

1. Provide a concrete scenario in which linear hashing performs significantly better than extensible hashing. Provide another in which linear hashing performs significantly worse than extensible hashing. The scenario must include a log of events (insertions, lookups, rearrangements, etc.) that occur over time, along with complete cost calculations. You may focus on disk operations, but try to account for CPU and main memory costs as well.
2. Generalize the ideas in your answer to Question 1 to characterize the class of scenarios in which one method outperforms the other. Your characterization should make it easy (or at least significantly easier) for a programmer to decide which method to use when she is familiar with the application.
3. For a query of your choice, exhibit two real query plans that differ in efficiency by at least an order of magnitude on the same database instance. By real query plan, we mean one generated by PostgreSQL. Use tables of reasonable size, so that you do not compare, for example, 0.001 second with 0.01 second! Repeat this exercise for Oracle. Hints: Your query can be as complex as you wish (but try to generate the simplest example you can). You are free to create indexes that are available to only one plan.

For this and the next two questions, in addition to your written answer, include in your submission traces of *psql* and *sqlplus* sessions that illustrate your examples in action. Please name your files in an easily decipherable manner and include a README file that explains them briefly.

4. Repeat the exercise of Question 3 with the additional restriction that both plans must have the same set of indexes available to them. However, it is permissible for the plans to operate on different database instances *provided* the instances agree on the cardinality of each relation and on the size of the query result. You must allow the database system to recompute statistics after you modify the instance.
5. Repeat the exercise of Question 3, with the additional restriction that the same set of indexes be available to both plans. However, you are permitted to submit different expressions of your query to the database system (as long as they are logically equivalent).

**Submission** Please submit your written responses electronically in plain text, Latex, or PDF form, along with the trace files as a single package. Follow the general submission procedure described in HW1, replacing *hw1* with *hw2* in the filename.