

Distributed Database Systems

© 2001 M. Tamer Özsu and Patrick Valduriez

Outline

- Introduction
- Background
- Distributed DBMS Architecture
- Distributed Database Design (Briefly)
- Distributed Query Processing (Briefly)
- Distributed Transaction Management (Extensive)
- Building Distributed Database Systems (RAID)
- Mobile Database Systems
- Privacy, Trust, and Authentication
- Peer to Peer Systems

Instructor Introduction

Bharat Bhargava

Professor of Computer Sciences, Purdue University

West Lafayette, IN 47907

Phone: 765-494-6013, Email: bbshail@purdue.edu

Bharat Bhargava is a professor of the Department of Computer Science with a courtesy appointment in the School of Electrical & Computer Engineering at Purdue University. Professor Bhargava is conducting research in security and privacy issues in distributed systems. This involves host authentication and key management, secure routing and dealing with malicious hosts, adaptability to attacks, and experimental studies. Related research is in formalizing evidence, trust, and fraud. Applications in e-commerce and transportation security are being tested in a prototype system. Based on his research in reliability, he is studying vulnerabilities in systems to assess threats to large organizations. He has developed techniques to avoid threats that can lead to operational failures. The research has direct impact on nuclear waste transport, bio-security, disaster management, and homeland security. These ideas and scientific principles are being applied to the building of peer-to-peer systems, cellular assisted mobile ad hoc networks, and to the monitoring of QoS-enabled network domains.

Instructor Introduction

In the 1988 IEEE Data Engineering Conference, he and John Riedl received the best paper award for their work on "A Model for Adaptable Systems for Transaction Processing." Professor Bhargava is a Fellow of the Institute of Electrical and Electronics Engineers and of the Institute of Electronics and Telecommunication Engineers. In 1999, he received the IEEE Technical Achievement Award for a major impact of his decade long contributions to foundations of adaptability in communication and distributed systems. He has been awarded the charter Gold Core Member distinction by the IEEE Computer Society for his distinguished service. He received Outstanding Instructor Awards from the Purdue chapter of the ACM in 1996 and 1998. He has graduated the largest number of Ph.D students in CS department and is active in supporting/mentoring minority students. In 2003, he was inducted in the Purdue's Book of Great Teachers.

He serves on seven editorial boards of international journals. He also serves the IEEE Computer Society on Technical Achievement award and Fellow committees. Professor Bhargava is the founder of the IEEE Symposium on Reliable and Distributed Systems, IEEE conference on Digital Library, and the ACM Conference on Information and Knowledge Management.

For details please see <http://www.cs.purdue.edu/people/bb>

Course Introduction

- This course will deal with the fundamental issues in large distributed systems which are motivated by the computer networking and distribution of processors, and control. The theory, design, implementation, and performance of large systems will be discussed. Concurrency, Consistency, Integrity, Reliability, Privacy, and Security in distributed systems will be included.
- Advanced features of the course include research related to Mobile Data Management, Streaming databases, and Peer to Peer systems.

Slides

- Most of the slides are taken from <http://www.cs.ualberta.ca/~database/ddbook.html> based on the main textbook by Tamer Oszu and Patrick Valduriez
- Some slides have been excluded from various chapters
- Some slides have been updated by Prof. Bhargava
- Some slides have been prepared by Prof. Bhargava
- Some slides are based on research papers that are available at the website as reading materials
- <https://www.cs.purdue.edu/homes/bb/cs542-22Spr/>

Reading materials

□ Textbooks

- Principles of Distributed Database Systems, Prentice Hall, Tamer Özsu and Patrick Valduriez (Maint Text)
- Concurrency Control and Reliability in Distributed Systems, Van Nostrand and Reinhold Publishers, Bharat Bhargava (Ed.), 1987 (Out of Print)
- Transaction Processing: Concepts and Techniques, Morgan Kaufmann, Jim Gray and Andreas Reuter, 1992.

