



Carnegie Mellon Univ.
Dept. of Computer Science
15-415 - Database Applications

Query Optimization
(R&G ch. 15; Sys. R q-opt paper)



Overview - detailed

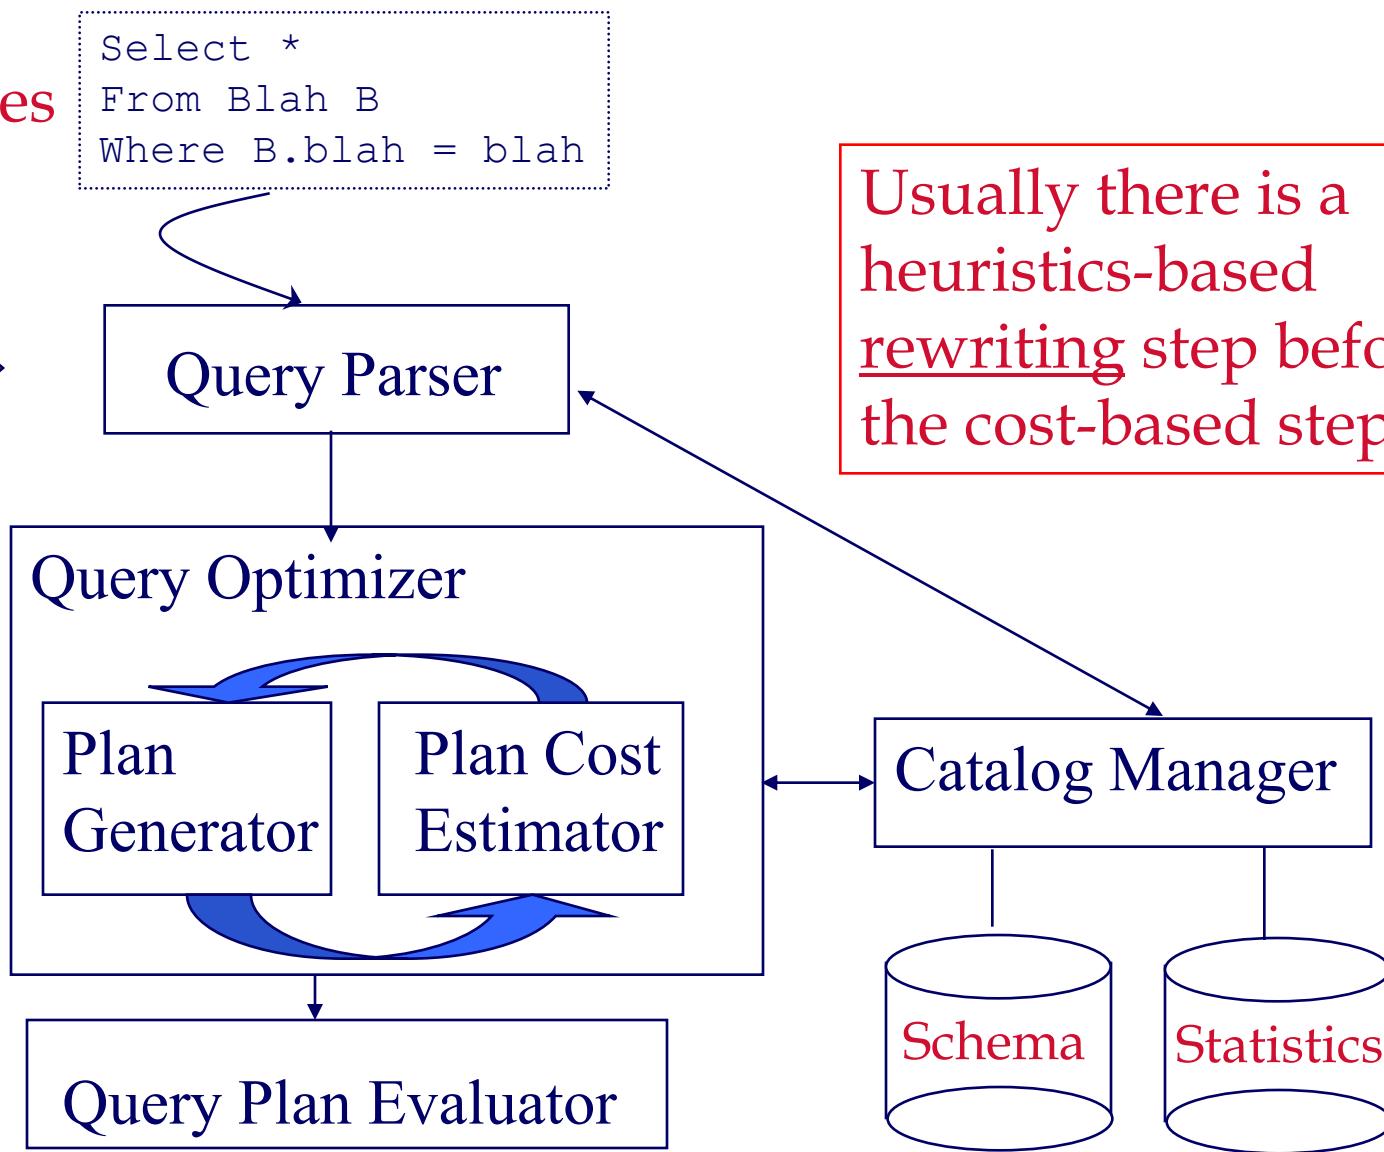
- Why q-opt?
- Equivalence of expressions
- Cost estimation
- Plan generation
- Plan evaluation



Cost-based Query Sub-System

Queries

```
Select *\nFrom Blah B\nWhere B.blah = blah
```





Why Q-opt?

- SQL: ~declarative
- good q-opt \rightarrow big difference
 - eg., seq. Scan vs
 - B-tree index, on $P=1,000$ pages



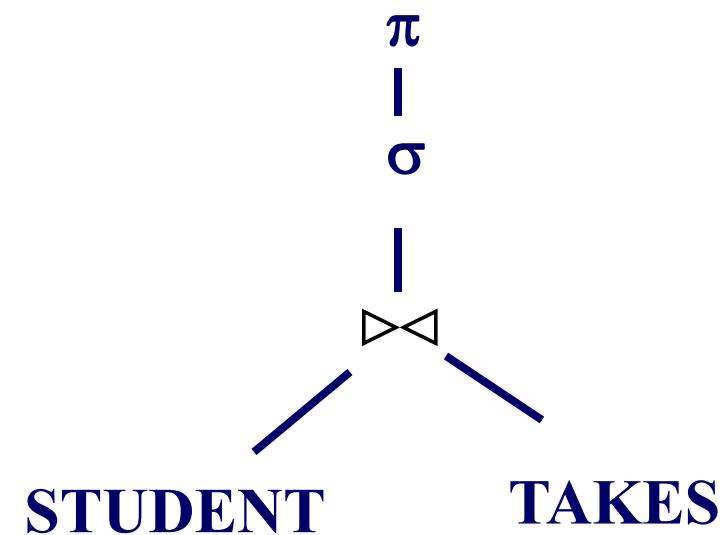
Q-opt steps

- bring query in internal form (eg., parse tree)
- ... into ‘canonical form’ (syntactic q-opt)
- generate alt. plans
- estimate cost; pick best



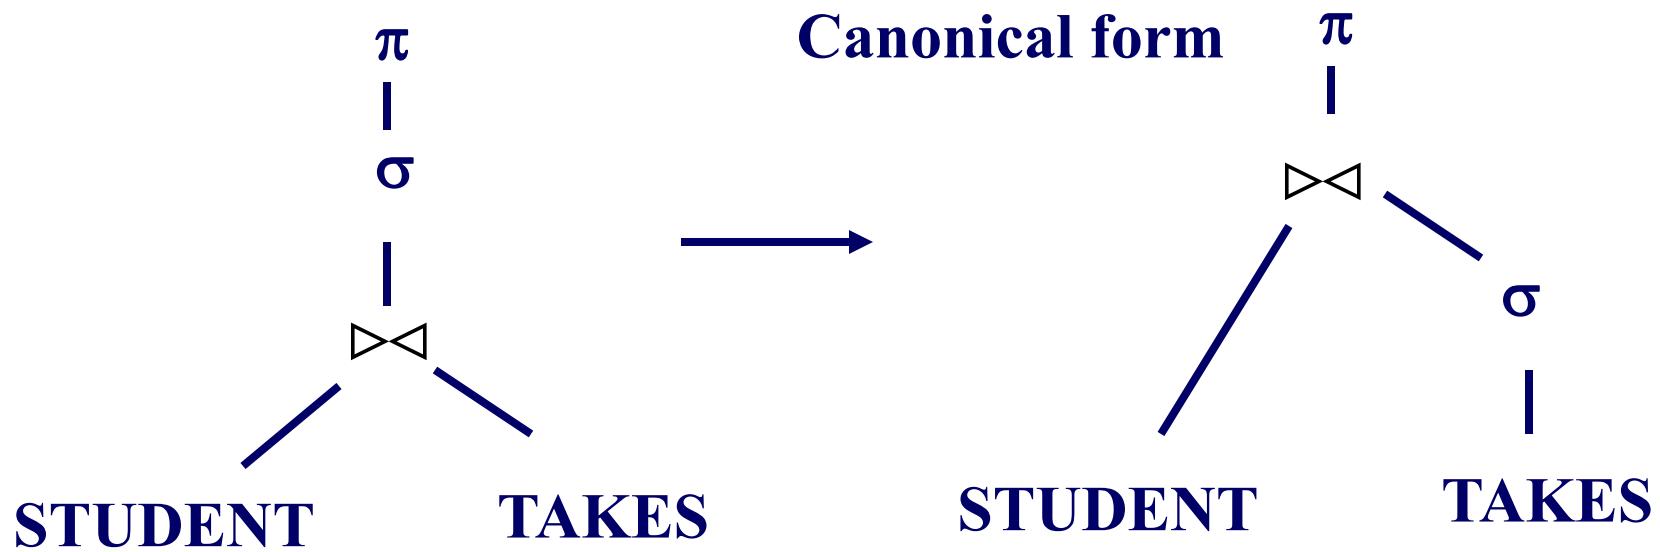
Q-opt - example

```
select name  
from STUDENT, TAKES  
where c-id='415' and  
STUDENT.ssn=TAKES.ssn
```



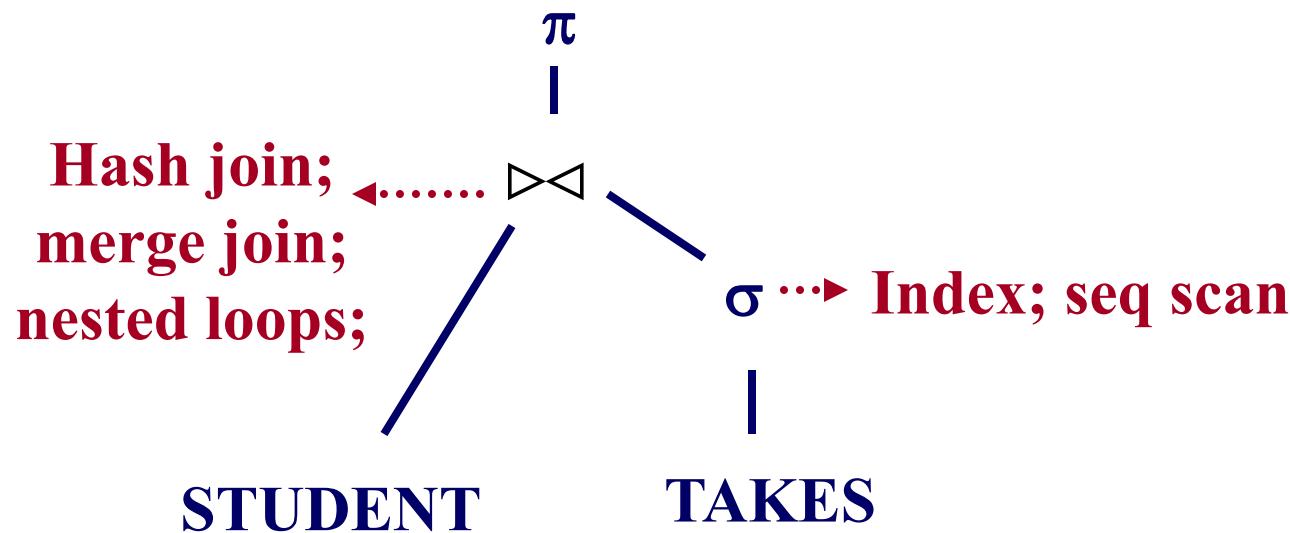


Q-opt - example





Q-opt - example





Overview - detailed

- Why q-opt?
- **Equivalence of expressions**
- Cost estimation
- ...



Equivalence of expressions

- A.k.a.: syntactic q-opt
- in short: perform selections and projections early
- More details: see transf. rules in text



Equivalence of expressions

- Q: How to prove a transf. rule?

$$\sigma_P(R1 \bowtie R2) \stackrel{?}{=} \sigma_P(R1) \bowtie \sigma_P(R2)$$

- A: use RTC, to show that LHS = RHS, eg:

$$\sigma_P(R1 \cup R2) \stackrel{?}{=} \sigma_P(R1) \cup \sigma_P(R2)$$



Equivalence of expressions

$$\sigma_P(R1 \cup R2) \stackrel{?}{=} \sigma_P(R1) \cup \sigma_P(R2)$$

$t \in LHS \Leftrightarrow$

$$t \in (R1 \cup R2) \wedge P(t) \Leftrightarrow$$

$$(t \in R1 \vee t \in R2) \wedge P(t) \Leftrightarrow$$

$$(t \in R1 \wedge P(t)) \vee (t \in R2 \wedge P(t)) \Leftrightarrow$$



Equivalence of expressions

$$\sigma_P(R1 \cup R2) \stackrel{?}{=} \sigma_P(R1) \cup \sigma_P(R2)$$

...

$$(t \in R1 \wedge P(t)) \vee (t \in R2 \wedge P(t)) \Leftrightarrow$$

$$(t \in \sigma_P(R1)) \vee (t \in \sigma_P(R2)) \Leftrightarrow$$

$$t \in \sigma_P(R1) \cup \sigma_P(R2) \Leftrightarrow$$

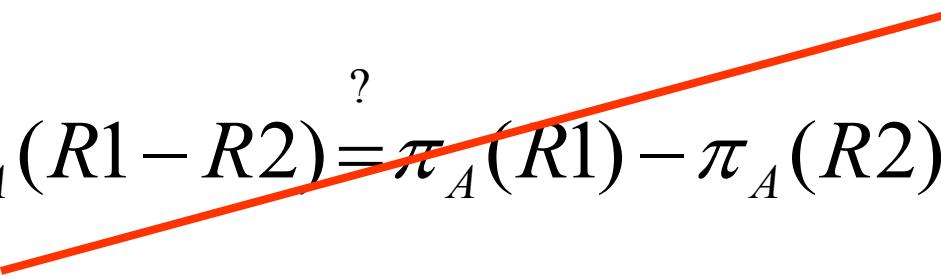
$$t \in RHS$$

QED



Equivalence of expressions

- Q: how to disprove a rule??

$$\pi_A(R1 - R2) \stackrel{?}{=} \pi_A(R1) - \pi_A(R2)$$




Equivalence of expressions

- Selections
 - perform them early
 - break a complex predicate, and push
$$\sigma_{p_1 \wedge p_2 \wedge \dots \wedge p_n}(R) = \sigma_{p_1}(\sigma_{p_2}(\dots \sigma_{p_n}(R))\dots)$$
 - simplify a complex predicate
 - ('X=Y and Y=3') -> 'X=3 and Y=3'



Equivalence of expressions

- Projections
 - perform them early (but carefully...)
 - Smaller tuples
 - Fewer tuples (if duplicates are eliminated)
 - project out all attributes except the ones requested or required (e.g., joining attr.)



Equivalence of expressions

- Joins
 - Commutative , associative

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

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

- Q: n-way join - how many diff. orderings?



Equivalence of expressions

- Joins - Q: n-way join - how many diff. orderings?
- A: Catalan number $\sim 4^n$
 - Exhaustive enumeration: too slow.



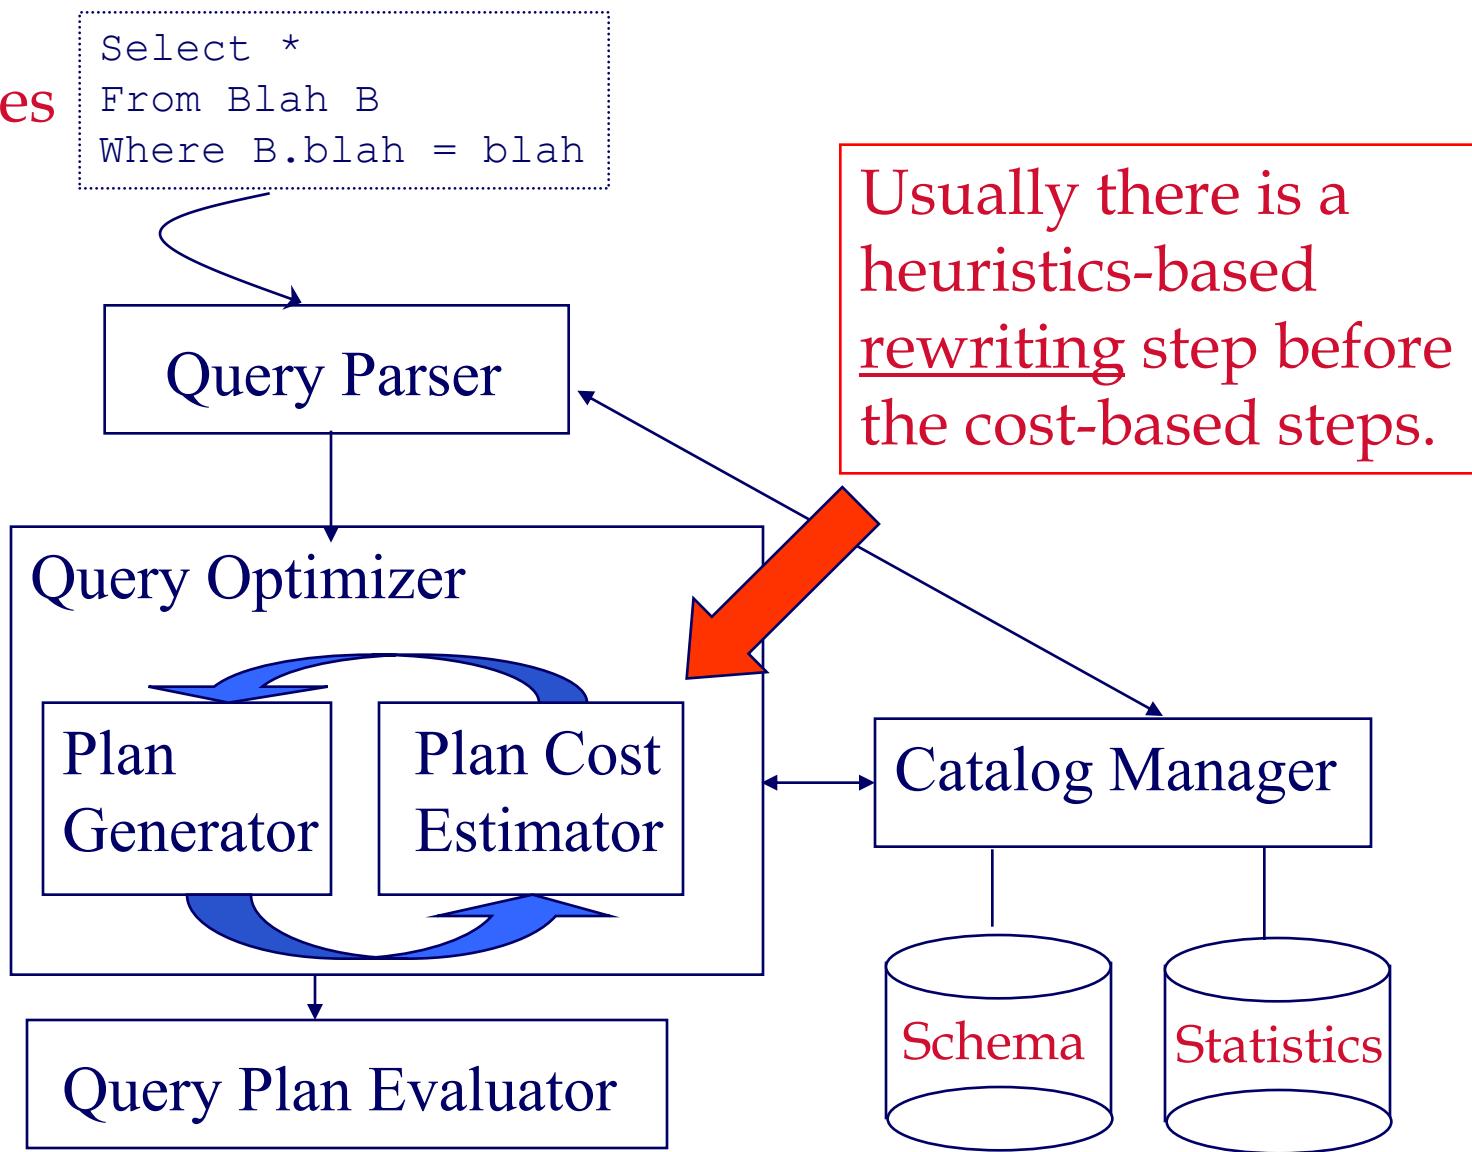
Q-opt steps

- bring query in internal form (eg., parse tree)
- ... into ‘canonical form’ (syntactic q-opt)
- generate alt. plans
- **estimate cost**; pick best



Cost-based Query Sub-System

Queries





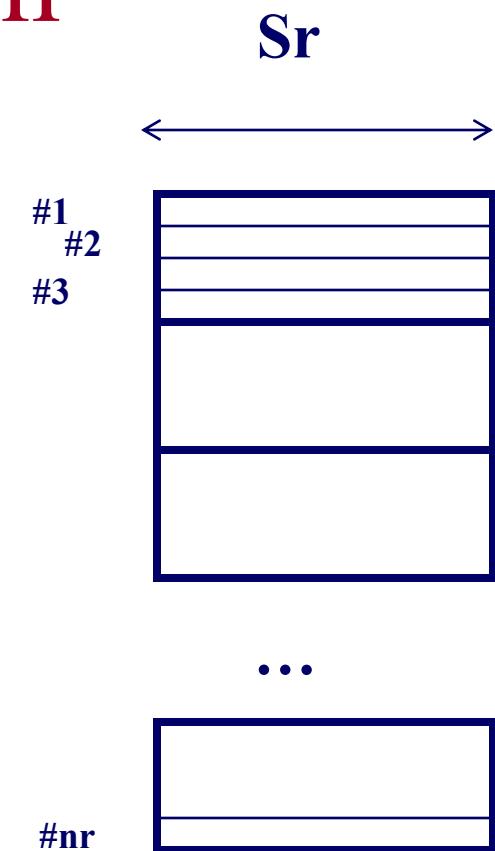
Cost estimation

- Eg., find ssn's of students with an 'A' in 415 (using seq. scanning)
- How long will a query take?
 - CPU (but: small cost; decreasing; tough to estimate)
 - Disk (mainly, # block transfers)
- How many tuples will qualify?
- (what statistics do we need to keep?)



