

CS 525: Advanced Database Organisation



07: Query Processing Overview

Boris Glavic

Slides: adapted from a [course](#) taught by [Hector Garcia-Molina](#), Stanford InfoLab



Query Processing

Q \rightarrow Query Plan



Query Processing

Q \rightarrow Query Plan

Focus: Relational System

- Others?



Example

Select B,D

From R,S

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

R	A	B	C
a	1	10	
b	1	20	
c	2	10	
d	2	35	
e	3	45	

S	C	D	E
	10	x	2
	20	y	2
	30	z	2
	40	x	1
	50	y	3

R	A	B	C	S	C	D	E
a	1	10			10	x	2
b	1	20			20	y	2
c	2	10			30	z	2
d	2	35			40	x	1
e	3	45			50	y	3

Answer

B	D
2	x

- How do we execute query?

One idea

- Do Cartesian product
- Select tuples
- Do projection




RXS	R.A	R.B	R.C	S.C	S.D	S.E
a	1	10	10	x	2	
a	1	10	20	y	2	
.						
.						
C	2	10	10	x	2	
.						
.						

RXS	R.A	R.B	R.C	S.C	S.D	S.E
a	1	10	10	x	2	
a	1	10	20	y	2	
.						
.						
Bingo! →	C	2	10	10	x	2
Got one...	.					
.						


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

Ex: Plan I



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



Ex: Plan I



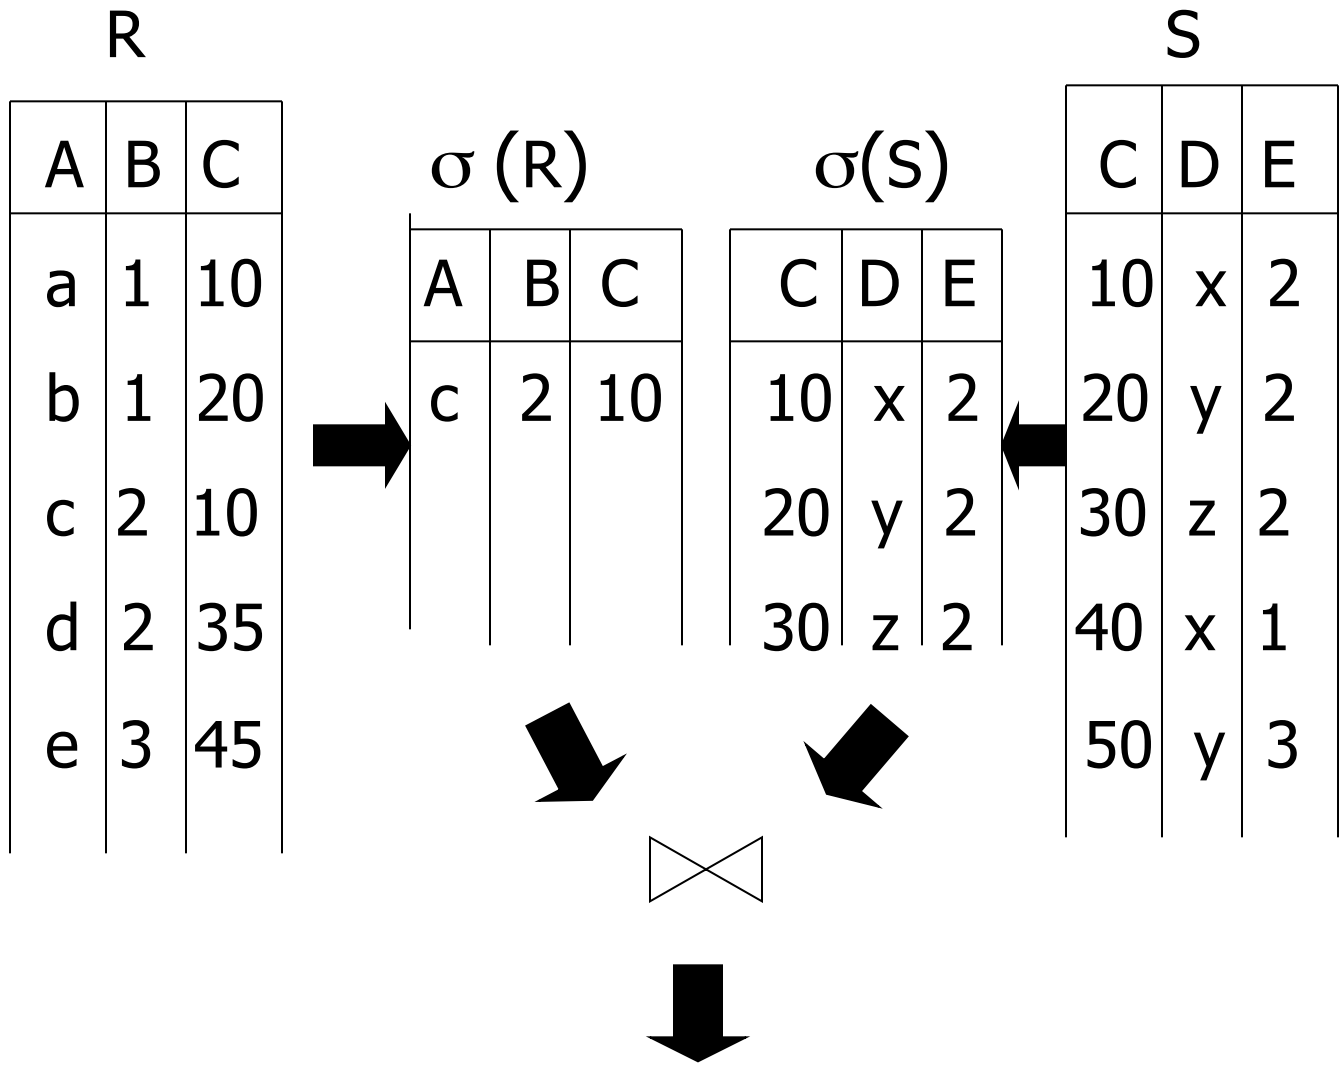
OR: $\Pi_{B,D} [\sigma_{R.A="C" \wedge S.E=2 \wedge R.C=S.C} (RXS)]$

