



Principles of Distributed Database Systems (2nd Edition)

By M. Tamer Ozsu, Patrick Valduriez

Download now

Read Online 

Principles of Distributed Database Systems (2nd Edition) By M. Tamer Ozsu, Patrick Valduriez

Provides a comprehensive treatment of technical problems of distributed database systems from a holistic viewpoint. **KEYTOPICS:** Explores the development of distributed database management systems—focusing on concepts and technical issues.

 [Download Principles of Distributed Database Systems \(2nd Ed ...pdf](#)

 [Read Online Principles of Distributed Database Systems \(2nd ...pdf](#)

Principles of Distributed Database Systems (2nd Edition)

By *M. Tamer Ozsu, Patrick Valduriez*

Principles of Distributed Database Systems (2nd Edition) By *M. Tamer Ozsu, Patrick Valduriez*

Provides a comprehensive treatment of technical problems of distributed database systems from a holistic viewpoint. **KEYTOPICS:** Explores the development of distributed database management systems—focusing on concepts and technical issues.

Principles of Distributed Database Systems (2nd Edition) By *M. Tamer Ozsu, Patrick Valduriez*
Bibliography

- Sales Rank: #2625139 in Books
- Brand: Brand: Prentice Hall
- Published on: 1999-01-29
- Original language: English
- Number of items: 1
- Dimensions: 9.40" h x 1.42" w x 7.20" l, 2.55 pounds
- Binding: Paperback
- 666 pages

 [Download Principles of Distributed Database Systems \(2nd Ed ...pdf](#)

 [Read Online Principles of Distributed Database Systems \(2nd ...pdf](#)

Editorial Review

From the Inside Flap

PREFACE TO THE SECOND EDITION Many things have changed since the publication of the first edition of this book in 1991. At the time, we reported projections that, by 1998, centralized database managers (DBMSs) would be an antique curiosity" and most organizations would move towards distributed database managers. Distribution was slowly starting and "client/server" had just started to enter our daily jargon. These systems were generally multiple client/single server systems in which the distribution was mostly in terms of functionality, not data. If multiple servers were used, clients were responsible for managing the connections to these servers. Thus, transparency of access was not widely supported, and each client had to "know" the location of the required data. The distribution of data among multiple servers was very primitive; systems did not support fragmentation or replication of data. Systems of the time were homogeneous" in that each system could manage only data that were stored in its own database, with no linkage to other repositories. Things have changed dramatically since then. Many vendors are much closer to achieving true distribution in their development cycle. Client/server systems remain the preferred solution in many cases, but they are much more sophisticated. For example, today's client/server systems provide significant transparency in accessing data from multiple servers, support distributed transactions to facilitate transparency, and execute queries over (horizontally) fragmented data. Further, new systems implement both synchronous and asynchronous replication protocols, and many vendors have introduced gateways to access other databases. In addition, significant achievements have taken place in the development and deployment of parallel database servers. Object database managers have entered the marketplace and have found a niche market in some classes of applications which are inherently distributed. In parallel with these developments in the database system front, there have been phenomenal changes in the computer networking infrastructure that supports these systems. The relatively slow (10Mbit/sec) Ethernet has been replaced as the de facto local area network standard by much faster networks (FDDI or switched Ethernet) operating at around 100Mbit/sec, and broadband networks (particularly the ATM technology) have been deployed for both local area and wide area networking. These networks, coupled with very low overhead networking protocols, such as SCI, reduce the differences between local area and wide area networks (other than latency considerations) and potentially eliminate the network as the major performance bottleneck. This, in turn, requires us to review our system development assumptions and performance tuning criteria. Use of the Internet which is basically a heterogeneous network with links of varying capacities and capabilities has exploded. There is clearly a technology push/application pull in effect with respect to distributed DBMS development: new applications are requiring changes in DBMS capabilities, and new technological developments are making these changes possible. With these developments, it was time to prepare a revised second edition of the book. In the process, we have retained the fundamental characteristics and key features of the book as outlined in the Preface to the first edition. However, the material has been heavily edited. Every chapter has been revised some in fundamental ways, others more superficially. The major changes are the following: 1. The query processing/optimization chapters (Chapters 7{9) have been revised to focus on the techniques employed in commercial systems. New algorithms, such as randomized search strategies, are now included. 2. The transaction management chapters (Chapters 10{12) now include material on advanced transaction models and work flows. 3. Chapter 13, which focused on the relationship of distributed DBMSs and distributed operating systems, has been dropped and some of the material is incorporated into the relevant chapters. 4. The first edition contained a chapter (Chapter 15) which discussed current issues at the time parallel DBMSs, distributed knowledge-base systems (mainly deductive DBMSs), and distributed object DBMSs. In the intervening years, two of these topics have matured and become major forces in their own rights, while the third (deductive databases) has not achieved the same prominence. In this edition, we devote

