

CS 2550: Principles of Database Systems

Department of Computer Science, University of Pittsburgh

Course Reference Number (CRN): 21095

Term: Spring 2012 (12-2 or 2124)

When: Monday & Wednesday, 11:30 am – 12:45 pm (regular)
Friday 11:30 am – 12:45 pm (make-up)

Where: Room 5313, Sennott Square Building (SENSQ)

Instructor: Prof. Panos K. Chrysanthis

Email: panos@cs.pitt.edu

Web: <http://www.cs.pitt.edu/~panos>

Office: 6421 Sennott Square

Phone: 412-624-8924

Office hours: Monday, 2:30 – 4:00 pm
Wednesday, 2:30 – 4:00 pm
by appointment

Teaching Assistants: Panickos Neophytou

Email: panickos@cs.pitt.edu

Office: 6414 Sennott Square

Phone: 412-624-8443

Office hours: Tue & Thu: 3:15 – 4:15 pm
Fridays: 2:00 – 3:00 pm
by appointment

Course Description: The principle objective of this course is to provide an in-depth knowledge of Database Management Systems design. Important aspects of distributed database systems and new data processing paradigms (such as data streams) will also be covered.

Prerequisite: Knowledge of data structures and files, basic operating systems concepts, and exposure to data models (i.e., CS1555 or its equivalent). Working knowledge in Java (or C/C++) is assumed.

Class Web Page: <http://db.cs.pitt.edu/courses/cs2550/spring2012>

All handouts and class notes will be published on the class web page. You are expected to check this page frequently (at least twice a week).

Note on Email Communication: In order to receive the highest priority, you must include the keyword `cs2550` in the subject line of your email messages. (This rule applies to all email messages, that are sent to the instructor, the teaching assistant, or both: cs2550-staff@cs.pitt.edu).

Textbooks: *Database Systems: The Complete Book* (2nd Edition). Hector Garcia-Molina, Jeffrey D. Ullman, Jennifer Widom, 2009. (ISBN-10: 0-13-187325-3; ISBN-13: 978-0-13-187325-4)

Reference: *Transactional Information Systems*. Weikum and Vossen, Morgan Kaufmann, 2002.
Concurrency and Recovery in Database Systems. Bernstein, Hadzilacos and Goodman, Addison-Wesley, 1987 [URL: <http://research.microsoft.com/pubs/ccontrol/> (in PDF)]
Transaction Processing: Concepts and Techniques. Cray and Reuter, Morgan Kauffman, 1993

[please turn over]

Course Grading:

Assignments & Project	40%	
Midterm Exam	30%	Wednesday, February 29, 11:30 am – 12:45 pm
Final Exam	30%	Wednesday, April 25, 11:30 am – 12:45 pm
Participation	5%	bonus

Grading Policy: Unless explicitly noted otherwise, the work in this course is to be done independently. Discussions with other students on the assignments should be limited to understanding the statement of the problems. Cheating in any way, including giving your work to someone else will result in an F for the course and a report to the appropriate University authority.

Marks can be appealed up to two weeks after they have been posted, after that no appeals will be considered.

Submission & Late Policy: All written assignments must be submitted electronically and there is no late submission. An assignment which is late will be accepted *only* under special circumstances with the instructor's permission prior to its deadline. In such a case, the instructor will determine any penalty in a fair manner.

Make-up Policy: Students are expected to take both midterm and final exams. Make-up exams will only be given in the event of a medical situation or an emergency, and only if this is documented and the instructor is notified *immediately if in advance is not possible*. Missing an exam will result in a zero for the exam.

Students with Disabilities: If you have a disability for which you are or may be requesting an accommodation, you are encouraged to contact both your instructor and Disability Resources and Services, 216 William Pitt Union, 412-648-7890 or 412-383-7355 (TTY) as early as possible in the term. DRS will verify your disability and determine reasonable accommodations for this course.

Religious Observance: In order to accommodate the observance of religious holidays, students should inform the instructor of any such days within the first two weeks of the term by email.

Outline: A detailed reading guide will be published on the web page, along with the class notes. Although we will follow the basic outline of the textbooks, we will re-order the material to improve the flow. Additional material will be introduced as needed.

Time permitting, we will cover the following topics:

1. Introduction to Database Systems & Data Models
2. Physical Database Organization
3. Transactions, Serializability Theory
4. Locking Schedulers, Deadlocks
5. Non-Locking Schedulers
6. Multiversion Concurrency Control
7. Recovery Concepts and Techniques
8. Query Processing and Optimization
9. Introduction to Distributed Database systems
10. Distributed Concurrency Control and Recovery
11. Extended Transactions and Workflows
12. Data Streams