

CSC2509 Spring 2006

Data Management Systems

This course covers principles of database management systems, especially relational systems, from the perspective of the system developer and database theorist.

Lectures: Thursdays 3-5pm, BA 1240.

Instructor: Anthony Bonner, BA 4268, 978-7441, bonner@cs.toronto.edu

Office Hours: by appointment.

Teaching Assistant: TBA

Web Page: Accessible via my home page, www.cs.toronto.edu/~bonner

Grading Scheme: Four assignments, 25% each. No late assignments will be accepted.

Prerequisites: A facility for mathematics and programming. Experience with Unix and text editing. Some knowledge of C would be helpful.

References:

1. Anthony J. Bonner. *CSC2509s Lecture Slides, Parts I and II*. Available on the course web page. Required.
2. J. D. Ullman, *Principles of Database and Knowledge-Base Systems, Volumes I and II*, Computer Science Press, 1988/89. (*Useful for this course*)
3. Bernstein, Hadzilacos and Goodman, *Concurrency Control and Recovery in Database Systems*, Addison-Wesley, 1987. (*Useful for this course*)
4. Chamberlain, *A Complete Guide to DB2 Universal Database*, Morgan Kaufmann. (*Useful for this course*)
5. Kernighan and Ritchie, *The C Programming Language*, Prentice Hall. (*Useful for this course if you don't know C*)
6. J.N. Clarke, *A Student's Guide to CDF*. (*Short introduction to Unix.*) Available at the University Book Store.
7. Lewis, Bernstein and Kifer, *Databases and Transaction Processing*, Addison Wesley.
8. Silberschatz, Korth and Sudarshan, *Database System Concepts*, WCB McGraw-Hill.
9. Ramakrishnan and Gehrke, *Database Management Systems*. McGraw-Hill.
10. O'Neil, *Database Principles, Programming, Performance*. Morgan Kaufmann.

Topics

The Relational Model: Relations and relational algebra.

Application Programming: SQL, Embedded SQL.

Query Optimization: Algebraic optimization, join order, estimating query sizes.

Distributed Databases: Query processing, semi joins, communication cost.

Transaction Management: Transaction models, concurrency control, livelock and deadlock, serializability, protocols, timestamps, crash recovery.

Database Design Theory: Functional dependencies, Armstrong's axioms, normal forms, decomposition algorithms.