

Introduction to Database Systems

Syllabus

Web Page

<http://www.eecs.northwestern.edu/~pdinda/db>

Instructor

Peter A. Dinda
Technological Institute L463
pdinda@northwestern.edu
Office hours: Thursdays, 2-5pm, or by appointment

Teaching assistants

Besim Avci
Tech L580
besim@u.northwestern.edu
Office hours: Tuesdays and Wednesdays, 11am-1pm, or by appointment

Brandon Staunton
Wilkinson Lab
brandonstaunton2015@u.northwestern.edu
Office hours: Mondays and Fridays, 10:30am-12:30pm, or by appointment

Location and Time

Lecture: Technological Institute, Tech A110, 4-4:50pm, MWF

Prerequisites

Required	EECS 214 (311) or equivalent data structures course
Required	EECS 213 or equivalent computer systems course
Recommended	Familiarity with concepts from discrete math such as set theory (EECS 212/310 for example)
Recommended	Some familiarity with Perl or other scripting language

Textbook and other readings

Hector Garcia-Molina, Jeffrey D. Ullman, Jennifer D. Widom, *Database Systems: The Complete Book*, 2nd Edition, Prentice Hall, 2009. (Textbook - Required)

- An in-depth introduction to databases and database implementation

Phillip Greenspun, *SQL for Web Nerds*, <http://philip.greenspun.com/sql/>.
(Required, but available for free on the web)

- A great introduction to RDBMS systems from the perspective of a web application developer.

Joe Celko, *SQL for Smarties: Advanced SQL Programming*, 4th edition, Morgan Kaufman, 2010. (Useful)

- A collection of wisdom on how working developers get useful things done in SQL.

Tom Christiansen, brian d foy, Larry Wall, Jon Orwant, *Programming Perl*, 4th Edition, O'Reilly and Associates, 2012. (Useful)

- This is the bible for the Perl language

Eben Hewitt, *Cassandra: The Definitive Guide*, O'Reilly and Associates, 2010. (Interesting)

- This describes the Cassandra distributed database, but is a bit dated

Objectives, framework, philosophy, and caveats

This course introduces the underlying concepts behind data modeling and database systems using relational database management systems (RDBMS), the structured query language (SQL), and web applications (Perl/JavaScript/CGI) as examples.

You will learn:

- How to model your data using the entity-relationship model
- How to design a normalized schema in the relational data model
- How to implement your schema using SQL
- How to keep your data consistent and safe with your schema using the ACID properties that a modern RDBMS gives you
- How to query your data using SQL
- How to interface to a modern RDBMS from a modern programming language.
- How such interfaces are used to create web applications
- How an RDBMS provides quick access to your data using indices, and how indices are implemented.
- How an RDBMS manages the storage hierarchy.
- How an RDBMS optimizes and execute your queries using the relational algebra, the theoretical underpinning of database systems.
- How an RDBMS implements transactions.
- Special topics: database security, including SQL injection attacks, and, if time: NoSQL/distributed databases, CAP theorem.

The textbook I have chosen is actually a combination of two books, an introduction to the concepts and use of databases and an introduction to the implementation of RDBMS systems. We will cover mostly the former. However, this is a very useful and essentially timeless book to have on your bookshelf for both elements. At the beginning of the course, we will also use a

very practical, and highly irreverent, free introductory book on relational databases and web applications. The idea is to dive in quickly and learn how to use a database as the core of a web application, and then to back up and consider data modeling, query modeling, and database systems more deeply.

This is a learn-by-doing kind of class. You will dive right in and modify a small database-based, mobile, geolocating web application. Next you will design and implement your own database-backed web application for financial portfolio management. Finally, you'll implement a B+Tree index data structure, a common index structure used in many database engines. The majority of the programming in this class will be from scratch. We will use SQL, Perl, JavaScript, and C++ on Linux systems.

Projects

At the beginning of the course, I will provide you with a simple web application that implements a mobile, map-based view of political candidates, committees, and contributors. This application is based on an Oracle database and provides a web interface using a combination of client-side JavaScript and a CGI application written in Perl that talks to the database via DBI. This is a very common form of web application. You will learn how this application works, and then you will extend it in several ways, focusing on the database backend. The goal is to immediately introduce you to SQL right away using a substantial dataset, namely the Federal Election Commission's disclosure database from 1980 to the present. This project will take three weeks.

The second project is focused on developing a simple financial portfolio manager that tracks a user's investments, and allows the user to "mine" historical financial data in several ways. I will give you a set of requirements and access to about 10 years of stock price data, and you will design and implement a database-backed web-based system. This project will take four weeks.

The third project is to build a B+Tree data structure. B+Trees are common on-disk (as opposed to in-memory) data structures used in relational database systems and many other systems. I will provide you with a framework, starter code, and a test harness.

Homework

There will be three homework sets that will be periodically assigned to help you improve your understanding of the material. These will focus on the entity relationship model, the relational model, and relational algebra.

Exams

There will be a midterm exam and a final exam. The midterm exam will take place in the evening outside of class. The final exam will not be cumulative.

Grading

15%	Project A: Dry-run project (“Red, White, and Blue”)
20%	Project B: Portfolio manager project
15%	Project C: B+Tree project
10%	Homework
20%	Midterm
20%	Final

Grades are not a competition.

Late Policy

For each calendar day after the due date for a homework or a lab, 10% is lost. After 1 day, the maximum score is 90%, after 2 days, 80%, etc, for a maximum of 10 days.

