

Query Processing

- $Q \rightarrow$ Query Plan

Example:

Select B,D

From R,S

Where $R.A = \text{"c"} \wedge S.E = 2 \wedge R.C = S.C$


How do we execute query?

One idea

- Do Cartesian product
- Select tuples
- Do projection

Relational Algebra - can be used to describe plans...

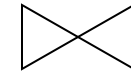

Plan I



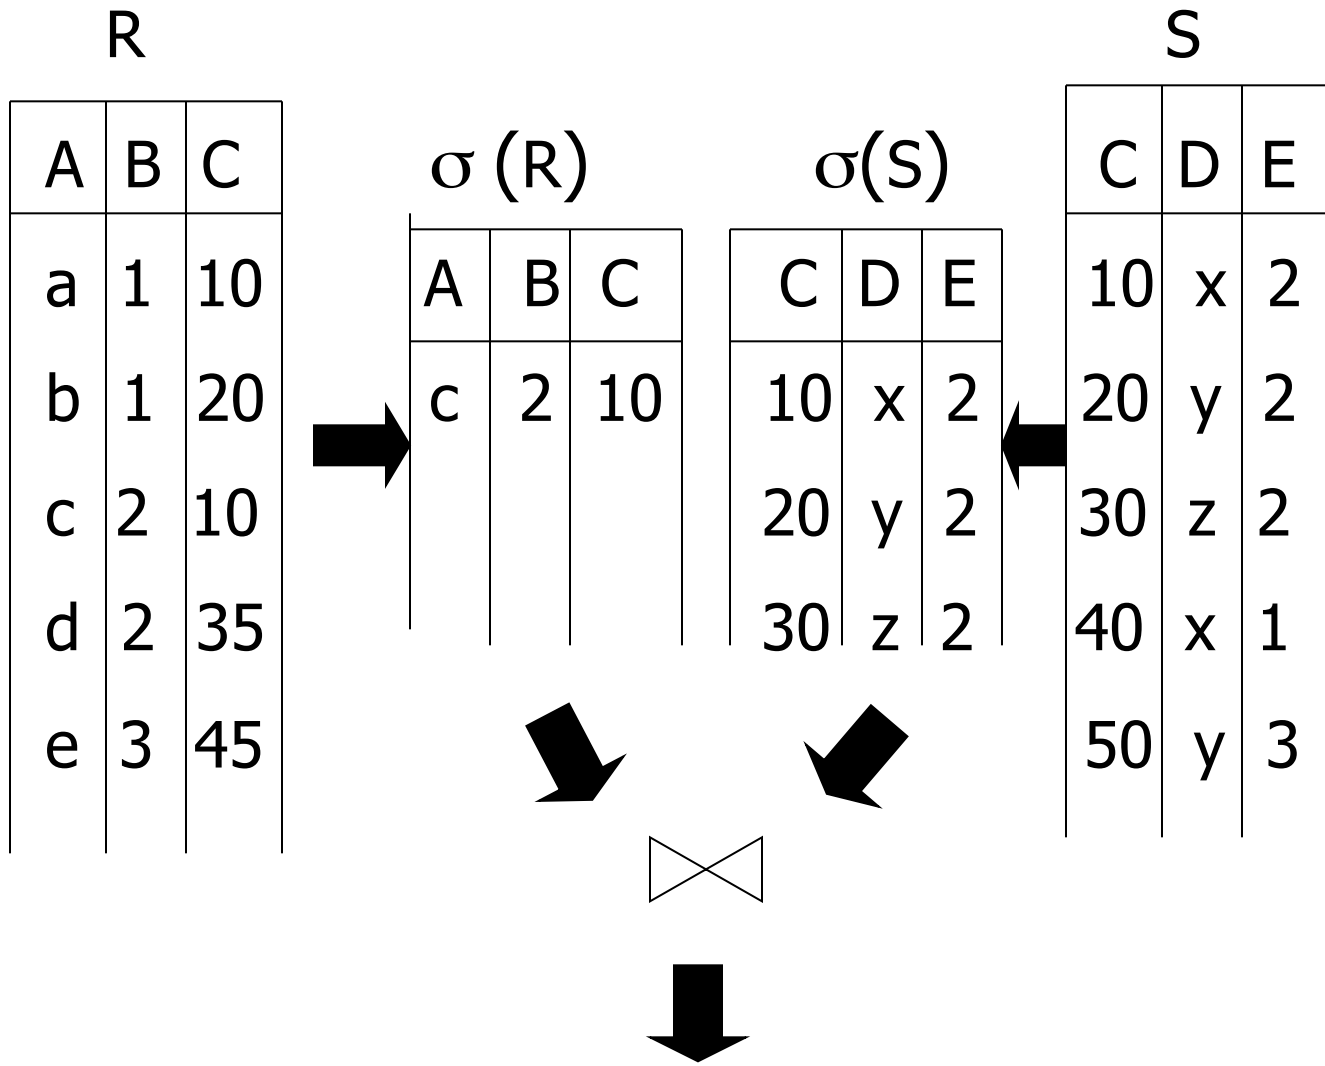
OR: $\Pi_{B,D} [\sigma_{R.A="c" \wedge S.E=2 \wedge R.C=S.C} (RXS)]$

Another idea:

Plan II



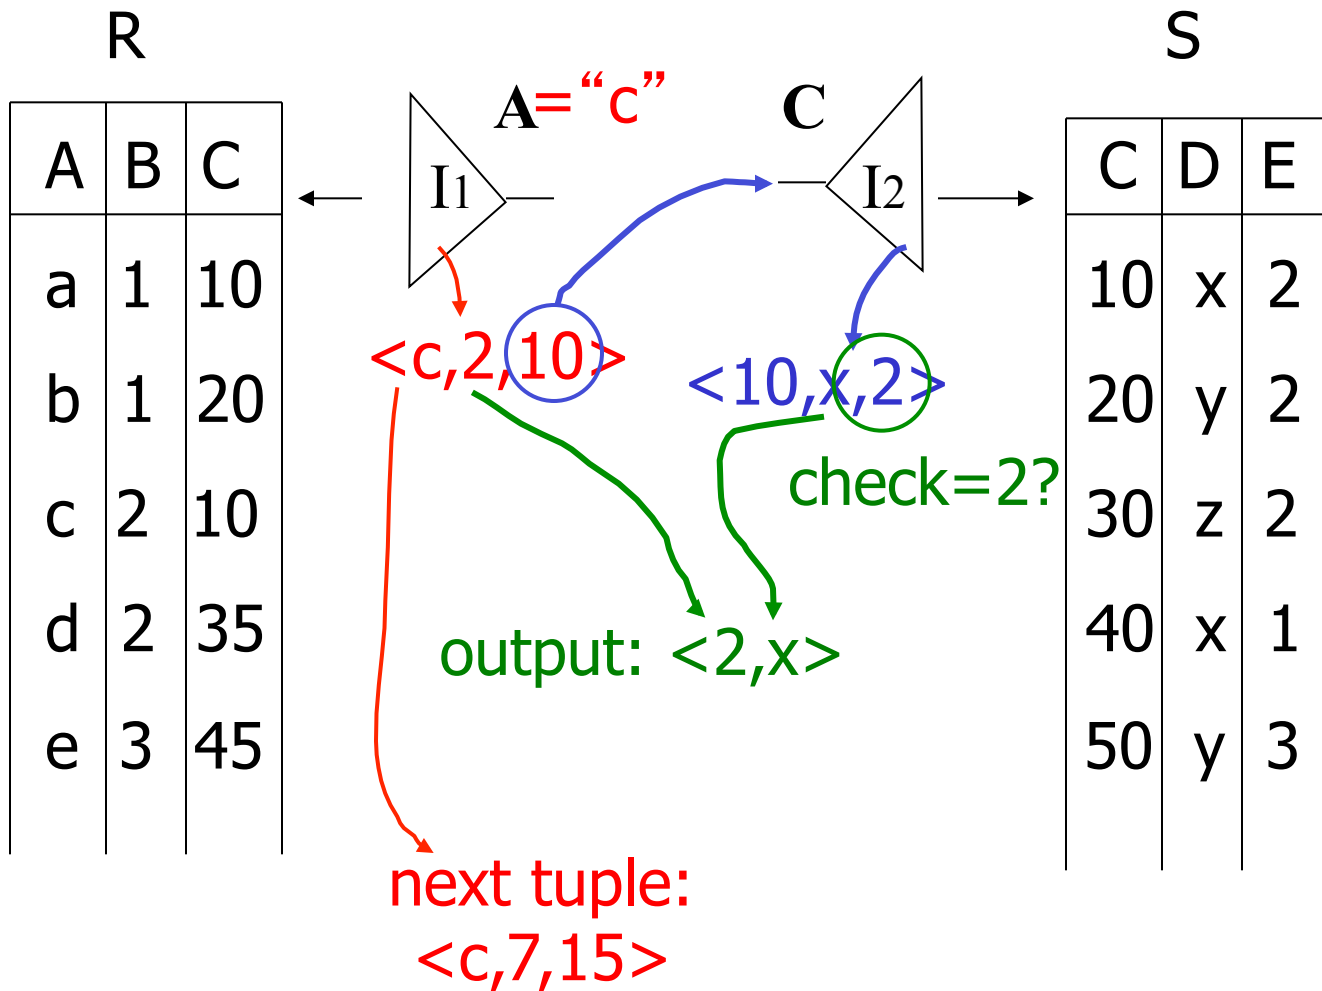
natural join



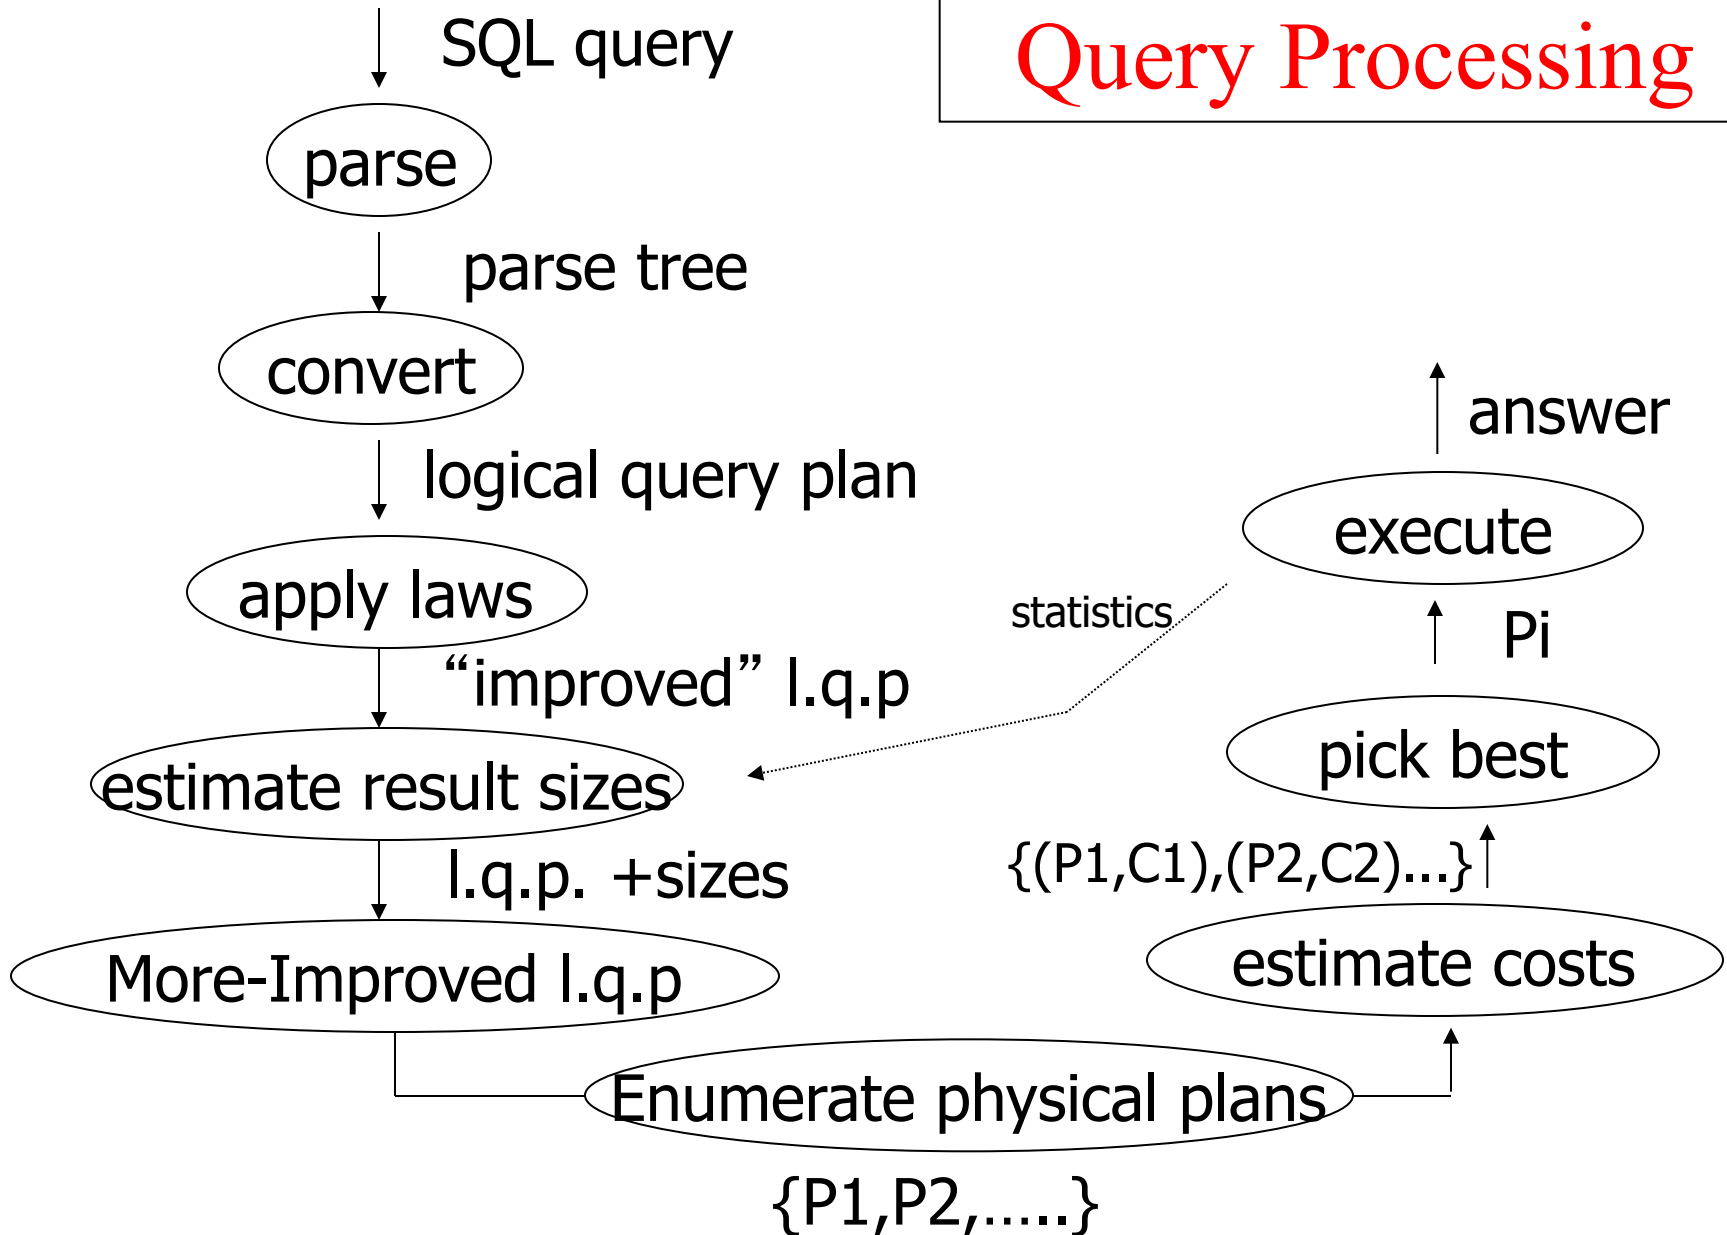
Plan III

Use R.A and S.C Indexes

- (1) Use R.A index to select R tuples with R.A = “c”
- (2) For each R.C value found, use S.C index to find matching tuples
- (3) Eliminate S tuples S.E \neq 2
- (4) Join matching R,S tuples, project B,D attributes and place in result



Query Processing



Key Steps in Query Processing

Logical query plan(s)

→ **Physical query plans**

→ **“Best” Physical plan.**

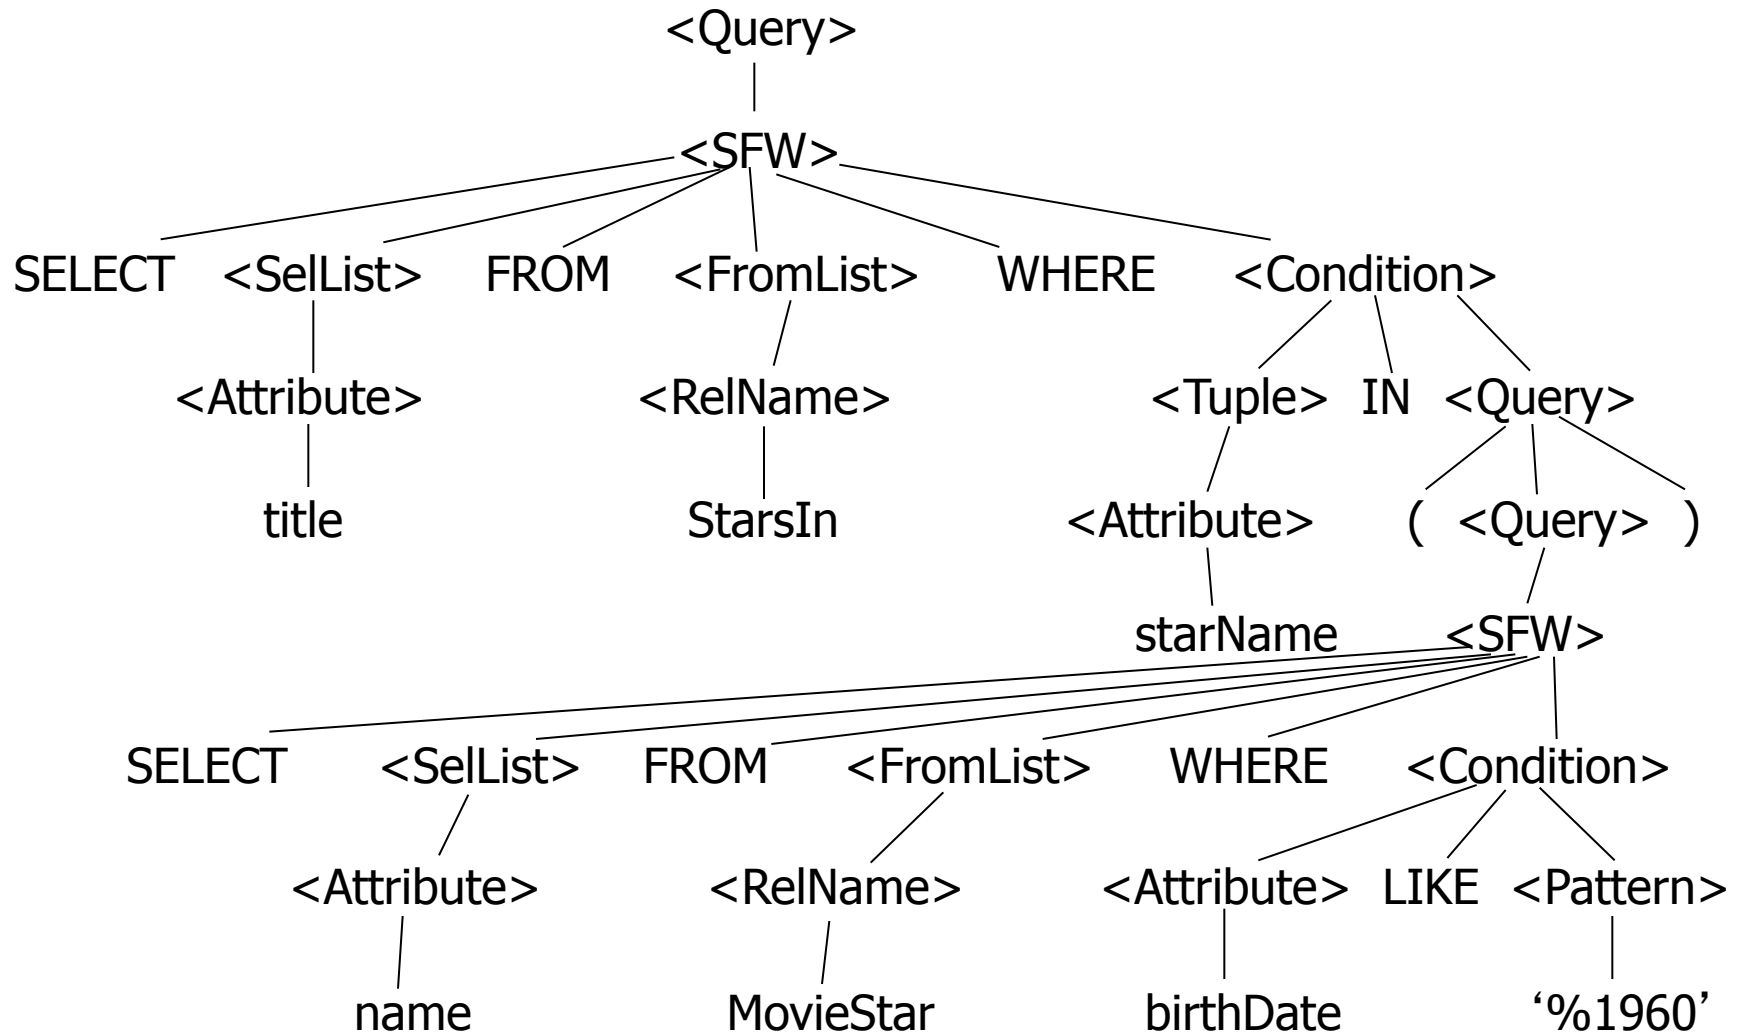
- Improving logical/physical plans requires **cost estimation** – which in turn requires **size estimation** of intermediate results.
- **Size estimation techniques**: Keep data statistics, and use certain models (later slides).