Cost estimation

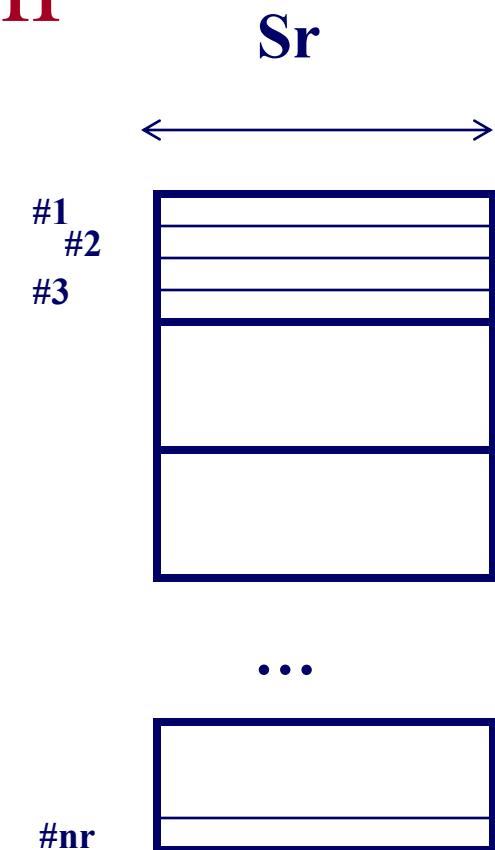
- Statistics: for each relation ‘r’ we keep
 - nr : # tuples;
 - Sr : size of tuple in bytes





Cost estimation

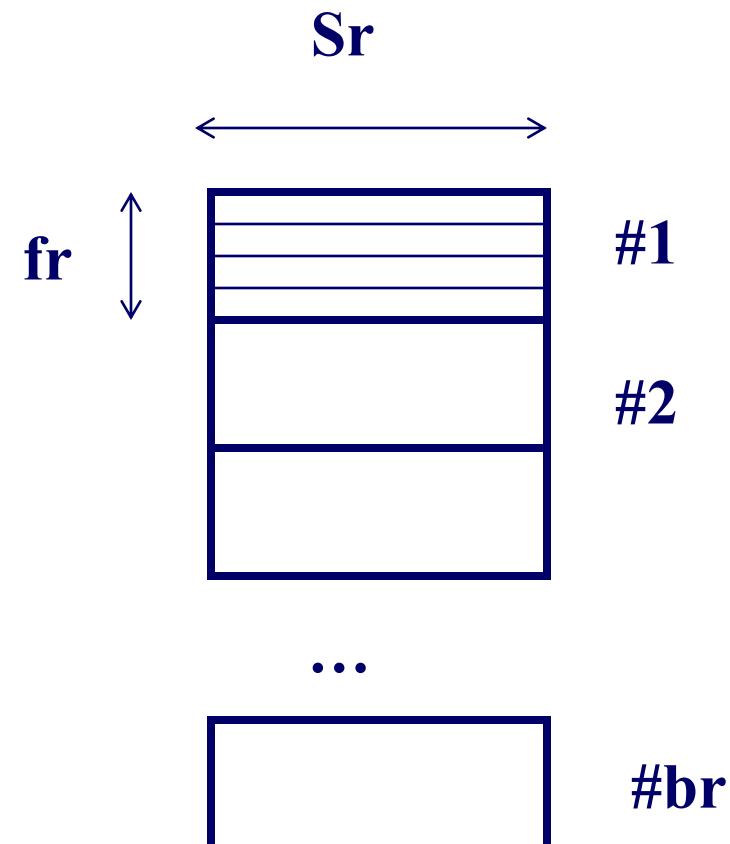
- Statistics: for each relation ‘r’ we keep
 - ...
 - $V(A,r)$: number of distinct values of attr. ‘A’
 - (recently, histograms, too)





Derivable statistics

- blocking factor =
max# records/block
 $(=??$)
- br: # blocks $(=??$)
- $SC(A,r)$ = selection cardinality = avg# of records with $A=given$
 $(=??$)





Derivable statistics

- blocking factor = max# records/block (= B/S_r ; B : block size in bytes)
- br : # blocks (= $nr / (\text{blocking-factor})$)



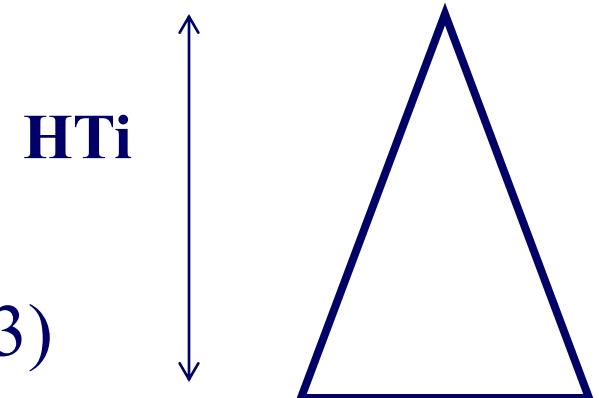
Derivable statistics

- $SC(A,r) = \text{selection cardinality} = \text{avg\# of records with } A=\text{given}$ ($= nr / V(A,r)$)
(assumes uniformity...) – eg: 10,000 students, 10 colleges – how many students in SCS?



Additional quantities we need:

- For index ‘i’:
 - f_i : average fanout ($\sim 50-100$)
 - HT_i : # levels of index ‘i’ ($\sim 2-3$)
 - $\sim \log(\#entries)/\log(f_i)$
 - LBi : # blocks at leaf level





Statistics

- Where do we store them?
- How often do we update them?



Q-opt steps

- bring query in internal form (eg., parse tree)
- ... into ‘canonical form’ (syntactic q-opt)
- generate alt. plans
- **estimate cost**; pick best



Selections

- we saw simple predicates (A=constant; eg., ‘name=Smith’)
- how about more complex predicates, like
 - ‘salary > 10K’
 - ‘age = 30 **and** job-code=“analyst” ’
- what is their selectivity?



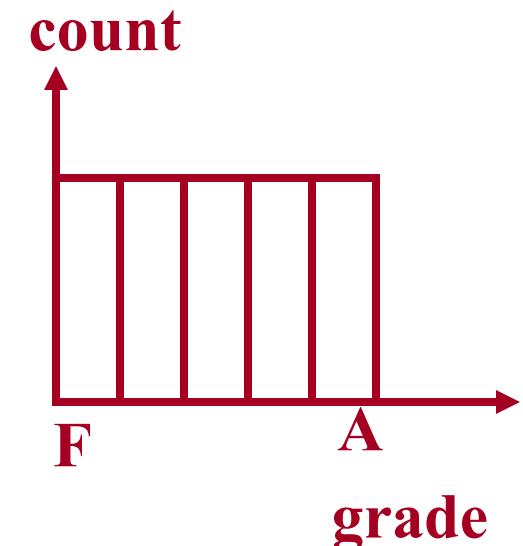
Selections – complex predicates

- selectivity $\text{sel}(P)$ of predicate P :
 == fraction of tuples that qualify
$$\text{sel}(P) = \text{SC}(P) / \text{nr}$$



Selections – complex predicates

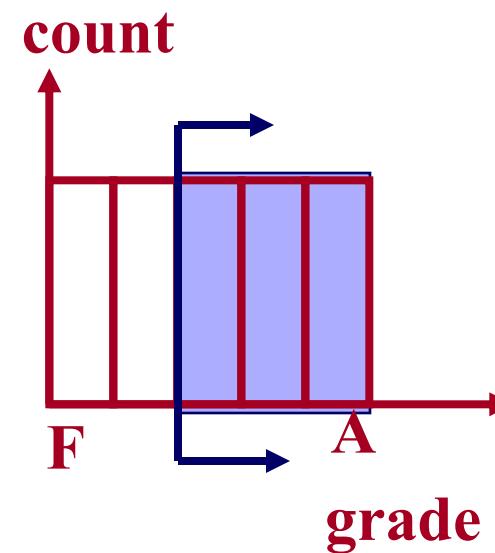
- eg., assume that $V(\text{grade}, \text{TAKES})=5$ distinct values
- simple predicate P : $A=\text{constant}$
 - $\text{sel}(A=\text{constant}) = 1/V(A,r)$
 - eg., $\text{sel}(\text{grade}='B') = 1/5$
- (what if $V(A,r)$ is unknown??)





Selections – complex predicates

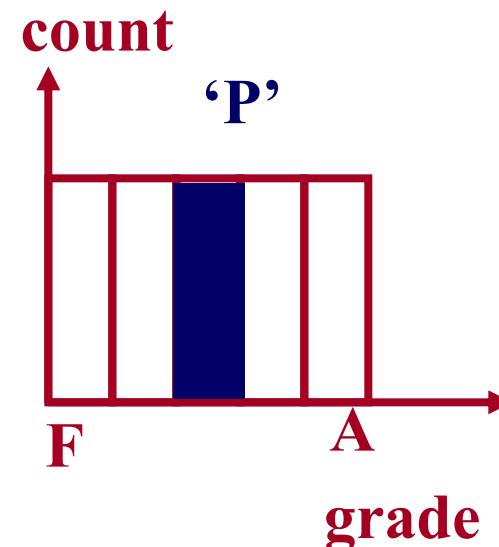
- range query: $\text{sel}(\text{ grade} \geq \text{'C'})$
 - $\text{sel}(A > a) = (A_{\max} - a) / (A_{\max} - A_{\min})$





Selections - complex predicates

- negation: $\text{sel}(\text{ grade } \neq \text{'C'})$
 - $\text{sel}(\text{ not } P) = 1 - \text{sel}(P)$
 - (Observation: selectivity \approx probability)

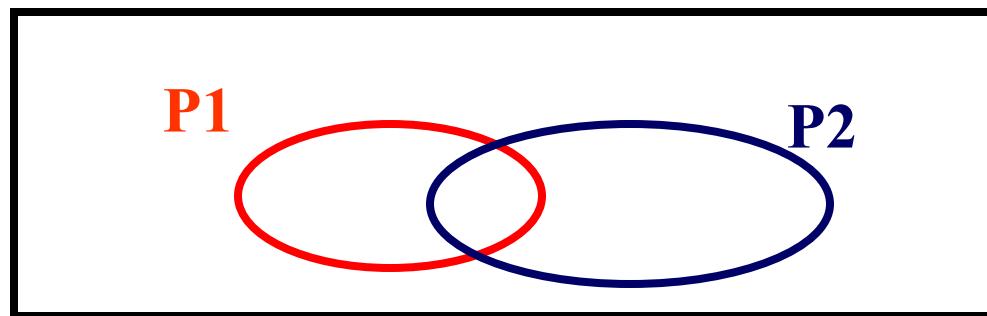




Selections – complex predicates

conjunction:

- $\text{sel}(\text{ grade} = \text{'C'} \text{ and course} = \text{'415'})$
- $\text{sel}(P1 \text{ and } P2) = \text{sel}(P1) * \text{sel}(P2)$
- INDEPENDENCE ASSUMPTION

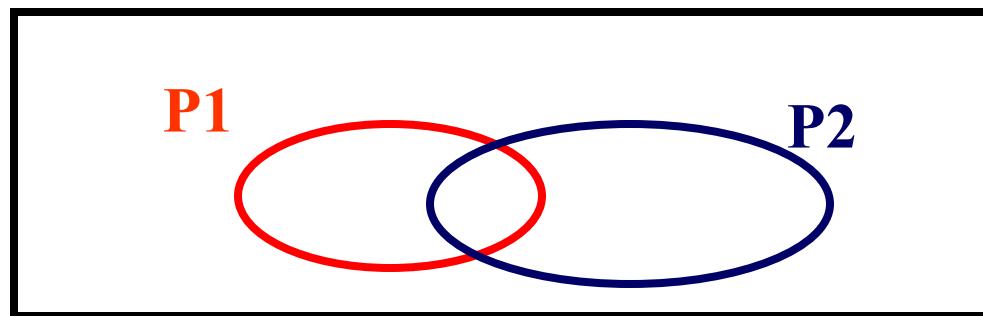




Selections – complex predicates

disjunction:

- $\text{sel}(\text{ grade} = \text{'C'} \text{ or course} = \text{'415'})$
- $\text{sel}(P1 \text{ or } P2) = \text{sel}(P1) + \text{sel}(P2) - \text{sel}(P1 \text{ and } P2)$
- $= \text{sel}(P1) + \text{sel}(P2) - \text{sel}(P1)*\text{sel}(P2)$
- INDEPENDENCE ASSUMPTION, again



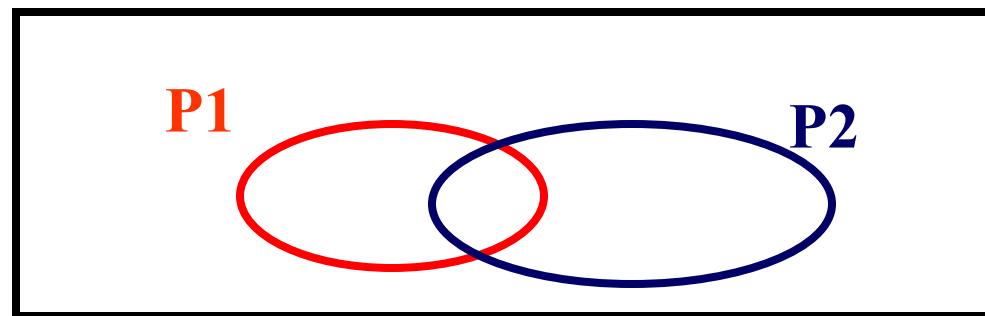


Selections – complex predicates

disjunction: in general

$$\text{sel}(P_1 \text{ or } P_2 \text{ or } \dots P_n) =$$

$$1 - (1 - \text{sel}(P_1)) * (1 - \text{sel}(P_2)) * \dots * (1 - \text{sel}(P_n))$$





Selections – summary

- $\text{sel}(A=\text{constant}) = 1/V(A,r)$
- $\text{sel}(A > a) = (A_{\max} - a) / (A_{\max} - A_{\min})$
- $\text{sel}(\text{not } P) = 1 - \text{sel}(P)$
- $\text{sel}(P_1 \text{ and } P_2) = \text{sel}(P_1) * \text{sel}(P_2)$
- $\text{sel}(P_1 \text{ or } P_2) = \text{sel}(P_1) + \text{sel}(P_2) - \text{sel}(P_1) * \text{sel}(P_2)$
- $\text{sel}(P_1 \text{ or } \dots \text{ or } P_n) = 1 - (1 - \text{sel}(P_1)) * \dots * (1 - \text{sel}(P_n))$

- UNIFORMITY and INDEPENDENCE ASSUMPTIONS



Result Size Estimation for Joins

- Q: Given a join of R and S, what is the range of possible result sizes (in #of tuples)?
 - Hint: what if $R_{\text{cols}} \cap S_{\text{cols}} = \emptyset$?
 - $R_{\text{cols}} \cap S_{\text{cols}}$ is a key for R (and a Foreign Key in S)?



