], without revealing i to any of the servers.
- A coded k-server PIR protocol P^{*}(Q^{*}, A^{*}, C^{*}) that emulates a conventional k-server PIR protocol P(Q, A, C).

Note: The emulation property of $\mathcal{P}^*(\mathcal{Q}^*, \mathcal{A}^*, \mathcal{C}^*)$ can be formally defined.

Coded *k***-server PIR: Definition**

Definition: Coded k-server PIR scheme

A **coded** *k***-server PIR scheme with** *s* **parts and** *m* **shares** consists of the following ingredients:

- A binary string *x* of length *n*, called the **database**, that is partitioned into *s* **parts** *x*₁, *x*₂,...,*x*_s, each of length *n*/s.
- Coded shares c₁, c₂,..., c_m of length n/s, where c_j is a linear function of x₁, x₂,..., x_s for all j ∈ [m], stored in m non-communicating servers S₁, S₂,..., S_m.
- A user Alice who wishes to retrieve x_i for some i ∈ [n], without revealing i to any of the servers.
- A coded k-server PIR protocol P^{*}(Q^{*}, A^{*}, C^{*}) that emulates a conventional k-server PIR protocol P(Q, A, C).

Note: The emulation property of $\mathcal{P}^*(\mathcal{Q}^*, \mathcal{A}^*, \mathcal{C}^*)$ can be formally defined.



Theorem 1: Storage overhead of coded PIR

The storage overhead of a coded *k*-server PIR scheme with *s* parts and *m* coded shares is m/s.





- Why does the bit retrieval in the example work? Why does everything nicely cancel out?
- For which values of *m*, *s*, and *k* do coded *k*-server PIR schemes with *s* parts and *m* shares exist?
- What about their communication complexity?
- How small can we make the storage overhead ratio *m*/*s*?



- Why does the bit retrieval in the example work? Why does everything nicely cancel out?
- For which values of *m*, *s*, and *k* do coded *k*-server PIR schemes with *s* parts and *m* shares exist?
- What about their communication complexity?
- How small can we make the storage overhead ratio *m*/*s*?



- Why does the bit retrieval in the example work? Why does everything nicely cancel out?
- For which values of *m*, *s*, and *k* do coded *k*-server PIR schemes with *s* parts and *m* shares exist?
- What about their communication complexity?
- How small can we make the storage overhead ratio *m*/*s*?



- Why does the bit retrieval in the example work? Why does everything nicely cancel out?
- For which values of *m*, *s*, and *k* do coded *k*-server PIR schemes with *s* parts and *m* shares exist?
- What about their communication complexity?
- How small can we make the storage overhead ratio *m*/*s*?

So far, we have seen a general definition, and a single example of a coded PIR scheme with 4 parts and 8 shares that conforms to this definition.



- Why does the bit retrieval in the example work? Why does everything nicely cancel out?
- For which values of *m*, *s*, and *k* do coded *k*-server PIR schemes with *s* parts and *m* shares exist?
- What about their communication complexity?
- How small can we make the storage overhead ratio *m*/*s*?

To answer these questions, let us begin by revisiting the example.

In the encoding equations (\star) of the example, the 8 coded shares are computed from the four database parts x_1, x_2, x_3, x_4 as follows:

$$c_1 = x_1,$$
 $c_3 = x_3,$ $c_5 = x_1 + x_2,$ $c_7 = x_3 + x_4$
 $c_2 = x_2,$ $c_4 = x_4,$ $c_6 = x_2 + x_3,$ $c_8 = x_4 + x_1$

Rewrite these equations in matrix form:

$$(c_1, c_2, c_3, c_4, c_5, c_6, c_7, c_8) = (x_1, x_2, x_3, x_4) \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 \end{bmatrix}$$

Observe that each part x_1 , x_2 , x_3 , x_4 of the database can be recovered from the coded shares in k = 3 different ways. Explicitly:

$$\begin{array}{rcl} x_1 &=& c_1 &=& c_5 + c_2 &=& c_8 + c_4 \\ x_2 &=& c_2 &=& c_5 + c_1 &=& c_6 + c_3 \\ x_3 &=& c_3 &=& c_6 + c_3 &=& c_7 + c_4 \\ x_4 &=& c_4 &=& c_7 + c_3 &=& c_8 + c_1 \end{array}$$

In the encoding equations (\star) of the example, the 8 coded shares are computed from the four database parts x_1, x_2, x_3, x_4 as follows:

$$c_1 = x_1,$$
 $c_3 = x_3,$ $c_5 = x_1 + x_2,$ $c_7 = x_3 + x_4$
 $c_2 = x_2,$ $c_4 = x_4,$ $c_6 = x_2 + x_3,$ $c_8 = x_4 + x_1$

Rewrite these equations in matrix form:

$$(c_1, c_2, c_3, c_4, c_5, c_6, c_7, c_8) = (x_1, x_2, x_3, x_4) \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 \end{bmatrix}$$

Observe that each part x_1 , x_2 , x_3 , x_4 of the database can be recovered from the coded shares in k = 3 different ways. Explicitly:

In the encoding equations (\star) of the example, the 8 coded shares are computed from the four database parts x_1, x_2, x_3, x_4 as follows:

$$c_1 = x_1,$$
 $c_3 = x_3,$ $c_5 = x_1 + x_2,$ $c_7 = x_3 + x_4$
 $c_2 = x_2,$ $c_4 = x_4,$ $c_6 = x_2 + x_3,$ $c_8 = x_4 + x_1$

Rewrite these equations in matrix form:

$$(c_1, c_2, c_3, c_4, c_5, c_6, c_7, c_8) = (x_1, x_2, x_3, x_4) \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 \end{bmatrix}$$

Observe that each part x_1 , x_2 , x_3 , x_4 of the database can be recovered from the coded shares in k = 3 different ways. Explicitly:

In the encoding equations (\star) of the example, the 8 coded shares are computed from the four database parts x_1, x_2, x_3, x_4 as follows:

$$c_1 = x_1,$$
 $c_3 = x_3,$ $c_5 = x_1 + x_2,$ $c_7 = x_3 + x_4$
 $c_2 = x_2,$ $c_4 = x_4,$ $c_6 = x_2 + x_3,$ $c_8 = x_4 + x_1$

Rewrite these equations in matrix form:

$$(c_1, c_2, c_3, c_4, c_5, c_6, c_7, c_8) = (x_1, x_2, x_3, x_4) \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 \end{bmatrix}$$

Observe that each part x_1 , x_2 , x_3 , x_4 of the database can be recovered from the coded shares in k = 3 different ways. Explicitly:

Definition: *k***-server PIR matrix**

Let e_i denote the binary unit vector with 1 in position *i* and zeros elsewhere. An $s \times m$ binary matrix *G* is said to have **property** P_k if for all $i \in [s]$ there exist *k* disjoint sets of columns of *G* that add to e_i . A matrix that has property P_k is also said to be a *k*-server PIR matrix.

Definition: *k***-server PIR matrix**

Let e_i denote the binary unit vector with 1 in position *i* and zeros elsewhere. An $s \times m$ binary matrix *G* is said to have **property** P_k if for all $i \in [s]$ there exist *k* disjoint sets of columns of *G* that add to e_i . A matrix that has property P_k is also said to be a *k*-server PIR matrix.



Definition: *k***-server PIR matrix**

Let e_i denote the binary unit vector with 1 in position *i* and zeros elsewhere. An $s \times m$ binary matrix *G* is said to have **property** P_k if for all $i \in [s]$ there exist *k* disjoint sets of columns of *G* that add to e_i . A matrix that has property P_k is also said to be a *k*-server PIR matrix.



Definition: *k***-server PIR matrix**

Let e_i denote the binary unit vector with 1 in position *i* and zeros elsewhere. An $s \times m$ binary matrix *G* is said to have **property** P_k if for all $i \in [s]$ there exist *k* disjoint sets of columns of *G* that add to e_i . A matrix that has property P_k is also said to be a *k*-server PIR matrix.



Definition: *k***-server PIR matrix**

Let e_i denote the binary unit vector with 1 in position *i* and zeros elsewhere. An $s \times m$ binary matrix *G* is said to have **property** P_k if for all $i \in [s]$ there exist *k* disjoint sets of columns of *G* that add to e_i . A matrix that has property P_k is also said to be a *k*-server PIR matrix.



Definition: *k***-server PIR matrix**

Let e_i denote the binary unit vector with 1 in position *i* and zeros elsewhere. An $s \times m$ binary matrix *G* is said to have **property** P_k if for all $i \in [s]$ there exist *k* disjoint sets of columns of *G* that add to e_i . A matrix that has property P_k is also said to be a *k*-server PIR matrix.



Definition: *k***-server PIR matrix**

Let e_i denote the binary unit vector with 1 in position *i* and zeros elsewhere. An $s \times m$ binary matrix *G* is said to have **property** P_k if for all $i \in [s]$ there exist *k* disjoint sets of columns of *G* that add to e_i . A matrix that has property P_k is also said to be a *k*-server PIR matrix.



Note: This is the encoding matrix for the PIR scheme in our example.

Definition: *k*-server PIR code

A binary linear code \mathbb{C} of length *m* and dimension *s* will be called a *k*-server **PIR code** if there exists a generator matrix *G* for \mathbb{C} with property P_k .