full chapters to parallel DBMSs (Chapter 13) and distributed object DBMSs (Chapter 14), and have dropped deductive DBMSs. 5. Following the same approach, we introduce a new chapter devoted to current issues (Chapter 16). This chapter now includes sections on data warehousing (from a distributed data management perspective), World Wide Web and databases, push-based technologies, and mobile DBMSs. 6. The chapter on multidatabase systems (Chapter 15 in the current edition) has been revised to include a discussion of general interoperability issues and distributed object platforms such as OMA/CORBA and DCOM/OLE. We are quite satisfied with the result, which represents a compromise between our desire to address new and emerging issues, and maintain the main characteristic of the book in addressing the principles of distributed data management. Certain chapters, in particular Chapters 15 and 16, require further depth, but those will be topics of future editions. The guide to reading the book, introduced in the Preface to the first edition, is still valid in general terms. However, we now discuss, in Chapter 3, the relationship between distributed DBMSs and the new networking technologies. Thus, this chapter no longer serves simply as background and should be read (at least the relevant sections) following Chapter 1. We have set up a Web site to communicate with our readers. The site is at cs.ualberta.ca/~database/distdb.html. This site contains presentation slides that accompany the book as well as other information regarding the book's use as a textbook. Many colleagues have helped with the revisions. Maggie Dunham and Nandid Soparkar provided detailed and early comments on the overall structure and content of the book. Maggie also provided input for the mobile database management section (Section 16.4). Ioannis Nikolaidis helped immensely with the revisions to Chapter 3 he made us rewrite that chapter three times. Jari Veijaleinen provided many exercises which have been incorporated into this edition. Esther Pacitti provided input for replication protocols. Peter Trianta_llou provided material on this topic as well. Alexander Thomasian's input for performance evaluation work was invaluable, as was Elliot Moss's critical review of the nested transaction discussion in the transaction processing chapters. Mukesh Singhal advised us of the new advances in distributed deadlock management. Luc Bouganim contributed significantly to Chapter 13 on parallel DBMSs. Ken Barker and Kamalakar Karlapalem provided the material that formed the basis of distributed object database design in Chapter 14. Kaladhar Voruganti wrote the first draft of the architectural and system issues sections of Chapter 14. Randal Peters read Chapter 14 and forced us to revise many parts of it. The distributed garbage collection section of that chapter is based on a draft provided by Laurent Amsaleg and Michael Franklin. Amit Sheth provided input on the revised outline for Chapter 15. Asuman Dogac read the complete chapter and provided feedback. Mokrane Bouzeghoub and Eric Simon helped on the data warehouse section of Chapter 16. Dana Florescu, Alon Levy, Ioana Manolescu and Anthony Tomasic provided input for research prototypes in the section on Web and databases in Chapter 16. The material in push-based technologies section was reviewed (a number of times) by Stan Zdonik and Mike Franklin. Both of them also provided significant feedback about the characterization of data delivery alternatives. We are indebted to all of them, as well as to those who helped with the original edition of the book and whom we cite in the Preface to the First Edition. Many other colleagues have asked questions and provided suggestions over the years; unfortunately, we have not kept their names. Our thanks to everyone who has provided input. We look forward to receiving more suggestions on the second edition. We have had very good luck with our editors at Prentice Hall. Our current editor, Alan Apt, and our development editor, Sondra Chavez, have been tremendously helpful in both pushing us forward and providing the necessary institutional support. Our production editors, Ed DeFelippis and Irwin Zucker, have managed the production process so that the production of earlier chapters could proceed in parallel with our writing of the later chapters. This allowed the revised edition to be ready within one year. Stephen Lee, as our copy editor, made the entire text significantly more readable. Anne Nield helped us in many ways editing chapters, correcting the text and keeping us organized. Paul Iglinski wrote a number of scripts that helped immensely with cleaning up the b