Result Size Estimation for Joins

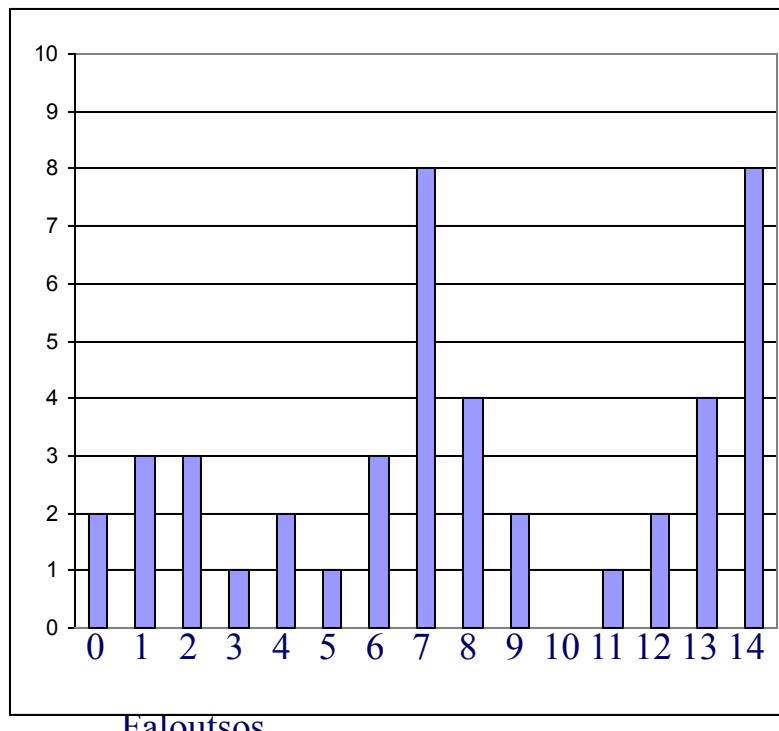
- General case: $R_{\text{cols}} \cap S_{\text{cols}} = \{A\}$ (and A is key for neither)
 - match each R-tuple with S-tuples
$$\begin{aligned}\text{est_size} &\sim \text{NTuples}(R) * \text{NTuples}(S) / \text{NKeys}(A, \mathbf{S}) \\ &\sim nr * ns / V(A, S)\end{aligned}$$
 - symmetrically, for S:
$$\begin{aligned}\text{est_size} &\sim \text{NTuples}(R) * \text{NTuples}(S) / \text{NKeys}(A, \mathbf{R}) \\ &\sim nr * ns / V(A, R)\end{aligned}$$
 - Overall:
$$\text{est_size} = \text{NTuples}(R) * \text{NTuples}(S) / \text{MAX}\{\text{NKeys}(A, \mathbf{S}), \text{NKeys}(A, \mathbf{R})\}$$



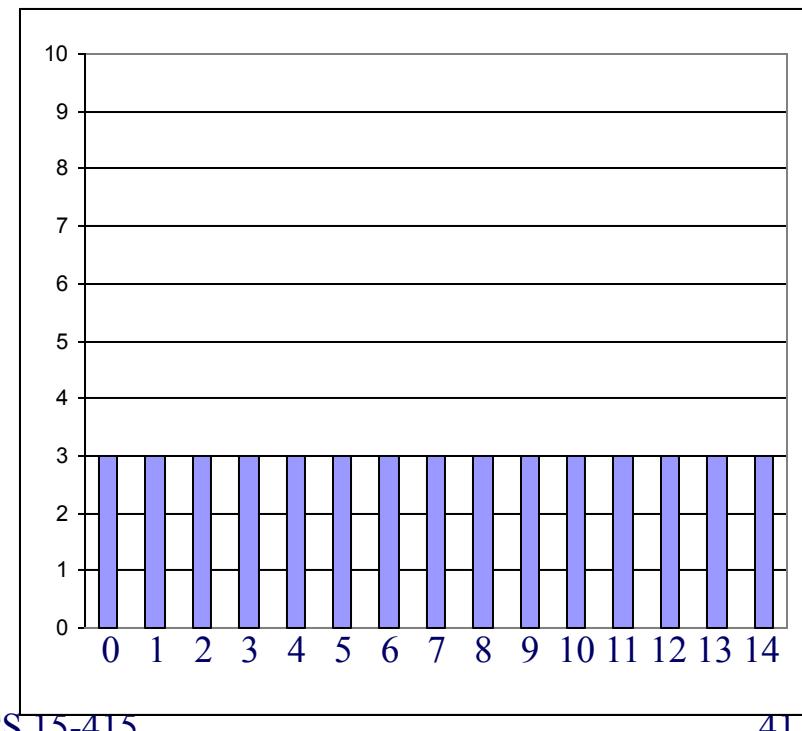
On the Uniform Distribution Assumption

- Assuming uniform distribution is rather crude

Distribution D



Uniform distribution approximating D

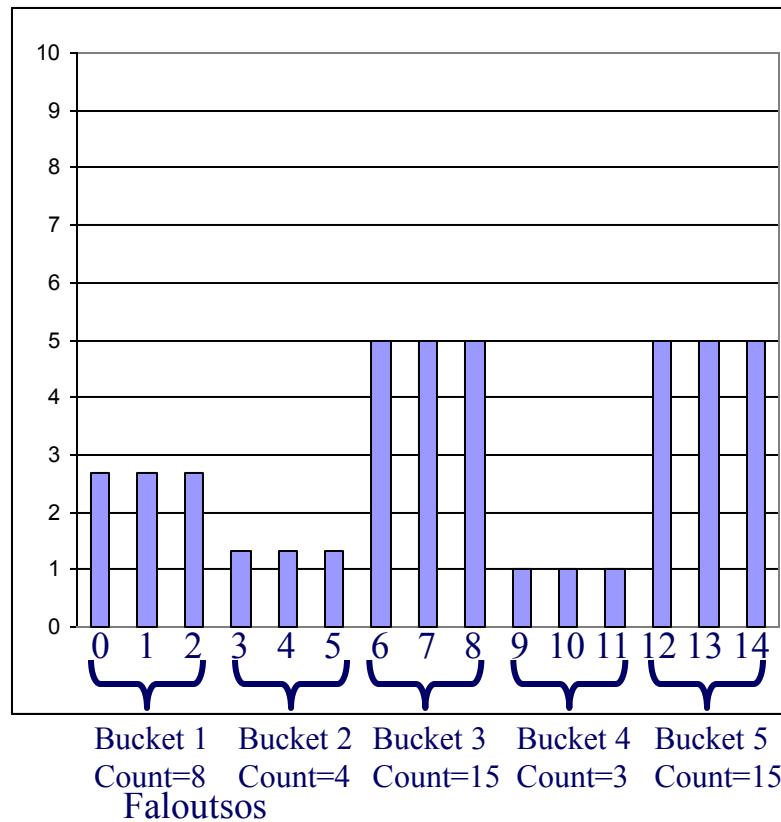




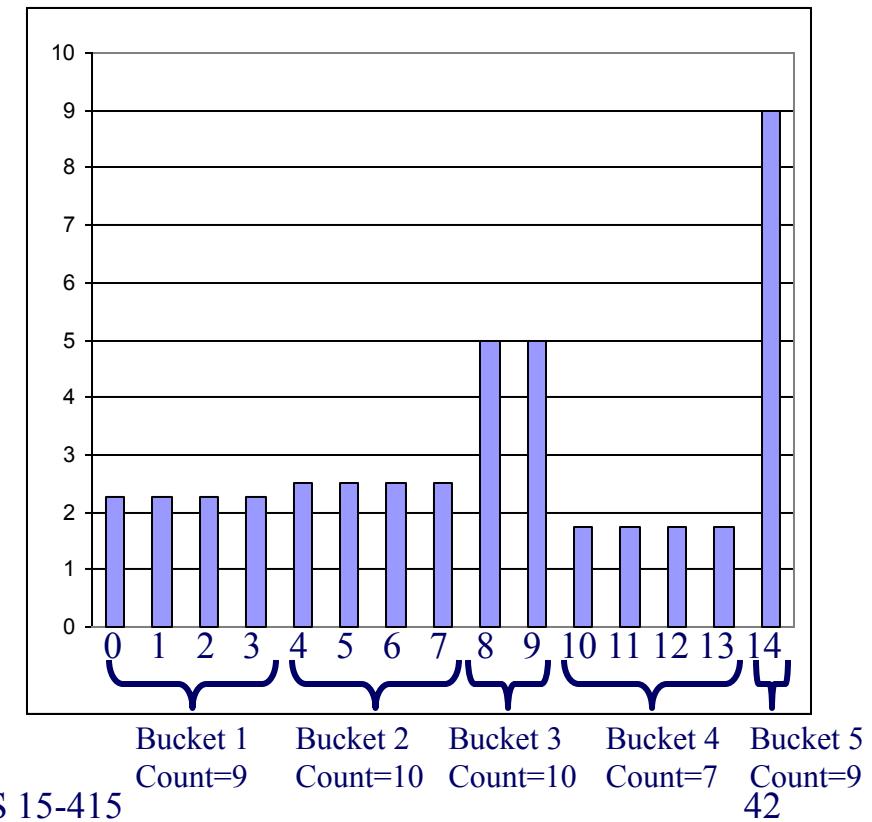
Histograms

- For better estimation, use a *histogram*

Equiwidth histogram



Equidepth histogram ~ quantiles





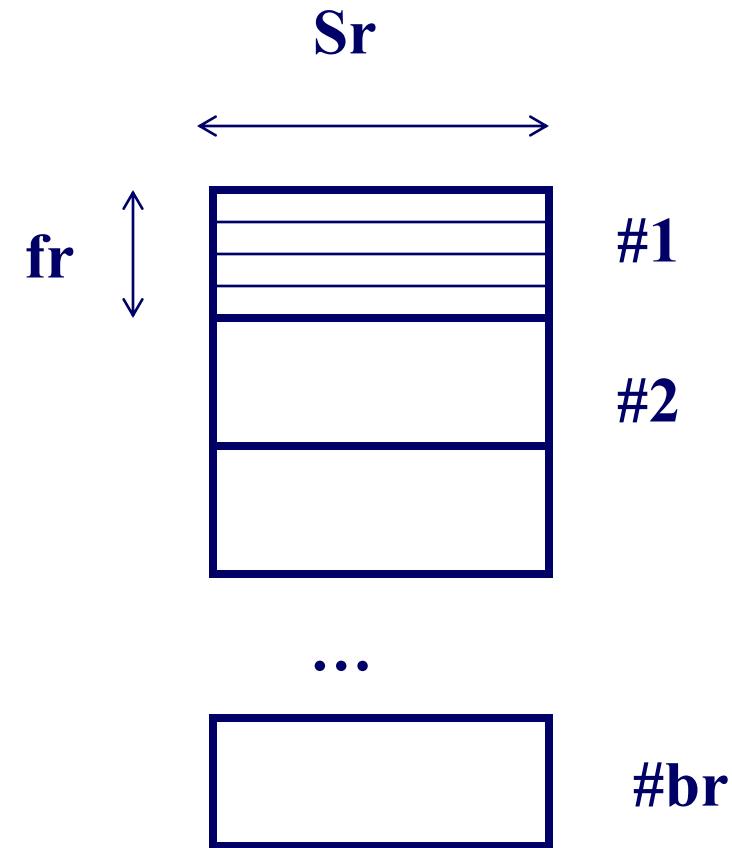
Q-opt steps

- bring query in internal form (eg., parse tree)
- ... into ‘canonical form’ (syntactic q-opt)
- **generate alt. plans**
 - **single relation**
 - multiple relations
- estimate cost; pick best



plan generation

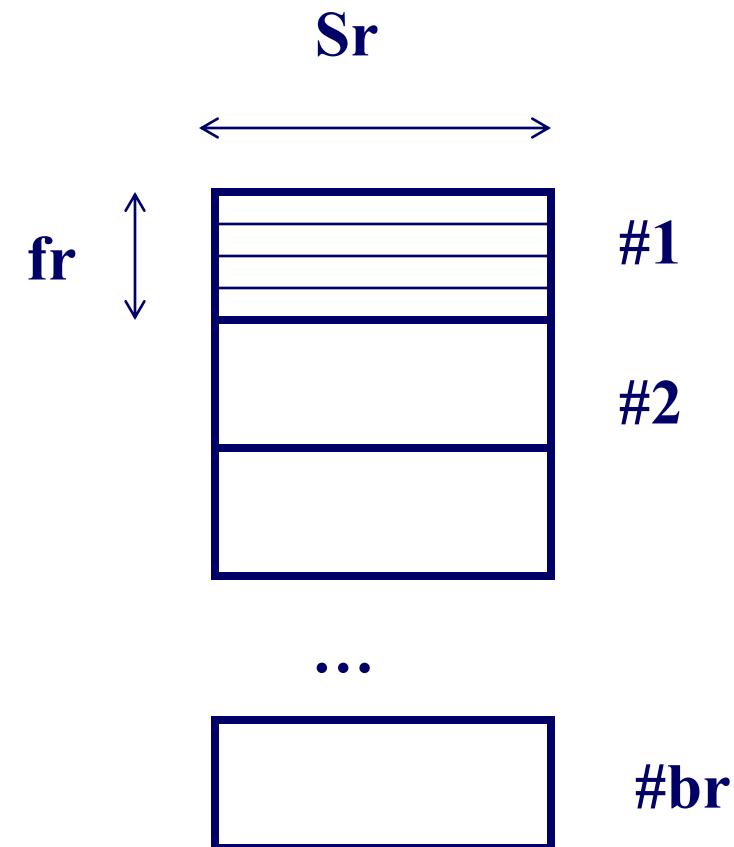
- Selections – eg.,
select *
from TAKES
where grade = ‘A’
- Plans?





plan generation

- Plans?
 - seq. scan
 - binary search
 - (if sorted & consecutive)
 - index search
 - if an index exists

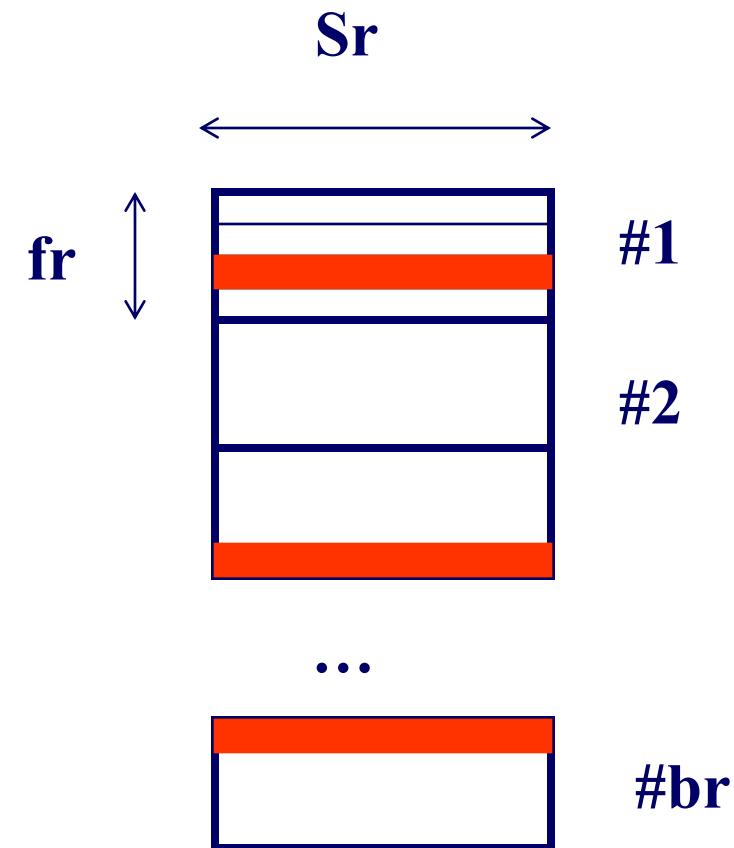




plan generation

seq. scan – cost?

- br (worst case)
- $br/2$ (average, if we search for primary key)



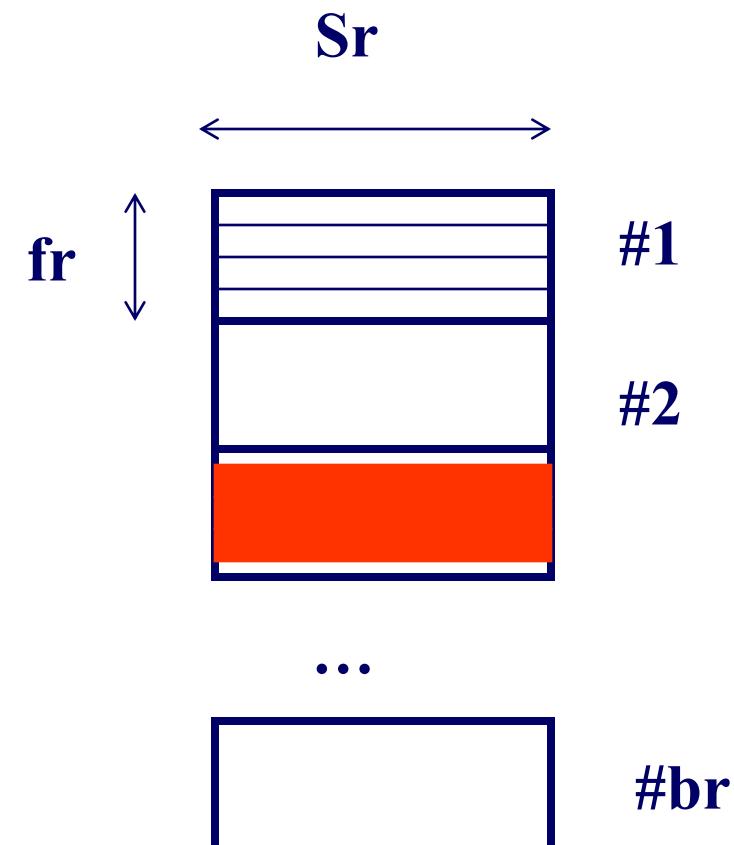


plan generation

binary search – cost?

if sorted and consecutive:

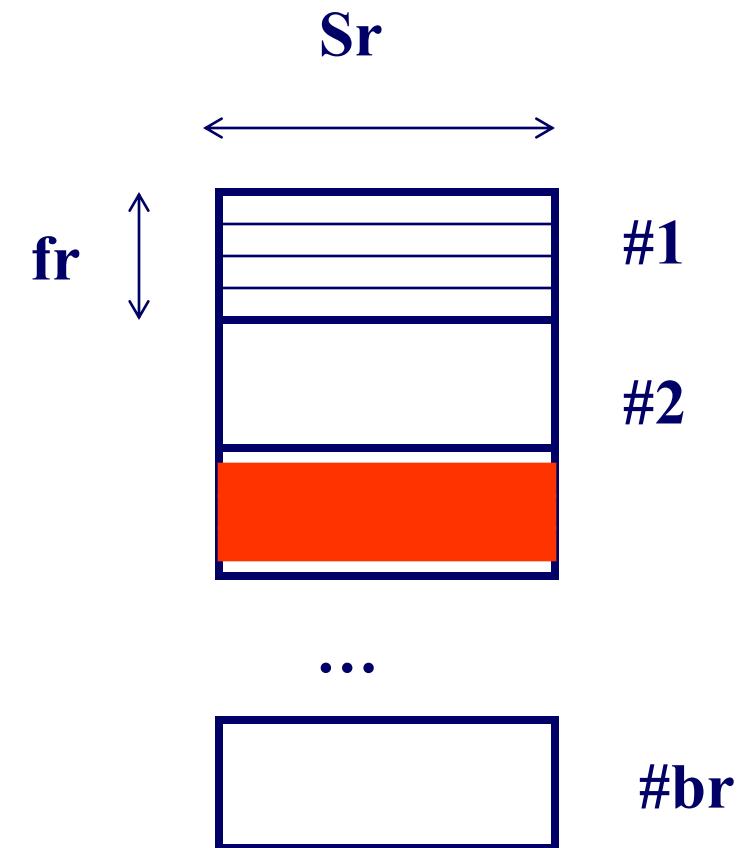
- $\sim \log(br) +$
- $SC(A,r)/fr$ (=blocks spanned by qual. tuples)





plan generation

estimation of selection cardinalities $SC(A,r)$:
non-trivial – we saw it earlier





plan generation

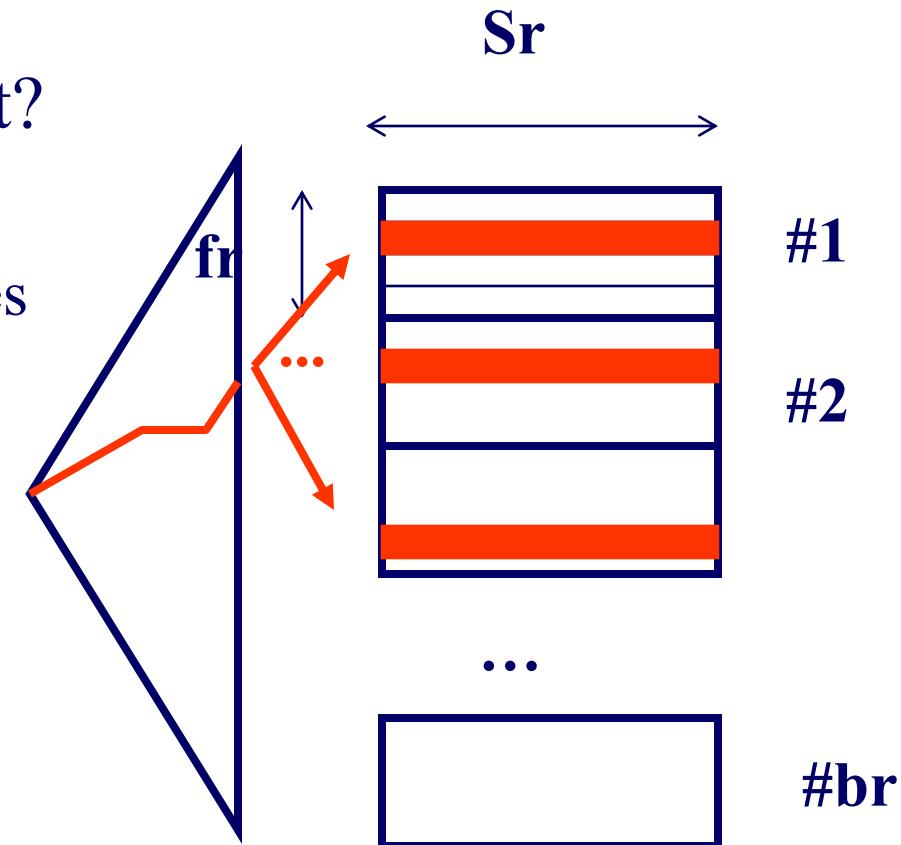
method#3: index – cost?