Recovery equations from PIR codes

Lemma 2: Disjoint recovery sets

Let \mathbb{C} be a *k*-server PIR code and let *G* be an $s \times m$ generator matrix for \mathbb{C} with property P_k . Let c = xG be the encoding of a message $x = (x_1, x_2, ..., x_s)$. Then for all $i \in [s]$, there exist *k* disjoint recovery sets $\mathcal{R}_1, \mathcal{R}_2, ..., \mathcal{R}_k$ such that

$$x_i = \sum_{j \in \mathcal{R}_1} c_j = \sum_{j \in \mathcal{R}_2} c_j = \cdots = \sum_{j \in \mathcal{R}_k} c_j$$

Proof. Let $g_1, g_2, ..., g_m$ denote the columns of *G*. Then c = xG can be written in terms of the inner products of these columns with *x*, as follows:

$$c = (c_1, c_2, \ldots, c_m) = (\langle x, g_1 \rangle, \langle x, g_2 \rangle, \ldots, \langle x, g_m \rangle)$$

Now suppose that for some set of indices $\mathcal{R} = \{j_1, j_2, ..., j_r\} \subseteq [m]$, the corresponding columns of *G* add to the unit vector e_i . Then

$$c_{j_1}+\cdots+c_{j_r} = \langle x,g_{j_1}\rangle + \cdots + \langle x,g_{j_r}\rangle = \langle x,g_{j_1}+\cdots+g_{j_r}\rangle = \langle x,e_i\rangle = x_i$$

It follows from the above that the recovery sets $\mathcal{R}_1, \mathcal{R}_2, \ldots, \mathcal{R}_k \subseteq [m]$, are simply the indices of the disjoint sets of columns of *G* that add up to e_i .

Recovery equations from PIR codes

Lemma 2: Disjoint recovery sets

Let \mathbb{C} be a *k*-server PIR code and let *G* be an $s \times m$ generator matrix for \mathbb{C} with property P_k . Let c = xG be the encoding of a message $x = (x_1, x_2, ..., x_s)$. Then for all $i \in [s]$, there exist *k* disjoint recovery sets $\mathcal{R}_1, \mathcal{R}_2, ..., \mathcal{R}_k$ such that

$$x_i = \sum_{j \in \mathcal{R}_1} c_j = \sum_{j \in \mathcal{R}_2} c_j = \cdots = \sum_{j \in \mathcal{R}_k} c_j$$

Proof. Let $g_1, g_2, ..., g_m$ denote the columns of *G*. Then c = xG can be written in terms of the inner products of these columns with x, as follows:

$$\boldsymbol{c} = (c_1, c_2, \ldots, c_m) = (\langle \boldsymbol{x}, \boldsymbol{g}_1 \rangle, \langle \boldsymbol{x}, \boldsymbol{g}_2 \rangle, \ldots, \langle \boldsymbol{x}, \boldsymbol{g}_m \rangle)$$

Now suppose that for some set of indices $\mathcal{R} = \{j_1, j_2, ..., j_r\} \subseteq [m]$, the corresponding columns of *G* add to the unit vector e_i . Then

$$c_{j_1}+\cdots+c_{j_r} = \langle \mathbf{x}, \mathbf{g}_{j_1} \rangle + \cdots + \langle \mathbf{x}, \mathbf{g}_{j_r} \rangle = \langle \mathbf{x}, \mathbf{g}_{j_1}+\cdots+\mathbf{g}_{j_r} \rangle = \langle \mathbf{x}, \mathbf{e}_i \rangle = x_i$$

It follows from the above that the recovery sets $\mathcal{R}_1, \mathcal{R}_2, \ldots, \mathcal{R}_k \subseteq [m]$, are simply the indices of the disjoint sets of columns of *G* that add up to e_i .

Construction of coded PIR schemes

Theorem 3: Coded PIR schemes from PIR codes

Suppose there exists a *k*-server PIR code \mathbb{C} of length *m* and dimension *s* and *ak*-server linear PIR protocol $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$. Then there exists a coded PIR scheme with *s* parts and *m* shares along with the corresponding coded PIR protocol.

Proof. Let *G* be a generator matrix for \mathbb{C} with property P_k . Then the coded shares are computed from the database parts $x_1, x_2, ..., x_s$ as follows:

$$(c_1, c_2, \ldots, c_m) = (x_1, x_2, \ldots, x_s) G$$

Assume Alice wishes to read the *i*-th bit from the ℓ -th part of the database, namely the bit $x_{\ell,i}$ for some $i \in [n/s]$. She will proceed as follows.

- Alice invokes the query algorithm of $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ to generate *k* randomized queries $q_1, q_2, \ldots, q_k := \mathcal{Q}(k, n/s; i)$.
- ② She next finds *k* disjoint recovery sets $\mathcal{R}_1, \mathcal{R}_2, \ldots, \mathcal{R}_k \subseteq [m]$ such that

$$x_\ell = \sum_{j \in \mathcal{R}_1} c_j = \sum_{j \in \mathcal{R}_2} c_j = \cdots = \sum_{j \in \mathcal{R}_k} c_j$$

Such sets exist by Lemma 2. They are used to determine how to assign the queries q_1, q_2, \ldots, q_k to the servers S_1, S_2, \ldots, S_m .

Construction of coded PIR schemes

Theorem 3: Coded PIR schemes from PIR codes

Suppose there exists a *k*-server PIR code \mathbb{C} of length *m* and dimension *s* and *ak*-server linear PIR protocol $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$. Then there exists a coded PIR scheme with *s* parts and *m* shares along with the corresponding coded PIR protocol.

Proof. Let *G* be a generator matrix for \mathbb{C} with property P_k . Then the coded shares are computed from the database parts $x_1, x_2, ..., x_s$ as follows:

 $(\boldsymbol{c}_1,\boldsymbol{c}_2,\ldots,\boldsymbol{c}_m) = (\boldsymbol{x}_1,\boldsymbol{x}_2,\ldots,\boldsymbol{x}_s) G$

Assume Alice wishes to read the *i*-th bit from the ℓ -th part of the database, namely the bit $x_{\ell,i}$ for some $i \in [n/s]$. She will proceed as follows.

- Alice invokes the query algorithm of $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ to generate *k* randomized queries $q_1, q_2, \ldots, q_k := \mathcal{Q}(k, n/s; i)$.
- **2** She next finds *k* disjoint recovery sets $\mathcal{R}_1, \mathcal{R}_2, \ldots, \mathcal{R}_k \subseteq [m]$ such that

$$\mathbf{x}_{\ell} = \sum_{j \in \mathcal{R}_1} c_j = \sum_{j \in \mathcal{R}_2} c_j = \cdots = \sum_{j \in \mathcal{R}_k} c_j$$

Such sets exist by Lemma 2. They are used to determine how to assign the queries q_1, q_2, \ldots, q_k to the servers S_1, S_2, \ldots, S_m .
Proof of main theorem ...continued

③ Let $\mathcal{R} = \mathcal{R}_1 \cup \mathcal{R}_2 \cdots \cup \mathcal{R}_k$ be the union of the *k* recovery sets. For each *j* ∈ \mathcal{R} , Alice finds the unique *t* ∈ [*k*] such that *j* ∈ \mathcal{R}_t and sets $q_j^* = q_t$. For *j* ∉ \mathcal{R} , the query q_j^* can be set arbitrarily (say $q_j^* = q_1$), since the response from \mathcal{S}_j will be ignored. Alice sends the queries to servers as follows:

$$(\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_m) \leftarrow (q_1^*, q_2^*, \ldots, q_m^*)$$

Note: The privacy of the queries $q_1^*, q_2^*, \ldots, q_m^*$ is inherited from the original PIR protocol $\mathcal{P}(\mathcal{Q}, \mathcal{A}, \mathcal{C})$ being emulated.

Alice collects the answers $a_j = \mathcal{A}(k, j; c_j, q_j^*) = \mathcal{A}(k, t; c_j, q_t)$ from the servers, for all $j \in \mathcal{R}$, and computes:

$$a'_{t} \stackrel{\text{def}}{=} \sum_{j \in \mathcal{R}_{t}} \mathcal{A}(k, t; \boldsymbol{c}_{j}, q_{t}) = \mathcal{A}\left(k, t; \sum_{j \in \mathcal{R}_{t}} \boldsymbol{c}_{j}, q_{t}\right) = \mathcal{A}(k, t; \boldsymbol{x}_{\ell}, q_{t})$$

for t = 1, 2, ..., k, where the first equality follows from the linearity of the answer algorithm A and the second from the recovery equations for x_{ℓ} .

Alice completes the retrieval by invoking the reconstruction algorithm of the emulated protocol P(Q, A, C) as follows:

 $\mathcal{C}(k,n/s;i,a'_1,\ldots,a'_k) = \mathcal{C}(k,n/s;i,\mathcal{A}(k,1;x_\ell,q_1),\ldots,\mathcal{A}(k,k;x_\ell,q_k)) = x_{\ell,i}$

What about communication cost?

In order to reduce storage overhead, we emulate a conventional PIR protocol \mathcal{P} by a coded PIR protocol \mathcal{P}^* . How much do we pay in communication complexity?

