



## Exercises Data Warehousing View Materialization

- An important technique to speed up analytical queries is by pre-computing and materializing aggregations. Consider the following lattice of views that can be requested by the user, along with the number of rows in each view.



A is the view representing the base relation. The edges indicate the relation “can be computed from.”

- Suppose that only the top-view A has been materialized. Select three additional views from the views B, C, D, E, F, G, H, I, and J to materialize. Apply the greedy method described by *Harinarayan, Rajaraman, and Ullman* in their seminal paper “Implementing Data Cubes Efficiently” (SIGMOD 1996)
  - What benefit gives the additional materialization of these three views under the cost model introduced by these authors?
- Consider a data cube with dimensional attributes **Product** and **Date** and measure **Total Sales**. The following hierarchy is present on the dimension **Date** (No hierarchy on **Product**):



- Suppose that the cube is dense; i.e., for every product  $p$  and date  $d$  there is a tuple  $(p, d, t)$  in the base relation with  $t > 0$ . Furthermore, there are 100 products and 3 years of data (1095 days; no leap year, 157 weeks) in the database. Determine the sizes of the different views.
  - Apply the greedy algorithm to select 2 views to materialize.
  - What is the total benefit of materializing those two views?
- Imagine a National Bureau of Statistics holding demographic data of the inhabitants of its country. For each person the following attributes are recorded: *City*, *Age*, *Ethnicity*, and *Income*. The bureau wants to make aggregations of this data available for other governmental agencies as well. The agencies can query the data by selecting any subset of the first three attributes. For any combination of the selected attributes the average income of people with this particular combination will be returned. For example, if an agency selects *City* and *Age*, it will get for every city and age, the average income of the people of that age in that city.  
For reasons of simplicity we assume that for every city, age, and ethnicity there is at least one person that has this combination. There are 400 cities, 100 age values, and 10 ethnicity types in the database. It is decided that the data will be stored in a database system, but for performance reasons some query answers will be materialized in advance. Every time new data arrives these materialized views are updated.
    - Suppose 2 views can be materialized. Which views should be materialized? Explain the reasoning behind your answer.

- (b) After a few months the bureau realizes that the system's performance is suboptimal. As a first step towards solving this problem, an analysis is made of what is the frequency of the different queries. It turns out that not all 8 combinations of the attributes *City*, *Age*, and *Ethnicity* are equally likely to be requested. The following table summarizes the results of the analysis (every row represents how often one selection of attributes is being requested. The symbol  $\checkmark$  in the column for attribute *A* in a row indicates that *A* is in the selection represented by that row):

<i>City</i>	<i>Age</i>	<i>Ethnicity</i>	<i>Frequency</i>
$\checkmark$	$\checkmark$	$\checkmark$	5%
$\checkmark$	$\checkmark$		20%
$\checkmark$		$\checkmark$	30%
$\checkmark$			30%
	$\checkmark$	$\checkmark$	5%
	$\checkmark$		5%
		$\checkmark$	5%
			0%

Hence, for example, it turns out that the average income over all people (no attributes selected) is never requested and that, e.g., the average income per city and the average income per city and ethnicity type are disproportionately frequent (together they form 60% of the requests!).

Should the views to materialize be changed based on this information? If so, which views should be materialized instead? Explain your answer.

- (c) What additional benefit can be expected if it is decided to materialize also a third view? Use the settings described in (b).