AC-04-17

Outline

Distributed Data Management

COMP-614

Bettina Kemme

Introduction:

While for a long time, sophisticated distributed data management was only implemented in large companies like banks and the stock exchange, the prevalence of distributed infrastructure in smaller companies, and the advanced electronic communication between different organizations, has rapidly lead to the development of new forms of distributed information systems. On the local scale, this means that the data and applications within an organization are distributed and replicated over many computers in order to provide scalability and fault-tolerance. On the global scale, this means that different organizations expose data or services to outside users, e.g., through web based or programmatic interfaces. In the extreme case, each individual can provide data or services to the rest of the world, as is done, e.g., in music sharing environments. These new developments require sophisticated solutions to provide data consistency, easy, efficient and authorized access to data, and data availability despite failures. Furthermore, tools and abstraction mechanisms are needed to hide the complexity of the underlying infrastructure.

This course will introduce and discuss in depth several of the issues related to distributed data management, both from a theoretical and practical point of view. The objective is to present the state-of-the-art in distributed data management, to give an overview of the problems and their nature, how they can be solved and how these solutions are implemented in practice.

Description:

There are many important issues in distributed data management.

- There exist a wide range of architectural design alternatives for distributed information systems, from homogeneous system over federated architectures to peer-to-peer databases.
- Data consistency is a challenging task in distributed environments. Transaction management addresses these issues.
- Data replication and data caching is used to provide scalability, fast local access and fault-tolerance. Many different approaches exist tailored for the specific needs of the application.

- Querying distributed data requires to find and collect data from different resources, and post-process it in order to present to the user.
- Data in different data sources might have different formats, making it necessary to match and reformat the data. This is the task of schema integration.
- Many advanced topics like authorization and access control, communication support for data intensive tasks, or load balancing, are options to further explore if time allows.

Since this is a graduate course, the exact schedule might considerably change from year to year, and different emphasis might be given to different topics. For instance, one year the main focus might be data consistency, data replication and caching. In another year, the focus might lie on distributed query processing. Also, the instructor might choose one year to focus on a single architecture and discuss the topics only from the perspective of this architecture, e.g., peer-to-peer systems, mobile databases, clustered databases, etc.

Below is a schedule that covers all topics considering different architectures with a slight emphasis on consistency and replication. It is designed for a course layout in which the first 9 to 10 weeks, the topics described above are discussed one by one. This includes an introduction to each topic, the presentation of basic algorithms and techniques, and their application in practice. The remaining weeks will look at specific sub-topics in more detail, discussing recent advances in research and industry. Since students are supposed to provide presentations during class, some of these last lectures will be given by students, depending on enrollment numbers.

Overview (0.5 week)
Architectures (1.5 weeks)
Data Consistency (3 weeks)

- Advanced transaction models
- Advanced concurrency control
- Distributed recovery

Data Replication and Caching (2-3 weeks)
Distributed Query Processing (1 week)
Schema Integration (1 week)
Advanced Topics, Discussion of Specific Solutions to above Topics (3-4 weeks)

Reading Material:

The course is based on several textbooks and research papers. Additionally, lecture notes will be provided for some of the lectures. Chapters from textbooks and survey papers will provide an overview of the different topics, deepen the understanding of the material or help in reading up the background information. Out of the papers, each student is expected to read a subset of them. The reading list, especially the papers, is subject to regular adjustments and replacements in order to reflect the focus of the course, and newest research achievements.

Books

- Distributed Systems - Principles and Paradigms, by A.S. Tanenbaum and M. van Steen, Prentice Hall, 2002. or
- Distributed Systems - Concepts and Design, 3rd ed., G. Coulouris, J. Dollimore, and T. Kindberg, Addison Wesley, 2001.
- M.T. Özsu and P. Valduriez, Principles of Distributed Database Systems, 2nd edition, Prentice-Hall, 1999.
- P. Bernstein and E. Newcomer, Principles of Transaction Processing Morgan Kaufmann, 1997
- G. Weikum and G. Vossen, Transactional Information Systems Morgan Kaufmann, 2002
- M. Buretta, Data Replication , Wiley, 1997

Research papers will be selected from well established conferences and journals in the relevant fields.

Pre-requisites:

COMP-421 *Database Systems* and one of the following: COMP 512 *Distributed Systems,* COMP 435 *Basics of Computer Networks,* COMP 535 *Computer Networks. 1*

Evaluation:

The evaluation scheme will be as follows:

- 10% class participation
- 30% assignments
- 60% project
  - 20% survey report
  - 20% presentation of survey or specific paper
  - 20% research report

Assignments can be either a set of problems to be solved, written and/or oral summaries of research papers, or paper critiques.
The project is an in-depth study of one research area. It consists of three deliverables. A survey report presents an overview of recent research advances within a particular area (15-20 pages, 12pt, 1.5 space). A presentation sometime during the term will present this report to the other students, a short presentation at the end of the term will present the own research results. The third delivery is a research report (15-20 pages, 12pt, 1.5 space). There are three alternatives.

- A research proposal for a relatively complex problem. It should be sufficiently detailed to give an indication that the student knows how to solve the problem.
- A solution to a small-size, more specialized problem.

- An implementation and evaluation of algorithms presented in some research paper.

A note on academic integrity

McGill University values academic integrity. Therefore all students must understand the meaning and consequences of cheating, plagiarism and other academic offences under the Code of Student Conduct and Disciplinary Procedures (see http://www.mcgill.ca/integrity for more information).