Example

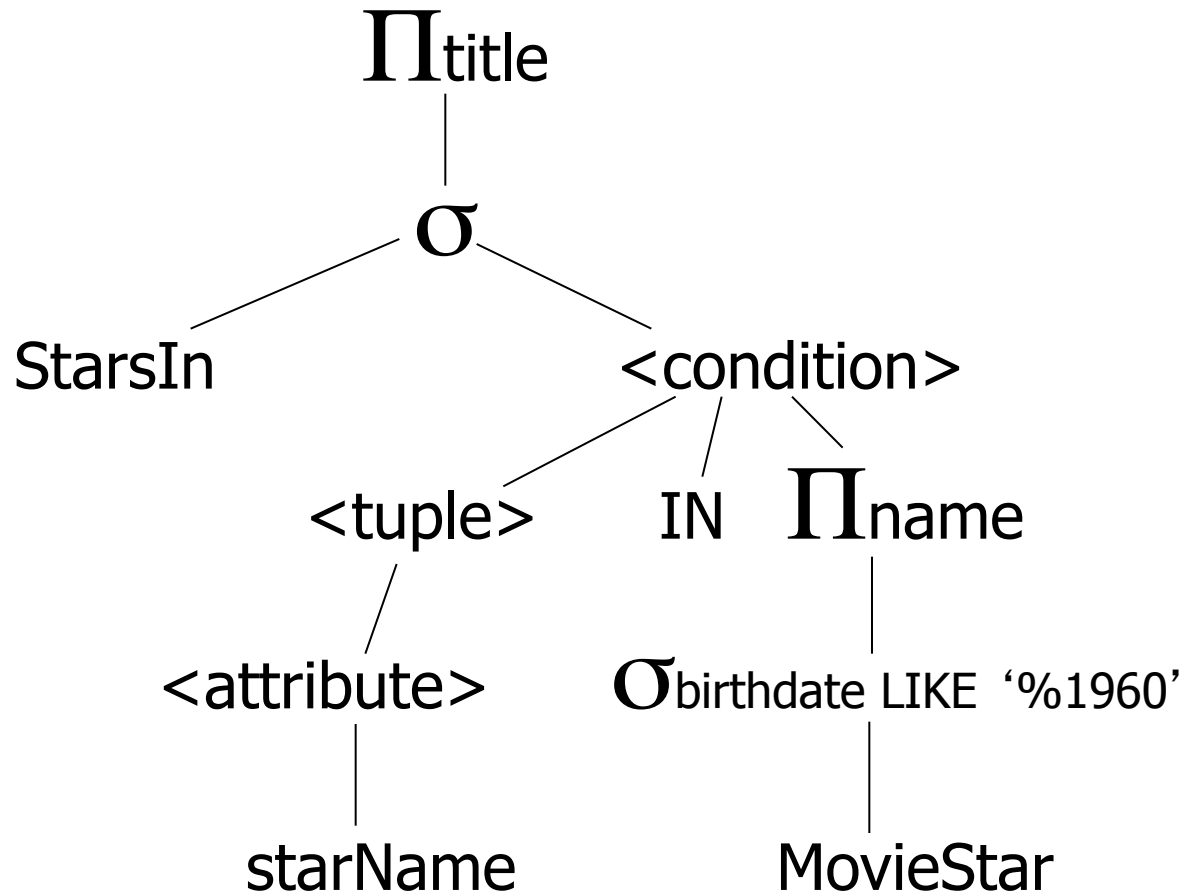
```
SELECT title
FROM StarsIn
WHERE starName IN (SELECT name
                    FROM MovieStar
                    WHERE birthdate LIKE '%1960');
```

(Find the movies with stars born in 1960)


Example: Parse Tree




Example: Generating Relational Algebra



Example: Logical Query Plan

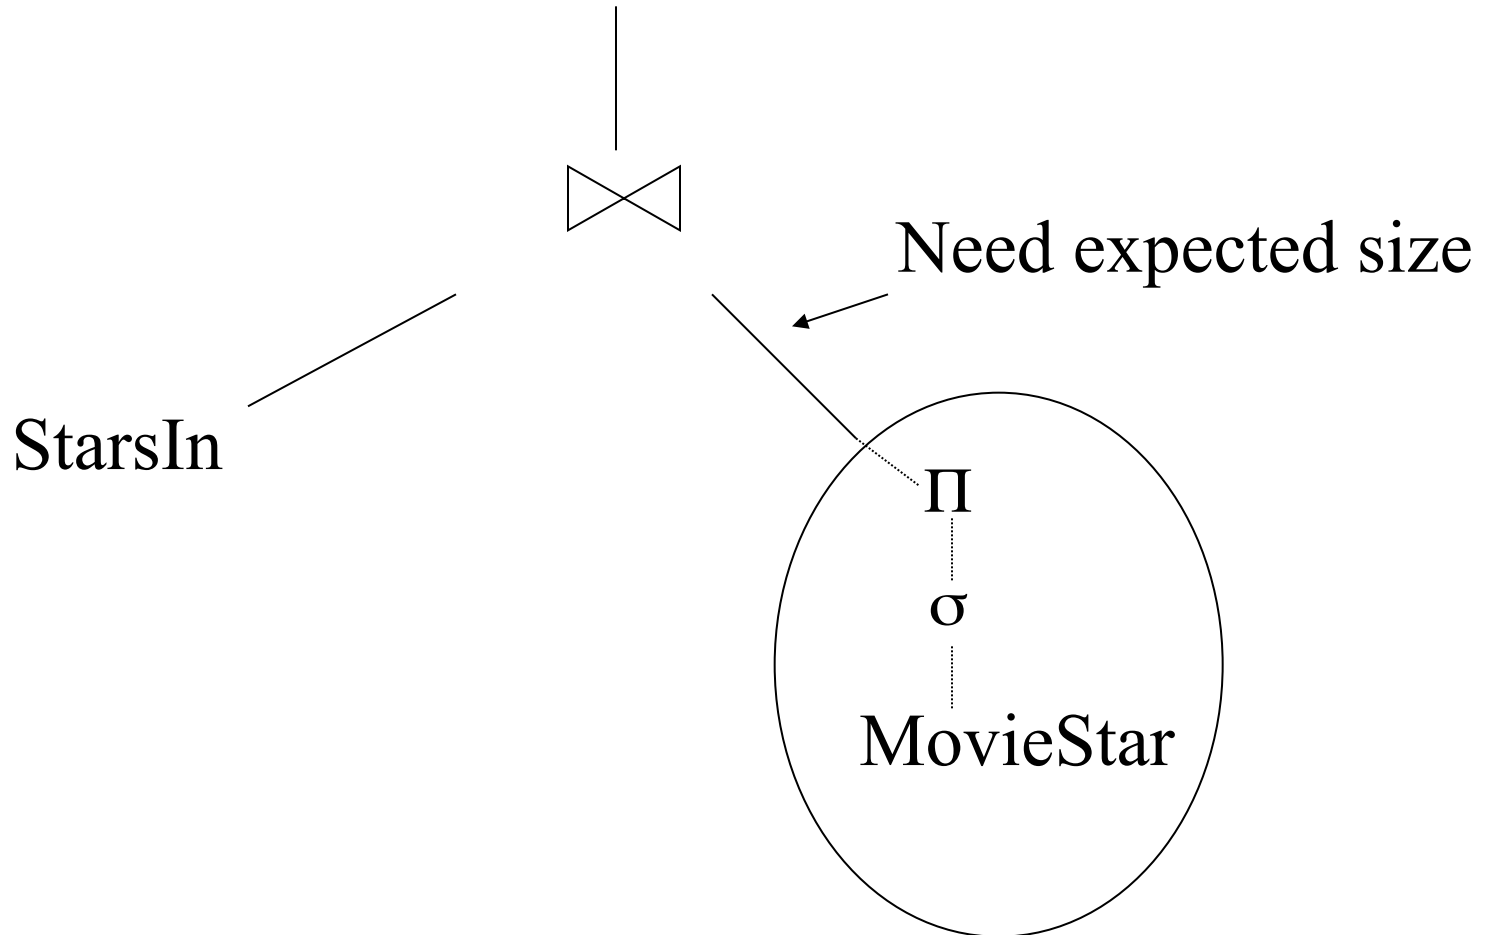


Example: Improved Logical Query Plan

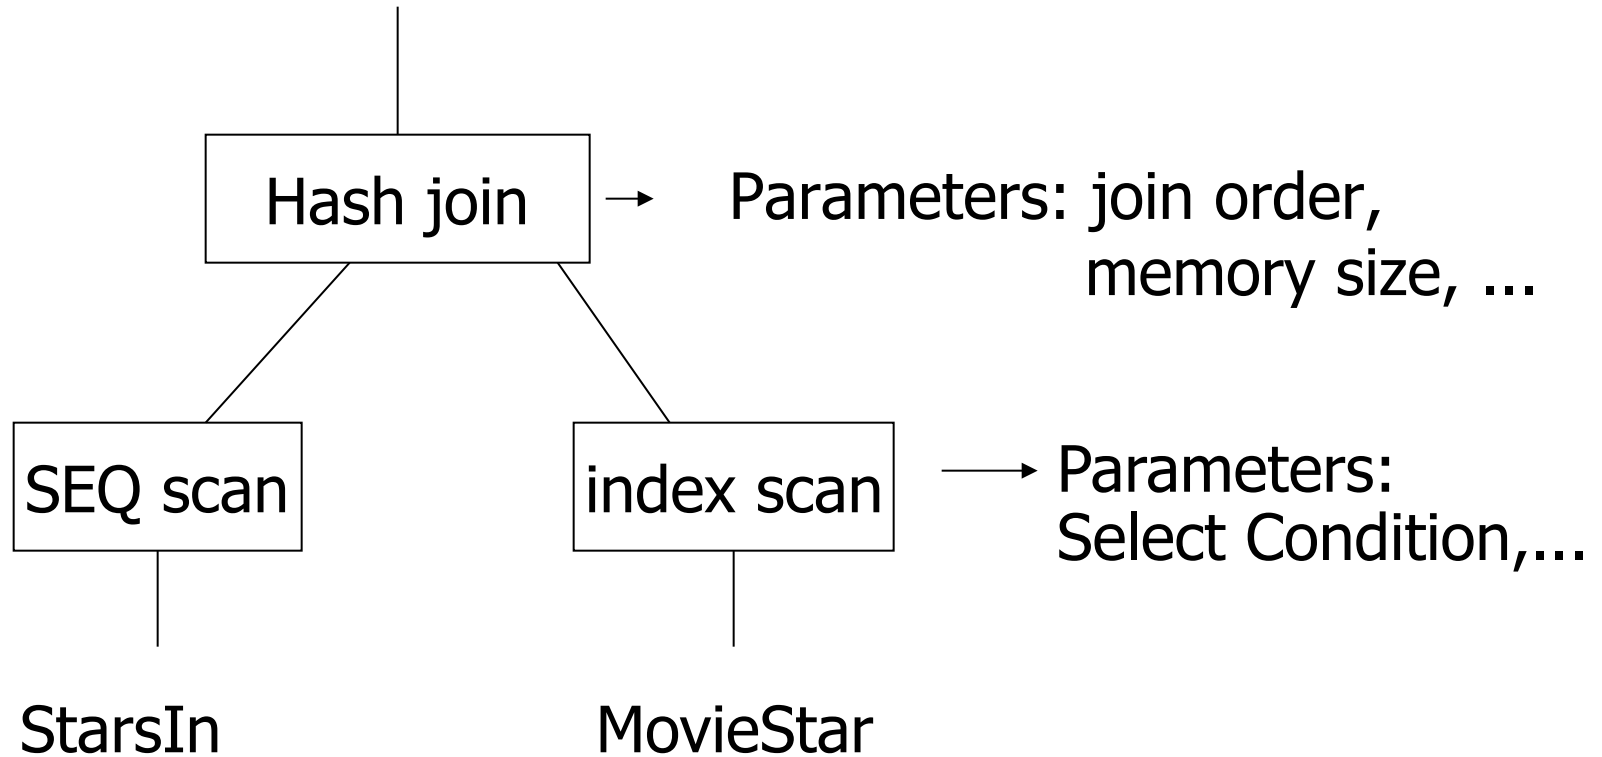


Question:
Push project to
StarsIn?