- levels of index +
- blocks w/ qual. tuples

case#1: primary key

case#2: sec. key –
clustering index

case#3: sec. key – non-
clust. index





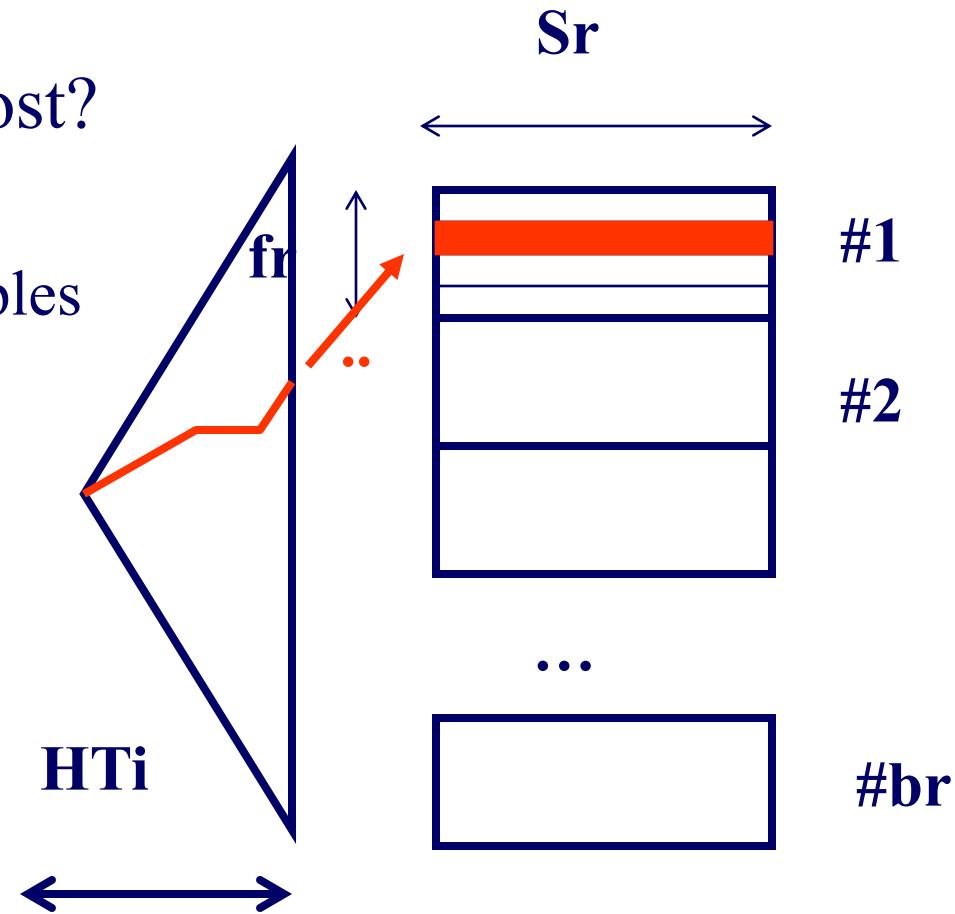
plan generation

method#3: index – cost?

- levels of index +
- blocks w/ qual. tuples

case#1: primary key – cost:

$HT_i + 1$





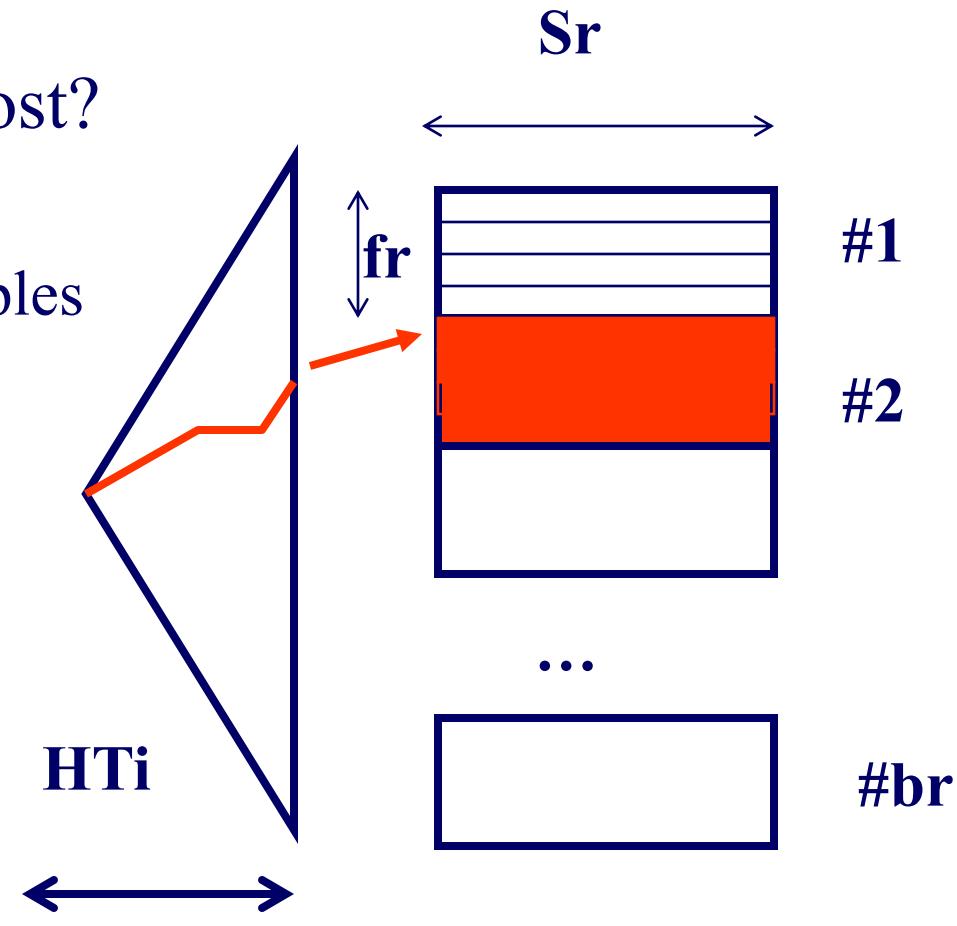
plan generation

method#3: index – cost?

- levels of index +
- blocks w/ qual. tuples

case#2: sec. key –
clustering index

$HTi + SC(A,r)/fr$





plan generation

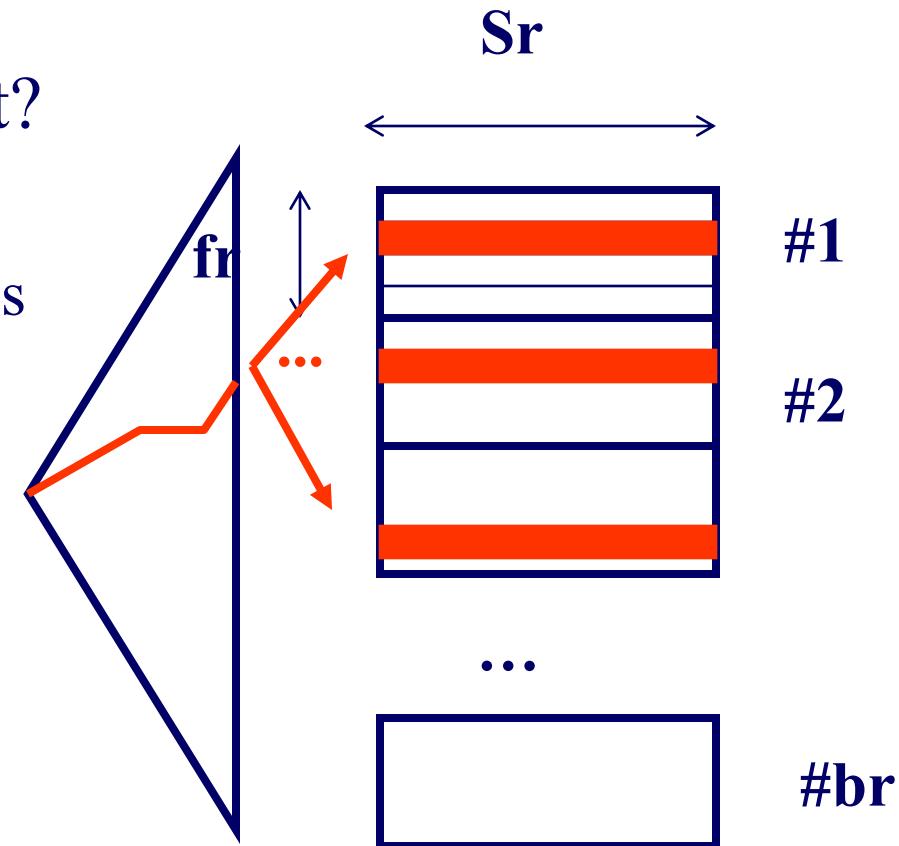
method#3: index – cost?

- levels of index +
- blocks w/ qual. tuples

case#3: sec. key – non-clust. index

HTi + SC(A,r)

(actually, pessimistic...)





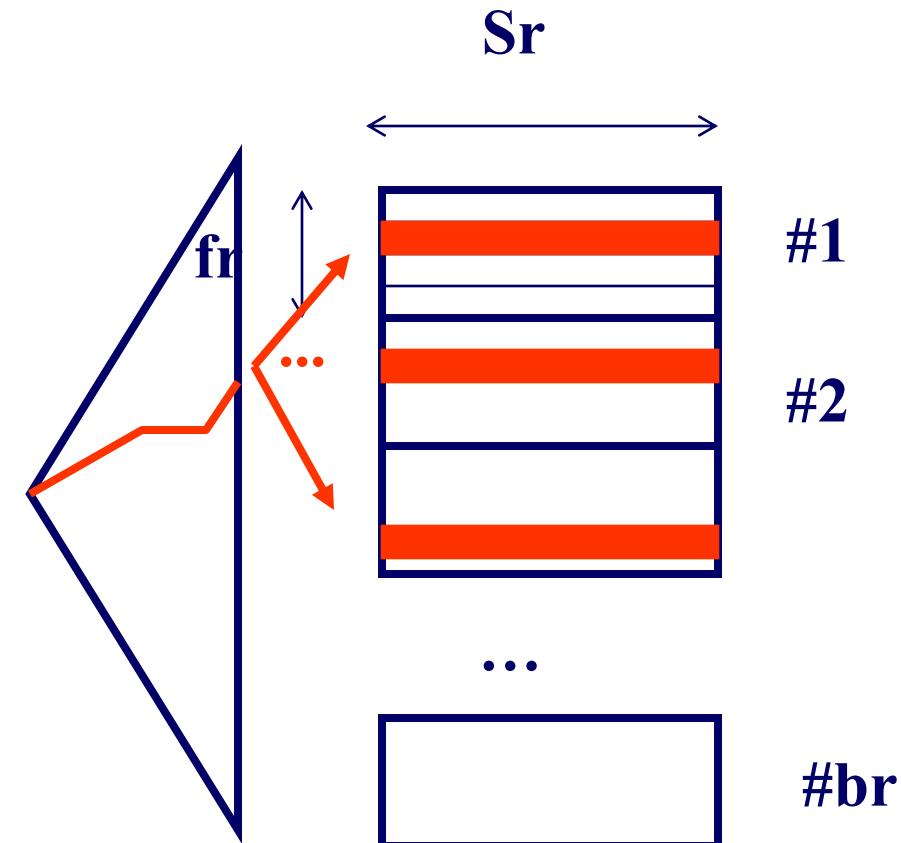
plan generation

method#3: index – cost?

- levels of index +
- blocks w/ qual. tuples

(actually, pessimistic...)

better estimates:
Cardenas' formula)



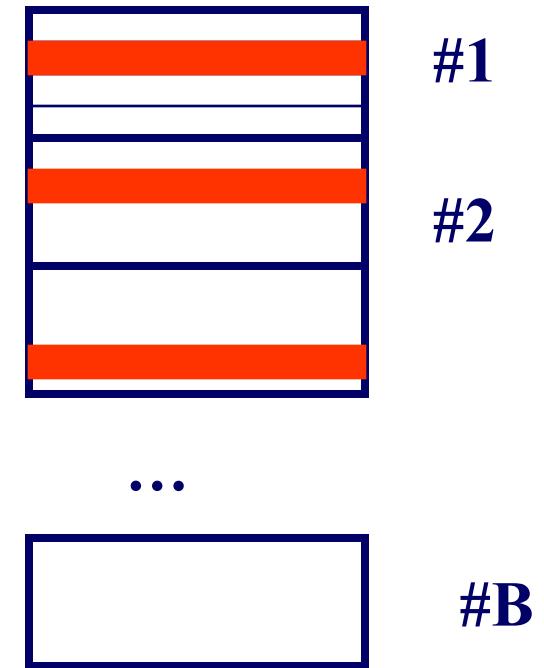


Cardena's formula

- q: # qual records
- Q: # qual. blocks
- N: # records total
- B: # blocks total
- $Q=??$



*Alfonso Cardenas
(UCLA)*

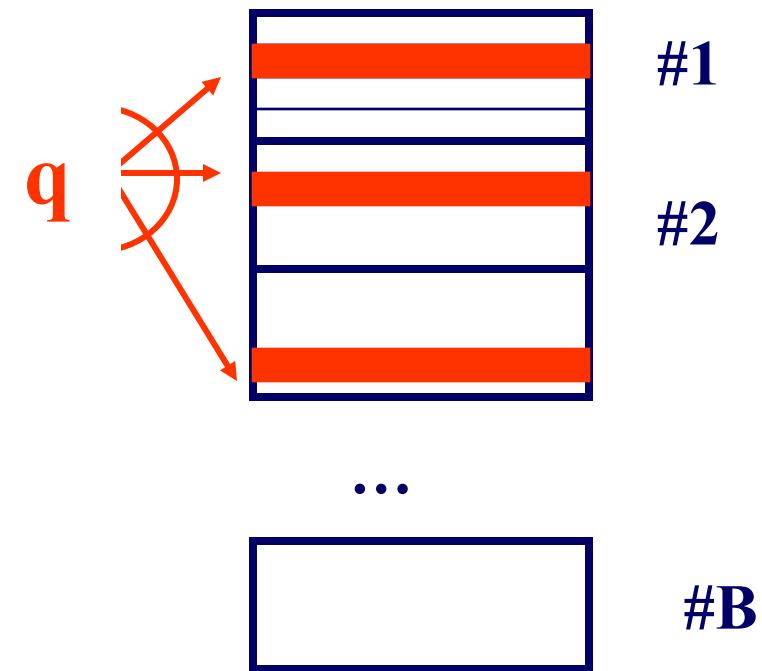




Cardena's formula

- Pessimistic:

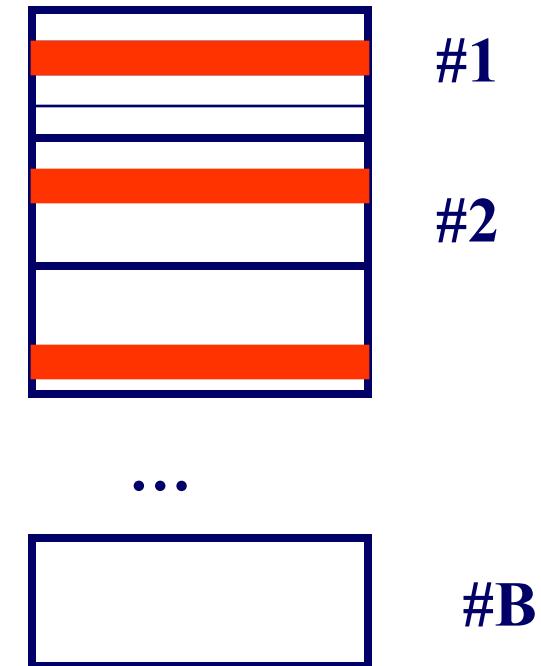
- $Q = q$





Cardena's formula

- Pessimistic:
 - $Q = q$
- More realistic
 - $Q = q$ if $q \leq B$
 - $Q = B$ otherwise

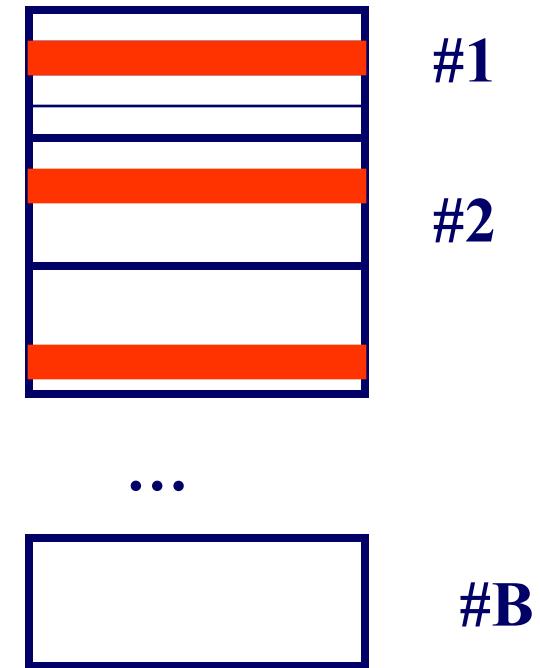




Cardena's formula

- Cardenas' formula

$$Q = B \left[1 - \left(1 - 1/B \right)^q \right]$$





Plans for single relation - summary

- no index: scan (dup-elim; sort)
- with index:
 - single index access path
 - multiple index access path
 - sorted index access path
 - index-only access path



Citation

- P. G. Selinger, M. M. Astrahan, D. D. Chamberlin, R. A. Lorie, and T. G. Price.
Access path selection in a relational database management system. In SIGMOD Conference, pages 23--34, 1979.



Frequently cited database publications

<http://www.informatik.uni-trier.de/~ley/db/about/top.html>

#	Publication
608	Peter P. Chen: The Entity-Relationship Model - Toward a Unified View of Data. ACM Trans. Database Syst. 1(1): 9-36(1976)
580	E. F. Codd: A Relational Model of Data for Large Shared Data Banks. Commun. ACM 13(6): 377-387(1970)
371	Patricia G. Selinger, Morton M. Astrahan, Donald D. Chamberlin, Raymond A. Lorie, Thomas G. Price: Access Path Selection in a Relational Database Management System. SIGMOD Conference 1979: 23-34
366	Jeffrey D. Ullman: Principles of Database and Knowledge Base Systems, Volume I. Computer Science Press 1988, ISBN 0-7167-8158-1
...	...



Statistics for Optimization

- NCARD (T) - cardinality of relation T in tuples
- TCARD (T) - number of pages containing tuples from T
- $P(T) = \text{TCARD}(T)/(\# \text{ of non-empty pages in the segment})$
 - If segments only held tuples from one relation there would be no need for $P(T)$
- ICARD(I) - number of distinct keys in index I
- NINDEX(I) - number of pages in index I



Predicate Selectivity Estimation

attr = value	$F = 1/ICARD(attr\ index)$ – if index exists $F = 1/10$ otherwise
attr1 = attr2	$F = 1/max(ICARD(I1),ICARD(I2))$ or $F = 1/ICARD(Ii)$ – if only index i exists, or $F = 1/10$
vall < attr < val2	$F = (value2-value1)/(high\ key-low\ key)$ $F = 1/4$ otherwise
expr1 or expr2	$F = F(expr1)+F(expr2)-F(expr1)*F(expr2)$
expr1 and expr2	$F = F(expr1) * F(expr2)$
NOT expr	$F = 1 - F(expr)$



Costs per Access Path Case

Unique index matching equal predicate	$1+1+W$
Clustered index I matching ≥ 1 preds	$F(\text{preds}) * (\text{NINDEX}(I) + \text{TCARD}) + W * \text{RSICARD}$
Non-clustered index I matching ≥ 1 preds	$F(\text{preds}) * (\text{NINDEX}(I) + \text{NCARD}) + W * \text{RSICARD}$
Segment scan	$\text{TCARD}/P + W * \text{RSICARD}$



Q-opt steps

- bring query in internal form (eg., parse tree)
- ... into ‘canonical form’ (syntactic q-opt)
- generate alt. plans
 - single relation
 - **multiple relations**
 - Main idea
 - Dynamic programming – reminder
 - Example
- estimate cost; pick best



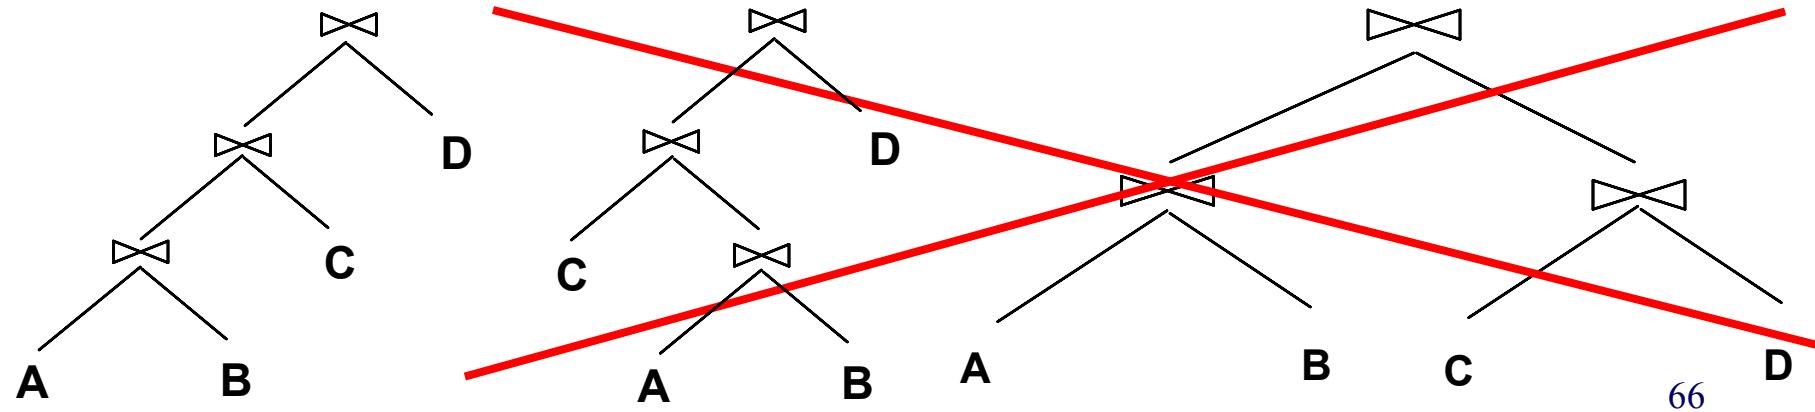
n-way joins

- $r_1 \text{ JOIN } r_2 \text{ JOIN } \dots \text{ JOIN } r_n$
- typically, break problem into 2-way joins
 - choose between NL, sort merge, hash join, ...