Another idea:

Plan II



natural join



Plan III

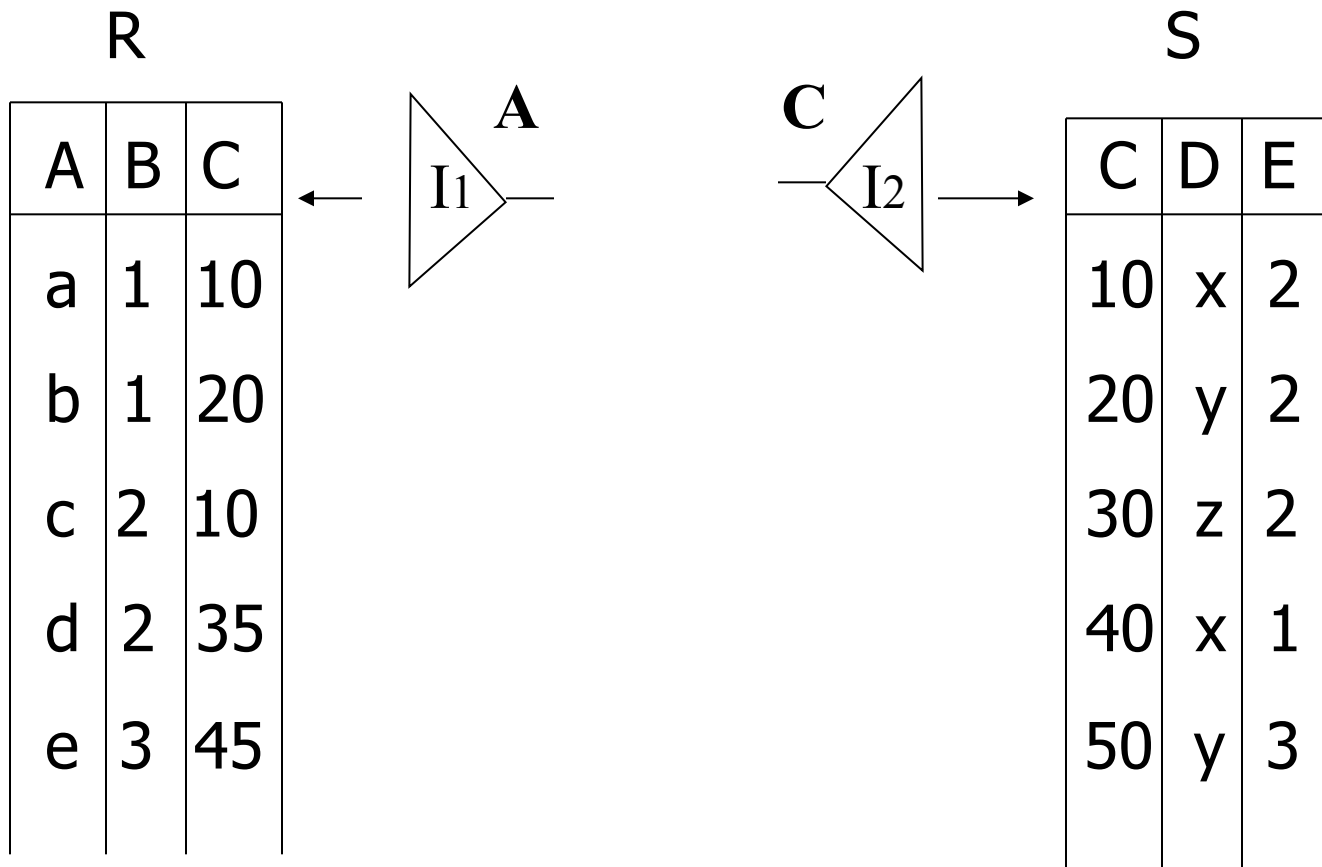
Use R.A and S.C Indexes


- (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


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



