

Principles of Database Systems (CS6083)

Course Prerequisites: Good programming skills, familiarity with basics of operating systems (incl. file systems, caching, locking and concurrency), and knowledge of basic algorithms and data structures (incl. sorting, heaps, search trees, and hashing). Formally, you need to have taken a course on data structures or algorithms, and a course on operating systems, at the graduate or undergraduate level, or have permission from the instructor.

Instructor: The course is taught by Prof. Torsten Suel (email torsten.suel@nyu.edu). Phone: 646 997-3354. Office: Room 856 in 370 Jay Street.

Time and Location: Mondays from 3:20pm to 5:50pm in room 9.011 in 2 Metrotech Center.

Office Hours: Mondays 2-3pm in my office, or by appointment. The course TAs will also announce their own office hours.

Course Webpage: On NYU Classes. All announcements will also be sent via the course mailing list to your NYU student email account.

Required Text: A. Silberschatz, H. Korth, S. Sudarshan, *Database System Concepts*, 6th Edition, McGraw-Hill, 2009. (You may also use the 4th or 5th or 7th edition if those are cheaper or easier to get.)

Grading Policy: Problem sets: 20%. Project: 20%. Midterm or tests: 20%. Final exam: 40%.

Final Exam Date: The final exam will be held on Monday, May 18, at 3:20pm. Location to be announced.

General Information: The course aims to give a broad introduction to relational database systems, including the relational data model, query languages, index and file structures, query processing and optimization, concurrency and recovery, transaction management, and database design, plus optional material if time permits. The workload of the course is expected to be above average, and class attendance is strongly recommended.

Policy on Academic Dishonesty: Please see the policy on academic dishonesty on our school's website at <https://engineering.nyu.edu/campus-and-community/student-life/office-student-affairs/policies/student-code-conduct>. Common examples of misconduct include cheating, fabrication, plagiarism, and/or unauthorized collaboration. Students are expected to work on their own, with the possible exception of group projects if allowed by the Professor. Students may verbally discuss work with other individuals either in the class or outside the class, but they may not reuse code, results, or text received or retrieved from any source unless clearly disclosed in their submissions. Any student who is found to be violating this policy will be given a failing grade for the work and will be reported to the authorities, including the CSE Department's student records, as described in the University's Student Code.

Moses Center Statement of Disability: If you are a student with a disability who is requesting accommodations, please contact New York University's Moses Center for Students with Disabilities (CSD) at 212-998-4980 or mosescsd@nyu.edu. You must be registered with CSD to receive accommodations. For more information, see <http://www.nyu.edu/students/communities-and-groups/students-with-disabilities.html>. The Moses Center is located at 726 Broadway on the 2nd and 3rd floors.

Excused Absences: The detailed NYU Tandon policies and procedures on excused absences can be found at <https://engineering.nyu.edu/campus-and-community/student-life/office-student-affairs/policies#chapter-id-30199>. In short, an absence can be excused if you have missed no more than 10 days of school. If an illness or special circumstance has caused you to miss more than two weeks of school, please refer to the section labeled Medical Leave of Absence. Students may request special accommodations for an absence to be excused in the following cases: (1) Medical reasons, (2) death in the immediate family, (3) personal qualified emergencies (documentation must be provided), and (4) religious expression or practice. Deanna Rayment, deanna.rayment@nyu.edu, is the Coordinator of Student Advocacy, Compliance, and Student Affairs, and handles excused absences. She is located in 5 MTC, LC240C, and can assist you should it become necessary.

Programming Assignments: In addition to written problems, there will also be one or two programming assignments as part of the homework, plus a two-part project using a database system of your choice:

1. Programming Assignments: How to use a DBMS, create schemas, pose queries, and program in PHP.
2. Project Part I: Designing a complex schema and queries for a given application scenario.
3. Project Part II: implementing a web-based application, typically with PHP and Oracle, mySQL, or another database system, based on the schema and scenario from Part I.

Course Outline (tentative)

1. Introduction and Basics (Chapters 1 to 2 and Chapter 6) (*weeks 1-2*)
 - (a) Overview
 - (b) Relational Model
 - (c) Relational Algebra
 - (d) Relational Calculus
2. Relational Queries (Chapters 3 to 5) (*weeks 3-4*)
 - (a) Basic SQL
 - (b) Advanced SQL
 - (c) Integrity Constraints
3. Database and Application Design (Chapters 7 to 9) (*weeks 5-6*)
 - (a) Intro to the ER Model
 - (b) Relational Design and Normal Forms
 - (c) Application Development

Midterm Exam: (around mid to late March, covering material up to Chapter 7)

4. Storage and Indexing (Chapters 10 and 11) (*weeks 8-9*)
 - (a) Disk Models
 - (b) External Sorting
 - (c) Disk and File Organization
 - (d) Indexing
5. Query Processing (Chapters 12 and 13) (*week 10-11*)
 - (a) Query Evaluation
 - (b) Query Optimization
6. Transaction Management (Chapters 14 to 16) (*weeks 12-13*)
 - (a) Transactions
 - (b) Concurrency Control
 - (c) Recovery