Queries Over Multiple Relations

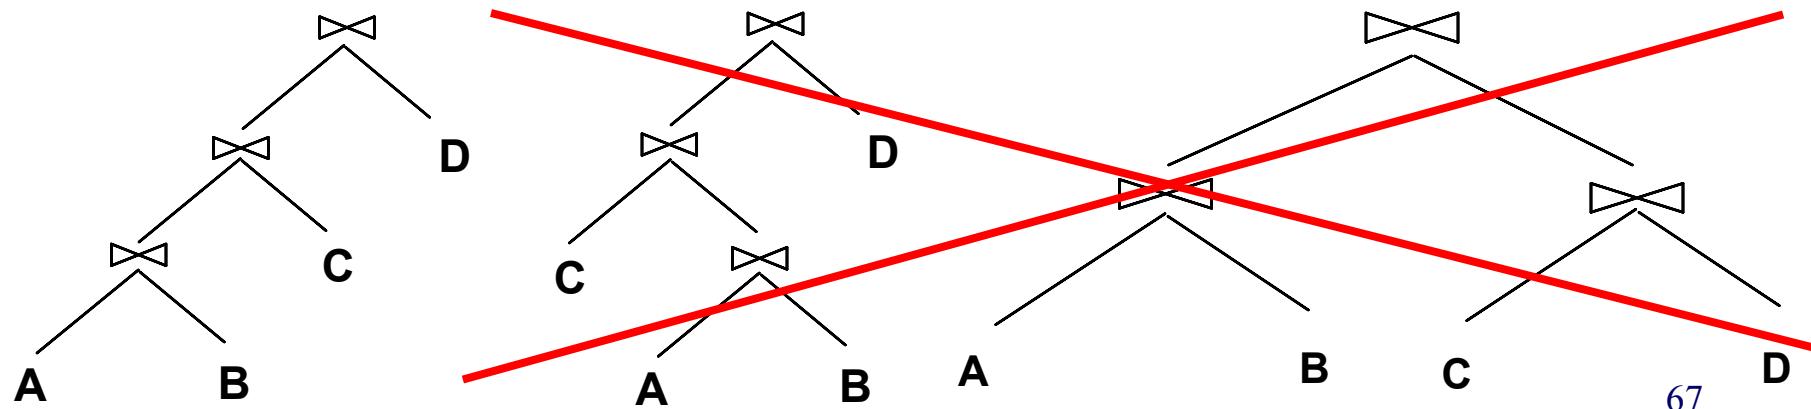
- As number of joins increases, number of alternative plans grows rapidly → *need to restrict search space*
- Fundamental decision in System R: only left-deep join trees are considered. Advantages?





Queries Over Multiple Relations

- As number of joins increases, number of alternative plans grows rapidly → *need to restrict search space*
- Fundamental decision in System R: only left-deep join trees are considered. Advantages?
 - fully pipelined* plans.
 - Intermediate results not written to temporary files.
 - Not all left-deep trees are fully pipelined (e.g., SM join).





Queries over Multiple Relations

- Enumerate the orderings (= left deep tree)
- enumerate the plans for each operator
- enumerate the access paths for each table

Dynamic programming, to save cost estimations

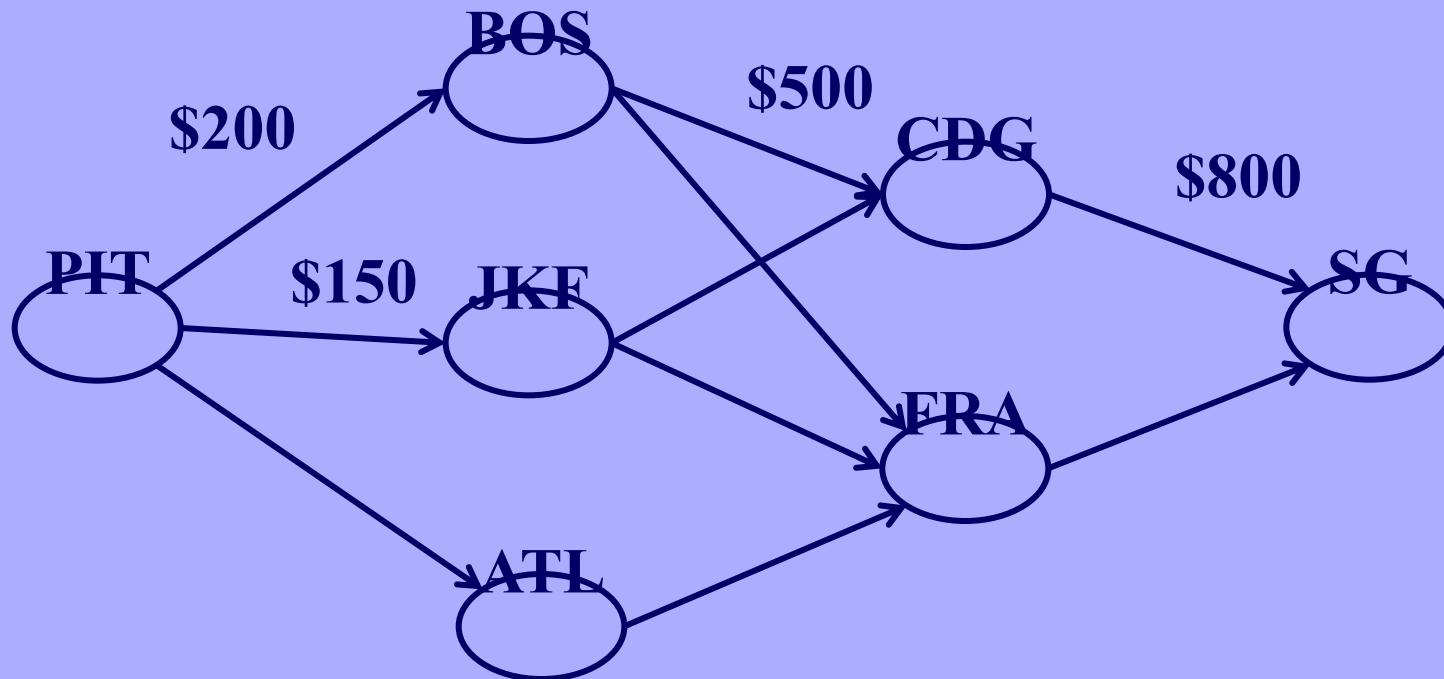


Q-opt steps

- bring query in internal form (eg., parse tree)
- ... into ‘canonical form’ (syntactic q-opt)
- generate alt. plans
 - single relation
 - multiple relations
 - Main idea
 - **Dynamic programming – reminder**
 - Example
- estimate cost; pick best



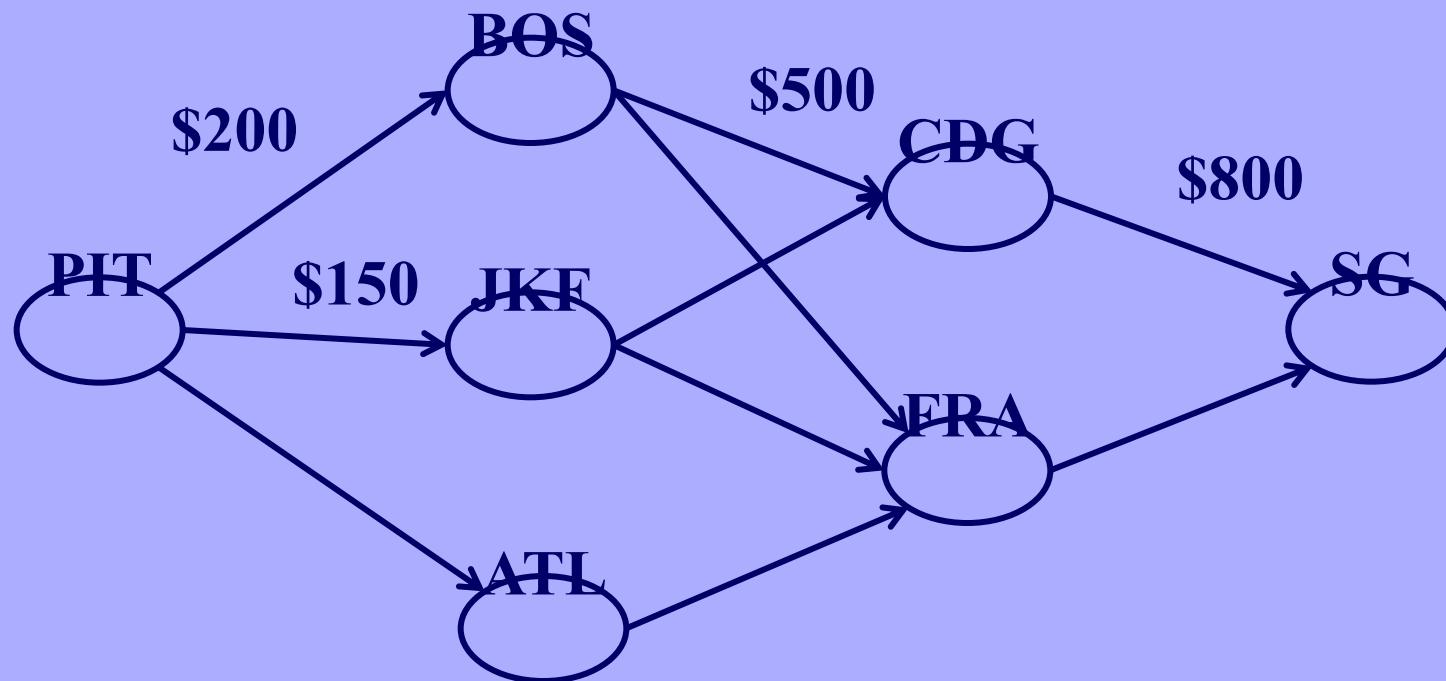
(Reminder: Dynamic Programming)



Cheapest flight PIT -> SG ?



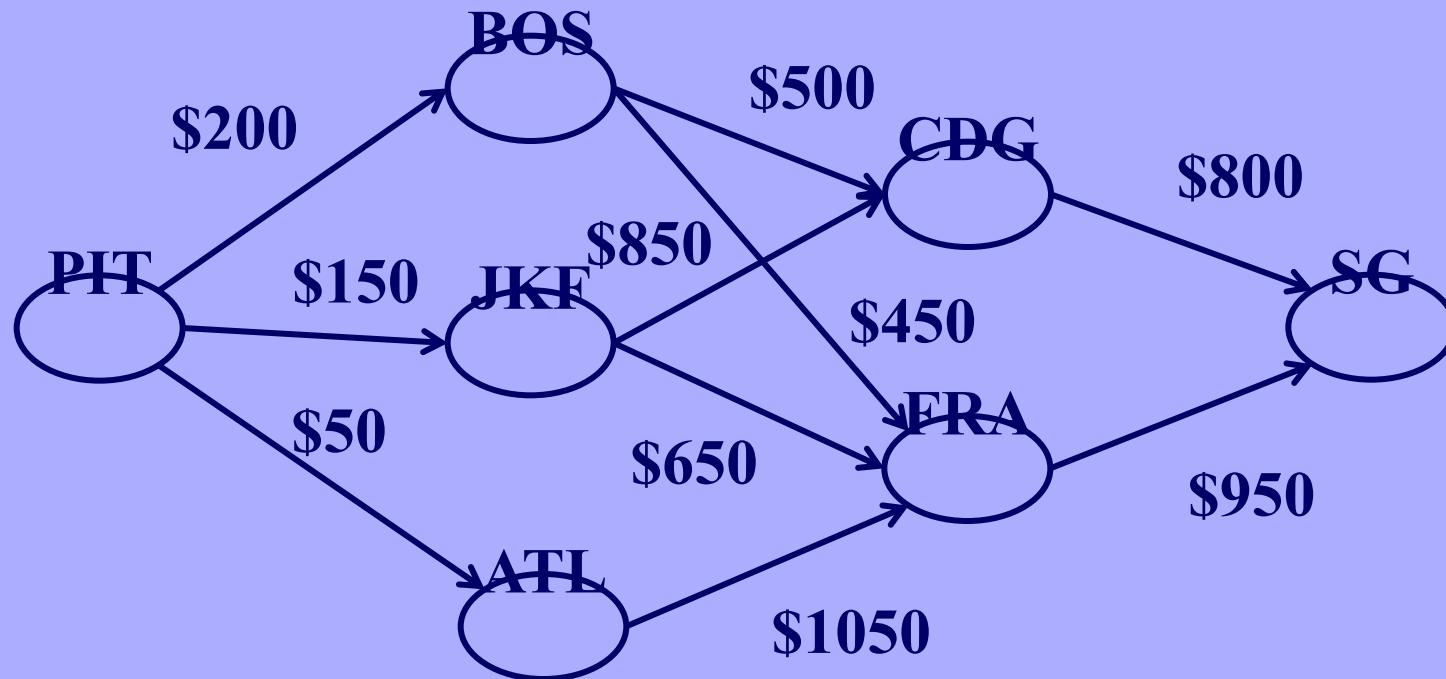
(Reminder: Dynamic Programming)



Assumption: NO package deals: cost CDG->SG
is always \$800, no matter how reached CDG



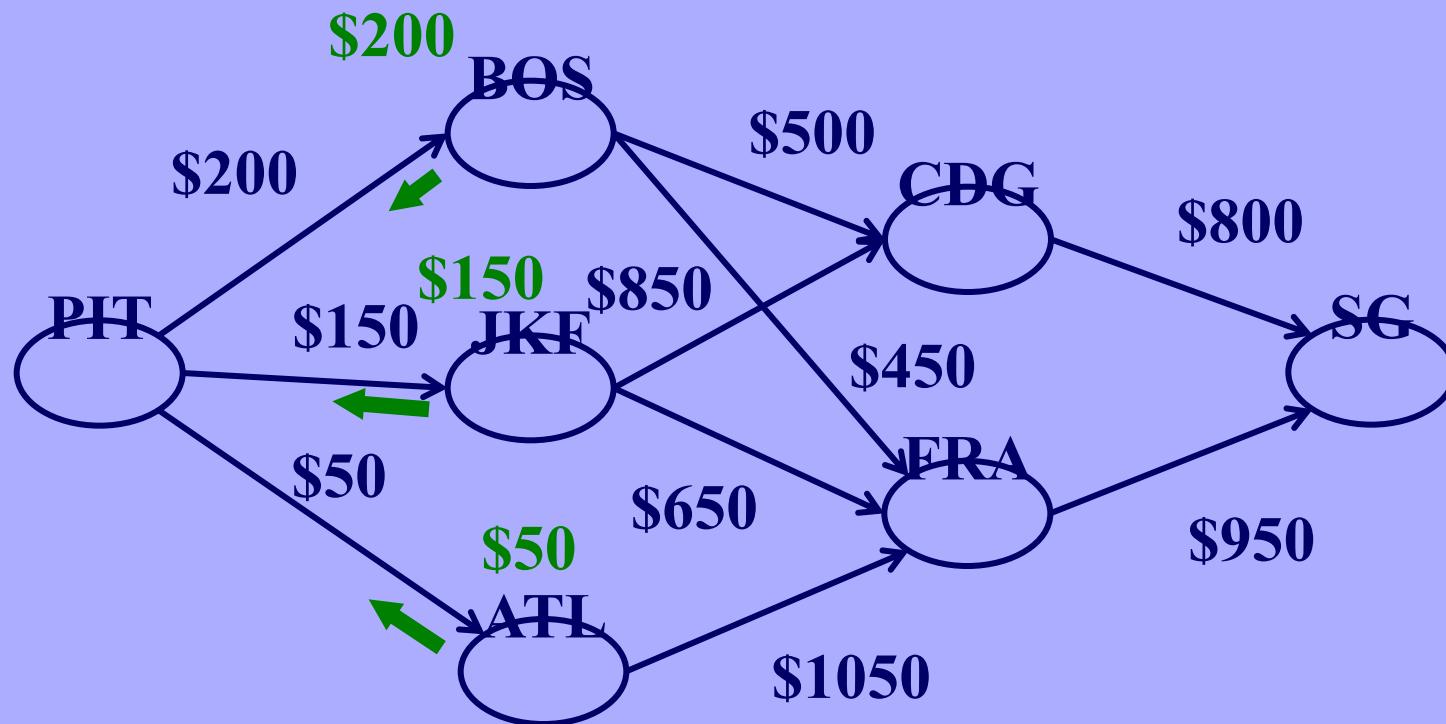
(Reminder: Dynamic Programming)



Solution: compute partial optimal, left-to-right:



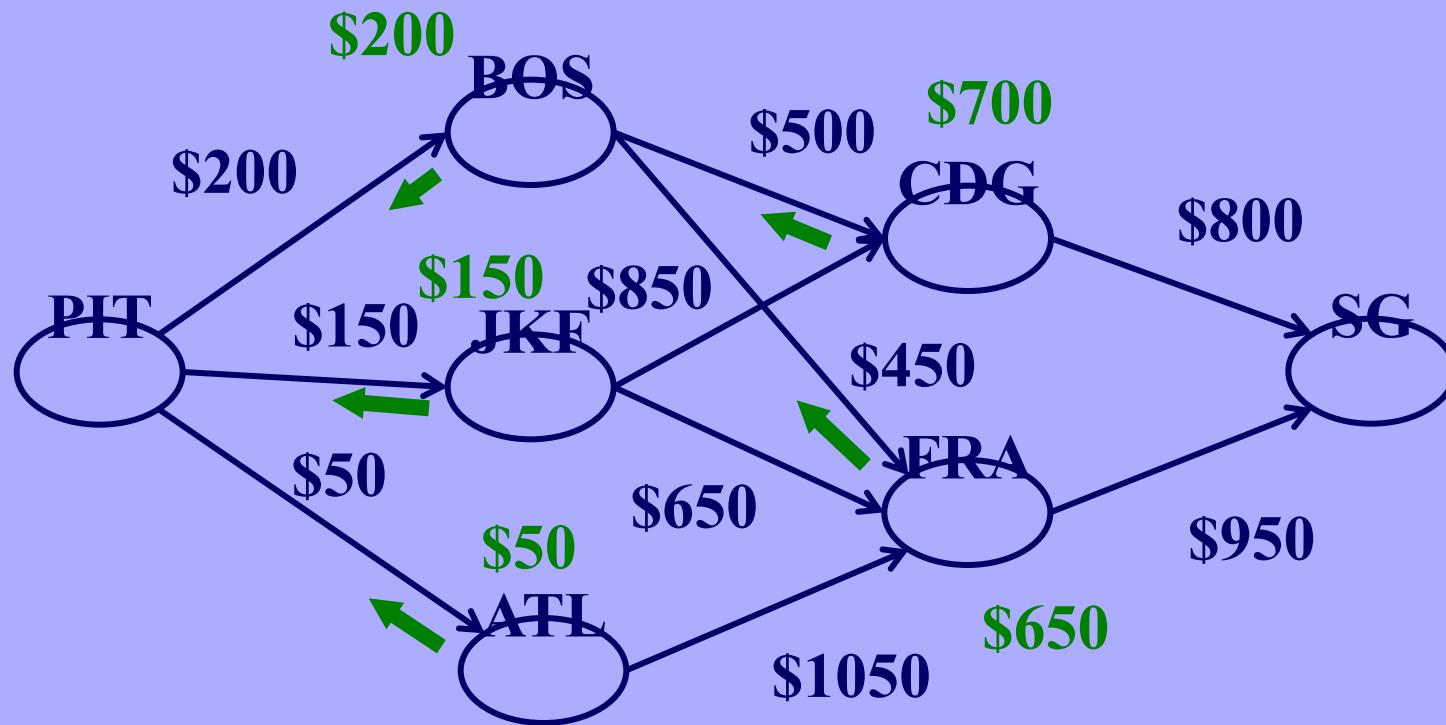
(Reminder: Dynamic Programming)