From the Back Cover

In the Second Edition of this best-selling distributed database systems text, the authors address new and emerging issues in the field while maintaining the key features and characteristics of the First Edition. The text has been revised and updated to reflect changes in the field. This comprehensive text focuses on concepts and technical issues while exploring the development of distributed database management systems (DBMS). Principles of Distributed Database Systems presents distributed database systems within the framework of distributed data processing in general, rather than as a problem in isolation.

NEW TO THIS EDITION

- The relationship of distributed DBMSs with the new networking technologies is discussed.
- The query processing/optimization chapters now focus on techniques employed in commercial systems and include new algorithms such as randomized search strategies.
- Discussion of advanced transaction models and workflows has been added to the transaction management chapters.
- Full chapters are devoted to parallel DBMSs and distributed object DBMSs.
- Current issues are discussed in a new chapter, including sections on data warehousing, world wide web and databases, push-based technologies, and mobile DBMSs.
- General interoperability issues and distributed object platforms such as OMA/CORBA and DCOM/OLE have been added to the multidatabase systems chapter.
- The authors' web site contains presentation slides, helpful information for instructors, and direct communication with the authors. The url is <http://www.cs.ualberta.ca/~database/distdb.html>.

Excerpt. © Reprinted by permission. All rights reserved.

PREFACE TO THE SECOND EDITION

Many things have changed since the publication of the first edition of this book in 1991. At the time, we reported projections that, by 1998, centralized database managers (DBMSs) would be an antique curiosity" and most organizations would move towards distributed database managers. Distribution was slowly starting and "client/server" had just started to enter our daily jargon. These systems were generally multiple client/single server systems in which the distribution was mostly in terms of functionality, not data. If multiple servers were used, clients were responsible for managing the connections to these servers. Thus, transparency of access was not widely supported, and each client had to "know" the location of the required data. The distribution of data among multiple servers was very primitive; systems did not support fragmentation or replication of data. Systems of the time were homogeneous" in that each system could manage only data that were stored in its own database, with no linkage to other repositories.

Things have changed dramatically since then. Many vendors are much closer to achieving true distribution in their development cycle. Client/server systems remain the preferred solution in many cases, but they are much more sophisticated. For example, today's client/server systems provide significant transparency in accessing data from multiple servers, support distributed transactions to facilitate transparency, and execute queries over (horizontally) fragmented data. Further, new systems implement both synchronous and asynchronous replication protocols, and many vendors have introduced gateways to access other databases. In addition, significant achievements have taken place in the development and deployment of parallel database servers. Object database managers have entered the marketplace and have found a niche market in some classes of applications which are inherently distributed.

In parallel with these developments in the database system front, there have been phenomenal changes in the computer networking infrastructure that supports these systems. The relatively slow (10Mbit/sec) Ethernet has been replaced as the de facto local area network standard by much faster networks (FDDI or switched Ethernet) operating at around 100Mbit/sec, and broadband networks (particularly the ATM technology) have

been deployed for both local area and wide area networking. These networks, coupled with very low overhead networking protocols, such as SCI, reduce the differences between local area and wide area networks (other than latency considerations) and potentially eliminate the network as the major performance bottleneck. This, in turn, requires us to review our system development assumptions and performance tuning criteria. Use of the Internet which is basically a heterogeneous network with links of varying capacities and capabilities has exploded.