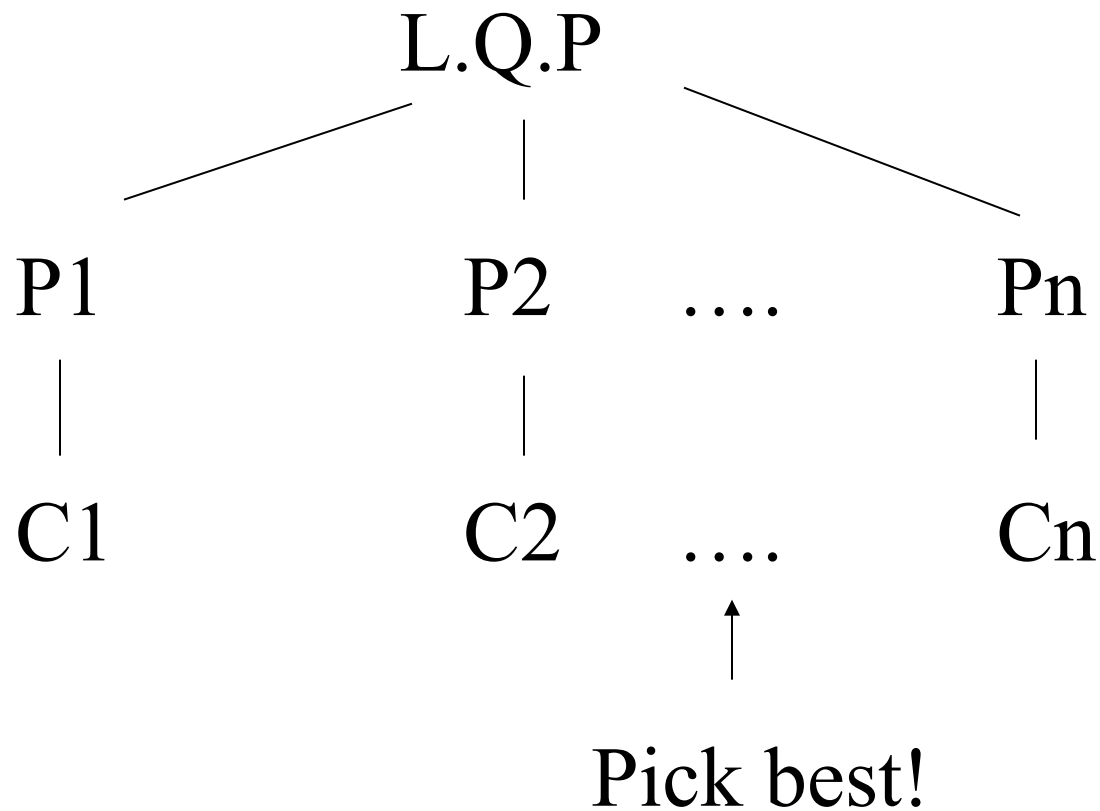
Example: Estimate Result Sizes



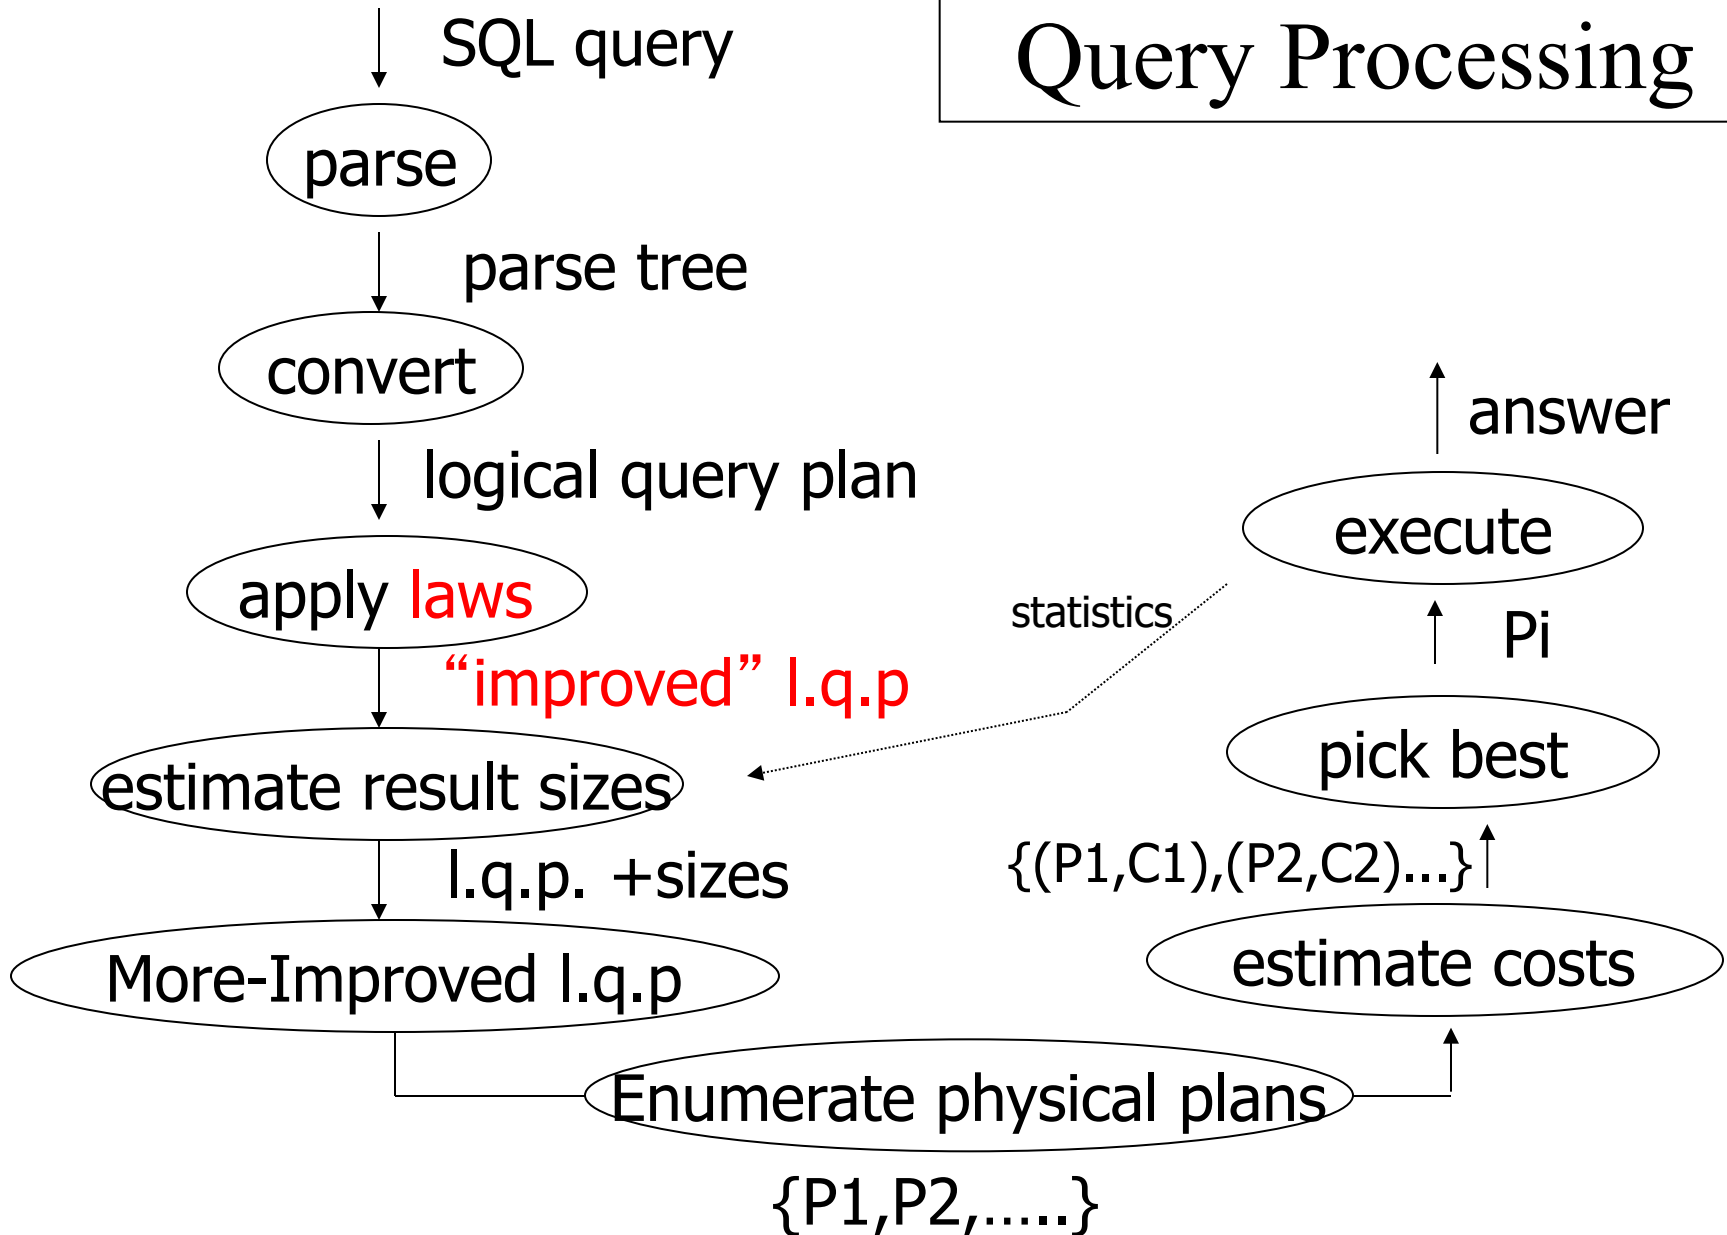
Example: One Physical Plan



Example: Estimate costs



Query Processing



Rules: Joins, products, union

$$R \bowtie S = S \bowtie R$$

$$(R \bowtie S) \bowtie T = R \bowtie (S \bowtie T)$$

$$R \times S = S \times R$$

$$(R \times S) \times T = R \times (S \times T)$$

$$R \cup S = S \cup R$$

$$R \cup (S \cup T) = (R \cup S) \cup T$$

Rules: Selects

$$\sigma_{p_1 \wedge p_2}(R) = \sigma_{p_1} [\sigma_{p_2}(R)]$$

$$\sigma_{p_1 \vee p_2}(R) = [\sigma_{p_1}(R)] \cup [\sigma_{p_2}(R)]$$

Rules: σ + \bowtie combined

Let p = predicate with only R attribs

q = predicate with only S attribs

m = predicate with only R,S attribs

$$\sigma_p (R \bowtie S) = [\sigma_p (R)] \bowtie S$$

$$\sigma_q (R \bowtie S) = R \bowtie [\sigma_q (S)]$$

Derived Rules: $\sigma + \bowtie$ combined

$$\sigma_{p \wedge q} (R \bowtie S) = [\sigma_p (R)] \bowtie [\sigma_q (S)]$$

$$\sigma_{p \wedge q \wedge m} (R \bowtie S) =$$
$$\sigma_m \left[(\sigma_p R) \bowtie (\sigma_q S) \right]$$

$$\sigma_{p \vee q} (R \bowtie S) =$$
$$\left[(\sigma_p R) \bowtie S \right] \cup \left[R \bowtie (\sigma_q S) \right]$$