□ Supplemental Textbook

- Principles of Database Systems. Computer Science Press, Rockville, J. D. Ullman, 1982

□ Research papers

- From the reading list on the course web page

<https://www.cs.purdue.edu/homes/bb/cs542-22Spr/>

Assignments and Grading Policy

- Non programming assignments: 4 (every 2-3 weeks: 20% of grade)
- Mid Term and Final Exams: 2 (20% of grade each)
- Project: 30% of grade
- Class contributions: 10% of grade
- $> 94\% = A$, $90-93\% = A-$
- Same scale for B and C grade.
- If you do not do any part of the project or home works or score less than 75%, a C grade is possible.
- To pass the qualifier, a grade of A- in class and 80% in qualifier questions is expected.

Distributed Database Systems

- Computer network (communication system)
- Database systems
- Users (programs, transactions)

Examples:

Distributed INGRES (UC-Berkley)

SDD-1 (Computer Corporation of America)

DB2 and System R* (IBM)

SIRIUS – DELTA (INRIA, France)

RAID (Purdue)

Distributed Database Systems

Computer Networks:

Ethernet

ATM

FDDI

ARPANET

BITNET

Internet2

...

Communications:

UDP/IP

TCP/IP

ISO

User Interaction:

SQL

Transaction

Fundamental References

- Bharat Bhargava (Ed.), *Concurrency Control and Reliability in Distributed Systems*, Van Nostrand and Reinhold Publishers, 1987.
- A. Helal, A. Heddaya, and B. Bhargava, *Replication Techniques in Distributed Systems*, Klumer Academic Publishers, 1996.
- J. Gray and A. Reuter. *Transaction Processing - Concepts and Techniques*. Morgan Kaufmann, 1993.
- M.T. Özsu and P. Valduriez. *Principles of Distributed Database Systems*, 2nd edition. Prentice Hall, 1999.
- S. Ceri and G. Pelagatti. *Distributed Databases - Principles and Systems*. McGraw Hill, 1984.
- D.A. Bell and J.B. Grimson. *Distributed Database Systems*. Addison-Wesley, 1992.

Fundamental References (see Website)

- B. Bhargava, *Building Distributed Database Systems*.
- B. Bhargava and John Riedl, *The Raid Distributed Database System*, IEEE Trans on Software Engineering, 15(6), June 1989.
- B. Bhargava, *Concurrency Control in Database Systems*, IEEE Trans on Knowledge and Data Engineering, 11(1), Jan.-Feb. 1999
- B. Bhargava and John Riedl, *A Model for Adaptable Systems for Transaction Processing*, IEEE Transactions on Knowledge and Data Engineering, 1(4), Dec 1989.
- B. Bhargava and M. Annanalai, *A framework for communication software and measurements for digital library*, Journal of Multimedia systems, 2000.
- B. Bhargava and C. Hua. *A Causal Model for Analyzing Distributed Concurrency Control Algorithms*, IEEE Transactions on Software Engineering, SE-9, 470-486, 1983.
- E. Mafla, and B. Bhargava, *Communication Facilities for Distributed Transaction Processing Systems*, IEEE Computer, 24(8), 1991.
- Y. Zhang and B. Bhargava, *WANCE: Wide area network communication emulation systems*, IEEE workshop on Parallel and Distributed Systems, 1993.
- G. Ding and B. Bhargava, *Peer-to-peer File-sharing over Mobile Ad hoc Networks*, in the First International Workshop on Mobile Peer-to-Peer Computing, Orlando, Florida, March 2004
- M. Hefeeda, A. Habib, B. Botev, D. Xu, B. Bhargava, *PROMISE: Peer-to-Peer Media Streaming Using CollectCast*, In Proc. of ACM Multimedia 2003, 45-54, Berkeley, CA, November 2003.

