Today's class: Data Cubes; §§ 10.{6,7}.
Next class: Data Cubes paper, XSLT; § 12.3.

1. List the members of your group below. Underline your name.

2. Using illustrative examples of queries, explain the key difference between OLTP and OLAP.

3. Example 10.28 in the textbook suggests a pseudo-table Days the dimension table for the date attribute of the fact table Sales. What foreign-key constraint connects the two? How is the pseudo-table realized in PostgreSQL?

4. Is the second query of Example 10.30 in the textbook is likely to produced the desired result? Explain.

5. For the above example, suppose we wish to partly automate the task of detecting poor sales by defining them to be any aggregated sales figures that less than 20% of the average comparable sales figures. Provide a SQL query that finds poor sales.

6. If F is a fact table, what are lower and upper bounds on the ratio $|\text{CUBE}(F)|/|F|$? Explain.

7. Define a view `SalesA` that presents the data from `Sales` aggregated in the manner suggested by the first two paragraphs of Example 10.31.

8. Is it possible to create a view that is equivalent to the view `SalesCube` of Example 10.32 without using any data-cube features of SQL (such as `with cube`)?