

**University of Helsinki/ Department of Computer Science
Database Management, Course exam 16.12.2004 / H. Laine**

Write the name of the course, date of exam, your name, your date of birth and your signature on each separate answer paper.

Use a separate answer paper for each task.

1. a) Explain briefly the main differences in B+ -tree and ISAM structures. (6p)
- b) Present the principle of the B+ -tree insert algorithm. (6p)
- c) A file is implemented as a B+ -tree. How many records at least and at most may be stored in a B+ tree of height 3 when leaf nodes may contain at most 20 records and index nodes may contain at most 200 links. (6p)

2. Let's consider the tables

```
Product(productnumber, productname, weight, price,..)
PurchaseOrder(Ordernumber, client->customer, date_of_order,
              date_of_delivery,...)
Customer(customernumber, name,..., something),
Orderrow(ordeNum->Purchaseorder, rownumber, product->product, amount)
```

Database uses 4 KB page size. The sizes of the table are the following:

Table	Rows	Rows on page in average
Product	4000	10
PurchaseOrder	40000	10
Customer	2000	5
Orderrow	160000	100

There are about 1000 orders that have not been delivered, yet. They have null value in column date_of_delivery. Orders are distributed evenly among customers and products. All tables are stored as piles (unordered sequential). They have secondary indexes on primary keys and on foreign keys (foreign keys are indicated with '->' symbol). You may assume that each index page may contain at most about 400 entries.

Consider the query

```
select customernumber,ordernumber,productnumber,productname, price
from purchaseorder, customer, orderrow, product
where customernumber=12345 and
      date_of_delivery is null and
      purchaseorder.client=customer.customernumber and
      purchaseorder.ordernumber=orderrow.orderNum and
      orderrow.product=product.productnumber
order by customernumber, product.productnumber;
```

- a) Estimate the number of rows in the result of the query. (6p)
 - b) Define the most efficient way of executing the query. How many disk accesses are needed? Motivate your answer. You may not use more than 50 pages of main memory. (12p)
3. a) How do read and write locks differ from each other? (4p)
 - b) Explain the principle of strict 2 phase locking, 2PL. (5p)
 - c) Explain the concept of deadlock. How can deadlocks be prevented or opened? (5p)