

Ch 14 More on evaluating Relational Operators

Considers in enough detail to implement as well as estimate cost

Running Example

Reserves (sid: int, bid: int, day: date, rname: string)
tuple size 40, bfr=100, # of pages 1000

Sailors (sid: int, sname: ~~string~~^{string}, rating: ~~int~~^{int}, age: ~~int~~^{real})
tuple size 50, bfr=80, # of pages=500

In determining cost, we

- (1) considers I/Os only
- (2) ignore cost of final output
(since will be same independent of algorithm for evaluation)

Selection

Suppose had

```
SELECT *
```

```
FROM Reserves R
```

```
WHERE R.rname='Joe'
```

called a simple ^{selection} query b/c don't have "and", "or" in where

No index, unsorted Data

$\sigma_{R.attr op value}(R)$ - to evaluate need to scan relation
so would typically need 1000 I/Os.

No index, Sorted Data

$\sigma_{R.attr op value}(R)$ - ~~binary search~~
binary search + scan to output all matches
since in our case ignoring output cost, cost would be $\log_2 1000 \approx 10$ I/Os.

B+ Tree Index

If a clustered B+ index is on $R.attr$ and $op \neq "="$ then best strategy for $\sigma_{R.attr op value}(R)$ is to use the index. Still ok if $op = "="$ but a better choice would be to have a hash index.

To evaluate find the first index entry that pts to a qualifying tuple then scan adjacent leaf pages until have returned all satisfying values.

~~If clustered~~

Suppose 10% of tuples satisfying condition then cost is

$\approx 1000 \times 10 + \text{height of index}$

≈ 100 I/Os

usually small

↖ here are paying attention to output

Note: would be 1000 if unclustered

Hash Index, Equality Selection
cost will be cost to find correct bucket.
+ # of I/Os to return correct tuples from that
bucket (might be in several pages)

So 1072 I/Os to find bucket

General Selects

Selection w/o Disjunction

Could ① do file scan to retrieve tuples
or use an index that matches some
of conjuncts to scan records then
cut from output according to non-primary
conjuncts

② try to use several indexes

Idea sort rids satisfying conditions
represent using bit map
and maps

Selection w/ Disjunction