 $U(\mathcal{P};n) \stackrel{\text{def}}{=} \frac{Worst\text{-}case \text{ total number of bits uploaded}}{by \text{ a protocol } \mathcal{P} \text{ for a database of length } n}$ $D(\mathcal{P};n) \stackrel{\text{def}}{=} Worst\text{-}case \text{ total number of bits downloaded}}$

 $D(\mathcal{P}; n) \stackrel{\text{def}}{=} Worst-case \text{ total number of bits downloaded} \\ by a \text{ protocol } \mathcal{P} \text{ for a database of length } n$

Theorem 4: Communication complexity of coded PIR

Suppose there exists a k-server PIR code \mathbb{C} of length m and dimension s. Then any linear k-server PIR protocol \mathcal{P} can be emulated by a coded PIR protocol \mathcal{P}^* with s parts and m shares, having communication complexity:

 $U(\mathcal{P}^*;n) \leq \frac{m}{k} U(\mathcal{P};n/s) + m \log k \text{ and } D(\mathcal{P}^*;n) \leq \frac{m}{k} D(\mathcal{P};n/s)$

Proof. On the upload side, the number of queries increases from *k* to *m*, but each query is shorter as it is generated by Q(k, n/s; i) rather than Q(k, n; i). On the download side, the number of answers also increases from *k* to *m*.

What about communication cost?

In order to reduce storage overhead, we emulate a conventional PIR protocol \mathcal{P} by a coded PIR protocol \mathcal{P}^* . How much do we pay in communication complexity?

 $U(\mathcal{P};n) \stackrel{\text{def}}{=} Worst-case \text{ total number of bits uploaded} \\ by a \text{ protocol } \mathcal{P} \text{ for a database of length } n$

 $D(\mathcal{P}; n) \stackrel{\text{def}}{=} \begin{array}{c} \text{Worst-case total number of bits downloaded} \\ by a \text{ protocol } \mathcal{P} \text{ for a database of length } n \end{array}$

Theorem 4: Communication complexity of coded PIR

Suppose there exists a *k*-server PIR code \mathbb{C} of length *m* and dimension *s*. Then any linear *k*-server PIR protocol \mathcal{P} can be emulated by a coded PIR protocol \mathcal{P}^* with *s* parts and *m* shares, having communication complexity:

 $U(\mathcal{P}^*;n) \leq \frac{m}{k} U(\mathcal{P};n/s) + m \log k \text{ and } D(\mathcal{P}^*;n) \leq \frac{m}{k} D(\mathcal{P};n/s)$

Proof. On the upload side, the number of queries increases from *k* to *m*, but each query is shorter as it is generated by Q(k, n/s; i) rather than Q(k, n; i). On the download side, the number of answers also increases from *k* to *m*.

We have shown that:







- Why does the bit retrieval in the example work? Why does everything cancel out?
- For which *m*, *s*, and *k* do coded *k*-server PIR schemes with *s* parts and *m* shares exist?
- What about the communication complexity of coded PIR schemes?

We have shown that:





- Why does the bit retrieval in the example work? Why does everything cancel ou?
- For which *m*, *s*, and *k* do coded *k*-server PIR schemes with *s* parts and *m* shares exist?
- What about the communication complexity of coded PIR schemes?

We have shown that:





We have shown that:

- Why does the bit retrieval in the example work? Why does everything cancel ou?
- For which *m*, *s*, and *k* do coded *k*-server PIR schemes with *s* parts and *m* shares exist?
- What about the communication complexity of coded PIR schemes?

How small can we make the storage overhead ratio *m/s*?

existing *k*-server linear PIR protocol *P* +

k-server PIR code \mathbb{C} of length *m* and dimension *s*

coded *k*-server PIR protocol \mathcal{P}^* with storage overhead m/s

New problem: High-rate PIR codes

According to our construction, coded PIR schemes exist whenever PIR codes exist. The storage overhead of such coded PIR schemes is completely determined by the rate of the underlying PIR code.

Open Problem: Given positive integers *s* and *k*, determine the smallest *m* such that there exists a *k*-server PIR code of length *m* and dimension *s*.

 $M(s,k) \stackrel{\text{def}}{=} Shortest possible length m of$ a k-server PIR code of dimension s

 $D(s,k) \stackrel{\text{def}}{=} Smallest possible redundancy of$ a k-server PIR code of dimension s

With this notation:

storage overhead

$$\frac{M(s,k)}{s} = 1 + \frac{\rho(s,k)}{s}$$

New problem: High-rate PIR codes

According to our construction, coded PIR schemes exist whenever PIR codes exist. The storage overhead of such coded PIR schemes is completely determined by the rate of the underlying PIR code.

Open Problem: Given positive integers *s* and *k*, determine the smallest *m* such that there exists a *k*-server PIR code of length *m* and dimension *s*.

 $M(s,k) \stackrel{\text{def}}{=} Shortest possible length m of a k-server PIR code of dimension s$

 $\rho(s,k) \stackrel{\text{def}}{=} Smallest possible redundancy of$ a k-server PIR code of dimension s

With this notation:

storage overhead
$$= \; rac{M(s,k)}{s} \; = \; 1 \; + \; rac{
ho(s,k)}{s}$$

New problem: High-rate PIR codes

According to our construction, coded PIR schemes exist whenever PIR codes exist. The storage overhead of such coded PIR schemes is completely determined by the rate of the underlying PIR code.

Open Problem: Given positive integers *s* and *k*, determine the smallest *m* such that there exists a *k*-server PIR code of length *m* and dimension *s*.

 $M(s,k) \stackrel{\text{def}}{=} Shortest possible length m of a k-server PIR code of dimension s$

 $\rho(s,k) \stackrel{\text{def}}{=} \begin{array}{c} \text{Smallest possible redundancy of} \\ a k - server PIR code of dimension s \end{array}$

With this notation:

storage overhead = $\frac{M}{-}$

$$\frac{(s,k)}{s} = 1 + \frac{\rho(s)}{s}$$



We have converted a PIR problem to a coding theory problem!

Optimal solution for two servers

For k = 2, the coding-theory problem is trivial. The single parity-check code of dimension *s* is a 2-server PIR code, and therefore:

$$M(s,2) = s+1$$

$$\rho(s,2) = 1$$

Why is this true? The encoding of each message $x = (x_1, x_2, ..., x_s)$ consists of appending an overall parity bit

$$c = x_1 + x_2 + \cdots + x_s$$

Thus for all $i \in [s]$, we have $x_i = x_1 + \cdots + x_{i-1} + c + x_{i+1} + \cdots + x_s$. This corresponds to two disjoint recovery sets $\mathcal{R}_1 = \{i\}$ and $\mathcal{R}_2 = [s+1] \setminus \{i\}$.

Optimal solution for two servers

For k = 2, the coding-theory problem is trivial. The single parity-check code of dimension *s* is a 2-server PIR code, and therefore:

$$M(s,2) = s+1$$

$$\rho(s,2) = 1$$

Why is this true? The encoding of each message $x = (x_1, x_2, ..., x_s)$ consists of appending an overall parity bit

$$c = x_1 + x_2 + \cdots + x_s$$

Thus for all $i \in [s]$, we have $x_i = x_1 + \cdots + x_{i-1} + c + x_{i+1} + \cdots + x_s$. This corresponds to two disjoint recovery sets $\mathcal{R}_1 = \{i\}$ and $\mathcal{R}_2 = [s+1] \setminus \{i\}$.

Optimal solution for two servers

For k = 2, the coding-theory problem is trivial. The single parity-check code of dimension *s* is a 2-server PIR code, and therefore:

$$M(s,2) = s+1$$

$$\rho(s,2) = 1$$

Why is this true? The encoding of each message $x = (x_1, x_2, ..., x_s)$ consists of appending an overall parity bit

$$c = x_1 + x_2 + \cdots + x_s$$

Thus for all $i \in [s]$, we have $x_i = x_1 + \cdots + x_{i-1} + c + x_{i+1} + \cdots + x_s$. This corresponds to two disjoint recovery sets $\mathcal{R}_1 = \{i\}$ and $\mathcal{R}_2 = [s+1] \setminus \{i\}$.

Theorem 5: PIR without storage overhead

For all $\varepsilon > 0$ it is possible to achieve information-theoretic PIR with communication complexity $n^{o(1)}$ by storing at most $(1 + \varepsilon)n$ bits.

Proof. Take $s = 1/\varepsilon$, and combine our results for k = 2 with the results of Dvir-Gopi on 2-server PIR with subpolynomial communication.

Open Problem: Can we achieve information-theoretic PIR with low communication cost *without doubling the number of bits we need to store?*

Open Problem: Can we achieve information-theoretic PIR with low communication cost *without doubling the number of bits we need to store?*

Any reason to go on... Why not stop here?

Open Problem: Can we achieve information-theoretic PIR with low communication cost *without doubling the number of bits we need to store?*

Any reason to go on... Why not stop here?

• As the number of servers *k* grows, the **communication complexity improves** dramatically:

$$n^{O\left(\sqrt{\frac{\log\log n}{\log n}}\right)} \xrightarrow{k \text{ large}} \text{polylog}(n)$$

Open Problem: Can we achieve information-theoretic PIR with low communication cost *without doubling the number of bits we need to store?*

Any reason to go on... Why not stop here?