Solution: compute partial optimal, left-to-right:



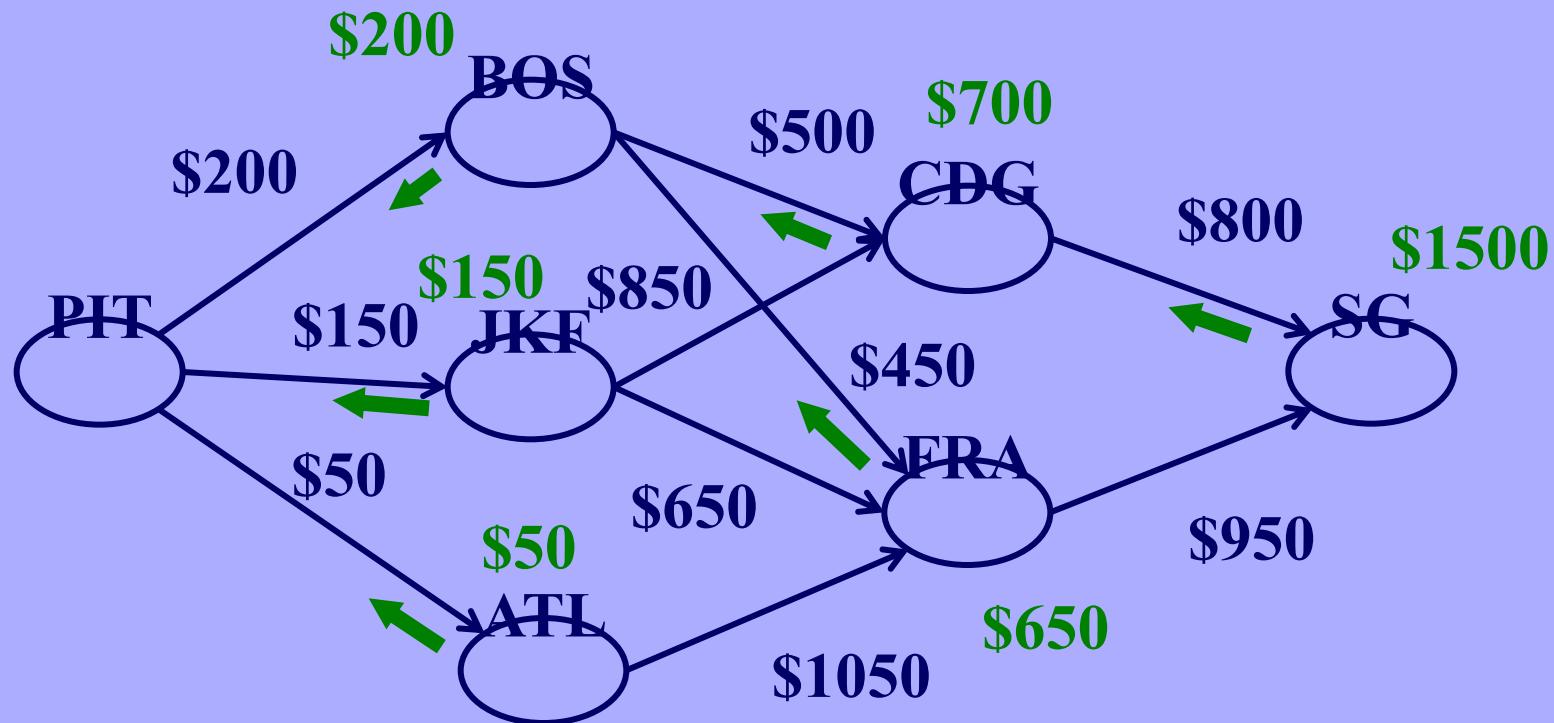
(Reminder: Dynamic Programming)



Solution: compute partial optimal, left-to-right:



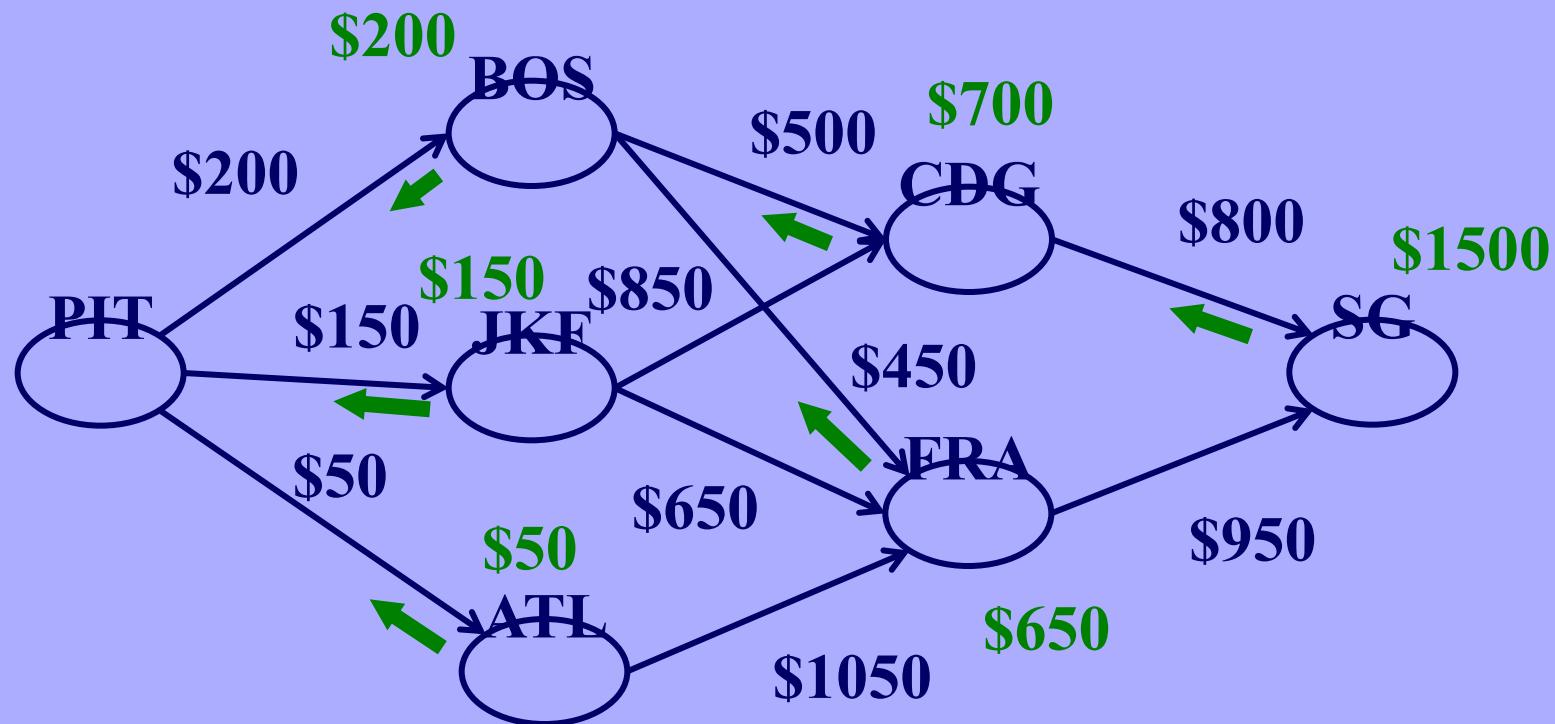
(Reminder: Dynamic Programming)



Solution: compute partial optimal, left-to-right:



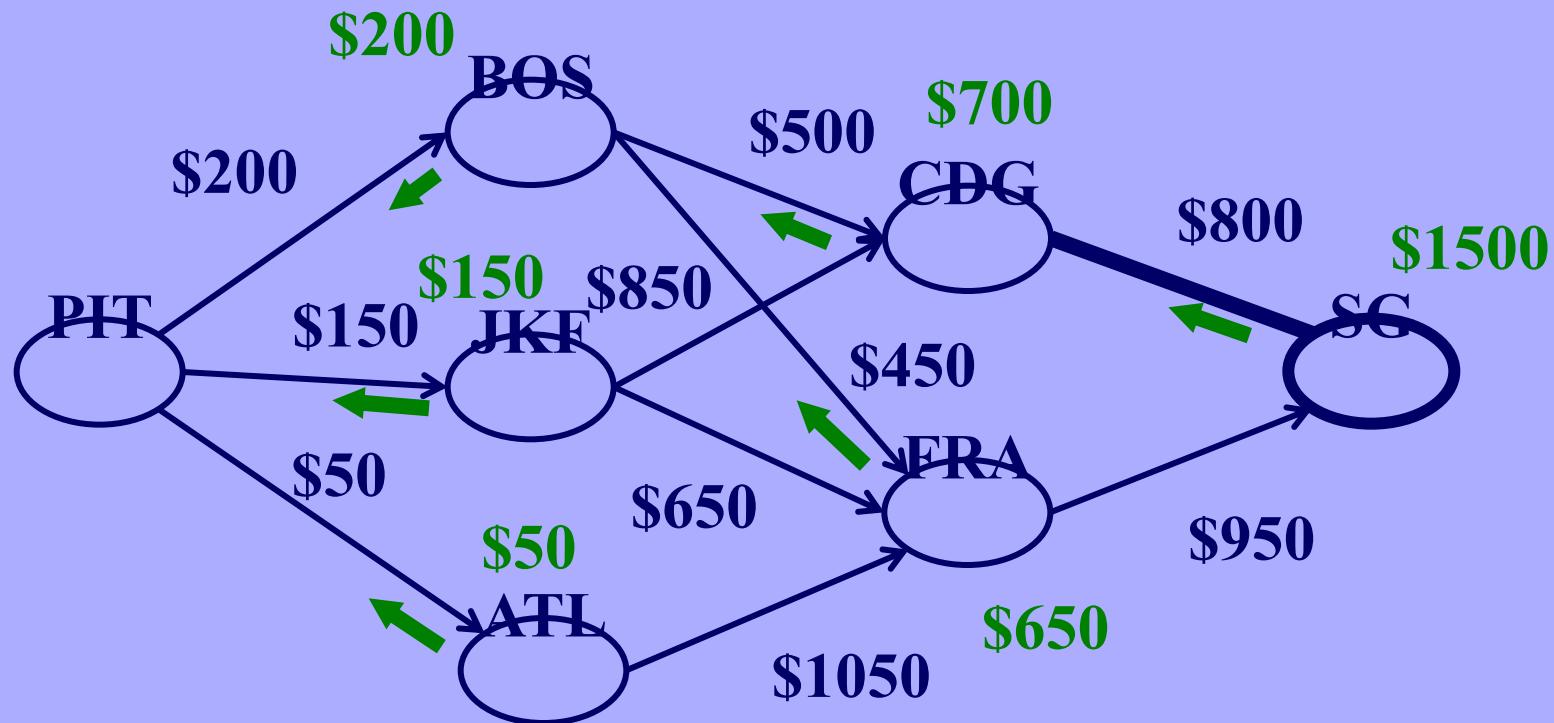
(Reminder: Dynamic Programming)



So, best price is \$1,500 – which legs?



(Reminder: Dynamic Programming)

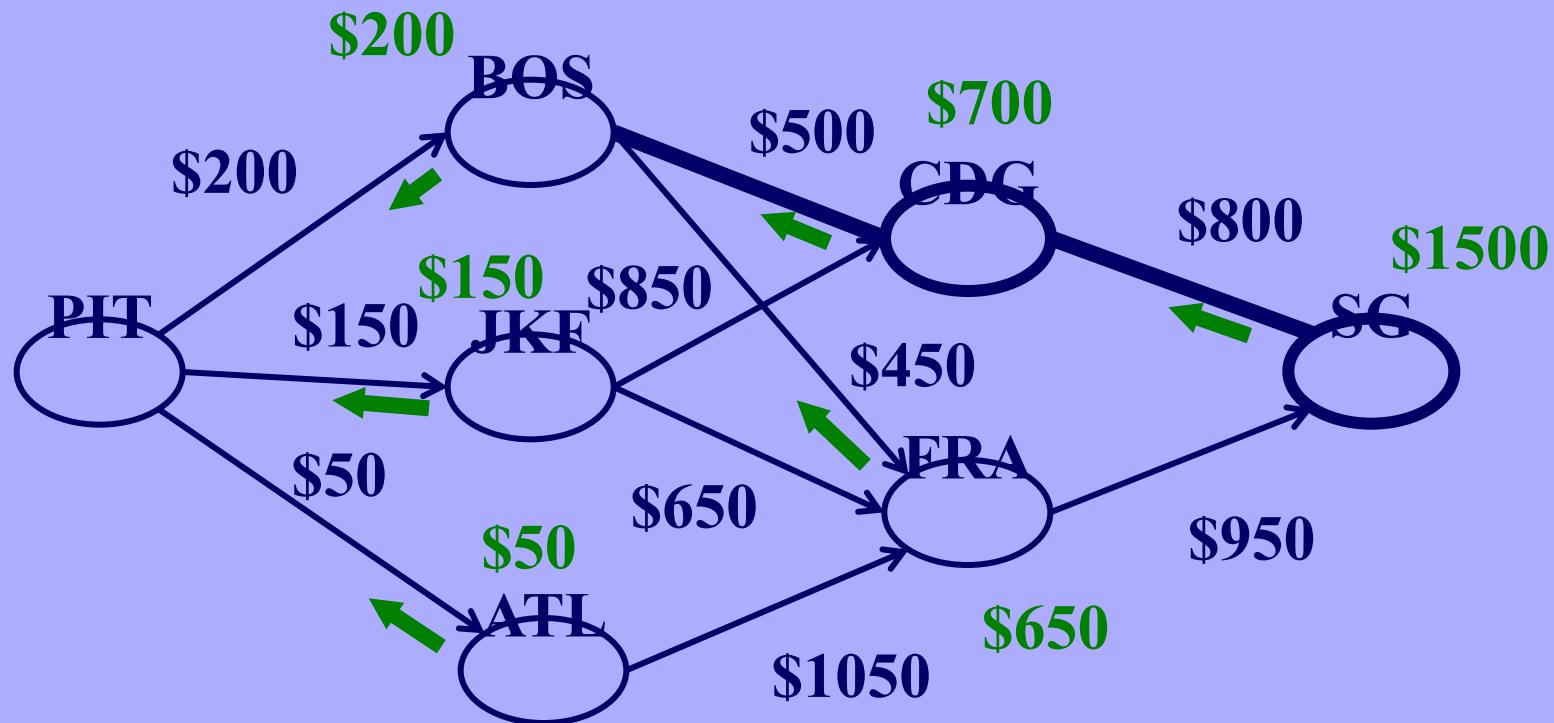


So, best price is \$1,500 – which legs?

A: follow the winning edges, backwards



(Reminder: Dynamic Programming)

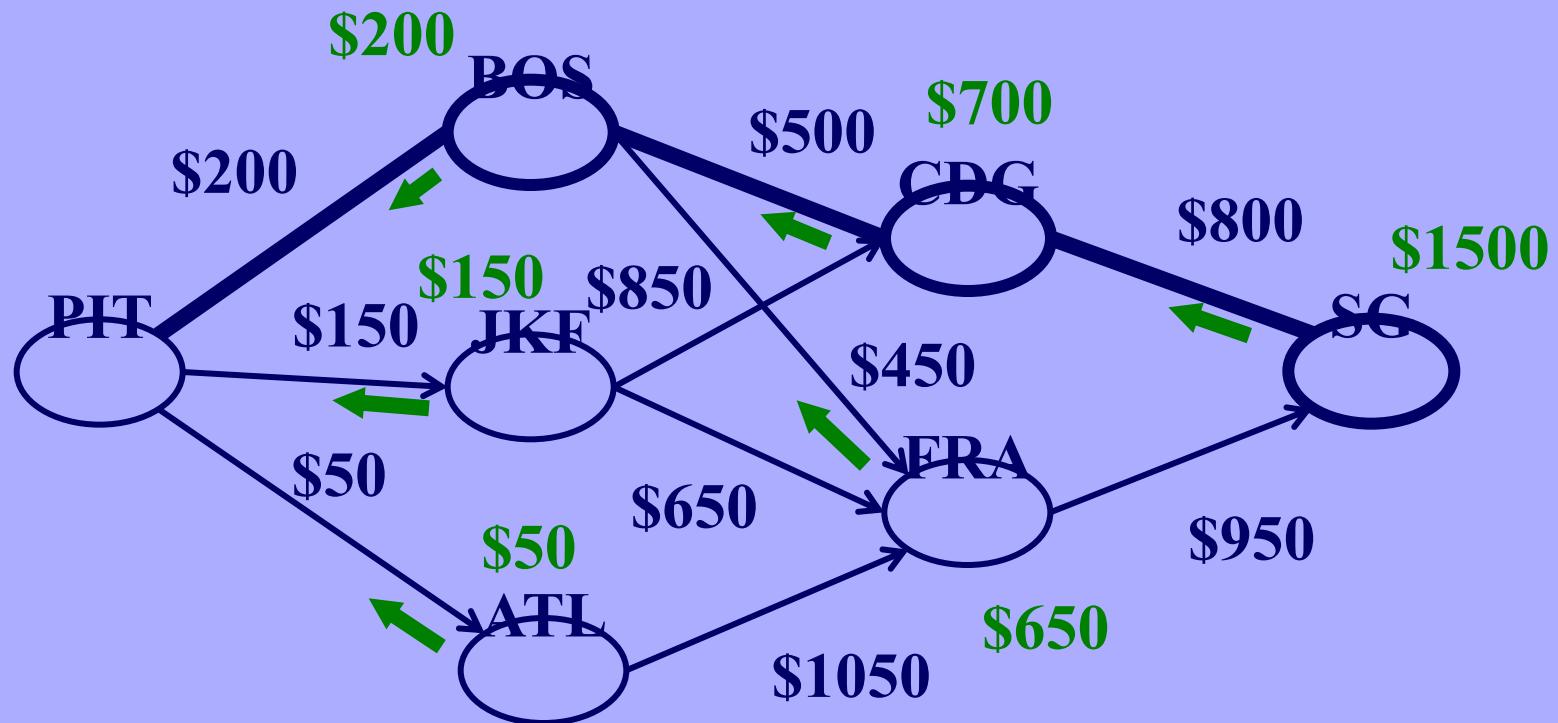


So, best price is \$1,500 – which legs?

