

CSE544

Data Management

Lecture 11

Query Optimization – Part 2

Announcements

- Project:
 - Feedback on proposals posted
 - Milestone due on Monday 2/26
- Wednesday 2/14: review 3 due
- Monday 2/19: no lecture (holiday)
- Wednesday 2/21: no lecture
- Friday 2/23: makeup lecture, HW3 due

Query Optimization

Three major components:

1. Search space last week
2. Cardinality and cost estimation today
3. Plan enumeration algorithms today+Wed.

Cardinality Estimation

Problem: given statistics on base tables and a query, estimate size of the answer

Very difficult, because:

- Need to do it very fast
- Need to use very little memory

Cardinality Estimation

- Traditional method:
 - Simple probability space on the set of possible tuples
- Alternate methods:
 - Based on ML
 - Based on sampling

Traditional Cardinality Estimation

Statistics on Base Data

- Number of tuples (cardinality) $T(R)$
- Number of physical pages $B(R)$
- Indexes, number of keys in the index $V(R,a)$

Statistics on Base Data

- Number of tuples (cardinality) $T(R)$
- Number of physical pages $B(R)$
- Indexes, number of keys in the index $V(R,a)$
- Histogram on single attribute (1d)
- Histogram on two attributes (2d)

Computed periodically, often using sampling

Probability Space

```
SELECT *                --- attrs A1, A2, ..., An
FROM R1, R2, ..., Rm
WHERE condition
```

Probability Space

```
SELECT *                --- attrs A1, A2, ..., An
FROM R1, R2, ..., Rm
WHERE condition
```

Probability space:

- Outcomes: $R_1 \times R_2 \times \dots \times R_m$
- Probability: $p(A_1, A_2, \dots, A_n)$

Cardinality estimate: $Est = p(\text{condition}) \cdot T(R_1) \cdots T(R_m)$

Probability Space

```
SELECT *           --- attrs A1, A2, ..., An
FROM R1, R2, ..., Rm
WHERE condition
```

Probability space:

- Outcomes: $R_1 \times R_2 \times \dots \times R_m$
- Probability: $p(A_1, A_2, \dots, A_n)$

Cardinality estimate: $Est = p(\text{condition}) \cdot T(R_1) \dots T(R_m)$

Goal: compute $p(\text{condition})$

Assumptions

- Uniformity
- Independence
- Containment of values
- Preservation of values

Selectivity Factors

$p(pred)$ is called selectivity factor θ

$$\text{Est}(\sigma_{pred}(R)) = \theta_{pred} * T(R)$$

$$\text{Est}(R \bowtie_{A=B} S) = \theta_{A=B} * T(R) * T(S)$$

Selectivity Factors

Uniformity assumption

Equality:

- $\theta_{A=c} = 1/V(R,A)$

$$\sigma_{A=c}(R)$$

Selectivity Factors

Uniformity assumption

Equality:

- $\theta_{A=c} = 1/V(R,A)$

$$\sigma_{A=c}(R)$$

Range:

- $\theta_{c1 < A < c2} = (c2 - c1) / (\max(R,A) - \min(R,A))$

$$\sigma_{c1 < A < c2}(R)$$

Selectivity Factors

Uniformity assumption

Equality:

- $\theta_{A=c} = 1/V(R,A)$

$$\sigma_{A=c}(R)$$

Range:

- $\theta_{c1 < A < c2} = (c2 - c1) / (\max(R,A) - \min(R,A))$

$$\sigma_{c1 < A < c2}(R)$$

Independence assumption

- $\theta_{\text{pred1 and pred2}} = \theta_{\text{pred1}} * \theta_{\text{pred2}} = 1/V(R,A) * 1/V(R,B)$

$$\sigma_{A=c \text{ and } B=d}(R)$$

Supplier(sid, sname, scity, sstate)

T(Supplier) = 100000
V(Supplier, scity) = 2000
V(Supplier, sstate) = 50

Example

```
SELECT *  
FROM Supplier  
WHERE scity = 'Seattle'
```

Est = ????

Supplier(sid, sname, scity, sstate)

T(Supplier) = 100000
V(Supplier, scity) = 2000
V(Supplier, sstate) = 50

Example

```
SELECT *  
FROM Supplier  
WHERE scity = 'Seattle'
```

$$\text{Est} = 1/2000 * 100000 = 50$$

Supplier(sid, sname, scity, sstate)

T(Supplier) = 100000
V(Supplier, scity) = 2000
V(Supplier, sstate) = 50

Example

```
SELECT *  
FROM Supplier  
WHERE scity = 'Seattle'
```

$$\text{Est} = 1/2000 * 100000 = 50$$

```
SELECT *  
FROM Supplier  
WHERE sstate = 'WA'
```

$$\text{Est} = 1/50 * 100000 = 2000$$

Supplier(sid, sname, scity, sstate)

T(Supplier) = 100000
V(Supplier, scity) = 2000
V(Supplier, sstate) = 50

Example

```
SELECT *  
FROM Supplier  
WHERE scity = 'Seattle'
```

$$\text{Est} = 1/2000 * 100000 = 50$$

```
SELECT *  
FROM Supplier  
WHERE sstate = 'WA'
```

$$\text{Est} = 1/50 * 100000 = 2000$$

```
SELECT *  
FROM Supplier  
WHERE scity = 'Seattle'  
and sstate = 'WA'
```

$$\text{Est} = \text{????}$$

Supplier(sid, sname, scity, sstate)

$T(\text{Supplier}) = 100000$
 $V(\text{Supplier}, \text{scity}) = 2000$
 $V(\text{Supplier}, \text{sstate}) = 50$

Example

```
SELECT *  
FROM Supplier  
WHERE scity = 'Seattle'
```

$$\text{Est} = 1/2000 * 100000 = 50$$

```
SELECT *  
FROM Supplier  
WHERE sstate = 'WA'
```

$$\text{Est} = 1/50 * 100000 = 2000$$

```
SELECT *  
FROM Supplier  
WHERE scity = 'Seattle'  
and sstate = 'WA'
```

$$\text{Est} = 1/2000 * 1/50 * 100000 = 1$$

Independence

Supplier(sid, sname, scity, sstate)

T(Supplier) = 100000
V(Supplier, scity) = 2000
V(Supplier, sstate) = 50

Example

```
SELECT *
FROM Supplier
WHERE scity = 'Seattle'
```

$$\text{Est} = 1/2000 * 100000 = 50$$

```
SELECT *
FROM Supplier
WHERE sstate = 'WA'
```

$$\text{Est} = 1/50 * 100000 = 2000$$

Independence

```
SELECT *
FROM Supplier
WHERE scity = 'Seattle'
and sstate = 'WA'
```

$$\text{Est} = 1/2000 * 1/50 * 100000 = 1$$

Very wrong
Why?

Selectivity Factors

Join

- $\theta_{R.A=S.B} = 1 / (\text{MAX}(V(R,A), V(S,B)))$

Why? Will explain next...

Selectivity Factors

Containment of values: if $V(R,A) \subseteq V(S,B)$, then the set of A values of R is included in the set of B values of S

- Note: this indeed holds when A is a foreign key in R, and B is a key in S

Selectivity Factors

Assume $V(R,A) \leq V(S,B)$

- Tuple t in R joins with $T(S)/V(S,B)$ tuples in S
- Hence $\text{Est}(R \bowtie_{A=B} S) = T(R) T(S) / V(S,B)$

Selectivity Factors

Assume $V(R,A) \leq V(S,B)$

- Tuple t in R joins with $T(S)/V(S,B)$ tuples in S
- Hence $\text{Est}(R \bowtie_{A=B} S) = T(R) T(S) / V(S,B)$

In general:

- $\text{Est}(R \bowtie_{A=B} S) = T(R) T(S) / \max(V(R,A), V(S,B))$
- $\theta_{R.A=S.B} = 1 / (\max(V(R,A), V(S,B)))$

Supplier(sid, sname, scity, sstate)
Supply(sid, pno, quantity)

T(Supplier) = 100000
V(Supplier, sid) = 100000

Example

```
SELECT *  
FROM Supplier x, Supply y  
WHERE x.sid = y.sid
```

T(Supply) = 5000000
V(Supply, sid) = 80000

What is the **exact** output size?

Supplier(sid, sname, scity, sstate)
Supply(sid, pno, quantity)

T(Supplier) = 100000
V(Supplier, sid) = 100000

Example

T(Supply) = 5000000
V(Supply, sid) = 80000

```
SELECT *  
FROM Supplier x, Supply y  
WHERE x.sid = y.sid
```

What is the **exact** output size?

T(Supply)=5000000

Supplier(sid, sname, scity, sstate)
Supply(sid, pno, quantity)

T(Supplier) = 100000
V(Supplier, sid) = 100000

Example

T(Supply) = 5000000
V(Supply, sid) = 80000

```
SELECT *  
FROM Supplier x, Supply y  
WHERE x.sid = y.sid
```

What is the **exact** output size?

T(Supply)=5000000

$$Est(Supplier \bowtie Supply) = \frac{T(Supplier)T(Supply)}{\max(V(Supplier, sid), V(Supply, sid))}$$

Supplier(sid, sname, scity, sstate)
Supply(sid, pno, quantity)

T(Supplier) = 100000
V(Supplier, sid) = 100000

Example

T(Supply) = 5000000
V(Supply, sid) = 80000

```
SELECT *  
FROM Supplier x, Supply y  
WHERE x.sid = y.sid
```

What is the **exact** output size?

T(Supply)=5000000

$$Est(Supplier \bowtie Supply) = \frac{T(Supplier)T(Supply)}{\max(V(Supplier, sid), V(Supply, sid))}$$

Always larger
Why?

Supplier(sid, sname, scity, sstate)
Supply(sid, pno, quantity)

T(Supplier) = 100000
V(Supplier, sid) = 100000

Example

T(Supply) = 5000000
V(Supply, sid) = 80000

```
SELECT *  
FROM Supplier x, Supply y  
WHERE x.sid = y.sid
```

What is the **exact** output size?

T(Supply)=5000000

$$Est(Supplier \bowtie Supply) = \frac{T(Supplier)T(Supply)}{\max(V(Supplier, sid), V(Supply, sid))}$$

$$= \frac{T(Supplier)T(Supply)}{V(Supplier, sid)}$$

Always larger
Why?

Supplier(sid, sname, scity, sstate)
Supply(sid, pno, quantity)

T(Supplier) = 100000
V(Supplier, sid) = 100000

Example

T(Supply) = 5000000
V(Supply, sid) = 80000

```
SELECT *  
FROM Supplier x, Supply y  
WHERE x.sid = y.sid
```

What is the **exact** output size?

T(Supply)=5000000

$$Est(Supplier \bowtie Supply) = \frac{T(Supplier)T(Supply)}{\max(V(Supplier, sid), V(Supply, sid))}$$

$$= \frac{T(Supplier)T(Supply)}{V(Supplier, sid)} = T(Supply)$$

Always larger
Why?

Final Assumption

Preservation of values:

For any other attribute C:

- $V(R \bowtie_{A=B} S, C) = V(R, C)$ or
- $V(R \bowtie_{A=B} S, C) = V(S, C)$
- This is needed higher up in the plan


Computing the Cost of a Plan

- Estimate **cardinalities** bottom-up
- Estimate **cost** by using estimated cardinalities
- Example next...

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Logical Query Plan 1



```

SELECT sname
FROM Supplier x, Supply y
WHERE x.sid = y.sid
      and y.pno = 2
      and x.scity = 'Seattle'
      and x.sstate = 'WA'

```

T(Supply) = 10000
 B(Supply) = 100
 V(Supply, pno) = 2500

T(Supplier) = 1000
 B(Supplier) = 100
 V(Supplier, scity) = 20
 V(Supplier, state) = 10

M=11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)


Logical Query Plan 1

Estimated
(why?)

$\sigma_{pno=2 \wedge scity='Seattle' \wedge sstate='WA'}$

T = 10000

π_{sname}



sid = sid

Supply

Supplier

```
SELECT sname
FROM Supplier x, Supply y
WHERE x.sid = y.sid
and y.pno = 2
and x.scity = 'Seattle'
and x.sstate = 'WA'
```

T(Supply) = 10000
B(Supply) = 100
V(Supply, pno) = 2500

T(Supplier) = 1000
B(Supplier) = 100
V(Supplier, scity) = 20
V(Supplier, state) = 10

M=11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Estimated
(why?)


