

The Matrix Mechanism: Optimizing Linear Counting Queries Under Differential Privacy

Gerome Miklau

Univ. of Massachusetts, Amherst

Joint work with:

Chao Li *Univ. of Massachusetts, Amherst*

Andrew McGregor *Univ. of Massachusetts, Amherst*

Michael Hay *Cornell University*

Vibhor Rastogi *Yahoo! Research*

Dan Suciu *University of Washington*




June 30, 2011

Data analysis under differential privacy

- The **differential guarantee** for participants in a data set:
 - Information released about a private data set is virtually indistinguishable whether or not a participant's data is included.
- Resistant to informed adversaries.
- Precise (public) error bounds on private output.

A central open question: what are **utility-optimal** mechanisms satisfying differential privacy?


Differentially private mechanisms



An optimal mechanism is known for answering a single query

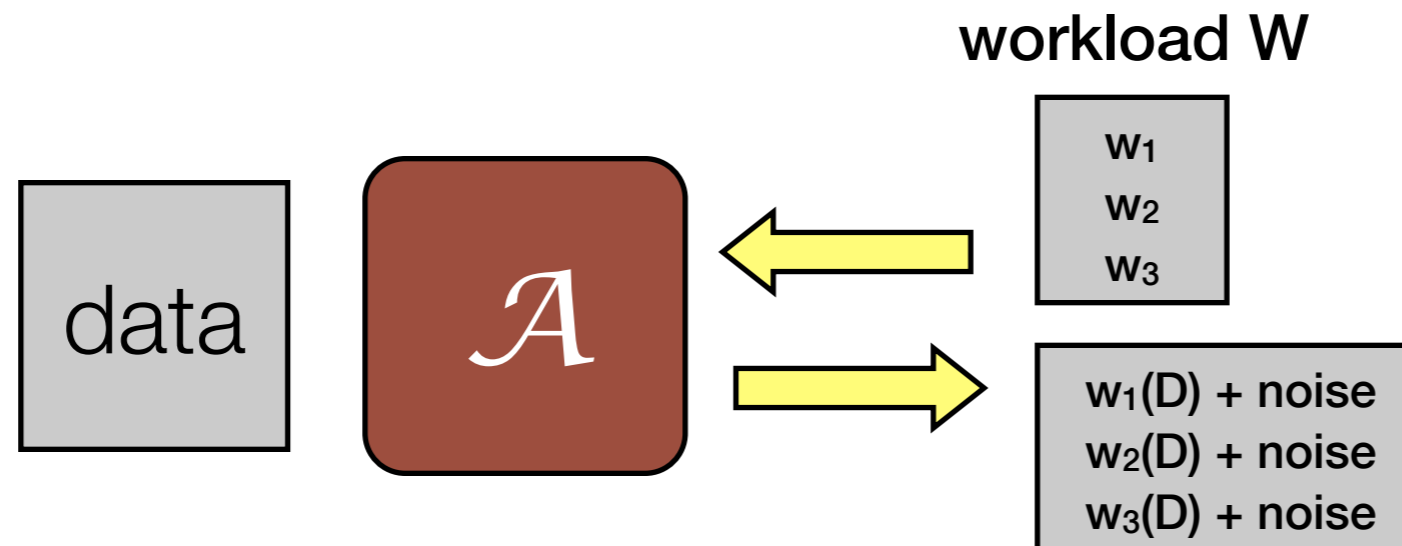
[Ghosh, 2009]

Differentially private mechanisms



An optimal mechanism is known for answering a single query

[Ghosh, 2009]



This mechanism is often sub-optimal for multiple queries

Query workloads

Query workloads

- A query workload is a set of **linear counting queries**, known ahead of time.
 - may include predicate counting queries, range queries, data cubes, sets of marginals, CDFs, etc...

Query workloads

- A query workload is a set of **linear counting queries**, known ahead of time.
 - may include predicate counting queries, range queries, data cubes, sets of marginals, CDFs, etc...
- How do query workloads arise ?
 - ... from decomposing a more complex data analysis task into a set of queries.
 - ... from multiple users accessing sensitive data, each issuing one or more queries.
 - ... from uncertainty about the eventual query answers needed--design workload to include all queries possibly of interest.

Query workloads

- A query workload is a set of **linear counting queries**, known ahead of time.
 - may include predicate counting queries, range queries, data cubes, sets of marginals, CDFs, etc...
- How do query workloads arise ?
 - ... from decomposing a more complex data analysis task into a set of queries.
 - ... from multiple users accessing sensitive data, each issuing one or more queries.
 - ... from uncertainty about the eventual query answers needed--design workload to include all queries possibly of interest.
- Our output can be treated as a **synthetic data set**; one which is designed to provide particularly accurate answers for the given workload queries.

Privacy definitions & mechanisms

Privacy definitions & mechanisms

- Differential privacy

Privacy definitions & mechanisms

- Differential privacy

A randomized algorithm \mathcal{A} provides (ϵ, δ) -**differential privacy** if:
for all neighboring databases D and D' , and
for any set of outputs S :

$$\Pr[\mathcal{A}(D) \in S] \leq e^\epsilon \Pr[\mathcal{A}(D') \in S] + \delta$$

Privacy definitions & mechanisms

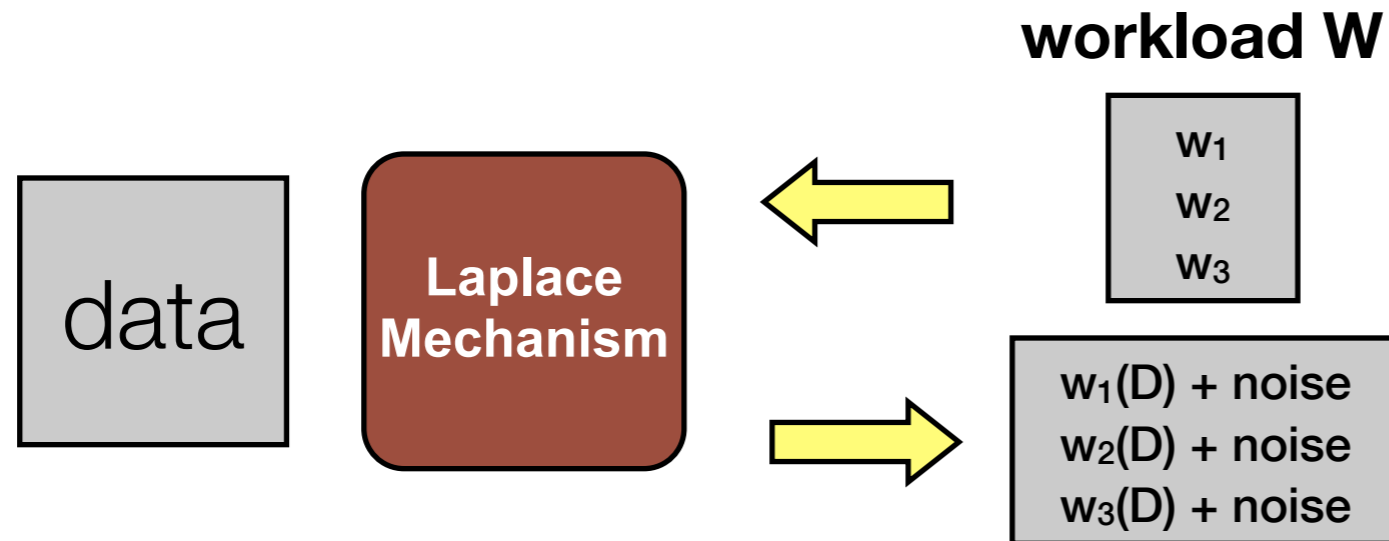
- Differential privacy

A randomized algorithm \mathcal{A} provides (ϵ, δ) -**differential privacy** if:
for all neighboring databases D and D' , and
for any set of outputs S :

$$Pr[\mathcal{A}(D) \in S] \leq e^\epsilon Pr[\mathcal{A}(D') \in S] + \delta$$

- if $\delta=0$, standard ϵ -differential privacy
 - **Laplace(0,b) noise where $b=\|q\|_1/\epsilon$**
- if $\delta>0$, approximate (ϵ, δ) -differential privacy:
 - **Gaussian(0, σ) noise where $\sigma=\|q\|_2 f(\delta)/\epsilon$**

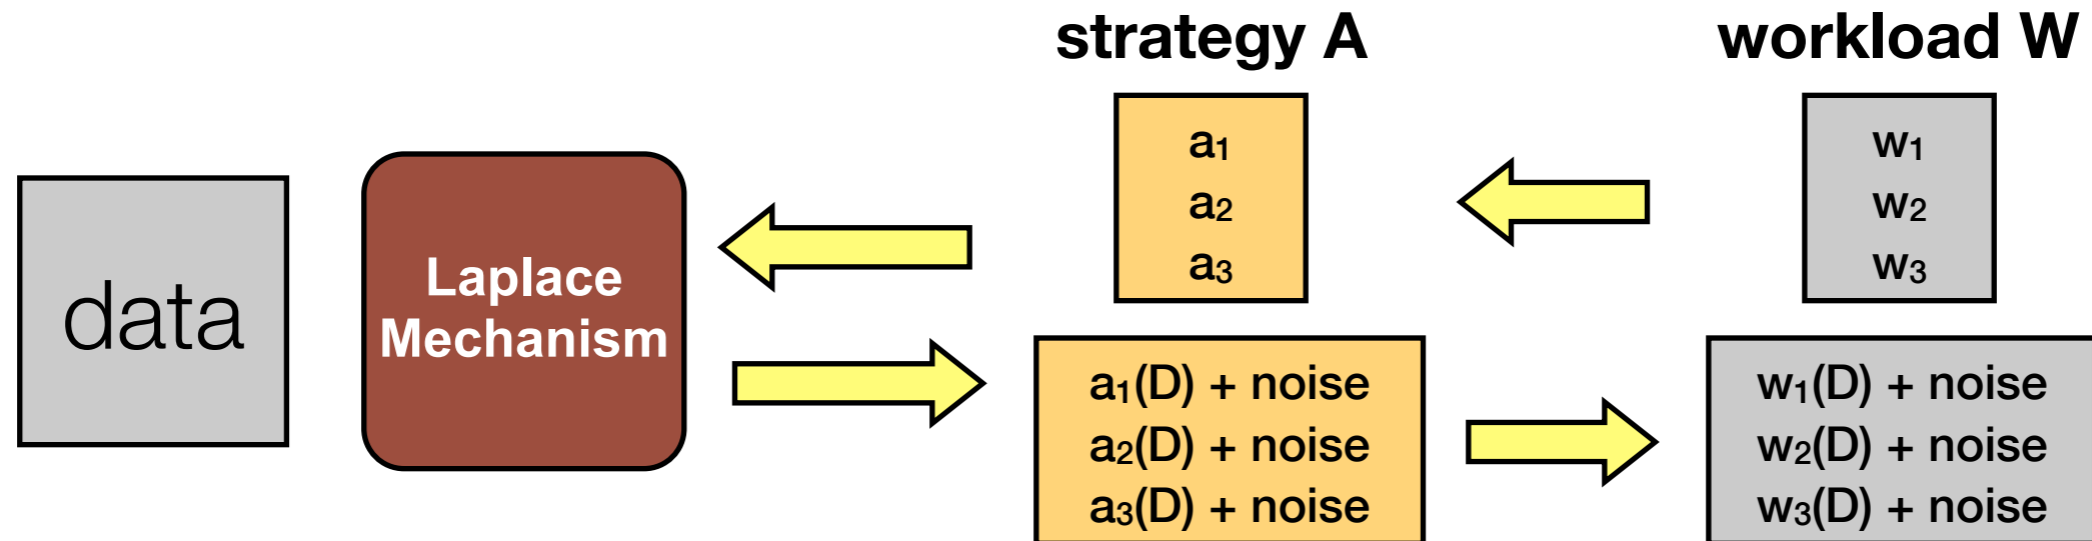
Main approach



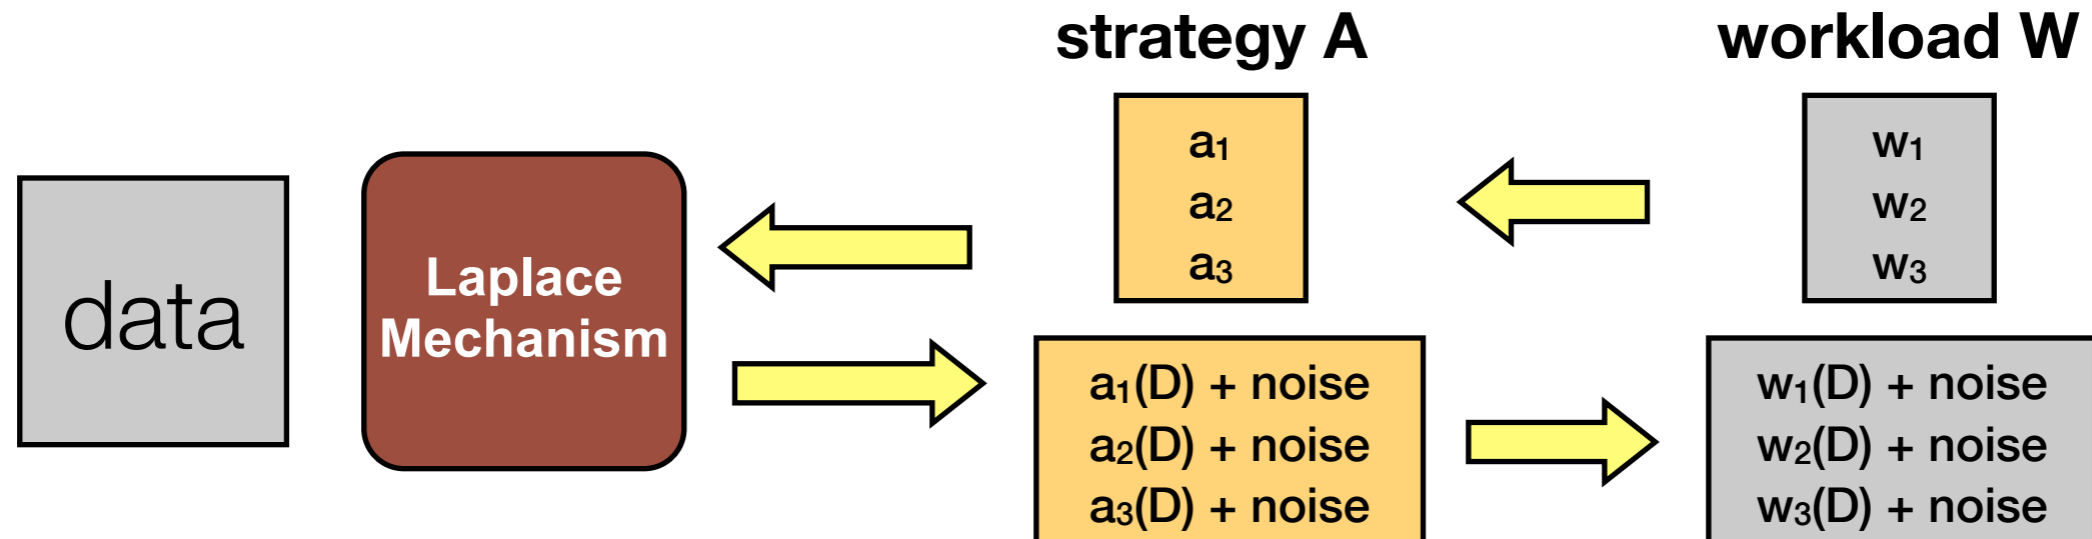
Main approach



Main approach

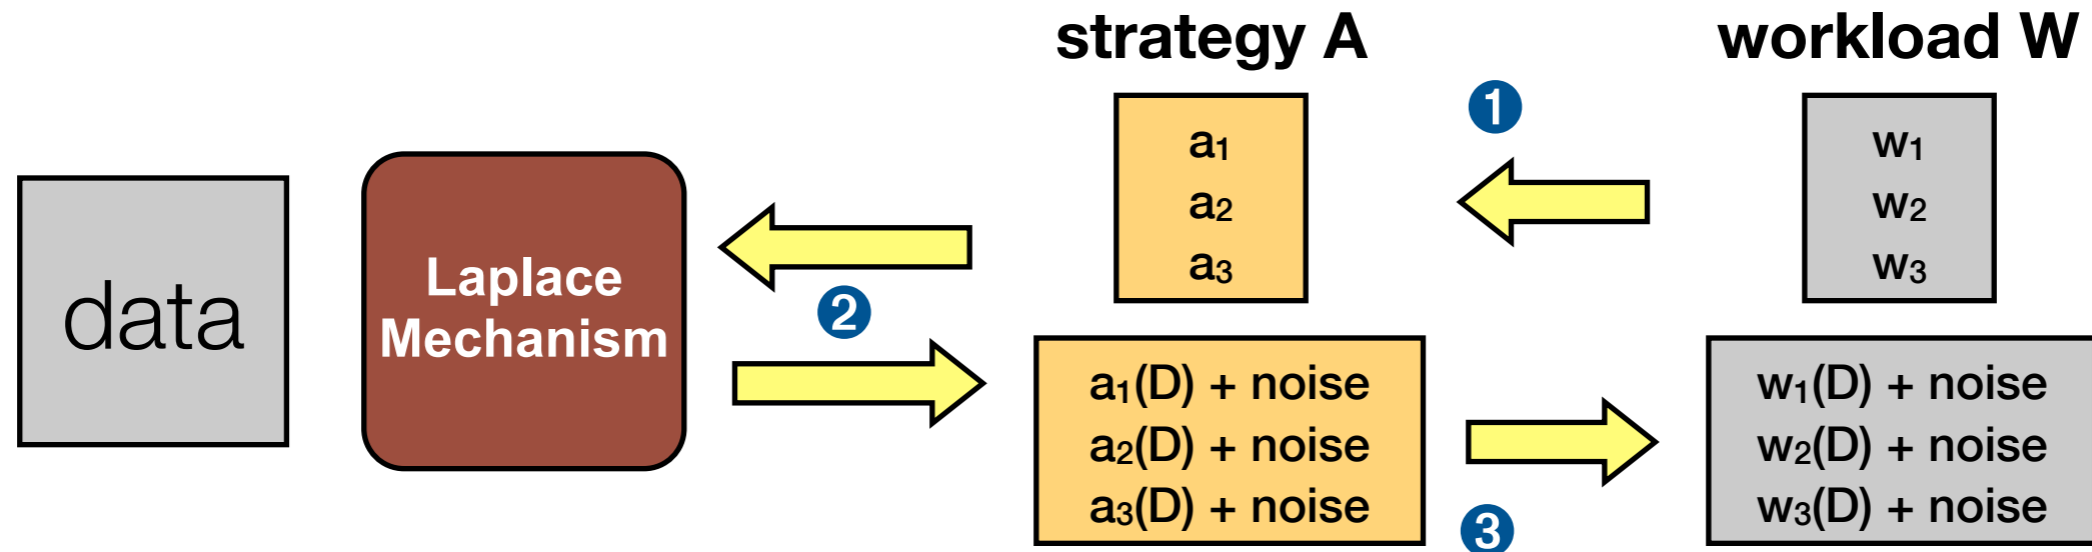


Main approach



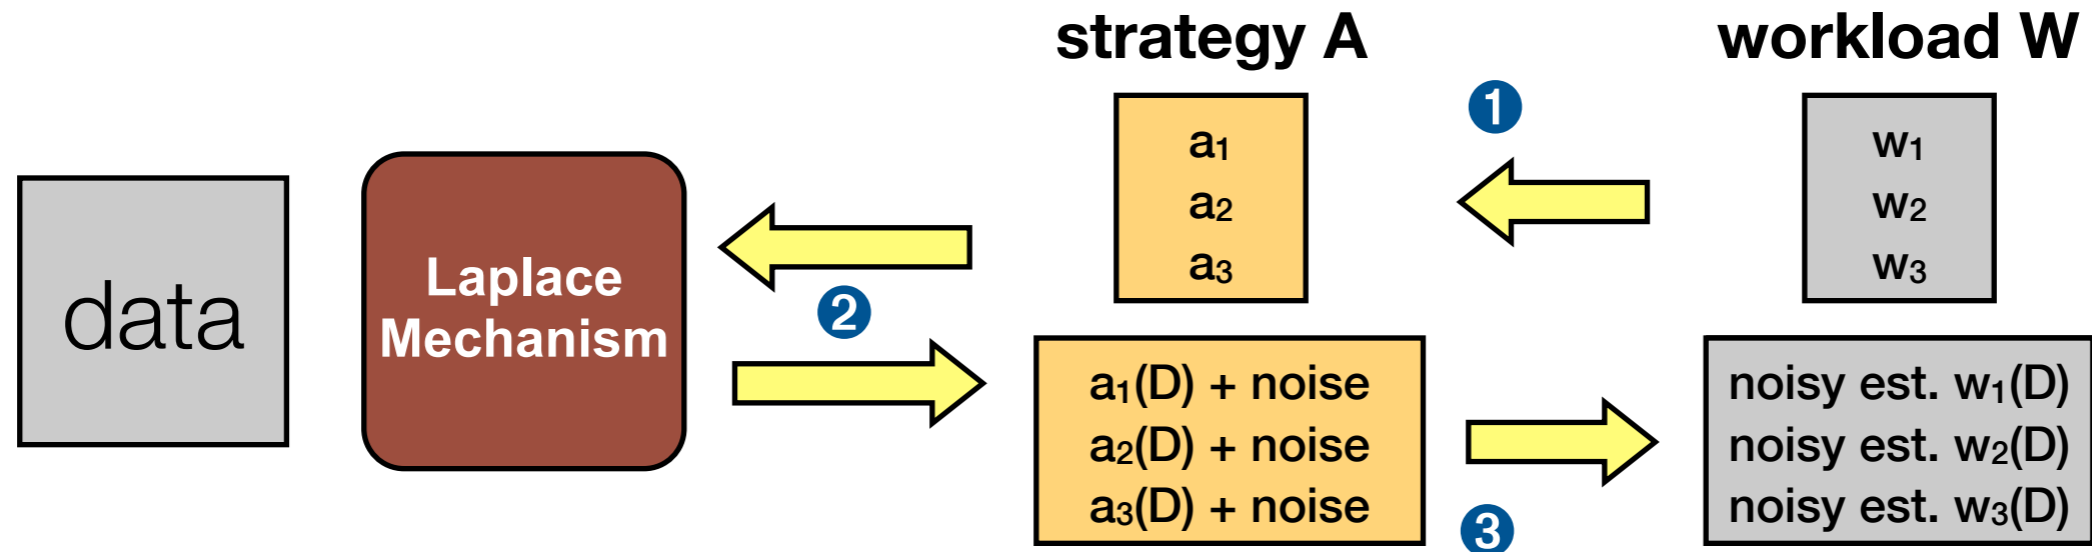
- 1 **(Design)** Choose a set of queries **A** (the strategy)
- 2 **(Apply Laplace)** Use the Laplace mechanism to answer **A**
- 3 **(Derivation)** Compute each query in **W** using answers to **A**

Main approach



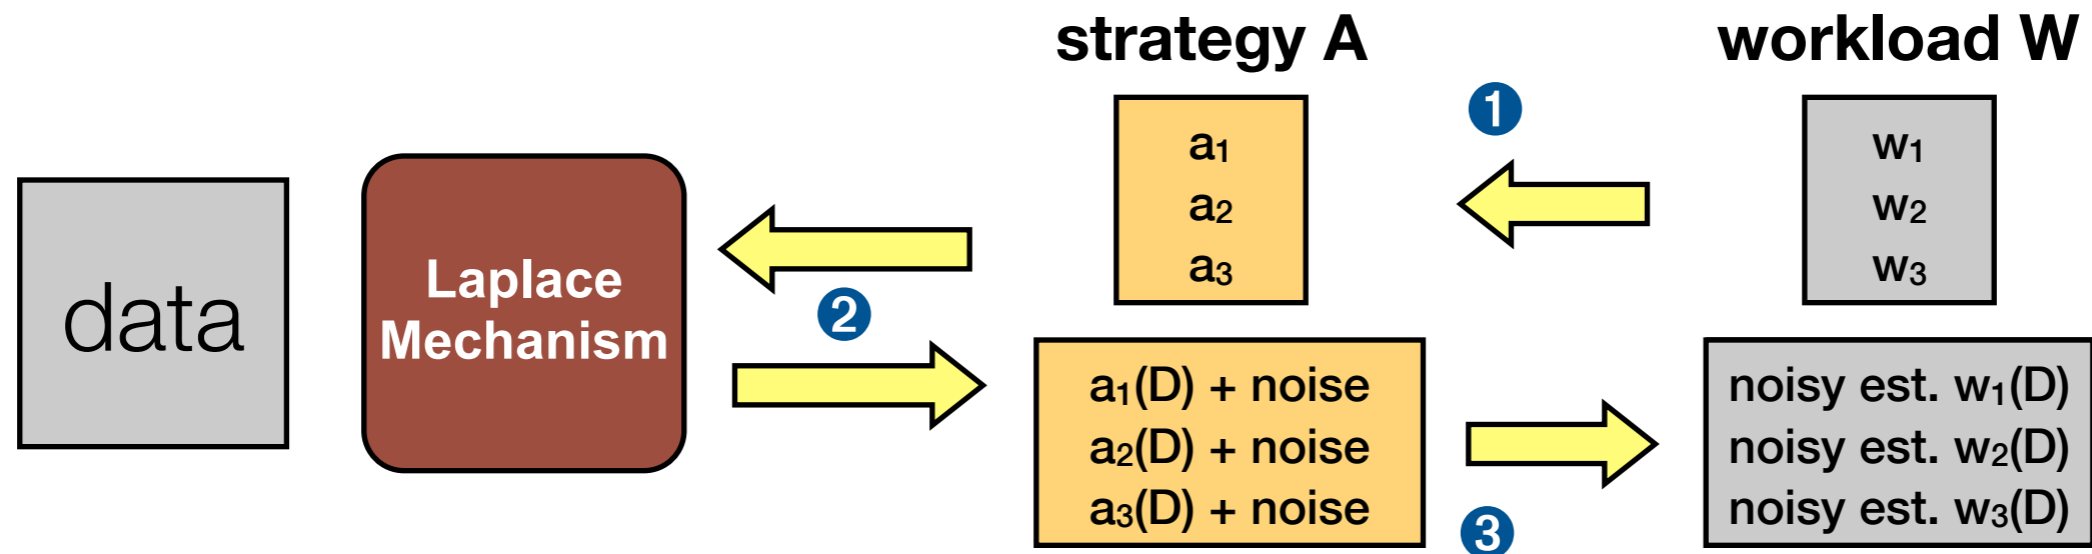
- ① **(Design)** Choose a set of queries A (the strategy)
- ② **(Apply Laplace)** Use the Laplace mechanism to answer A
- ③ **(Derivation)** Compute each query in W using answers to A

