

# CS 145 – Introduction to Databases

## Course Information

### Course Content

This course provides the student with an introduction to the design of databases and the use of database management systems for applications. We will cover object-oriented (O/O) and entity-relationship (E/R) approaches to database design. We will cover the relational data model, mappings from O/O and E/R designs to relations, relational design issues, and how one uses a relational database management system to implement a database. Abstract query languages such as relational algebra will be discussed, along with extensive coverage of SQL—the standard language for creating, querying, and modifying relational databases. We will cover a variety of other issues important to database designers and users, including views, integrity constraints, triggers, transactions, and security. Finally, time permitting, we will consider advanced database management features such as object-relational systems and data warehousing.

This course is aimed at database design and use, and the implementation of database applications. It does *not* cover the implementation of database management systems. This topic is covered in the follow-on courses CS 245 (formerly CS 245A) and CS 346 (formerly CS 245B).

### Time and Place

Mondays and Wednesdays 11:00–12:15, Gates B01 (Hewlett-Packard Auditorium)

### Course Personnel

Instructor Jennifer Widom  
Office: Gates Building room 422  
Phone: 723-7690  
Office hours: Mondays 12:30–1:30, Wednesdays 12:30–2:00

T.A. Roy Goldman  
Office: Gates Building room 508  
Phone: 725-3545  
Office hours: TBA  
E-mail: [cs145ta@cs.stanford.edu](mailto:cs145ta@cs.stanford.edu)

T.A. Jun Yang  
Office: Gates Building room 505  
Phone: 723-4214  
Office hours: TBA  
E-mail: [cs145ta@cs.stanford.edu](mailto:cs145ta@cs.stanford.edu)

Secretary Sharon Lambeth  
Office: Gates Building room 419  
Phone: 725-3358  
E-mail: [lambeth@db.stanford.edu](mailto:lambeth@db.stanford.edu)

## Prerequisites

- CS 107 Programming Paradigms
- CS 109 Introduction to Computer Science

These prerequisites are a guide but not an absolute requirement. CS 145 should be accessible to anyone with a solid Computer Science foundation that includes a reasonable amount of programming (as in CS 107) and knowledge of basic Computer Science theory (as in CS 109).

## Books

The following two books are required for the course. The first book accompanies the lecture material, while the second book serves as a reference for SQL.

- *A First Course in Database Systems*  
J.D. Ullman and J. Widom, Prentice Hall, 1997
- *A Guide to The SQL Standard* (4th edition)  
C.J. Date and H. Darwen, Addison-Wesley, 1997

In addition, the following book is recommended for further details on the ODMG standard.

- *The Object Database Standard: ODMG 2.0*  
R.G.G. Cattell et al., Morgan Kaufmann, 1997

One copy each of these three books will be on reserve at the Math and Computer Science library. You may also find it useful to have a book devoted specifically to the Oracle relational database management system, which you will be using for your programming assignments. Numerous Oracle books and manuals exist and they change rapidly, so we will not recommend a specific one—you can easily browse them at the Stanford Bookstore. We will select one or two to put on reserve at the Math & CS library.

## Assignments

There will be seven written assignments and five programming assignments. Written and programming assignments will be given on Wednesdays and will be due the following Wednesday. There will be an assignment due every full week except the first one and the week of the midterm exam. Written assignments should be turned in during class or at the course secretary's office: Gates Building room 419. Programming assignments will be turned in electronically.

## Late Policy

**On-campus students:** Written assignments are due at 5:00 PM on Wednesdays. Programming assignments must be *received* by the TA's before Wednesday midnight. Assignments less than 24 hours late will be penalized 10%. Assignments more than 24 but less than 48 hours late will be penalized 30%. No assignments will be accepted more than 48 hours late—i.e., no written assignments will be accepted after 5:00 PM on Fridays, and all programming assignments must be received by the TA's before Friday midnight.

**SITN students:** The late policy for programming assignments is the same as for on-campus students as specified above. For written assignments, assignments due on Wednesday must be

timestamped by the Thursday courier. Late assignments timestamped by the Friday courier will be penalized 20%. Assignments timestamped later than Friday will not be accepted. Assignments not sent by SITN courier are subject to the late policy for on-campus students specified above. No late written assignments are permitted for remote SITN students.

**For emergencies:** Since emergencies do arise, each student is allocated four “chits”. Each chit may be used to turn in a single assignment (written or programming) up to 24 hours late with no penalty, or two chits for up to 48 hours late. For SITN students, two chits are required to turn in a written assignment one day late. As with penalized late work, no assignments will be accepted after Friday. We’ll keep track of your chits automatically.

*This late policy will be strictly enforced.*

## Examinations

The midterm exam will be held on Monday May 4 in class at the regular time. The final exam will be held on Wednesday June 10 at 8:30 AM (location to be announced). Alternative or make-up exams will *not* be given. All on-campus and SITN students must attend the midterm and final exams on the Stanford campus at the scheduled time and place.

## Grading and Honor Code

Written assignments	15%
Programming assignments	35%
Midterm exam	20%
Final exam	30%

Under the Honor Code at Stanford, each of you is expected to submit your own work in this course, including written assignments, programming assignments, and exams. On many occasions when working on assignments it is useful to ask others (the instructor, the TA’s, or other students) for hints or debugging help, or to talk generally about the written problems or programming strategies. Such activity is both acceptable and encouraged, but you must indicate on both your written and your programming assignments any assistance you received. Any assistance received that is not given proper citation will be considered a violation of the Honor Code. In any event, you are responsible for understanding and being able to explain on your own all written and programming solutions that you submit. The course staff will pursue aggressively all suspected cases of Honor Code violations, and they will be handled through official University channels.

