

INLS 523-004 Database Systems I

Schedule: Monday and Wednesday, 10:10 - 11:25 in Bingham 101

Instructor: Eric Chernoff, eric_chernoff@unc.edu, 919-744-8383

Course Website: https://sakai.unc.edu/portal/site/very_exciting_db_systems_class

Office hours: By appointment; usually available right after class. Will respond to email and messages within 24 hours.

Course Overview

The goal for this course is to be a pleasant introduction to the wonderful world of database systems. By the end of the semester, we will have answered the following questions:

1. What are databases, when are they needed, and how they are used?
2. What is a relational database and SQL?
3. What is a document database and NoSQL?
4. How do you use SQL to retrieve information from a relational database? How do you use SQL to create and update this data?
5. What are data types and data normalization and why are they so important? Is it possible for an adjunct professor to make this topic interesting?
6. How do you write really fast SQL queries, and why wasn't that covered in question 4?
7. How do programmers work with databases? How do they enrage database admins?
8. What are some key points on database security and high availability?

We will be using MySQL to explore these topics. You will learn a lot about MySQL and MariaDB, but the goal is to learn the concepts behind those implementations.

Mandatory Required Necessities

Before participating in this exciting course, all must students must have:

- Completed the required prerequisite INLS 161.
- A proper laptop with a SSH client.
- A copy of Murach's MySQL 2nd Edition (Joel Murach, ISBN 978-1-890774-82-0)
- Access to Sakai to complete the quizzes and exams.

Grading / Expectations

While I am friendly and often entertaining, this is a very important course for your education and career. We will have fun, but everyone is expected to attend class, be on-time, participate, leave smart devices locked, and treat one other respectfully.

Each class meeting will include a lecture, demonstrations, hopefully an exercise, and time for questions and discussion. There will be readings assigned for most class meetings, and we will randomly have 10 quizzes on the readings.

The grading scale will be:

A = 90-100	B = 80 - 89.9	C = 70 - 79.9	D = 60 - 69.9	F = 0 - 59.9
------------	---------------	---------------	---------------	--------------

Your grade will be based on:

Reading Quizzes	25%	10 quizzes for selected readings. Your bottom two scores will be discarded (including any non-submissions)
Weekly assignments	35%	Submitted via Sakai. 10% daily penalty if late.
Mid-term exam	20%	Multiple-choice
Final exam	20%	Multiple-choice

This course is about learning key information and concepts, so grading will be objective.

You are expected to attend this class, and it will be very difficult to pass without attending.

Our conduct in this class will be governed by the UNC Student Code of Conduct and UNC Honor Code. Information on both codes is available here: <https://studentconduct.unc.edu/>

Please note that audio and video recording is not allowed without the instructor's written permission.

Course / Topic Schedule

Readings and Assignments will be due on the date of class.

Date	Class Plan
Wed, 8/22	Introduce this fantastic course, meet our wonderful classmates, and discuss the concept of databases. Brief history of database technology, then discuss how databases are used in the (real) world.
Mon, 8/27	<p>Things that deserve capitalization for emphasis: Relational Database Concepts, Semantic Modeling and Entity Relationships.</p> <p>Reading: Murach's, pages 4-21</p>
Wed, 8/29	<p>The class where we talk about Document Databases and NoSQL</p> <p>Reading: https://ils.unc.edu/courses/2018_fall/inls523_004/nosql.pdf → OK to skip the 'Indexing' bits, but please read the comments under the article. Assignment: Data Collection Project</p>
Mon, 9/3	LABOR DAY, NO CLASS
Wed, 9/5	<p>Review of relational databases then meeting our main attraction: Structured Query Language, also known as SQL.</p> <p>Reading: Murach's, pages 22-33, 66-69</p>
Mon, 9/10	<p>SQL Operators: <code>SELECT info FROM instructor WHERE will_be_on_exam = 'Yes' AND (importance > 10 OR is_boring = 'No');</code></p> <p>Reading: Murach's, pages 74-103</p>
Wed, 9/12	CLASS CANCELED DUE TO FLORENCE
Mon, 9/17	CLASS CANCELED DUE TO FLORENCE
Wed, 9/19	<p>You too can be a SQL guru: Clauses, Order-by, Limit, Grouping (Aggregation)</p> <p>Reading: Murach's, pages 104-109 and 174-183 Assignment: Basic Querying Tasks</p>
Mon, 9/24	<p>Calculating the right time to talk about date/time functions.</p> <p>Reading: Murach's, pages 262-267</p>