Rules: π, σ combined

Let x = subset of R attributes

z = attributes in predicate P
(subset of R attributes)

$$\pi_x[\sigma_p(R)] = \pi_x \left\{ \sigma_p \left[\overset{\pi_{xz}}{\cancel{\pi_x}}(R) \right] \right\}$$

Rules: σ , U combined

$$\sigma_p(R \cup S) = \sigma_p(R) \cup \sigma_p(S)$$

$$\sigma_p(R - S) = \sigma_p(R) - S = \sigma_p(R) - \sigma_p(S)$$

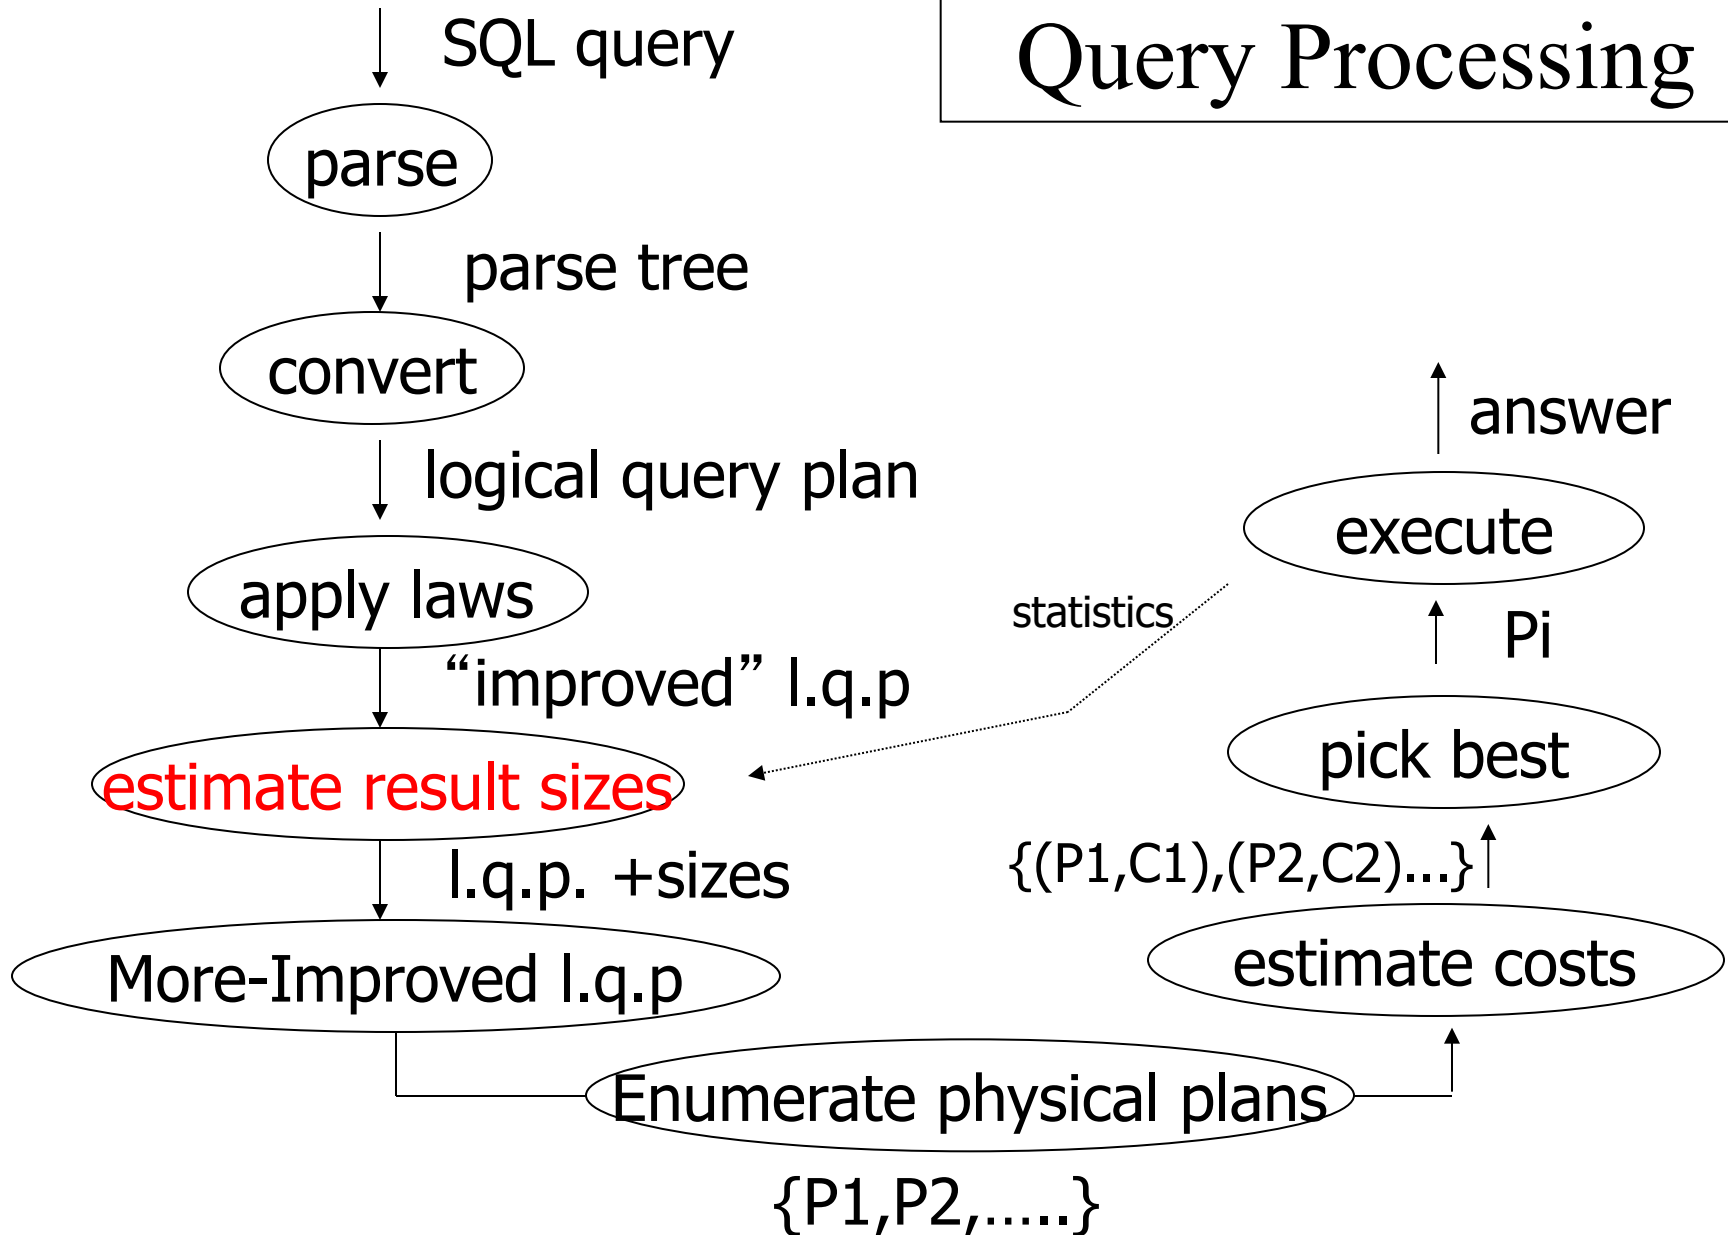
Which are “good” transformations?

- $\sigma_{p_1 \wedge p_2} (R) \rightarrow \sigma_{p_1} [\sigma_{p_2} (R)]$
- $\sigma_p (R \bowtie S) \rightarrow [\sigma_p (R)] \bowtie S$
- $R \bowtie S \rightarrow S \bowtie R$
- $\pi_x [\sigma_p (R)] \rightarrow \pi_x \{ \sigma_p [\pi_{xz} (R)] \}$

Bottom line:

- No transformation is always good
- Usually good: early selections

Query Processing



Estimating result size

- Keep statistics for relation R
 - $T(R)$: # **tuples** in R
 - $S(R)$: # **bytes** in each R tuple
 - $B(R)$: # **blocks** to hold all R tuples
 - $V(R, A)$: # **distinct values** in R for attribute A

Example

R

A	B	C	D
cat	1	10	a
cat	1	20	b
dog	1	30	a
dog	1	40	c
bat	1	50	d

A: 20 byte string

B: 4 byte integer

C: 8 byte date

D: 5 byte string

$$T(R) = 5 \quad S(R) = 37$$

$$V(R,A) = 3$$

$$V(R,C) = 5$$

$$V(R,B) = 1$$

$$V(R,D) = 4$$

Size estimates for $W = R1 \times R2$

$$T(W) = T(R1) \times T(R2)$$

$$S(W) = S(R1) + S(R2)$$

Size estimate for $W = \sigma_{A=a}(R)$

$$S(W) = S(R)$$

$$T(W) = ?$$

Example

R

A	B	C	D
cat	1	10	a
cat	1	20	b
dog	1	30	a
dog	1	40	c
bat	1	50	d

$$V(R,A)=3$$

$$V(R,B)=1$$

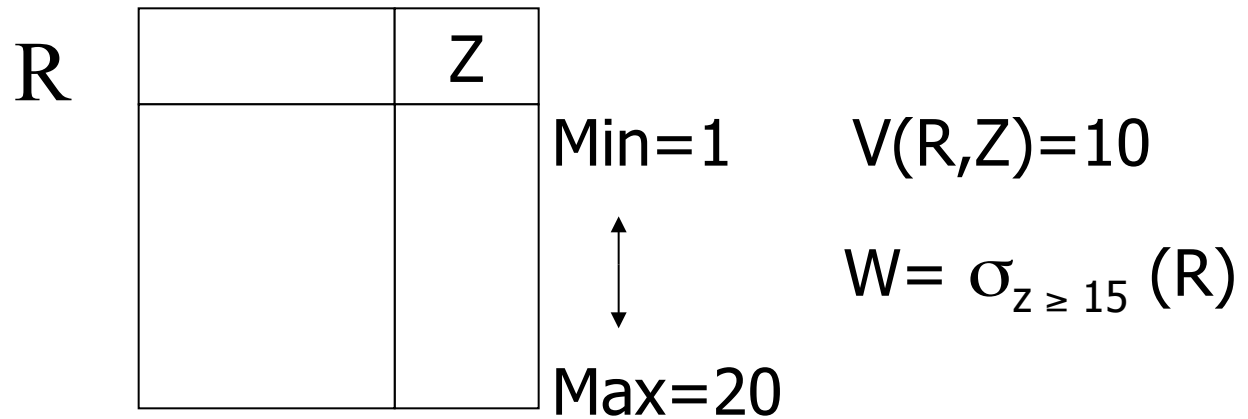
$$V(R,C)=5$$

$$V(R,D)=4$$

$$W = \sigma_{z=\text{val}}(R); \quad T(W) = \frac{T(R)}{V(R,Z)}$$

(under assumption of uniform distribution of values of
Z over $V(R,Z)$)

$$W = \sigma_{z \geq \text{val}} (R). \quad T(W)?$$



$$f = \frac{20-15+1}{20} = \frac{6}{20} \quad (\text{fraction of range})$$

$$T(W) = f \times T(R)$$

Size estimate for $W = R1 \bowtie R2$

Let x = attributes of $R1$

y = attributes of $R2$

Case 1

$$\underline{X \cap Y = \emptyset}$$

Same as $R1 \times R2$

Case 2

$$\underline{W = R1 \bowtie R2 \quad X \cap Y = A}$$


R1	A	B	C	R2	A	D

Assumption (containment of values):