Logical Query Plan 1

T < 1

Π_{sname}

$\sigma_{pno=2 \wedge scity='Seattle' \wedge sstate='WA'}$

T = 10000



sid = sid

Supply

Supplier

```
SELECT sname
FROM Supplier x, Supply y
WHERE x.sid = y.sid
      and y.pno = 2
      and x.scity = 'Seattle'
      and x.sstate = 'WA'
```

T(Supply) = 10000
B(Supply) = 100
V(Supply, pno) = 2500


T(Supplier) = 1000
B(Supplier) = 100
V(Supplier, scity) = 20
V(Supplier, state) = 10

M=11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Logical Query Plan 2



```

SELECT sname
FROM Supplier x, Supply y
WHERE x.sid = y.sid
      and y.pno = 2
      and x.scity = 'Seattle'
      and x.sstate = 'WA'

```

T(Supply) = 10000
 B(Supply) = 100
 V(Supply, pno) = 2500

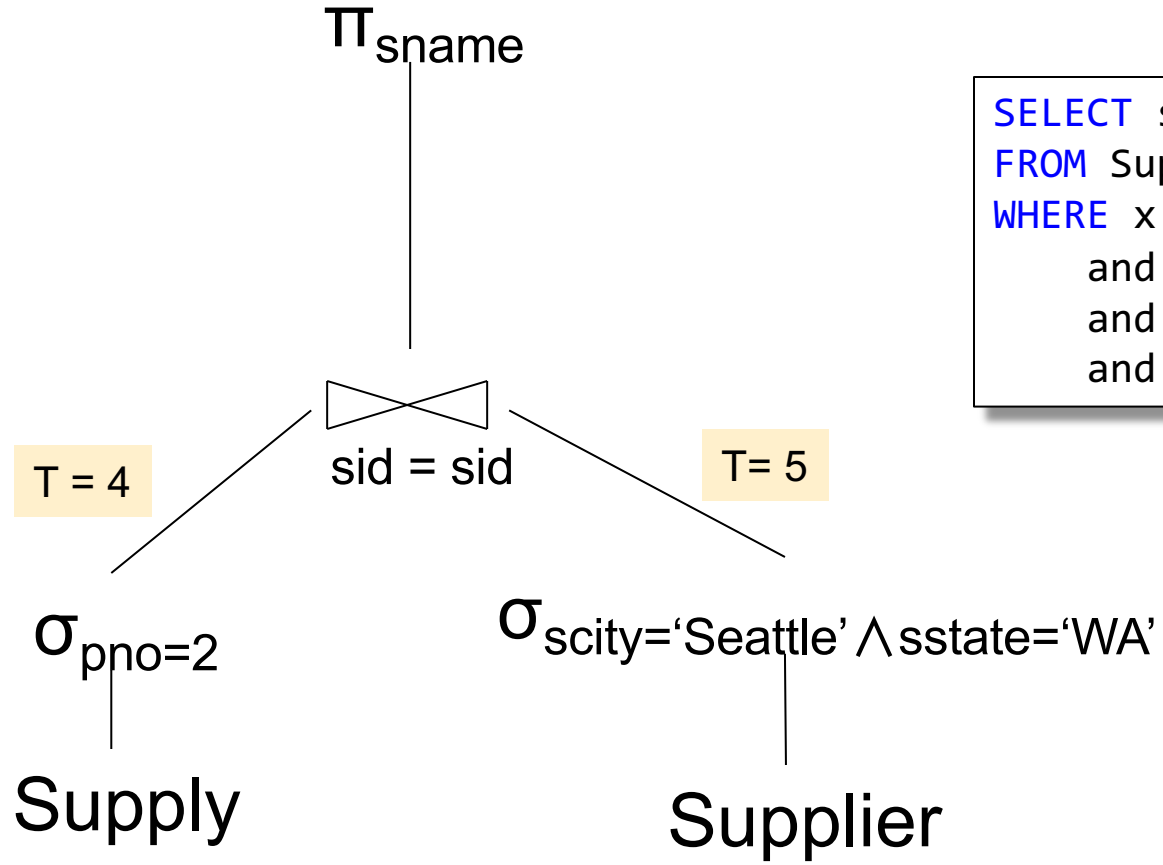
T(Supplier) = 1000
 B(Supplier) = 100
 V(Supplier, scity) = 20
 V(Supplier, state) = 10

M=11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Logical Query Plan 2



```

SELECT sname
FROM Supplier x, Supply y
WHERE x.sid = y.sid
      and y.pno = 2
      and x.scity = 'Seattle'
      and x.sstate = 'WA'

```

$T(\text{Supply}) = 10000$
 $B(\text{Supply}) = 100$
 $V(\text{Supply}, pno) = 2500$

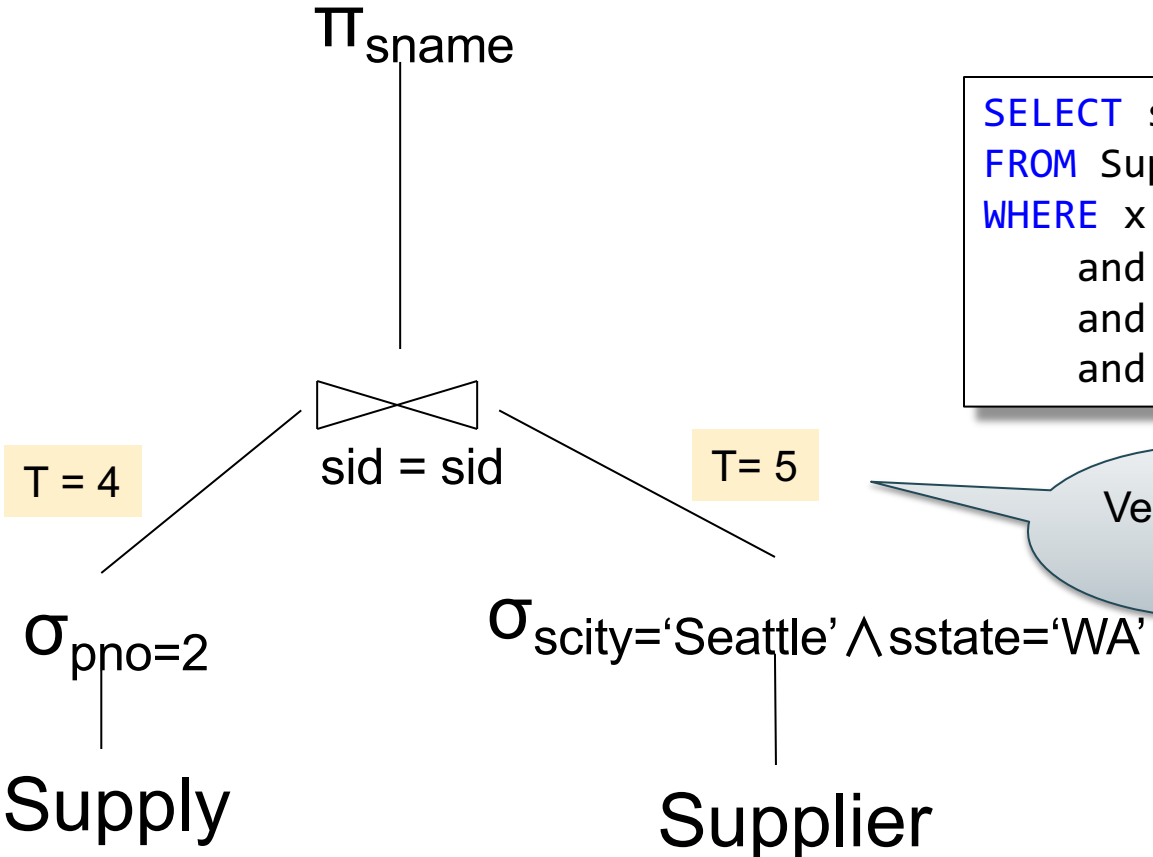
$T(\text{Supplier}) = 1000$
 $B(\text{Supplier}) = 100$
 $V(\text{Supplier}, scity) = 20$
 $V(\text{Supplier}, state) = 10$

M=11

Supplier(sid, sname, scity, sstate)
Supply(sid, pno, quantity)

Logical Query Plan 2

```
SELECT sname
FROM Supplier x, Supply y
WHERE x.sid = y.sid
      and y.pno = 2
      and x.scity = 'Seattle'
      and x.sstate = 'WA'
```



Very wrong!
Why?

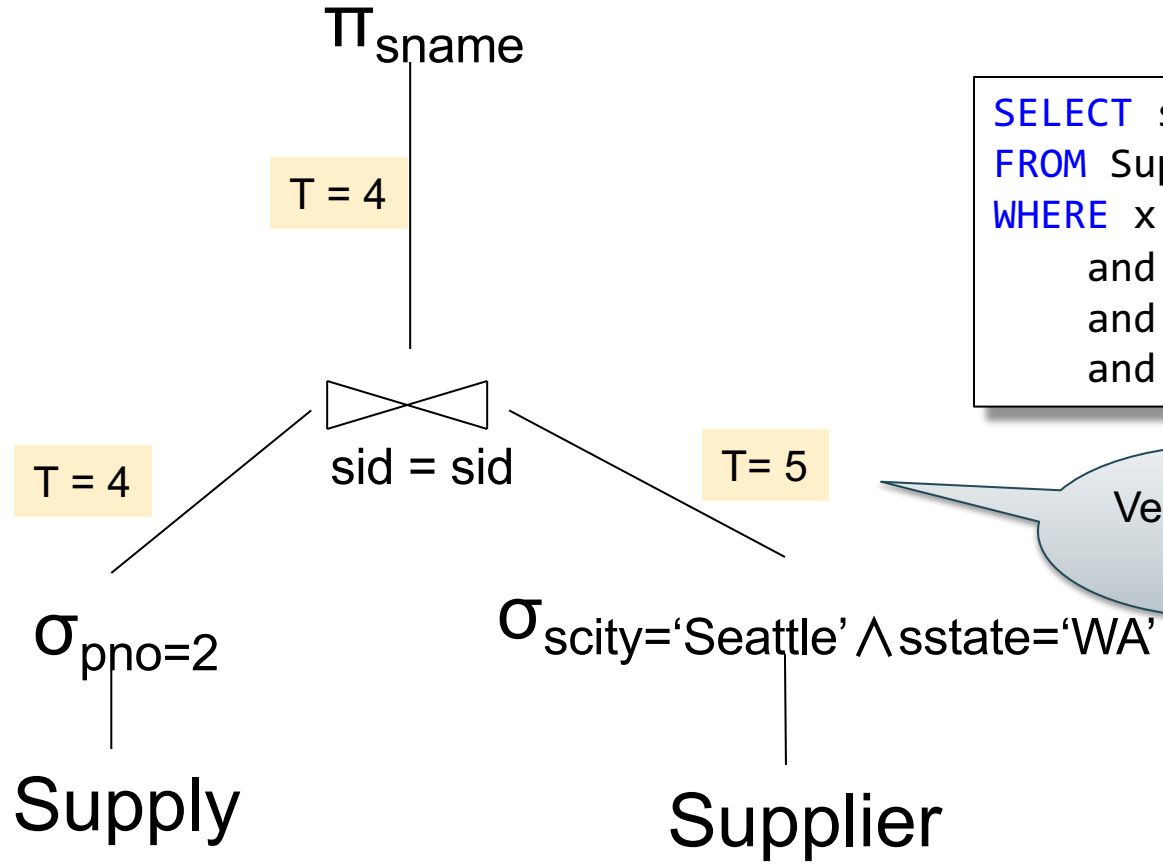
T(Supply) = 10000
B(Supply) = 100
V(Supply, pno) = 2500

T(Supplier) = 1000
B(Supplier) = 100
V(Supplier, scity) = 20
V(Supplier, state) = 10

M=11

Supplier(sid, sname, scity, sstate)
Supply(sid, pno, quantity)

Logical Query Plan 2



```

SELECT sname
FROM Supplier x, Supply y
WHERE x.sid = y.sid
      and y.pno = 2
      and x.scity = 'Seattle'
      and x.sstate = 'WA'
  
```

Very wrong!
Why?