Wed, 9/26	<p>Reading is no fun without writing: SQL for creating, updating and deleting data.</p> <p>Reading: Murach's, pages 152 - 163 Assignment: Intermediate Querying Tasks</p>
Mon, 10/1	<p>Schema schemes: Column Data Types and SQL for defining tables (DDL)</p> <p>Reading: Murach's, pages 216-231 and 316-327</p>
Wed, 10/3	<p>Simple tricks for great magicians: Table Joins along with Keys and Indexes.</p> <p>Reading: Murach's, pages 114 - 139 and 328-329 Assignment: Aggregate Querying Tasks</p>
Mon, 10/8	<p>Don't be sad -- your favorite SELECT can live forever! Creating Table Views (virtual tables) Bonus: How to write slow SQL: Unions and sub-selects</p> <p>Reading: Murach's, pages 188 - 207 and 354-369</p>
Wed, 10/10	<p>The Four or Maybe Seven Forms of Database Normalization. Also, why does it need to be fast?</p> <p>Readings: Murach's, pages 278-303 https://blog.udemy.com/normalization-in-database-with-example/ Assignment: Date and Time Calculations</p>
Mon, 10/15	<p>Let's review for the mid-term. Bring your questions.</p>
Wed, 10/17	<p>MID-TERM EXAM</p>
Mon, 10/22	<p>We need to talk about thinking about very large datasets. Plus: The Old Man's Guide to Optimization and normalization in The Real World. Creating faster tables and queries.</p> <p>Reading: https://dev.mysql.com/doc/refman/8.0/en/optimize-overview.html</p>
Wed, 10/24	<p>The boring stuff that saves the world: Data sanity and integrity.</p> <p>Reading: https://www.lifewire.com/common-database-mistakes-4140757</p>
Mon, 10/29	<p>Wisdom from the Oracle: Storage engines, ACID compliance and Database Transactions</p> <p>Reading: Murach's, pages 402 - 413</p>
Wed, 10/31	<p>MySQL Workbench and other GUI cheats, um, tools.</p> <p>Reading: Murach's, pages 42 -63 Assignment: Creating Normalized Tables</p>
Mon, 11/5	<p>SQL is code too: All above Stored Procedures and Triggers</p>

	Reading: Murach's, pages 416 - 439
Wed, 11/7	The glamorous DBA Life: Server tuning, monitoring logs, and backing up databases Reading: Murach's, pages 462 - 473, 528 - 537 Assignment: Creating an Optimized Table View
Mon, 11/12	Vertical scaling isn't just for accordions: Database replication and clustering and high availability
Wed, 11/14	The most important thing we forget to remember: Database security: user privileges and network controls Reading: Murach's, pages 494 - 515 Assignment: Advanced Querying Tasks
Mon, 11/19	Things your boss will ask about: Database security: Encryption and SQL Injection Reading: https://www.owasp.org/index.php/SQL_Injection and https://en.wikipedia.org/wiki/Data_at_rest
Wed, 11/21	THANKSGIVING, NO CLASS
Mon, 11/26	ORM's and other ways developers drive DBAs mad. Readings: https://www.fullstackpython.com/object-relational-mappers-orms.html https://dev.to/bertilmuth/orm-vs-sql-2bbf
Wed, 11/28	When relational databases are the wrong answer: Storing files, logging, task queues, and more true confessions. Reading: https://dzone.com/articles/which-is-better-saving-files-in-database-or-in-file Assignment: Database Design Exercise (25% of Assignment grading)
Mon, 12/3	Additional advanced concepts for DBA's, developers, and sysadmins. Readings: https://blog.timescale.com/why-sql-beating-nosql-what-this-means-for-future-of-data-time-series-database-348b777b847a
Wed, 12/5	Let's review for the final exam. Bring your questions.
Fri, 12/14	8:00AM - FINAL EXAM