

```
1 SELECT Umsatz
2 FROM Verkaeufe, Ort, Zeit, Produkt
3 WHERE Produkt.id= Verkaeufe.Produkt_id AND
4       Produkt.Produktgruppe='VW' AND
5       Ort.id= Verkaeufe.shop_id AND
6       Ort.Region='Magdeburg' AND
7       Zeit.id= Verkaeufe.day_id AND
8       Zeit.Jahr='2004' OR
9       Zeit.Jahr='2005' OR
10      Zeit.Jahr='2006' AND
11      Zeit.Monat= '12';
```

Figure 1: DWH-Query.

1. What is a join index?
2. Classify the index structures of the last exercise!
3. Classify and define aggregate functions.
4. What is an Iceberg-Cube?
How can cubes in general be calculated efficiently?
5. Discuss the approach of common DBMS query optimizers. Are the approaches sensible for data warehouses?
 - (a) A data warehouse schema is given, that associates one fact (Verkaeufe) to 3 Dimensions (Zeit,Ort,Produkt). Furthermore, the following meta-data are given:
 - 50.000.000 tupels are inside the fact table.
 - The time dimension contains 10 years (20 days per month).
 - There are 50 product groups each having 20 products.
 - There 50 locations with 100 car-salers each.

The sales are distributed uniformly for all dimensions. Which execution plans are proposed by a common database optimizer regarding the query of Figure 1? Which optimal execution plan is not proposed by standard dbms optimizers?