• As the number of servers *k* grows, the **communication complexity improves** dramatically: $\begin{array}{c} & & \\$

- Steiner systems and *t*-designs
- majority-logic decodable codes
- local codes with availability

- multiset batch codes
- bipartite graphs of girth 6
- constant-weight codes

Open Problem: Can we achieve information-theoretic PIR with low communication cost *without doubling the number of bits we need to store?*

Any reason to go on... Why not stop here?

• As the number of servers *k* grows, the **communication complexity improves** dramatically: $\begin{array}{c} & & \\$

- Steiner systems and *t*-designs
- majority-logic decodable codes
- local codes with availability

- multiset batch codes
- bipartite graphs of girth 6
- constant-weight codes

Suppose that k = 3 and $s = \sigma^2$ for some $\sigma \in \mathbb{Z}$. Arrange the σ^2 message bits in the form of a $\sigma \times \sigma$ *square*. To every message, we append 2σ parity bits given by:

$$c_i = x_{i,1} + x_{i,2} + \dots + x_{i,\sigma} \quad \text{for } i \in [\sigma]$$

$$c'_j = x_{1,j} + x_{2,j} + \dots + x_{\sigma,j} \quad \text{for } j \in [\sigma]$$

$$\begin{array}{c|c|c} x_{1,1}\cdots x_{1,j}\cdots x_{1,\sigma} & c_1\\ \vdots & \vdots & \vdots \\ x_{i,1}\cdots x_{i,j}\cdots x_{i,\sigma} & c_i\\ \vdots & \vdots & \vdots \\ x_{\sigma,1}\cdots x_{\sigma,j}\cdots x_{\sigma,\sigma} & c_{\sigma} \\ \hline \hline c_1'\cdots c_j'\cdots c_{\sigma}' \end{array}$$

Suppose that k = 3 and $s = \sigma^2$ for some $\sigma \in \mathbb{Z}$. Arrange the σ^2 message bits in the form of a $\sigma \times \sigma$ *square*. To every message, we append 2σ parity bits given by:

$$c_i = x_{i,1} + x_{i,2} + \dots + x_{i,\sigma} \quad \text{for } i \in [\sigma]$$

$$c'_j = x_{1,j} + x_{2,j} + \dots + x_{\sigma,j} \quad \text{for } j \in [\sigma]$$



Then for each message bit $x_{i,j}$ we have **three disjoint re**covery equations given by $x_{i,j}$ itself and:

 $x_{i,1} + \dots + x_{i,j-1} + c_i + x_{i,j+1} + \dots + x_{i,\sigma} = x_{1,j} + \dots + x_{i-1,j} + c'_j + x_{i+1,j} + \dots + x_{\sigma,j}$

Suppose that k = 3 and $s = \sigma^2$ for some $\sigma \in \mathbb{Z}$. Arrange the σ^2 message bits in the form of a $\sigma \times \sigma$ *square*. To every message, we append 2σ parity bits given by:

$$c_i = x_{i,1} + x_{i,2} + \dots + x_{i,\sigma} \quad \text{for } i \in [\sigma]$$

$$c'_j = x_{1,j} + x_{2,j} + \dots + x_{\sigma,j} \quad \text{for } j \in [\sigma]$$

 $\begin{array}{c|cccc} x_{1,1}\cdots x_{1,j}\cdots x_{1,\sigma} & c_1\\ \vdots & \vdots & \vdots \\ x_{i,1}\cdots x_{i,j}\cdots x_{i,\sigma} & c_i\\ \vdots & \vdots & \vdots \\ x_{\sigma,1}\cdots x_{\sigma,j}\cdots x_{\sigma,\sigma} & c_{\sigma}\\ \hline c_1'\cdots c_j'\cdots c_{\sigma}' \end{array}$

Then for each message bit $x_{i,j}$ we have **three disjoint re**covery equations given by $x_{i,j}$ itself and:

$$x_{i,1} + \dots + x_{i,j-1} + c_i + x_{i,j+1} + \dots + x_{i,\sigma} = x_{1,j} + \dots + x_{i-1,j} + c'_j + x_{i+1,j} + \dots + x_{\sigma,j}$$

More generally, we arrange σ^{k-1} message bits in the form of a (k-1)-*dimensional hypercube* and append a parity bit to each of its $(k-1)\sigma^{k-2}$ columns. This proves:

 $M(s,k) = s + (k-1) \left\lceil \sqrt[k-1]{s} \right\rceil^{k-2}$

 $\rho(s,k) \leqslant (k{-}1)^{\left\lceil k{-}1 \sqrt{s} \right\rceil}^{k-2}$

Suppose that k = 3 and $s = \sigma^2$ for some $\sigma \in \mathbb{Z}$. Arrange the σ^2 message bits in the form of a $\sigma \times \sigma$ *square*. To every message, we append 2σ parity bits given by:

$$c_i = x_{i,1} + x_{i,2} + \dots + x_{i,\sigma} \quad \text{for } i \in [\sigma]$$

$$c'_j = x_{1,j} + x_{2,j} + \dots + x_{\sigma,j} \quad \text{for } j \in [\sigma]$$



Then for each message bit $x_{i,j}$ we have **three disjoint recovery equations** given by $x_{i,j}$ itself and:

$$x_{i,1} + \dots + x_{i,j-1} + c_i + x_{i,j+1} + \dots + x_{i,\sigma} = x_{1,j} + \dots + x_{i-1,j} + c'_j + x_{i+1,j} + \dots + x_{\sigma,j}$$

More generally, we arrange σ^{k-1} message bits in the form of a (k-1)-*dimensional hypercube* and append a parity bit to each of its $(k-1)\sigma^{k-2}$ columns. This proves:

$$M(s,k) = s + (k-1) \left[\sqrt[k-1]{s}\right]^{k-2}$$

$$\rho(s,k) \leqslant (k-1) \left\lceil \sqrt[k-1]{s} \right\rceil^{k-2}$$

It follows that $\lim_{s\to\infty} M(s,k)/s = 1$ for all fixed $k \ge 2$. Therefore, we have proved:



Corollary 6: Multiple-server PIR without storage overhead

For all fixed $k \ge 2$ and all $\varepsilon > 0$, there exist *k*-server coded PIR schemes that store at most $(1 + \varepsilon)n$ bits.

Majority-logic decoding originated with the work of Reed and Massey over 50 years ago. 100s of papers in the 1960s and 1970s... now forgotten.

Majority-logic decoding originated with the work of Reed and Massey over 50 years ago. 100s of papers in the 1960s and 1970s... now forgotten.

Definition: Majority-logic decodable codes

A linear code \mathbb{C} of length n is **majority-logic decodable** with parameter J iff for each position $i \in [n]$, there exist J **parity-checks orthogonal on this position**:



Majority-logic decoding originated with the work of Reed and Massey over 50 years ago. 100s of papers in the 1960s and 1970s... now forgotten.



Majority-logic decoding: Given *y*, evaluate the *J* orthogonal parity-checks for each position *i*.

Majority-logic decoding originated with the work of Reed and Massey over 50 years ago. 100s of papers in the 1960s and 1970s... now forgotten.



Majority-logic decoding: Given *y*, evaluate the *J* orthogonal parity-checks for each position *i*. Then:

• $y_i = c_i$ and $\leq t$ other errors \implies at least J - t checks evaluate to 0

Majority-logic decoding originated with the work of Reed and Massey over 50 years ago. 100s of papers in the 1960s and 1970s... now forgotten.



Majority-logic decoding: Given *y*, evaluate the *J* orthogonal parity-checks for each position *i*. Then:

- $y_i = c_i$ and $\leq t$ other errors \implies at least J t checks evaluate to 0
- $y_i \neq c_i$ and $\leq t$ other errors \implies at least J t checks evaluate to 1

Majority-logic decoding originated with the work of Reed and Massey over 50 years ago. 100s of papers in the 1960s and 1970s... now forgotten.



Majority-logic decoding: Given *y*, evaluate the *J* orthogonal parity-checks for each position *i*. Then:

• $y_i = c_i$ and $\leq t$ other errors \implies at least J - t checks evaluate to 0 • $y_i \neq c_i$ and $\leq t$ other errors \implies at least J - t checks evaluate to 1

There is an error at position *i* iff a *majority of the J checks* evaluate to 1.

Lemma 7: PIR codes from majority-logic codes

Let \mathbb{C} be a majority-logic decodable code with parameter *J*. Then \mathbb{C} is also a *k*-server PIR code with k = J + 1.



Proof. It is easy to see that a **systematic generator matrix** *G* for \mathbb{C} has property P_k with k = J + 1. Since *G* is systematic, the column in position *i* is e_i .



Lemma 7: PIR codes from majority-logic codes

Let \mathbb{C} be a majority-logic decodable code with parameter *J*. Then \mathbb{C} is also a *k*-server PIR code with k = J + 1.

Proof. It is easy to see that a *systematic generator matrix G* for \mathbb{C} has property P_k with k = J + 1. Since *G* is systematic, the column in position *i* is e_i .



Lemma 7: PIR codes from majority-logic codes

Let \mathbb{C} be a majority-logic decodable code with parameter J. Then \mathbb{C} is also a *k*-server PIR code with k = J + 1.

