

These questions complement continuing class discussion of the paper describing *AQuery*<sup>1</sup> and Graefe's survey.<sup>2</sup> The focus of this exercise is studying the query plans generated by the PostgreSQL optimizer with and without various indexes for the simple microbenchmark database instance discussed earlier.

1. List the members of your group below. Underline your name.
  
2. Recall the *best-profit* and *packet-grouping* queries (Section 1.1) discussed earlier. Provide two different standard SQL expressions of those queries.

---

<sup>1</sup>Alberto 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).

<sup>2</sup>Goetz Graefe, "Query evaluation techniques for large databases," *ACM Computing Surveys* 25/2 (1993).

3. Based on your experimentation with the queries of Question 2 in PostgreSQL and other systems, depict the generated query plans in the absence of any indexes. Explain the physical operators chosen and comment on the reasons for their selection. Note the estimated and actual run times based on your study.

4. Repeat Question 3 in the presence of the most profitable indexes you were able to determine.

5. Compare the results of Questions 3 and 4 with those in the AQuery paper.