

NAME _____

In questions that follow, **SHOW ALL WORK FOR PARTIAL CREDIT!** In the short-answer questions that follow, you do not **HAVE** to cite statements in the text that support your answer; however, you should **KNOW** of such statements (they always exist) to be sure of your answers!

1. (20 pts.) For each 4 point question below, circle True or False. Be careful!

- (a) To grant a user access to a view, it is necessary to first grant access to the underlying base tables used in the view. **True**
False
- (b) It is possible to add a NOT NULL constraint to a column of a Table using an ALTER TABLE statement in DB2 UDB. **True** **False**
- (c) There is no way to create a constraint that will require each row in the PRODUCTS table to have at least one row in the ORDERS table referencing it (i.e., with the same pid value). **True** **False**
- (d) One can create a view V in Oracle that permits a user with full privileges on the view V to update some columns of V and not others. **True** **False**
- (e) In Oracle, a Bitmap index K4X on column K4 with int values 1 to 4, on a table with 5,000 rows, will require about 10,000 bytes of disk space, assuming all Bitmaps are verbatim (uncompressed). **True** **False**

2. (15 pts, 5 pts each.) Consider two tables T1 and T2, where T1 has three int columns C1, C2, and C3, none defined NOT NULL, with a constraint clause: primary key (C1, C2, C3).

(a) Is it possible for C1 to be NULL on any row of T1? (Circle one) **YES** **NO**

(b) Write the FOREIGN KEY clause you would use in creating table T2 with columns D1, D2, and D3, to make them (as a group) reference C1, C2, and C3 in T1 and give the clause a constraint name .

(c) Assume in ORACLE that **no** ON DELETE clause is added to the foreign key definition of part (b); if a Delete command were issued to delete a row in T1 with (C1, C2, C3) values that were referenced by (D1, D2, D3) values in a row of T2, what would happen as a result of this Delete command?

3. (13 pts) Suppose you need to provide your CAP customer data to a user philb, but you don't want him to see the customer discounts.

(a) (8 pts) Give the SQL statements (including a create view statement) needed to allow philb to run queries on your customer table data but not be able to access the discnt column.

(b) (5 pts) Is the view you created in (a) an updatable view? Explain why or why not.

4. (12 pts, 4 pts each) For each of the following column and query descriptions on a large table (10,000,000 rows), list an appropriate index type (*B-tree*, *Hash Index*, or *Bitmap index*) that is appropriate to use. Provide the BEST CHOICE in each question, assuming performance is VERY important in each case.

(a) Unique column values, no range retrievals by this column (all unique row retrievals):

Index type(s):

(b) Non-unique column values, three or four rows for each value, range retrievals common:

Index type(s):

(c) Only a few (about 25) column values; retrieve by single column value or occasionally by range.

Index type(s):

5. (18 pts) In DB2 (4 KByte pages), consider a table `prospects` of 10,000,000 rows with 100 bytes each, and assume the following command has been issued to load the `prospects` table.

```
create table prospects (prid char(8) not null unique primary key, . . . <other columns>
    pctfree 50;
```

Show calculations in all questions that follow.

(a) (5 pts) What is the number of 4 KByte pages on which the table `prospects` sits.

(b) (8 pts) Assume we are performing a select statement like this:

```
select prid from prospects where hobby = 'chess';
```

Assume chess is a valid value for hobby, and about 1/50 of the rows will be retrieved. (How many rows?)

Give a "Darts in Slots" formula to estimate how many total pages this set of rows will sit on. First give the formula in terms of N and M , and then in terms of real numbers that arise in this problem. **You do not** have to calculate the right answer, just leave it as a formula.

(c) (5 pts) If we have memory buffer for disk pages that will contain 100,000 pages, and we are running the query in part (b) in the morning, with no competition for buffer pages, will any page that's been read in once ever fall out of buffer because of another page being read in for the query in part (b)? Ignore index pages, just count pages in table `prospects`. HINT: the correct power of e in the formula of part (b) has value 0.67. Explain your reasoning.

