King Fahd University of Petroleum and Minerals

Information and Computer Science Department

ICS 541: Database Design & Implementation

DO NOT OPEN UNTIL INSTRUCTED TO DO SO!!!!

Write clearly, precisely, and briefly!!

ID:							
Name:							

Grades		
Section	**Max**	**Scored**
A	**8**	
B	**12**	
C	**16**	
D	**6**	
E	**8**	
TOTAL	**50**	

A. Questions from the paper of home work 3 [8 points]

Assume the following SQL statements in answering the following questions:

```
CREATE TABLE STUDENT ( Id CHAR(7),
                       Name VARCHAR(100),
                       FullTime    BOOLEAN,
                       DOB         DATE);

INSERT into Student values("1234567", "AAA", TRUE, "16-10-1986");
INSERT into Student values("2345678", "BBBB", TRUE, "17-11-1987");
INSERT into Student values("3456789", "CCCCC", TRUE, "18-12-1988");

SELECT * FROM  Student WHERE Fulltime = TRUE;
```

1. Draw what the PAX, DSM, and NSM block or blocks look like.
2. Explain in detail why the above SELECT statement will run faster in PAX than in the NSM or DSM.

B. Questions from Chapter 12. [12 pts]

Use the SQL statements in the previous question in answering the following questions.

1. Draw the layout of the first record inserted, when fields are required to start at multiples of 4. (Assume DATE takes 7 bytes and don't forget record header)

2. Assume the three inserted records reside in the same block. Draw how the records will be organized in the block.. Don't forget block header.

3. Explain what pointer swizzling mean.

4. Breifly mention two strategies we can use to determine when to swizzle pointers.

C. Questions from Chapter 13. [16 pts]

I. Suppose an index block can hold a maximum of 99 keys and 100 pointers and a data block can hold a maximum of 20 data records. Also assume that average B+-tree node is 66% full. If a table of one million data records is organized using B+-tree and the size of one disk block is 4k, then **[12 pts**]

1. What is the number of data records?
2. What is the number of index blocks?
3. What is the depth of the B+tree?
4. How much disk accesses do you need to answer an exact match query?

II. Explain the difference between linear and extensible hashing **[4pts]**

D. **Questions from Chapter 14** **[6 pts]**

1. Encode the following bit-vector using run-length encoding

   000000110000000000010100000

E. **Questions from Chapter 15** **[8 pts]**

Suppose B(R) = B(S) = 5000 and M = 1000. Calculate the disk I/O of a sort based join.