Proof. It is easy to see that a **systematic generator matrix** G for C has property P_k with k = J + 1. Since G is systematic, the column in position i is e_i .

$$G = \begin{bmatrix} \mathcal{R}_{1} & \mathbf{e}_{i} & \mathcal{R}_{2} & \cdots & \mathcal{R}_{J-1} & \mathcal{R}_{J} \end{bmatrix}$$

$$\stackrel{11...11}{\underset{i}{\overset{1}{\underset{i}{\underset{i}{\atop 1}}}} \cdots \\ \underset{i}{\overset{1}{\underset{i}{\atop 1}}} \stackrel{11...1}{\underset{i}{\atop 1}} } J \text{ codewords of } \mathbb{C}^{\perp}$$

Thus $\{e_i\}$ and $\mathcal{R}_1, \mathcal{R}_2, \ldots, \mathcal{R}_I$ are *disjoint sets of columns* of *G* that add to e_i .

Lemma 7: PIR codes from majority-logic codes

Let \mathbb{C} be a majority-logic decodable code with parameter *J*. Then \mathbb{C} is also a *k*-server PIR code with k = J + 1.



Numerous algebraic constructions of cyclic majority-logic decodable codes are known. For example, Reed-Muller codes, BCH codes, and other codes invariant under the *group of affine permutations:*

$$\boldsymbol{\alpha}^i \mapsto \boldsymbol{\beta} \boldsymbol{\alpha}^i + \boldsymbol{\gamma} \quad \text{for all } i = 0, 1, \dots, 2^m - 2 \text{ and } \boldsymbol{\beta}, \boldsymbol{\gamma} \in \operatorname{GF}(2^m)$$

T. Kasami, S. Lin, and W.W. Peterson, Some results on cyclic codes which are invariant under the affine group, *Information and Control*, vol. 2, pp. 475–496, November 1968.

Theorem: Doubly-transitive majority-logic codes

Let $n = 2^{2ab} - 1$ and let \mathbb{C} be a binary cyclic code of length n and co-dimension $(2^{b+1}-1)^a - 1$. Then \mathbb{C} is majority-logic decodable with parameter $J = 2^a + 1$.

Lemma 7: PIR codes from majority-logic codes

Let \mathbb{C} be a majority-logic decodable code with parameter *J*. Then \mathbb{C} is also a *k*-server PIR code with k = J + 1.



Numerous algebraic constructions of cyclic majority-logic decodable codes are known. For example, Reed-Muller codes, BCH codes, and other codes invariant under the *group of affine permutations:*

$$\boldsymbol{\alpha}^i \mapsto \boldsymbol{\beta} \boldsymbol{\alpha}^i + \boldsymbol{\gamma} \quad \text{for all } i = 0, 1, \dots, 2^m - 2 \text{ and } \boldsymbol{\beta}, \boldsymbol{\gamma} \in \operatorname{GF}(2^m)$$

T. Kasami, S. Lin, and W.W. Peterson, Some results on cyclic codes which are invariant under the affine group, *Information and Control*, vol. 2, pp. 475–496, November 1968.

Theorem: Doubly-transitive majority-logic codes

Let $n = 2^{2ab} - 1$ and let \mathbb{C} be a binary cyclic code of length n and co-dimension $(2^{b+1}-1)^a - 1$. Then \mathbb{C} is majority-logic decodable with parameter $J = 2^a + 1$.

As a corollary to this theorem and Lemma 7, whenever the number of servers is of the form $k = 4, 6, 10, ..., 2^a + 2$, we have: $\rho(s, k) = O(\sqrt{s})$

Construction from certain set systems

Definition: Almost disjoint *k*-covers

Let $A = \{A_1, A_2, ..., A_r\}$ be a collection of subsets of [s]. We say that A is a *k*-cover of [s] if every $i \in [s]$ belongs to at least *k* of the subsets in A. We say that these subsets are **almost disjoint** if any two of them intersect in at most one element.

Given any collection $\mathcal{A} = \{A_1, A_2, ..., A_r\}$ of subsets of [s], we construct a **systematic** (s + r, s) **linear code** $\mathbb{C}(\mathcal{A})$ as follows. To each message $x = (x_1, x_2, ..., x_s)$, we **append** *r* **parity bits** given by:

$$c_1 = \sum_{j \in A_1} x_j, \ c_2 = \sum_{j \in A_2} x_j, \ \cdots, \ c_r = \sum_{j \in A_r} x_j$$

Construction from certain set systems

Definition: Almost disjoint *k*-covers

Let $\mathcal{A} = \{A_1, A_2, ..., A_r\}$ be a collection of subsets of [s]. We say that \mathcal{A} is a *k*-cover of [s] if every $i \in [s]$ belongs to at least *k* of the subsets in \mathcal{A} . We say that these subsets are **almost disjoint** if any two of them intersect in at most one element.

Given any collection $\mathcal{A} = \{A_1, A_2, ..., A_r\}$ of subsets of [s], we construct a **systematic** (s + r, s) **linear code** $\mathbb{C}(\mathcal{A})$ as follows. To each message $\mathbf{x} = (x_1, x_2, ..., x_s)$, we **append** *r* **parity bits** given by:

$$c_1 = \sum_{j \in A_1} x_j$$
, $c_2 = \sum_{j \in A_2} x_j$, \cdots , $c_r = \sum_{j \in A_r} x_j$
Definition: Almost disjoint *k*-covers

Let $\mathcal{A} = \{A_1, A_2, ..., A_r\}$ be a collection of subsets of [s]. We say that \mathcal{A} is a *k*-cover of [s] if every $i \in [s]$ belongs to at least *k* of the subsets in \mathcal{A} . We say that these subsets are **almost disjoint** if any two of them intersect in at most one element.

Given any collection $\mathcal{A} = \{A_1, A_2, \dots, A_r\}$ of subsets of [s], we construct a **systematic** (s + r, s) **linear code** $\mathbb{C}(\mathcal{A})$ as follows. To each message $\mathbf{x} = (x_1, x_2, \dots, x_s)$, we **append** *r* **parity bits** given by:

$$c_1 = \sum_{j \in A_1} x_j, \quad c_2 = \sum_{j \in A_2} x_j, \quad \cdots \quad , c_r = \sum_{j \in A_r} x_j$$

Lemma 8: PIR codes from almost disjoint k-covers

Suppose that $A = \{A_1, A_2, ..., A_r\}$ is a (k-1)-cover of [s] and the sets in A are almost disjoint. Then the resulting (s + r, s) code $\mathbb{C}(A)$ is a *k*-server PIR code.

Definition: Almost disjoint *k*-covers

Let $\mathcal{A} = \{A_1, A_2, ..., A_r\}$ be a collection of subsets of [s]. We say that \mathcal{A} is a *k*-cover of [s] if every $i \in [s]$ belongs to at least *k* of the subsets in \mathcal{A} . We say that these subsets are **almost disjoint** if any two of them intersect in at most one element.

Given any collection $\mathcal{A} = \{A_1, A_2, \dots, A_r\}$ of subsets of [s], we construct a **systematic** (s + r, s) **linear code** $\mathbb{C}(\mathcal{A})$ as follows. To each message $\mathbf{x} = (x_1, x_2, \dots, x_s)$, we **append** *r* **parity bits** given by:

$$c_1 = \sum_{j \in A_1} x_j, \quad c_2 = \sum_{j \in A_2} x_j, \quad \cdots \quad , c_r = \sum_{j \in A_r} x_j$$

Lemma 8: PIR codes from almost disjoint k-covers

Suppose that $A = \{A_1, A_2, ..., A_r\}$ is a (k-1)-cover of [s] and the sets in A are almost disjoint. Then the resulting (s + r, s) code $\mathbb{C}(A)$ is a *k*-server PIR code.

Proof. Given $i \in [s]$, find k - 1 subsets in A that contain i. W.l.o.g., suppose these subsets are $A_1, A_2, \ldots, A_{k-1}$. Let $A'_j = A_j \setminus \{i\}$ for all j. Then the sets $A'_1, A'_2, \ldots, A'_{k-1}$ are disjoint. These sets give rise to k disjoint recovery equations:

$$x_i = c_1 + \sum_{j \in A'_1} x_j = c_2 + \sum_{j \in A'_2} x_j = \cdots = c_{k-1} + \sum_{j \in A'_{k-1}} x_j$$

Definition: Almost disjoint *k*-covers

Let $\mathcal{A} = \{A_1, A_2, ..., A_r\}$ be a collection of subsets of [s]. We say that \mathcal{A} is a *k*-cover of [s] if every $i \in [s]$ belongs to at least *k* of the subsets in \mathcal{A} . We say that these subsets are **almost disjoint** if any two of them intersect in at most one element.

Given any collection $\mathcal{A} = \{A_1, A_2, ..., A_r\}$ of subsets of [s], we construct a **systematic** (s + r, s) **linear code** $\mathbb{C}(\mathcal{A})$ as follows. To each message $\mathbf{x} = (x_1, x_2, ..., x_s)$, we **append** *r* **parity bits** given by:

$$c_1 = \sum_{j \in A_1} x_j, \quad c_2 = \sum_{j \in A_2} x_j, \quad \cdots \quad , c_r = \sum_{j \in A_r} x_j$$

Lemma 8: PIR codes from almost disjoint k-covers

Suppose that $A = \{A_1, A_2, ..., A_r\}$ is a (k-1)-cover of [s] and the sets in A are almost disjoint. Then the resulting (s + r, s) code $\mathbb{C}(A)$ is a *k*-server PIR code.