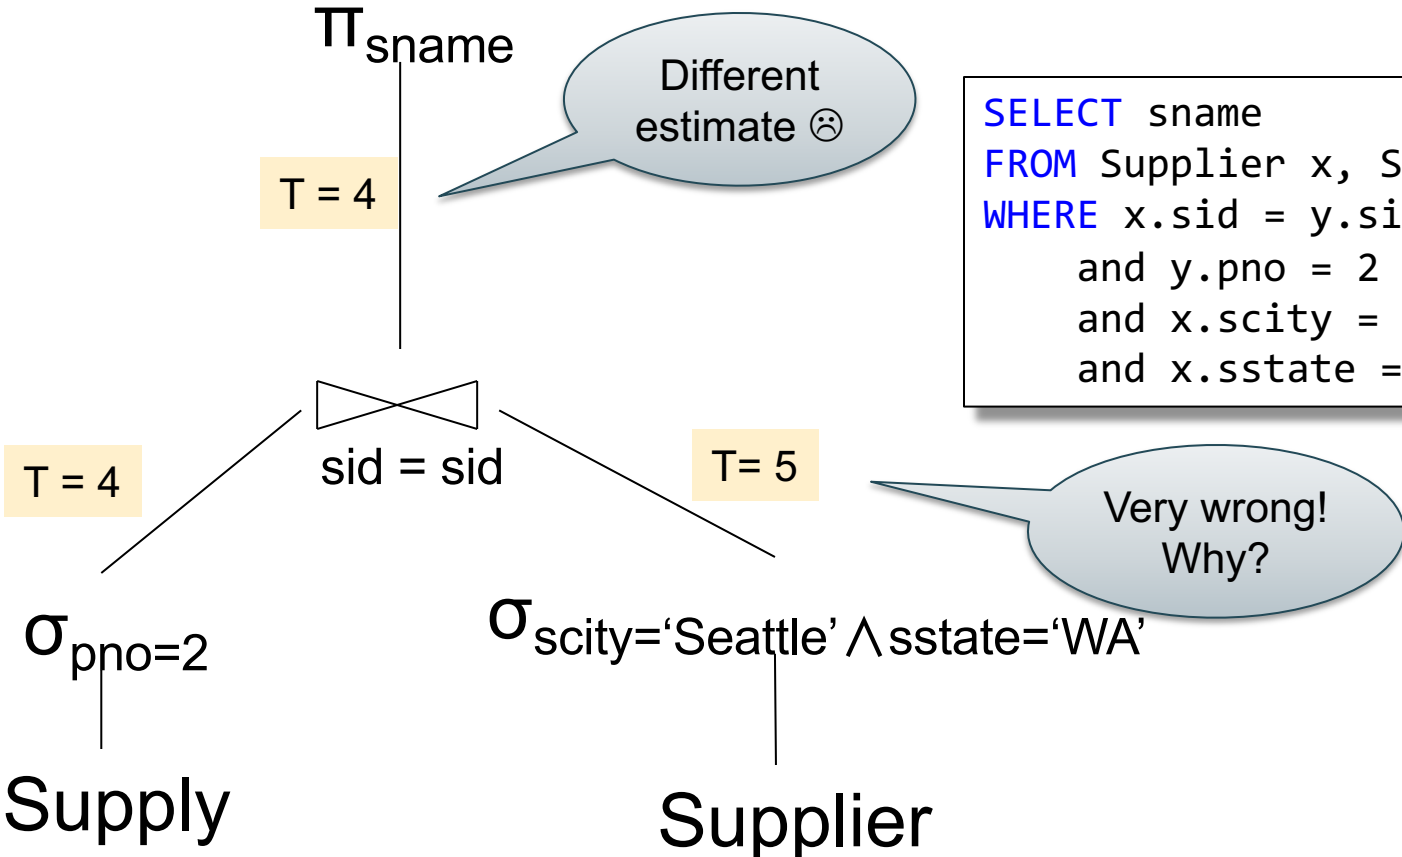
T(Supply) = 10000
 B(Supply) = 100
 V(Supply, pno) = 2500

T(Supplier) = 1000
 B(Supplier) = 100
 V(Supplier, scity) = 20
 V(Supplier, state) = 10

M=11

Supplier(sid, sname, scity, sstate)
 Supply(sid, pno, quantity)

Logical Query Plan 2



```

SELECT sname
FROM Supplier x, Supply y
WHERE x.sid = y.sid
      and y.pno = 2
      and x.scity = 'Seattle'
      and x.sstate = 'WA'
  
```

T(Supply) = 10000
 B(Supply) = 100
 V(Supply, pno) = 2500

T(Supplier) = 1000
 B(Supplier) = 100
 V(Supplier, scity) = 20
 V(Supplier, state) = 10

M=11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Physical Plan 1

Π_{sname}


T < 1

$\sigma_{pno=2 \wedge scity='Seattle' \wedge sstate='WA'}$

T = 10000

Total cost:

$B(\text{Supply}) + B(\text{Supply})B(\text{Supplier}) / (M - 1)$



sid = sid

Block nested loop join

Scan

Supply

Scan

Supplier

T(Supply) = 10000
B(Supply) = 100
V(Supply, pno) = 2500

T(Supplier) = 1000
B(Supplier) = 100
V(Supplier, scity) = 20
V(Supplier, state) = 10

M = 11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Physical Plan 1

Π_{sname}


T < 1

$\sigma_{pno=2 \wedge scity='Seattle' \wedge sstate='WA'}$

T = 10000

Total cost:

$$B(\text{Supply}) + \frac{B(\text{Supply})B(\text{Supplier})}{(M-1)} = 100 + \frac{100 \cdot 100}{10} = 1100$$



sid = sid

Block nested loop join

Scan

Supply

Scan

Supplier

T(Supply) = 10000
B(Supply) = 100
V(Supply, pno) = 2500

T(Supplier) = 1000
B(Supplier) = 100
V(Supplier, scity) = 20
V(Supplier, state) = 10

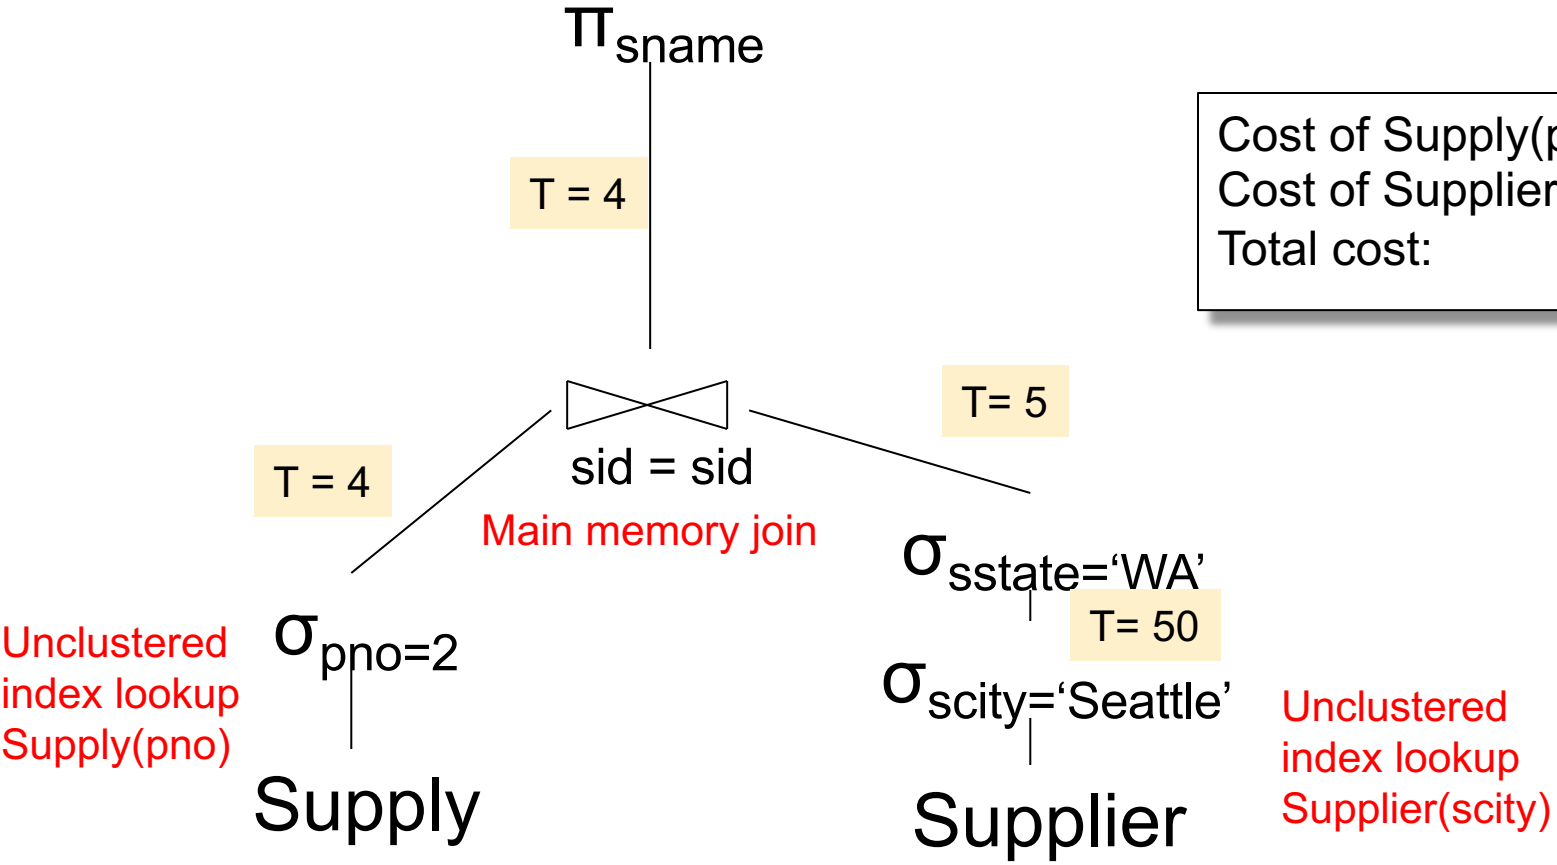
M=11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Physical Plan 2

Cost of Supply(pno) =
Cost of Supplier(scity) =
Total cost:



T(Supply) = 10000
B(Supply) = 100
V(Supply, pno) = 2500

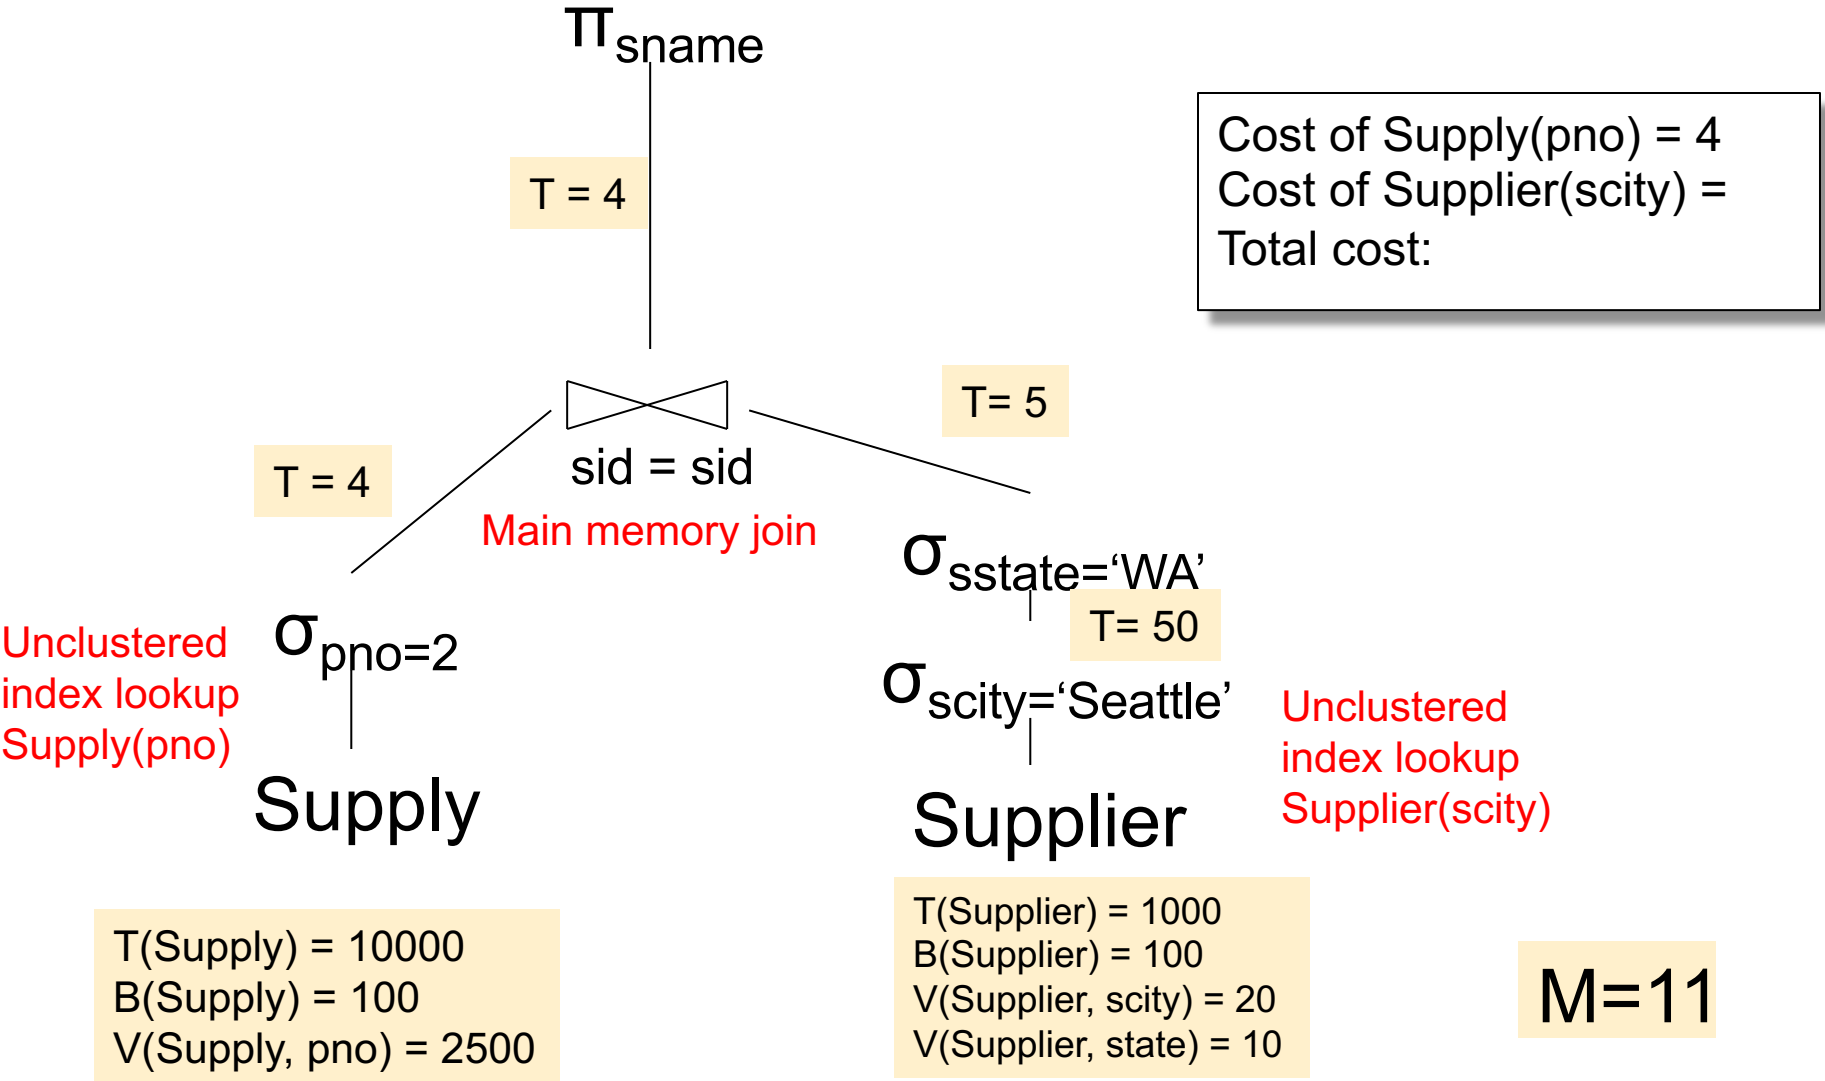
T(Supplier) = 1000
B(Supplier) = 100
V(Supplier, scity) = 20
V(Supplier, state) = 10

M=11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

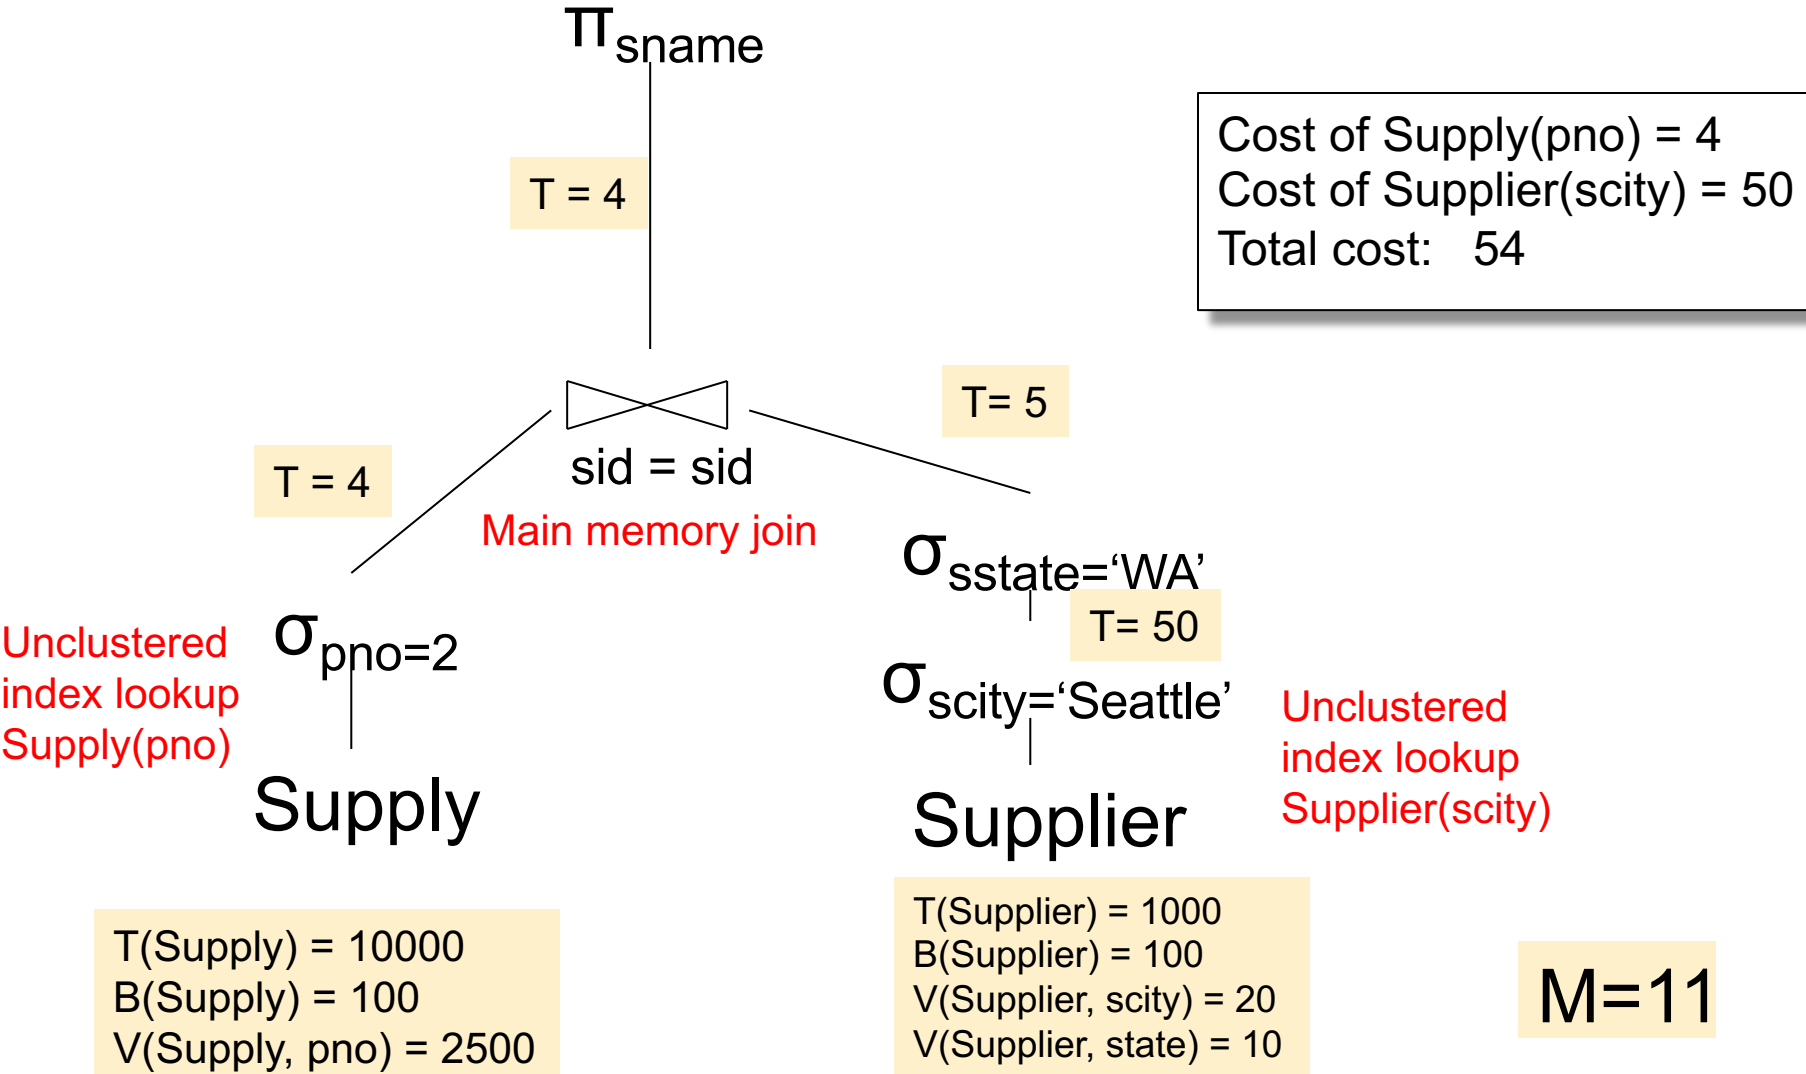
Physical Plan 2



Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Physical Plan 2



Cost of Supply(pno) = 4
 Cost of Supplier(scity) = 50
 Total cost: 54

T = 4

T = 4

T = 5

T = 50

Unclustered index lookup
Supplier(scity)

$T(\text{Supply}) = 10000$
 $B(\text{Supply}) = 100$
 $V(\text{Supply}, pno) = 2500$

$T(\text{Supplier}) = 1000$
 $B(\text{Supplier}) = 100$
 $V(\text{Supplier}, scity) = 20$
 $V(\text{Supplier}, state) = 10$

M = 11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Physical Plan 3


T = 4

Π_{sname}

$\sigma_{scity='Seattle' \wedge sstate='WA'}$

Cost of Supply(pno) =
Cost of Index join =
Total cost:

T = 4



sid = sid

Clustered
Index join

Unclustered
index lookup
Supply(pno)

$\sigma_{pno=2}$

Supply

Supplier

T(Supply) = 10000
B(Supply) = 100
V(Supply, pno) = 2500

T(Supplier) = 1000
B(Supplier) = 100
V(Supplier, scity) = 20
V(Supplier, state) = 10

M=11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Physical Plan 3


T = 4

Π_{sname}

$\sigma_{scity='Seattle' \wedge sstate='WA'}$

Cost of Supply(pno) = 4
Cost of Index join =
Total cost:

T = 4



sid = sid

Clustered
Index join

Unclustered
index lookup
Supply(pno)

$\sigma_{pno=2}$

Supply

Supplier

T(Supply) = 10000
B(Supply) = 100
V(Supply, pno) = 2500

T(Supplier) = 1000
B(Supplier) = 100
V(Supplier, scity) = 20
V(Supplier, state) = 10

M=11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Physical Plan 3


T = 4

Π_{sname}

$\sigma_{scity='Seattle' \wedge sstate='WA'}$

Cost of Supply(pno) = 4
Cost of Index join = 4
Total cost: 8

T = 4



sid = sid

Clustered
Index join

Unclustered
index lookup
Supply(pno)

$\sigma_{pno=2}$

Supply

Supplier

T(Supply) = 10000
B(Supply) = 100
V(Supply, pno) = 2500

T(Supplier) = 1000
B(Supplier) = 100
V(Supplier, scity) = 20
V(Supplier, state) = 10

