

SE305 Database System Technology

Assignment 5 (due Oct 23, 2018)

- (20') Give an example of a relational-algebra expression and a query-processing strategy in each of the following situations:
 - MRU is preferable to LRU.
 - LRU is preferable to MRU.
- (20') In the BCNF decomposition algorithm, suppose you use a functional dependency $a \rightarrow b$ to decompose a relation schema $r(a, b, g)$ into $r_1(a, b)$ and $r_2(a, g)$.
 - What primary and foreign-key constraint do you expect to hold on the decomposed relations?
 - Give an example of an inconsistency that can arise due to an erroneous update, if the foreign-key constraint were not enforced on the decomposed relations above.
 - When a relation is decomposed into 3NF, what primary and foreign key dependencies would you expect will hold on the decomposed schema?
- (20') Consider the data and parity-block arrangement on four disks depicted in Figure. The B_i s represent data blocks; the P_i s represent parity blocks. Parity block P_i is the parity block for data blocks B_{4i-3} to B_{4i} . What, if any, problem might this arrangement present?

Disk 1	Disk 2	Disk 3	Disk 4
B_1	B_2	B_3	B_4
P_1	B_5	B_6	B_7
B_8	P_2	B_9	B_{10}
\vdots	\vdots	\vdots	\vdots

- (21') Consider the deletion of record 5 from the file of Figure. Compare the relative merits of the following techniques for implementing the deletion:
 - Move record 6 to the space occupied by record 5, and move record 7 to the space occupied by record 6.
 - Move record 7 to the space occupied by record 5.
 - Mark record 5 as deleted, and move no records.

record 0	10101	Srinivasan	Comp. Sci.	65000
record 1	12121	Wu	Finance	90000
record 2	15151	Mozart	Music	40000
record 11	98345	Kim	Elec. Eng.	80000
record 4	32343	El Said	History	60000
record 5	33456	Gold	Physics	87000
record 6	45565	Katz	Comp. Sci.	75000
record 7	58583	Califieri	History	62000
record 8	76543	Singh	Finance	80000
record 9	76766	Crick	Biology	72000
record 10	83821	Brandt	Comp. Sci.	92000

5. (19') It is important to be able to quickly find out if a block is present in the buffer, and if so where in the buffer it resides. Given that database buffer sizes are very large, what (in-memory) data structure would you use for the above task?.