A: follow the winning edges, backwards



(Reminder: Dynamic Programming)

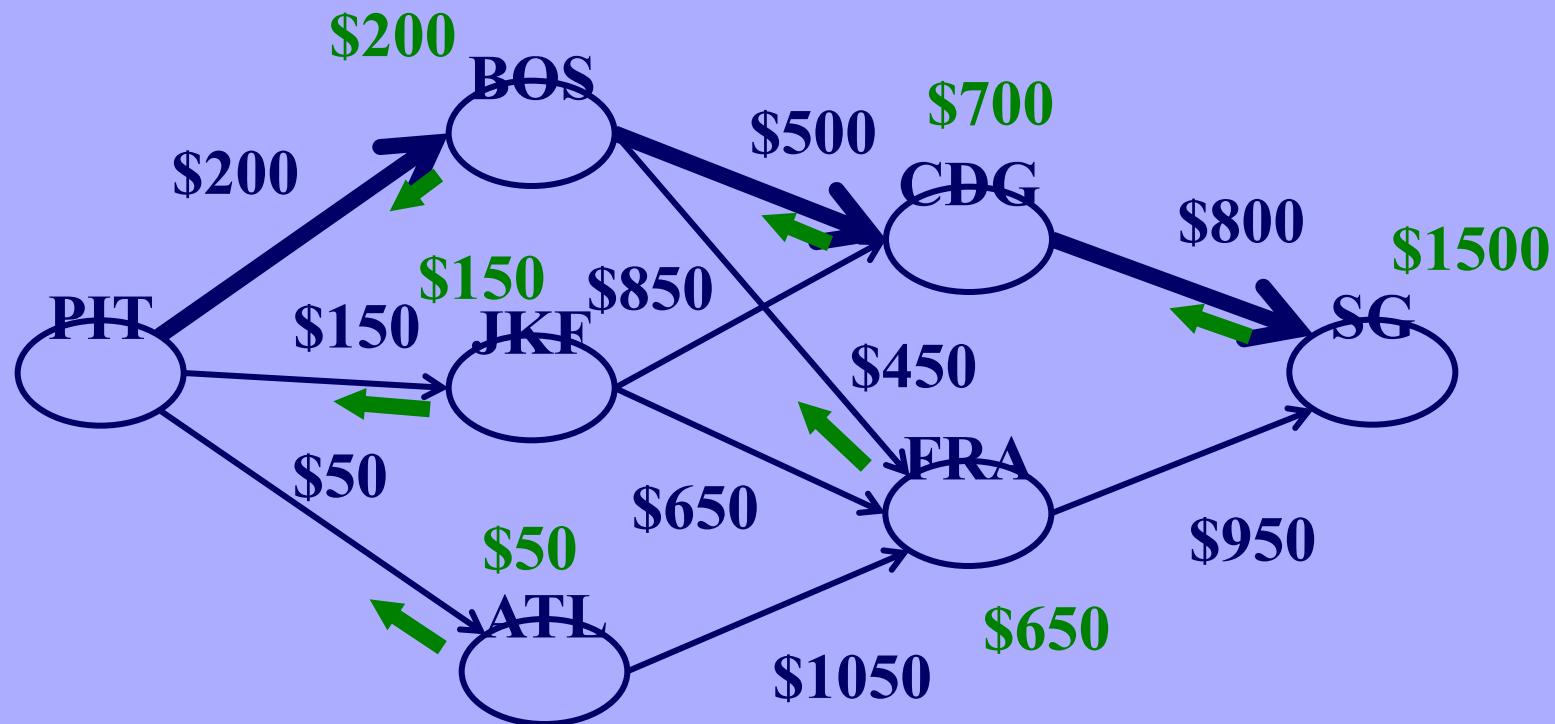


So, best price is \$1,500 – which legs?

A: follow the winning edges, backwards



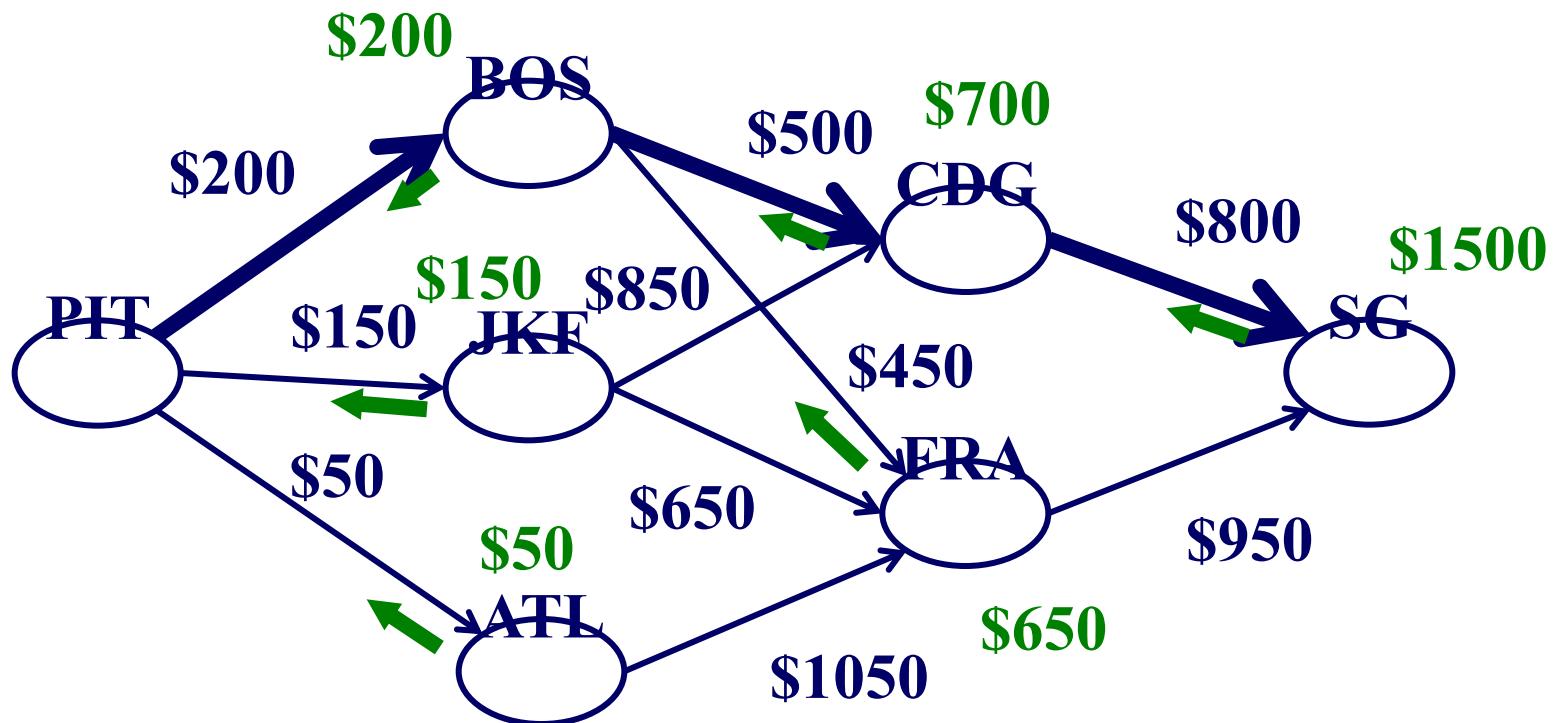
(Reminder: Dynamic Programming)



Q: what are the states, costs and arrows, in q-opt?



(Reminder: Dynamic Programming)



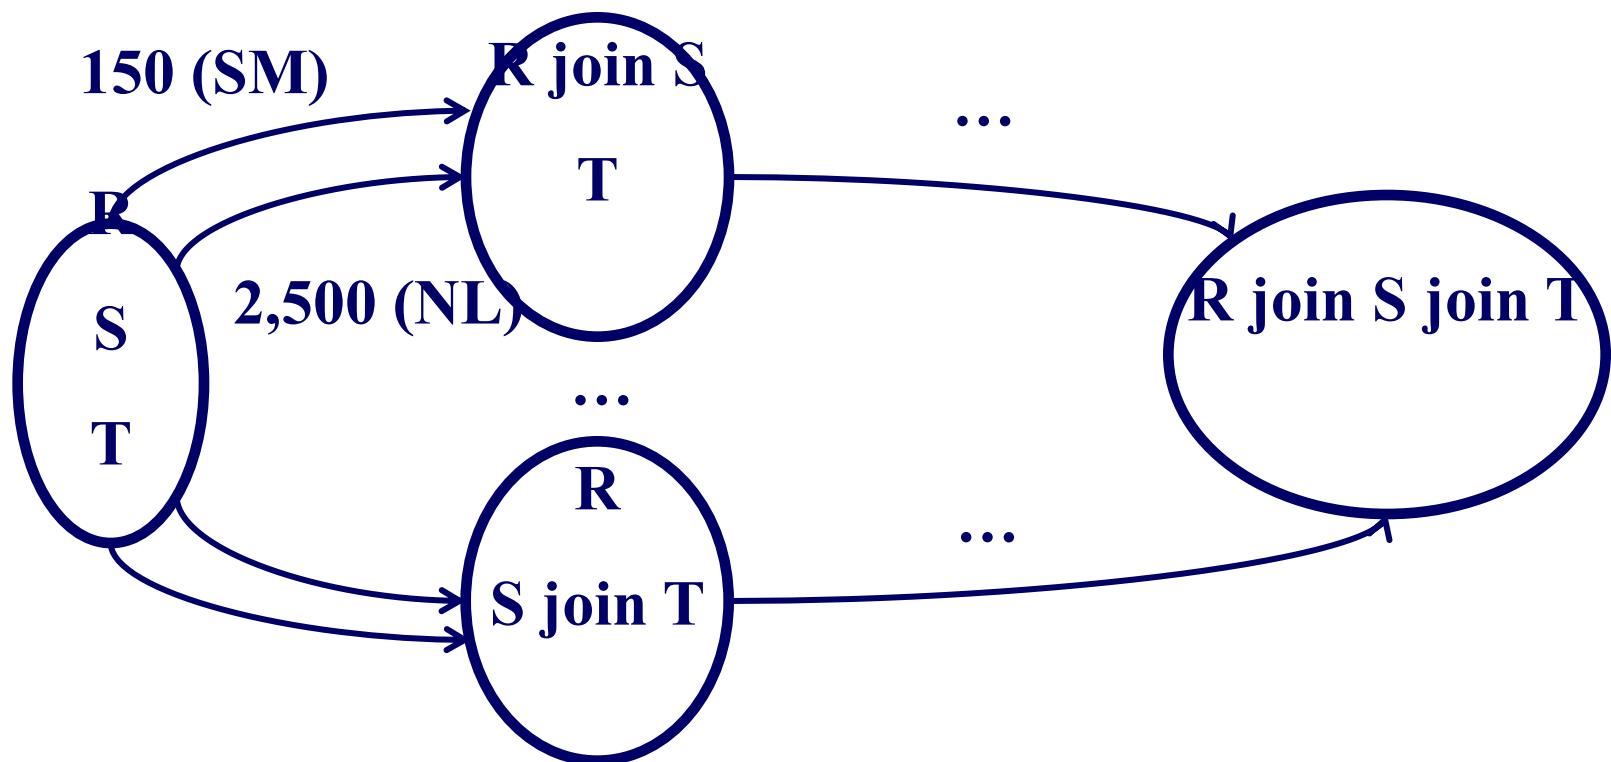
Q: what are the states (and costs and arrows), in q-opt?

A: set of intermediate result tables



Q-opt and Dyn. Programming

- E.g., compute $R \text{ join } S \text{ join } T$





Q-opt and Dyn. Programming

- Details: how to record the fact that, say R is sorted on R.a? or that the user requires sorted output?
- A:
- E.g., consider the query
select *
from R, S, T
where R.a = S.a and S.b = T.b
order by R.a



Q-opt and Dyn. Programming

- Details: how to record the fact that, say R is sorted on R.a? or that the user requires sorted output?

- A: record orderings, in the state
- E.g., consider the query

select *

from R, S, T

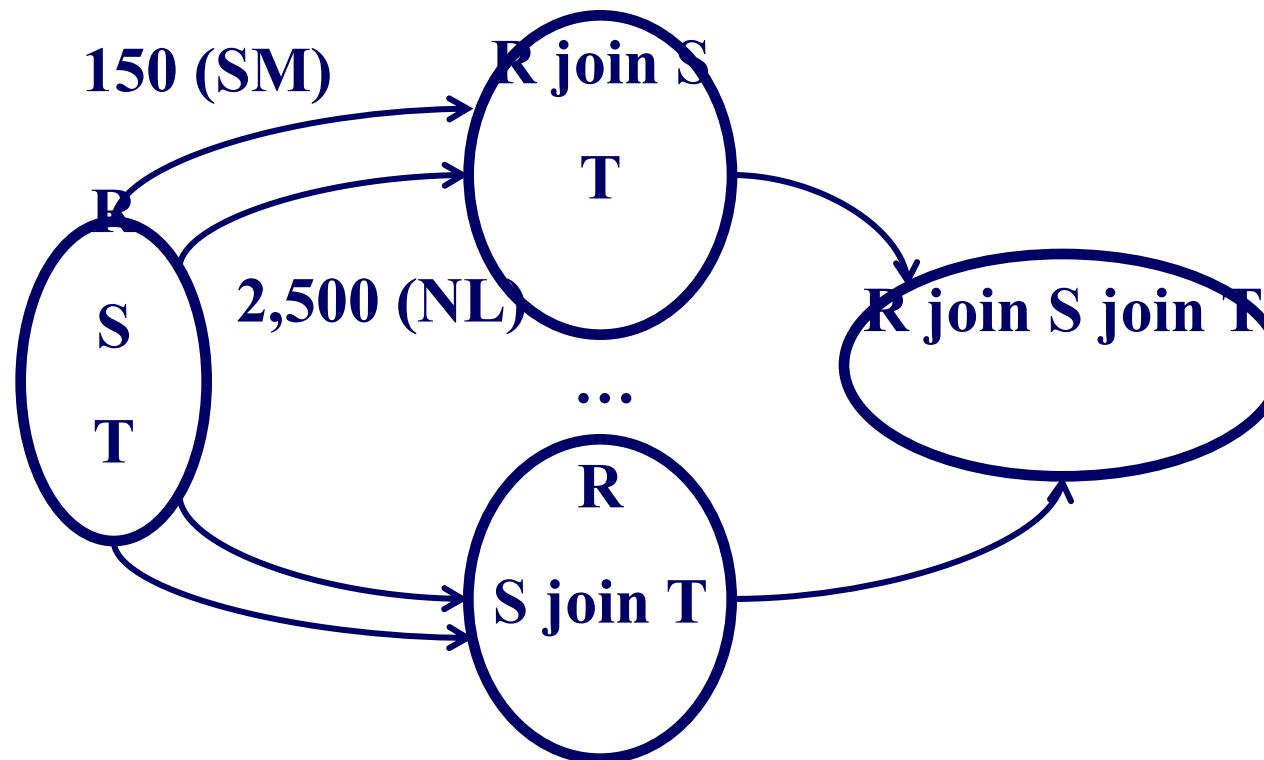
where R.a = S.a and S.b = T.b

order by R.a



Q-opt and Dyn. Programming

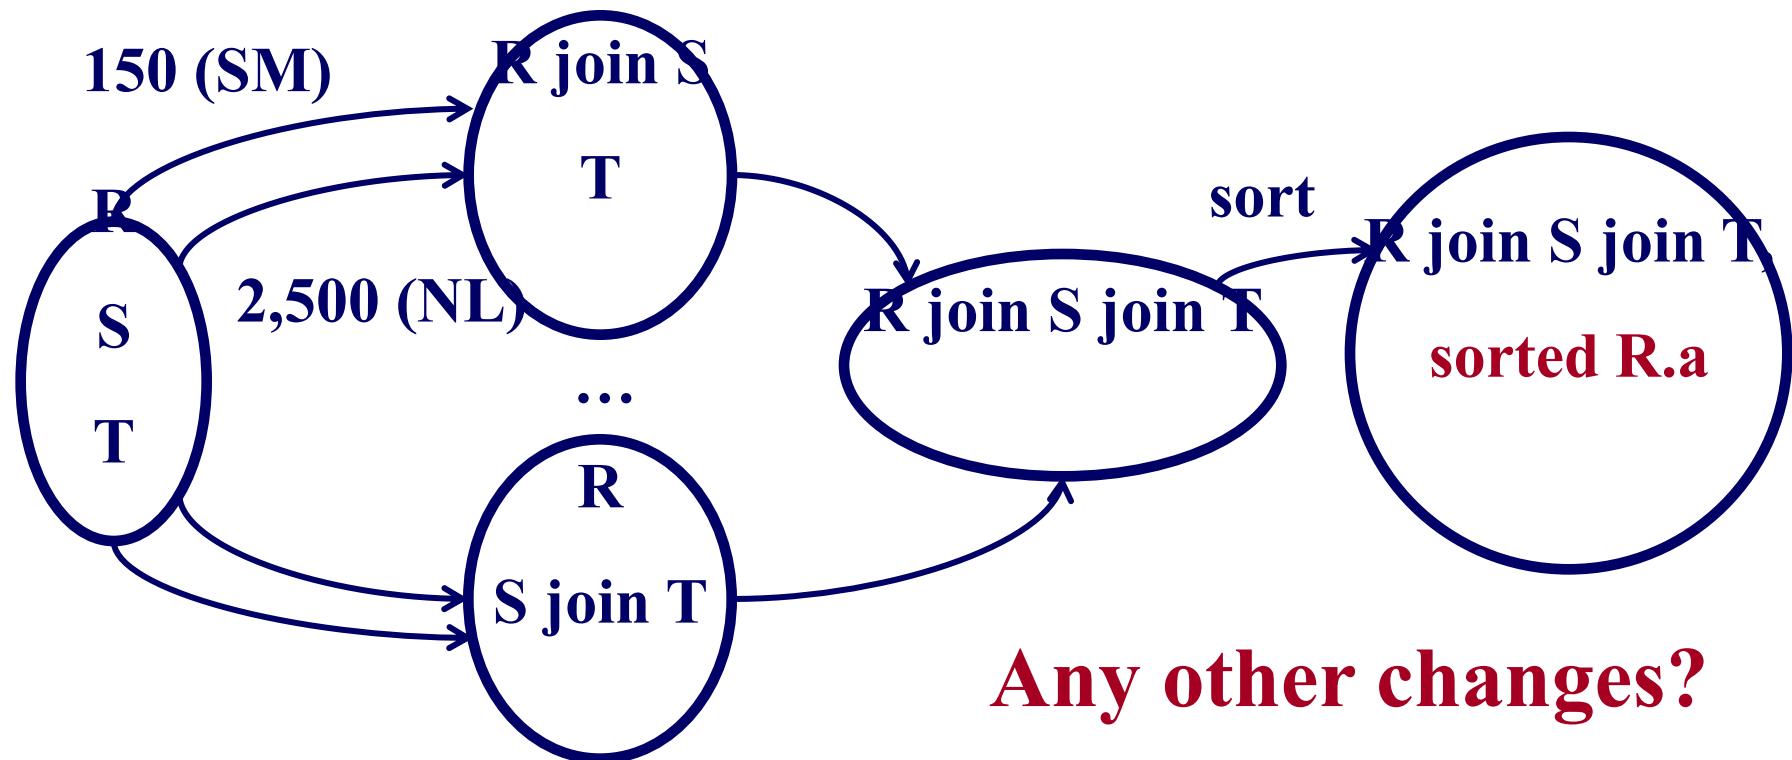
- E.g., compute $R \text{ join } S \text{ join } T$ order by R.a