M=11

Discussion

- We considered only IO cost; systems also need to estimate the CPU cost
- We assumed that all index pages were in memory: sometimes we need to add the cost of fetching index pages from disk

Histograms

- 1d Histograms mitigate uniformity assumption
- 2d Histograms mitigate independence assumption

1d-Histograms

- Histogram on R.A refines $T(R)$, $V(R,A)$
- Each bucket: $T(\text{bucket})$, $V(\text{bucket}, A)$

1d-Histograms

Employee(ssn, name, age)

$T(\text{Employee}) = 25000$, $V(\text{Employee}, \text{age}) = 50$

$\sigma_{\text{age}=48}(\text{Employee}) = ?$

1d-Histograms

Employee(ssn, name, age)

$T(\text{Employee}) = 25000$, $V(\text{Employee}, \text{age}) = 50$

$\sigma_{\text{age}=48}(\text{Employee}) = ?$

Estimate: $T(\text{Employee}) / V(\text{Employee}, \text{age}) = 500$

1d-Histograms

Employee(ssn, name, age)

$T(\text{Employee}) = 25000$, $V(\text{Employee}, \text{age}) = 50$

$\sigma_{\text{age}=48}(\text{Employee}) = ?$

Estimate: $T(\text{Employee}) / V(\text{Employee}, \text{age}) = 500$

Age:	0..20	20..29	30-39	40-49	50-59	> 60
T =	200	800	5000	12000	6500	500

1d-Histograms

Employee(ssn, name, age)

$T(\text{Employee}) = 25000$, $V(\text{Employee}, \text{age}) = 50$

$\sigma_{\text{age}=48}(\text{Employee}) = ?$

Estimate: $T(\text{Employee}) / V(\text{Employee}, \text{age}) = 500$

Age:	0..20	20..29	30-39	40-49	50-59	> 60
T =	200	800	5000	12000	6500	500

Assume $V = 10$

1d-Histograms

Employee(ssn, name, age)

$T(\text{Employee}) = 25000$, $V(\text{Employee}, \text{age}) = 50$

$\sigma_{\text{age}=48}(\text{Employee}) = ?$

~~Estimate: $T(\text{Employee}) / V(\text{Employee}, \text{age}) = 500$~~

Age:	0..20	20..29	30-39	40-49	50-59	> 60
T =	200	800	5000	12000	6500	500

Estimate: $12000/10 = 1200$

Assume $V = 10$

1d-Histograms

Employee(ssn, name, age)

$T(\text{Employee}) = 25000$, $V(\text{Employee}, \text{age}) = 50$

$\sigma_{\text{age}=48}(\text{Employee}) = ?$

~~Estimate: $T(\text{Employee}) / V(\text{Employee}, \text{age}) = 500$~~

Age:	0..20	20..29	30-39	40-49	50-59	> 60
T =	200	800	5000	12000	6500	500
V =	3	10	7	6	5	4

Estimate: $12000/10 = 1200$

1d-Histograms

Employee(ssn, name, age)

$T(\text{Employee}) = 25000$, $V(\text{Employee, age}) = 50$

$\sigma_{\text{age}=48}(\text{Employee}) = ?$

~~Estimate: $T(\text{Employee}) / V(\text{Employee, age}) = 500$~~

Age:	0..20	20..29	30-39	40-49	50-59	> 60
T =	200	800	5000	12000	6500	500
V =	3	10	7	6	5	4

Estimate: ~~$12000/10 = 1200$~~ $12000/6 = 2000$

Types of 1d-Histograms

- Eq-Width
- Eq-Depth
- Compressed: store outliers separately
- V-Optimal histograms

Employee(ssn, name, age)

Types of 1d-Histograms

Eq-width:

Age:	0..20	20..29	30-39	40-49	50-59	> 60
Tuples	200	800	5000	12000	6500	500

Eq-depth:

Age:	0..32	33..41	42-46	47-52	53-58	> 60
Tuples	1800	2000	2100	2200	1900	1800

Compressed: store separately highly frequent values: (48,1900)

V-Optimal Histogram

- $domain(A) = \{v_1, \dots, v_n\}$
- Have budget of $b+1$ buckets for R.A:
$$-\infty < d_1 < d_2 < \dots < d_b < \infty$$
- Choose d_1, \dots, d_b to minimize error of all queries $\sigma_{A=v}(R)$, for $v \in Domain(A)$

V-Optimal Histogram

- Error:

$$\sum_{v \in \text{Domain}(A)} \left(|\sigma_{A=v}(R)| - \text{est}_{\text{Hist}}(\sigma_{A=v}(R)) \right)^2$$

- Bucket boundaries = $\text{argmin}_{\text{Hist}}(\text{Error})$
- Dynamic programming

Discussion: 1d-Histograms

- All systems support some forms
- Histograms need to be small to reside in main memory: 1000 – 10000 buckets
- Recomputed periodically, often from samples. E.g Postgres ANALYSE

2d-Histograms

```
SELECT *  
FROM Supplier  
WHERE scity = 'Seattle'  
and sstate = 'WA'
```

	AL-AR	CA-FL	...	TX-WA	WV-WY
Ab..Co					
...					
Sa...Tu					

Independence only
needed within the bucket

T(bucket),
V(bucket,scity)
V(bucket,sstate)

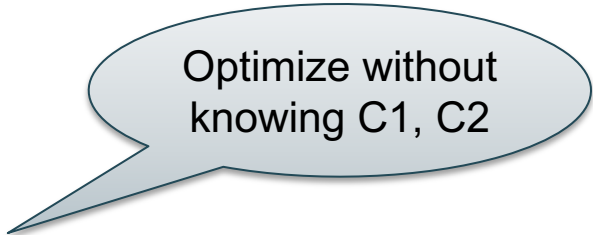
Issues with 2d-Histograms

- Limited # of dividers: $\sqrt{1000} - \sqrt{10000}$
- Too many 2d-histograms: $n(n-1)/2$
- Problem: how do we estimate $P(A,B,C)$ given all 1d and 2d-histograms?
- Few systems support 2d-histograms

Yet Another Difficulty


- SQL Queries issued from applications
- Optimized once using **prepare** statement, executed repeatedly
- Constants in the query are not know at optimization time

```
select
  o_year, sum(case when nation = 'BRAZIL' then volume else 0 end) / sum(volume)
from
  (select YEAR(o_orderdate) as o_year,
    l_extendedprice * (1 - l_discount) as volume,
    n2.n_name as nation
  from part, supplier, lineitem, orders,
    customer, nation n1, nation n2, region
  where p_partkey = l_partkey and s_suppkey = l_suppkey
    and l_orderkey = o_orderkey and o_custkey = c_custkey
    and c_nationkey = n1.n_nationkey
    and n1.n_regionkey = r_regionkey
    and r_name = 'AMERICA'
    and s_nationkey = n2.n_nationkey
    and o_orderdate between '1995-01-01'
    and '1996-12-31'
    and p_type = 'ECONOMY ANODIZED STEEL'
    and s_acctbal ≤ C1 and l_extendedprice ≤ C2 ) as all_nations
group by o_year order by o_year
```



Optimize without
knowing C1, C2

QueryTemplate Plan Diag Reduced Plan Diag Comp Cost Diag Comp Card Diag Exec Cost Diag Exec Card Diag Sel Log
Plan Diagram QTD: DB2_9_opp_U_100_q0_30ap1 # of Plans: 76



Different optimal plans for different C1, C2

Min Est Cost: 8.26E5
Max Est Cost: 1.05E6
Min Est Card: 5.98E-2
Max Est Card: 9.08E0

Parameter → Operator Diff
Regenerate Diagram
Reset View

Gini Coeff: 0.83

P1	29.60 %
P2	17.69 %
P3	8.47 %
P4	4.73 %
P5	4.19 %
P6	4.02 %
P7	2.85 %
P8	2.49 %
P9	2.43 %
P10	2.38 %
P11	2.38 %
P12	1.63 %
P13	1.56 %
P14	1.30 %
P15	1.27 %
P16	1.21 %
P17	1.06 %
P18	0.91 %
P19	0.82 %
P20	0.76 %
P21	0.71 %
P22	0.71 %
P23	0.71 %
P24	0.62 %
P25	0.58 %

Traditional CE -- Summary

- Used by all DBMS
- Assumptions: uniformity, independence,...
- 1d-histograms; sometimes also 2d
- Often underestimate (why?)

Alternate CE Methods: ML and Sampling

ML-Based CE

- A trend in research: use some ML model to do cardinality estimation
- Main hope: remove independence and uniformity assumptions
- Data-driven or Query-driven

Data-Driven Estimators

- Train a generative ML model that represents $p(A,B,C,\dots)$ from the database instance
- Map query Q to hyperrectangle
- Problem: space is over all attributes
 - “Full outer join”
 - Lots of complications to make this work

Query-Driven Estimators

- Training set is a set of pairs $(Q, |Q|)$
- All predicates in the *WHERE* condition need to be featurized, embedded
- Train a discriminative model for $\text{Est}(Q)$

Discussion

Problems with ML-based CE:

- Large model size: 1MB – 1GB
- Inference time varies
- Updates, change in skewness, correlations
- Not explainable

Consensus: **not ready for production**

Sampling-Based CE

- Main idea: use a sample of the database to estimate the output size
- New probability space: random choices of the sampler v.s. data distribution
- Offline Sampling or Online Sampling

Offline Sampling

- Compute uniform sample $R_{sample} \subseteq R$
- Horwitz-Thompson:

$$Est(Q(DB)) = \frac{|R|}{|R_{sample}|} \cdot \frac{|S|}{|S_{sample}|} \cdots |Q(DB_{sample})|$$

- The missing tuple problem: $Q(DB_{sample}) = 0$
- High variance

Online Sampling

- To estimate $R \bowtie S$:
 - Sample a tuple from R
 - Sample a **matching tuple** from S
 - Repeat
- Need to use index on S
- **WanderJoin**: next slides show how to estimate $R \bowtie S \bowtie T$ using random walk