Cheating

Since cheaters are mostly hurting themselves, we do not have the time or energy to hunt them down. We much prefer that you act collegially and help each other to learn the material and to solve development problems than to have you live in fear of our wrath and not talk to each other. Nonetheless, if we detect blatant cheating, we will deal with the cheaters as per Northwestern guidelines.

Schedule

Note that the schedule is subject to change due to travel and other factors. I will announce schedule and due-date changes via email. If you do not receive a welcome email from me, please let me know.

Lecture	Date	Topics	Readings	Homework and Project
1	9/24	Class mechanics Introductory material, Web applications, client/server, and three-tier	G UW 1, 9.1, 9.3.1,9.3.2; PG preface + 1	Project A (RWB) out
<i>If you're unfamiliar with Unix, now would be a good time to view the Unix introduction video available from the course web site.</i>				
2	9/26	SQL in a nutshell, Start walk through of RWB (SQL)	PG 1-7, Perl HO, JS HO, WOT HO	Note: you might find PG 10 useful reading
3	9/29	How web applications work. Apache, CGI, Perl, JavaScript, DBI, RDBMS, SQL in a nutshell, continue walk through of RWB (SQL)	PG 1-7, Perl HO, JS HO, WOT HO, G UW 9.3.9	

4	10/1	Returning to the big picture: Relational and distributed databases, Data modeling, transactions/ACID, queries, abstracting storage+indices,	G UW 1; PG preface + 1	
5	10/3	Back to the nitty gritty: Perl	Perl HO	
6	10/6	Walk through RWB (Perl)	Perl HO	
7	10/8	Walk through RWB (Perl) <i>Instructor may be out of town</i>	Perl HO	
8	10/10	Database security topics or catchup <i>Instructor may be out of town</i>	G UW 10.1 (although lecture will focus elsewhere)	<i>Optional SQL Injection Attack Challenge</i>
9	10/13	NoSQL/Distributed Databases		
10	10/15	Data models and Data modeling: Why? Start Entity-Relationship: Entity sets, attributes, relationships, ER diagrams, instances, multiplicity, roles, multiway	G UW 2.1, 4.1-4.4	Project A (RWB) in. Project B (Portfolio) out HW 1 out
11	10/17	Entity-Relationship Model: conversion to binary relationships, subclassing, design principles	G UW 4.1-4.4	
12	10/20	Entity-Relationship Model: constraints, weak entity sets	G UW 4.1-4.4	
13	10/22	Relational Data Model: basics, translating from ER to relational	G UW 2.2, 2.3, 4.5	HW 1 in HW 2 out
14	10/24	Relational Data Model: basics, translating from ER to relational	G UW 2.2, 2.3, 4.5	
15	10/27	Relational Data Model: subclasses, functional dependencies	G UW 4.6, 3.1-3.2	
16	10/29	Relational Data Model: Schema design and normal forms	G UW 3.3-3.5, 3.6.6	
17	10/31	Relational Data Model: Multivalued dependencies	G UW 3.6	
18	11/3	Slack time or Midterm Review		HW 2 in

<i>Midterm Exam: Tuesday, 11/4, 6-8pm location TBD</i>				
<i>Midterm Exam will cover Lectures 1-18</i>				
19	11/5	Relational Algebra: Sets: union, intersection, difference, selection, projection, Cartesian product, and cross, inner, outer, left, right joins	G UW 2.4, 5.1-5.2	HW 3 out
20	11/7	Relational Algebra: Bags, equivalent expressions, some extended operators	G UW 5.1-5.2	
21	11/10	Relational Algebra: grouping, constraints, data-mining	G UW 5.1-5.2, 2.5	
22	11/12	Advanced SQL: strings, regular expressions, date/time, nulls, 3-valued logic, explain plan, subqueries in/exists/>all/>any, correlation	G UW 6	Project B (Portfolio) in Project C (B+Tree) out
23	11/14	Advanced SQL: insert/update/delete, multi-statement transactions using PL/SQL; create schemas: bit-fields, decimal, blob; drop, alter; indexes; views	G UW 6, 7, 8	
24	11/17	Advanced SQL: Constraints, Triggers, systems aspects.	G UW 6, 7, 8	HW 3 in
25	11/19	Implementation: Storage and Representing Data	G UW 13	
26	11/21	Implementation: Indexes, Btrees	G UW 14.1, 14.2	
27	11/24	Implementation: Indexes, Hashes	G UW 14.3	
27a	11/26	<i>(probably no class)</i>		
<i>Thanksgiving Break (6pm, 11/25; classes resume on 12/1)</i>				
28	12/1	Implementation: Indexes, Bitmaps	G UW 14.7	
29	12/3	Implementation: Transactions (Logging, Locking)	G UW 17.1-17.4, 18.1-18.3	
30	12/5	Implementation: Transactions (Logging, Locking)	G UW 17.1-17.4, 18.1-18.3	Project C (B+Tree) in
<i>Final Exam, Thursday, 12/11, 7-9pm, in our classroom.</i>				
<i>Final Exam will Cover Lectures 19-30</i>				

PG = Phillip Greenspun, *SQL for Web Nerds*

GUW = Hector Garcia-Molina, Jeffrey D. Ullman, Jennifer D. Widom, *Database Systems: The Complete Book*

Perl HO = *Perl in a Nutshell* handout

JS HO = *JavaScript Model in a Nutshell* handout

WOT HO = *Using Databases in the Web of Things Environment* handout