

Today's class: Data Cubes, early paper<sup>1</sup>; §§ 10.{6,7}.

Next class: synthesis and review, XSLT; § 12.3.

1. List the members of your group below. Underline your name.
2. If  $F$  is a fact table, what are lower and upper bounds on the ratio  $|\text{CUBE}(F)|/|F|$ ? Explain.
3. Define a view `SalesA` that presents the data from `Sales` aggregated in the manner suggested by the first two paragraphs of Example 10.31.
4. 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`)?

---

<sup>1</sup>Jim Gray et al., "Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab, and Sub-Totals," *Data Mining and Knowledge Discovery* 1 (1997).

5. The abstract refers to SQL aggregation queries producing zero- or one-dimensional aggregates. Is there a simple test to determine which? Explain.
6. Explain the comment on “creating  $2^N$  aggregation columns” (bottom of page 34) in the context of the example of Table 3. Generalize.
7. Explain how to produce a spreadsheet table analogous to Table 4 using *LibreOffice Calc*.
8. Depict a likely mapping of the query of page 36 to logical and physical plans. Later, compare your work with the plans generated by PostgreSQL.
9. Provide a precise description of the query mentioned in the penultimate paragraph of Section 2 (page 38).
10. Devise and perform experiments to evaluate the claims made in the last paragraph of Section 2 on a current SQL implementation.