

DBMS Models and implementation
Instructor: Sharma Chakravarthy
SQL Query Plan Analysis on Oracle

Made available on: 11/9/2025
Submit by: 12/8/2025 (11:59 PM) No demo for this project (no extension)
Submit to: Canvas (1 zipped folder containing all the files/sub-folders)
Weight: 7% of total
Total Points: 100

Problem specification: Use the IMDb database on omega whose schema details were given for project Ia. You will use the same databases on Omega: imdb00 and imdb01

Caution: This is NOT a toy database as it has Millions of rows in each table and hence your queries can produce tens of thousands of rows as output! Hence, you need to think before posing a query.

What to do for this project:

- I. Choose 2 queries that you implemented for project Ia.
- II. For each of the 2 queries and the choices of indexes you have used, generate the query plans using the EXPLAIN statement of Oracle, and compare and discuss these plans with respect to the indexes chosen as well as the one (imdb00) where there are no indexes. You will do this as was done in project Ia by running the query on imdb00 as is and with index choices on imdb01 (by using the same hints). Understand
 - a. How/whether the query plans changed based on index choice
 - b. The costs associated with the plan and how they have changed due to change in query plans
- III. Refer to [EXPLAIN PLAN](https://docs.oracle.com/cd/B19306_01/server.102/b14211/ex_plan.htm#i16971) for using the explain statement of Oracle. Please refer to https://docs.oracle.com/cd/B19306_01/server.102/b14211/ex_plan.htm#i16971 for understanding and interpreting the plans generated. Explain statement can generate output into a table which can be printed and included/analyzed.

This project is meant to give an idea of how query plans change based on availability of indexes. This is done automatically by the query optimizer when indexes are created by the DBA. Here you are forcing index choices to see how the query optimizer uses them. **Pay attention to change in join algorithms and how indexes are used for the condition evaluation in the where clause.**

1. **Project Report:** Please include (at least) the following sections in a **REPORT.{doc, docx}** file that you will turn in with your code:
 - i. **Overall Status**

Two Queries you have chosen and their index choices. Query plans for the base query and all index variants used.
 - ii. **Analysis Results:**
 - All English queries, SQL equivalents that use Hint and explain clauses should be in one output file (named Fall-2025_5331_<Team_no>.sql)
 - An output file (named Fall-2025_5331_<Team_no>.txt) that has the above and NEATLY FORMATTED output of each query along with the query plan. You may have to print the relation in which query plan is stored.
 - For both of the above, you should use the comment capability (--) in SQL to write English queries and use set echo on and spool <filename> and spool off for getting everything in one file. I will briefly go over this in the class.
 - Compare the query plans obtained to analyze whether they differ and how they differ. Explain how they compute the same result even though the plans are different.
 - This is best done as a table where rows are alternative index choices and columns are queries. add an analysis column to discuss query variants for each query
 - iii. Your analysis file can be separate from the above file. Name it as Fall-2025-5331-<team-no>-ana.docx
 - iv. **File Descriptions**

List the files you have created and *briefly* explain their major functions and/or data structures.
 - v. **Division of Labor**

Describe how you divided the work, i.e. which group member did what. Please also include how much time each of you spent on this project. (This has no impact on your grade whatsoever; we will only use this as feedback in planning future projects -- so be honest!)
 - vi. **Logical errors and how you handled them:**

List at least 3 logical errors you encountered during writing of SQL query. Pick those that challenged you. This will provide us some insights into how we can improve the description and forewarn students for future assignments.

2. What to submit:

- After you are satisfied that your query does exactly what the project requires, you may turn it in for grading. Please submit your project report with your project.
- You will turn in one zipped file containing your
 - a. **Source code (English queries, SQL queries, their output)**
 - b. **Report that includes analysis results for the chosen queries: cost computations using oracle units, comparison spreadsheets etc.**
- All of the above files should be placed in a single zipped folder named as - 'Fall-2025-project-1b_team_<TEAM_NO>'. **Only one zipped folder should be uploaded using canvas.**
- You can submit your zip file at most 3 times. The latest one (based on timestamp) will be used for grading. So, be careful in what you turn in and when!
- **Only one person per group should turn in the zip file!**

3. Grading rubric:

Project 1b will be graded using the following rubric:

a. Correctness of the SQL Queries and Output (from project 1a)	10
b. Query Plan generation using Hint and EXPLAIN statement for each variant	40
c. Explanation of query plans and their analysis	30
Overall Completeness of Report <i>Status, File Descriptions, Division of Labor, Logical Errors</i>	10
TOTAL	100