$V(R1,A) \leq V(R2,A) \Rightarrow$ Every A value in R1 is in R2

$V(R2,A) \leq V(R1,A) \Rightarrow$ Every A value in R2 is in R1

Computing $T(W)$ when $V(R1,A) \leq V(R2,A)$



1 tuple matches with $\frac{T(R2)}{V(R2,A)}$ tuples...

$$\text{so } T(W) = \frac{T(R2) \times T(R1)}{V(R2, A)}$$

- $V(R1,A) \leq V(R2,A) \quad T(W) = \frac{T(R2) T(R1)}{V(R2,A)}$

- $V(R2,A) \leq V(R1,A) \quad T(W) = \frac{T(R2) T(R1)}{V(R1,A)}$

[A is common attribute]

In general, $W = R1 \bowtie R2$

$$T(W) = \frac{T(R2) T(R1)}{\max \{ V(R1,A), V(R2,A) \}}$$

In all cases:

$$S(W) = S(R1) + S(R2) - S(A) \leftarrow \begin{array}{l} \text{size of attribute A} \end{array}$$

Similarly, we can estimate sizes of:

$\Pi_{AB} (R)$; $\sigma_{A=a \wedge B=b} (R)$;

$R \bowtie S$ with common attributes. A, B, C ;

Union, intersection, diff, Sec. 16.4.7

Note: for complex expressions, need intermediate T,S,V results.

$$\text{E.g. } W = [\underbrace{\sigma_{A=a}(R1)}] \bowtie R2$$

Treat as **relation U**

$$T(U) = T(R1)/V(R1,A)$$

$$S(U) = S(R1)$$

$$V(U,*) = ?? \quad (\text{Needed for } T(W))$$

To estimate V_s

E.g., $U = \sigma_{A=a}(R1)$

Say R1 has attributes A,B,C,D

$$V(U, A) = 1 \quad (\text{exact})$$

$$V(U, B) = V(R1, B) \quad (\text{guess})$$

$$V(U, C) = V(R1, C) \quad (\text{guess})$$

$$V(U, D) = V(R1, D) \quad (\text{guess})$$

For Joins $U = R1(A,B) \bowtie R2(A,C)$

$$V(U,A) = \min \{ V(R1, A), V(R2, A) \}$$

$$V(U,B) = V(R1, B)$$

$$V(U,C) = V(R2, C)$$

[called “preservation of value sets” assumption]

Example:

$$Z = R1(A,B) \bowtie R2(B,C) \bowtie R3(C,D)$$

R1	T(R1) = 1000	V(R1,A)=50	V(R1,B)=100
R2	T(R2) = 2000	V(R2,B)=200	V(R2,C)=300
R3	T(R3) = 3000	V(R3,C)=90	V(R3,D)=500

Partial Result: $U = R \bowtie S$

$$T(U) = \frac{1000 \times 2000}{200}$$

$$V(U,A) = 50$$

$$V(U,B) = 100$$

$$V(U,C) = 300$$

$$\underline{Z = U \bowtie R3}$$

$$T(Z) = \frac{1000 \times 2000 \times 3000}{200 \times 300}$$

$$V(Z,A) = 50$$

$$V(Z,B) = 100$$

$$V(Z,C) = 90$$

$$V(Z,D) = 500$$

Recap

- Estimating size of results is an “art”
- Computing Statistics: Scan, Sample.
- Also, statistics must be kept up to date...
- **Other Statistics:**
 - Histograms (equi-height, equi-width, most-freq.)
- **Next:** Using size estimates to improve an LQP.

Improve LPQ with Size

- Size estimates can also help improve the LQP.
 - But, need some heuristic to estimate costs.
 - Good heuristic: Cost = Sum of intermediate-result sizes.
 - **Next Slide:** Example.

Ex: Improve LQP with Sizes

- $R(a,b)$: $T(R)=5000$, $V(R,a)=50$, $V(R,b)=100$
- $S(b,c)$: $T(S) = 2000$, $V(S,b)=200$, $V(S,c)=100$.

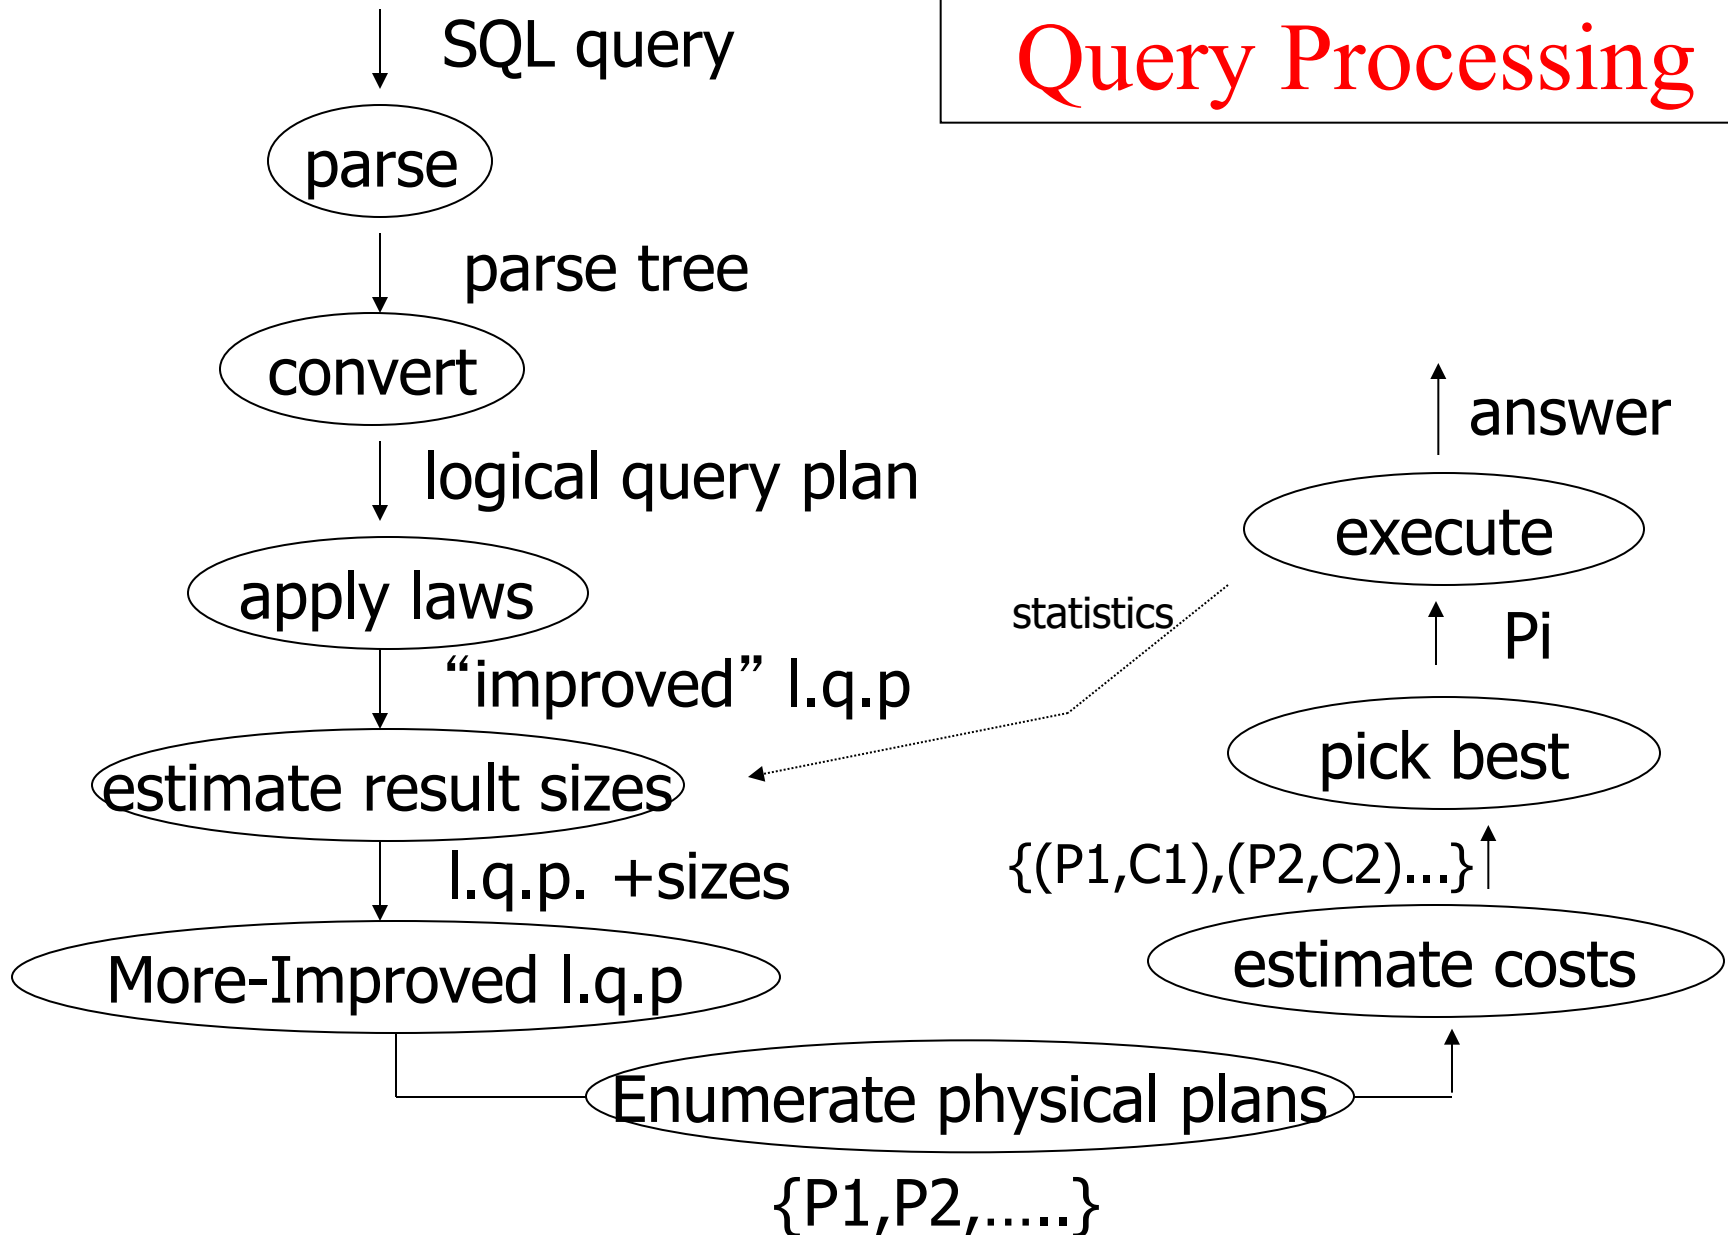
$\delta(\sigma_p(R \text{ JOIN } S))$. Where p is $(R.a=20)$.