Main approach



- ① **(Design)** Choose a set of queries A (the strategy)
- ② **(Apply Laplace)** Use the Laplace mechanism to answer A
- ③ **(Derivation)** Compute each query in W using answers to A

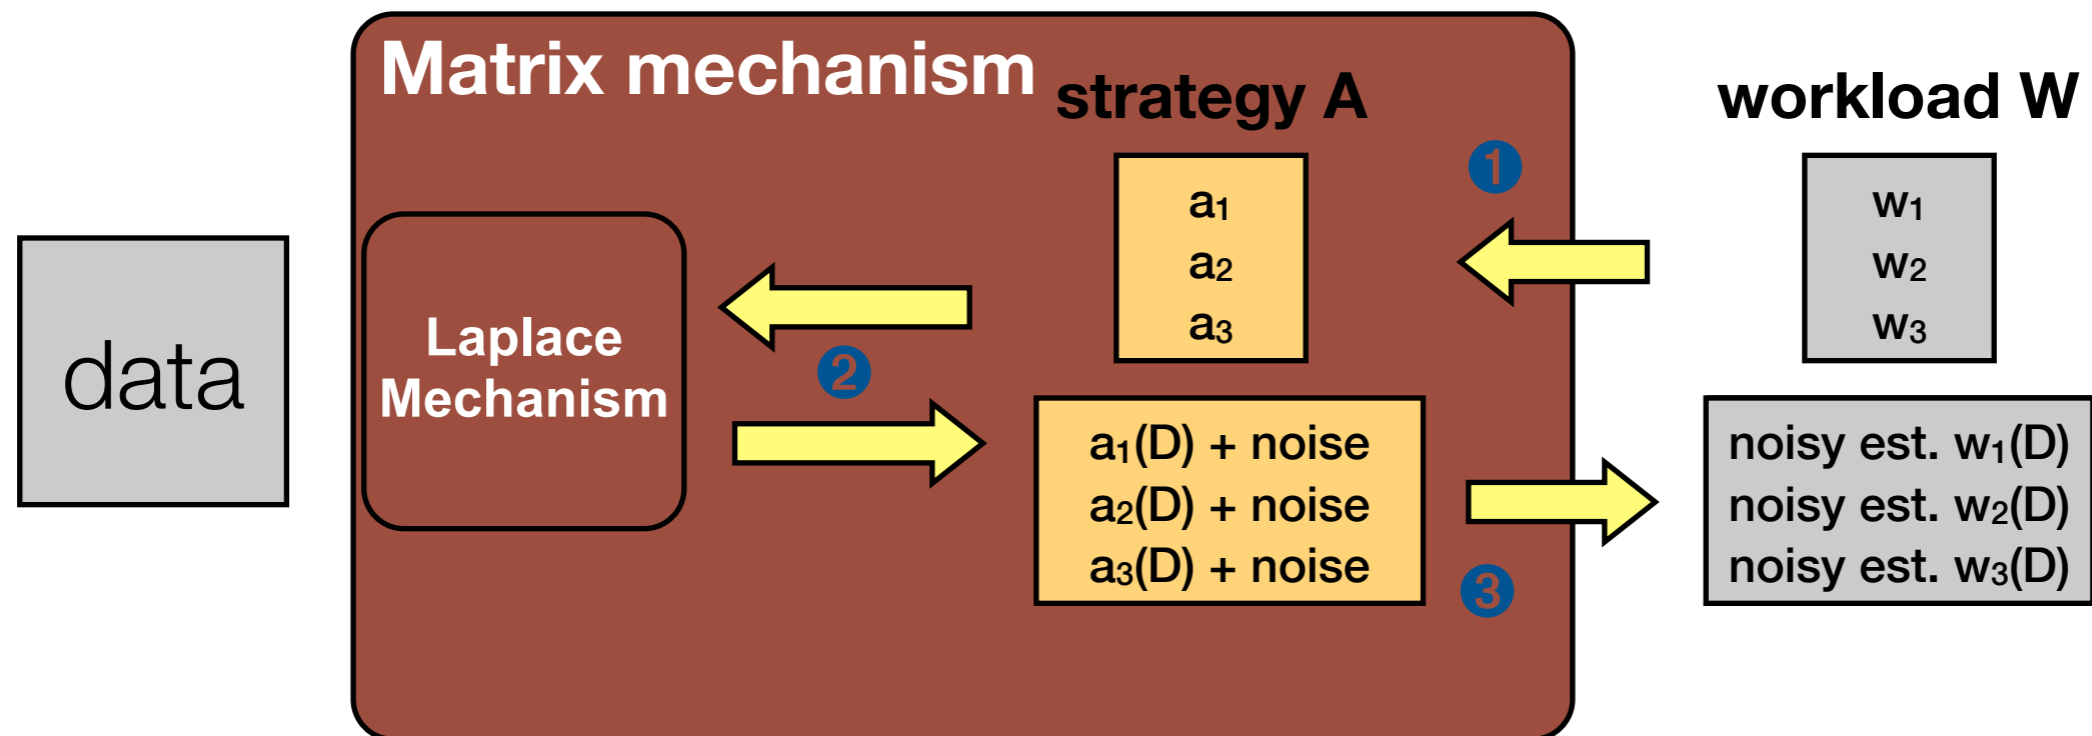
Main approach



- 1 (Design) Choose a set of queries **A** (the strategy)
- 2 (Apply Laplace) Use the Laplace mechanism to answer **A**
- 3 (Derivation) Compute each query in **W** using answers to **A**

Opportunity: choose **A** to minimize the error for queries in **W**.

Main approach



- 1 (Design) Choose a set of queries **A** (the strategy)
- 2 (Apply Laplace) Use the Laplace mechanism to answer **A**
- 3 (Derivation) Compute each query in **W** using answers to **A**

Opportunity: choose **A** to minimize the error for queries in **W**.

Outline

1. Case study: answering 1-dim range queries
2. The matrix mechanism -- formal description.
3. The matrix mechanism -- new tools & techniques.
4. Conclusion

Frequency representation of the database

name	gender	grade
Alice	Female	91
Bob	Male	84
Carl	Male	82
Dave	Male	97
Edwina	Female	88
Faith	Female	78
Ghita	Female	85
...

Relational database

{gender, grade}

gender	grade	count
Male	100	10
Male	99	13
Male	98	5
Male	97	7
...
Female	100	15
Female	99	21
Female	98	4
Female	97	14
Female	96	9

Frequency vector

x_1
x_2
x_3
x_4
x_5
x_6
x_7
x_8
...
x_n

x

Adding/removing tuple changes component frequency exactly 1

Frequency vector length n , where n is "domain size" important part of the discussion

Frequency representation of the database

name	gender	grade
Alice	Female	91
Bob	Male	84
Carl	Male	82
Dave	Male	97
Edwina	Female	88
Faith	Female	78
Ghita	Female	85
...

Relational database

Frequency vector

x

Adding/removing
tuple changes
component
frequency
exactly 1

Frequency vector
length n, where
"domain size"
important part
the discussion

Frequency representation of the database

name	gender	grade
Alice	Female	91
Bob	Male	84
Carl	Male	82
Dave	Male	97
Edwina	Female	88
Faith	Female	78
Ghita	Female	85
...

Relational database

{grade}

grade	count
90-100	10
80-90	23
70-80	16
60-70	3

Frequency vector

x ₁
x ₂
x ₃
x ₄

x

Adding/removing tuple changes component frequency exactly 1

Frequency vector length n, where "domain size" is an important parameter in the discussion

Answering all range queries

Goal: answer all **range-count queries** over x

$$\text{AllRange} = \{ w \mid w = x_i + \dots + x_j \text{ for } 1 \leq i \leq j \leq n \}$$

workload W

w ₁	range(x ₁ ,x ₄)
w ₂	range(x ₁ ,x ₃)
w ₃	range(x ₂ ,x ₄)
w ₄	range(x ₁ ,x ₂)
w ₅	range(x ₂ ,x ₃)
w ₆	range(x ₃ ,x ₄)
w ₇	range(x ₁ ,x ₁)
w ₈	range(x ₂ ,x ₂)
w ₉	range(x ₃ ,x ₃)
w ₁₀	range(x ₄ ,x ₄)

x ₁	+	x ₂	+	x ₃	+	x ₄
x ₁	+	x ₂	+	x ₃		
		x ₂	+	x ₃	+	x ₄
x ₁	+	x ₂				
		x ₂	+	x ₃		
				x ₃	+	x ₄
x ₁						
		x ₂				
				x ₃		
						x ₄

Answering all range queries

Goal: answer all **range-count queries** over x

$$\text{AllRange} = \{ w \mid w = x_i + \dots + x_j \text{ for } 1 \leq i \leq j \leq n \}$$

workload W

w_1	$\text{range}(x_1, x_4)$	x_1	+	x_2	+	x_3	+	x_4
w_2	$\text{range}(x_1, x_3)$	x_1	+	x_2	+	x_3		
w_3	$\text{range}(x_2, x_4)$			x_2	+	x_3	+	x_4
w_4	$\text{range}(x_1, x_2)$	x_1	+	x_2				
w_5	$\text{range}(x_2, x_3)$			x_2	+	x_3		
w_6	$\text{range}(x_3, x_4)$					x_3	+	x_4
w_7	$\text{range}(x_1, x_1)$	x_1						
w_8	$\text{range}(x_2, x_2)$			x_2				
w_9	$\text{range}(x_3, x_3)$					x_3		
w_{10}	$\text{range}(x_4, x_4)$							x_4

$X =$

10	23	16	3
----	----	----	---

Answering all range queries

Goal: answer all **range-count queries** over x

$$\text{AllRange} = \{ w \mid w = x_i + \dots + x_j \text{ for } 1 \leq i \leq j \leq n \}$$

workload W

w_1	range(x_1, x_4)	x_1	+	x_2	+	x_3	+	x_4	w_1	52
w_2	range(x_1, x_3)	x_1	+	x_2	+	x_3			w_2	49
w_3	range(x_2, x_4)			x_2	+	x_3	+	x_4	w_3	42
w_4	range(x_1, x_2)	x_1	+	x_2					w_4	33
w_5	range(x_2, x_3)			x_2	+	x_3			w_5	39
w_6	range(x_3, x_4)					x_3	+	x_4	w_6	19
w_7	range(x_1, x_1)	x_1							w_7	10
w_8	range(x_2, x_2)			x_2					w_8	23
w_9	range(x_3, x_3)					x_3			w_9	16
w_{10}	range(x_4, x_4)							x_4	w_{10}	3

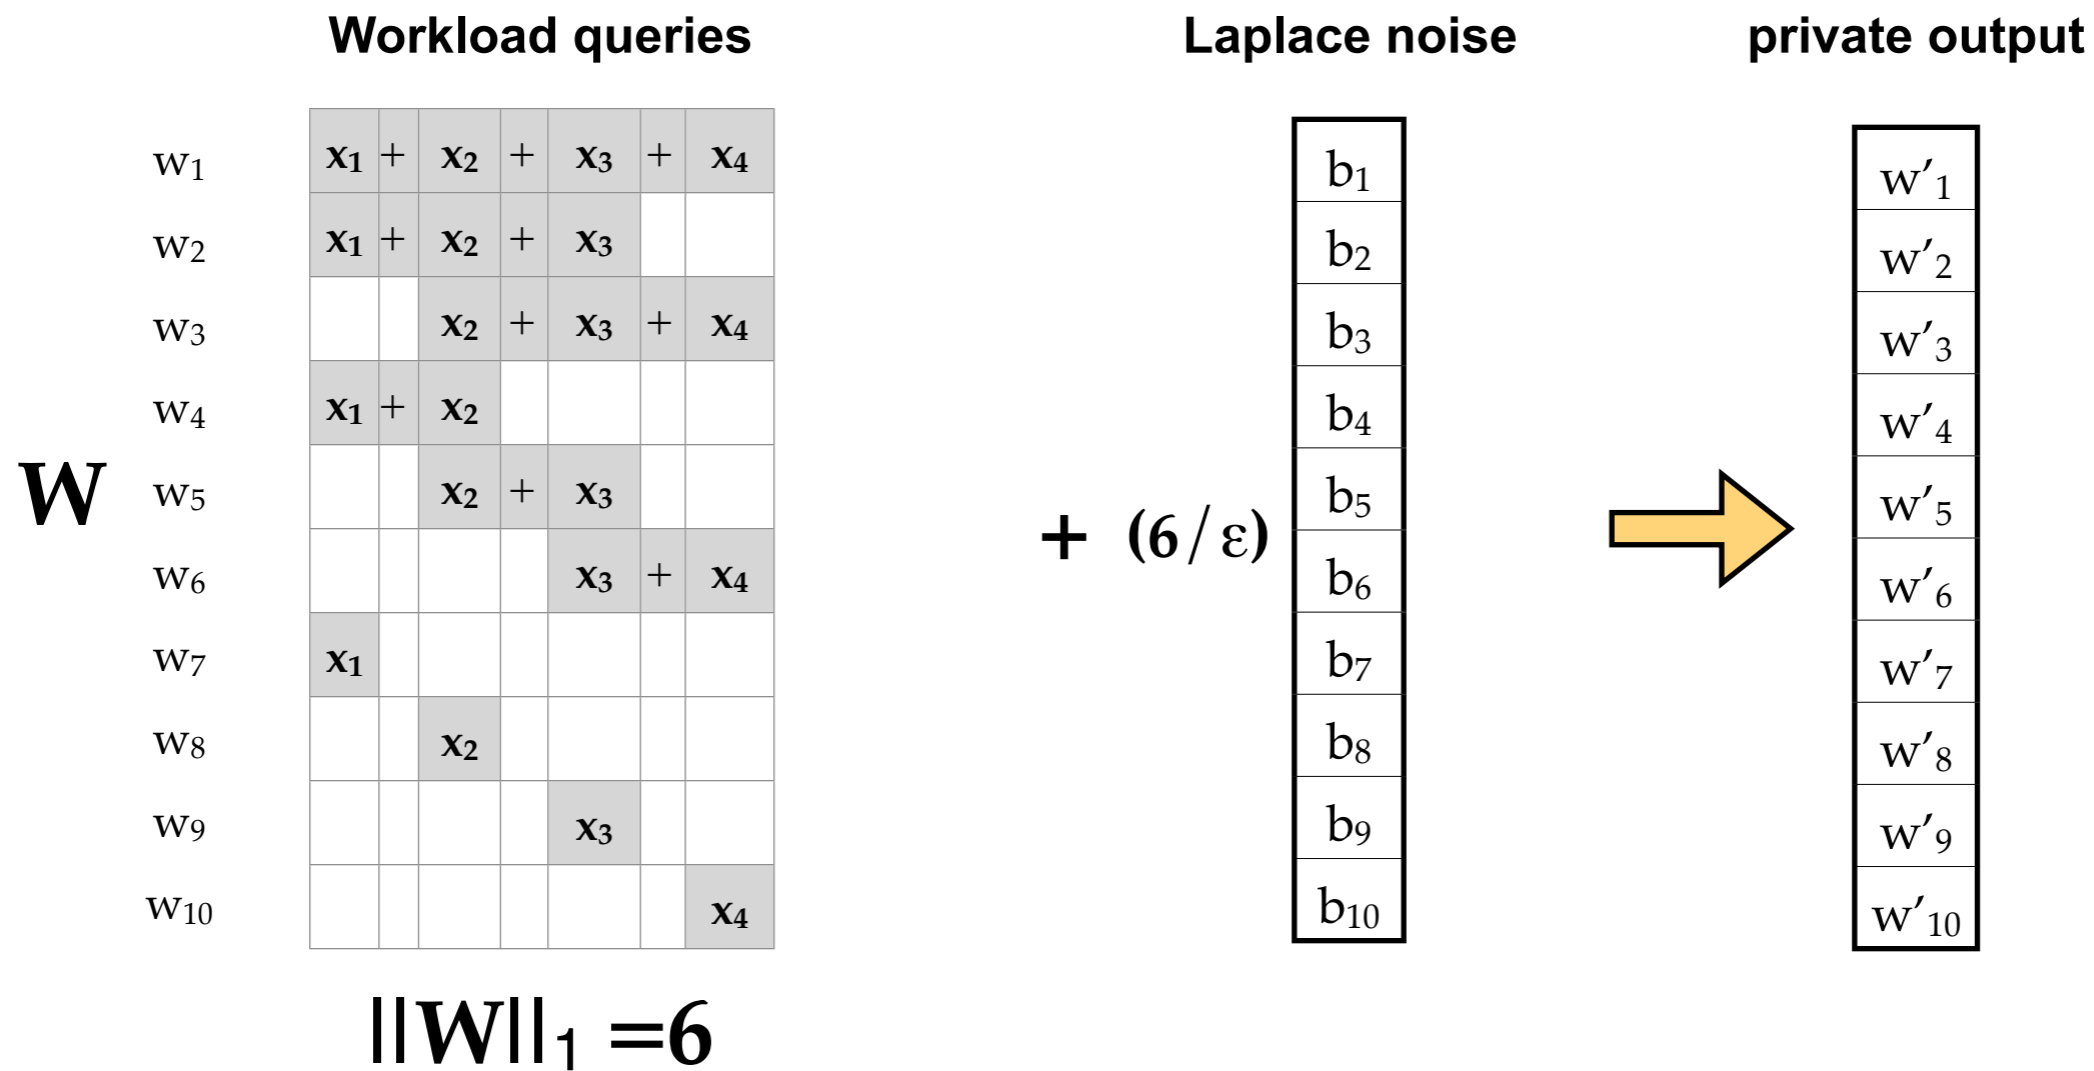
$x =$

10	23	16	3
----	----	----	---

Approach 1: basic Laplace mechanism

Two problems:

- high error
- inconsistency



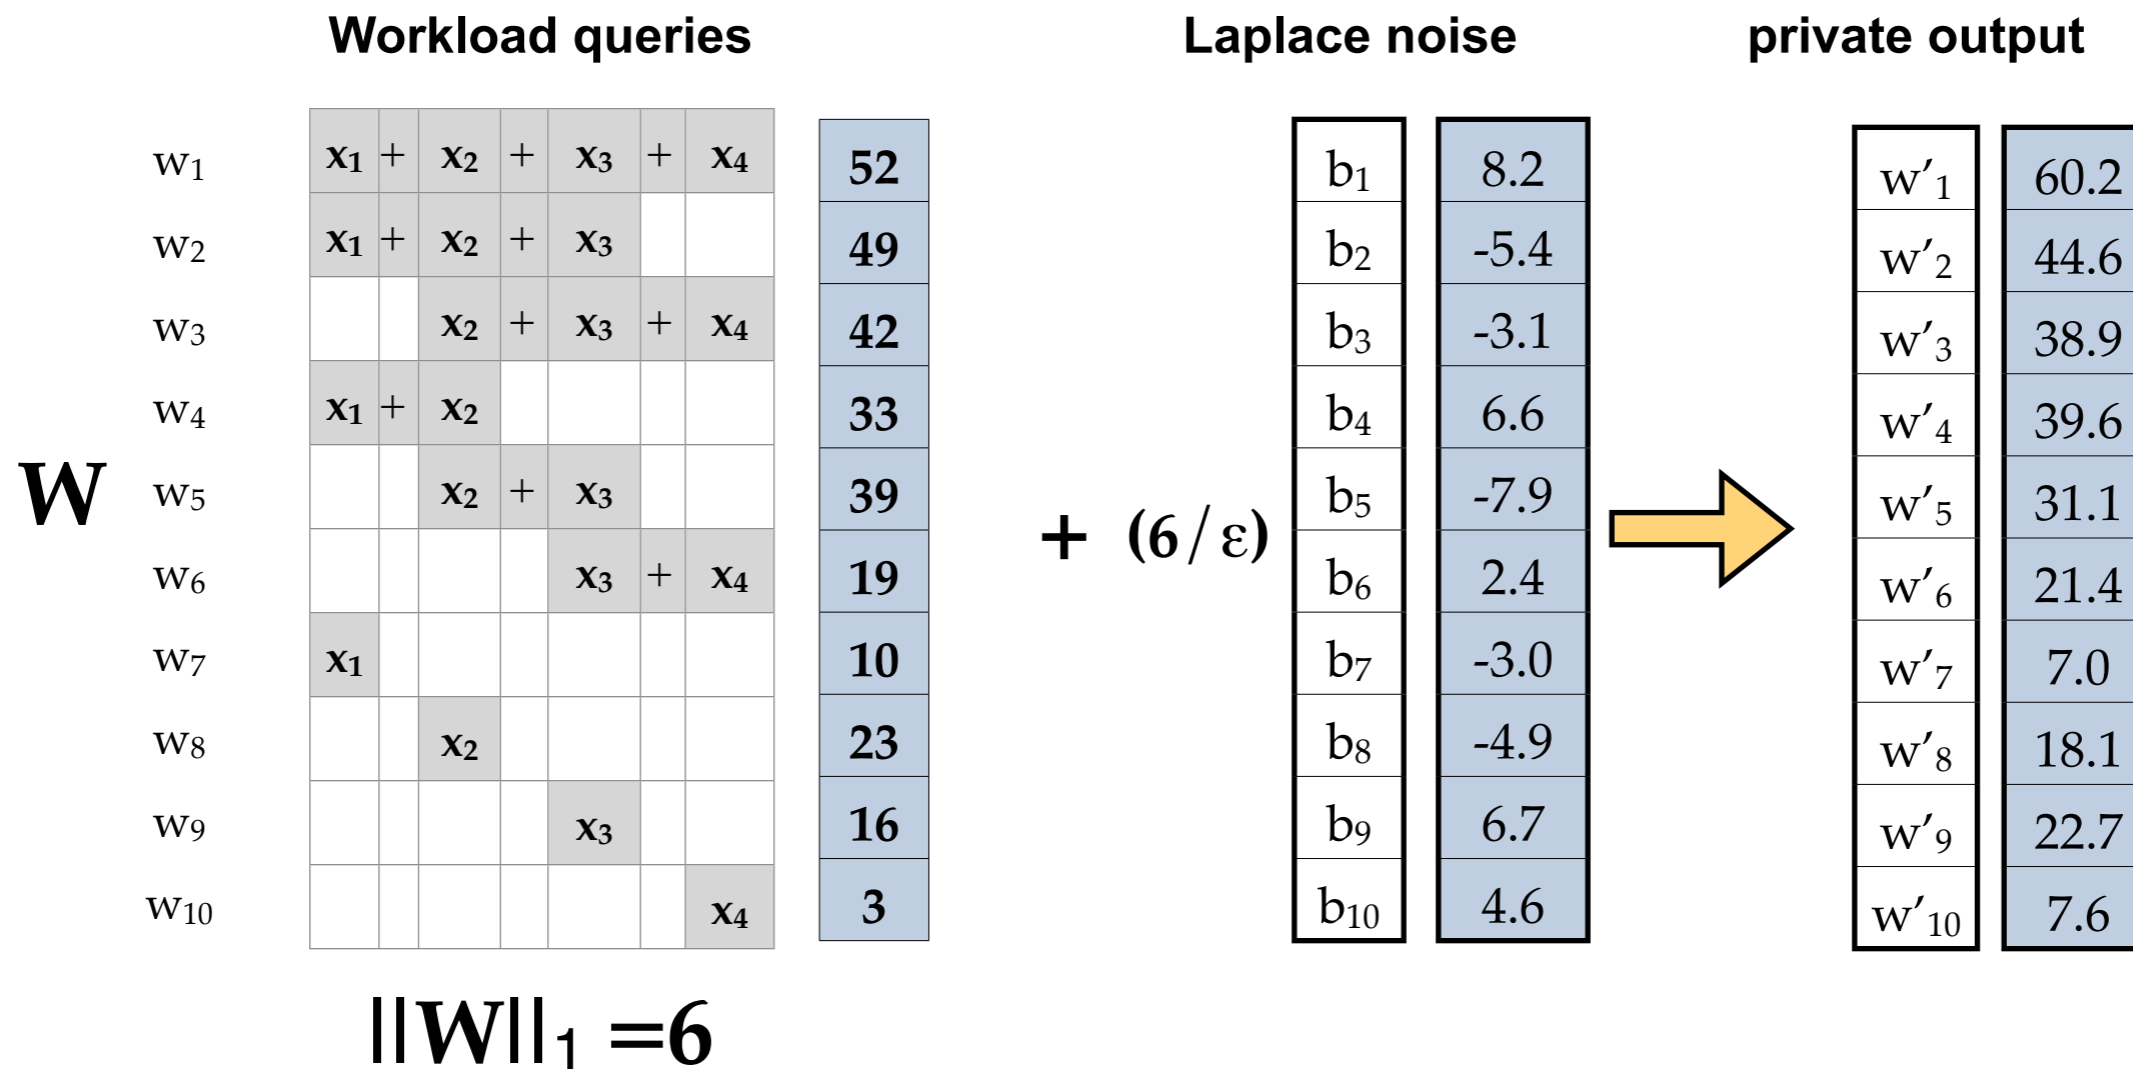
Explain sensi

Error is measured as variance

Approach 1: basic Laplace mechanism

Two problems:

