

BMI 544: Databases

Winter term, 2008

A continued intensive introduction to computer science, focusing on databases and database management systems. The main topics covered include database principles, relational databases, database design with Entity-Relationship modeling, the SQL query language, database optimization, and data warehousing. Individual database projects will be completed.

Instructor: Judith R. Logan, M.D., M.S.
Office: 503-494-5902 / FAX 503-494-4551
Cell: 503-789-5933
loganju@ohsu.edu

Required Textbook:

Database Design, Application Development, & Administration, 3rd edition
Michael V. Mannino
McGraw-Hill, 2007

* Note: No CD is required.

Grading:

35% homework assignments: there will be weekly reading assignments and homework assignments, similar to the schedule for BMI 540. Late assignments will lose 10% or more depending on number of days late unless prior permission obtained.

15% quizzes

20% final examination

30% project

Computing requirements:

Students will be required to use the Oracle SQL*Plus utility program. In order to do this, you will receive an account on the davinci server at OHSU. This will require you to have an sFTP program (such as Putty, available at <http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>) as well as a secure file transfer program (such as WinSCP, available at <http://sourceforge.net/projects/winscp/>). You will receive instructions for setting up your davinci account the first week of classes.

If you wish, you may also download a copy of Oracle so that you may work on your local computer. Oracle 10g Express Edition can be obtained without a licensing fee at <http://www.oracle.com/technology/software/index.html>. Note that this is resource-intensive. You may also use Microsoft Access, but be warned that there are differences in the SQL for Access and Oracle and you are responsible for knowing the Oracle version of SQL.

In addition, each student will need an application to create E-R diagrams. Your textbook Online Learning Center can be found at <http://www.mhhe.com/mannino>. You can download a program

called ER Assistant from that site which will support the diagramming conventions we will use. Other programs such as SmartDraw and Visio can also be used.

Goals:

- Understand the basic principles and purposes of databases.
- Have cursory knowledge of the historical database models and how they differ from the relational database model.
- Understand the structure and major tenets of relational databases.
- Be able to model data using Entity-Relationship diagrams and then convert those models into a relational database.
- Be able to optimize a database design through normalization to Boyce-Codd Normal Form
- Be able to write data definition, manipulation and query statements in SQL.
- Execute data definition, manipulation and query statements in SQL in an Oracle database using the Oracle utility SQL*Plus.
- Execute complex SQL queries using Oracle-specific functions and clauses.
- Understand principles of decision support using databases: data warehousing, OLAP, data mining techniques
- Understand the principles of transaction control and concurrency.
- Understand the basis for data definition, manipulation and query languages in relational algebra and its use in query optimization.
- Understand issues in physical database design and database tuning.

Professional Conduct Policy:

The professional conduct policy of the Graduate Studies Program can be found at <http://davinci.ohsu.edu/~loganju/BMI544/profconductpolicy.pdf>. While discussion between students and use of external resources are important learning tools, all homework assignments, quizzes, projects, and exams are expected to be the work of the enrolled student only. Any violation will result in zero points for that homework assignment, quiz, project or exam as well as referral to the professional conduct committee of DMICE.

Syllabus:

Week	Topics	Assignments
1: Jan 9 - 15	Data, databases and database management systems The relational model	Mannino, Chapters 1, 2, and 3 (3.1 – 3.2)
2: Jan 16 - 22	Relational algebra operators SQL*Plus SQL CREATE statement	Mannino, Chapters 3 (3.3-3.4) Assignment #1 due 1/16 Project, step 1 due 1/18
3: Jan 23 – 29	SQL: SELECT statements Use cases	Mannino, Chapter 4 Assignment #2 due 1/23
4: Jan 30 – Feb 5	Conceptual data modeling: Entity-Relationship diagrams Optional programming workshop:	Mannino, Chapter 5 Assignment #3 due 1/30 Project, step 2 due 2/1

	Database connectivity with java Quiz #1	
5: Feb 6 - 12	E-R diagrams (continued) Converting an ER Diagram to relational tables	Mannino, Chapter 6
6: Feb 13 – 19	Database optimization: Normalization of relational tables	Mannino, Chapter 7 Assignment #4 due 2/13 Project, step 3 due 2/15
7: Feb 20 – 26	Physical database design Optional programming workshop: Embedded SQL Quiz #2	Mannino, Chapter 8 Assignment #5 due 2/20
8: Feb 27 – Mar 4	Advanced SQL Views	Mannino, Chapter 9 and 10 Project, part 4 due 2/29
9: Mar 5 – 11	Database administration Transaction management	Mannino, Chapter 14, 15 Assignment #6 due 3/5
10: Mar 12 – 18	Data warehousing Optional programming workshop: Stored procedures and triggers Quiz #3	Mannino, Chapter 16 Assignment #7 due 3/12 Complete project report due 3/16
11: Exam week	The final exam will be posted on Monday, March 17 and must be completed by Saturday, March 22	

BMI 544 Term Project, Winter 2008:

The project for this term is for you to design a database-dependent application and implement the relational database required for that application. The application might be for a business that you are familiar with or for a household need, perhaps one of your hobbies. The data requirements must be simple enough to manage in a small project but still be challenging to you. **You cannot, however, use an Electronic Health Record as your project** – they are both too complex and too familiar to serve as good learning experiences.

There are 5 steps for the project and a final evaluation. You will be asked to submit the first four steps on set dates either to the course instructor or TA, to make sure that project requirements are met and to help with any problems. The final report consists of all documents from the five steps, but compiled in a professional manner. Note that at any step you can improve on your previous work...you will receive a grade for the complete project only.

1. Describe the application that you would like to design. This should include the purpose of the application and an informal description of the application functionality. List and discuss in detail

the data elements that must be captured in this application and stored in the database. The application may meet all of the data needs of a business or may be specific to a single need within that business. Turn in a text document with this description.

- a. Note that you should not describe the TABLES that will contain this data... table design comes later. For now, just make sure that you know what data points you want to capture.

2. Write three to five use cases that help us understand the interaction of the application with the database. Each use case should identify the actor(s) involved along with a step-by-step description of what the actor will accomplish when the use case follows the normal course of events and is successful. You may assume that if the user is a person, that person has authenticated to the application (i.e. don't include login scenarios). Turn in a text document which contains the project description and these use cases.

3. Draw one or more Entity-Relationship diagrams for your database design. Hopefully you will have 7-10 entities. Include a write-up discussing the options that you have for this design and why you chose the option that you did. The E-R diagrams must conform to the conventions that we have chosen to use in this class. The diagrams must be added to your document; with ER Assistant, you can cut and paste your diagrams into Word. Only one document should be turned in.

4. Create the relational schema for your database based on your E-R diagram. Then revise this schema to be normalized or, if you choose not to, explain why denormalized tables are preferable. Implement this in Oracle and populate the database with test data. Turn in a text file containing the relational schema and the data definition and manipulation statements that you used to create the tables and add data. .

- a. Beware of using common names for your tables since there will be multiple projects in this one database.
- b. Do not show the statements for all data entered, just a sample for each table.
- c. Please do NOT turn in screen captures, just the final statements

5. Write the SQL queries that support at least three of your use cases. Perform a screen capture showing the results when you run these queries against your database. Please edit the screen capture to delete all unnecessary text so that this will be legible to your readers.

Finally, add a conclusion that briefly sums up the project and any ideas you have for implementing it.