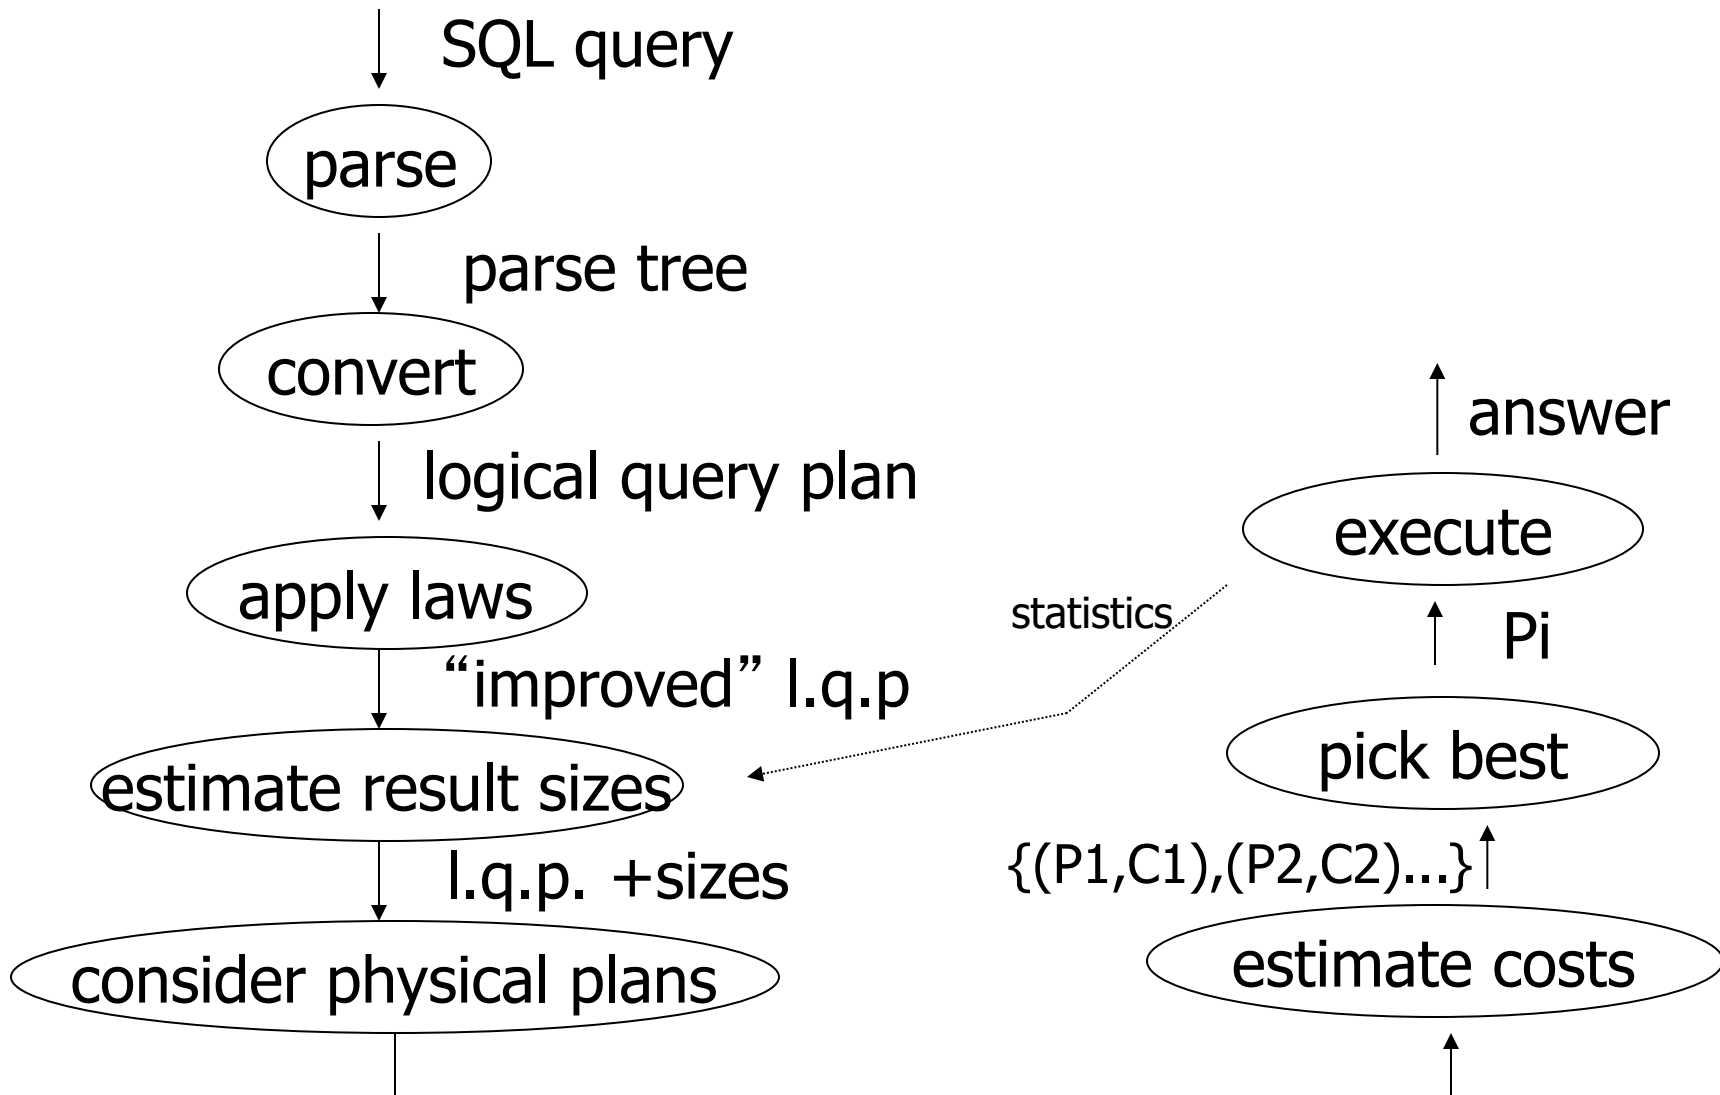








Overview of Query Optimization




Example: SQL query

```
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

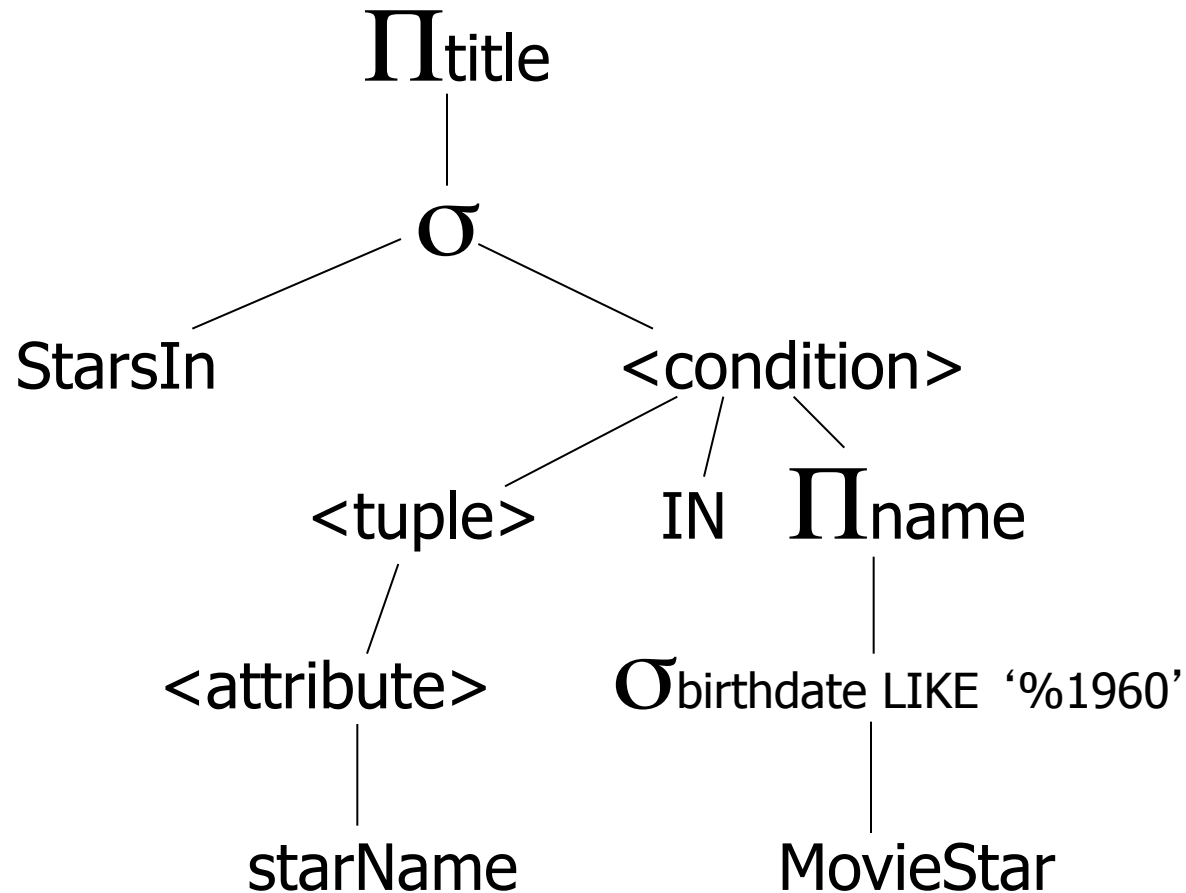