- high error
- inconsistency



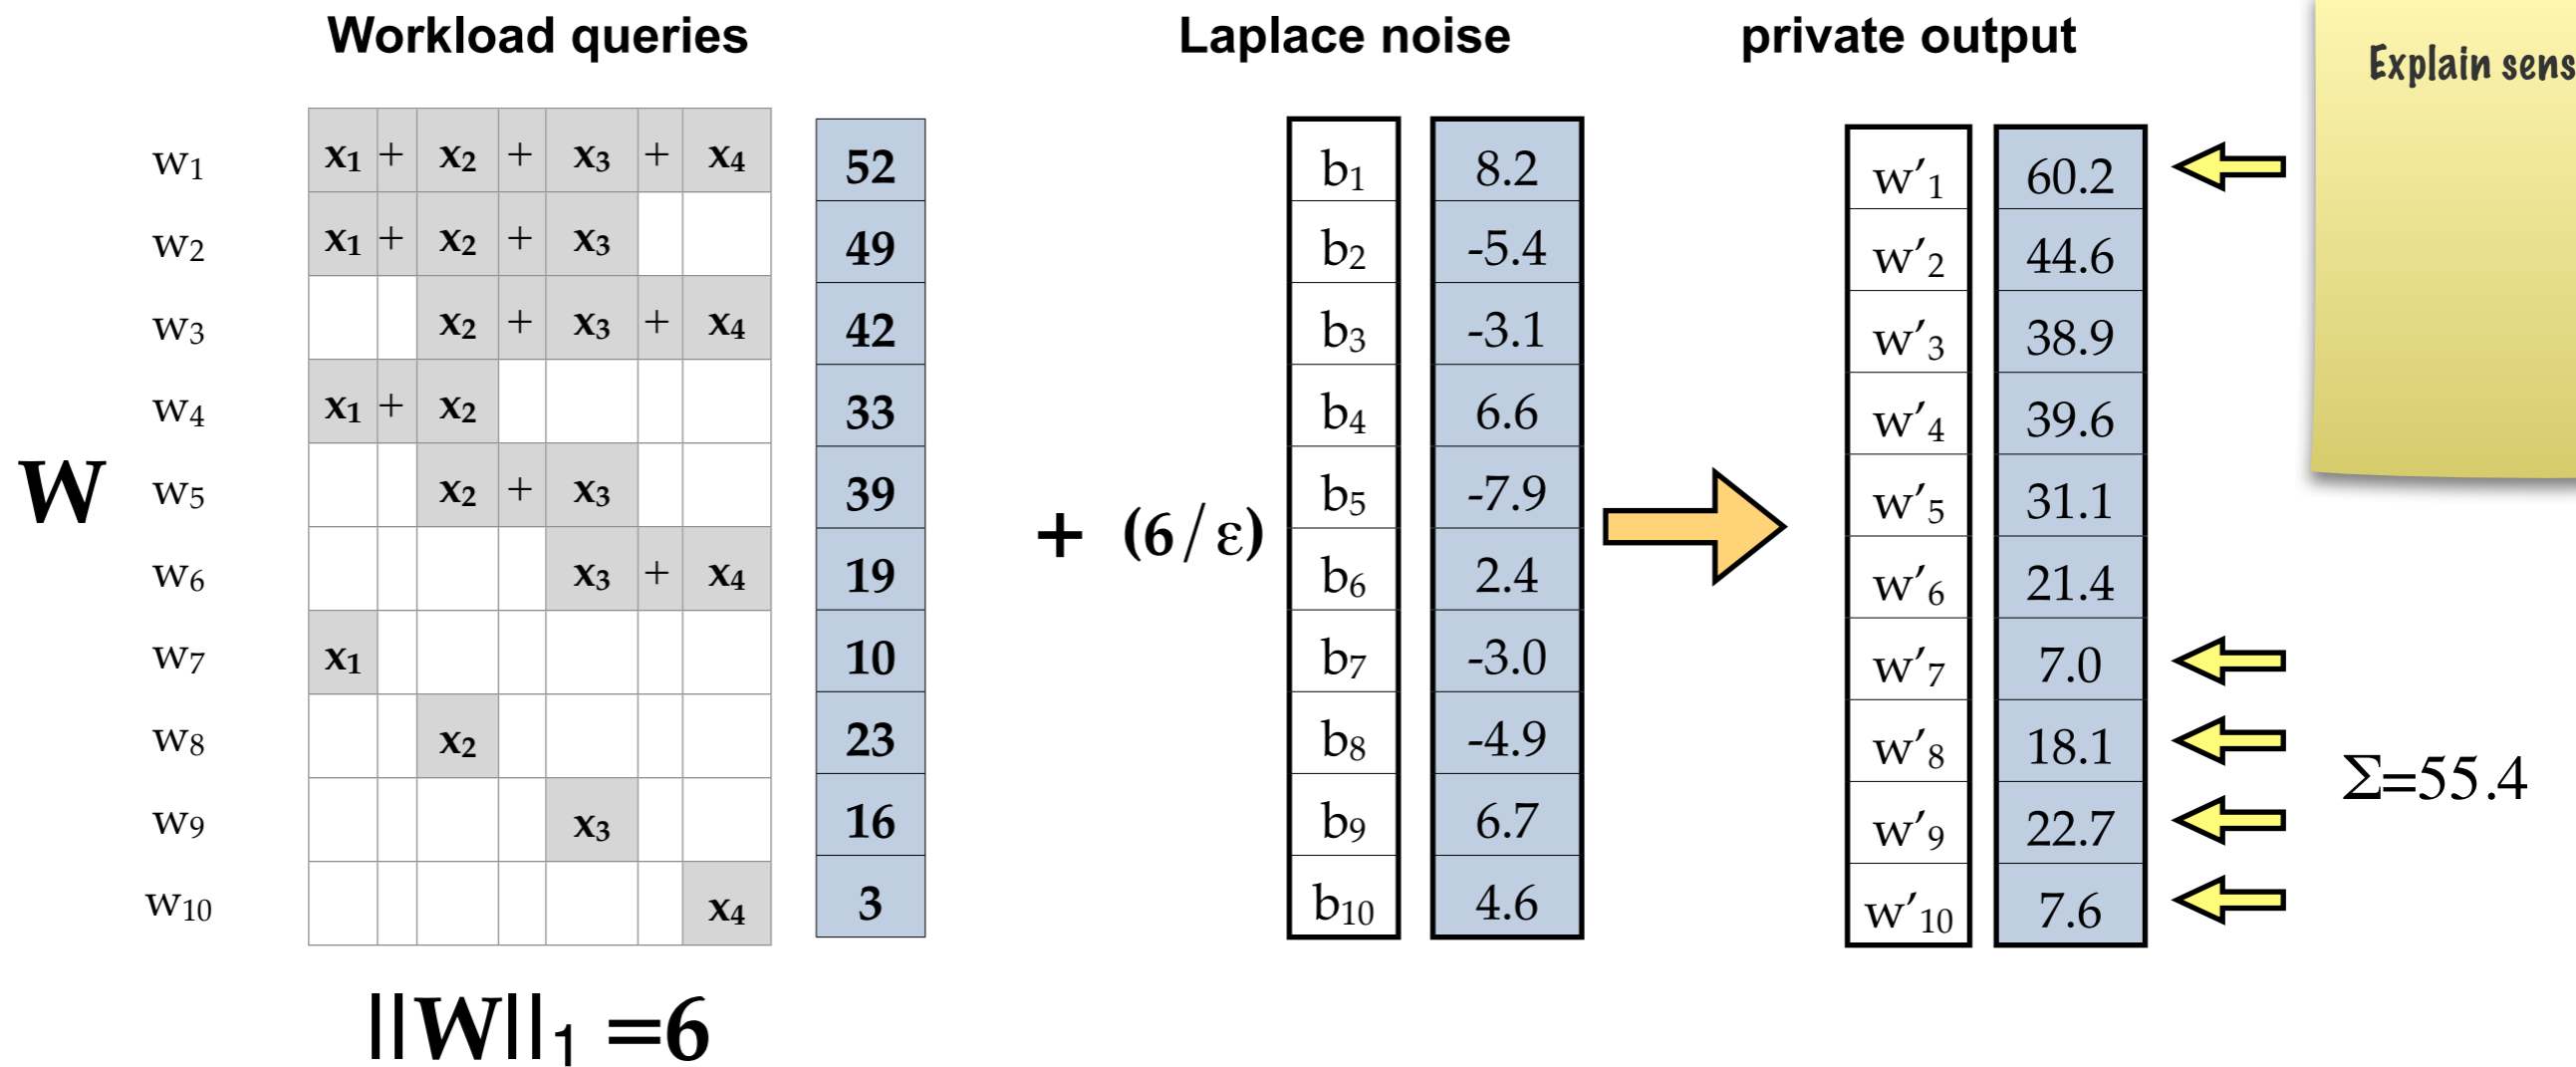
Explain sens

Error is measured as variance

Approach 1: basic Laplace mechanism

Two problems:

- high error
- inconsistency

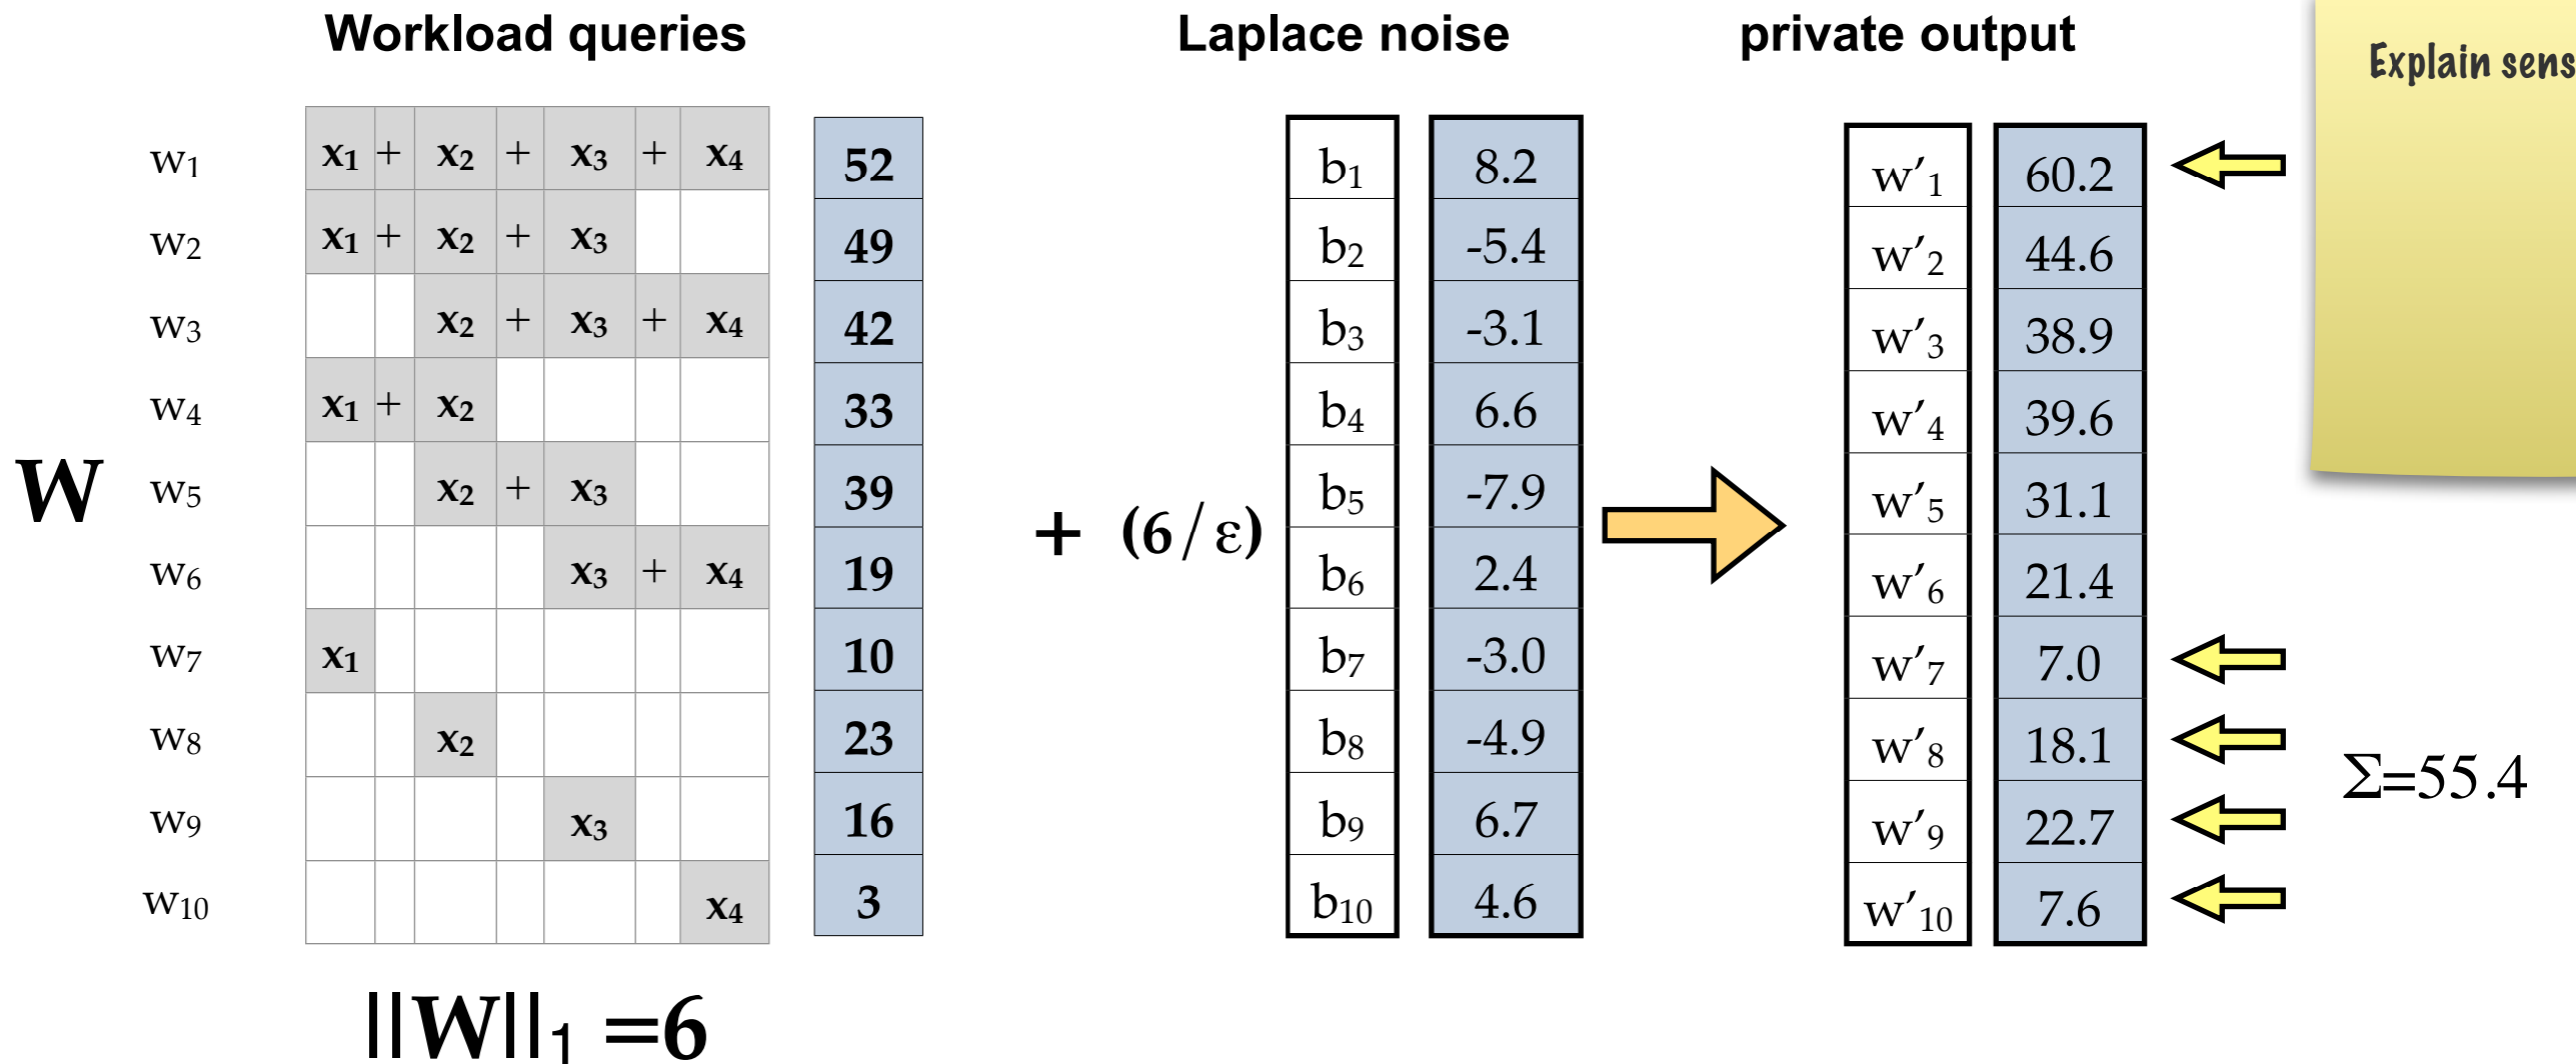


Error is measured as variance

Approach 1: basic Laplace mechanism

Two problems:

- high error
- inconsistency



Explain sensi

	n=4	n
Sensitivity W ₁	6	O(n ²)
Error per query	2(W ₁ /ε) ² = 72/ε ²	2(W ₁ /ε) ² = O(n ⁴)/ε ²

Error is measured as variance

Approach 2: noisy frequency counts

Use Laplace mechanism to get noisy estimates for each x_i .