Two choices:

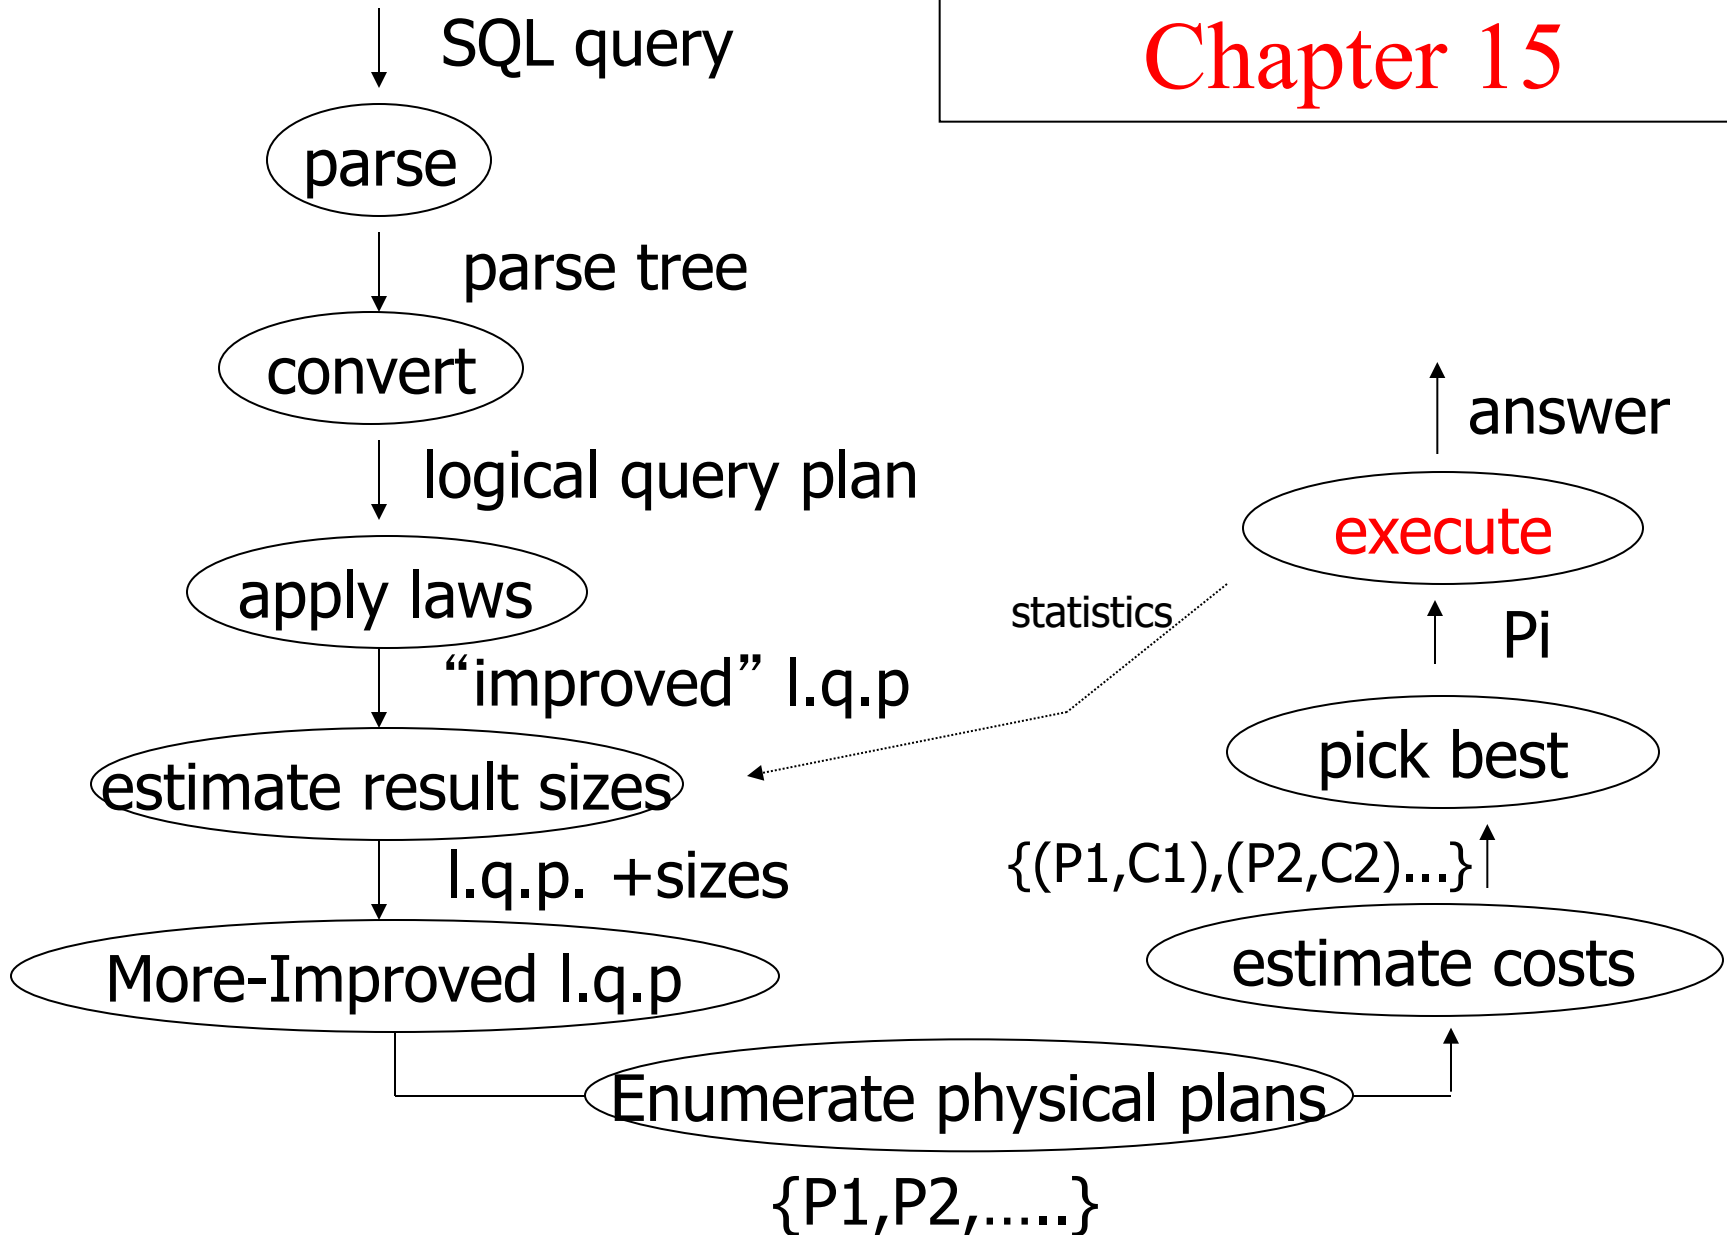
1. $\delta(\sigma_p(R)) \text{ JOIN } (\delta(S))$
2. $\delta(\sigma_p(R \text{ JOIN } S))$

Cost Model = Sum of the sizes of intermediate results. [1100 vs 1150]