Fundamental References (cont'd)

- Y. Lu, W. Wang, D. Xu, and B. Bhargava, *Trust-Based Privacy Preservation for Peer-to-peer*, in the 1st NSF/NSA/AFRL workshop on secure knowledge management (SKM), Buffalo, NY, Sep. 2004.
- B. Bhargava, Y. Zhang, and E. Mafla, Evolution of a communication system for distributed transaction processing in RAID, *Computing Systems*, 4(3), 1991.
- E. Pitoura and B. Bhargava, *Data Consistency in Intermittently Connected Distributed Systems*, *IEEE TKDE*, 11(6), 1999.
- E. Pitoura and B. Bhargava, *Maintaining Consistency of Data in Mobile Distributed Environments*, *ICDCS*, 1995.
- A. Zhang, M. Nodine, and B. Bhargava, *Global scheduling for flexible transactions in heterogeneous distributed database systems*, *IEEE TKDE*, 13(3), 2001.
- P. Bernstein and N. Goodman, *Concurrency Control in Distributed Database Systems*, *ACM Computer Survey*, 13(2), 1981.
- P. Bernstein, D. Shipman, and J. Rothnie, *Concurrency control in a system for distributed databases (SDD-1)*, *ACM Transactions on Database Systems*, 5(1), 1980.
- Jim Gray, *The Transaction Concept: Virtues and Limitations*, *VLDB*, 1981.
- H.T. Kung and John T. Robinson, *On Optimistic Methods for Concurrency Control*, *ACM Trans. Database Systems*, 6(2), 1981.

Fundamental References (cont'd)

- C. Papadimitriou, *The serializability of concurrent database updates*, Journal of the ACM, 26(4), 1979.
- D. Skeen, *A Decentralized Termination Protocol*, IEEE Symposium on Reliability in Distributed Software and Database Systems, July 1981.
- D. Skeen, *Nonblocking commit protocols*, ACM SIGMOD, 1981.
- D. Skeen and M Stonebraker, *A Formal Model of Crash Recovery in a Distributed System*, IEEE Trans. Software Eng. 9(3): 219-228, 1983.
- W. W. Chu, *Optimal File Allocation in Multiple Computer System*, IEEE Transaction on Computers, 885-889, October 1969.
- B. Bhargava and L. Lilien, *Private and Trusted Collaborations*, in Proceedings of Secure Knowledge Management (SKM), Amherst, NY, Sep. 2004.
- S. B. Davidson, *Optimism and consistency in partitioned distributed database systems*, ACM Transactions on Database Systems 9(3): 456-481, 1984.
- S. B. Davidson, H. Garcia-Molina, and D. Skeen, *Consistency in Partitioned Networks*, ACM Computer Survey, 17(3): 341-370, 1985.
- B. Bhargava, *Resilient Concurrency Control in Distributed Database Systems*, IEEE Trans. on Reliability, R-31(5): 437-443, 1984.
- Jr. D. Parker, et al., *Detection of Mutual Inconsistency in Distributed Systems*, IEEE Trans. on Software Engineering, SE-9, 1983.

Other References

□ Transaction Management:

- P.A. Bernstein and E. Newcomer. *Principles of Transaction Processing for the Systems Professional*, Morgan Kaufmann, 1997.
- P.A. Bernstein; V. Hadzilacos and N. Goodman. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley, 1987. (out of print)
- M. Buretta. *Data Replication*, Wiley, 1997.
- V. Kumar (ed.). *Performance of Concurrency Control Mechanisms in Centralized Database Systems*, Prentice Hall, 1996.
- V. Kumar and S.H. Son. *Database Recovery*, Kluwer, 1998.
- C.H. Papadimitriou. *The Theory of Concurrency Control*. Computer Science Press, 1986.

