UNIVERSITY OF MANNHEIM
School of Business Informatics and Mathematics

Data and Web Science Group
Prof. Dr. Heiko Paulheim

Database Technology

Exercise 7: Database Architectures

7.1. File Organization

a. Consider the deletion of records 5 from the file shown below. Compare the relative merits of the following techniques for implementing the deletion:
   1. Move record 6 to the space occupied by record 5, and move record 7 to the space occupied by record 6.
   2. Move record 7 to the space occupied by record 5.
   3. 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 3	22222	Einstein	Physics	95000
record 4	32343	El Said	History	60000
record 5	33456	Gold	Physics	87000
record 6	45565	Katz	Comp. Sci.	87000
record 7	58583	Califieri	History	62000

b. Show the structure of the file shown below after each of the following steps:
   1. Insert (24556, Turnamian, Finance, 98000).
   2. Delete record 2.
   3. Insert (34556, Thompson, Music, 67000).

header				
record 0	10101	Srinivasan	Comp. Sci.	65000
record 1				
record 2	15151	Mozart	Music	40000
record 3	22222	Einstein	Physics	95000
record 4				
record 5	33456	Gold	Physics	87000
record 6				
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
record 11	98345	Kim	Elec. Eng.	80000

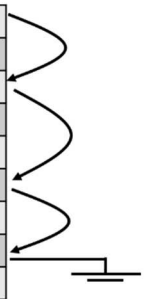

7.2. Table of data storage

List the physical storage media available on the computers you use routinely. Give the speed with which data can be accessed on each medium, the average amount of data which can be stored, an estimated price per GB and if the data needs power to be stored.

7.3. Buffer Replacement Strategies

Give a query-processing strategy for the relational-algebra expression $R_1 \bowtie R_2$ where:

a. MRU is preferable to LRU
b. LRU is preferable to MRU

7.4. Shared Structures

Instead of storing shared structures in shared memory, an alternative architecture would be to store them in the local memory of a special process, and access the shared data by interprocess communication with the process. What would be the drawback of such an architecture?

7.5. Speedup

Assume we have two computers A(1,2 GHz) and B (3,6 Ghz). Compute the speedup given that A needs 35 ms and B needs 15 ms for answering a query. Check if it is linear or sublinear.

7.6. Data Servers

Consider a database system based on a client–server architecture, with the server acting as a data server.

a. What is the effect of the speed of the interconnection between the client and the server on the choice between tuple and page shipping?
b. If page shipping is used, the cache of data at the client can be organized either as a tuple cache or a page cache. The page cache stores data in units of a page, while the tuple cache stores data in units of tuples. Assume tuples are smaller than pages. Describe one benefit of a tuple cache over a page cache.