

This exercise complements initial classroom discussion of data cubes based on the paper¹ introducing them, and the textbook.²

1. List the members of your group below. Underline your name.
2. The abstract refers to SQL aggregation queries producing zero- or one-dimensional aggregates. Is there a simple test to determine which? Explain.
3. Describe how the enhanced aggregation functions on page 33 of the paper may be expressed in current standard SQL. [Hint: Recall the examples from the AQuery paper.³]
4. Explain the comment on “creating 2^N aggregation columns” (bottom of page 34) in the context of the example of Table 3. Generalize.

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

²Jeffrey D. Ullman and Jennifer Widom, *A first course in database systems*, third edition (Upper Saddle River, New Jersey: Prentice-Hall, 2008), §§ 10.6–10.7.

³Alberto Lerner and Dennis Shasha, “AQuery: Query Language for Ordered Data, Optimization Techniques, and Experiments,” in *Proceedings of the 29th International Conference on Very Large Data Bases (VLDB)* (Berlin, Germany, 2003).

5. Explain how to produce a spreadsheet table analogous to Table 4 using *OpenOffice Calc*.
6. 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.
7. Provide a precise description of the query mentioned in the penultimate paragraph of Section 2 (page 38).
8. Devise and perform experiments to evaluate the claims made in the last paragraph of Section 2 on a current SQL implementation.