

Distributed Data Management

Please make sure you can answer the following correctly:

1. Which sharing architectures are usually used in DDBMS? What does distinguish them?
2. Define partitioning, fragmentation, and allocation.
3. When partitioning, what does it mean if the Completeness Rule is violated?
4. Which operation is used when performing Horizontal fragmentation?
5. How might a distributed database designed for a local area network (LAN) differ from one designed for a wide area network (WAN)?
6. When deciding for a good fragmentation and allocation, which criteria should be optimized?
7. What are the advantages of Horizontal Fragmentation?
8. Is it a good idea to use full replication? If yes, in which cases?
9. Consider the following relation that is fragmented horizontally by city

employee (name, address, salary, city)

Assume that each fragment has two replicas: one is in Berlin and the other one is stored locally at the corresponding city. Describe a processing strategy for the following queries entered at Braunschweig:

- a) Find all the employees names at Braunschweig.
- b) Find the average salary of all employees.
- c) Find the lowest paid employee at each of the following cities: Stuttgart, Manheim, and Frankfurt.
- d) Find the highest paid employee at the company.