

# Block Nested Join

Suppose smaller relation <sup>(say R)</sup> can be held in buffer w/ two pages left over.

Then can scan R once into memory and use remaining two pages to do join w/ S.

Total I/O cost  $M+N$  (much better)

Even if don't have enough memory to completely hold smaller rel<sup>n</sup> would try holding  $B-2$  pages of it, yielding:

For each block of  $B-2$  pages of R do:

For each page of S do

{

For all matching in-memory tuples  $r \in R\text{-block}$  &  $s \in S\text{-page}$  output  $(r, s)$

}

Now S is scanned  $\lceil \frac{M}{B-2} \rceil$  times  
So total cost is  $M + N \cdot \lceil \frac{M}{B-2} \rceil$

How do we do matching:

- ① Build a hashtable in memory for the block of R's page.
- ② Use this to look up matches for an S.

Suppose can hold in hash table 100 pages of R  
 Then total cost is  $1000 + 500 \cdot \lceil \frac{1000}{100} \rceil = 6000$

## Sort Merge Join

Algorithm for RAS<sub>asb</sub>

External merge sort R  
 External merge sort S } - result is partitioned  
 Scan R (as much as can into buffers) into set on which join value is constant  
 For each partition of R merge with corresponding partition of S

Cost If join condition is equality on Key  
 $2M \log M + 2N \log N + M + N$   
 Sort R      Sort S

For running example

consider  $\log_{10}$

$$2 \cdot 2 \cdot 1000 = 4000$$

$$2 \cdot 2 \cdot 500 = 2000$$

1000

500

7500

Hash Join (as with ~~any~~ hash based projection)  
 have condition  $B > \sqrt{F \cdot (M+N)}$

hash both relations on the join attribute  
 using same hash f<sub>R</sub>

Produce B-1 partition files containing row from R & S.

Now read in each partition apply a 2<sup>nd</sup> hash f<sub>S</sub> to B-1 partitions (hopefully < F-1 page size). Merge & output these single pages.  $(M+N) \cdot B$

# Overview of Transaction Management

What properties do we want of transactions?

ACID

- atomic - each transaction has all its operations applied to DB or none of them
- consistency - transaction takes DB from consistent state to consistent state
- isolation - <sup>action of</sup> a transaction should not be affected by concurrently running transactions.
- durability - affects <sup>of transaction</sup> should persist even if system crashes (recovery manager)

## Transactions & Schedules

transaction - a list of operations to be applied to DB. Ex withdraw \$40 bank account.

operations - usually reads or writes of DB objects. Written as  $R_T(O)$  or  $W_T(O)$ .

- also have commit (~~complete~~ successfully) and abort operations.  $Commit_T$ ,  $Abort_T$ .

## Assumptions

- transactions do not communicate w/ each other
- A db is a fixed collection of independent objects