CMPT 354
Database Management Systems

Dr Evgenia Ternovska
Associate Professor

Simon Fraser University

September 2017

What is a database?

A collection of data items, related to a specific enterprise, which is structured and organized so as to be more easily accessed, managed, and updated

What is a database?

A collection of data items, related to a specific enterprise, which is structured and organized so as to be more easily accessed, managed, and updated

What is a Database Management System (DBMS)?

- system software for creating and managing databases
- mediates interaction between end-users (or applications) and the database
- ensures that data is consistently organized and remains easily accessible

Goals of This Course

Databases are everywhere

- Electronic commerce, websites (e.g., Wordpress blogs)
- Banking applications, booking systems, even your laptop
- Data science, data analytics

Three types of users

1. End-users
2. Administrators
3. Implementors

Goals of this course

- Teach you to be **good end-users** (SQL is not enough)
- Provide you with **solid foundations** of how a DBMS works (and so understand something of the other two roles)

Topics To Be Covered

(not necessarily in this order)

- An introduction to data modelling.
- Database management systems and the relational model.
- Data definition, data manipulation, queries.
- Views, constraints, triggers.
- Query languages: relational algebra and relational calculus, logical foundations.
- Commercial languages (such as SQL and QBE).
- Data definition, data manipulation, triggers.
- Functional dependencies. Normal forms.
- Design theory. The entity-relationship model, E-R diagrams.

Topics To Be Covered (Cont.)

- Transaction processing and concurrency control, transaction management, serializability, locking.
- Database access from applications: embedded and dynamic SQL
- Storage and indexing: B-trees, hashing.
- Query evaluation and optimization: join strategies, query plans.
- Deductive databases: Datalog and recursive queries
- Incomplete data: null values and certain answers

Other topics may include some recent developments in databases, XML, constraint databases, and semistructured data.

Prerequisites

CMPT 225, (and MACM 101 or (ENSC 251 and ENSC 252)).

Some background in discrete mathematics:

- Set theory (sets, set operations, relations, orders)
- **Mathematical logic** (predicate logic, inference, satisfiability)
- Combinatorics (permutations, combinations, partitions)
- Graph theory (directed/undirected graphs, trees)
- Algorithm complexity analysis (Big-O notation)

Mathematical Logic

Essential to

- understand relational calculus (as well as Datalog)
- write correct SQL queries

We will have a brief introduction to logic at the level required for the course

- refresher for who studied logic before
- crash course for newcomers

Textbook(s)

Main text:

Ramakrishnan, Gehrke: **Database Management Systems** McGraw-Hill, 3rd edition

Most lectures will be following this textbook, however the order of topics might be slightly different

Further reading: Abiteboul, Vianu, Hull **Foundations of Databases** Addison-Wesley, 1995

- Mostly theoretical topics
- Out of print but freely available (for personal use only)

More books will be recommended as the course progresses

Contacts

Course Website

- Direct link:
  https://piazza.com/sfu.ca/fall2017/cmpt354/home
- Lecture notes, homework, assignments
- Cancellations and changes of location/date/time of classes

Piazza Forum

- **Preferred** channel of communication
- Course material + class discussions
- More on it on next slide ...

Email

- **Discouraged** (use Piazza instead)
- Write to ter@sfu.ca with subject line **exactly** [CMPT354]
- Use only your SFU student email (not e.g. gmail)

Piazza

This term we will be using Piazza for class discussion
Get help easily from classmates, the TA and myself

Rather than emailing questions, post your questions on Piazza

- You can post privately to instructors (TA and me)
- You can post anonymously to classmates

Find our class page at
https://piazza.com/sfu.ca/fall2017/cmpt354/home

Signup for the class (after receiving an email from the sytem)
with your SFU student email address (e.g., 1234567@sfu.ca)
by clicking the link in the email

Assessment

The course assessment will be based on (pretty standard):

- Homework: Pen and paper and SQL/programming at CSIL
- Midterm(s): Pen and paper (closed book)
- Final Exam: Pen and paper (closed book)

details will be specified next week