Other References

□ Interoperability:

- A.K. Elmagarmid, M. Rusinkiewicz, and A. Sheth (eds). *Management of Heterogeneous and Autonomous Database Systems*, Morgan Kaufmann, 1998.
- A. Bouguettaya, B. Benatallah, and A. Elmagarmid (eds.). *Interconnecting Heterogeneous Information Systems*, Kluwer, 1998.
- J. Siegel (ed.). *CORBA Fundamentals and Programming*, Wiley, 1996.
- K. Brockschmidt. *Inside OLE, 2nd edition*, Microsoft Press, 1995.
- K. Geiger. *Inside ODBC*, Microsoft Press, 1995.

Other References

□ Data Warehousing

- There are **many** books. A small sample:
 - W. Inmon. *Building the Data Warehouse*. John Wiley and Sons, 1992.
 - A. Berson and S.J. Smith. *Data Warehousing, Data Mining, and OLAP*. McGraw Hill, 1997.
- S. Chaudri and U. Dayal. Overview of Data Warehousing and OLAP Technology. *ACM SIGMOD Record*, March 1997, 26(1), pp. 65-74.
- *IEEE Q. Bull. Data Engineering, Special Issue on Materialised Views on Data Warehousing*, June 1995, 18(2).

Other References

- Parallel Database Servers:
 - P. Valduriez (ed). *Data Management and Parallel Processing*. Chapman and Hall, 1992.
 - M. Abdelguerfi and K-F. Wong. *Parallel Database Techniques*. IEEE Computer Society Press, 1988.
 - P. Valduriez. “Parallel Database Systems: Open Problems and New Issues,” *Parallel and Distributed Databases*, April 1993, 1(2): 137–165.
 - D. DeWitt and J. Gray. Parallel Database Systems: The Future of High-Performance Database Systems. *Communications of ACM*, June 1992, 35(6), 1992.

Other References

- Distributed Object Management:
 - E. Bertino and L. Martino. *Object-Oriented Database Systems*. Addison-Wesley, 1993.
 - A. Kemper and G. Moerkotte. *Object-Oriented Database Management*. Prentice-Hall, 1994.
 - A. Dogac, M.T. Özsu, A. Biliris, T. Sellis (ed.). *Advances in Object-Oriented Database Systems*. Springer-Verlag, 1994.
 - M.T. Özsu, U. Dayal and P. Valduriez (eds.). *Distributed Object Management*. Morgan Kaufman, 1994.
 - W. Kim (ed). *Modern Database Management – Object-Oriented and Multidatabase Technologies*. Addison-Wesley/ACM Press, 1994.
 - S. Zdonik and D. Maier (eds.). *Readings in Object-Oriented Database Systems*. Morgan Kaufmann, 1990.
 - R.G.G. Cattell. *Object Data Management*, 2nd edition. Addison-Wesley, 1994.

Other References

□ Mobile Databases

- A. Helal et al. *Any Time, Anywhere Computing*, Kluwer, 1999.
- T. Imielinski and H. Korth. *Mobile Computing*. Kluwer Publishers, 1996.
- E. Pitoura and G. Samaras. *Data Management for Mobile Computing*. Kluwer Publishers, 1998.
- T. Imielinski and B.R. Badrinath. Data Management Issues in Mobile Computing. *Communications of ACM*, October 1994, 37(10):18-28.
- M. H. Dunham and A. Helal. Mobile Computing and Databases: Anything New? *ACM SIGMOD Record*, December 1995, 24(4): 5-9.
- G. H. Forman and J. Zahorjan. The Challenges of Mobile Computing, *Computer*, April 1994, 27(4):38-47.

Other References

- Web Data Management
 - S. Abiteboul, P. Buneman, D. Suciu. *Data on the Web*, Morgan Kaufmann, 2000.
 - D. Florescu, A. Levy, and A. Mendelzon, *Database Technoques for the World Wide Web: A Survey*, *ACM SIGMOD Record*, 27(3): 59-74, 1998.
 - S. Bhowmick, S. Madria, and W. K. Ng, *Web Data Management: A Warehouse Approach*, Springer, 2003.