queries submitted


x_1			
	x_2		
		x_3	
			x_4

I

$$\|\mathbf{I}\|_1 = 1$$

Laplace noise

$$+ (1/\epsilon) \begin{bmatrix} b_1 \\ b_2 \\ b_3 \\ b_4 \end{bmatrix}$$

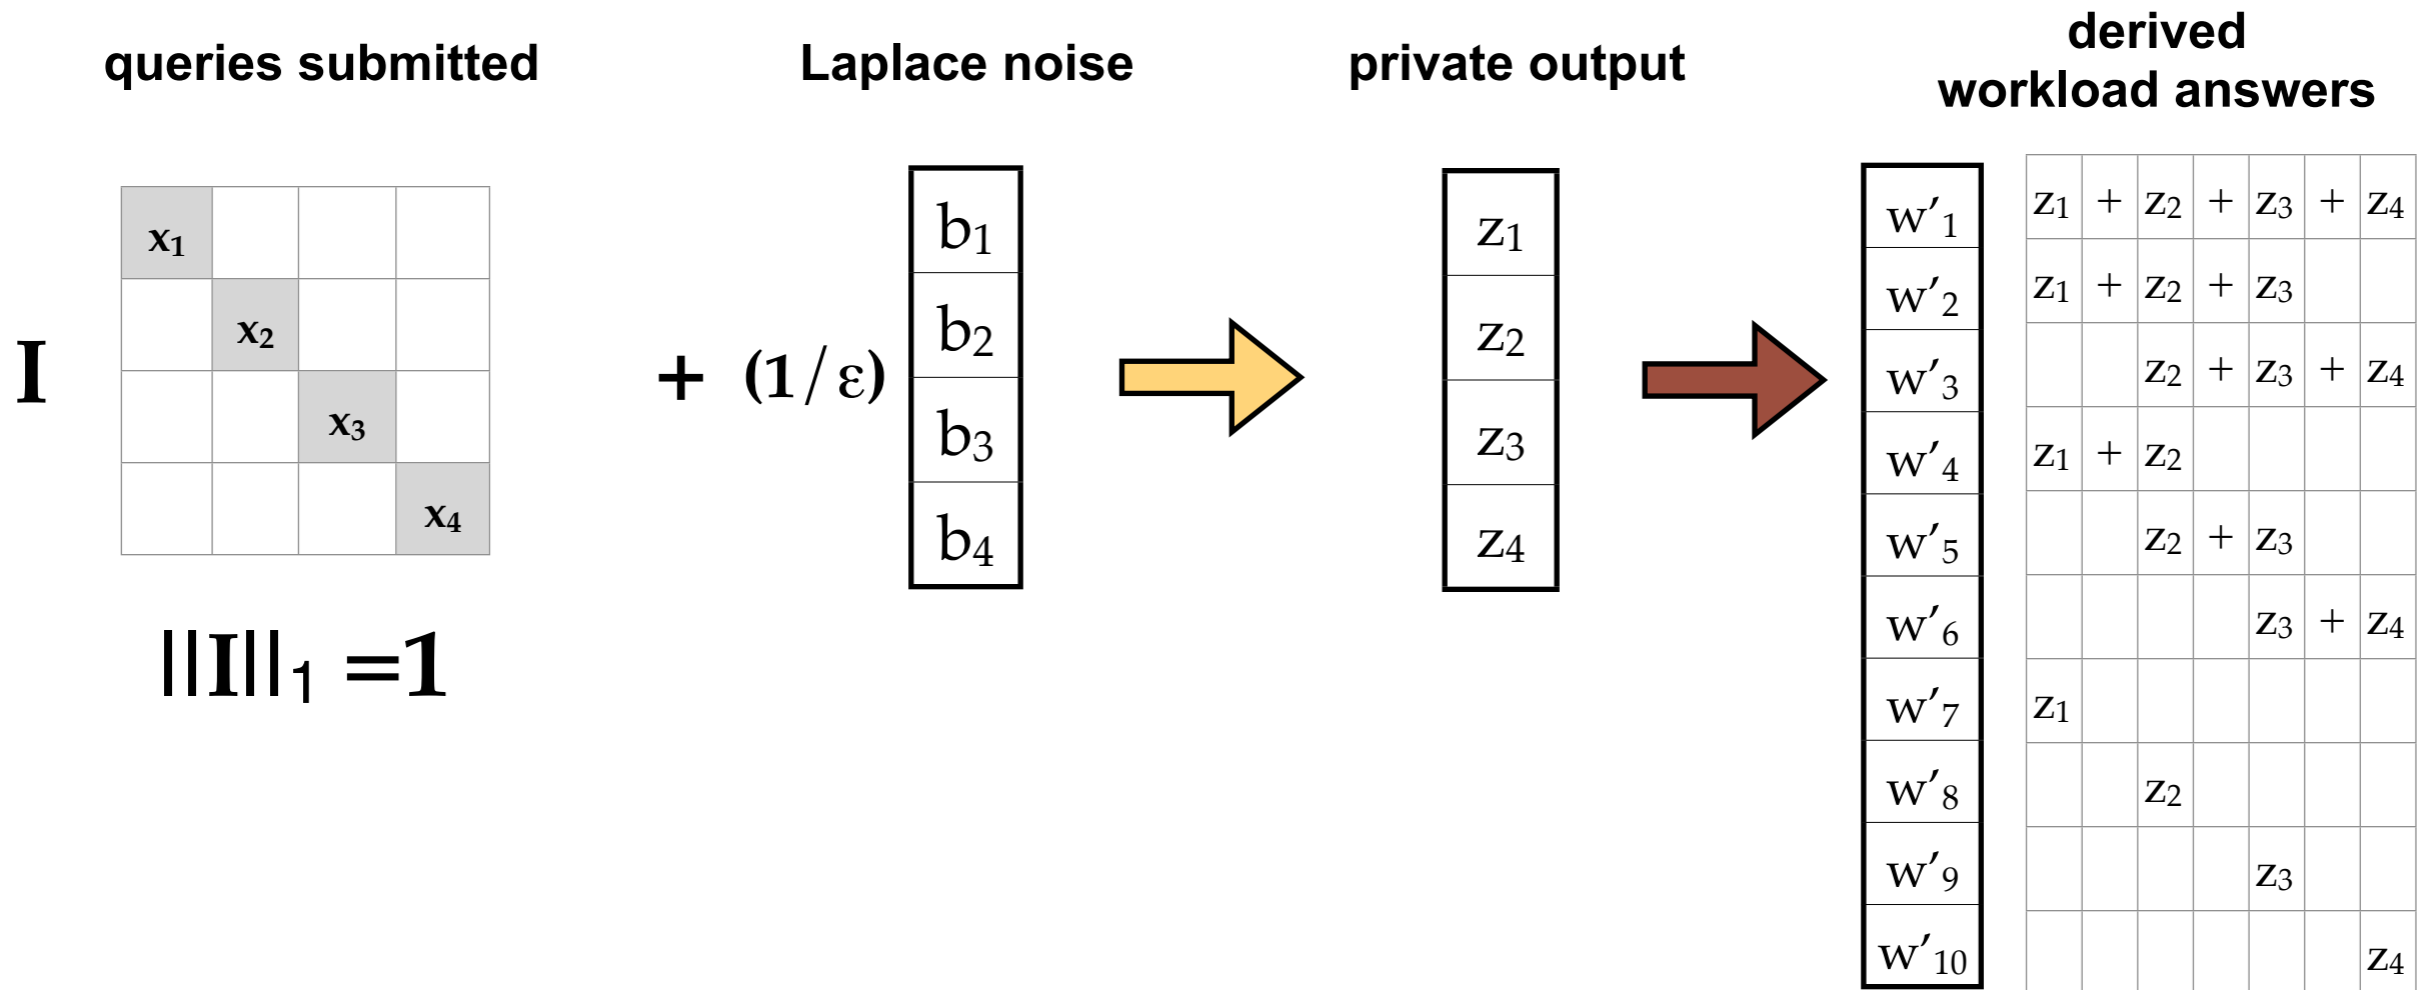


private output

$$\begin{bmatrix} z_1 \\ z_2 \\ z_3 \\ z_4 \end{bmatrix}$$

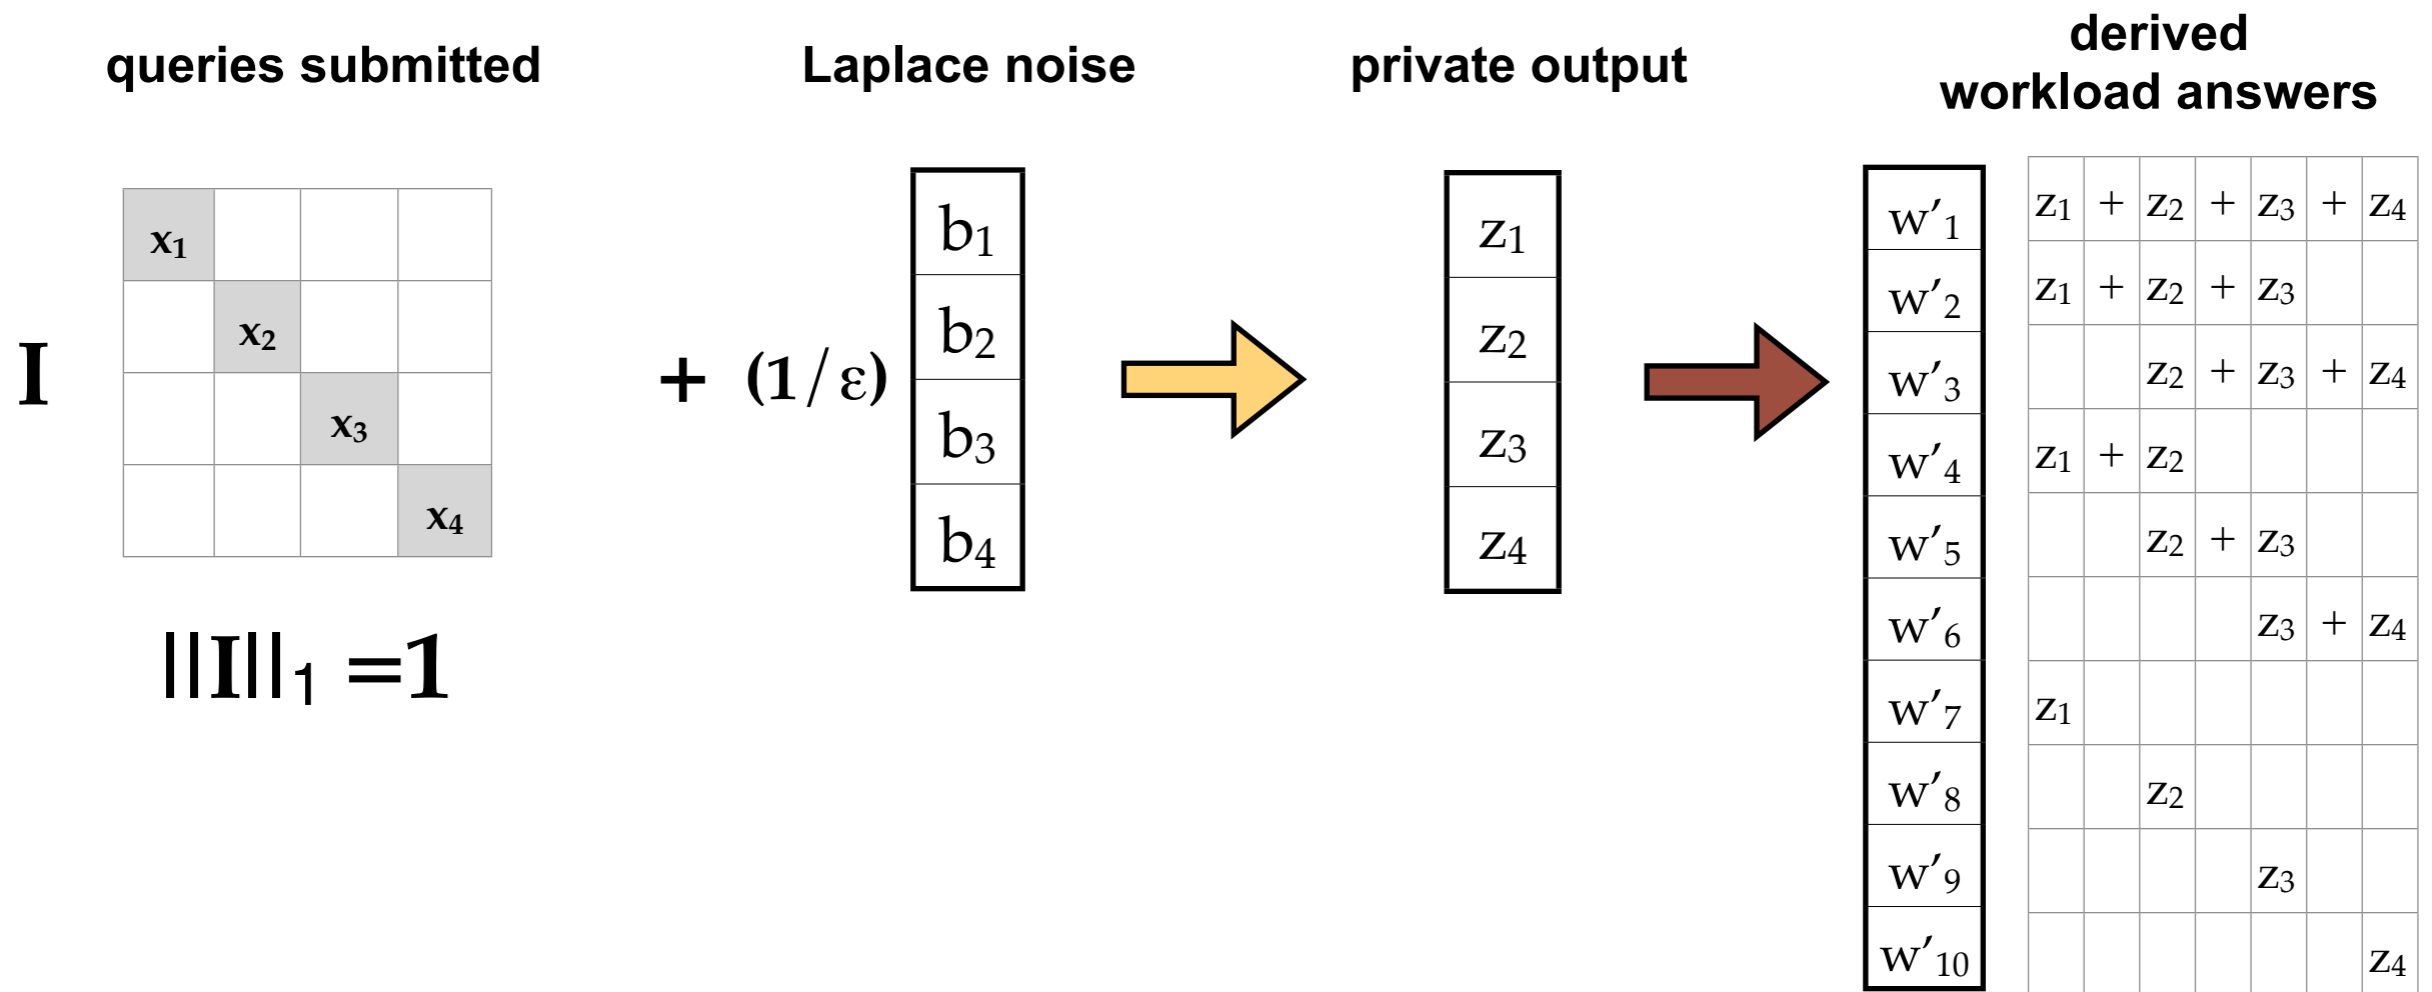
Approach 2: noisy frequency counts

Use Laplace mechanism to get noisy estimates for each x_i .



Approach 2: noisy frequency counts

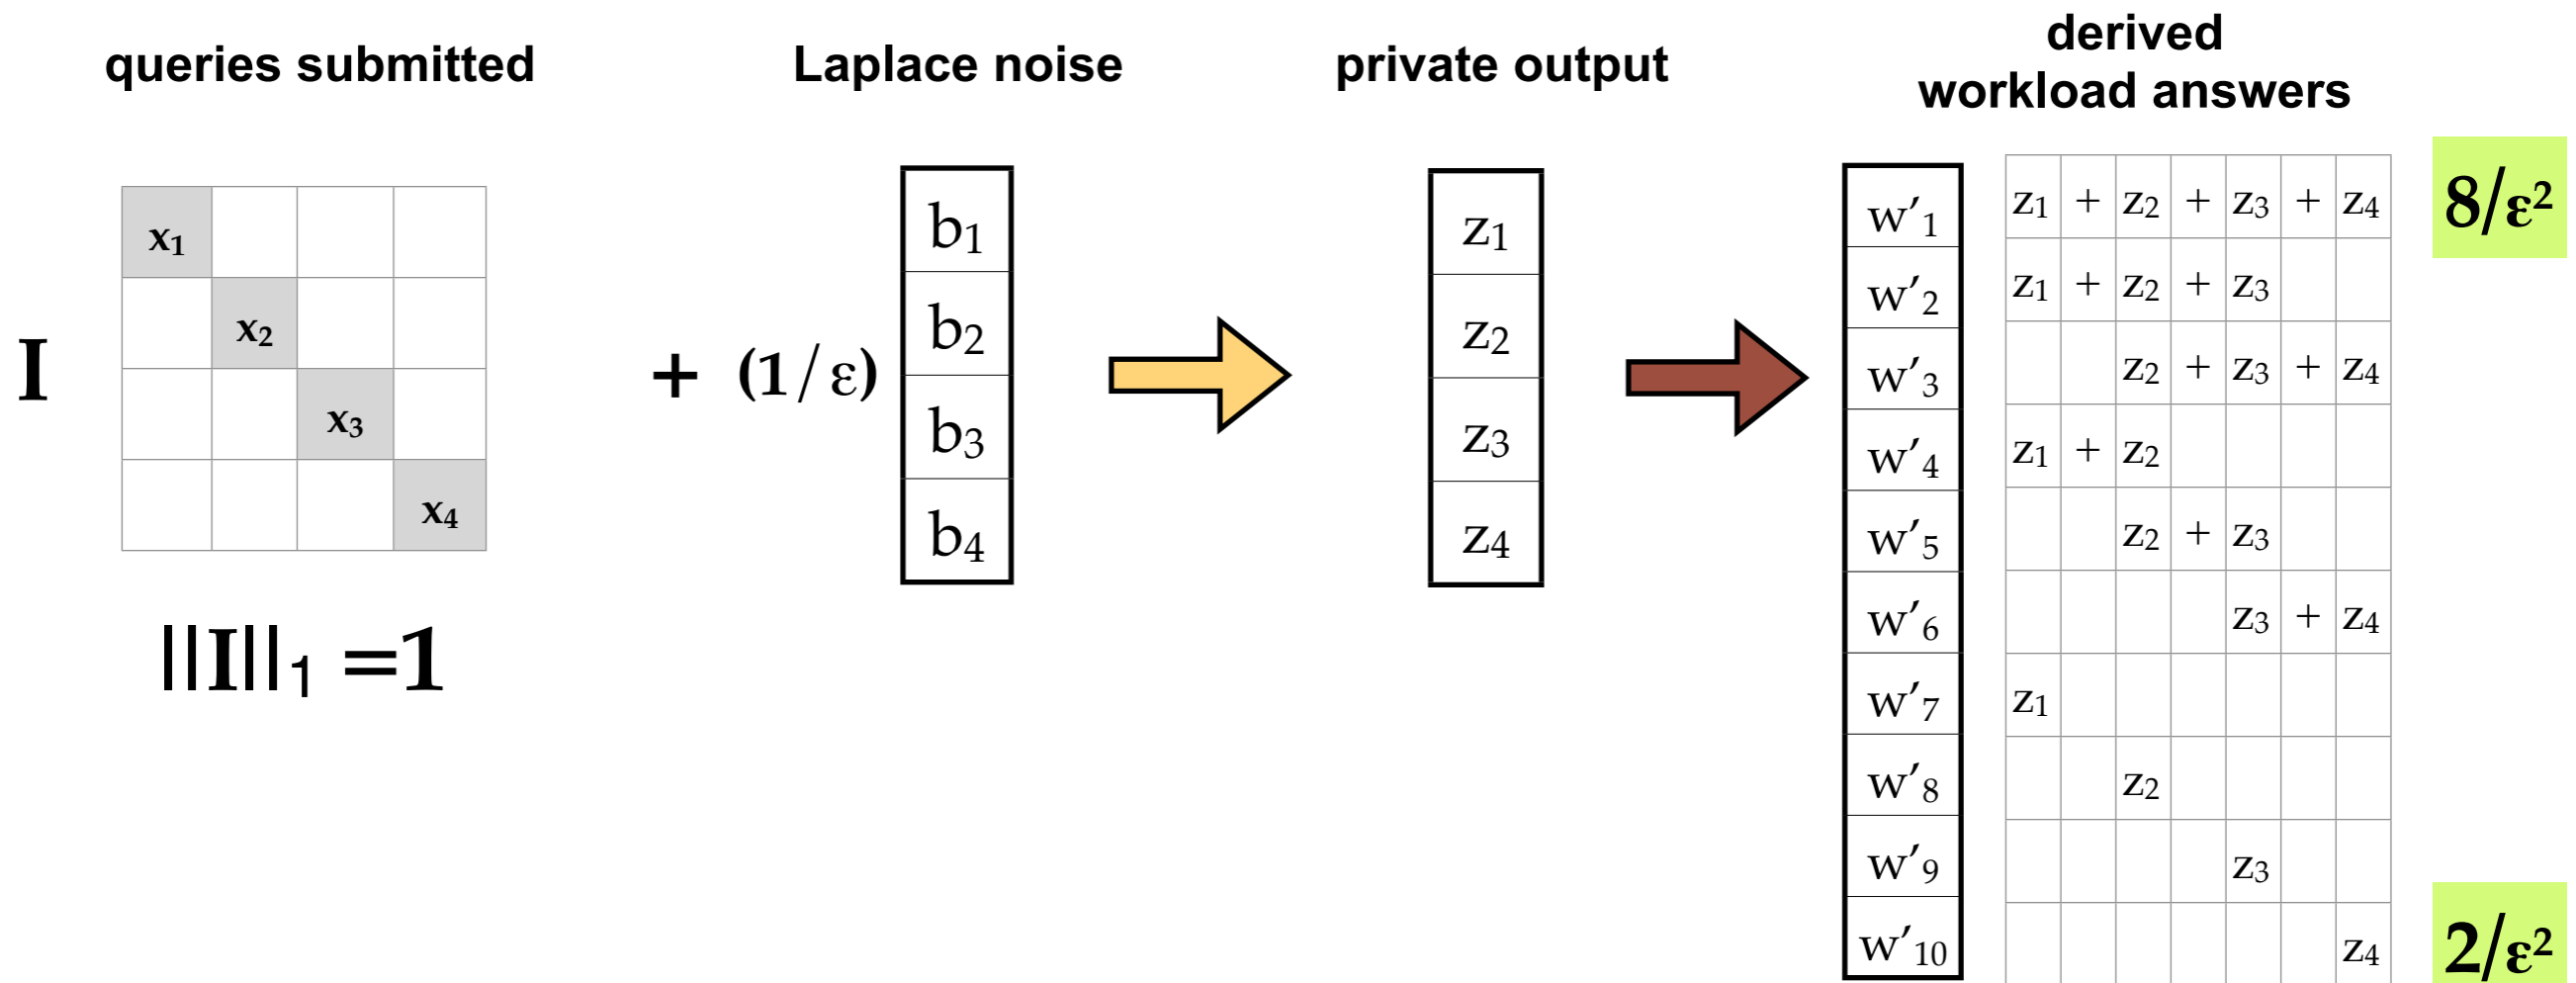
Use Laplace mechanism to get noisy estimates for each x_i .



For $w = \text{range}(x_i, x_j)$ $\text{Error}(w) = 2(j-i+1) / \epsilon^2$

Approach 2: noisy frequency counts

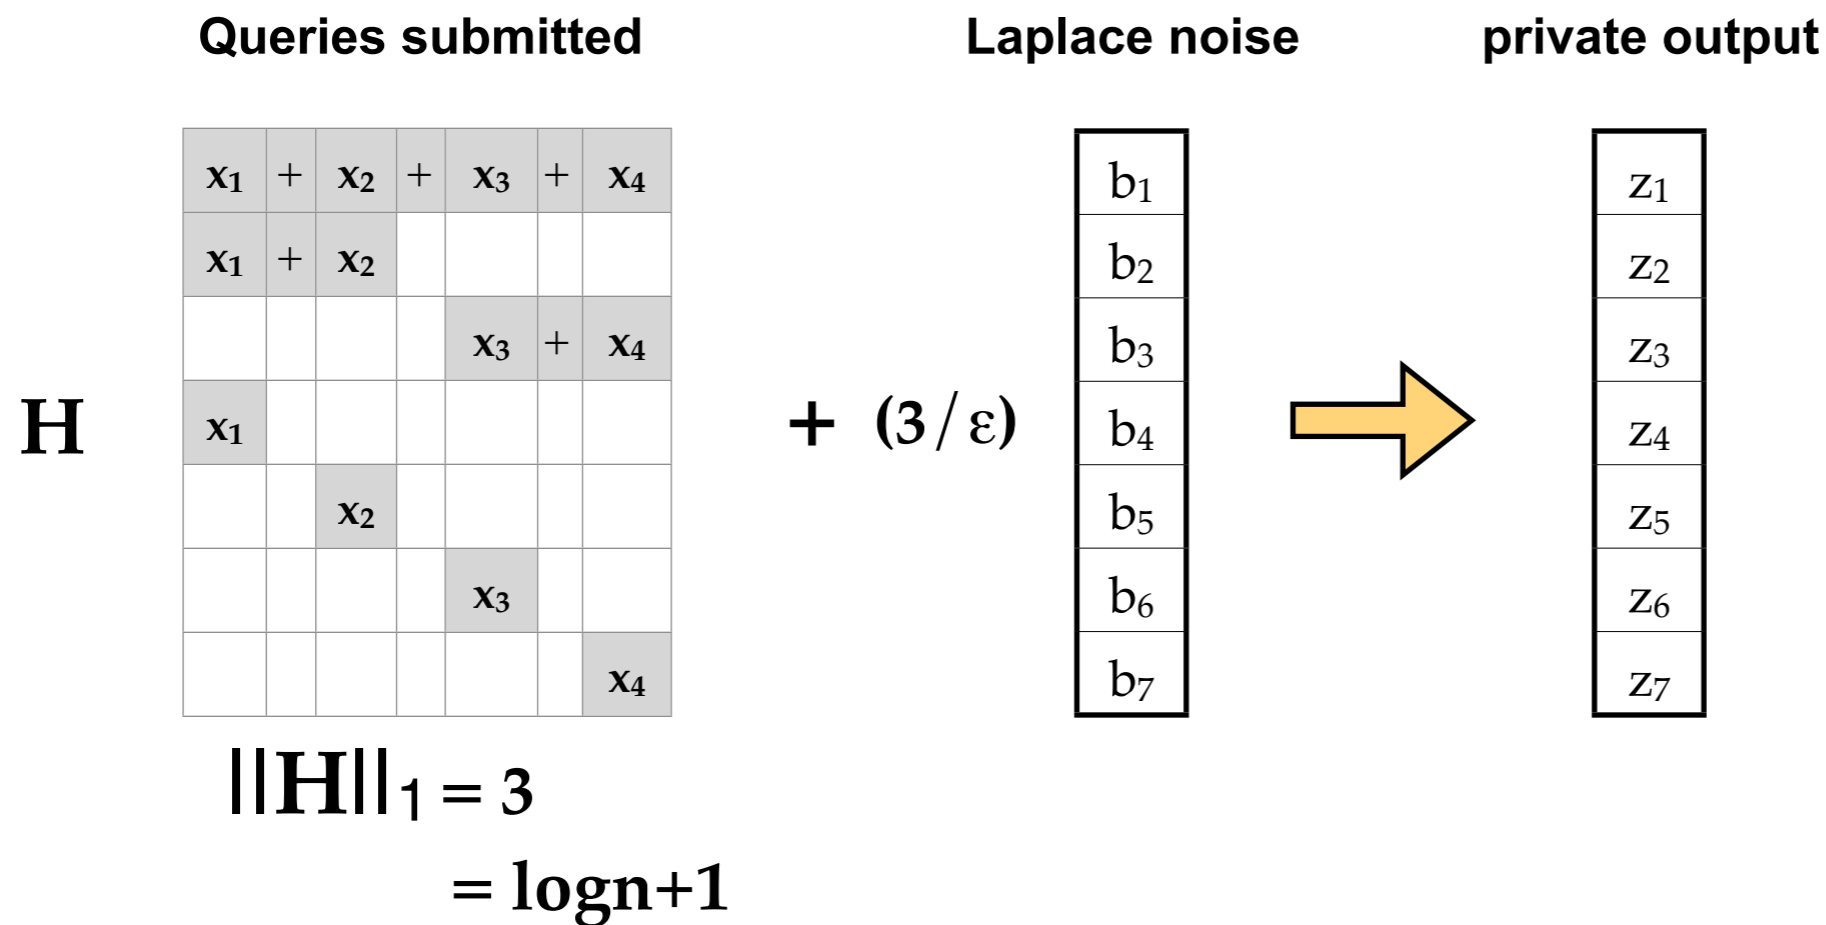
Use Laplace mechanism to get noisy estimates for each x_i .



For $w=\text{range}(x_i, x_j)$ $\text{Error}(w) = 2(j-i+1) / \epsilon^2$

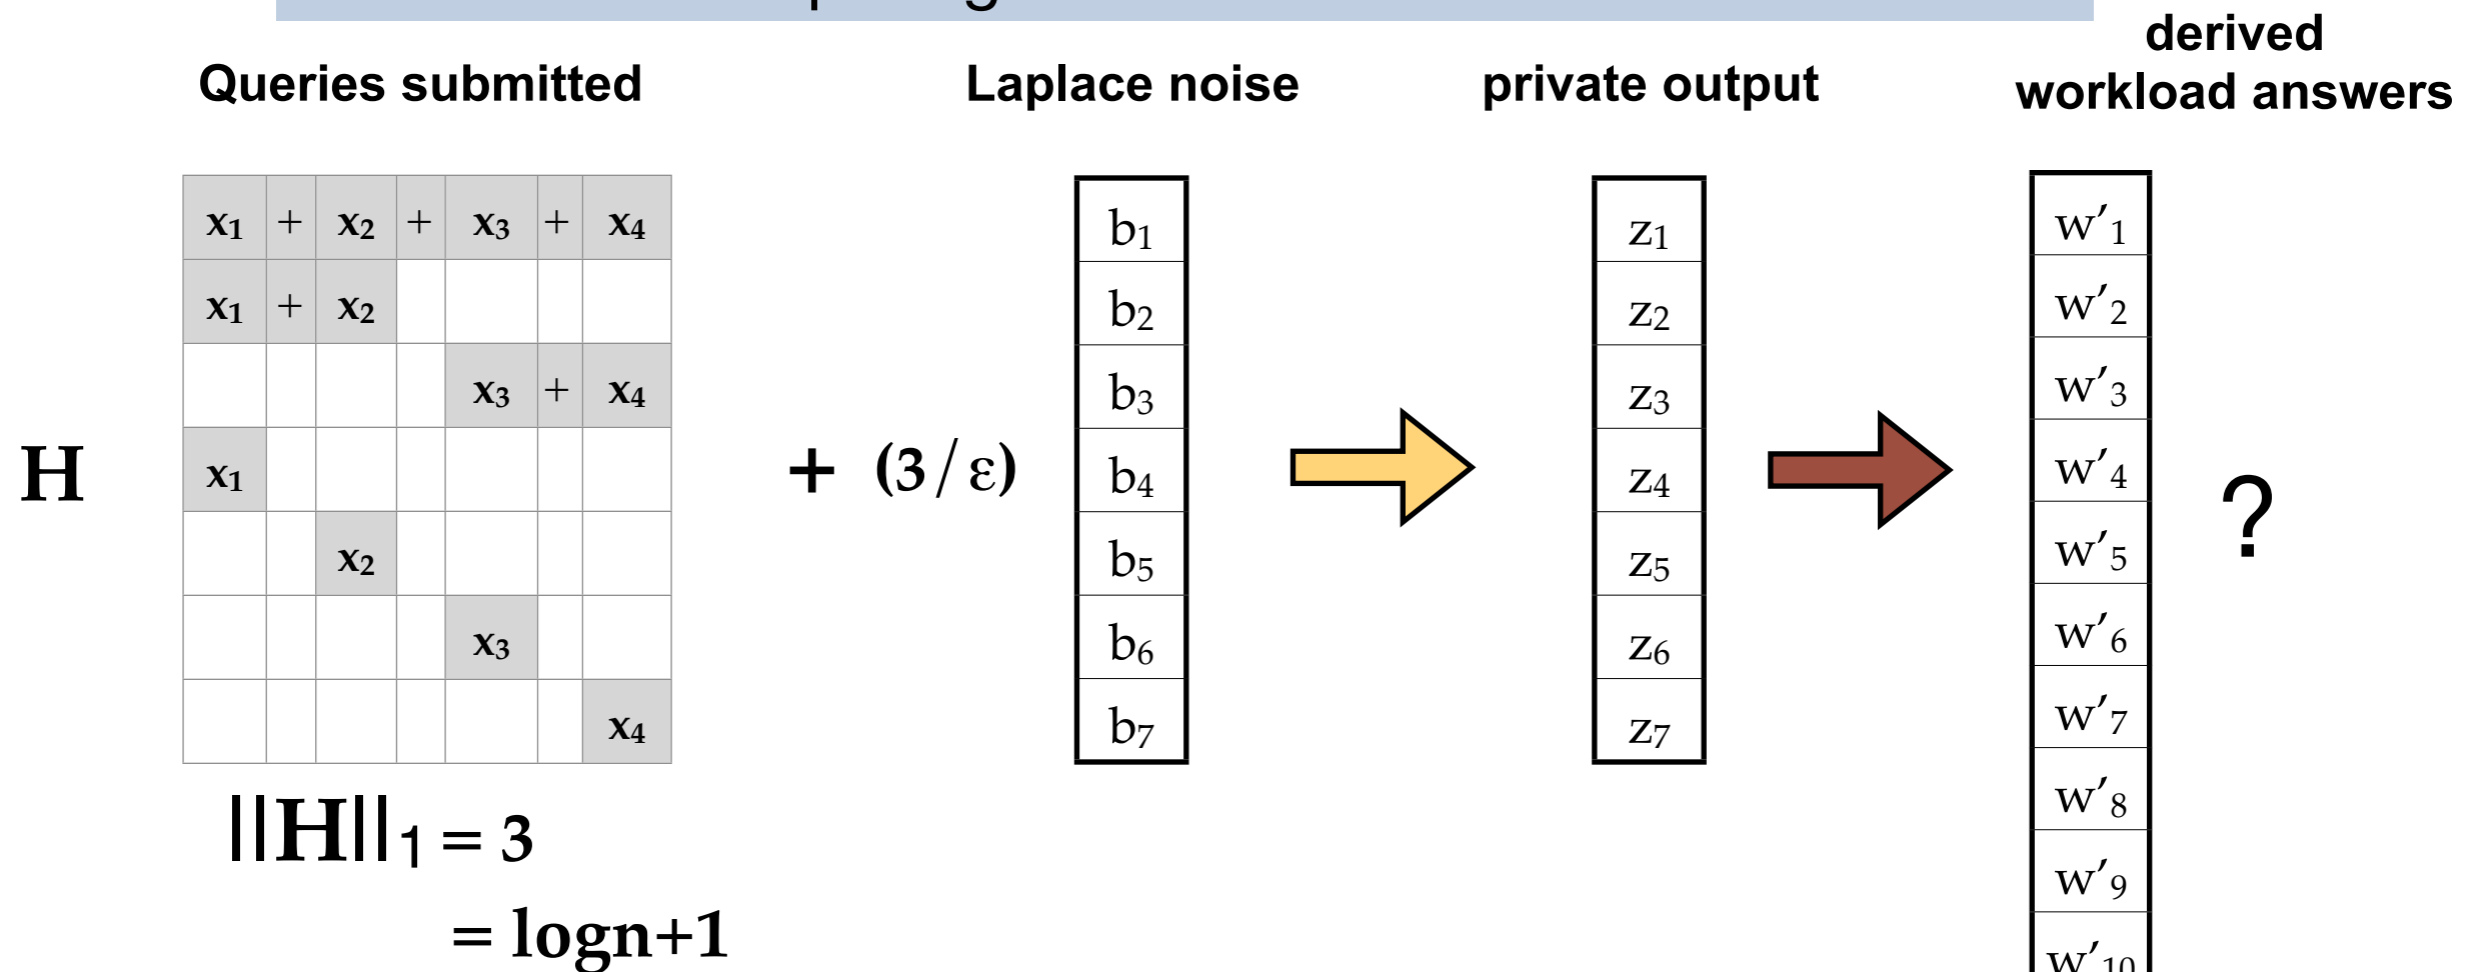
Approach 3: hierarchical queries

Hierarchical queries: recursively partition the domain, computing sums of each interval.