## Help Session

There will be a help session held most Mondays from 5:15–6:05 PM in Skilling 191. It will be televised live to SITN students on channel E1. The session will be conducted by one of the TA’s and is entirely optional. During the session, the TA will work problems and/or answer student questions, but the TA will not deliver a prepared lecture. If you attend the session (live or via television), please come armed with questions or discussion points.

## Web Page and Newsgroup

The course Web page is located at URL <http://www-leland.stanford.edu/class/cs145>. All handouts that are available on-line will be linked to the Web page in postscript format. In addition

to handouts, the Web page will contain other useful information regarding the course, such as a list of announcements and frequently-asked questions.

The course newsgroup is `su.class.cs145`. The newsgroup is accessible from any Unix machine on campus. Please check the course newsgroup frequently—we will assume that all students are reading the newsgroup at least every couple of days, and we will not hesitate to make important announcements via the newsgroup. If you are unable to read the newsgroup for some reason, please let the TA's know as soon as possible.

### E-mail and News Protocol

All questions or messages from students that could be of general interest to the class should be posted to the newsgroup. This includes, for example, any technical questions regarding lecture material, reading, homework problems, upcoming exams, etc. We very much encourage students in the class to post responses to questions. The TA's will also monitor these messages, and they will post responses to questions that have not previously been asked or answered. *Before sending a question, please do make sure that you've read all previous news messages and that your question has not yet been discussed.* To post a message, send the posting as a mail message to `cs145@cs`.

Private e-mail regarding the course should be sent to `cs145ta@cs.stanford.edu`. This address is only for non-technical questions applying to an individual student (e.g., my Oracle password seems to have changed, or I didn't get back my graded assignment). These messages will be read by course staff only. General technical questions sent to this address will be forwarded to the newsgroup. *One exception:* Remote SITN students tend to watch lectures and work on assignments up to a week later than everyone else. If you have a general technical question, you're sure it hasn't been answered already in the newsgroup, and you're well behind the on-campus students, please send it to `cs145ta@cs`. Clearly identify at the top of the message that you're a remote SITN student.

Please do not send any messages to individual TA's or to the instructor. Private e-mail addresses will not be monitored for course-related messages, so there may be significant delay. Messages appropriate for individuals or for the instructor will be routed from `cs145ta@cs` or `lambeth@cs` (the course secretary) appropriately.

### Programming

Programming assignments will use the Oracle relational database management system (version 7.3.2) and the C or C++ programming language. Oracle and C/C++ are available on the Sun "elaine", "adelbert", and "epic" workstations running Solaris (version 2.5.1). These machines are located on the second floor of Sweet Hall. To open an account on these machines, type `open` at the `login:` prompt and follow the instructions. SITN students can access the Sun workstations remotely via dial-in or telnet. We will assume that students are proficient already with Unix and C. (C++ is not required, but it is allowed for those who prefer to use it.)

### Using Your Own Workstation and Software

If you have one available, you may use a DBMS other than Stanford's Oracle system, or you may use a non-standard interface to Oracle (e.g., through ODBC or JDBC). Be warned, however: You are responsible for ensuring that the DBMS and/or interface you use supports all the same features as Oracle version 7.3.2. Those features include but are not limited to support for interactive and embedded full-feature SQL, indexes, transactions, authorization, views, constraints, and triggers.

If a student uses a different system and then discovers that it does not offer the features needed for the class, they will need to port their project to the Stanford Oracle system at that time. We will not be able to provide support or answer system-specific questions for any system or interface other than Stanford's Oracle C/C++ system.

### Special Note for SITN Students

While we understand the significant demands of taking a course while working, we would like to make it clear at the outset that no special exceptions can be made for SITN students. Local SITN students must follow the lectures, turn in the assignments, and attend the exams on exactly the same schedule as on-campus students. Please ensure in advance that you will not have any business trips or important deadlines that conflict with the course, since no exceptions will be made. Also note that like on-campus students, for emergencies you are allocated four unpenalized late days for assignments (see "Late Policy" above).

### Tentative Syllabus

Week	Date	Topic	Reading in Ullman/Widom
1	April 1	Intro. to course and subject	1
2	April 6 April 8	Object-oriented database design Entity-relationship model	2.1, 2.3–2.5 2.2–2.6
3	April 13 April 15	Relational model Relational database design	3.1–3.4 3.5–3.8
4	April 20 April 22	Relational database design (cont'd) Relational algebra	4.1
5	April 27 April 29	SQL SQL (cont'd)	5.1–5.7
6	May 4 May 6	<i>Midterm exam</i> Indexes, views	5.8
7	May 11 May 13	Views (cont'd), constraints Constraints (cont'd), triggers	6.1–6.5 6.6
8	May 18 May 20	Programming with SQL, transactions Transactions (cont'd), security	7.1–7.3 7.4
9	May 25 May 27	<i>Holiday – no class</i> Object-relational SQL3, OQL	8.1–8.3, 8.5–8.7
10	June 1 June 3	Topic TBD Topic TBD	
11	June 10	<i>Final exam</i>	