

Midterm Exam

CSE232A, Winter 2002

February 21, 2002

Name:

Brief Directions:

- Write clearly: First, you don't want me to spend the whole week grading, do you? Second, it's good for you to write clearly!
- Open books, notes, even databases...
- Good luck!

1 B+ Trees (20 points)

Consider a B+ tree where $n = 4$, i.e., the maximum number of keys in a node is 4 and the maximum number of pointers is 5 at internal nodes and 4 at leaf nodes. Assume that the B+ tree initially consists of a single node, which is both the root and the only leaf, that has the key 1.

1. **2 points** What is the minimum number of keys that may appear in a non-root internal node?
2. **2 points** What is the minimum number of keys that may appear in a non-root leaf node?
3. **8 points**¹ Consider the set of keys $S = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13\}$. Write down a sequence of inserting the keys of S such that at the end the resulting B+ tree has 3 levels and is as empty as possible, i.e., as many nodes as possible have the minimum number of nodes. Provide the B+ tree snapshots that correspond to the points right after node splits. (Use the white space at the end of the page and the back side of this page.)
4. **8 points**² Consider the set of keys $S' = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20\}$. Write down a sequence of inserting the keys of S' such that at the end the resulting B+ tree is as full as possible, i.e., as many nodes as possible have the maximum number of nodes. Provide the B+ tree snapshots that correspond to the points right after a node split. (Use the next blank page.)

Remarks

- Be consistent in the way you split nodes.
- Assume that there are no duplicate nodes in the internal leaves.

2 Extensible Hash Index (10 points)

Assume that each bucket of an extensible hash index can fit exactly two records (each record is a pair of hash key and pointer). Consider the following records, with the corresponding hash key values.

<i>Key</i>	<i>hash key</i>
<i>a</i>	0000
<i>b</i>	0001
<i>c</i>	0010
<i>d</i>	0011
<i>e</i>	0100
<i>f</i>	0110
<i>g</i>	1000

We insert the records in the order given above. Show the extensible hash index after all records have been inserted.

¹May be time consuming.

²May be time consuming

3 Algebra (10 points)

Consider the following two definitions of the semijoin operator \bowtie , which may or may not be equivalent.

- **Direct** (from page 255 of the textbook) The semijoin \bowtie of relations R and S , written $R \bowtie S$, is the bag of tuples t in R such that there is at least one tuple in S that agrees with t in all attributes that R and S have in common.
- **Indirect** Let us call $a(R)$ the list of attributes of R . Then, it is

$$R \bowtie S = \pi_{a(R)}(R \bowtie S)$$

where \bowtie stands for the natural join.

1. **5 points** Are the two definition equivalent? Assume bag semantics for the algebra. If the answer is yes provide proof, showing that, given arbitrary R and S , if a tuple t appears k times in the result of $R \bowtie S$ according to the direct definition, then the tuple t will appear k times in the result of $R \bowtie S$ according to the indirect definition. If the answer is no, provide an example with an R table and an S table and show $R \bowtie S$ for the direct and the indirect definition.
2. **5 points** Consider the indirect definition of \bowtie . Assume that the schema of P is $P(A, B, C, D)$ and the schema of T is $T(C, E)$. Prove the following, using the notation shown in the Appendix.

$$\pi_{A,B}\sigma_{D=5 \wedge E=6}(P \bowtie T) = \pi_{A,B}[(\pi_{A,B,C}\sigma_{D=5}P) \bowtie (\pi_C\sigma_{E=6}T)]$$

4 Query Processing (32 points)

Consider the relations $Cust(CID, Name, City)$, $Order(OID, CID, Date)$, and $LineItem(LID, OID, Product, Amount)$, where CID is a customer id and is a key for $Cust$, OID is an order id and is a key for $Order$, and LID is a line item id and is a key for $LineItem$. In addition the attribute CID of $Order$ is a foreign key referring to the CID of $Cust$, that is, for each CID c of $Order$ there is exactly one tuple of $Cust$ whose CID attribute is c . The OID of $LineItem$ is a foreign key referring to the OID of $Order$.

Assume the following statistics, where all numbers, except for *product*, correspond to millions.

$T(Cust) = 1$	$V(Cust, CID) = 1$
	$V(Cust, Name) = 0.5$
$T(Order) = 20$	$V(Order, OID) = 20$
	$V(Order, CID) = 1$
$T(LineItem) = 100$	$V(LineItem, LID) = 100$
	$V(LineItem, OID) = 20$
	$V(LineItem, Product) = 1000$

Consider the following SQL query, which returns the total amount for each product that “Jones” bought.

```

SELECT Product, SUM(Amount) AS Total
FROM Cust, Order, LineItem
WHERE Cust.CID = Order.CID AND Order.OID = LineItem.OID AND Cust.Name = 'Jones'
GROUPBY Product

```

1. **5 points** Write an algebra expression that uses two cartesian products \times , exactly one selection σ and the $SUM_{Product;Amount \rightarrow Total}$ operator and computes the SQL query.
2. **5 points** Show the series of transformations that transform the algebra expression of the previous question into an expression where the cartesian products have been replaced by joins and the selections are pushed as down (early) as possible.

Assume that the equation $\sigma_{R.A=S.A}(R \times S) = R \bowtie_{R.A=S.A} S$ is one of the transformation rules.

3. **5 points** Provide an additional (non-trivial) expression where the selections have been pushed as down (early) as possible, but the join order is different. No need to show the series of transformations that led you to this expression.
4. **10 points** So far you must have provided 2 algebra expressions with different join orders. For each one of them provide an estimate of the size of its intermediate results. Also, provide an estimate of the size of the final result. To save time, just put the size numbers next to the edges in the algebra expressions.
5. **7 points** Assume that you have got indices on all attributes. Consider the following simplification of the query:

```

SELECT Product, Amount
FROM Cust, Order, LineItem
WHERE Cust.CID = Order.CID AND Order.OID = LineItem.OID AND Cust.Name = 'Jones'

```

Apply the INGRES algorithm to get a join order. Indicate which is the small relation hyper-edge that you pick in each step of the algorithm and show the resulting plan.

A What will get you full points in proving equivalence of algebraic expressions

Assume that the exercise is to prove that

$$\sigma_{p \wedge q}(P \bowtie T) = (\sigma_p P) \bowtie (\sigma_q T)$$

where p refers to attributes of P only and q refers to attributes of T only.

The cleanest proofs are the ones where each step corresponds to application of one of the rules in the notes. In our example, it goes as follows:

$$\begin{aligned}\sigma_{p \wedge q}(P \bowtie T) &= \\ \sigma_p \sigma_q(P \bowtie T) &= \\ \sigma_p(P \bowtie \sigma_q T) &= \\ (\sigma_p P) \bowtie (\sigma_q T) &= \end{aligned}$$