Corollary 9: PIR codes from almost disjoint *k***-covers**

If there exists an almost disjoint (k-1)-cover of [s] with r sets, then $\rho(s,k) \leq r$.

Definition: Almost disjoint *k*-covers

Let $\mathcal{A} = \{A_1, A_2, ..., A_r\}$ be a collection of subsets of [s]. We say that \mathcal{A} is a *k*-cover of [s] if every $i \in [s]$ belongs to at least *k* of the subsets in \mathcal{A} . We say that these subsets are **almost disjoint** if any two of them intersect in at most one element.

Given any collection $\mathcal{A} = \{A_1, A_2, ..., A_r\}$ of subsets of [s], we construct a **systematic** (s + r, s) **linear code** $\mathbb{C}(\mathcal{A})$ as follows. To each message $\mathbf{x} = (x_1, x_2, ..., x_s)$, we **append** *r* **parity bits** given by:

$$c_1 = \sum_{j \in A_1} x_j, \quad c_2 = \sum_{j \in A_2} x_j, \quad \cdots \quad , c_r = \sum_{j \in A_r} x_j$$

Lemma 8: PIR codes from almost disjoint k-covers

Suppose that $A = \{A_1, A_2, ..., A_r\}$ is a (k-1)-cover of [s] and the sets in A are almost disjoint. Then the resulting (s + r, s) code $\mathbb{C}(A)$ is a *k*-server PIR code.

Corollary 9: PIR codes from almost disjoint *k***-covers**

If there exists an almost disjoint (k-1)-cover of [s] with r sets, then $\rho(s,k) \leq r$.



Where can we get almost disjoint k-covers or small size r?

Let *V* be a set with *r* elements, called **points**. A *Steiner system* S(2, q, r) is a collection \mathcal{B} of subsets of *V* of size *q*, called **blocks**, such that every pair of points is contained in exactly one block.

 Such a system is an example of a balanced incomplete block design.



Let *V* be a set with *r* elements, called **points.** A *Steiner system* S(2, q, r) is a collection \mathcal{B} of subsets of *V* of size *q*, called **blocks**, such that every pair of points is contained in exactly one block.

• Such a system is an example of a *balanced incomplete block design*.



Let *V* be a set with *r* elements, called **points.** A *Steiner system* S(2, q, r) is a collection \mathcal{B} of subsets of *V* of size *q*, called **blocks**, such that every pair of points is contained in exactly one block.

 Such a system is an example of a balanced incomplete block design.



Observation: There are $b = \binom{r}{2} / \binom{q}{2}$ blocks in \mathcal{B} and each point is contained in (r-1)/(q-1) of them. Moreover, any two blocks intersect in at most one point.

Let *V* be a set with *r* elements, called **points.** A *Steiner system* S(2, q, r) is a collection \mathcal{B} of subsets of *V* of size *q*, called **blocks**, such that every pair of points is contained in exactly one block.

• Such a system is an example of a *balanced incomplete block design*.



Observation: There are $b = \binom{r}{2} / \binom{q}{2}$ blocks in \mathcal{B} and each point is contained in (r-1)/(q-1) of them. Moreover, any two blocks intersect in at most one point. For more on Steiner systems and their properties, see the following papers:

G.D. Cohen anf P. Frankl, On generalized perfect codes and Steiner systems, Annals of Discrete Mathematics, 18, pp. 197–200, 1983.

G.D. Cohen and B. Montaron, Empilements parfaits de boules dans les espaces vectoriels binaires, *Compte Rendus de l'Academie des Sciences*, **288**, pp. 579–582, 1979.

B. Montaron and G.D. Cohen, Codes parfaits binaires a plusieurs rayons, *Revue du Centre d'Études Théoriques de la Détection et Communication*, **NS1979-2**, pp. 35–58, 1979.

Let *V* be a set with *r* elements, called **points.** A *Steiner system* S(2, q, r) is a collection \mathcal{B} of subsets of *V* of size *q*, called **blocks**, such that every pair of points is contained in exactly one block.

• Such a system is an example of a *balanced incomplete block design*.



Observation: There are $b = \binom{r}{2} / \binom{q}{2}$ blocks in \mathcal{B} and each point is contained in (r-1)/(q-1) of them. Moreover, any two blocks intersect in at most one point.

Conclusion: The blocks of a Steiner system S(2, q, s) form an almost disjoint (s-1)/(q-1)-cover of [s]. Therefore, when such Steiner systems exist, we have

$$p(s,k) \leq$$
number of blocks in $S(2,q,s) = \frac{s(s-1)}{q(q-1)} = \frac{s(k-1)^2}{s+k}$

where k = (s-1)/(q-1) + 1.

Let *V* be a set with *r* elements, called **points.** A *Steiner system* S(2, q, r) is a collection \mathcal{B} of subsets of *V* of size *q*, called **blocks**, such that every pair of points is contained in exactly one block.

• Such a system is an example of a *balanced incomplete block design*.



Observation: There are $b = \binom{r}{2} / \binom{q}{2}$ blocks in \mathcal{B} and each point is contained in (r-1)/(q-1) of them. Moreover, any two blocks intersect in at most one point.

Conclusion: The blocks of a Steiner system S(2, q, s) form an almost disjoint (s-1)/(q-1)-cover of [s]. Therefore, when such Steiner systems exist, we have

$$p(s,k) \leq$$
number of blocks in $S(2,q,s) = \frac{s(s-1)}{q(q-1)} = \frac{s(k-1)^2}{s+k}$

where k = (s-1)/(q-1) + 1. But, by Fisher's inequality (# *blocks* \geq # *points*), this gives $\rho(s,k) \leq s$ at best. We can do much better with Steiner systems!

Let *V* be a set with *r* elements, called **points.** A *Steiner system* S(2, q, r) is a collection \mathcal{B} of subsets of *V* of size *q*, called **blocks**, such that every pair of points is contained in exactly one block.

• Such a system is an example of a *balanced incomplete block design*.



Lemma 10: PIR codes from Steiner systems

Let S(2, q, r) be a Steiner system. For each $v \in V$, let $A_v \subset B$ be the set of blocks that contain v. Then the sets $\{A_v : v \in V\}$ form an almost disjoint q-cover of [b].

Proof. For any pair of points *u* and *v*, there is only one block that contains both. Hence $|A_v \cap A_u| = 1$, and the sets $\{A_v : v \in V\}$ are almost disjoint.

Let *V* be a set with *r* elements, called **points.** A *Steiner system* S(2, q, r) is a collection \mathcal{B} of subsets of *V* of size *q*, called **blocks**, such that every pair of points is contained in exactly one block.

• Such a system is an example of a *balanced incomplete block design*.



Lemma 10: PIR codes from Steiner systems

Let S(2, q, r) be a Steiner system. For each $v \in V$, let $A_v \subset B$ be the set of blocks that contain v. Then the sets $\{A_v : v \in V\}$ form an almost disjoint *q*-cover of [b].

Proof. For any pair of points *u* and *v*, there is only one block that contains both. Hence $|A_v \cap A_u| = 1$, and the sets $\{A_v : v \in V\}$ are almost disjoint.

By Wilson's theorem, a Steiner system S(2, q, r) exists for all sufficiently large r whenever (q-1)|(r-1) and q(q-1)|r(r-1). Combining this theorem with Lemma 10 and Corollary 9, we have:

$$\rho(s,k) = O\left(\sqrt{s}\right)$$

for all fixed k

Let $\mathcal{G} = (U,V;\mathcal{E})$ be a bipartite graph, with bipartition U,V and edge set \mathcal{E} . We consider the neighborhoods $N(v) = \{u \in U : (u,v) \in \mathcal{E}\}$ of vertices in V.

Lemma 11: PIR codes from bipartite graphs

If \mathcal{G} has **no** 4-cycles, then the neighborhoods of vertices in V, namely the set $\{N(v) : v \in V\}$, form an almost disjoint k-cover of U, where $k = \min_{u \in U} \deg(u)$.

Proof. Assume to the contrary that there are vertices $v_1, v_2 \in V$ such that $|N(v_1) \cap N(v_2)| \ge 2$. Let u_1, u_2 be some two vertices in $N(v_1) \cap N(v_2)$. Then the induced subgraph on $\{v_1, v_2, u_1, u_2\}$ is $K_{2,2}$ which is a 4-cycle in \mathcal{G} .

Let $\mathcal{G} = (U,V;\mathcal{E})$ be a bipartite graph, with bipartition U,V and edge set \mathcal{E} . We consider the neighborhoods $N(v) = \{u \in U : (u,v) \in \mathcal{E}\}$ of vertices in V.

Lemma 11: PIR codes from bipartite graphs

If \mathcal{G} has **no** 4-cycles, then the neighborhoods of vertices in V, namely the set $\{N(v) : v \in V\}$, form an almost disjoint k-cover of U, where $k = \min_{u \in U} \deg(u)$.

Proof. Assume to the contrary that there are vertices $v_1, v_2 \in V$ such that $|N(v_1) \cap N(v_2)| \ge 2$. Let u_1, u_2 be some two vertices in $N(v_1) \cap N(v_2)$. Then the induced subgraph on $\{v_1, v_2, u_1, u_2\}$ is $K_{2,2}$ which is a 4-cycle in \mathcal{G} .