R

A	B
...	1
...	2
...	1
...	3
...	4
...	2
...	4

R


A	B
...	1
...	2
...	1
...	3
...	4
...	2
...	4

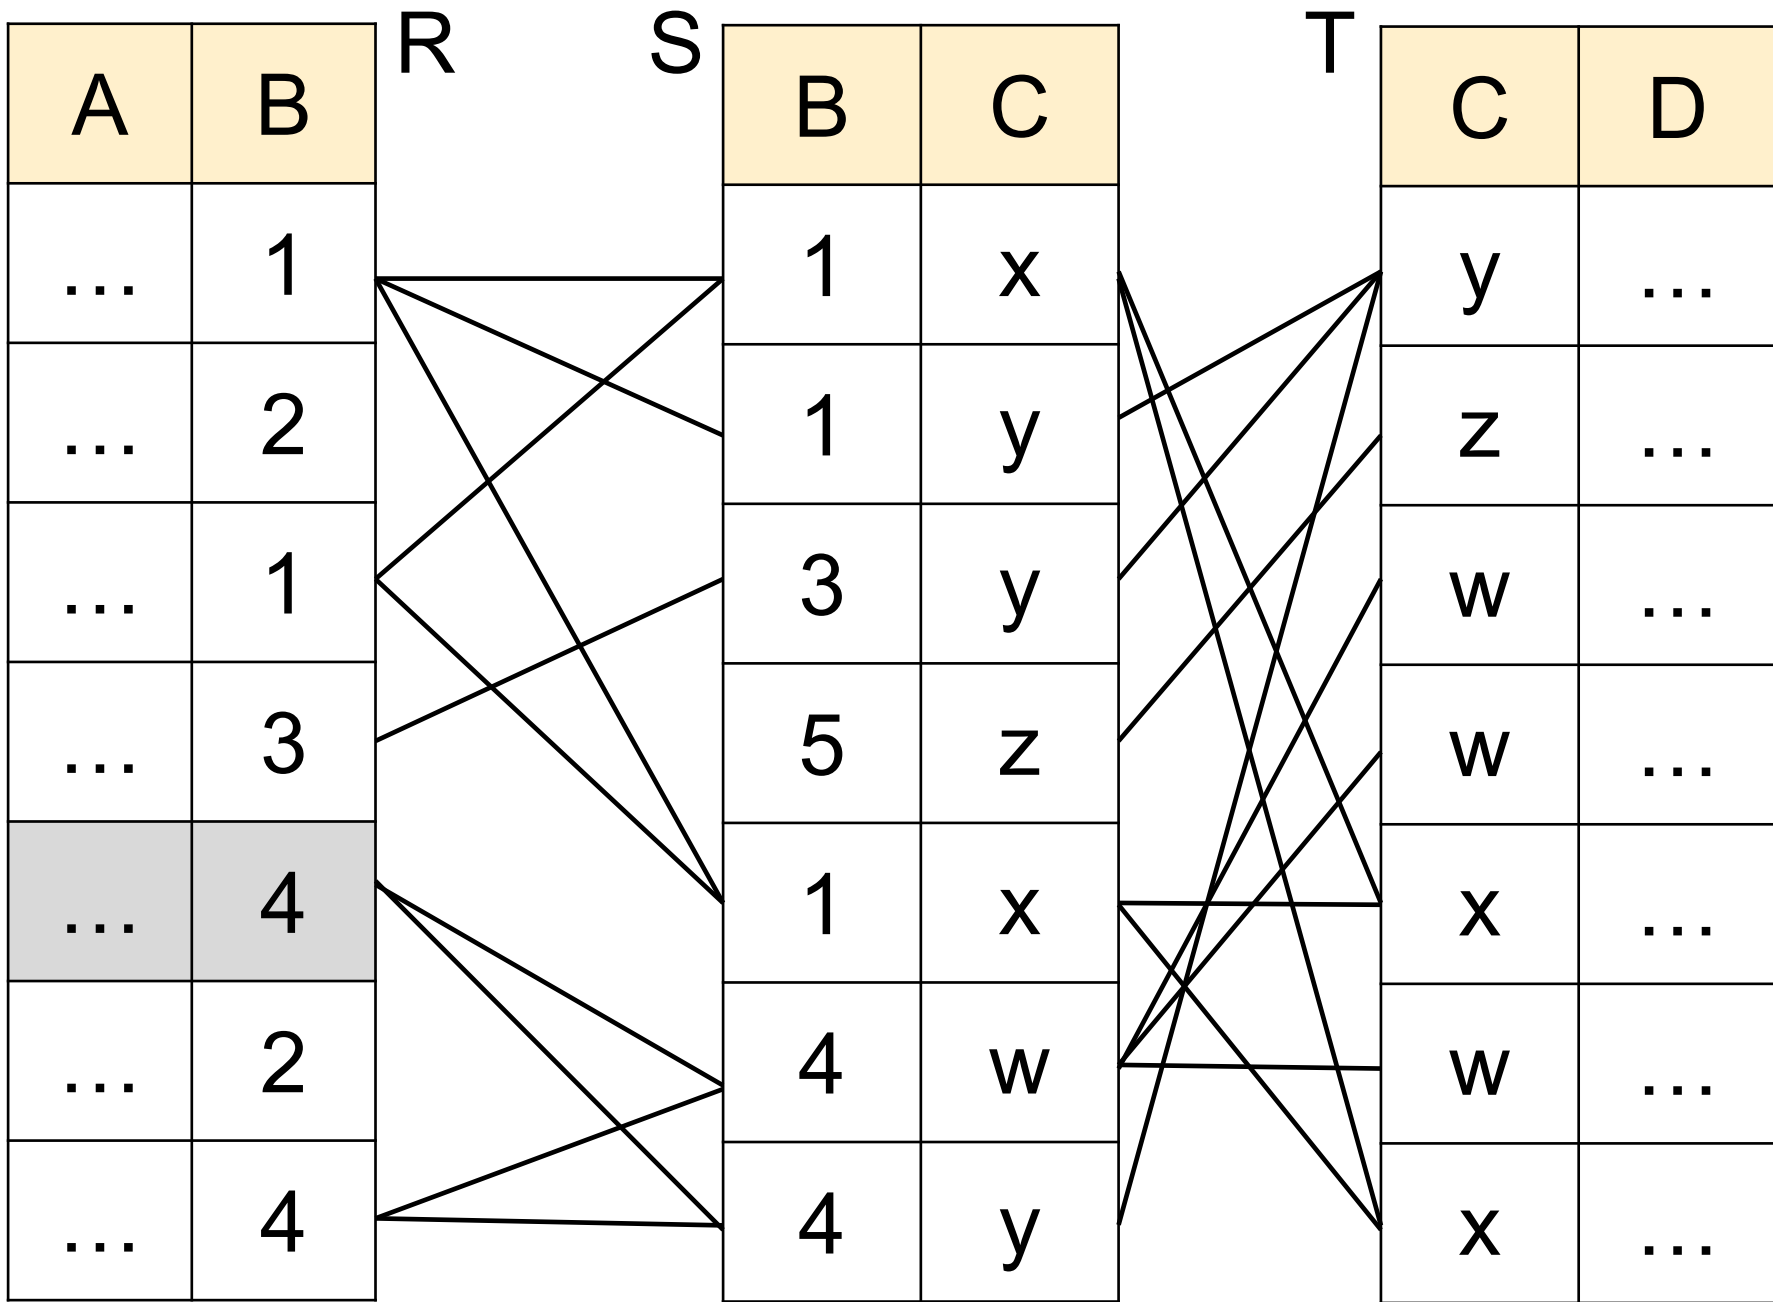
S


B	C
1	x
1	y
3	y
5	z
1	x
4	w
4	y


T

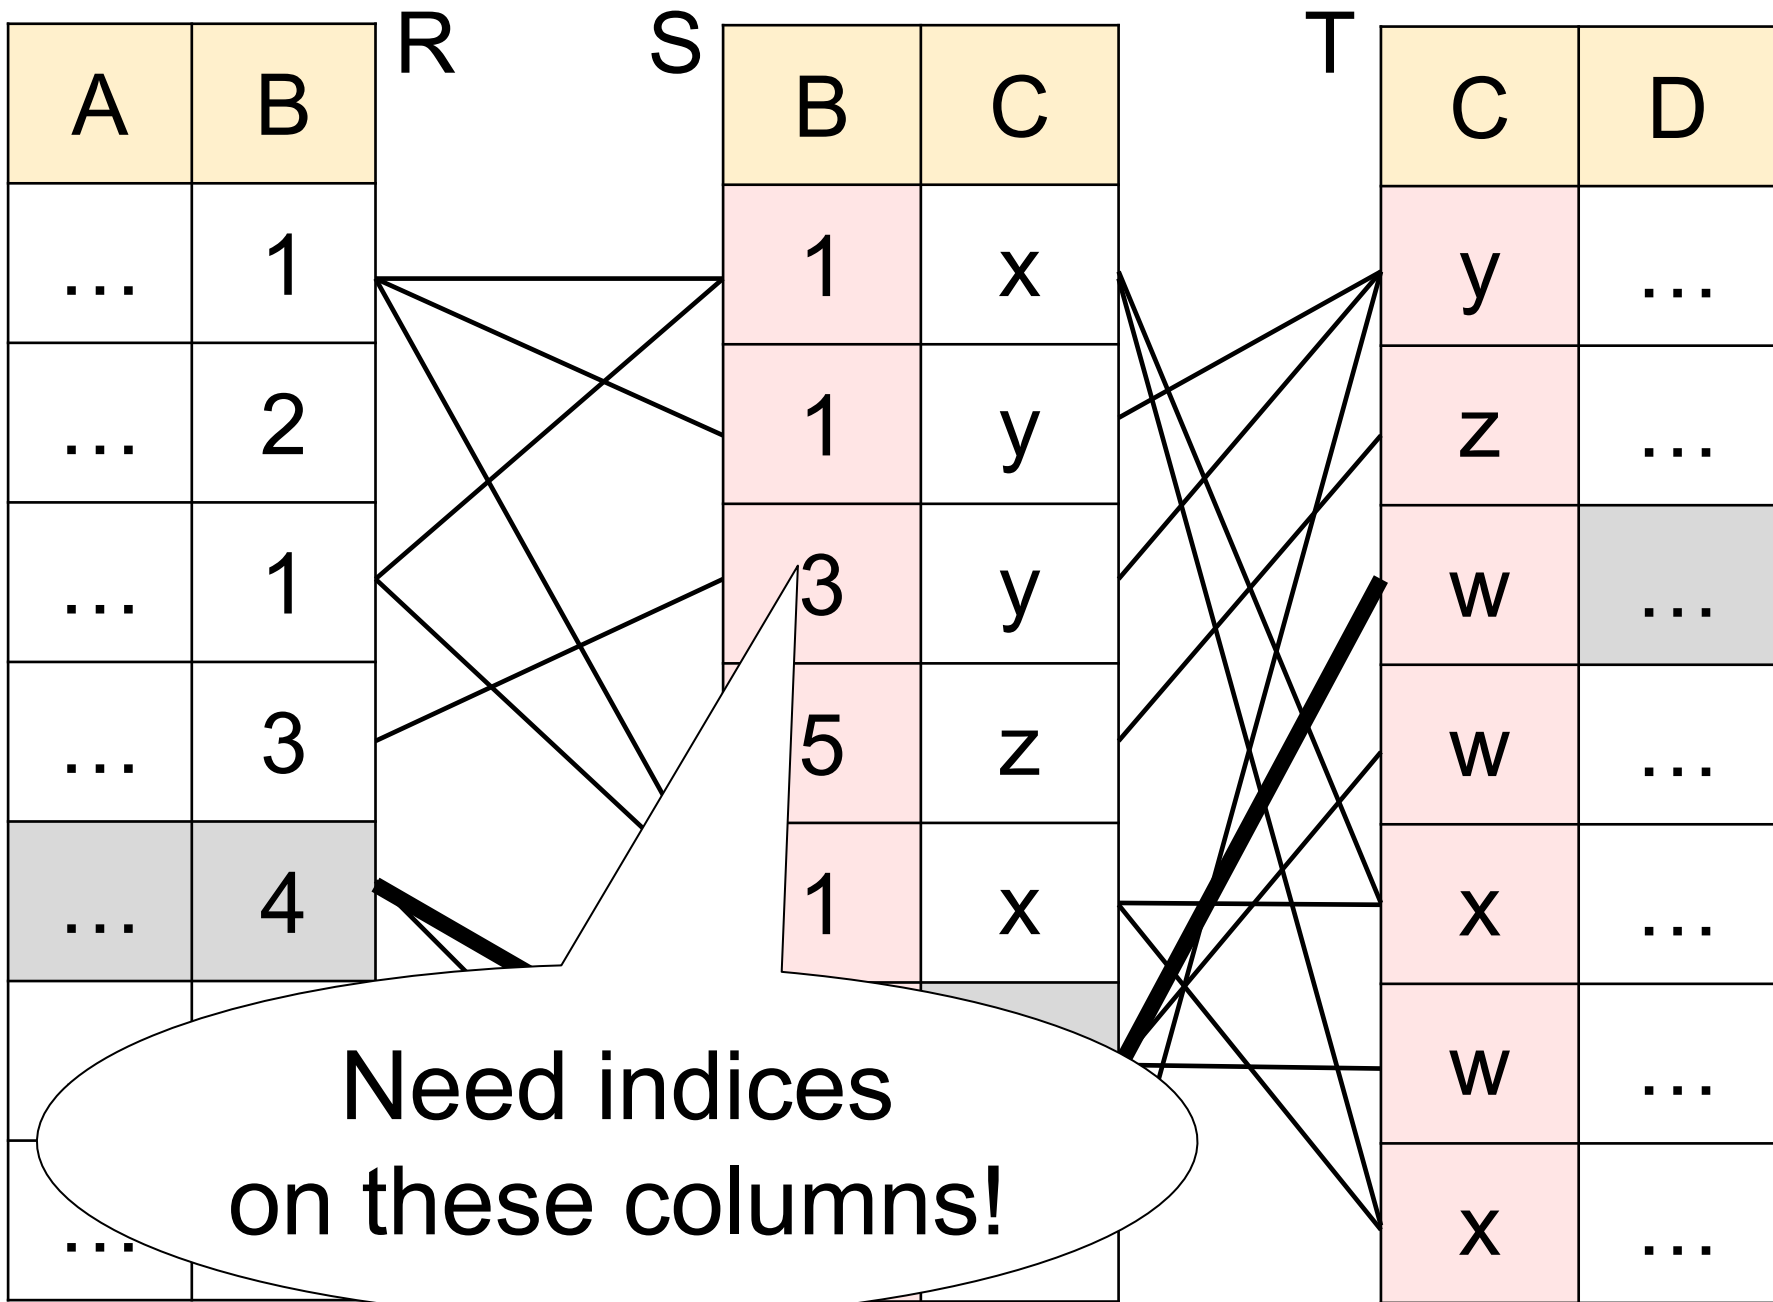
C	D
y	...
z	...
w	...
w	...
x	...
w	...
x	...