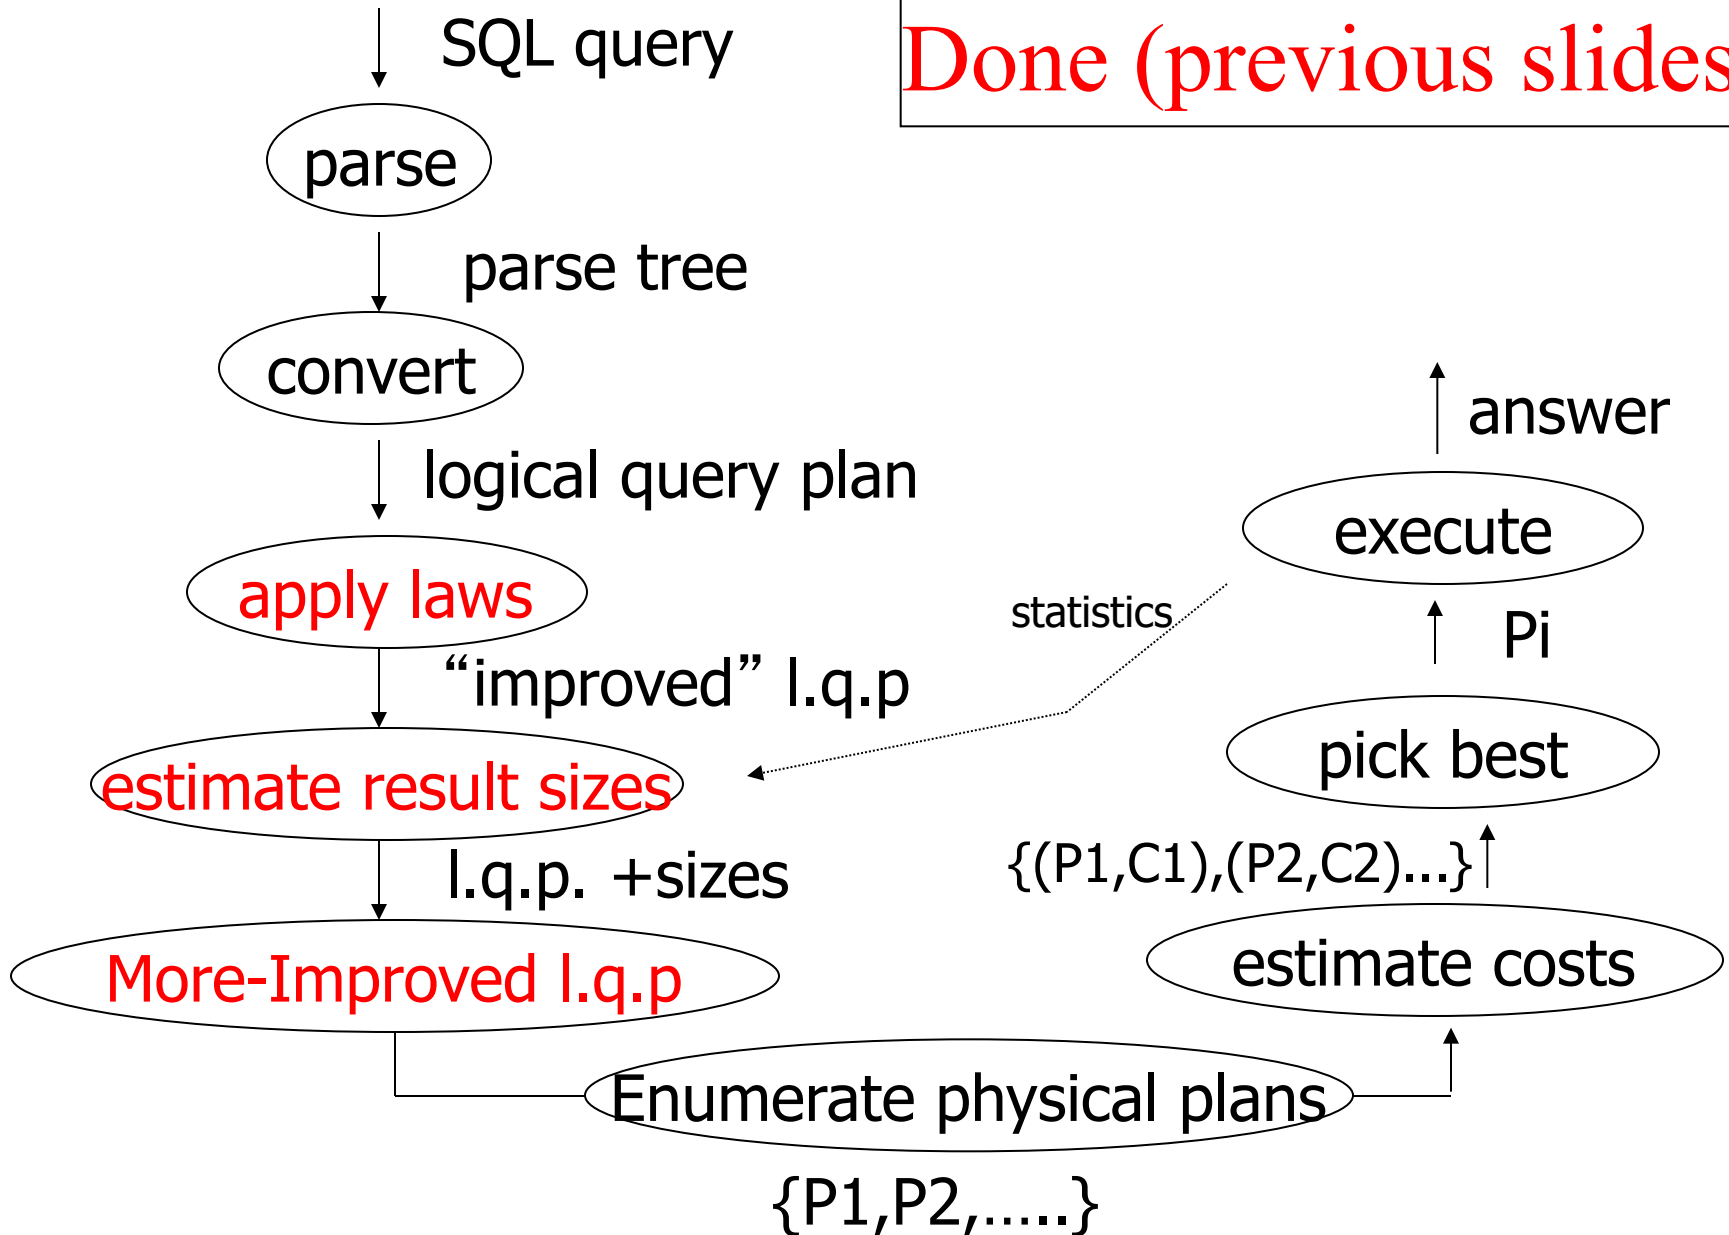
Query Processing



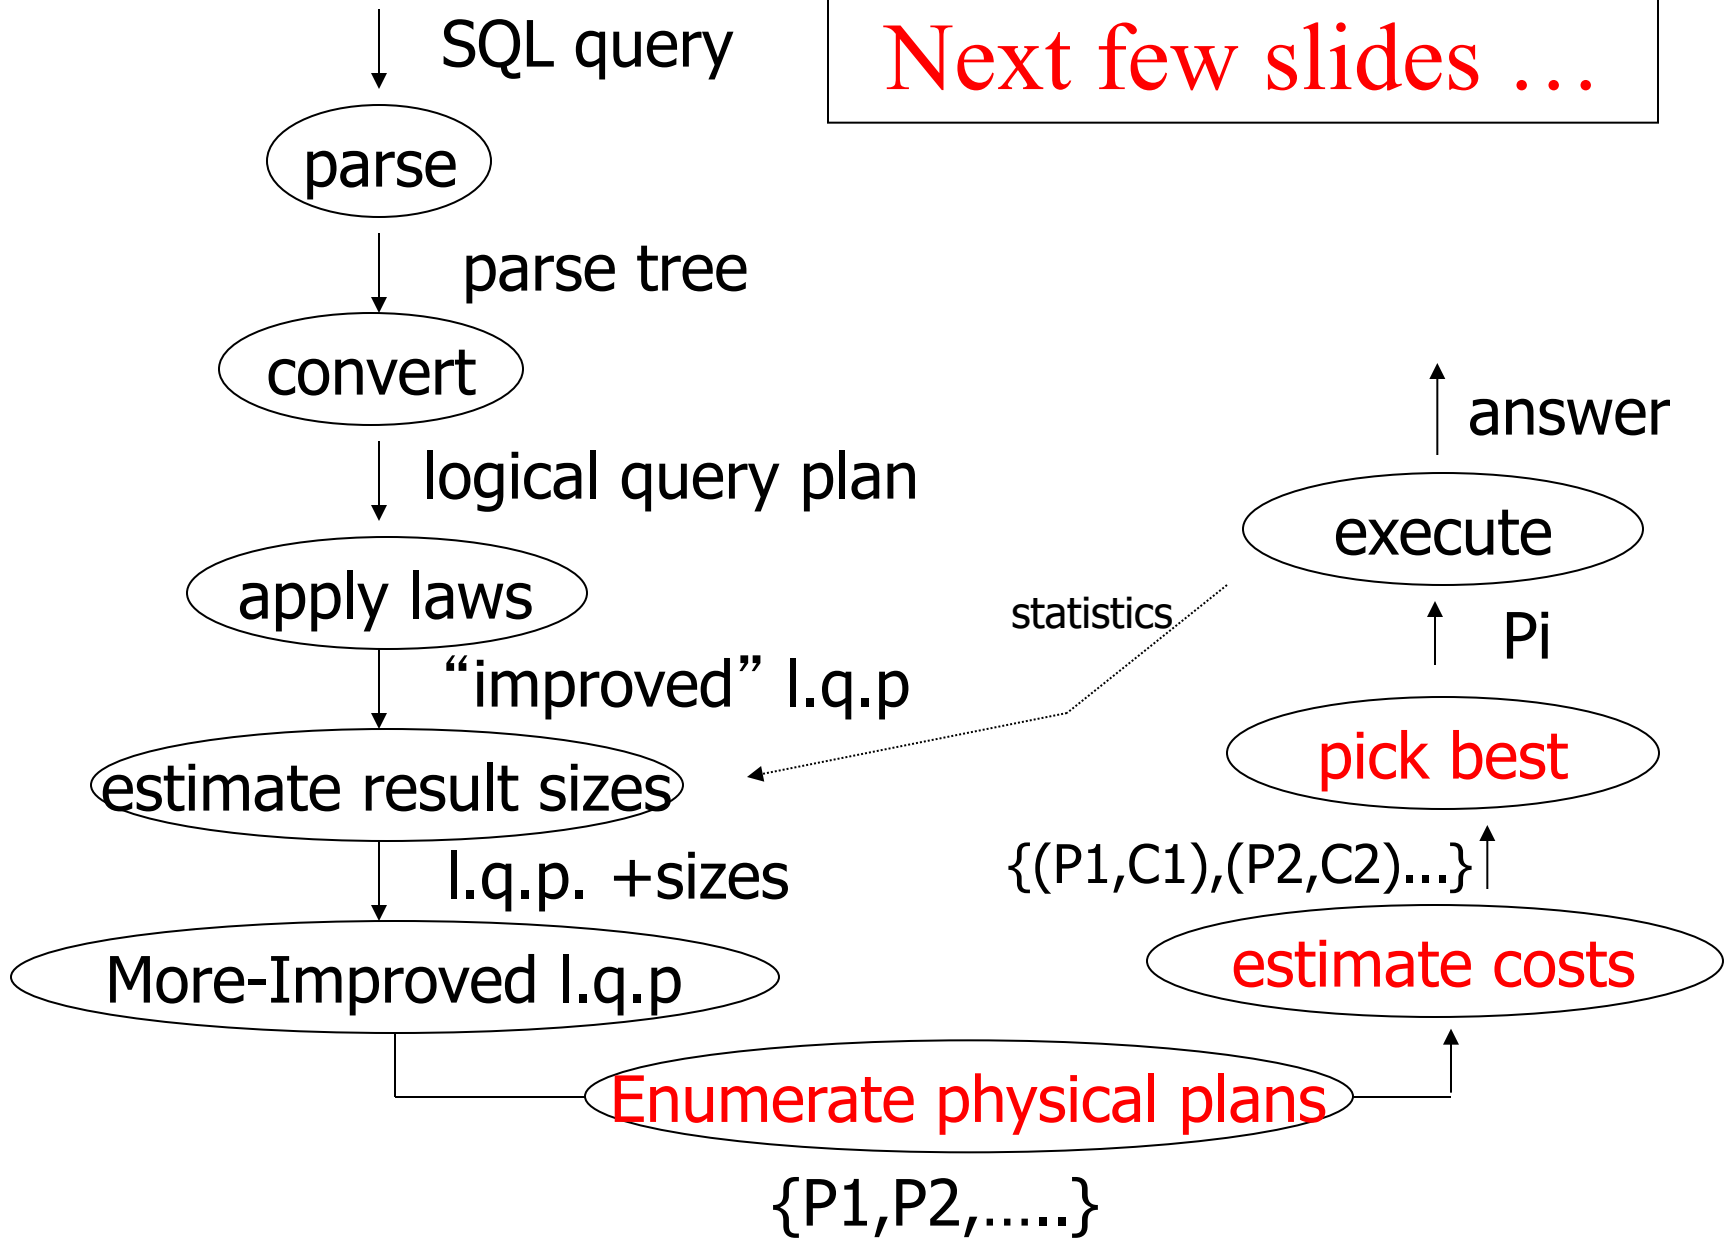
Chapter 15



Done (previous slides)



Next few slides ...




LQP vs. PQP

- LQP is a high-level expression-tree.
- PQP is (LQP + more execution details such as):
 - Order/grouping of join, union, intersection, etc.
 - Choice of join algorithm, etc.
 - Additional operators such as scanning, sorting, etc.
 - Intermediate steps between two operations. E.g., sorting, storage, pipelining, etc.

LQP → PQP

- Generating and comparing physical plans



PQP Enumeration Heuristics

1. Various techniques (details skipped):
 - Top-down
 - Bottom-up
 - Heuristics
 - Branch and Bound
 - Hill Climbing
 - Dynamic Programming
 - Selinger-style Optimization (Improved DP)

Join Ordering

- $R_1 \times R_2 \times R_3 \times \dots \times R_n$ (join of n tables).
- How many possible orderings?
- Join-selectivity.
- Left-deep, Right-deep trees.
- Dynamic Programming approach.

PQP Selection Example

Pipelining vs. Materialization

- $(R \text{ join } S) \text{ join } U$
- 5000, 10000, 10000 blocks respectively.
- $(R \text{ join } S)$ is of size k (we'll consider different k)
- We'll only use hash-joins
- Memory size: 101 blocks.

Example (contd)

- (R join S) join U.

R join S

- 100 buckets at most.
- R-buckets have 50 tuples each. So, need 51 buffers for the second pass.
- Pipeline: Use remaining 50 buffers to join the result with U.
 - If $k < 49$, then keep (R join S) in memory, and read U one block at a time.

Example (continued)

- If $k < 49$, then read U one block at a time, and do everything in main memory. Cost = $55k$.
- If $k > 49$, but < 5000 :
 - First, hash all tables. Hash U with 50 buckets.
 - When doing R join S , use the remaining 50 buffers to “hash” $(R$ join $S)$ result into the 50 buckets.
 - Each bucket of $(R$ join $S)$ is of size 100.
 - Then, do two-pass hash join of $(R$ join $S)$ with U .
 - Cost = $50000 + 15000 + (k + (10000 + k))$

Example – contd.

- If $k > 5000$:
 - Last step will require a 3-pass join, since each bucket of $(R \text{ join } S)$ as well as U is more than 100 blocks (buckets of U are 200 blocks each).
 - **Better:** No pipelining. Write $(R \text{ join } S)$. Then, do 2-pass hash-join with U (now the buckets of U are 100 blocks each).

$$\text{Cost} = 45000 + k + 3(10000 + k) = 75000 + 4k$$