Approach 3: hierarchical queries

Hierarchical queries: recursively partition the domain, computing sums of each interval.



Approach 3: hierarchical queries

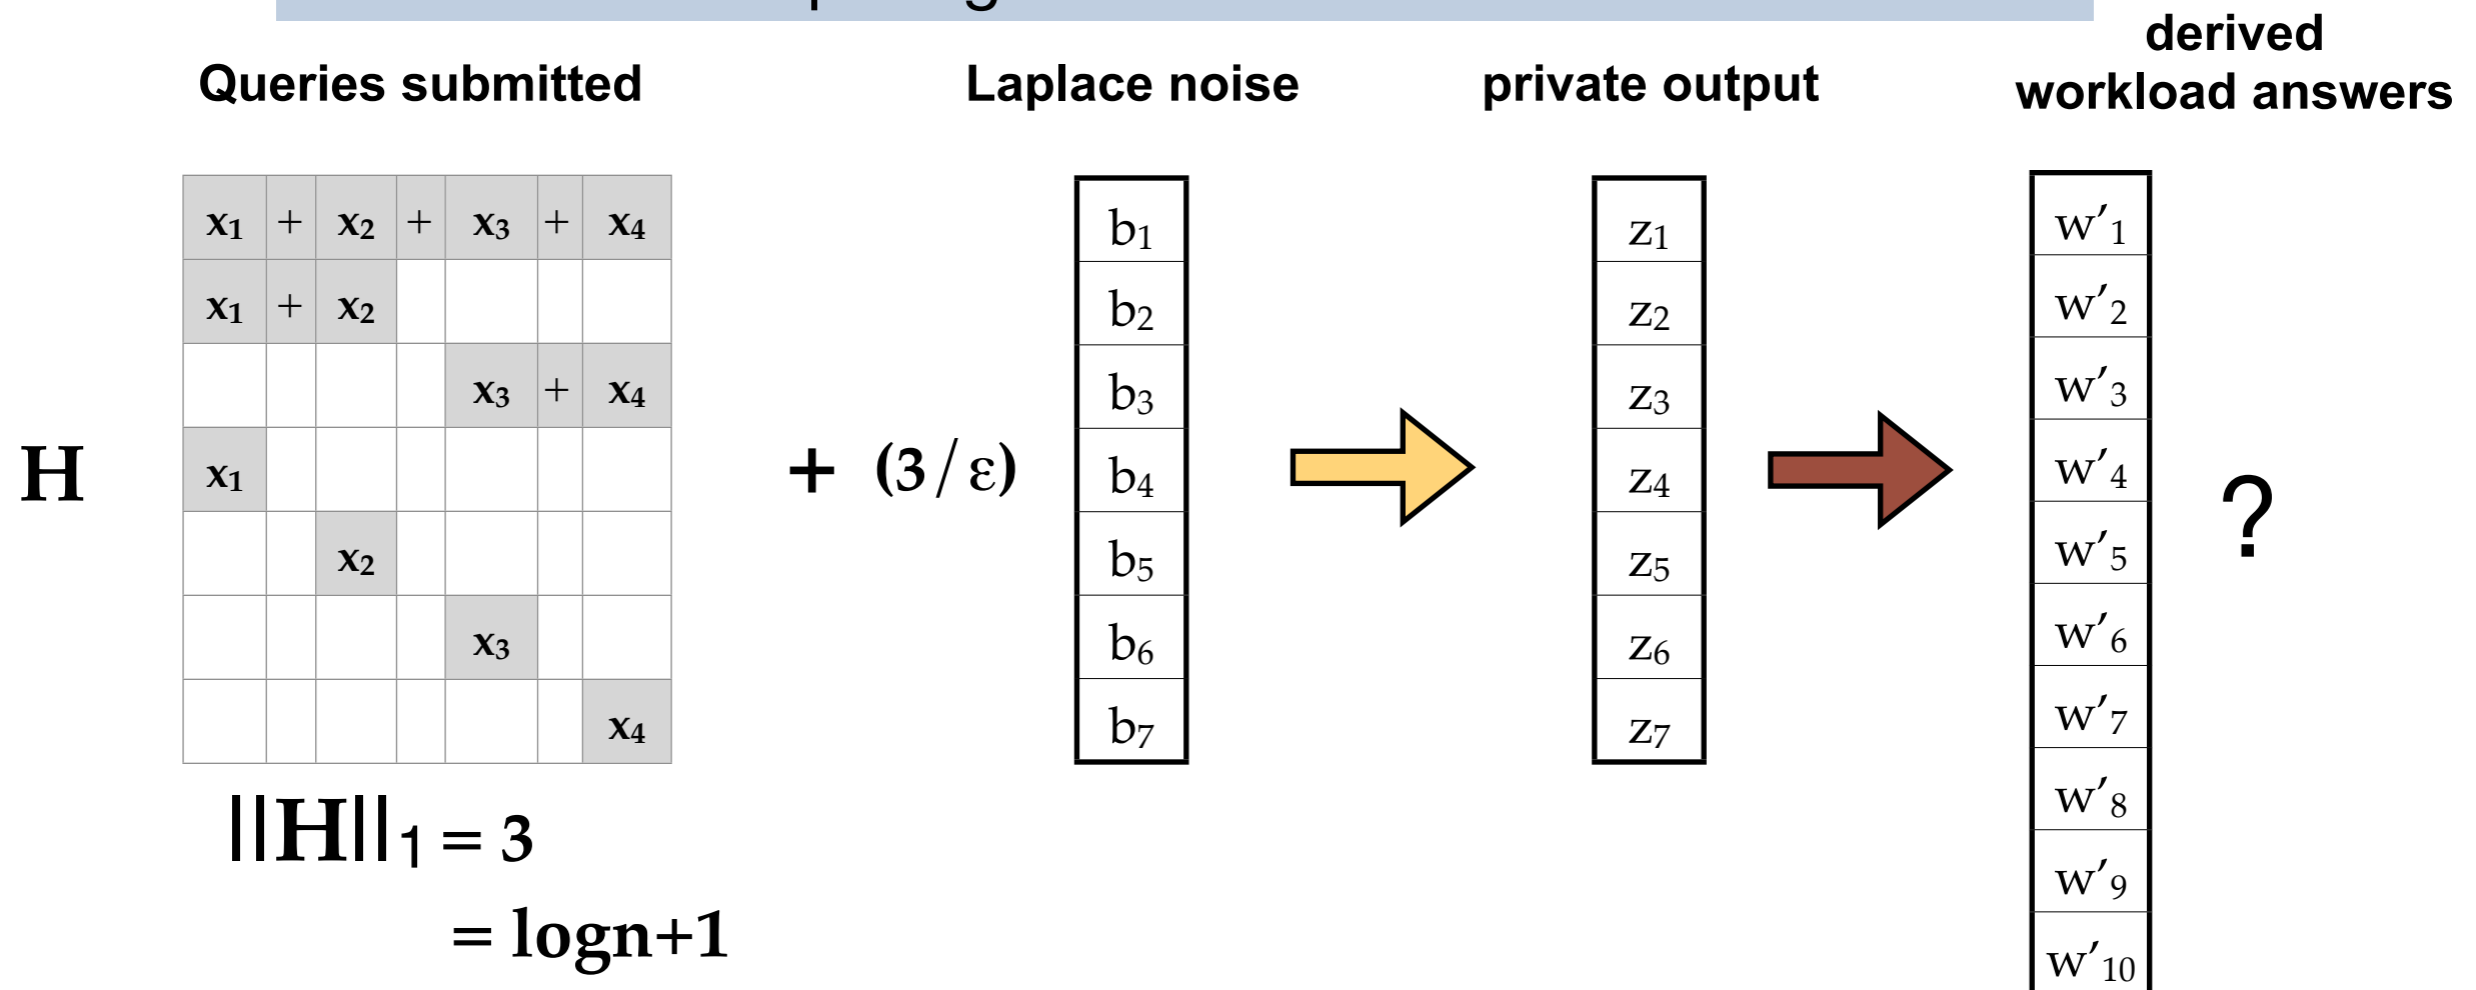
Hierarchical queries: recursively partition the domain, computing sums of each interval.



Possible estimates for query $\text{range}(x_2, x_3) = x_2 + x_3$

Approach 3: hierarchical queries

Hierarchical queries: recursively partition the domain, computing sums of each interval.




Possible estimates for query $\text{range}(x_2, x_3) = x_2 + x_3$

$$z_5 + z_6$$

Approach 3: hierarchical queries

Hierarchical queries: recursively partition the domain, computing sums of each interval.




Possible estimates for query $\text{range}(x_2, x_3) = x_2 + x_3$

$$z_5 + z_6$$

$$z_2 - z_4 + z_6$$

Approach 3: hierarchical queries

Hierarchical queries: recursively partition the domain, computing sums of each interval.



Possible estimates for query $\text{range}(x_2, x_3) = x_2 + x_3$


$$z_5 + z_6$$

$$z_2 - z_4 + z_6$$

$$z_1 - z_4 - z_7$$

Approach 3: hierarchical queries

Hierarchical queries: recursively partition the domain, computing sums of each interval.



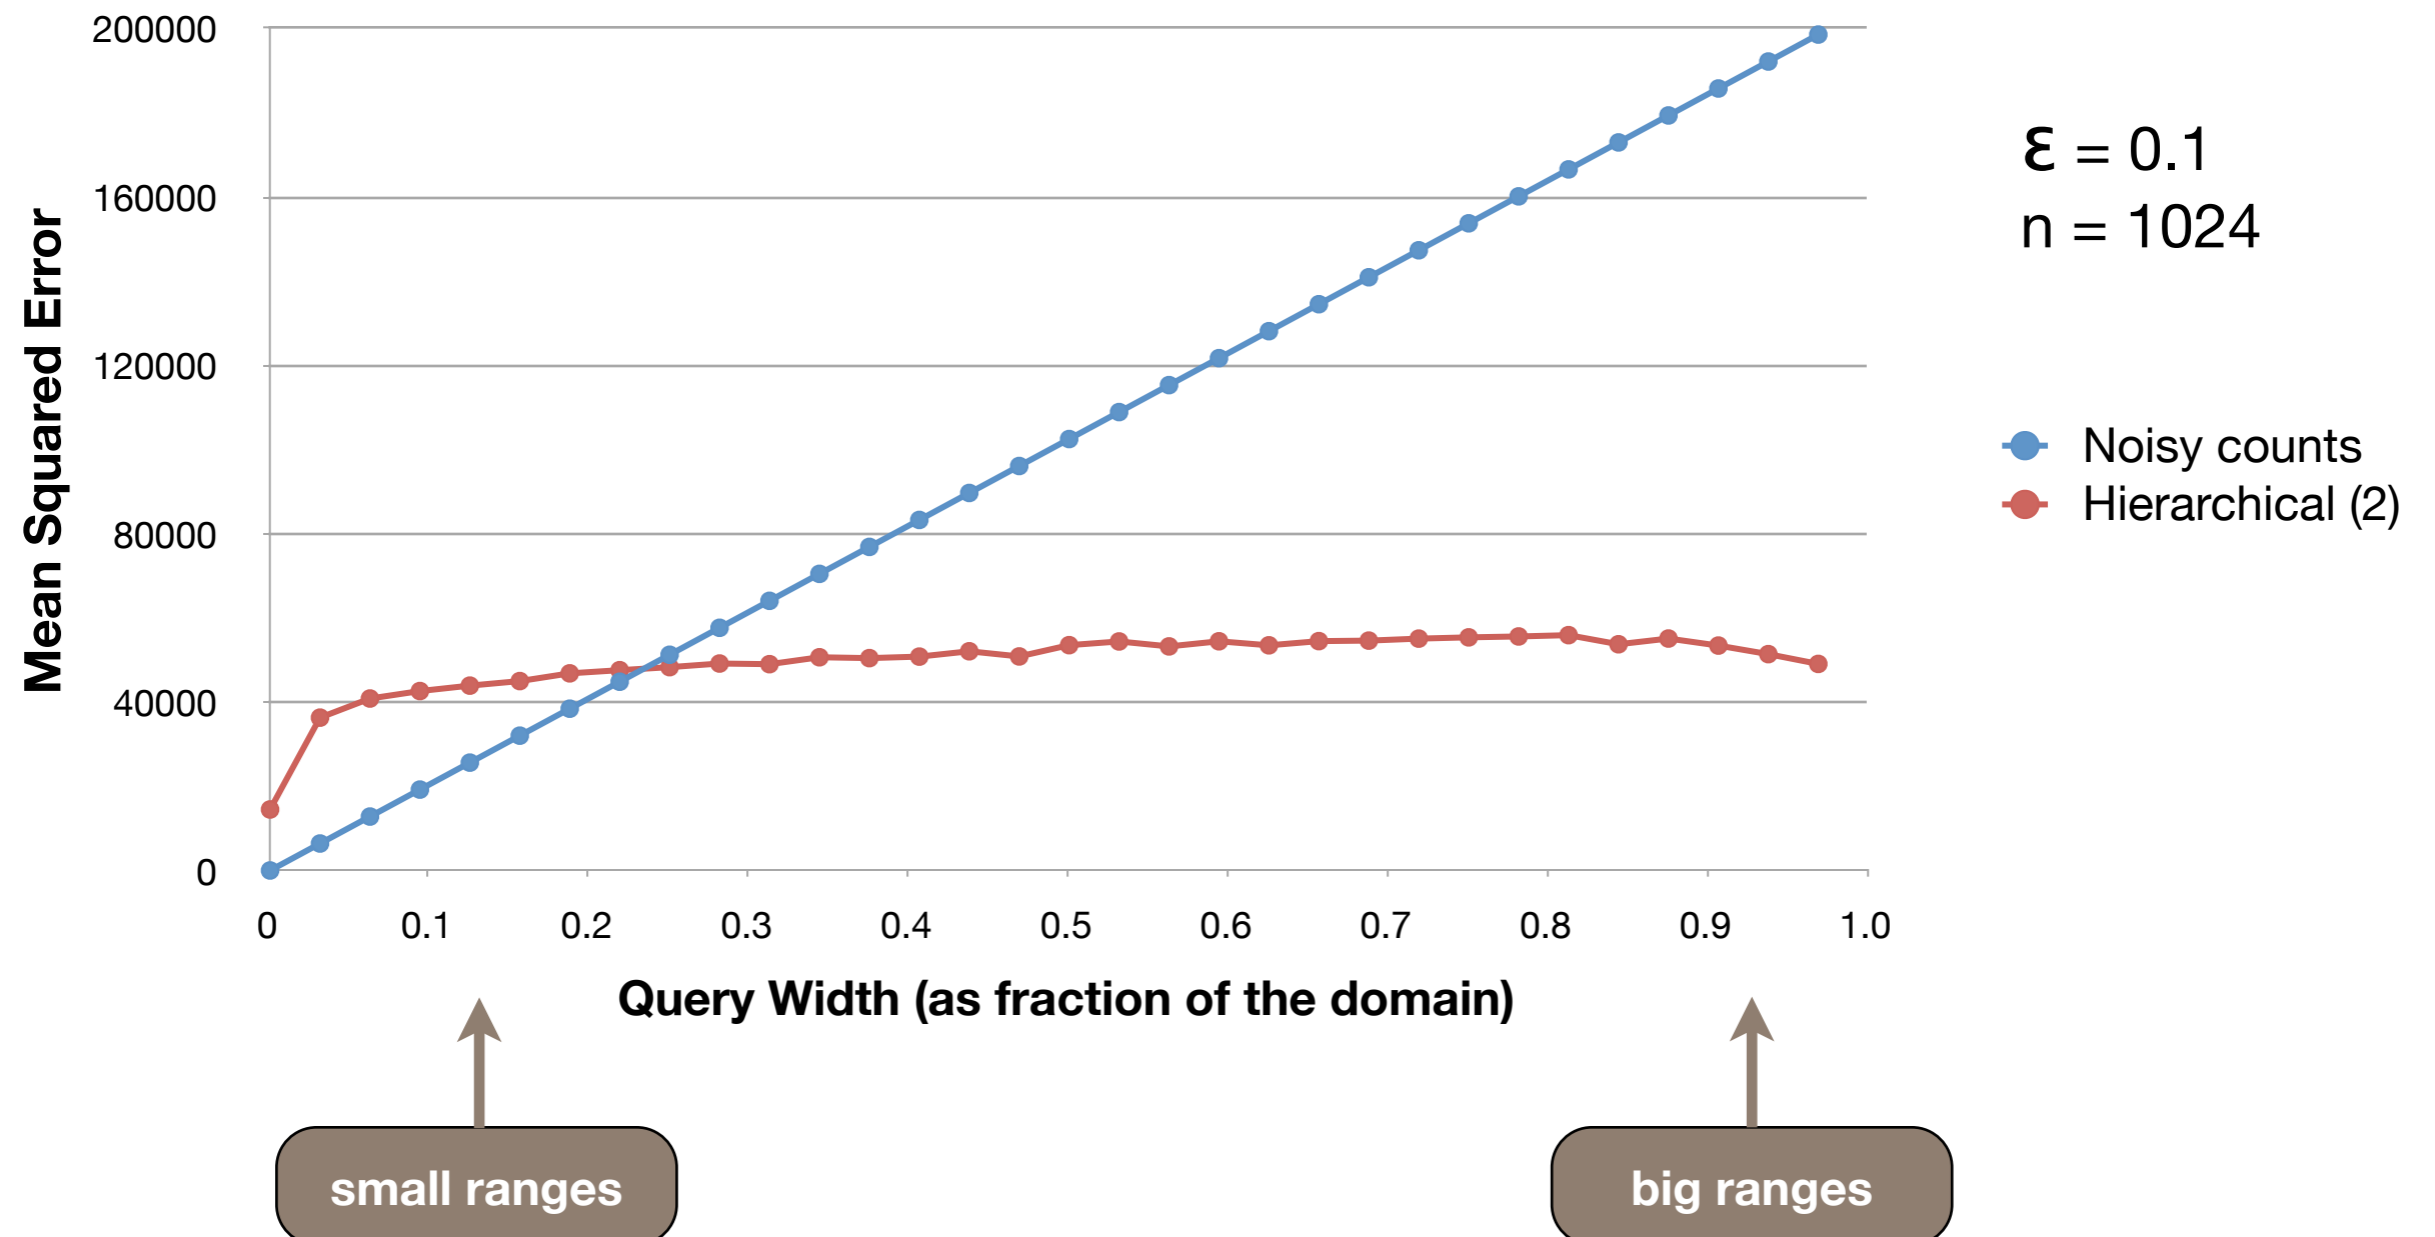
Possible estimates for query $\text{range}(x_2, x_3) = x_2 + x_3$

Least-squares
estimate

$$(6z_1 + 3z_2 + 3z_3 - 9z_4 + 12z_5 + 12z_6 - 9z_7) / 21$$

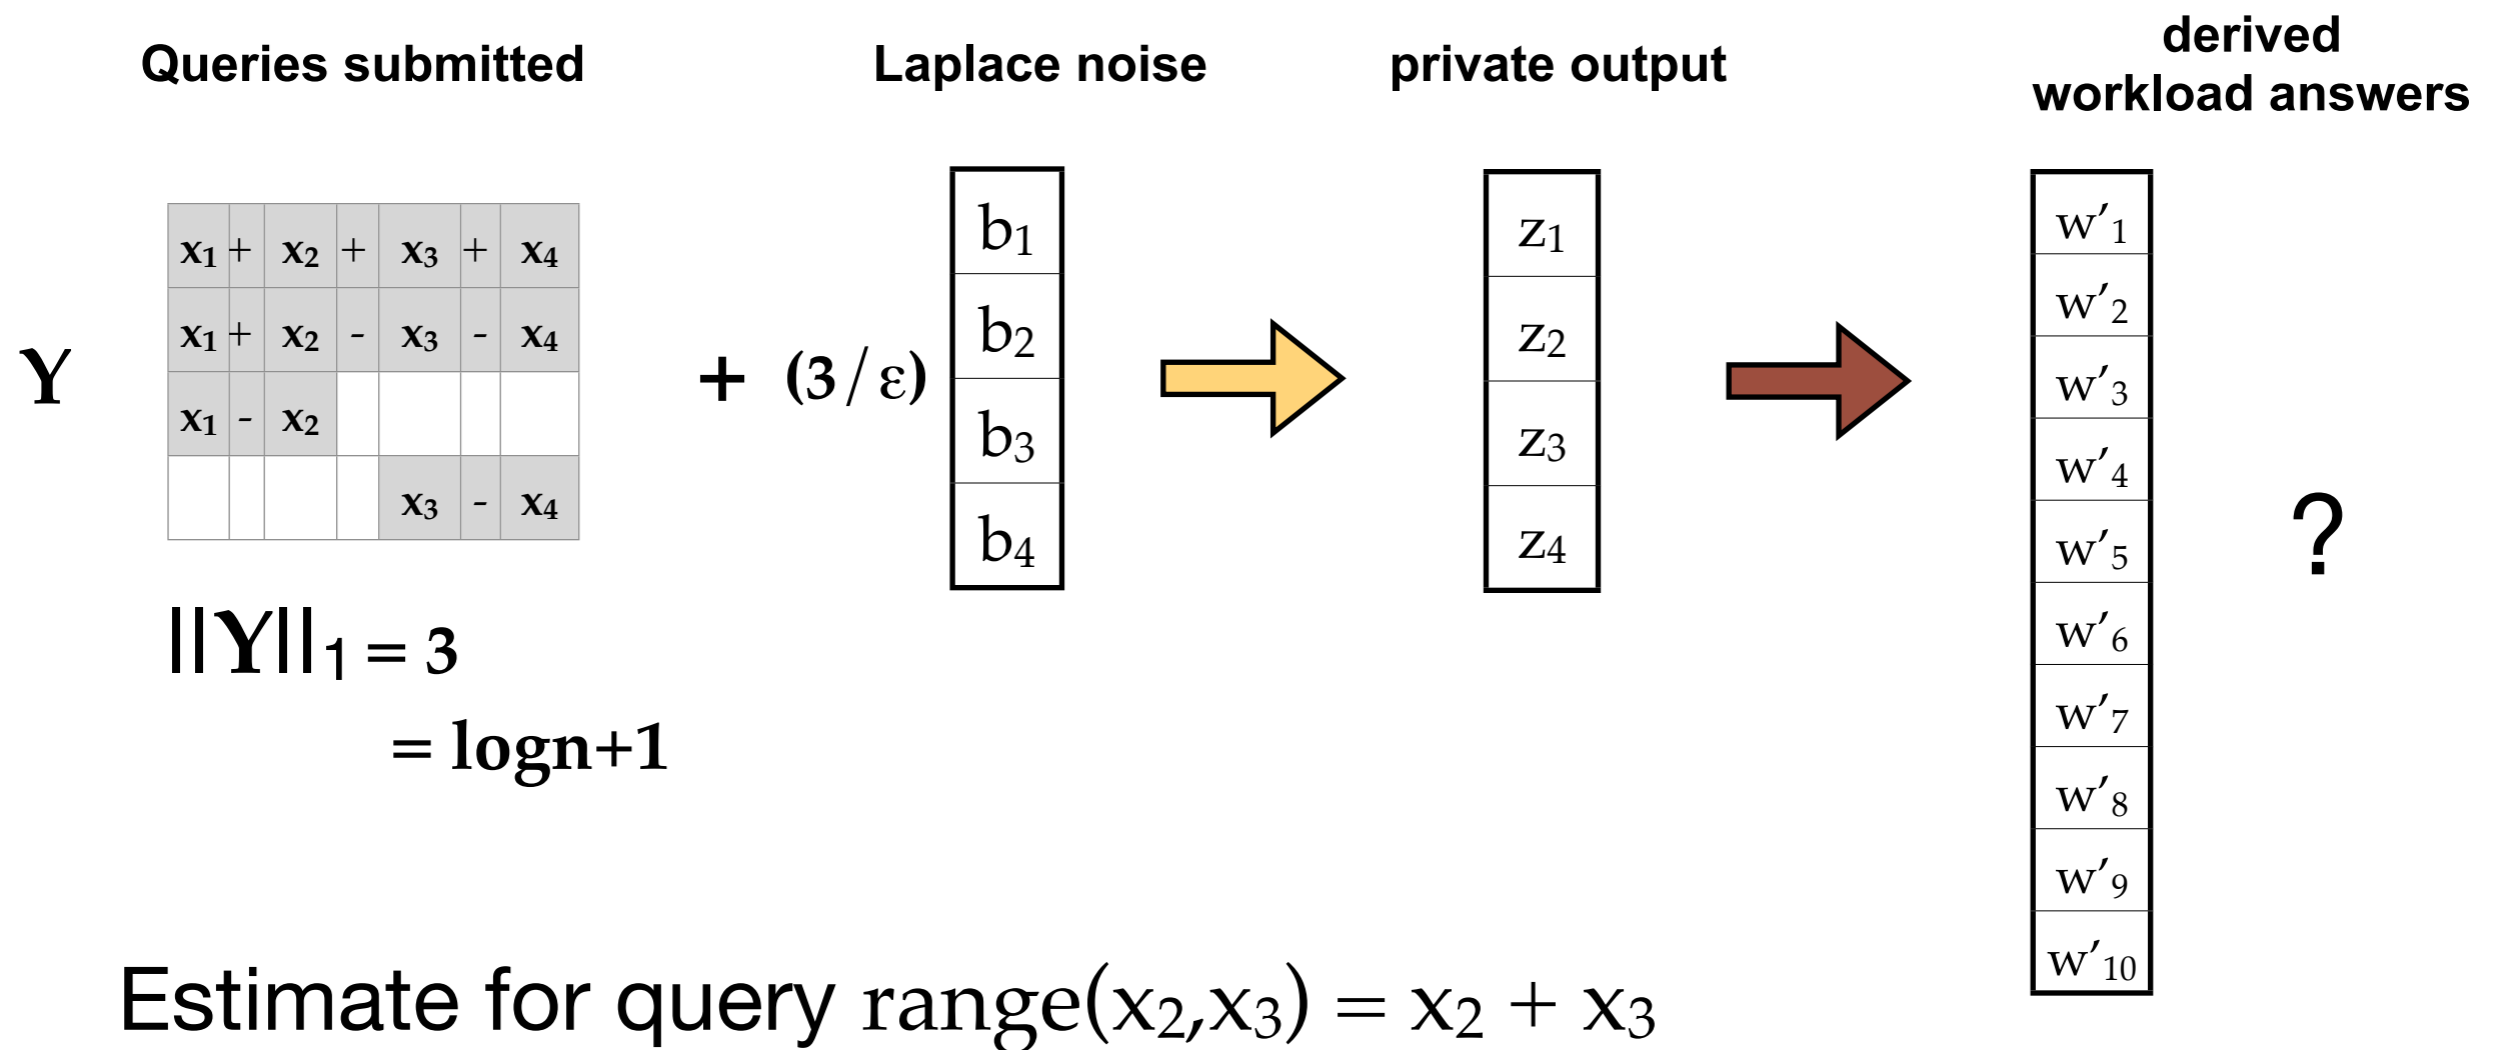
Error rates: workload of all range queries