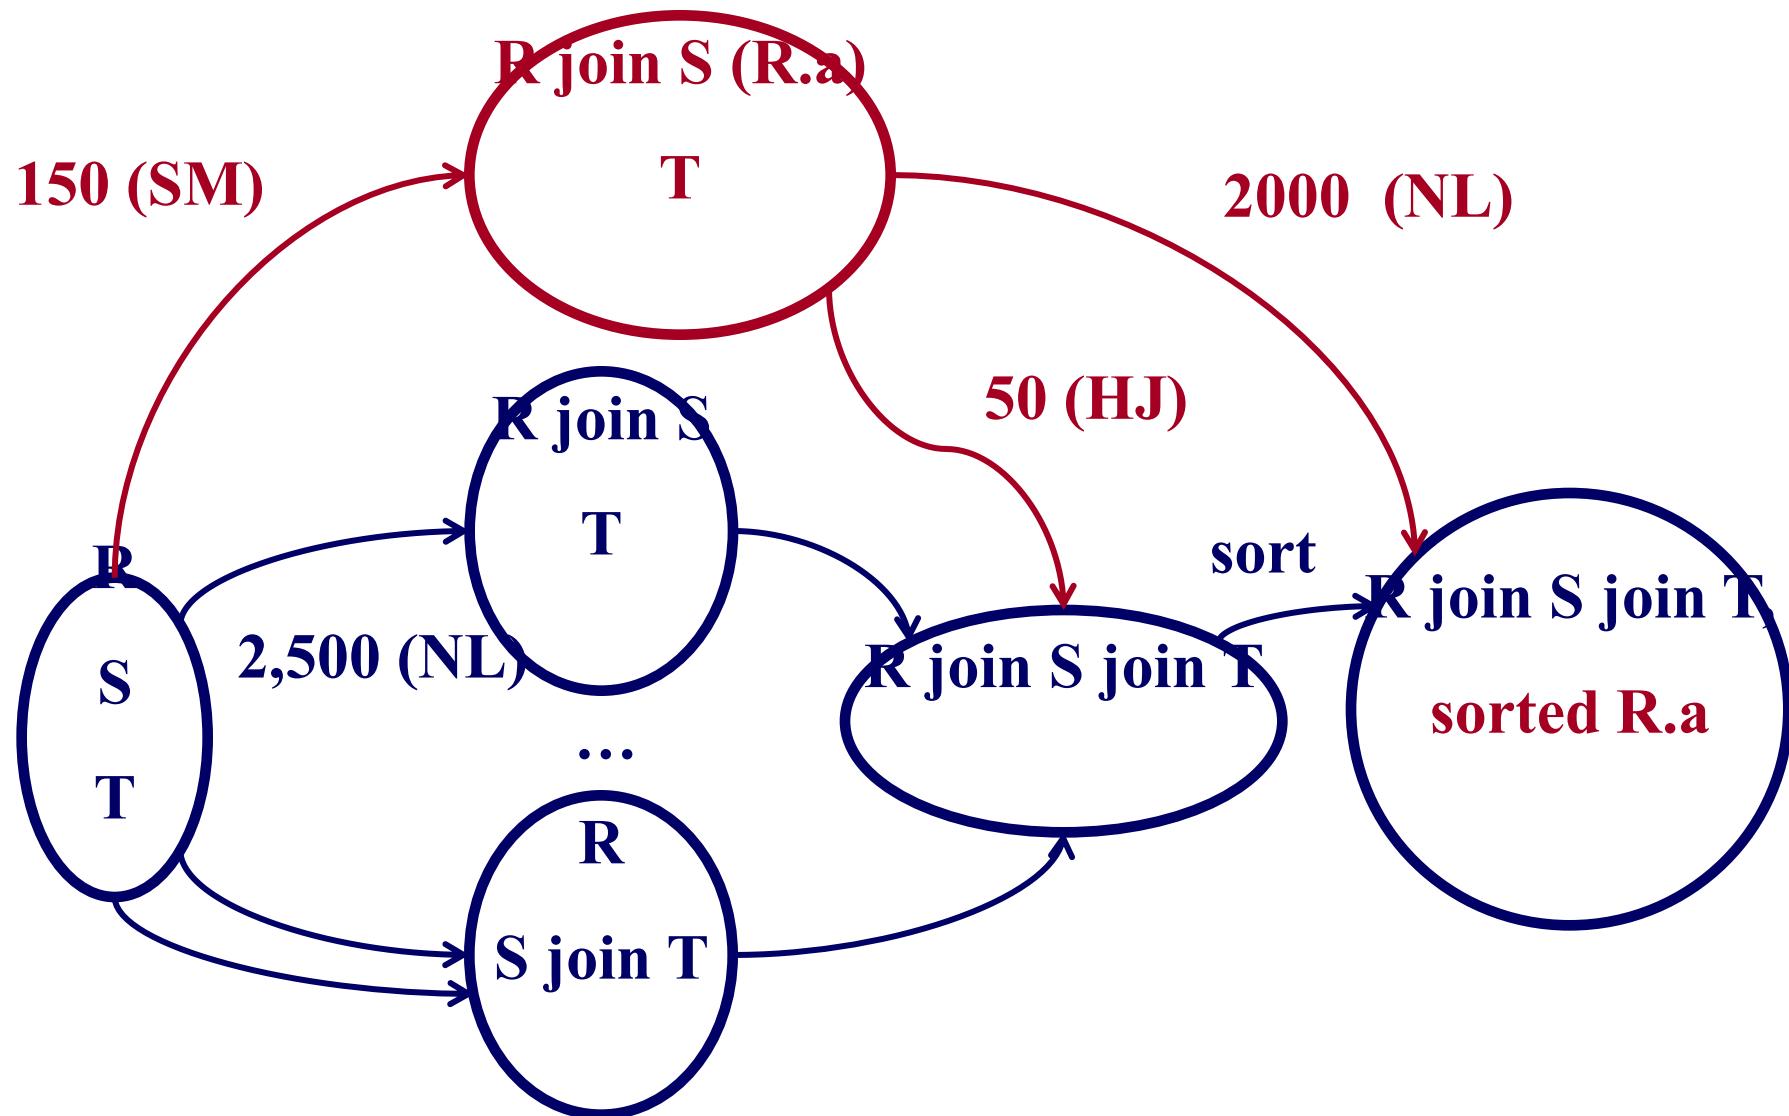
Q-opt and Dyn. Programming

- E.g., compute $R \text{ join } S \text{ join } T$ order by R.a





Q-opt and Dyn. Programming





Q-opt steps

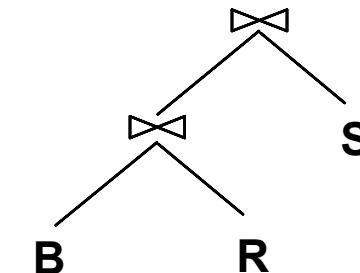
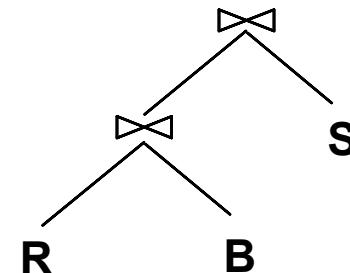
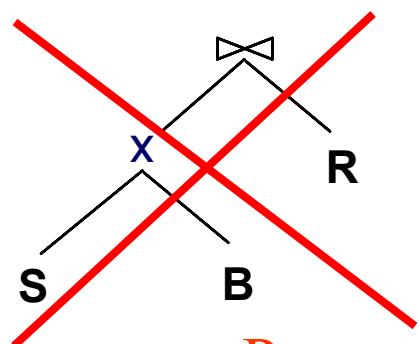
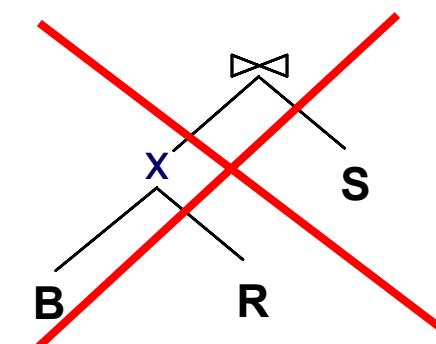
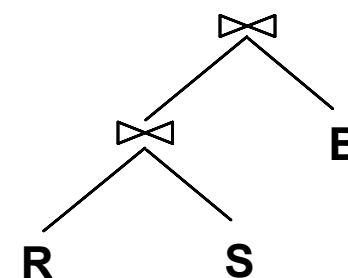
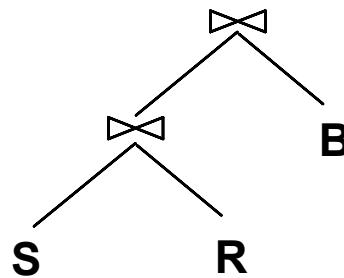
- bring query in internal form (eg., parse tree)
- ... into ‘canonical form’ (syntactic q-opt)
- generate alt. plans
 - single relation
 - multiple relations
 - Main idea
 - Dynamic programming – reminder
 - **Example**
- estimate cost; pick best



Candidate Plans

```
SELECT S.sname, B.bname, R.day  
FROM Sailors S, Reserves R, Boats B  
WHERE S.sid = R.sid AND R.bid = B.bid
```

1. Enumerate relation orderings:



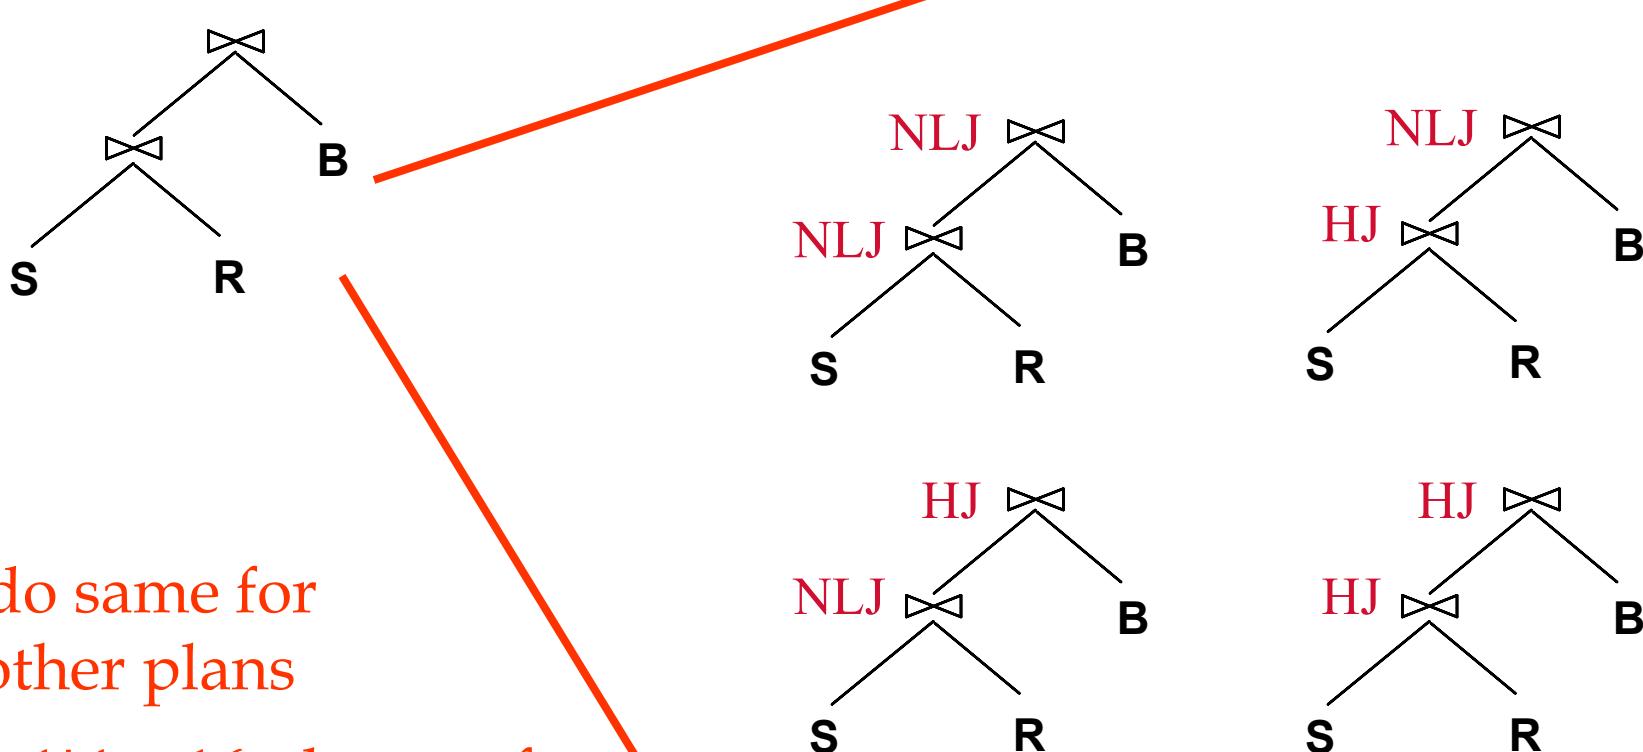
Prune plans with cross-products immediately!



Candidate Plans

```
SELECT S.sname, B.bname, R.day
FROM Sailors S, Reserves R, Boats B
WHERE S.sid = R.sid AND R.bid = B.bid
```

2. Enumerate join algorithm choices:



+ do same for
4 other plans

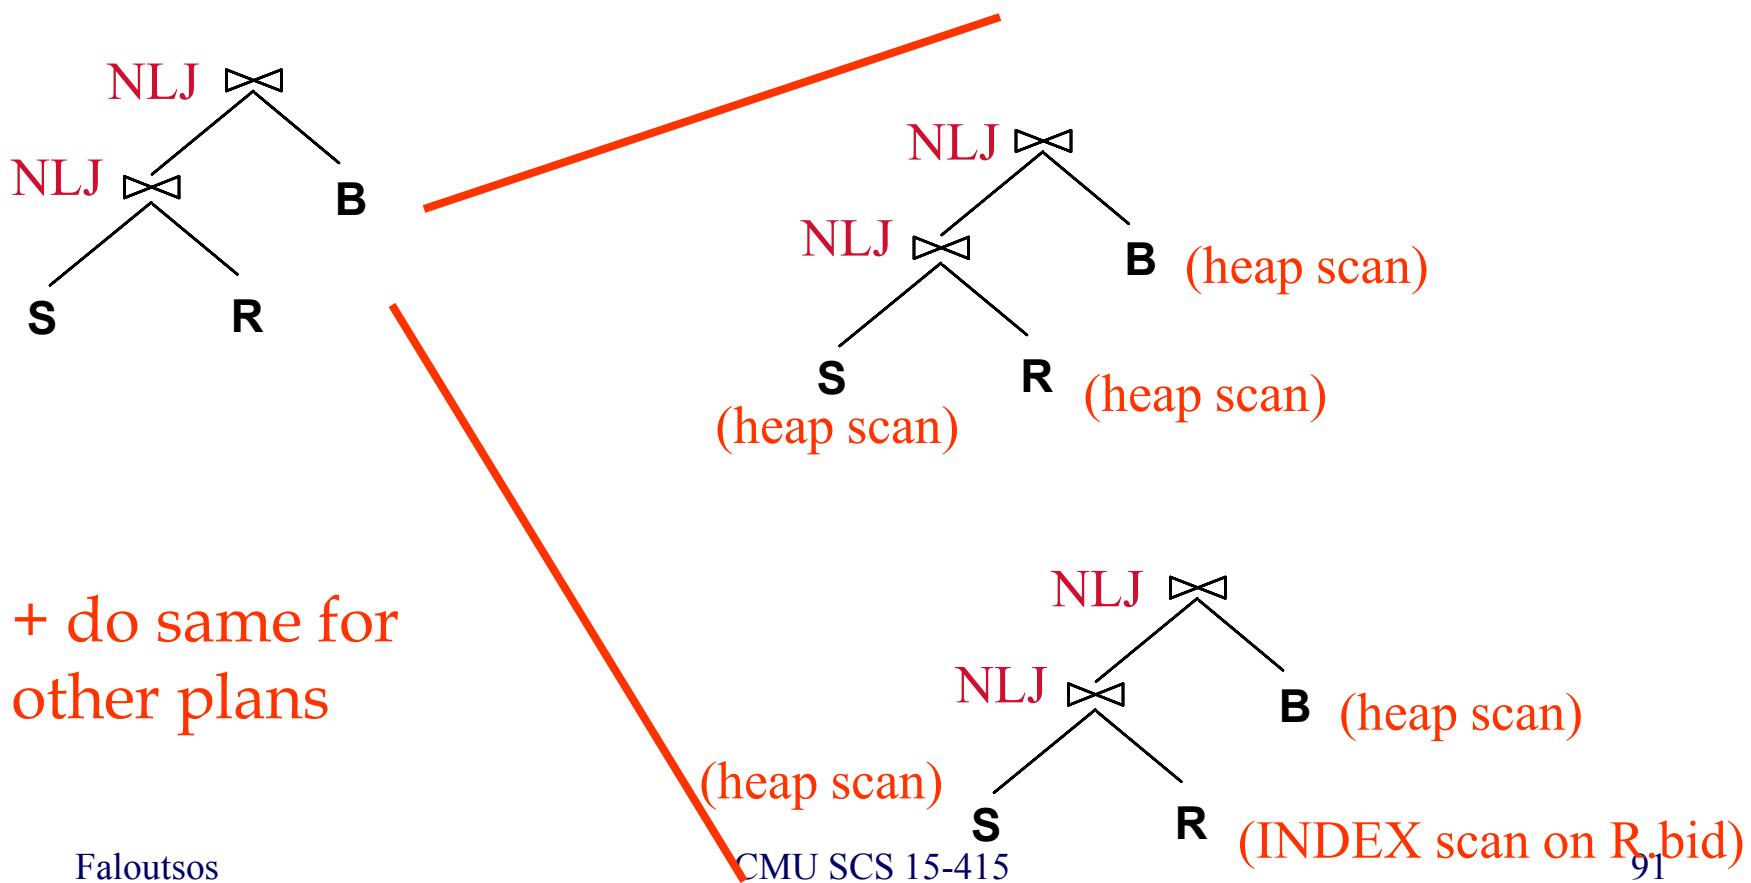
→ $4 \times 4 = 16$ plans so far..



Candidate Plans

```
SELECT S.sname, B.bname, R.day  
FROM Sailors S, Reserves R, Boats B  
WHERE S.sid = R.sid AND R.bid = B.bid
```

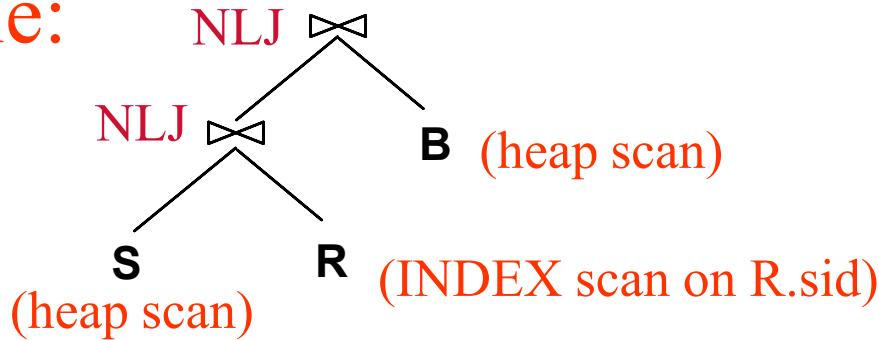
3. Enumerate access method choices:





Now estimate the cost of each plan

Example:





Q-opt steps

- bring query in internal form (eg., parse tree)
- ... into ‘canonical form’ (syntactic q-opt)
- generate alt. plans
 - single relation
 - multiple relations
 - **nested subqueries**
- estimate cost; pick best



Q-opt steps

- Everything so far: about a single query block



Query Rewriting

- Re-write nested queries
- to: **de-correlate** and/or **flatten** them



Example: Decorrelating a Query

```
SELECT S.sid  
FROM Sailors S  
WHERE EXISTS  
(SELECT *  
  FROM Reserves R  
 WHERE R.bid=103  
 AND R.sid=S.sid)
```

Equivalent uncorrelated query:
SELECT S.sid
FROM Sailors S
WHERE S.sid IN
(SELECT R.sid
 FROM Reserves R
 WHERE R.bid=103)

- **Advantage:** nested block only needs to be executed **once** (rather than once per S tuple)



Example: “Flattening” a Query

```
SELECT S.sid  
FROM Sailors S  
WHERE S.sid IN  
(SELECT R.sid  
  FROM Reserves R  
 WHERE R.bid=103)
```

Equivalent non-nested query:

```
SELECT S.sid  
FROM Sailors S, Reserves R  
WHERE S.sid=R.sid  
      AND R.bid=103
```

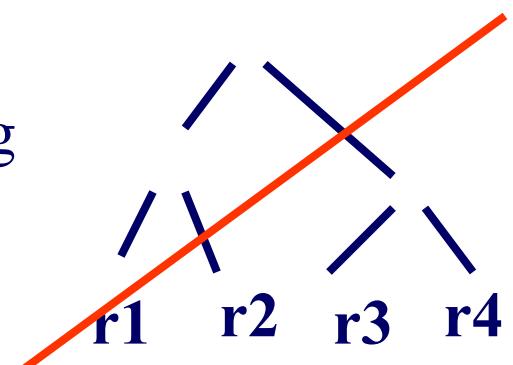
- **Advantage:** can use a join algorithm + optimizer can select among join algorithms & reorder freely



Structure of query optimizers:

System R:

- break query in query blocks
- simple queries (ie., no joins): look at stats
- n-way joins: left-deep join trees; ie., only one intermediate result at a time
 - pros: smaller search space; pipelining
 - cons: may miss optimal
- 2-way joins: NL and sort-merge





Structure of query optimizers:

More heuristics by Oracle, Sybase and
Starburst (-> DB2)

In general: q-opt is very important for large
databases.

(‘**explain select <sql-statement>**’ gives plan)



Q-opt steps

- bring query in internal form (eg., parse tree)
- ... into ‘canonical form’ (syntactic q-opt)
- generate alt. plans
- estimate cost; pick best



Conclusions

- Ideas to remember:
 - syntactic q-opt – do selections early
 - selectivity estimations (uniformity, indep.; histograms; join selectivity)
 - hash join (nested loops; sort-merge)
 - left-deep joins
 - dynamic programming