Review (1)

- We would like to sort the tuples of a relation R on a given key. The following is known about the relation: R contains 100,000 tuples. The size of a page on disk is 4000 bytes. The size of each R tuple is 400 bytes. R is clustered, i.e., each disk page holding R tuples is full of R tuples. The size of the sort key is 32 bytes. A record pointer is 8 bytes. Answer the following questions:
 - If we use a two pass sorting algorithm, what is the minimum amount of main memory (in terms of number of pages) required?

 - memory (in terms of number of pages) required? What is the cost of the two pags sorting algorithm in terms of number of disk I/Os? Include the cost of writing the sorted file to disk. Consider the following variant of the sorting algorithm. Instead of sorting the entire tuple, we just sort the (key, recordPointer) for each tuple. As in the conventional two pass sorting algorithm, we sort chunks of (key, recordPointer) in main memory and write the chunks to the tuple (from the original copy of R) and write the sorted relation to disk. What is the cost in terms of number of disk I/Os? Coaning all other parematers constant for what submer of tubes of tube arise in the variant
 - Keeping all other parameters constant, for what values of tuple size is the variant discussed above better (in the number of I/Os)?

Review (2)

- $\sqrt{|\mathbf{R}|} + 1 = 101$, where $|\mathbf{R}|$ denotes the size of R in pages
- 2 X 2 X |R| = 40000
- Memory required = 34 (an additional page is needed for the random access step in the second phase)
- This is an optimized version. The I/Os of the sorting scheme is 122000. This includes 10000 for initially reading R and constructing (key, recordPointer) pairs; 1000 I/Os for writing the the same from disk to merge the runs; 1000 to random access to retrieve the tuples pointed by the record pointer; and finally 10000 I/Os to write the sorted relation R to disk
- Assume that records are unspanned, then tuplesize > 2001















Sailor	r						Join Example						
sid	sn	ame	rati	ng	age		sid	bid	day	rname	:		
22	dustin		7	45.			31	101	10/11/96	lubber			
28	yu	рру	9		35.0		58	103	11/12/96	dustin			
31	lubber		8		55.5								
44	gu	рру	5		35.0								
58	rusty		1	0 35.0									
Query (join) output													
S	id	snam	e ratir		ng	<u>age</u>		bid	<u>day</u>	rname			
3	1	lubber		8		55.5		101	10/11/96	lubber			
58		rusty		10		35.0		103	11/12/96	dustin			
\$5208													

















CS5208

Example of Sort-Merge Join									
					sid	bid	<u>day</u>	rname	
<u>sid</u> 22 28 31 44 58	sname dustin yuppy lubber guppy	rating 7 9 8 5 10	age 45.0 35.0 55.5 35.0 35.0		28 28 31 31 31 58	103 103 101 102 101 103	12/4/96 11/3/96 10/10/96 10/12/96 10/11/96 11/12/96	guppy yuppy dustin lubber lubber dustin	
Cost	?	10	33.0	I	50	105	11/12/90	uustiii	
CS5208									

Examples of Block Nested Loops



#outer blocks = no. of pages in outer relation / block size

Example of Sort-Merge Join										
					sid	<u>bid</u>	<u>day</u>	rname		
<u>sid</u> 22 28 31 44 58	sname dustin yuppy lubber guppy rusty	rating 7 9 8 5	age 45.0 35.0 55.5 35.0 35.0		28 28 31 31 31 58	103 103 101 102 101 103	12/4/96 11/3/96 10/10/96 10/12/96 10/11/96	guppy yuppy dustin lubber lubber dustin		
$[58]$ [rusty 10 $[55.0]$ $[58]$ 103 $11/12/96$ dustin • Cost: $2M^*K_1 + 2N^*K_2 + (M+N)$ • K ₁ and K ₂ are the number of passes to sort R and S respectively • The cost of scanning, M+N, could be M*N (very unlikely!)										
CS5208										



















Examples of Index Nested Loops

Hash-index on sid of S (as inner):

- Scan R: 1000 page I/Os, 100*1000 tuples.
- For each R tuple: 1.2 I/Os to get data entry in index, plus 1 I/O to get (the exactly one) matching S tuple. Total: 220,000 I/Os.
- Hash-index on sid of R (as inner):
- Scan S: 500 page I/Os, 80*500 tuples.
- For each S tuple: 1.2 I/Os to find index page with data entries, plus cost of retrieving matching R tuples.
- Assuming uniform distribution, 2.5 reservations per sailor (100,000 / 40,000). Cost of retrieving them is 1 or 2.5 I/Os depending on whether the index is clustered.









CS5208



Two Approaches to General Selections

- First approach: Find the most selective access path, retrieve tuples using it, and apply any remaining terms that don't match the index:
 - Most selective access path: An index or file scan that we estimate will require the fewest page I/Os.
 - Terms that match this index reduce the number of tuples retrieved; other terms are used to discard some retrieved tuples, but do not affect number of tuples/pages fetched.
 - Consider day<8/9/94 AND bid=5 AND sid=3. A B+ tree index on day can be used; then, bid=5 and sid=3 must be checked for each retrieved tuple. Similarly, a hash index on <bid, sid> could be used; day<8/9/94 must then be checked.

Intersection of Rids

- <u>Second approach</u> (if we have 2 or more matching indexes (assuming leaf data entries are pointers):
 - · Get sets of rids of data records using each matching index.
 - Then intersect these sets of rids (we'll discuss intersection soon!)
 - · Retrieve the records and apply any remaining terms.
 - Consider day<8/9/94 AND bid=5 AND sid=3. If we have a B+ tree index on day and an index on sid, we can retrieve rids of records satisfying day<8/9/94 using the first, rids of recs satisfying sid=3 using the second, intersect, retrieve records and check bid=5.

CS5208



Set Operations Set Operations · Intersection and cross-product special cases of join. · Intersection and cross-product special cases of join. · Union (Distinct) and Difference similar. · Union (Distinct) and Difference similar. · Sorting based approach to union: · Sorting based approach to union: · Sort both relations (on combination of all attributes). · Sort both relations (on combination of all attributes). · Scan sorted relations and merge them. · Scan sorted relations and merge them. · Hash based approach to union: · Hash based approach to union? · Partition R and S using hash function h. · For each S-partition, build in-memory hash table (using h2), scan corr. R-partition and add tuples to table while discarding duplicates.

CS5208



Iterators for Implementation of Operators

- · Most operators can be implemented as an iterator
- An iterator allows a consumer of the result of the operator to get the result one tuple at a time
 - Open starts the process of getting tuples, but does not get a tuple. It initializes any data structures needed.
 - GetNext returns the next tuple in the result and adjusts the data structures as necessary to allow subsequent tuples to be obtained. It may calls GetNext one or more times on its arguments. It also signals whether a tuple was produced or there were no more tuples to be produced.
 - Close ends the iteration after all tuples have been obtained.

CS5208











- A virtue of relational DBMSs: queries are composed of a few basic operators; the implementation of these operators can be carefully tuned (and it is important to do this!).
- Many alternative implementation techniques for each operator; no universally superior technique for most operators.
- Must consider available alternatives for each operation in a query and choose best one based on system statistics, etc. This is part of the broader task of optimizing a query composed of several ops.