ϵ -differential privacy



Approach 4: wavelet queries

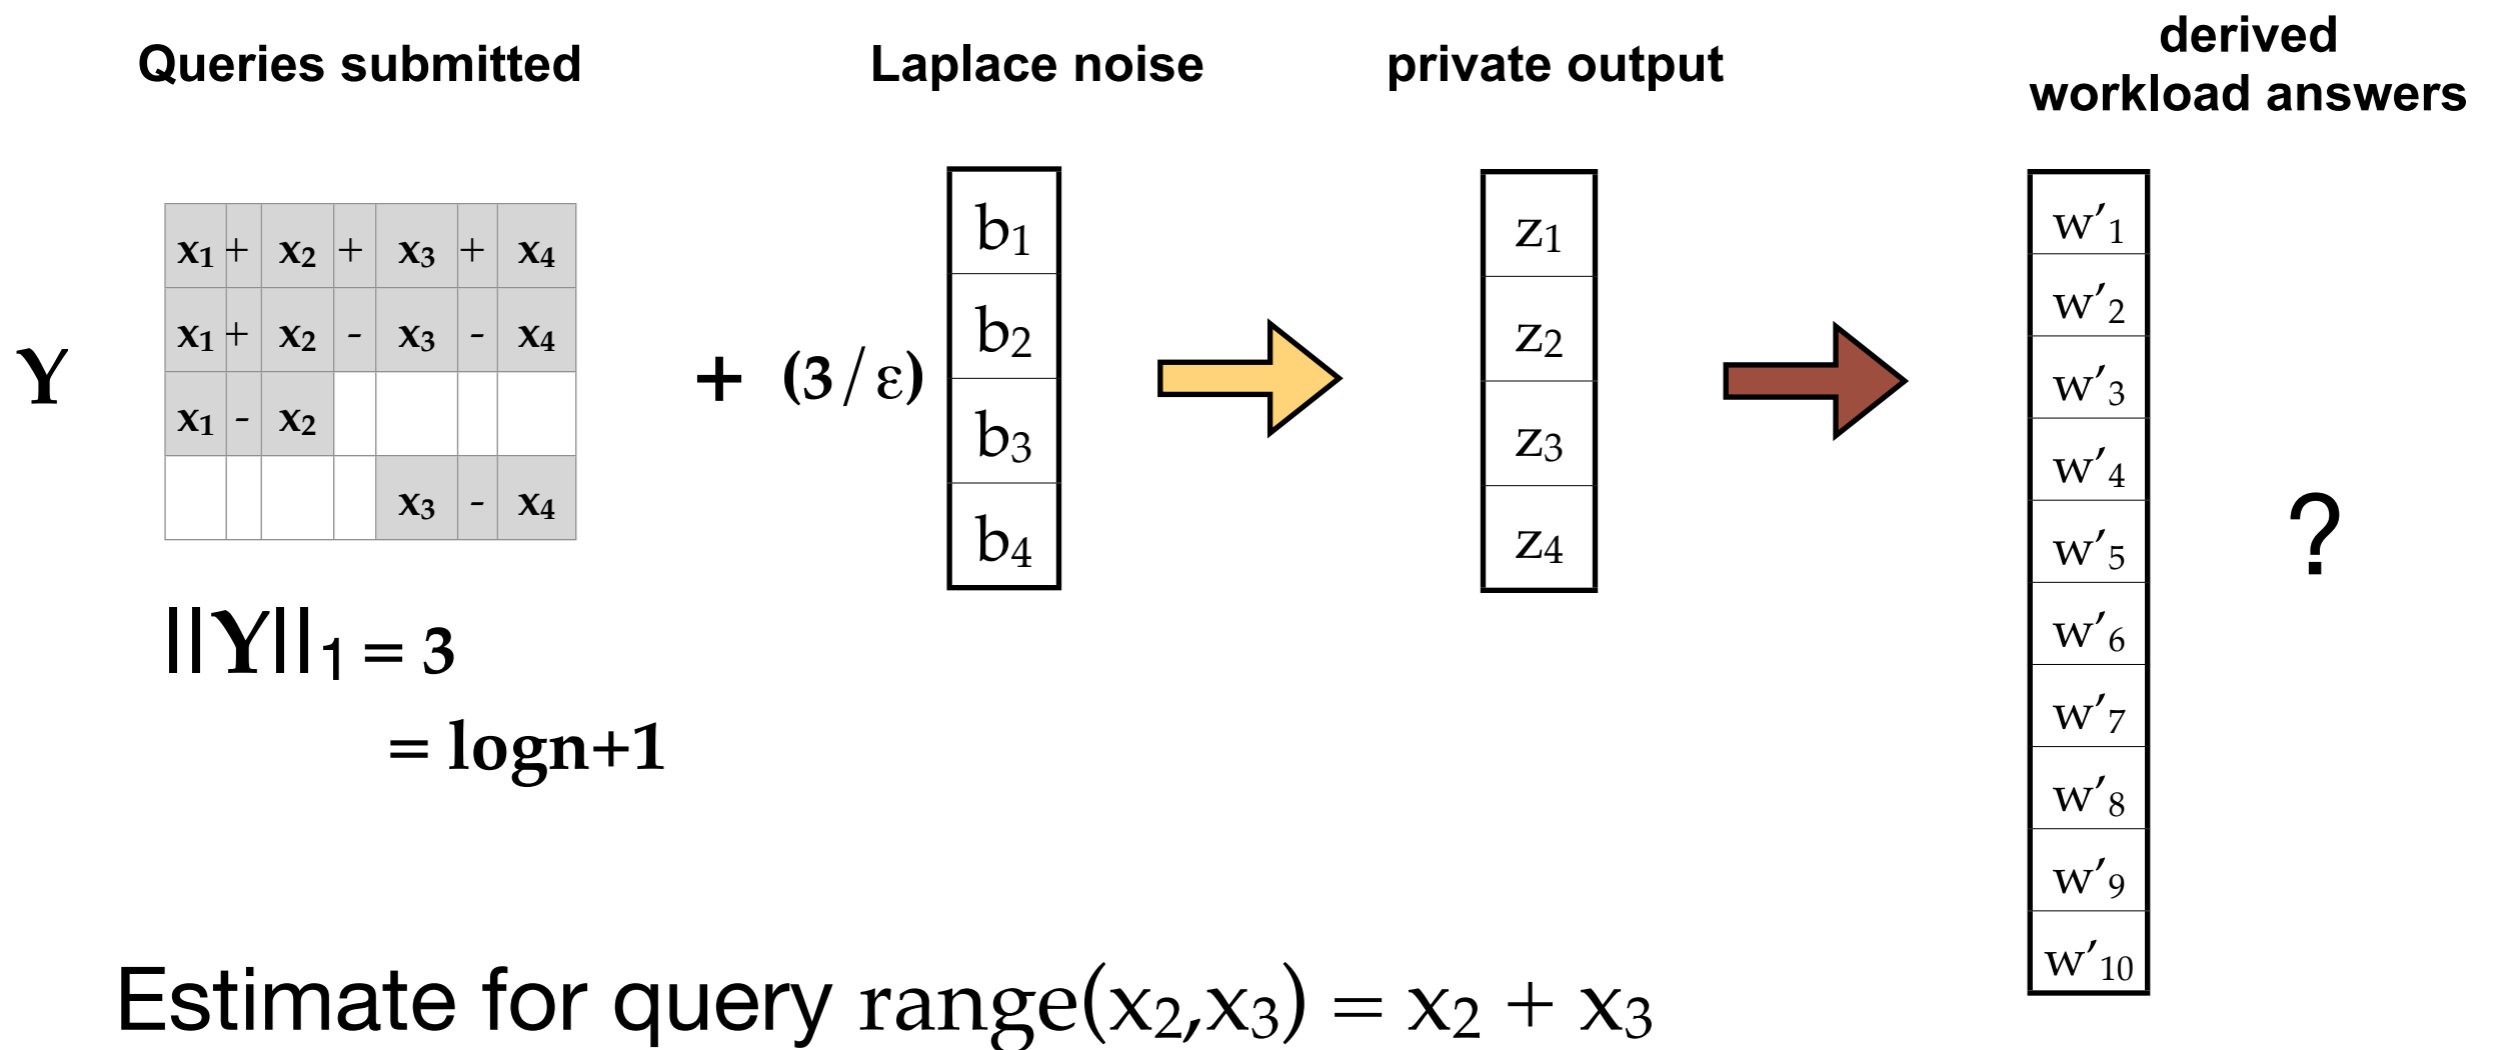
Wavelet queries: use Haar wavelet to get noisy summary of data.



$$.5z_1 + 0z_2 - .5z_3 + .5z_4$$

Approach 4: wavelet queries

Wavelet queries: use Haar wavelet to get noisy summary of data.



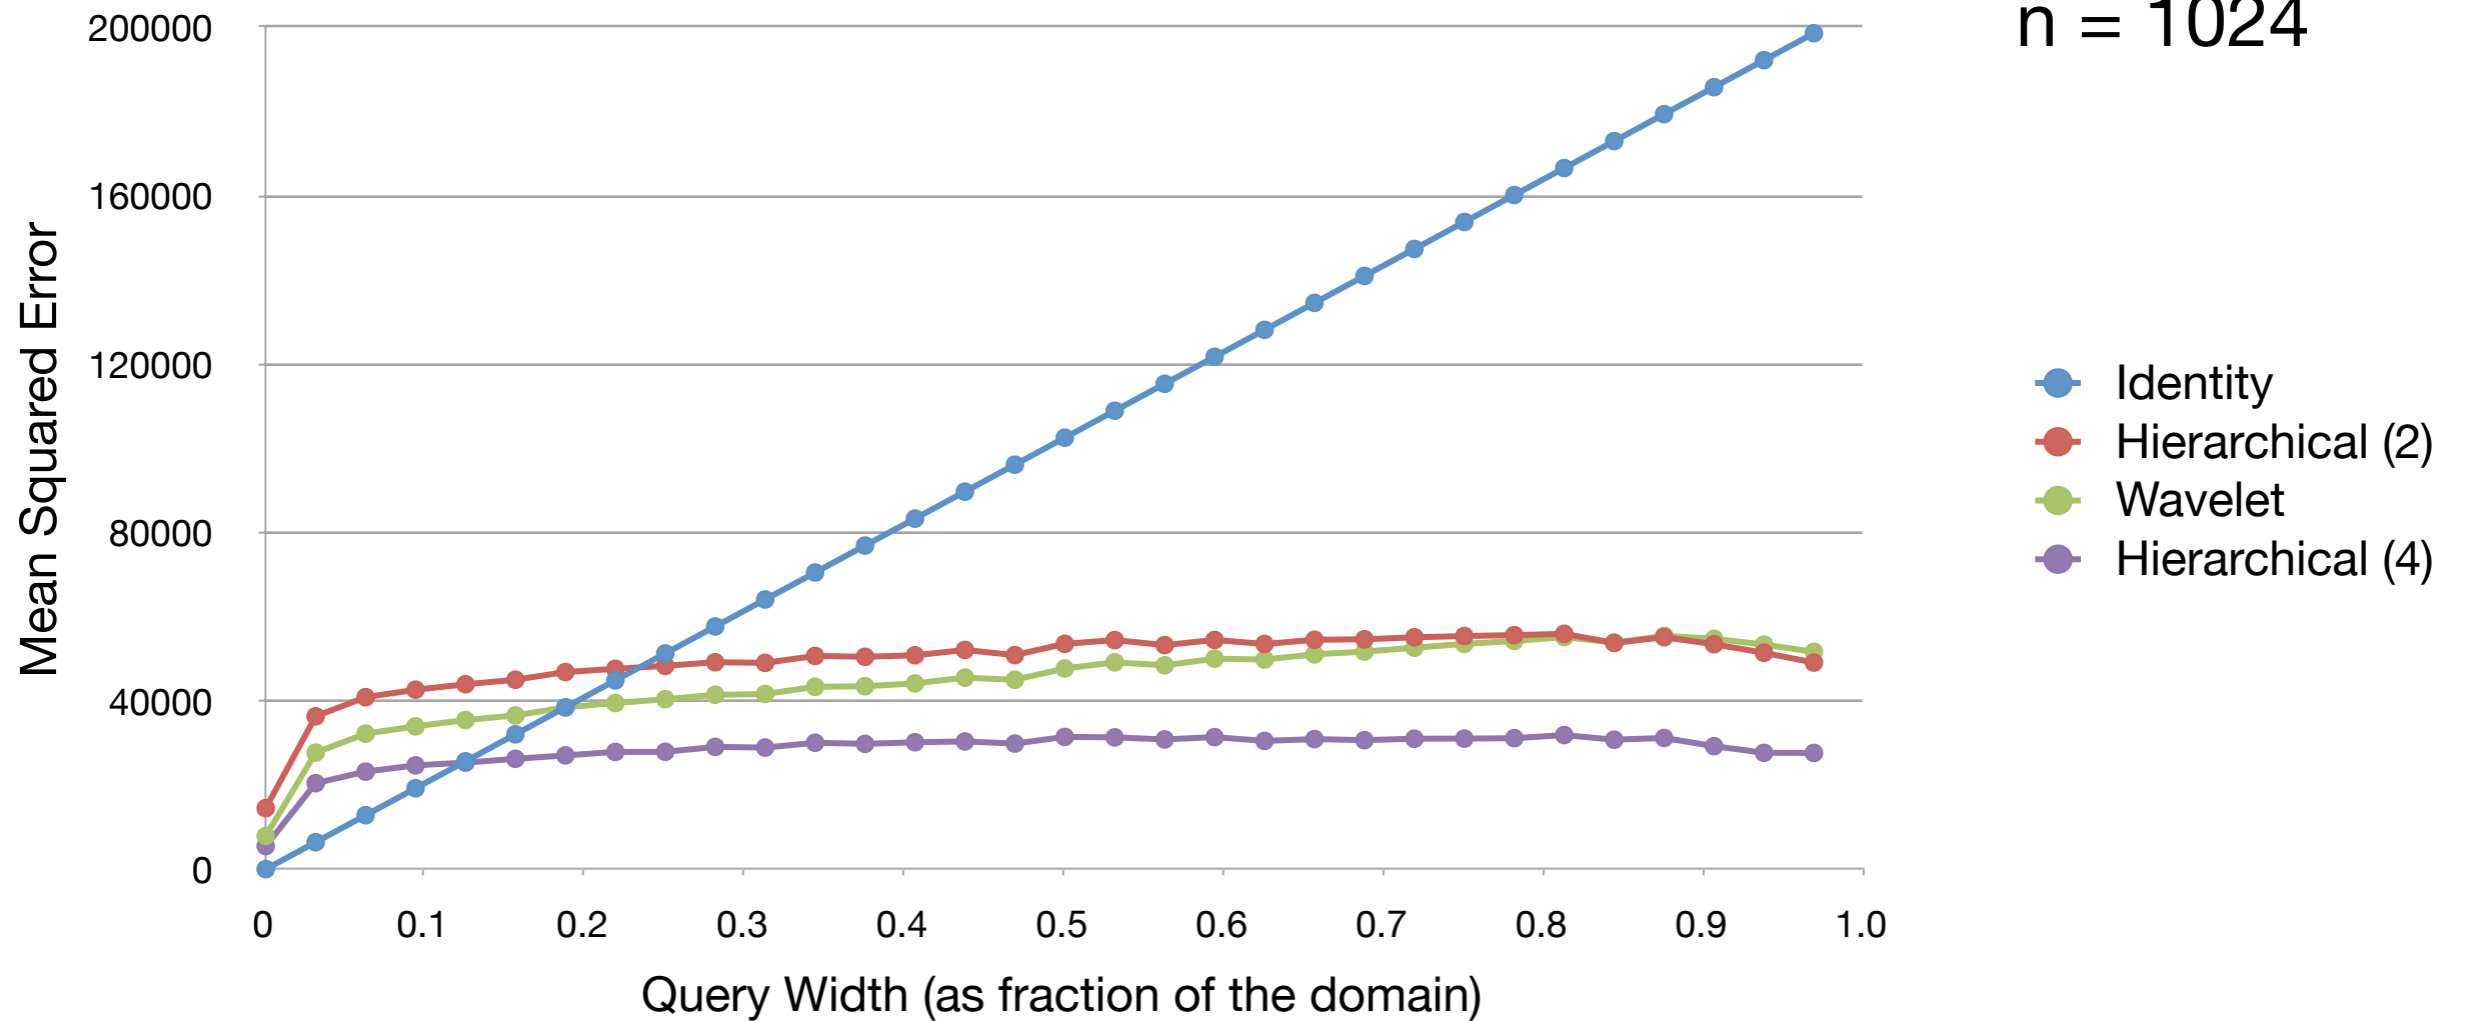
$$.5z_1 + 0z_2 - .5z_3 + .5z_4$$

Error: workload of all range queries

ϵ -differential privacy

$\epsilon = 0.1$

$n = 1024$



Strategies for workload of all range queries

Noisy counts

x_1						
		x_2				
				x_3		
						x_4

I

Very low sensitivity, but large ranges estimated badly.

Max/Avg
error

$$O(n/\epsilon^2)$$

Hierarchical

x_1	+	x_2	+	x_3	+	x_4
x_1	+	x_2				
				x_3	+	x_4
x_1						
		x_2				
				x_3		
						x_4

H

Low sensitivity, and all range queries can be estimated using no more than $\log n$ output entries.

$$O(\log^3 n / \epsilon^2)$$

Wavelet

x_1	+	x_2	+	x_3	+	x_4
x_1	+	x_2	-	x_3	-	x_4
x_1	-	x_2				
				x_3	-	x_4

Y

$$O(\log^3 n / \epsilon^2)$$

Questions raised

Questions raised

- Are these approaches optimal for all range queries?

Questions raised

- Are these approaches optimal for all range queries?
- What about other workloads?

Questions raised

- Are these approaches optimal for all range queries?
- What about other workloads?
- Big picture:
 - x values we cannot observe directly.
 - we can request noisy estimates of any linear function of the x values, at some cost.
 - what should we request to perform our task (i.e. answer workload queries) ?

Questions raised

- Are these approaches optimal for all range queries?
- What about other workloads?
- Big picture:
 - x values we cannot observe directly.
 - we can request noisy estimates of any linear function of the x values, at some cost.
 - what should we request to perform our task (i.e. answer workload queries) ?

Optimal
Experimental
Design

Outline

1. Case study: answering 1-dim range queries
2. The matrix mechanism -- formal description.
3. The matrix mechanism -- new tools & techniques.
4. Conclusion

Outline

1. Case study: answering 1-dim range queries
2. The matrix mechanism -- formal description.
3. The matrix mechanism -- new tools & techniques.
4. Conclusion

Linear counting queries

A **linear counting query** w computes a linear combination of the frequency vector counts:

$$w(\mathbf{D}) = w_1x_1 + w_2x_2 + \dots + w_nx_n \quad \text{each } w_i \in \mathbb{R}$$

1) Expressiveness of queries

2) Need to list ALL -- don't omit those to be derived

3) Can scale rows to error rates of each

Linear counting queries

1) Expressiveness of queries

2) Need to list ALL -- don't omit those to be derived

3) Can scale rows to error rates of each

A **linear counting query** w computes a linear combination of the frequency vector counts:

$$\mathbf{w}(\mathbf{D}) = w_1x_1 + w_2x_2 + \dots + w_nx_n \quad \text{each } w_i \in \mathbb{R}$$

... as a length n row vector:

$$\mathbf{w} = [w_1, w_2, w_3 \dots w_n]$$

The query result is:

$$\mathbf{w}\mathbf{x}$$

Linear counting queries

1) Expressiveness of queries

2) Need to list ALL -- don't omit those that can be derived

3) Can scale rows to control error rates of each

A **linear counting query** w computes a linear combination of the frequency vector counts:

$$w(\mathbf{D}) = w_1x_1 + w_2x_2 + \dots + w_nx_n \quad \text{each } w_i \in \mathbb{R}$$

... as a length n row vector:

$$w = [w_1, w_2, w_3 \dots w_n]$$

The query result is:

$$w\mathbf{x}$$

a set of linear counting queries is a matrix:

$$W = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 & -1 \end{bmatrix}$$

The query result is:

$$W\mathbf{x}$$

The sensitivity of a query matrix

- For two neighboring databases D and D', their frequency vectors x and x' will differ in one position, by exactly 1.

$$\begin{array}{c} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix} \\ \text{answers} \end{array} = \begin{array}{c} \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 & -1 \end{bmatrix} \\ \text{query matrix } \mathbf{W} \end{array} \times \begin{array}{c} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 \\ x_8 \\ x_9 \\ x_{10} \end{bmatrix} \\ \mathbf{x}' \end{array}$$

The value x_7 in the frequency vector \mathbf{x}' is highlighted in yellow and is equal to $+1$.

The sensitivity of a query matrix

- For two neighboring databases D and D', their frequency vectors x and x' will differ in one position, by exactly 1.

$$\begin{array}{c} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix} \\ \text{answers} \end{array} = \begin{array}{c} \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 & -1 \end{bmatrix} \\ \text{query matrix } \mathbf{W} \end{array} \times \begin{array}{c} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 \\ x_8 \\ x_9 \\ x_{10} \end{bmatrix} \\ \mathbf{x}' \end{array}$$

The diagram illustrates the relationship between answers, a query matrix, and frequency vectors. The query matrix \mathbf{W} is a 4x10 matrix. The 7th column of \mathbf{W} is highlighted in yellow, with values 1, 0, 1, and -1. The 7th element of the frequency vector \mathbf{x}' is highlighted in yellow, with a value of +1. This indicates that the frequency of the 7th item in the database D' is exactly 1 unit higher than in database D.

The sensitivity of a query matrix

- For two neighboring databases D and D', their frequency vectors x and x' will differ in one position, by exactly 1.

$$\begin{array}{c} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix} \\ \text{answers} \end{array} = \begin{array}{c} \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 & -1 \end{bmatrix} \\ \text{query matrix } \mathbf{W} \\ \|\mathbf{W}\|_1 = 4 \end{array} \times \begin{array}{c} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 + 1 \\ x_8 \\ x_9 \\ x_{10} \end{bmatrix} \\ \mathbf{x}' \end{array}$$

The sensitivity of a query matrix

- For two neighboring databases D and D' , their frequency vectors x and x' will differ in one position, by exactly 1.

$$\begin{array}{c}
 \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix} \\
 \text{answers}
 \end{array}
 =
 \begin{array}{c}
 \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\
 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\
 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
 1 & 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 & -1 \end{bmatrix} \\
 \text{query matrix } \mathbf{W} \\
 \|\mathbf{W}\|_1 = 4
 \end{array}
 \times
 \begin{array}{c}
 \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 + 1 \\ x_8 \\ x_9 \\ x_{10} \end{bmatrix} \\
 \mathbf{x}'
 \end{array}$$

The **L_1 sensitivity** of a query matrix is:
the maximum L_1 norm of the columns.

The sensitivity of a query matrix

- For two neighboring databases D and D' , their frequency vectors x and x' will differ in one position, by exactly 1.

$$\begin{array}{c}
 \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix} \\
 \text{answers}
 \end{array}
 =
 \begin{array}{c}
 \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\
 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\
 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
 1 & 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 & -1 \end{bmatrix} \\
 \text{query matrix } \mathbf{W} \\
 \|\mathbf{W}\|_1 = 4
 \end{array}
 \times
 \begin{array}{c}
 \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 + 1 \\ x_8 \\ x_9 \\ x_{10} \end{bmatrix} \\
 \mathbf{x}'
 \end{array}$$

The **L_1 sensitivity** of a query matrix is: the maximum L_1 norm of the columns.

(The **L_2 sensitivity** of a query matrix is: the maximum L_2 norm of the columns.)

Laplace mechanism (matrix notation)

$$\text{Laplace}(W, \mathbf{x}) = W\mathbf{x} + (\|W\|_1 / \epsilon)\mathbf{b}$$

Laplace mechanism (matrix notation)

$$\text{Laplace}(W, x) = Wx + (\|W\|_1 / \epsilon) \mathbf{b}$$

Diagram illustrating the Laplace mechanism in matrix notation:

- The private output vector $\begin{bmatrix} z_1 \\ z_2 \\ \vdots \\ z_m \end{bmatrix}$ is equal to the product of the workload matrix W and the data vector x , plus a noise vector b scaled by $(\|W\|_1 / \epsilon)$.
- The workload matrix W is labeled "workload of queries" and has m rows (queries) and n columns.
- The data vector x is labeled "data".
- The noise vector b is labeled "noise".

m independent samples from Laplace(1)

Laplace mechanism (matrix notation)

$$\text{Laplace}(W, x) = Wx + (\|W\|_1 / \epsilon) \mathbf{b}$$

The diagram illustrates the Laplace mechanism in matrix notation. It shows the private output vector \mathbf{z} , the workload matrix W , the data vector \mathbf{x} , and the noise vector \mathbf{b} .

