Foundations of Query Languages

Dr. Fang Wei

Lehrstuhl für Datenbanken und Informationssysteme
Universität Freiburg

SS 2011

Contents

- Conjunctive queries
- Query optimization, hypergraphs
- Database design, normal forms
- First order logic as query language
- Datalog
- ...

Main Subject of FQL

A deep understanding of various query languages: syntax, declarative and operational semantics

- How much resource (time, space) do we need for the computation of these semantics? $\Rightarrow$ Complexity
- What kind of properties can a given query language express?
- Is Q_1 more expressive than Q_2? $\Rightarrow$ Expressive power

A dream query language should have:

- lower complexity, and
- more expressive power

Main Subject of FQL

A deep understanding of various query languages: syntax, declarative and operational semantics

- How much resource (time, space) do we need for the computation of these semantics? $\Rightarrow$ Complexity
- What kind of properties can a given query language express?
- Is Q_1 more expressive than Q_2? $\Rightarrow$ Expressive power

A dream query language should have:

- lower complexity, and
- more expressive power

Main Subject of FQL

A deep understanding of various query languages: syntax, declarative and operational semantics

- How much resource (time, space) do we need for the computation of these semantics? $\Rightarrow$ Complexity
- What kind of properties can a given query language express?
- Is Q_1 more expressive than Q_2? $\Rightarrow$ Expressive power

A dream query language should have:

- lower complexity, and
- more expressive power

Main Subject of FQL

A deep understanding of various query languages: syntax, declarative and operational semantics

- How much resource (time, space) do we need for the computation of these semantics? $\Rightarrow$ Complexity
- What kind of properties can a given query language express?
- Is Q_1 more expressive than Q_2? $\Rightarrow$ Expressive power

A dream query language should have:

- lower complexity, and
- more expressive power

Main Subject of FQL

A deep understanding of various query languages: syntax, declarative and operational semantics

- How much resource (time, space) do we need for the computation of these semantics? $\Rightarrow$ Complexity
- What kind of properties can a given query language express?
- Is Q_1 more expressive than Q_2? $\Rightarrow$ Expressive power

A dream query language should have:

- lower complexity, and
- more expressive power

Main Subject of FQL

A deep understanding of various query languages: syntax, declarative and operational semantics

- How much resource (time, space) do we need for the computation of these semantics? $\Rightarrow$ Complexity
- What kind of properties can a given query language express?
- Is Q_1 more expressive than Q_2? $\Rightarrow$ Expressive power

A dream query language should have:

- lower complexity, and
- more expressive power

Prerequisite of the course

- You took the course Datenbanken und Informationssysteme, or
- You can prove that you have enough background knowledge which covers the above course
- Course on computational complexity could be a plus, but not a must
- Sufficient knowledge on English for reading and writing research papers

Course organization

- Mondays and Wednesdays at 14:00
- Exercises on Wednesday
- Book for reference: *Foundations of databases*, Serge Abiteboul, Richard Hull and Victor Vianu, 1995, Addison-Wesley. (out of print, you will get the PDF data link in the class)
- Other research papers will be posted online during the semester
- Reading assignment: course-related research papers will be posted online, and you are expected to write reviews on them.

The relational model

- A simple data model (relation)
- Declarative query language (SQL)
- Application users specify what answers a query should return, not how DBMS picks the best execution strategy based on availability of indexes, data/workload characteristics, etc.

Provides physical data independence

Physical data independence

- Applications should not need worry about how data is physically structured and stored
- Applications should work with a logical data model and declarative query language
- Leave the implementation details and optimization to DBMS
- The single most important reason behind the success of DBMS today

And a Turing Award for E. F. Codd

Structure of DBMS

- A typical DBMS has a layered architecture.
- The figure does not show the concurrency control and recovery components.
- This is one of several possible architectures; each system has its own variations.
- All the layers except for optimization and execution have to consider concurrent control and recovery.