Let $\mathcal{G} = (U,V;\mathcal{E})$ be a bipartite graph, with bipartition U,V and edge set \mathcal{E} . We consider the neighborhoods $N(v) = \{u \in U : (u,v) \in \mathcal{E}\}$ of vertices in V.

Lemma 11: PIR codes from bipartite graphs

If \mathcal{G} has **no** 4-cycles, then the neighborhoods of vertices in V, namely the set $\{N(v) : v \in V\}$, form an almost disjoint k-cover of U, where $k = \min_{u \in U} \deg(u)$.

Proof. Assume to the contrary that there are vertices $v_1, v_2 \in V$ such that $|N(v_1) \cap N(v_2)| \ge 2$. Let u_1, u_2 be some two vertices in $N(v_1) \cap N(v_2)$. Then the induced subgraph on $\{v_1, v_2, u_1, u_2\}$ is $K_{2,2}$ which is a 4-cycle in \mathcal{G} .

• Given *s* and *k*, we would like to construct a bipartite graph $\mathcal{G} = (U,V;\mathcal{E})$ with the following properties:

|U| = s $\min_{u \in U} \deg(u) = k - 1$ $girth(\mathcal{G}) \ge 6$

If we can do this, then $\rho(s,k) \leq |V|$ by Corollary 10. What is the least possible number of vertices in *V* for such a graph?

Let $\mathcal{G} = (U,V;\mathcal{E})$ be a bipartite graph, with bipartition U,V and edge set \mathcal{E} . We consider the neighborhoods $N(v) = \{u \in U : (u,v) \in \mathcal{E}\}$ of vertices in V.

Lemma 11: PIR codes from bipartite graphs

If \mathcal{G} has **no** 4-cycles, then the neighborhoods of vertices in V, namely the set $\{N(v) : v \in V\}$, form an almost disjoint k-cover of U, where $k = \min_{u \in U} \deg(u)$.

Proof. Assume to the contrary that there are vertices $v_1, v_2 \in V$ such that $|N(v_1) \cap N(v_2)| \ge 2$. Let u_1, u_2 be some two vertices in $N(v_1) \cap N(v_2)$. Then the induced subgraph on $\{v_1, v_2, u_1, u_2\}$ is $K_{2,2}$ which is a 4-cycle in \mathcal{G} .

• Given *s* and *k*, we would like to construct a bipartite graph $\mathcal{G} = (U,V;\mathcal{E})$ with the following properties:

$$|U| = s$$
 $\min_{u \in U} \deg(u) = k - 1$ $girth(\mathcal{G}) \ge 6$

If we can do this, then $\rho(s,k) \leq |V|$ by Corollary 10. What is the least possible number of vertices in *V* for such a graph? Using the best known results on bipartite cages, we get:

$$ho(s,k) = O\left(\sqrt{s}\right)$$
 for all

for all fixed k

Definition: Constant-weight codes

Let $A_2(n, d, w)$ be the number of codewords in the largest binary code \mathbb{C} of length n and minimum distance d such that all the codewords of \mathbb{C} have weight w.

Definition: Constant-weight codes

Let $A_2(n, d, w)$ be the number of codewords in the largest binary code \mathbb{C} of length n and minimum distance d such that all the codewords of \mathbb{C} have weight w.

To learn more about constant-weight codes and their properties, consult the following papers:

Ch. Bachoc, V. Chandar, G.D. Cohen, P. Solé, and A. Tchamkerten, On bounded weight codes, *IEEE Trans. Information Theory*, **57**, pp. 6780–6787, October 2011.

G.D. Cohen, P. Solé, and A. Tchamkerten, Heavy weight codes, *Proceedings IEEE International Symp. Information Theory*, pp. 1120–1124, Austin, TX., June 2010.

Definition: Constant-weight codes

Let $A_2(n, d, w)$ be the number of codewords in the largest binary code \mathbb{C} of length n and minimum distance d such that all the codewords of \mathbb{C} have weight w.

Observations:

• d = 2w if and only if any two codewords have *disjoint supports*:

Definition: Constant-weight codes

Let $A_2(n, d, w)$ be the number of codewords in the largest binary code \mathbb{C} of length n and minimum distance d such that all the codewords of \mathbb{C} have weight w.

Observations:

• d = 2w if and only if any two codewords have *disjoint supports*:

• d = 2w - 2 iff any two codewords intersect in at most one position:

Definition: Constant-weight codes

Let $A_2(n, d, w)$ be the number of codewords in the largest binary code \mathbb{C} of length n and minimum distance d such that all the codewords of \mathbb{C} have weight w.

Observations:

• d = 2w if and only if any two codewords have *disjoint supports*:

• d = 2w - 2 iff any two codewords *intersect in at most one position*:



Now let $s = A_2(n, 2w-2, w)$, and consider the $s \times n$ matrix having the codewords of **C** as its rows: n



As the weight of each row is *w*, columns form a *w*-cover of [*s*].

Definition: Constant-weight codes

Let $A_2(n, d, w)$ be the number of codewords in the largest binary code \mathbb{C} of length n and minimum distance d such that all the codewords of \mathbb{C} have weight w.

Observations:

• d = 2w if and only if any two codewords have *disjoint supports*:

• d = 2w - 2 iff any two codewords *intersect in at most one position*:



Now let $s = A_2(n, 2w-2, w)$, and consider the $s \times n$ matrix having the codewords of **C** as its rows: n



As the row supports are almost disjoint, the column supports are also almost disjoint.

Definition: Constant-weight codes

Let $A_2(n, d, w)$ be the number of codewords in the largest binary code \mathbb{C} of length n and minimum distance d such that all the codewords of \mathbb{C} have weight w.

Now let $s = A_2(n, 2w-2, w)$, and consider the $s \times n$ matrix having the codewords of \mathbb{C} as its rows: n



As the row supports are almost disjoint, the column supports are also almost disjoint.

Theorem 12: PIR codes from constant-weight codes

 $\rho(s,k) \leq \text{ the smallest } n \text{ such that } A_2(n, 2k-4, k-1) \geq s$ supports are almost disjoint (k-1)-cover of [s]

Definition: Constant-weight codes

Let $A_2(n, d, w)$ be the number of codewords in the largest binary code \mathbb{C} of length n and minimum distance d such that all the codewords of \mathbb{C} have weight w.

Now let $s = A_2(n, 2w-2, w)$, and consider the $s \times n$ matrix having the codewords of \mathbb{C} as its rows: n



As the row supports are almost disjoint, the column supports are also almost disjoint.

Theorem 12: PIR codes from constant-weight codes

 $\rho(s,k) \leq \text{ the smallest } n \text{ such that } A_2(n, 2k-4, k-1) \geq s$ supports are almost disjoint (k-1)-cover of [s]

For example, for k = 3 we conclude that $\rho(s, 3)$ is upper bounded by the smallest n such that $n(n-1) \ge 2s$. In general, we again have $\rho(s, k) = O(\sqrt{s})$.

number k of servers emulated

		2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
	1	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
	2	1	3	4	6	7	9	10	12	13	15	16	18	19	21	22
-	3	1	3	4	7	8	10	11	14	15	17	18	21	22	24	25
ť	4	1	4	5	7	8	10	11	15	16	19	20	22	23	25	26
a	5	1	4	5	7	8	13	14	17	18	20	21	23	24	25	26
d	6	1	4	5	7	8	14	15	18	19	21	22	28	29	32	33
e	7	1	5	6	7	8	15	16	20	21	22	23	30	31	35	36
3S	8	1	5	6	11	12	15	16	24	25	29	30	35	36	39	40
ق	9	1	5	6	12	13	15	16	25	26	30	31	37	38	40	41
D	10	1	5	6	13	14	15	16	26	27	31	32	39	40	41	42
a	11	1	6	7	13	14	21	22	30	31	37	38	39	40	41	42
σ	12	1	6	7	13	14	21	22	30	31	37	38	39	40	41	42
Ч.	13	1	6	7	13	14	21	22	30	31	37	38	39	40	41	42
•	14	1	6	7	14	15	21	22	30	31	37	38	39	40	41	42
S	15	1	6	7	15	16	21	22	30	31	37	38	39	40	41	42
er	16	1	7	8	16	17	21	22	30	31	45	46	51	52	55	56
ھ	17	1	7	8	16	17	21	22	30	31	46	47	55	56	60	61
B	18	1	7	8	16	17	21	22	30	31	47	48	56	57	61	62
Ξ	19	1	7	8	16	17	21	22	30	31	48	49	57	58	62	63
ā	20	1	7	8	16	17	21	22	30	31	49	50	58	59	63	64
	21	1	7	8	18	19	27	28	30	31	52	53	59	60	70	71
	22	1	8	9	18	19	28	29	30	31	53	54	61	62	71	72
	23	1	8	9	19	20	28	29	30	31	54	55	62	63	73	74
	24	1	8	9	19	20	28	29	30	31	55	56	63	64	74	75