The private output vector \mathbf{z} is shown as a column vector with elements z_1, z_2, \dots, z_m , labeled "private output".

The workload matrix W is shown as a matrix with m rows and n columns, labeled "workload of queries" and W . The rows are labeled "m queries" and the columns are labeled "n columns". The elements are $w_{11}, w_{12}, \dots, w_{1n}$ in the first row, $w_{21}, w_{22}, \dots, w_{2n}$ in the second row, and $w_{m1}, w_{m2}, \dots, w_{mn}$ in the m -th row.

The data vector \mathbf{x} is shown as a column vector with elements x_1, x_2, \dots, x_n , labeled "data" and \mathbf{x} .

The noise vector \mathbf{b} is shown as a column vector with elements b_1, b_2, \dots, b_m , labeled "noise" and \mathbf{b} . An orange arrow points from a box below to this vector, indicating that the noise is composed of m independent samples from $\text{Laplace}(1)$.

The equation is:
$$\begin{bmatrix} z_1 \\ z_2 \\ \vdots \\ z_m \end{bmatrix} = \begin{bmatrix} w_{11} & w_{12} & \dots & w_{1n} \\ w_{21} & w_{22} & \dots & w_{2n} \\ \vdots & \vdots & \vdots & \vdots \\ w_{m1} & w_{m2} & \dots & w_{mn} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} + (\|W\|_1 / \epsilon) \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_m \end{bmatrix}$$

$$\text{Error}(w) = 2 (\|W\|_1 / \epsilon)^2$$

m independent samples from $\text{Laplace}(1)$

The matrix mechanism

est-x can be viewed
as a synthetic
database.

Workload query
answers consistent

The matrix mechanism

- 1 **(Design)** Choose a **full rank** query strategy **A**

est-x can be viewed
as a synthetic
database.

Workload query
answers consistent

The matrix mechanism

- 1 **(Design)** Choose a **full rank** query strategy **A**
- 2 **(Apply Laplace)** Use the Laplace mechanism to answer **A**

est-x can be viewed
as a synthetic
database.

Workload query
answers consistent

The matrix mechanism

est- x can be viewed
as a synthetic
database.

Workload query
answers consistent

- 1 **(Design)** Choose a **full rank** query strategy **A**
- 2 **(Apply Laplace)** Use the Laplace mechanism to answer **A**

$$z = Ax + (\|A\|_1 / \epsilon)b$$

The matrix mechanism

est- \underline{x} can be viewed
as a synthetic
database.

Workload query
answers consistent

- 1 **(Design)** Choose a **full rank** query strategy \mathbf{A}
- 2 **(Apply Laplace)** Use the Laplace mechanism to answer \mathbf{A}

$$\mathbf{z} = \mathbf{A}\mathbf{x} + (\|\mathbf{A}\|_1 / \epsilon)\mathbf{b}$$

- 3 **(Derivation)** Compute estimate $\underline{\mathbf{x}}$ of \mathbf{x} using answers \mathbf{z} .

The matrix mechanism

est- \underline{x} can be viewed
as a synthetic
database.

Workload query
answers consistent

- 1 **(Design)** Choose a **full rank** query strategy \mathbf{A}
- 2 **(Apply Laplace)** Use the Laplace mechanism to answer \mathbf{A}

$$\mathbf{z} = \mathbf{Ax} + (\|\mathbf{A}\|_1 / \epsilon)\mathbf{b}$$

- 3 **(Derivation)** Compute estimate $\underline{\mathbf{x}}$ of \mathbf{x} using answers \mathbf{z} .
 - compute estimate $\underline{\mathbf{x}}$ of \mathbf{x} that minimizes squared error:

$$\|\mathbf{A}\underline{\mathbf{x}} - \mathbf{z}\|_2^2$$

The matrix mechanism

est- \underline{x} can be viewed
as a synthetic
database.

Workload query
answers consistent

- 1 **(Design)** Choose a **full rank** query strategy \mathbf{A}
- 2 **(Apply Laplace)** Use the Laplace mechanism to answer \mathbf{A}

$$\mathbf{z} = \mathbf{Ax} + (\|\mathbf{A}\|_1 / \epsilon)\mathbf{b}$$

- 3 **(Derivation)** Compute estimate $\underline{\mathbf{x}}$ of \mathbf{x} using answers \mathbf{z} .
 - compute estimate $\underline{\mathbf{x}}$ of \mathbf{x} that minimizes squared error:

$$\|\mathbf{Ax} - \mathbf{z}\|_2^2$$

$$\underline{\mathbf{x}} = \mathbf{A}^+ \mathbf{z}$$

where $\mathbf{A}^+ = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T$

The matrix mechanism

est- \underline{x} can be viewed as a synthetic database.

Workload query answers consistent

- 1 **(Design)** Choose a **full rank** query strategy \mathbf{A}
- 2 **(Apply Laplace)** Use the Laplace mechanism to answer \mathbf{A}

$$\mathbf{z} = \mathbf{Ax} + (\|\mathbf{A}\|_1 / \epsilon)\mathbf{b}$$

- 3 **(Derivation)** Compute estimate $\underline{\mathbf{x}}$ of \mathbf{x} using answers \mathbf{z} .
 - compute estimate $\underline{\mathbf{x}}$ of \mathbf{x} that minimizes squared error:

$$\|\mathbf{Ax} - \mathbf{z}\|_2^2$$

Thm: $\underline{\mathbf{x}}$ is unbiased and has the least variance among all linear unbiased estimators.

$$\underline{\mathbf{x}} = \mathbf{A}^+ \mathbf{z}$$

where $\mathbf{A}^+ = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T$

Strategy matrices for the range queries

Strategy matrices for the range queries

Identity

1	0	0	0
0	1	0	0
0	0	1	0
0	0	0	1

I

Hierarchical

1	1	1	1
1	1	0	0
0	0	1	1
1	0	0	0
0	1	0	0
0	0	1	0
0	0	0	1

H

Wavelet

1	1	1	1
1	1	-1	-1
1	-1	0	0
0	0	1	-1

Y

The matrix mechanism

Given a workload W , and any full-rank strategy matrix A , the following randomized algorithm is ϵ -differentially private:


$$\text{Matrix}_A(W, x) = Wx + (\|A\|_1 / \epsilon) WA^+ b \quad b = \text{Lap}(1)$$

The matrix mechanism

Given a workload W , and any full-rank strategy matrix A , the following randomized algorithm is ϵ -differentially private:

$$\mathbf{Matrix}_A(W, x) = Wx + (\|A\|_1 / \epsilon) WA^+ \mathbf{b} \quad \mathbf{b} = \text{Lap}(1)$$

instantiated with
strategy A

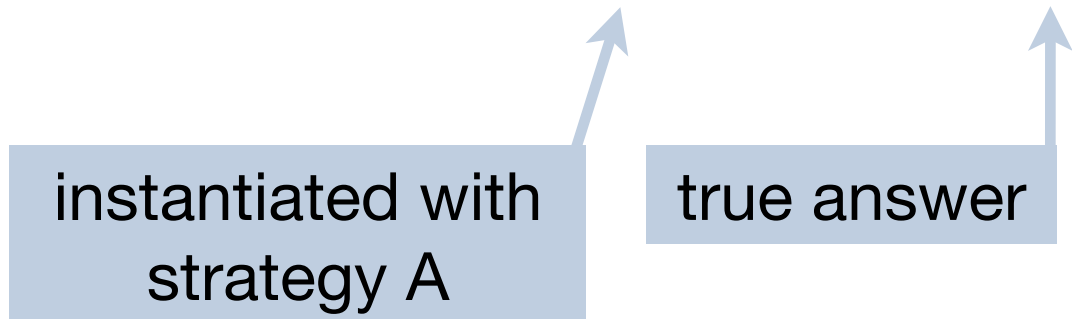


The matrix mechanism

Given a workload W , and any full-rank strategy matrix A , the following randomized algorithm is ϵ -differentially private:

$$\mathbf{Matrix}_A(W, x) = Wx + (\|A\|_1 / \epsilon) WA^+ \mathbf{b} \quad \mathbf{b} = \text{Lap}(1)$$

instantiated with
strategy A



true answer

The matrix mechanism

Given a workload W , and any full-rank strategy matrix A , the following randomized algorithm is ϵ -differentially private:

$$\text{Matrix}_A(W, x) = Wx + (\|A\|_1 / \epsilon) WA^+ b \quad b = \text{Lap}(1)$$

instantiated with
strategy A

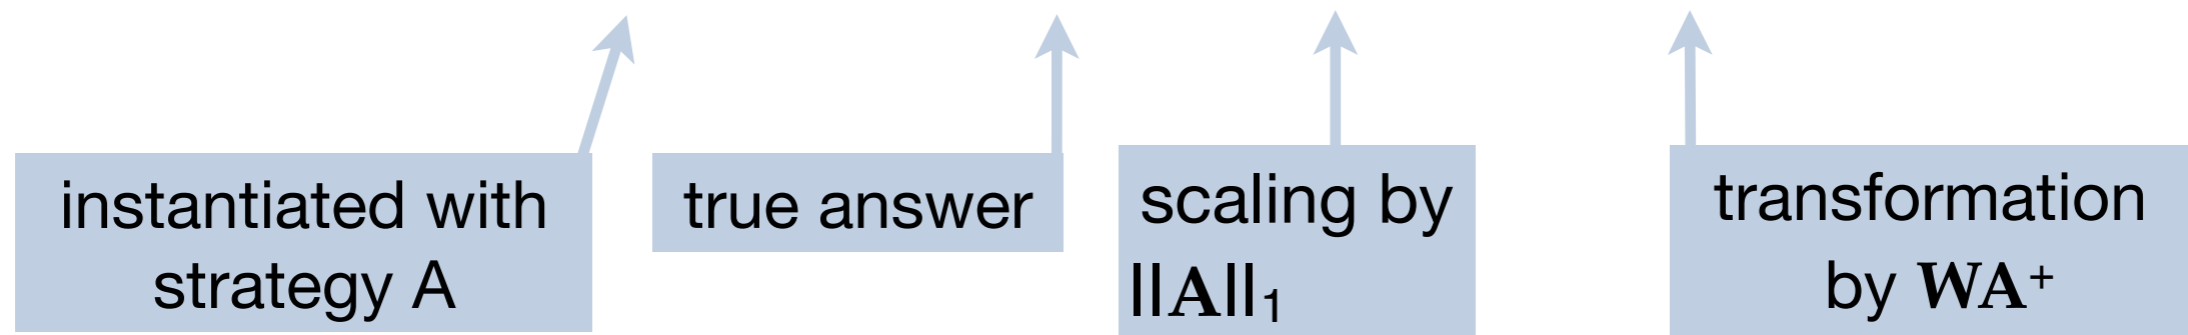
true answer

scaling by
 $\|A\|_1$

The matrix mechanism

Given a workload W , and any full-rank strategy matrix A , the following randomized algorithm is ϵ -differentially private:


$$\text{Matrix}_A(W, x) = Wx + (\|A\|_1 / \epsilon) WA^+ b \quad b = \text{Lap}(1)$$



The matrix mechanism

Given a workload W , and any full-rank strategy matrix A , the following randomized algorithm is ϵ -differentially private:

$$\text{Matrix}_A(W, x) = Wx + (\|A\|_1 / \epsilon) WA^+ b \quad b = \text{Lap}(1)$$



Compare with the Laplace mechanism:

$$\text{Laplace}(W, x) = Wx + (\|W\|_1 / \epsilon) b$$

NOTE: This error is completely independent of the input data!!

Error of the matrix mechanism

Given any full rank strategy \mathbf{A} , and any linear workload query \mathbf{w} , the error of the mechanism $\text{Matrix}_{\mathbf{A}}$ on query \mathbf{w} is:

$$\text{Error}_{\mathbf{A}}(\mathbf{w}) = (2 / \varepsilon^2) (\|\mathbf{A}\|_1)^2 \mathbf{w}(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{w}^T$$

NOTE: This error is completely independent of the input data!!

Error of the matrix mechanism

Given any full rank strategy \mathbf{A} , and any linear workload query \mathbf{w} , the error of the mechanism $\text{Matrix}_{\mathbf{A}}$ on query \mathbf{w} is:

$$\text{Error}_{\mathbf{A}}(\mathbf{w}) = (2 / \epsilon^2) (\|\mathbf{A}\|_1)^2 \mathbf{w}(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{w}^T$$

sensitivity term

NOTE: This error is completely independent of the input data!!

Error of the matrix mechanism

Given any full rank strategy \mathbf{A} , and any linear workload query \mathbf{w} , the error of the mechanism $\text{Matrix}_{\mathbf{A}}$ on query \mathbf{w} is:

$$\text{Error}_{\mathbf{A}}(\mathbf{w}) = (2 / \epsilon^2) (\|\mathbf{A}\|_1)^2 \mathbf{w}(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{w}^T$$

sensitivity term

“error profile” term

Profile equivalence

Definition: strategy matrices A and B are **profile equivalent** if $(A^T A) = (B^T B)$

Matrix_A **Error_A(w) = $(2 / \epsilon^2) (\|A\|_1)^2 w(A^T A)^{-1} w^T$**

Matrix_B **Error_B(w) = $(2 / \epsilon^2) (\|B\|_1)^2 w(B^T B)^{-1} w^T$**

If $(A^T A) = (B^T B)$ and $\|A\|_1 \leq \|B\|_1$ then **Matrix_A** has lower error than **Matrix_B** for **every** query.

Strategies equivalent to wavelet

1	1	1	1
1	1	-1	-1
1	-1	0	0
0	0	1	-1

Wavelet Y

$$\|Y\|_1 = 3$$

\equiv

1	1	0	0
0	0	1	1
1	0	0	0
0	1	0	0
0	0	1	0
0	0	0	1
1	0	0	0
0	1	0	0
0	0	1	0
0	0	0	1

Y'

$$\|Y'\|_1 = 3$$

\succ

1	1	0	0
0	0	1	1
$\sqrt{2}$	0	0	0
0	$\sqrt{2}$	0	0
0	0	$\sqrt{2}$	0
0	0	0	$\sqrt{2}$

Y''

$$\|Y''\|_1 = 2.414$$

Strategies equivalent to wavelet

1	1	1	1
1	1	-1	-1
1	-1	0	0
0	0	1	-1

Wavelet Y

$$\|Y\|_1 = 3$$

\equiv

1	1	0	0
0	0	1	1
1	0	0	0
0	1	0	0
0	0	1	0
0	0	0	1
1	0	0	0
0	1	0	0
0	0	1	0
0	0	0	1

Y'

$$\|Y'\|_1 = 3$$

\succ

1	1	0	0
0	0	1	1
$\sqrt{2}$	0	0	0
0	$\sqrt{2}$	0	0
0	0	$\sqrt{2}$	0
0	0	0	$\sqrt{2}$

Y''

$$\|Y''\|_1 = 2.414$$

Neither the hierarchical nor the wavelet strategy is **efficient**, i.e. there exist uniformly better strategies with matching error profiles.

Finding the optimal strategy

Objective: given workload W , find the query strategy A that minimizes the total error.

Error for a single query:

$$\text{Error}_A(\mathbf{w}) = (2 / \varepsilon^2)(\|A\|_1)^2 \mathbf{w}(A^T A)^{-1} \mathbf{w}^T$$

Total error for a workload of queries:

$$\begin{aligned} \text{TotalError}_A(\mathbf{w}) &= (2 / \varepsilon^2)(\|A\|_1)^2 \text{trace}(\mathbf{W}(A^T A)^{-1} \mathbf{W}^T) \\ &= (2 / \varepsilon^2)(\|A\|_1)^2 \text{trace}(\mathbf{W}^T \mathbf{W}(A^T A)^{-1}) \end{aligned}$$

Overview of problem solutions

Overview of problem solutions

	Objective	Problem Type
--	------------------	---------------------

Overview of problem solutions

	Objective	Problem Type
1	Given W , choose A to minimize TotalError_A(W)	SDP with rank constraints

Overview of problem solutions

	Objective	Problem Type
1	Given W , choose A to minimize $\text{TotalError}_A(W)$	SDP with rank constraints
2	Given $A^T A$, choose Q to minimize $\ A\ _1$	SDP with rank constraints

Overview of problem solutions

	Objective	Problem Type
1	Given W , choose A to minimize TotalError_A(W)	SDP with rank constraints
2	Given $A^T A$, choose Q to minimize $\ A\ _1$	SDP with rank constraints
3	Given W , choose A to minimize TotalError_A(W) under (ϵ, δ) -differential privacy	SDP

Outline

1. Case study: answering 1-dim range queries
2. The matrix mechanism -- formal description.
3. The matrix mechanism -- new tools & techniques.
4. Conclusion

Outline

1. Case study: answering 1-dim range queries
2. The matrix mechanism -- formal description.
3. The matrix mechanism -- new tools & techniques.
4. Conclusion

New techniques

- **Optimal error:** A lower bound on the error of the optimal strategy allows us to assess the quality of existing strategies and explore “workload error complexity”.
- **Efficient strategy selection:** the strategy selection problem can be approximately solved, resulting in strategy matrices customized to arbitrary workloads.
- **Inference for sparse datasets:** by imposing non-negativity constraints during inference, accuracy can be significantly improved. (But analysis of error is harder.)

Singular value bound

- Given workload \mathbf{W} , the optimal total error for \mathbf{W} is greater than or equal to the SVD bound.

THEOREM 3.3. (SINGULAR VALUE BOUND) *Given an $m \times n$ workload \mathbf{W} , let $\lambda_1, \lambda_2, \dots, \lambda_n$ be the singular values of \mathbf{W} .*

$$\min_{\mathbf{A}} \text{ERROR}_{\mathbf{A}}(\mathbf{W}) \geq P(\epsilon, \delta) \frac{1}{n} \left(\sum_{i=1}^n \lambda_i \right)^2,$$

where $P(\epsilon, \delta) = \frac{2 \log(2/\delta)}{\epsilon^2}$.

- Tight: bound is achievable for a certain class of workloads.
- Easy to compute.

Algorithm for efficient strategy selection

- Inspired by **optimal experimental design**
 - Given \mathbf{W} , choose a set of **basis queries** for the strategy:
 - $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ (the eigenvectors of \mathbf{W})
 - compute optimal scalars to minimize error

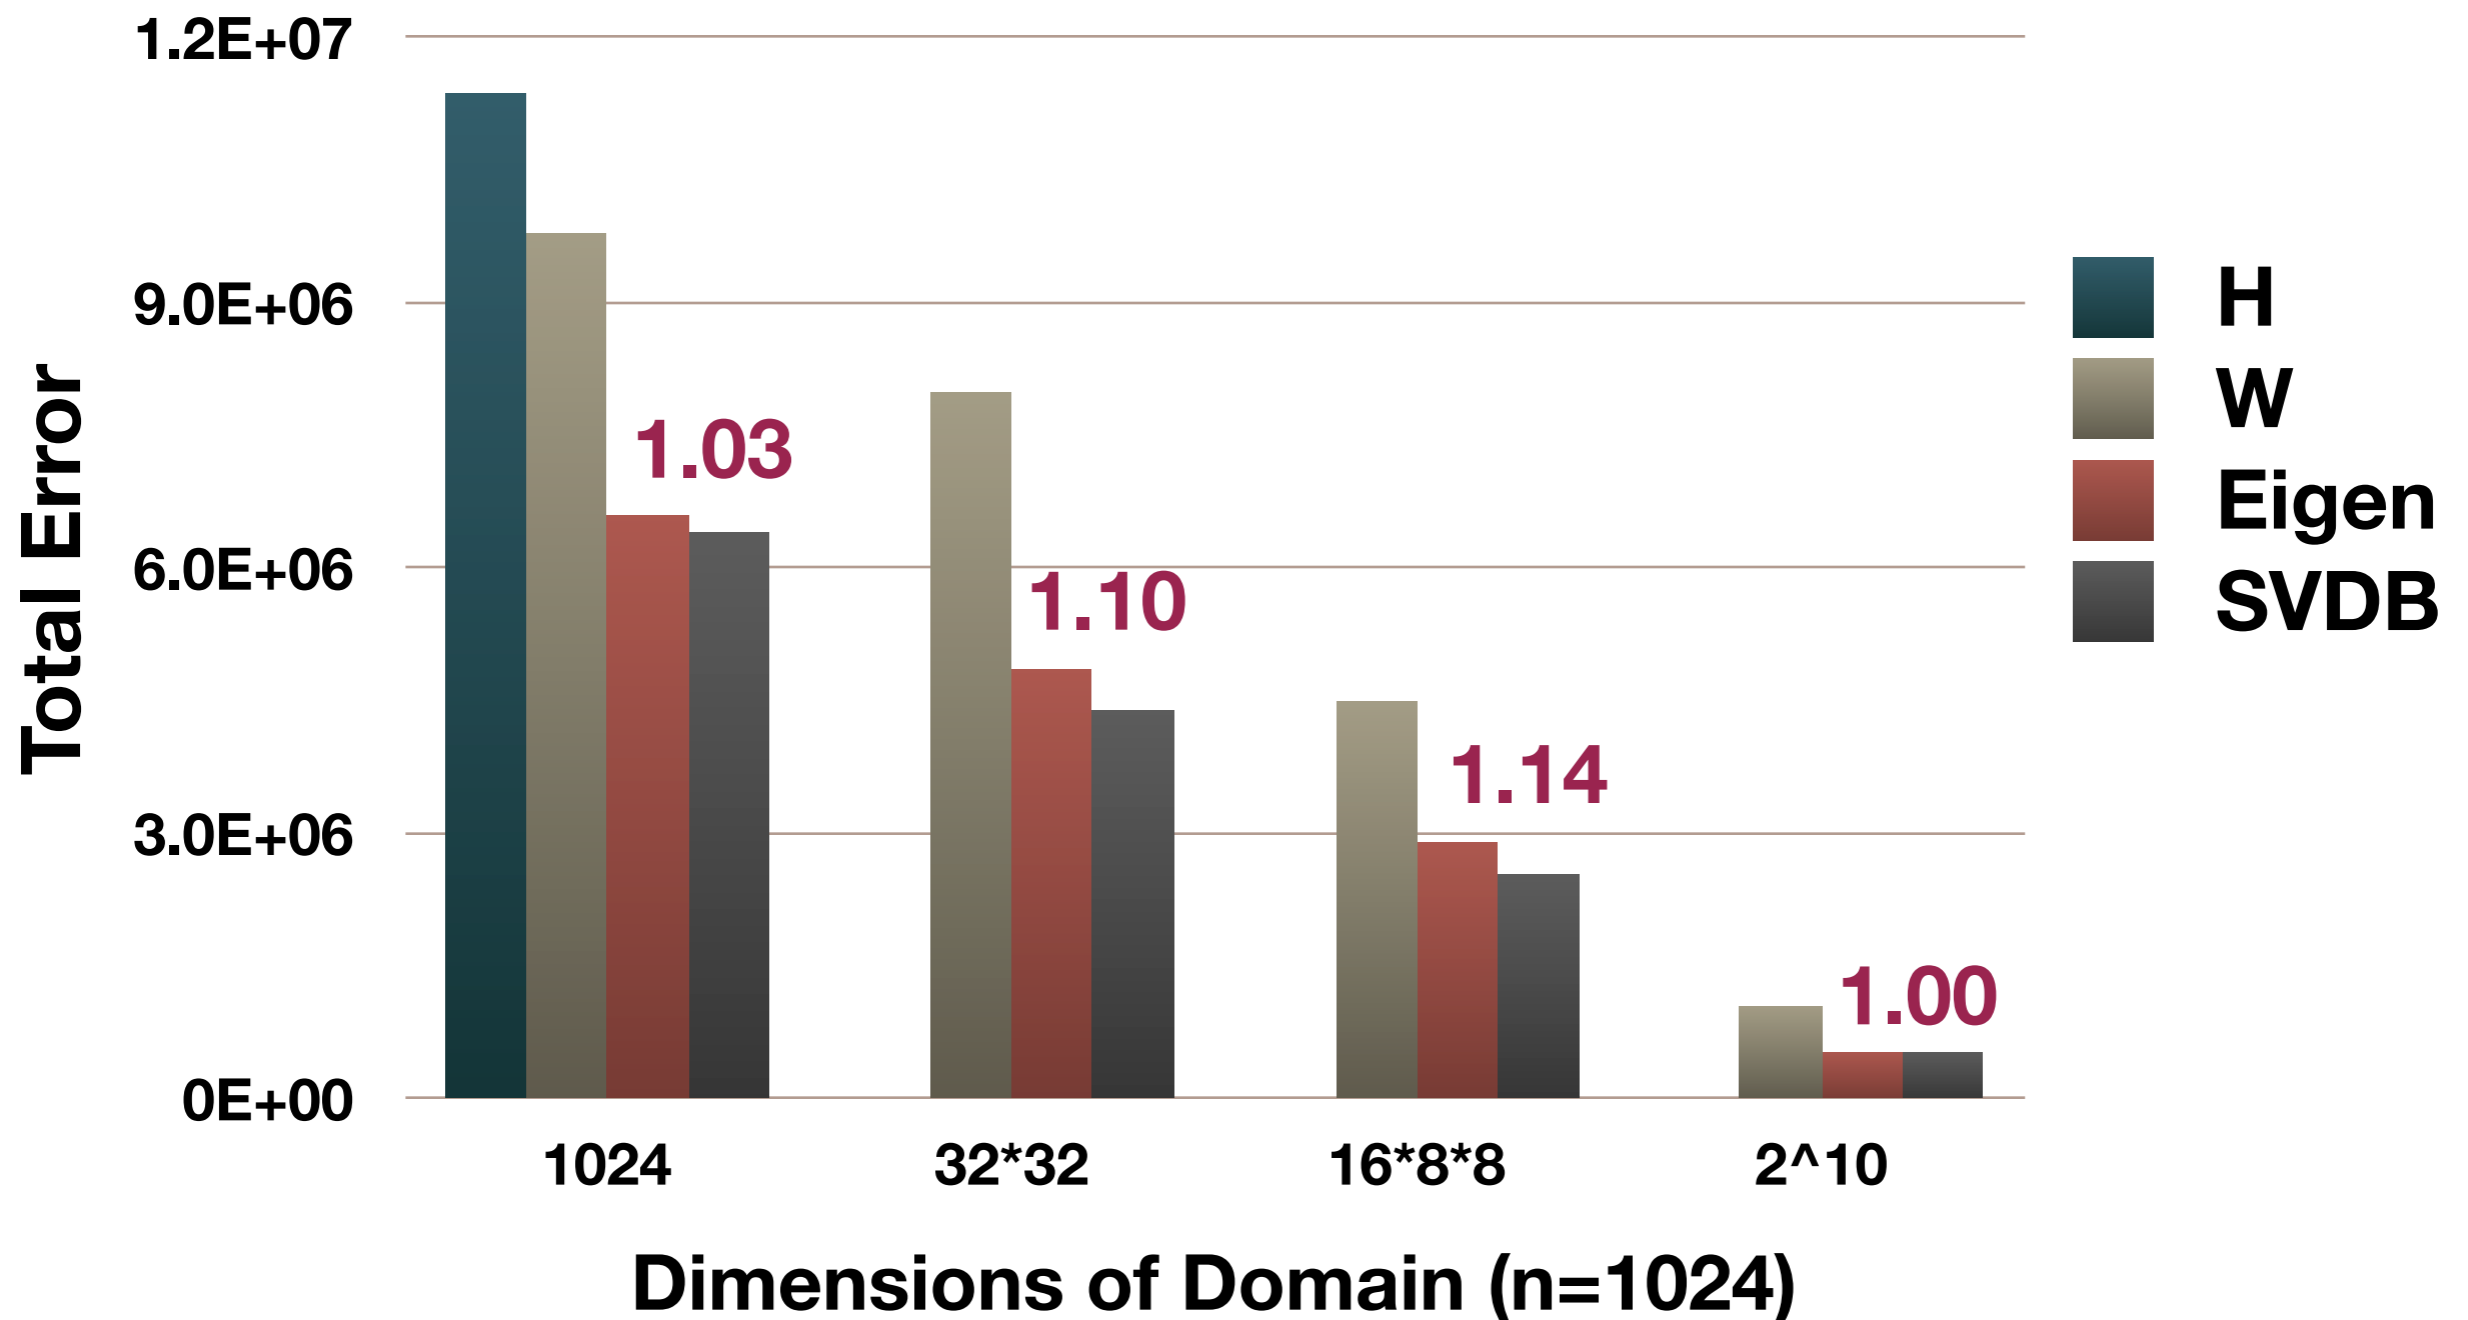
- Resulting strategy matrix is:

$$\mathbf{A} = \begin{bmatrix} c_1 \mathbf{v}_1 \\ c_2 \mathbf{v}_2 \\ \dots \\ c_n \mathbf{v}_n \end{bmatrix}$$

