

Exercise 11

Released: July 3, 2018 · Discussion: July 9, 2018

1 Fragmentation

1. Describe the principle of horizontal and vertical fragmentation. Discuss for which kind of queries which fragmentation technique might be beneficial?
2. The lecture discussed three correctness rules for fragmentations. Name them and discuss why they are beneficial for a distributed DBMS.
3. How can you calculate a horizontal or vertical fragmentation? How can the original relation be reconstructed from the partitions?

4. Given the following relations *Projects* and *Employees*:

PID	Title	Office	Budget	eid	Name	Proj.	Salary	Department
1	Aquarius	London	16000	628	J. Smith	1	58000	Research
2	Eridanus	Paris	21000	262	D. Miller	4	184000	Research
3	Centaurus	Paris	17000	381	P. Hanks	1	52000	Marketing
4	Andromeda	Rome	29000	725	D. Clark	3	55000	Development
5	Pegasus	London	23000	395	P. Jones	4	143000	Development
				738	S. Miles	2	38000	Sales

5. Divide *Projects* horizontally on the attribute *Office* and derive a horizontal fragmentation for *Employees*.
6. Now divide *Employees* horizontally on the attribute *Department* and derive the fragmentation for *Projects*
7. Discuss how the correctness criteria are affected by the choice of the fragmentation scheme in the previous two assignments.

2 Allocation

1. What is an allocation? Explain the connection between allocation and fragmentation.
2. For which criteria can a allocation be optimized?
3. Discuss the advantages and disadvantages of replication!
4. Explain the heuristics for calculating an allocation that were presented in the lecture.

3 Distributed Transactions

1. Describe the function of the Two-Phase Commit Protocol.
2. How can serializability be realized in a distributed database?