6. (22 pts.) Writing ESQL program. Consider a "secusts" table created in Oracle as follows:

```
create table secusts as
  select * from customers where city in ('Dallas', 'Fort Worth');
```

We assume customers and secusts both exist in a database and want you to keep secusts up to date by using ESQL application functions `custins()` and `custdel()` that are always called to insert or delete any row in customers (we'll ignore updates). The two ESQL functions, with arguments, are `custsins(custid, custname, custcity, custdiscnt)` and `custdel(custid)`. Write the logic (C code) below that you'd use. (You only need to write C code and Insert/Delete statements to change customers and secusts in these functions; don't worry about declarations of variables, `sqlca`, etc.). I'll start you off.

```
/* Start by writing the logic for the ESQL function custins(custid, custname, custcity, custdiscnt).
Use these same variables in your logic: begin by inserting the rows to customers */
```

```
custins(custid, custname, custcity, custdiscnt) -- include this three lines in your solution
{
  exec sql insert into customers
    values (:custid, :custname, :custcity, :custdiscnt);
```

```
/* Now you need to check if this customers row belongs in secusts (use strcmp() to compare
strings); if so, insert it in secusts. */
```

```
/* Next write the logic for the ESQL function custdel(custid); a bit of thought is required as to
how to determine if the row deleted is in secusts. */
```

Solutions to SP04 CS634 Exam 1

1. (a) FALSE (See last sentence on pg 443.) (b) FALSE (See pg. 422, Alter Table Statement, par 3, sentence starting on line 3; ORACLE has the capability to Modify a column Fig 7.7, but not DB2 UDB.) (c) TRUE (See statement of Exercise 7.4 (d)); (d) FALSE (check option doesn't necessarily keep you from updating some columns of a view -- only certain updates); (e) FALSE (2500 bytes)
2. (a) NO (See pg 414, first sentence.) (b) constraint constrname foreign key(D1, D2, D3) references T1 (c) The Delete would fail (See last sentence in Definition 7.1.3 on pg 414.)
3. (a) `create view philcusts as select cid, cname, city from customers;`
`grant select on philcusts to philb;`
 (b) Yes, by Figure 7.15, but of course philb can't update through it because the grant contained only the select privilege.
4. (a) Hash Index (B-tree is also OK, but Hash Index MUST be named.) (b) B-tree (Hash Index is NOT acceptable with range retrievals, Bitmap index is possible but not optimal.) (c) Bitmap index (B-tree is also OK but not optimal, Hash Index is NOT acceptable.)
5. (a) Since $pctfree = 50$, we fit 20 rows per page; 10 million rows / 20 rows/pg = 500,000 pgs (b) There will be $10,000,000/50 = 200,000$ rows retrieved. $N = 200,000$, $M = 500,000$, so number of pgs hit = $M(1 - e^{-N/M}) = 500000(1 - e^{-200000/500000}) = 500000(1 - e^{-0.4})$ (c) $500,000*(1 - 0.67) = 165,000$ is greater than 100,000 so pages WOULD BE replaced.
5. `custins(custid, custname, custcity, custdiscnt)`
`{`
`exec sql insert into customers`
`values (:custid, :custname, :custcity, :custdiscnt);`
`if (strcmp(custcity, "Dallas") == 0 || strcmp(custcity, "Fort Worth") == 0)`
`exec sql insert into secusts`
`values (:custid, :custname, :custcity, :custdiscnt);`
`} -- alternatively could retrieve count(*) from secusts where cid = :custid, and delete if == 1`
`custdel(custid)`
`{`
`exec sql select city into :custcity from customers where cid = :custid;`
`exec sql delete from customers where cid = :custid;`
`if (strcmp(custcity, "Dallas") == 0 || strcmp(custcity, "Fort Worth") == 0)`
`exec sql delete from customers where cid = :custid;`
`}`

Scaling of Numerical to Letter Grades, CS634 Exam 1, S04

F	D	C-	C	C+	B-	B	B+/A-	A	
			X		X	X	X X	X	X

^	^	^	^	^	^	^	^	^	^
30	40	50	60	70	80	90	100		