Redundancy $\rho(s, k)$ of the best-known PIR codes

number k of servers emulated

		2	3	4	5	6	7	8	9	10	11	12	13	14	15	16			
	1	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	1		
	2	1	3	4	6	7	9	10	12	13	15	16	18	19	21	22	1		
	3	1	3	4	7	8	10	11	14	15	17	18	21	22	24	25	1		
Ë	4	1	4	5	7	8	10	11	15	16	19	20	22	23	25	26			
a	5	1	4	5	7	8	13	14	17	18	20	21	23	24	25	26			
P	6	1	4	5	7	8	14	15	18	19	21	22	28	29	32	33	•	0.	
e	7	1	5	6	7	8	15	16	20	21	22	23	30	31	35	31	.h)		
S	8	1	5	6	11	12	15	16	24	25	29	30	35	36	30	× *	22		P
ŏ	9	1	5	6	12	13	15	16	25	26	30	31	37	38	4.4	S			
D	10	1	5	6	13	14	15	16	26	27	31	32	39		11	×-			
a	11	1	6	7	13	14	21	22	30	31	37	38	< · ·	5		12			
D	12	1	6	7	13	14	21	22	30	31	37	<u></u>	1	Y	.1	42			
H	13	1	6	7	13	14	21	22	30	31	25	~	Y)	5	41	42	I		
	14	1	6	7	14	15	21	22	30	31	. 6	21.		40	41	42			
s	15	1	6	7	15	16	21	22	30				- 39	40	41	42			
1	16	1	7	8	16	17	21	22	-	31	- 5	46	51	52	55	56	1		
6	17	1	7	8	16	17	21		1e	×.	±6	47	55	56	60	61	I		
	18	1	7	8	16	17	ĺ.,	~ 0	•	11	47	48	56	57	61	62			
	19	1	7	8	16	1	<u>~</u> 0) × _	J	31	48	49	57	58	62	63			
Ξ	20	1	7	8	16	11	7.1		30	31	49	50	58	59	63	64			
	21	1	7	8	18	~ *		28	30	31	52	53	59	60	70	71	1		
	22	1	8	9	18	1.	28	29	30	31	53	54	61	62	71	72			
	23	1	8	9	19	20	28	29	30	31	54	55	62	63	73	74			
	24	1	8	9	19	20	28	29	30	31	55	56	63	64	74	75			

Redundancy $\rho(s,k)$ of the best-known PIR codes

number *s* of database parts

number k of servers emulated

	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	2.00	3.00	4.00	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00	16.00
2	1.50	2.50	3.00	4.00	4.50	5.50	6.00	7.00	7.50	8.50	9.00	10.00	10.50	11.50	12.00
3	1.33	2.00	2.33	3.33	3.67	4.33	4.67	5.67	6.00	6.67	7.00	8.00	8.33	9.00	9.33
4	1.25	2.00	2.25	2.75	3.00	3.50	3.75	4.75	5.00	5.75	6.00	6.50	6.75	7.25	7.50
5	1.20	1.80	2.00	2.40	2.60	3.60	3.80	4.40	4.60	5.00	5.20	5.60	5.80	6.00	6.20
6	1.17	1.67	1.83	2.17	2.33	3.33	3.50	4.00	4.17	4.50	4.67	5.67	5.83	6.33	6.50
7	1.14	1.71	1.86	2.00	2.14	3.14	3.29	3.86	4.00	4.14	4.29	5.29	5.43	6.00	6.14
8	1.13	1.63	1.75	2.38	2.50	2.88	3.00	4.00	4.13	4.63	4.75	5.38	5.50	5.88	6.00
9	1.11	1.56	1.67	2.33	2.44	2.67	2.78	3.78	3.89	4.33	4.44	5.11	5.22	5.44	5.56
10	1.10	1.50	1.60	2.30	2.40	2.50	2.60	3.60	3.70	4.10	4.20	4.90	5.00	5.10	5.20
11	1.09	1.55	1.64	2.18	2.27	2.91	3.00	3.73	3.82	4.36	4.45	4.55	4.64	4.73	4.82
12	1.08	1.50	1.58	2.08	2.17	2.75	2.83	3.50	3.58	4.08	4.17	4.25	4.33	4.42	4.50
13	1.08	1.46	1.54	2.00	2.08	2.62	2.69	3.31	3.38	3.85	3.92	4.00	4.08	4.15	4.23
14	1.07	1.43	1.50	2.00	2.07	2.50	2.57	3.14	3.21	3.64	3.71	3.79	3.86	3.93	4.00
15	1.07	1.40	1.47	2.00	2.07	2.40	2.47	3.00	3.07	3.47	3.53	3.60	3.67	3.73	3.80
16	1.06	1.44	1.50	2.00	2.06	2.31	2.38	2.88	2.94	3.81	3.88	4.19	4.25	4.44	4.50
17	1.06	1.41	1.47	1.94	2.00	2.24	2.29	2.76	2.82	3.71	3.76	4.24	4.29	4.53	4.59
18	1.06	1.39	1.44	1.89	1.94	2.17	2.22	2.67	2.72	3.61	3.67	4.11	4.17	4.39	4.44
19	1.05	1.37	1.42	1.84	1.89	2.11	2.16	2.58	2.63	3.53	3.58	4.00	4.05	4.26	4.32
20	1.05	1.35	1.40	1.80	1.85	2.05	2.10	2.50	2.55	3.45	3.50	3.90	3.95	4.15	4.20
21	1.05	1.33	1.38	1.86	1.90	2.29	2.33	2.43	2.48	3.48	3.52	3.81	3.86	4.33	4.38
22	1.05	1.36	1.41	1.82	1.86	2.27	2.32	2.36	2.41	3.41	3.45	3.77	3.82	4.23	4.27
23	1.04	1.35	1.39	1.83	1.87	2.22	2.26	2.30	2.35	3.35	3.39	3.70	3.74	4.17	4.22
24	1.04	1.33	1.38	1.79	1.83	2.17	2.21	2.25	2.29	3.29	3.33	3.63	3.67	4.08	4.13

Storage overhead of the best-known PIR codes

number k of servers emulated

	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	2.00	3.00	4.00	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00	16.00
2	1.50	2.50	3.00	4.00	4.50	5.50	6.00	7.00	7.50	8.50	9.00	10.00	10.50	11.50	12.00
3	1.33	2.00	2.33	3.33	3.67	4.33	4.67	5.67	6.00	6.67	7.00	8.00	8.33	9.00	9.33
4	1.25	2.00	2.25	2.75	3.00	3.50	3.75	4.75	5.00	5.75	6.00	6.50	6.75	7.25	7.50
5	1.20	1.80	2.00	2.40	2.60	3.60	3.80	4.40	4.60	5.00	5.20	5.60	5.80	6.00	6.20
6	1.17	1.67	1.83	2.17	2.33	3.33	3.50	4.00	4.17	4.50	4.67	5.67	5.83	6.33	6.50
7	1.14	1.71	1.86	2.00	2.14	3.14	3.29	3.86	4.00	4.14	4.29	5.29	5.43	6.00	6.14
8	1.13	1.63	1.75	2.38	2.50	2.88	3.00	4.00	4.13	4.63	4.75	5.38	5.50	5.88	6.00
9	1.11	1.56	1.67	2.33	2.44	2.67	2.78	3.78	3.89	4.33	4.44	5.11	5.22	5.44	5.56
10	1.10	1.50	1.60	2.30	2.40	2.50	2.60	3.60	3.70	4.10	4.20	4.90	5.00	5.10	5.20
11	1.09	1.55	1.64	2.18	2.27	2.91	3.00	3.73	3.82	4.36	4.45	4.55	4.64	4.73	4.82
12	1.08	1.50	1.58	2.08	2.17	2.75	2.83	3.50	3.58	4.08	4.17	4.25	4.33	4.42	4.50
13	1.08	1.46	1.54	2.00	2.08	2.62	2.69	3.31	3.38	3.85	3.92	4.00	4.08	4.15	4.23
14	1.07	1.43	1.50	2.00	2.07	2.50	2.57	3.14	3.21	3.64	3.71	3.79	3.86	3.93	4.00
15	1.07	1.40	1.47	2.00	2.07	2.40	2.47	3.00	3.07	3.47	3.53	3.60	3.67	3.73	3.80
16	1.06	1.44	1.50	2.00	2.06	2.31	2.38	2.88	2.94	3.81	3.88	4.19	4.25	4.44	4.50
17	1.06	1.41	1.47	1.94	2.00	2.24	2.29	2.76	2.82	3.71	3.76	4.24	4.29	4.53	4.59
18	1.06	1.39	1.44	1.89	1.94	2.17	2.22	2.67	2.72	3.61	3.67	4.11	4.17	4.39	4.44
19	1.05	1.37	1.42	1.84	1.89	2.11	2.16	2.58	2.63	3.53	3.58	4.00	4.05	4.26	4.32
20	1.05	1.35	1.40	1.80	1.85	2.05	2.10	2.50	2.55	3.45	3.50	3.90	3.95	4.15	4.20
21	1.05	1.33	1.38	1.86	1.90	2.29	2.33	2.43	2.48	3.48	3.52	3.81	3.86	4.33	4.38
22	1.05	1.36	1.41	1.82	1.86	2.27	2.32	2.36	2.41	3.41	3.45	3.77	3.82	4.23	4.27
23	1.04	1.35	1.39	1.83	1.87	2.22	2.26	2.30	2.35	3.35	3.39	3.70	3.74	4.17	4.22
24	1.04	1.33	1.38	1.79	1.83	2.17	2.21	2.25	2.29	3.29	3.33	3.63	3.67	4.08	4.13

Storage overhead of the best-known PIR codes

Thank you for your attention!



Please send you queries...