

Tentamen Database Techniek

24 Maart 2009

Closed book exam. Motivate your answers.

Performance Benchmarking

1. 1.a (5pt) What is the most important requirements for any DBMS benchmark suite?
2. 1.b (5pt) What are the (6) most important performance characteristics?
3. 1.c (5pt) How can you deal with the widely different query response times and still give a balanced answer with a single number?

Storage, Indexing and Hashing

4. 2.a (5pt) Construct a B+tree with each node room for 4 pointers. Starting with an empty tree, what is the result after insertion of the following sequence: (2,3,5,7,11,17,19,23,29,31)
5. 2.b (5 pt) In how many ways, and how, can the B+-tree be used to solve a range query? (e.g. `select * from R where a > 12 and a < 40`)
6. 2.c (5pt) What are the differences between a KD-tree and a B+-tree?
7. 2.d (10pt) Give a short description of the Cracking approach to database indexing.

Distributed Database design

8. 3.a (5pt) What is the correctness criteria for horizontal fragmentation?
9. 3.b (5pt) Which query optimizer heuristic rewrite rule is specifically helpful in horizontally fragmented databases?
10. 3.c (5pt) Which query patterns benefit from hash-partitioned data distribution?
11. 3.d (5 pt) How can you handle data skew in a partitioned database setup?

Relational Operations

12. 4.a (10pt) Give a short description of the semijoin algorithm.

- 13. 4.b (10pt) There are at least three fundamental different algorithms to implement a natural join. Give a short description of each.

Query Execution

- 14. 5 (10pt) In what way is the MonetDB execution engine different from all other database engines. What is the weakest point in this design?

Wisconsin design

- 15. (10 pt) What were the three most important decisions for your Wisconsin implementation? How did they affect the performance of relational select and join operators?