Approximately optimal error rates

All range queries

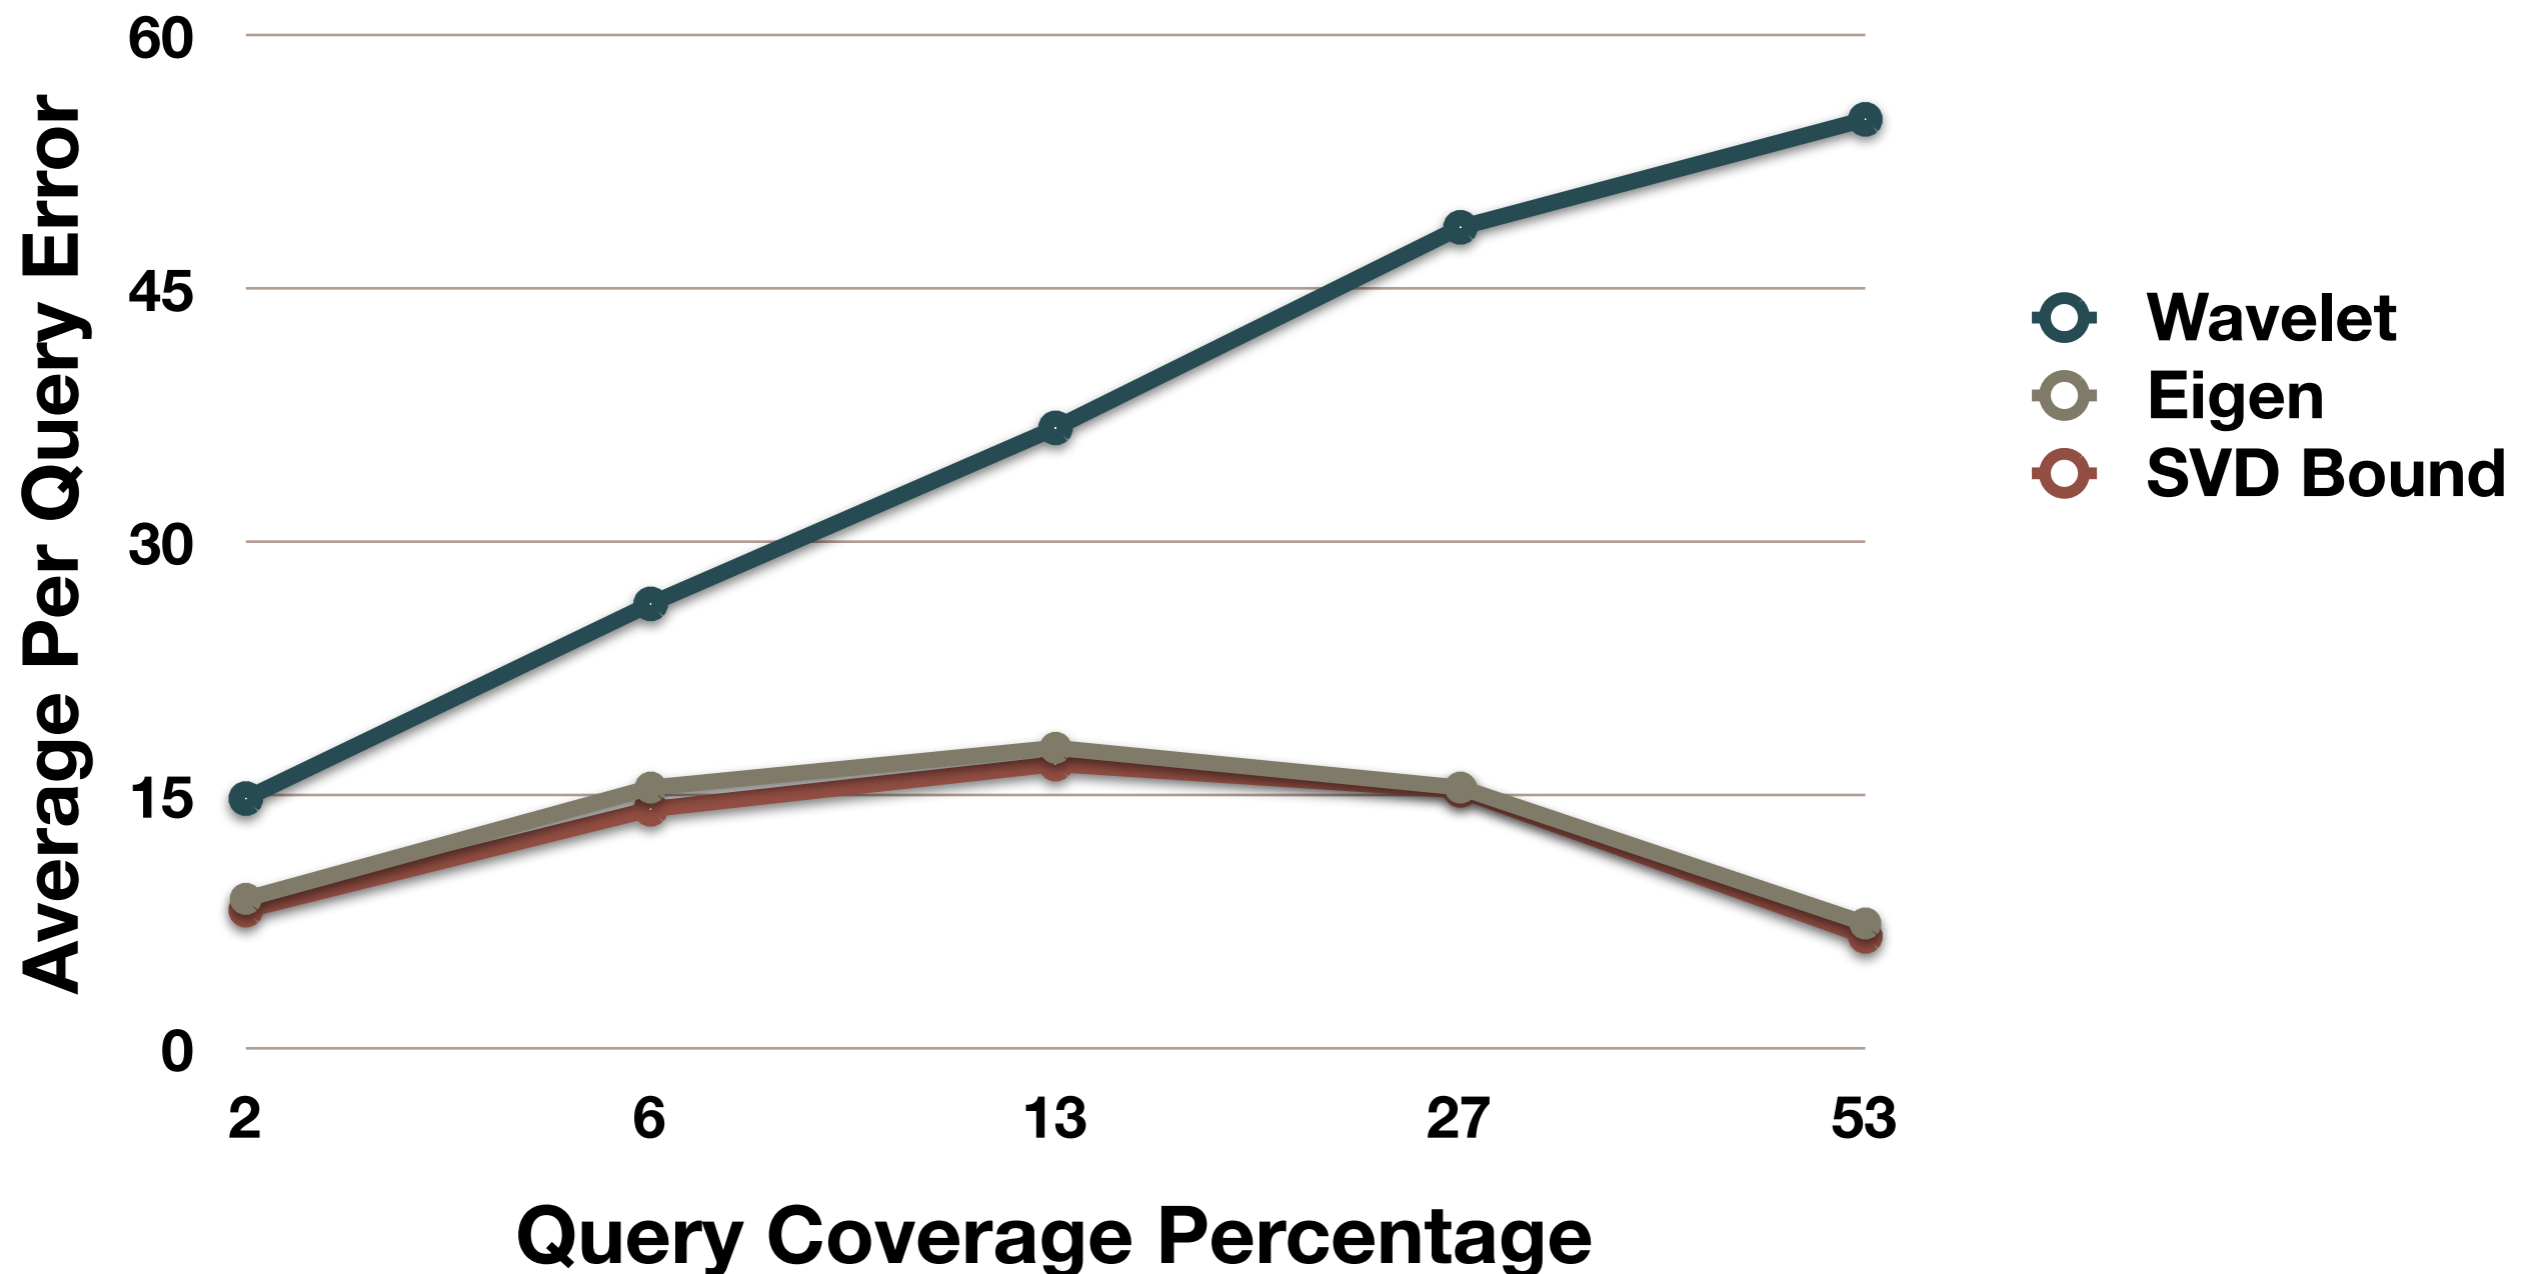
(ϵ, δ) -differential privacy



Customizing the strategy to the workload

Subsets of the range queries

(ϵ, δ) -differential privacy



Non-negative least squares

Non-negative least squares

- ① **(Design)** Choose a **full rank** query strategy \mathbf{A}
- ② **(Apply Laplace)** Use the Laplace mechanism to answer \mathbf{A}

$$\mathbf{z} = \mathbf{Ax} + (\|\mathbf{A}\|_1 / \varepsilon)\mathbf{b}$$

- ③ **(Inference)** Compute estimate $\underline{\mathbf{x}}$ of \mathbf{x} using answers \mathbf{z} .

Non-negative least squares

- 1 **(Design)** Choose a **full rank** query strategy \mathbf{A}
- 2 **(Apply Laplace)** Use the Laplace mechanism to answer \mathbf{A}

$$\mathbf{z} = \mathbf{Ax} + (\|\mathbf{A}\|_1 / \varepsilon)\mathbf{b}$$

- 3 **(Inference)** Compute estimate $\underline{\mathbf{x}}$ of \mathbf{x} using answers \mathbf{z} .
 - Non-negative least squares: compute estimate $\underline{\mathbf{x}}$ of \mathbf{x} that minimizes squared error: $\|\mathbf{Ax} - \mathbf{z}\|_2^2$

where each $\underline{x}_i > 0$

Non-negative least squares

- 1 **(Design)** Choose a **full rank** query strategy \mathbf{A}
- 2 **(Apply Laplace)** Use the Laplace mechanism to answer \mathbf{A}

$$\mathbf{z} = \mathbf{Ax} + (\|\mathbf{A}\|_1 / \varepsilon)\mathbf{b}$$

- 3 **(Inference)** Compute estimate $\underline{\mathbf{x}}$ of \mathbf{x} using answers \mathbf{z} .
 - Non-negative least squares: compute estimate $\underline{\mathbf{x}}$ of \mathbf{x} that minimizes squared error: $\|\mathbf{Ax} - \mathbf{z}\|_2^2$

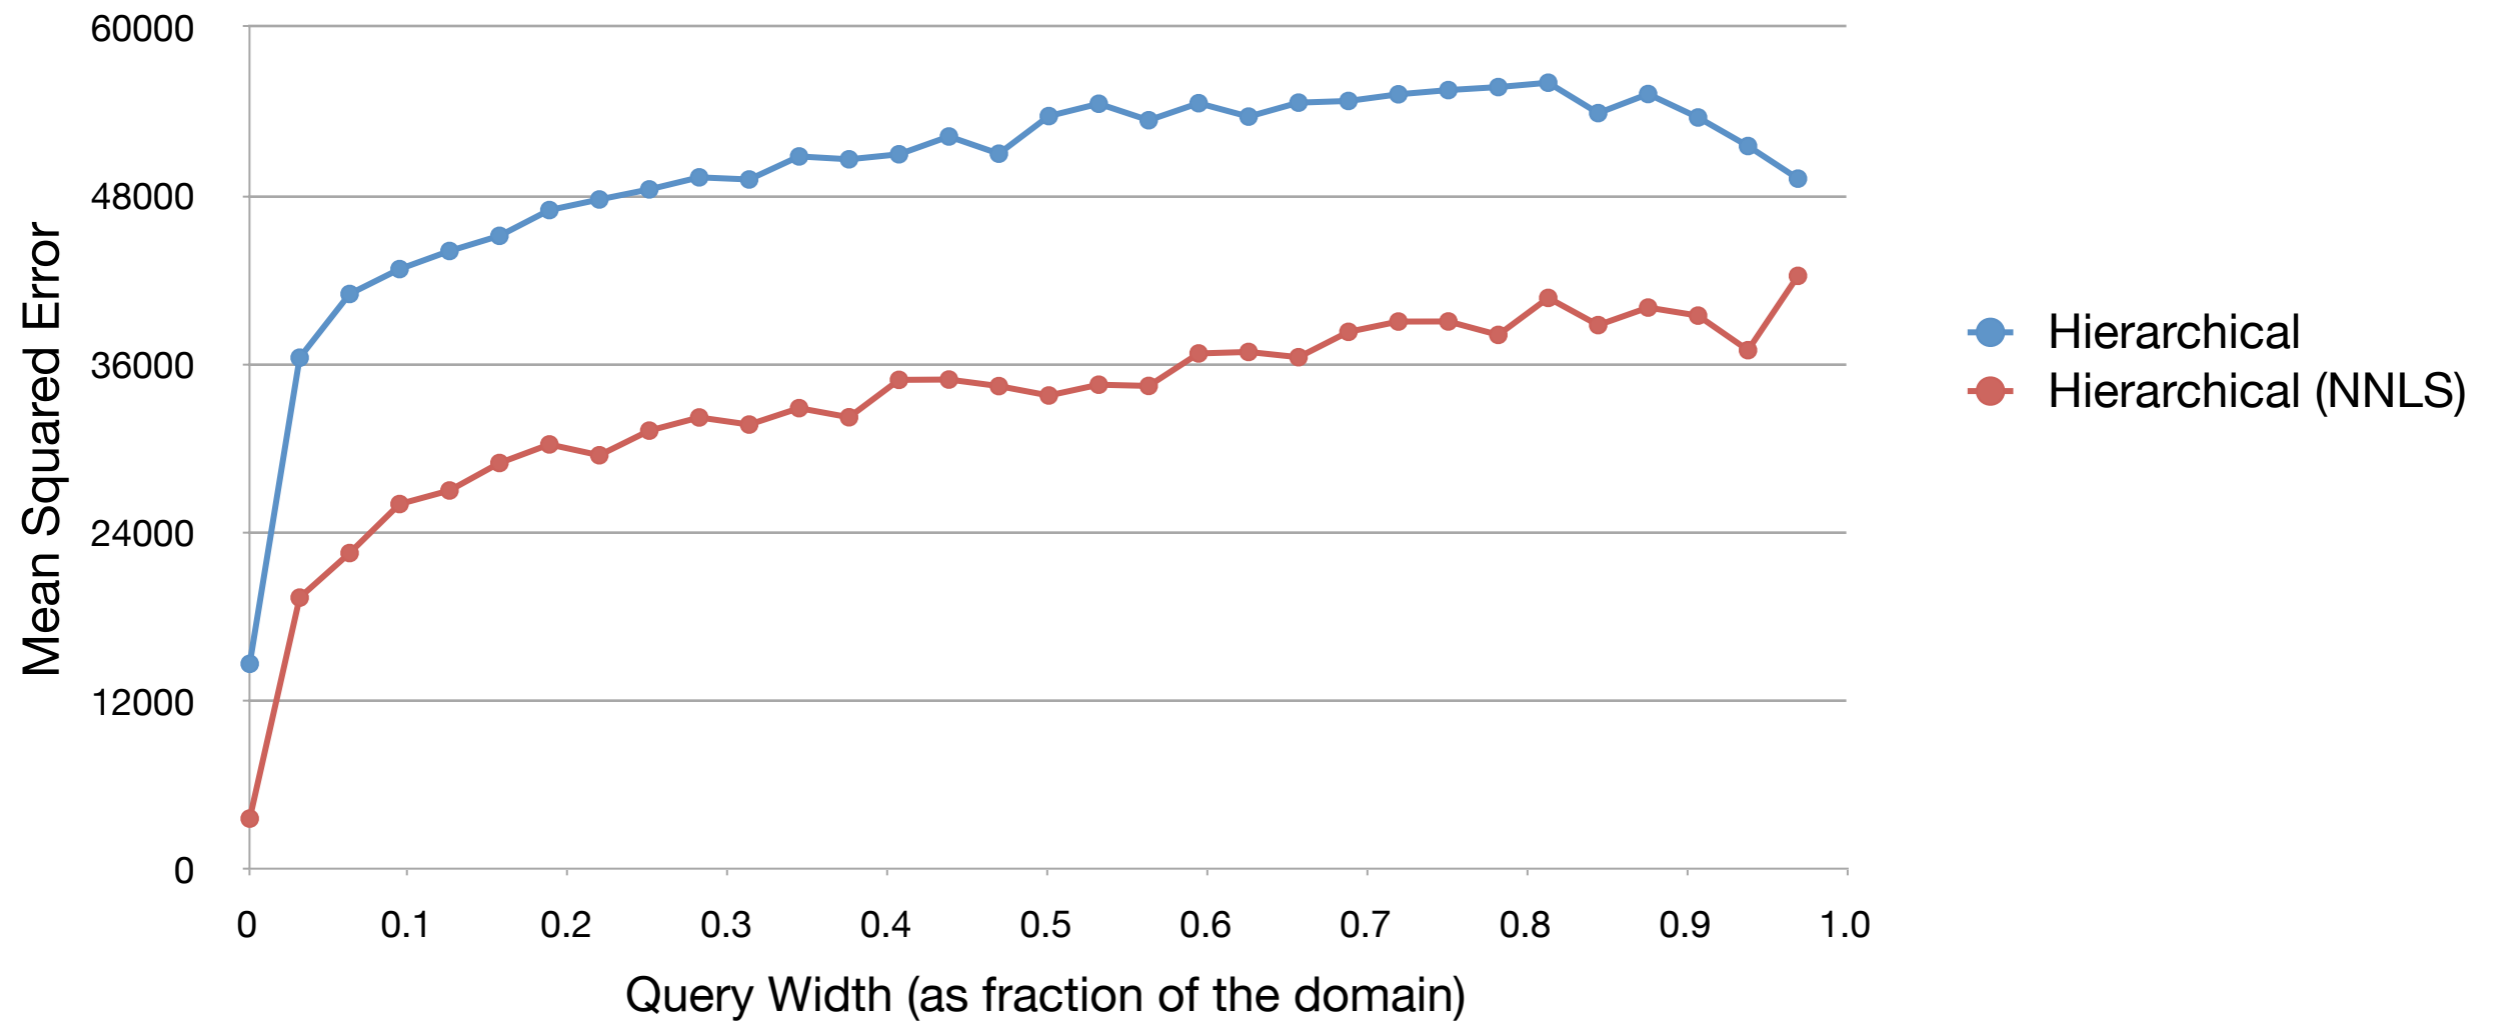
where each $\underline{x}_i > 0$

Effectiveness of non-negative least squares depends of properties of the data, epsilon, and \mathbf{A} .

Error: all range queries, non-negative least squares

Synthetic sparse dataset

Epsilon = 0.1
n = 1024



Outline

1. Case study: answering 1-dim range queries
2. The matrix mechanism -- formal description.
3. The matrix mechanism -- new tools & techniques.
4. Conclusion

Outline

1. Case study: answering 1-dim range queries
2. The matrix mechanism -- formal description.
3. The matrix mechanism -- new tools & techniques.
4. Conclusion

Summary and conclusions

- The Matrix mechanism generalizes the Laplace & Gaussian mechanisms, improving accuracy by exploiting correlation in the workload and reducing sensitivity.
- Two recent techniques for range queries are instances of the matrix mechanism; neither is optimal, but they are close.
- One strategy does not fit all workloads: adapting the strategy to the workload is essential to achieving low error.
- It is possible to compute the optimal strategy in $O(n^8)$ time, and approximately optimal strategies in $O(n^4)$.

Open questions

- The matrix mechanism is data-independent. What are the trade-offs for data-dependent approaches?
- What makes one workload “harder” to answer than another? How can we reliably measure workload error complexity?
- How do our results compare with lower bounds for differentially private output. (Our optimal strategies result in the least error for this particular mechanism, not necessarily the lowest error possible.)
- Can we avoid the computational dependence on the domain size n , without sacrificing accuracy?
- How do we analyze the error resulting from non-negative least squares?

Questions?

Project page and implementation of the Matrix Mechanism:

<http://bit.ly/ituyOt>

Additional details on our work may be found here:

- **[Li, ArXiv 2011]** C. Li and G. Miklau. Efficient Batch Query Answering Under Differential Privacy. CoRR abs/1103.1367, 2011.
- **[Li, PODS 2010]** C. Li, M. Hay, V. Rastogi, G. Miklau, and A. McGregor. Optimizing Linear Counting Queries Under Differential Privacy. Principles of Database Systems (PODS) 2010.
- **[Hay, PVLDB 10]** M. Hay, V. Rastogi, G. Miklau, and D. Suciu. Boosting the accuracy of differentially-private queries through consistency. Proceedings of the VLDB Endowment (PVLDB), 2010.

References

[Ghosh, 2009] A. Ghosh, T. Roughgarden, and M. Sundararajan. Universally utility-maximizing privacy mechanisms. In Symposium on Theory of computing (STOC), pages 351–360, 2009.

[Xiao, 2010] X. Xiao, G. Wang, and J. Gehrke. Differential privacy via wavelet transforms. In International Conference on Data Engineering, 2010.

[Barak, 2007] B. Barak, K. Chaudhuri, C. Dwork, S. Kale, F. McSherry, and K. Talwar. Privacy, accuracy, and consistency too: a holistic solution to contingency table release. In Principles of Database Systems (PODS) 2007.

[Xiao, 2011] Xiaokui Xiao, Gabriel Bender, Michael Hay, and Johannes Gehrke. iReduct: Differential privacy with reduced relative errors. SIGMOD, 2011.

[Lawson, 1987] C. L. Lawson and R. J. Hanson, Solving least squares Problems, Prentice Hall, 1987.