

Homework 1 Data Storage

Due: *Thursday, May 1, in-class*

Submission. Please submit your solutions on May 1, in-class. I encourage everyone to use text-processing software for the solutions and use Power-Point, xfig or other graphical tools to draw pictures/diagrams required in this homework. (I will accept handwritten solutions, but I discourage them. This policy is mainly to ensure expedient and error-free grading).

Problem 1

You will consider two different schemas for storage of data on disk. For each schema, you have to provide descriptions of (a) the page header and (b) records for each of the relational tables described below. You also have to specify both the page header size and the size of the record, and determine how many records can be stored on a single disk page.

In both problems, INT has size 4 and FLOAT has size 8 (bytes).

Storage Schema A.

The size of a disk page is 2Kb. You reserve 64 bytes for the constant-size part of the page header. The file organization is *heap*. You store a bitmap of available slots on the disk page in the header (this is the variable size portion of the page header. The total size of the page header is the sum of the sizes of both parts of the header). The organization of the constant-size part of the header is left up to you, but you need to ensure that the header contains all necessary information to successfully retrieve/maintain records.

All records are constructed by concatenating the field representations of individual attributes. There are no record headers, and you resolve the constraints on the starting positions of the fields by padding. Each record must start at an offset divisible by 8, which entails that the size of the page header must also be divisible by 8.

Deleted records are marked in the bitmap, and nowhere else. You can assume that the relation schema is stored in the file header page, i.e., your pages need not keep this information.

Storage Schema B.

The size of a disk page is 1Kb. You reserve 64 bytes for the page header. The file organization is *sequential*. You use *sliding* to reclaim space during deletions, which means that the free space on your pages can occur only at the end of the block. Your header must contain all information necessary for maintenance of data.

All records are constructed by concatenating the field representations of individual attributes. In addition, an 8-byte header is attached to each record. It includes two timestamp fields - one for the last read access and one - for the last write access to the record. You resolve the constraints on the starting positions of the fields by reordering the fields in the record, and adding and additional padding, if needed to the end of the record, to ensure that all records start with the offsets divisible by 8.

Relations

Below is a list of relations you have to use. Note, that these are standalone relations, NOT a list of relations forming a coherent database.


1. Employee(Id INT, SSN INT, Name VARCHAR(60), Department INT, Salary FLOAT, Bonus FLOAT, StartYear INT, Comment CHAR(80)).
2. Team(Id INT, Name CHAR(20), City CHAR(20), Owner INT, Division CHAR(20)).
3. Survey(RespondentId INT, Sex CHAR(1), Income INT, Age INT, Race INT, Q1 INT, Q2 INT, Q3 INT, Q4 INT, Q5 INT).
4. BookMarks(Id INT, URL CHAR(256), DateAdded DATE, Category INT, Note VARCHAR(200), LastVisited DATE).
5. Transcript((Id INT), FirstName CHAR(20), LastName CHAR(20), College CHAR(5), Status CHAR(2), Course CHAR(7), Semester INT, Year INT, Grade INT, CurrentGPA FLOAT).
6. Transaction(Id INT, DateTimePosted DATE, DateTimeOccurred DATE, CardNo CHAR(16), Location CHAR(40), Vendor INT, Amount FLOAT, FeeRate FLOAT, Fee FLOAT).

Questions

For each storage schema, and for each relation, do the following:

1. Describe the structure of the disk record storing a tuple from the relation. Draw a diagram (see below). Determine the size of the record.
2. Describe the structure of the page header (draw a diagram), determine its size (if not constant).
3. Determine the number of records that can be stored on a single disk page.

Notes. Assume DATE is represented as a single INT. For each schema, your descriptions of the page header must be consistent from relation to relation (in fact, you are allowed to specify it just once, and only show the size of the page header for each relation). A diagram of a record and/or header should look similar to the following diagram (created for Schema A storage for the relation R(A INT, B CHAR(2), C CHAR(7), D FLOAT, E VARCHAR(8), F INT)).



Submission Note. Please number your answers as follows: RelationNo Schema.Question. This way 1A.2 represents the answer to the question 2, for relation 1 using storage schema A.

Problem 2

Consider a database consisting of three data files, P, T and R. Each file consists of a number of disk pages (blocks). We refer to block number i from file F as Fi , e.g., P1 is the first page of file P, T20 is the 20th page of file T and so on.

The DBMS controlling the database is equipped with a buffer manager **B**. The size of the buffer space is **4 disk blocks**.

Consider the following possibilities for **B**.

- B1 : LRU (Least Recently Used) buffer manager. This manager uses the least recently used page and flushes it to disk, when a new page needs to be brought from disk.
- B2 : FIFO (First In, First Out) buffer manager. This manager flushes the oldest page in the buffer.
- B3 : Simple Clock buffer manager. This buffer manager uses the simple clock algorithm (starting with buffer 1 of the buffer space) to determine the next buffer to be flushed.

The query execution layer of the DBMS uses four operations to control data access:

- Read(Pageld): results in the page with the given Pageld being transferred to the buffer.
- Write(Pageld): results in the page with the given Pageld being marked as ready for transfer back to disk. If the page is NOT in the buffer, it is **transferred** back to the buffer, using the buffer manager's management strategy, and is marked *ready* afterwards.
- Pin(Pageld): given page is *pinned* in the buffer.
- Unpin(Pageld): given page is *unpinned* in the buffer.

All buffer managers first look to flush back to disk any pages marked as *ready* by the Write() command. If no such pages found, the respective buffer management strategy is engaged. (flushing a *ready* page does not affect the position of the pointer in the clock algorithms).

For each of the following two sequences of commands and each of the buffer management strategies outlined above, show the state of the buffer after each 8 operations. (marked as “**checkpoint**” in the sequence). For each of the four buffers in the buffer space, show the id of the page stored in it, and whether the *ready* and/or *pin* flags have been set. For the clock-based buffer managers show current counter values, and the position of the clock “hand”. For LRU and FIFO managers, show the timestamps (use the position in the sequence below as the timestamp).

In all cases assume you start with empty buffer space. For prioritized clock buffer manager, assume that all pages from T and P are added with counter values set to 1, while pages from R are added with counter values set to 2.

Sequence 1

```
Read(T1);
Read(T2);
Read(P2);
Pin(P2);
Read(R3);
Read(T2);
Read(T1);
Read(R4);
    Checkpoint;
Write(R4);
Unpin(P2);
Write(P2);
Read(T3);
Pin(T3);
Write(R3);
```

```
Read(R4);
Read(R5);
    Checkpoint;
Unpin(T3);
Write(T3);
Read(T1);
Read(T2);
Write(P2);
Pin(T1);
Read(R6);
Read(R7);
    Checkpoint;
```

Sequence 2

```
Read(P1);
Read(R1);
Read(R2);
Read(R3);
Read(R4);
Read(P2);
Read(T2);
Read(P1);
    Checkpoint;
Read(R3);
Pin(P1);
Read(R4);
Read(R5);
Read(T1);
Read(T2);
Write(T2);
Read(T4);
    Checkpoint;
Read(R1);
Unpin(P1);
Read(R2);
Read(R3);
Read(T2);
Read(T3);
Read(P1);
Read(T2);
    Checkpoint;
```