TDBD15 - Advanced Data Models and Systems

Tuesday March 26, 2002 Michael Minock

What is an "Advanced Data Model"?

Aside from the *conceptual* Extended Entity-Relationship model (EER), the models we talk about will be *representational*.

A representational model:

- has a well defined logical (or graphical) characterization
- provides declarative (or at least abstract) access/update operations
- (possibly) models behaviors
- provides *some* implementation direction
- hides storage and low-level implementation details

Some Candidates

- **The Relational Model**
- Hierarchical Data Model
- Network Data Model
- Semantic Networks
- **The Object-Relational Model**
- **The Nested Relational Data Model**
- **Object Data Model**
- **XML**
- **Datalog**
- Forward-Chaining Rule-bases (Production Systems)
- **Description Logics**
- **Conditional Knowledge-bases**
- Ontologies
- ...

What is a System? (in this class)

A system:

- implements a data model ± features
- facilitates sharing of data (knowledge) among systems using the same model (physical data independence)
- supports easy integration with tools built around the model
- marketed to support a (wide) class of real world problems

Some systems

- (Relational) Oracle, Sybase, Access, MySQL, DB2, SQLServer, ...
- (Object Relational)Oracle 8, **PostgrSQL**, Informix, ...
- (Object) Objectivity, Orion, O2, Object Store, ...
- (Description Logics) KL-one, **Classic**, LOOM, Back, Crack, ...
- (Deductive) **LDL**, NAIL, CORAL, ...
- (XML) XQL, XML-QL ...
- (Hierarchical) IMS, ...
- (Network) IDS, ...
- (Production System) Ops-5, CLIPS,...

How are we going to study all of this?

We focus on each data model in **bold** above.

Once motivated, we shall give a more or less formal definition of the model.

For the **bold-underlined** models above, we shall enjoy practicing these concepts on a system.

Note that we shall always try to employ the most direct technique to fit requirements – many times bringing us back to the relational model. (Occam's razor applied to data modeling)

Through this we shall discover what each of our candidate models can offer by uncovering canonical problems that they solve well.

Also we consider extensions to the relational model to handle special representation problems: space and time.

We shall also look into handling uncertainty and data mining.

So this class assumes an understanding of the relational model.

But How Exactly?

Lectures Continued

What is expected of you?

Final Projects: Students may work in groups of up to four persons. The students must successfully propose, in writing and in person, a project of joint interest by the 13th lecture of class. Projects will be presented (with demonstrations) at the end of the course. This will be 40% of the grade.

Final Exam: This will be based on the problems that shall be discussed during the course of the term. This will be worth 50% of the grade.

The other 10% of the grade will be based on 3 systems exercises.

I will award bonus points occasionally. But to get your bonus points you need to get 50% on the final exam.

Late Work

None of the work is obligatory. You could take just the exam - but you would have to score perfect to pass the course.

Late work will lose points.

See the syllabus for more details.

Project

- Proposal Phase
  - 350 words maximum
  - All members present to discuss project
  - (10 BONUS POINTS!) - for each member of the first three groups with a proposal accepted.
- Interaction
  - See me early if there are problems!
  - Accepted proposals posted to net
  - Code sharing between groups encouraged
- Presentation
  - Written report
  - 20 minute talk
  - System demonstration
- Evaluation

More details will be on the web page

Exam

I will, along the way, propose some paper exercises.

Variants will show up on the exam.

Do the readings and understand the slides.

There will be an exam review lecture.