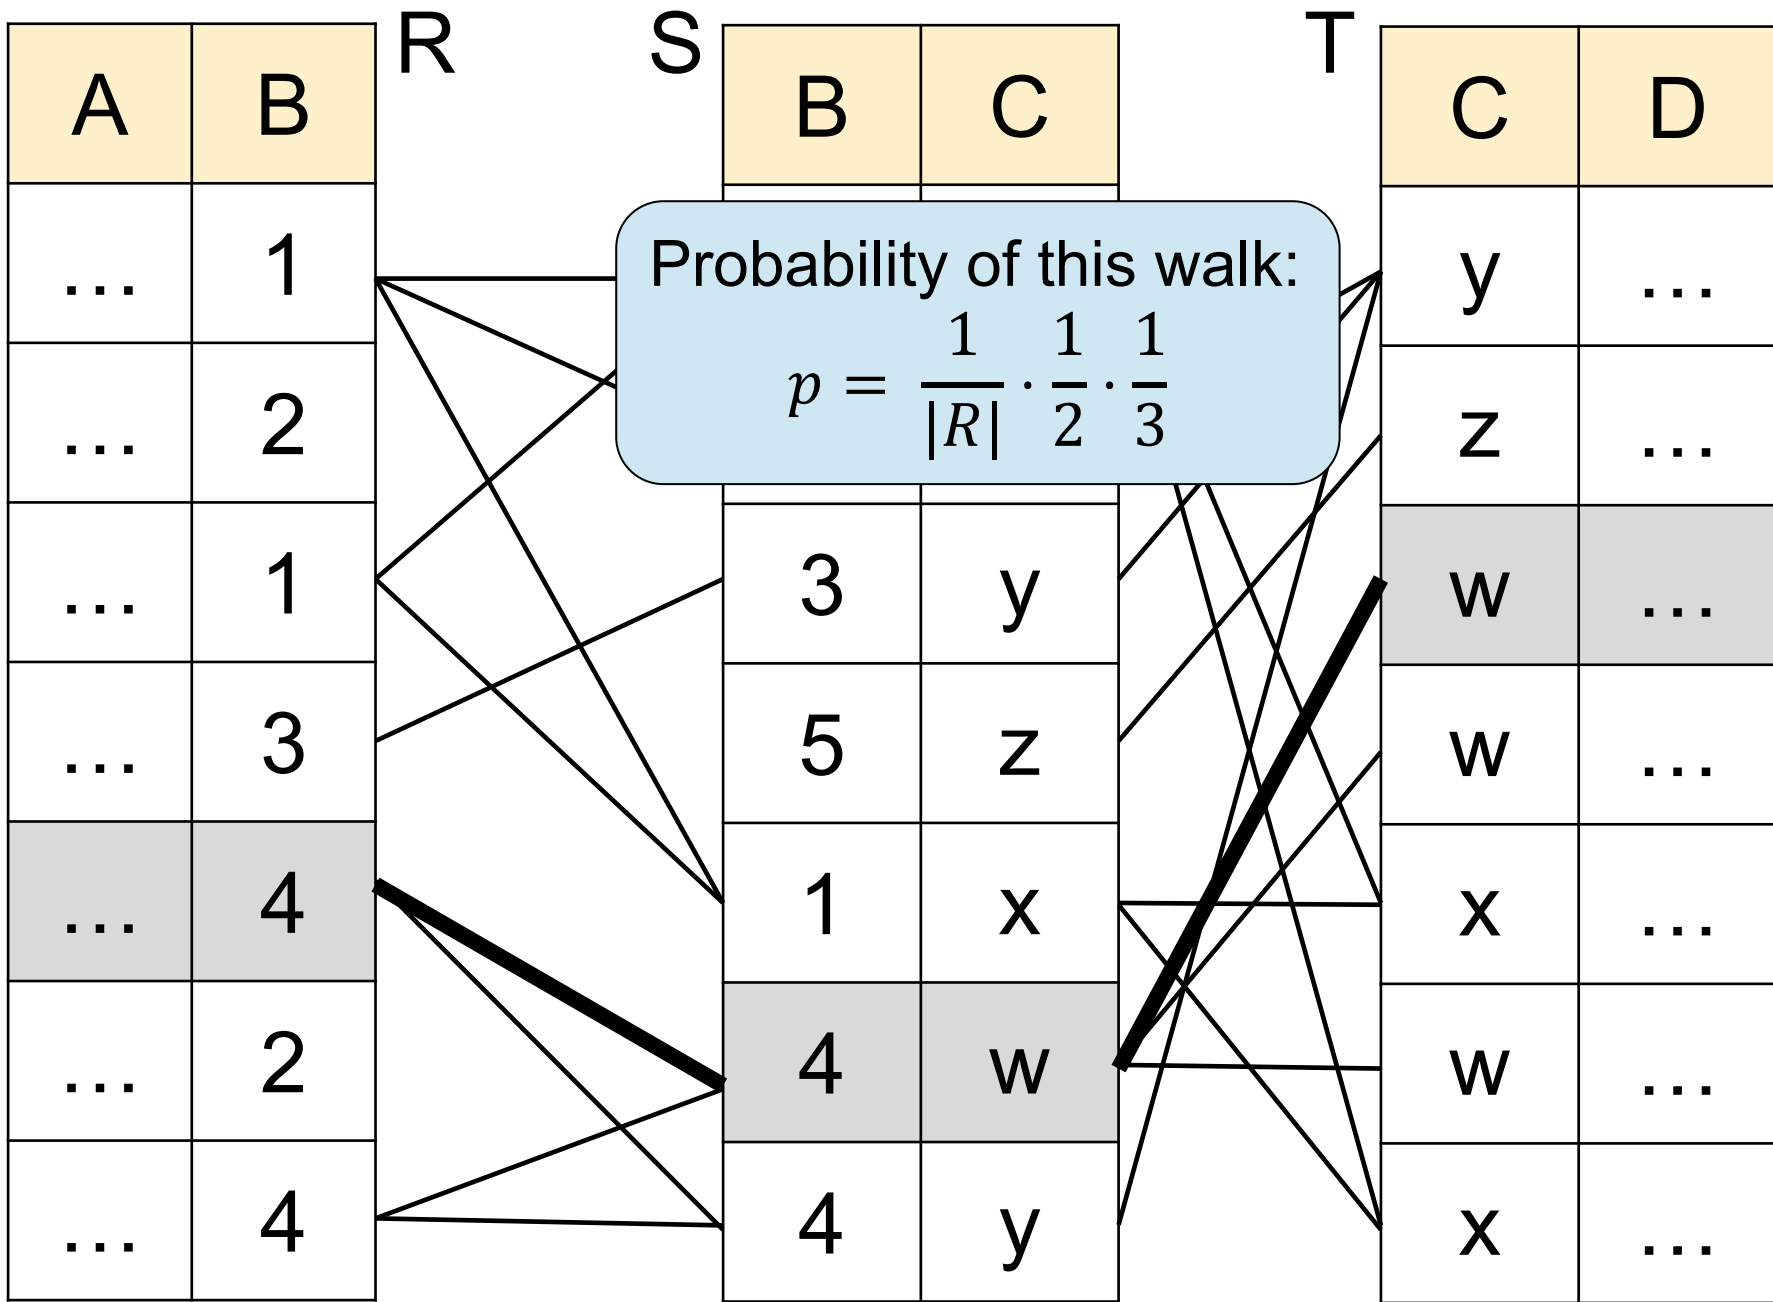


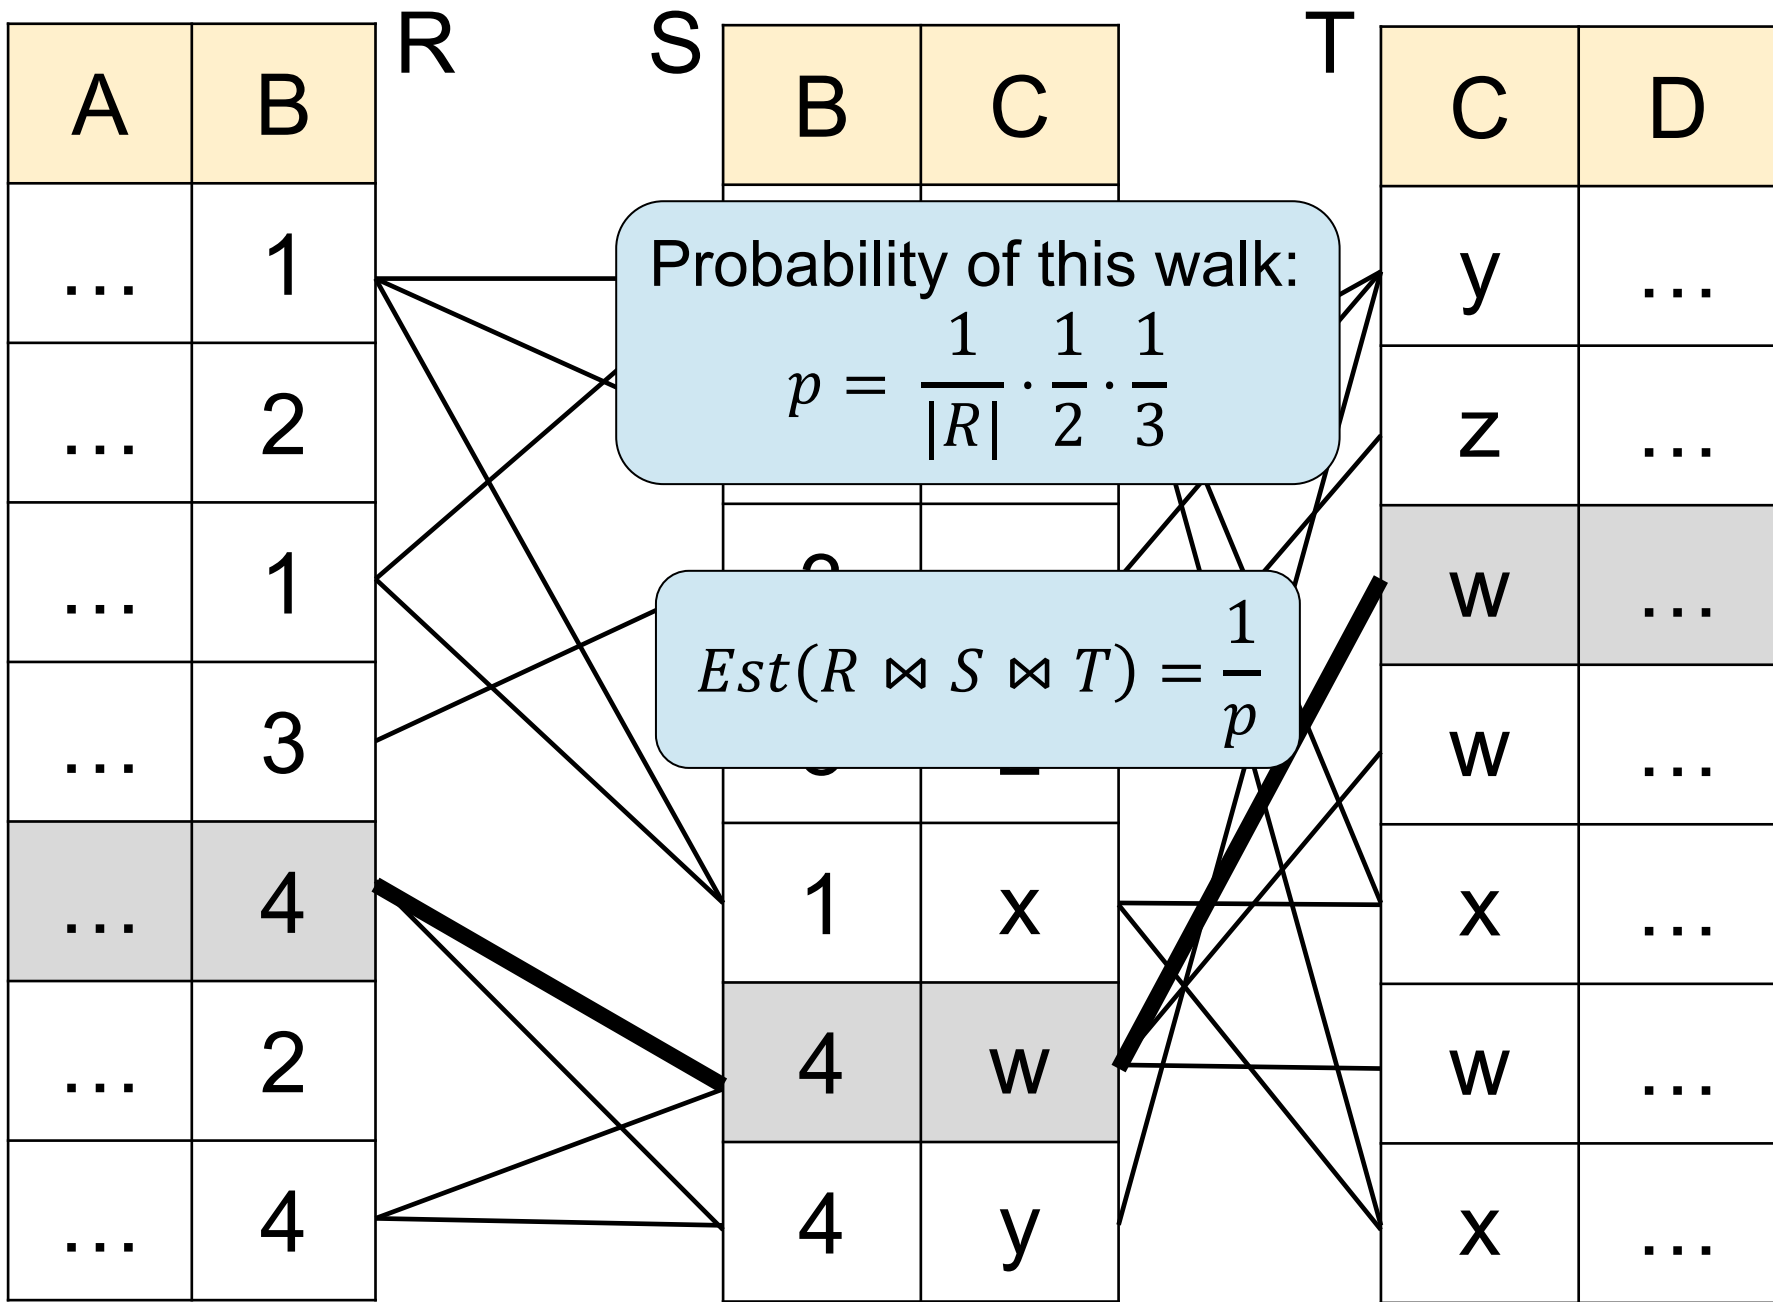


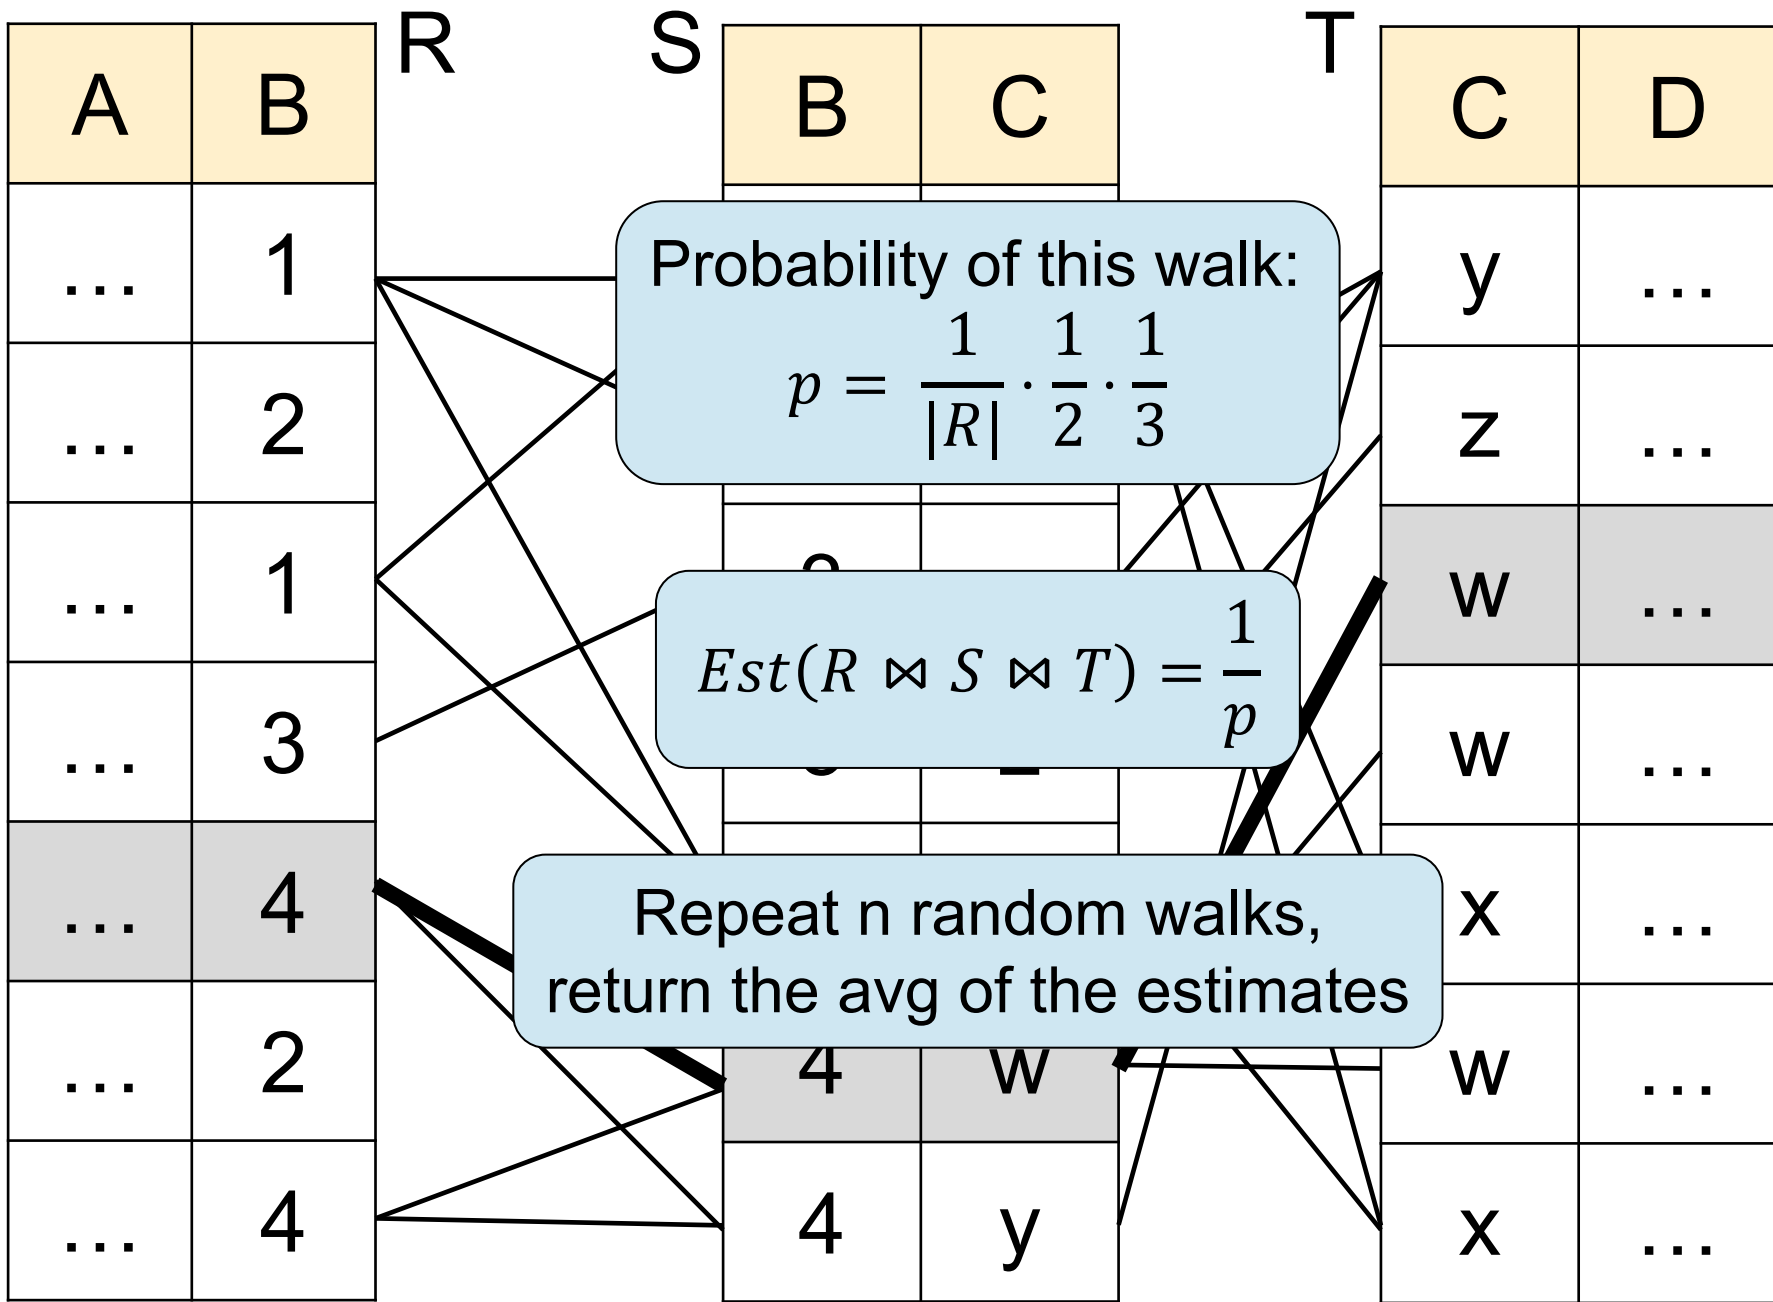












Discussion

- Sampling based CE have best precision
- However:
 - Offline samples need to be too large
 - Online samples require access to indices
- Some systems use offline sample for single-table predicates only

Query Optimization

Three major components:

1. Search space last week
2. Cardinality and cost estimation today
3. Plan enumeration algorithms today+Wed.

Two Types of Optimizers

- Heuristic-based optimizers
- Cost-based optimizers (next)

Two Types of Plan Enumeration Algorithms

- Dynamic programming (in class)
 - Based on System R [Selinger 1979]
 - *Join reordering algorithm*
- Rule-based algorithm (will not discuss)
 - Database of rules (=algebraic laws)
 - Usually: dynamic programming

System R Optimizer

For each subquery $Q \subseteq \{R_1, \dots, R_n\}$, compute best plan:

- Step 1: $Q = \{R_1\}, \{R_2\}, \dots, \{R_n\}$
- Step 2: $Q = \{R_1, R_2\}, \{R_1, R_3\}, \dots, \{R_{n-1}, R_n\}$
- ...
- Step n: $Q = \{R_1, \dots, R_n\}$

Details

For each subquery $Q \subseteq \{R_1, \dots, R_n\}$ store:

- Estimated Size(Q)
- A best plan for Q: Plan(Q)
- The cost of that plan: Cost(Q)

} One plan
for each
“interesting
order”

Details

Step 1: single relations $\{R_1\}, \{R_2\}, \dots, \{R_n\}$

- Consider all possible access paths:
 - Sequential scan, or
 - Index 1, or
 - Index 2, or
 - ...
- Keep optimal plan for each “interesting order”

Details

Step $k = 2 \dots n$:

For each $Q = \{R_{i_1}, \dots, R_{i_k}\}$

- For each $j=1, \dots, k$:
 - Let: $Q' = Q - \{R_{i_j}\}$
 - Let: $Plan(Q') \bowtie R_{i_j} \quad Cost(Q') + CostOf(\bowtie)$
- $Plan(Q), Cost(Q) =$ cheapest of the above
 - Keep separate optimal for “interesting orders”

Discussion

- All database systems implement Selinger's algorithm for join reorder
- For other operators (group-by, aggregates, difference): rule-based
- Many search strategies beyond dynamic programming

Final Discussion

- Query optimizer = critical part of DBMS
- Search space + Size est + Algorithm
- Ideal: find “optimal” plan
- In practice: avoid “very bad plans”
- Successful because:
 - RA is a set-at-a-time language
 - RA is order-independent
- Next time:
How good are they?; New approaches