Fig. 7.15: An expression using a two-argument σ , midway between a parse tree and relational algebra

Example: Logical Query Plan





Fig. 7.18: Applying the rule for IN conditions


Example: Improved Logical Query Plan



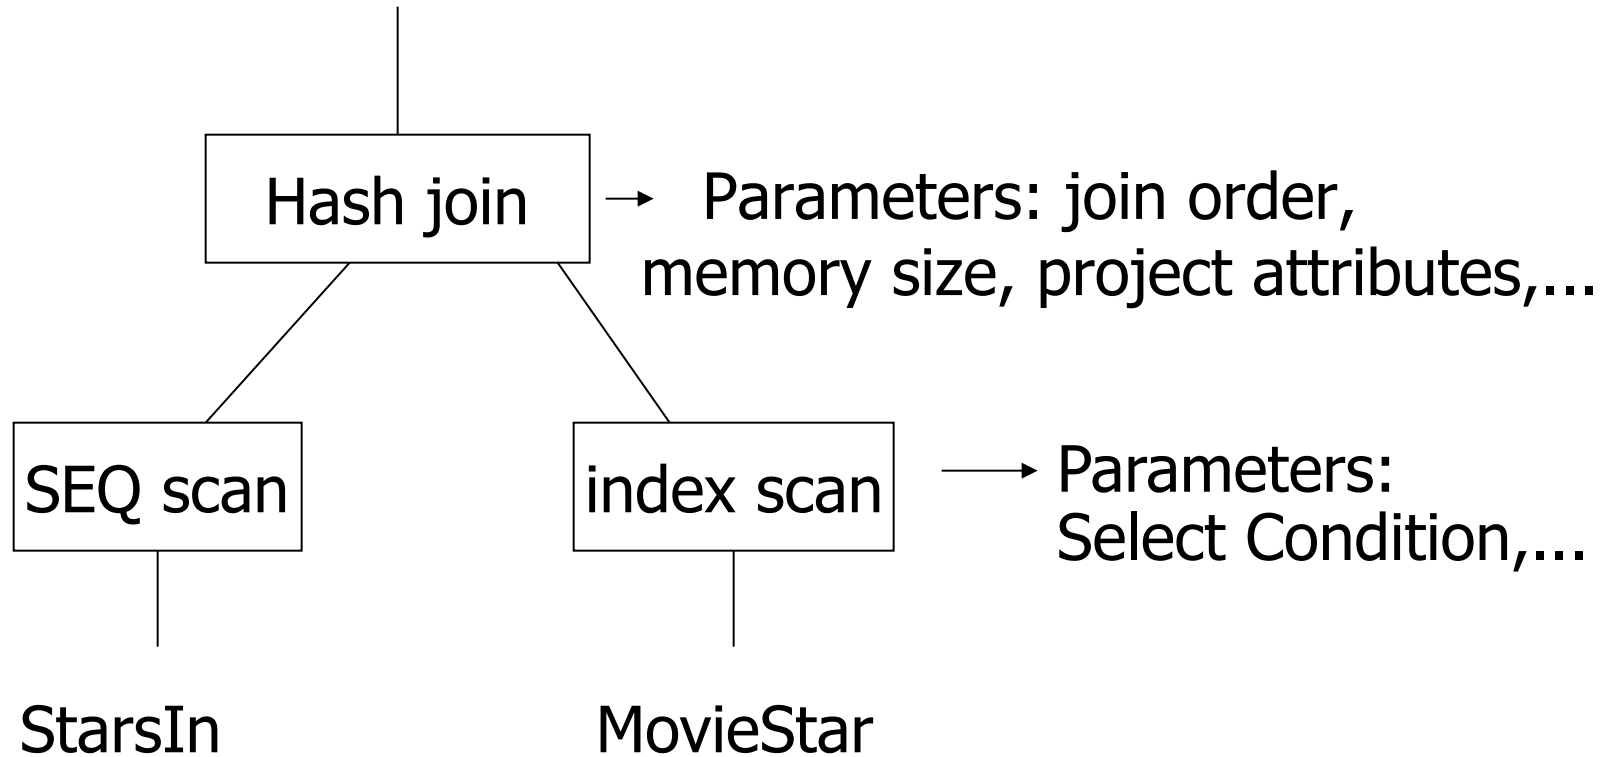
Question:
Push project to
StarsIn?

Fig. 7.20: An improvement on fig. 7.18.

Example: Estimate Result Sizes



Example: One Physical Plan



Example: Estimate costs