There is clearly a technology push/application pull in effect with respect to distributed DBMS development: new applications are requiring changes in DBMS capabilities, and new technological developments are making these changes possible. With these developments, it was time to prepare a revised second edition of the book. In the process, we have retained the fundamental characteristics and key features of the book as outlined in the Preface to the first edition. However, the material has been heavily edited. Every chapter has been revised some in fundamental ways, others more superficially. The major changes are the following:

1. The query processing/optimization chapters (Chapters 7{9) have been revised to focus on the techniques employed in commercial systems. New algorithms, such as randomized search strategies, are now included.
2. The transaction management chapters (Chapters 10{12) now include material on advanced transaction models and work flows.
3. Chapter 13, which focused on the relationship of distributed DBMSs and distributed operating systems, has been dropped and some of the material is incorporated into the relevant chapters.
4. The first edition contained a chapter (Chapter 15) which discussed current issues at the time parallel DBMSs, distributed knowledge-base systems (mainly deductive DBMSs), and distributed object DBMSs. In the intervening years, two of these topics have matured and become major forces in their own rights, while the third (deductive databases) has not achieved the same prominence. In this edition, we devote full chapters to parallel DBMSs (Chapter 13) and distributed object DBMSs (Chapter 14), and have dropped deductive DBMSs.
5. Following the same approach, we introduce a new chapter devoted to current issues (Chapter 16). This chapter now includes sections on data warehousing (from a distributed data management perspective), World Wide Web and databases, push-based technologies, and mobile DBMSs.
6. The chapter on multidatabase systems (Chapter 15 in the current edition) has been revised to include a discussion of general interoperability issues and distributed object platforms such as OMA/CORBA and DCOM/OLE.

We are quite satisfied with the result, which represents a compromise between our desire to address new and emerging issues, and maintain the main characteristic of the book in addressing the principles of distributed data management. Certain chapters, in particular Chapters 15 and 16, require further depth, but those will be topics of future editions.

The guide to reading the book, introduced in the Preface to the first edition, is still valid in general terms. However, we now discuss, in Chapter 3, the relationship between distributed DBMSs and the new networking technologies. Thus, this chapter no longer serves simply as background and should be read (at least the relevant sections) following Chapter 1.

We have set up a Web site to communicate with our readers. The site is at

<http://www.cs.ualberta.ca/~database/distdb.html>. This site contains presentation slides that accompany the book as well as other information regarding the book's use as a textbook.

Many colleagues have helped with the revisions. Maggie Dunham and Nandid Soparkar provided detailed and early comments on the overall structure and content of the book. Maggie also provided input for the mobile database management section (Section 16.4). Ioannis Nikolaidis helped immensely with the revisions to Chapter 3 he made us rewrite that chapter three times. Jari Veijaleinen provided many exercises which have been incorporated into this edition. Esther Pacitti provided input for replication protocols. Peter Trianta_llou provided material on this topic as well. Alexander Thomasian's input for performance evaluation work was invaluable, as was Elliot Moss's critical review of the nested transaction discussion in the transaction processing chapters. Mukesh Singhal advised us of the new advances in distributed deadlock management. Luc Bouganim contributed significantly to Chapter 13 on parallel DBMSs. Ken Barker and Kamalakar Karlappalem provided the material that formed the basis of distributed object database design in Chapter 14. Kaladhar Voruganti wrote the first draft of the architectural and system issues sections of Chapter 14. Randal Peters read Chapter 14 and forced us to revise many parts of it. The distributed garbage collection section of that chapter is based on a draft provided by Laurent Amsaleg and Michael Franklin. Amit Sheth provided input on the revised outline for Chapter 15. Asuman Dogac read the complete chapter and provided feedback. Mokrane Bouzeghoub and Eric Simon helped on the data warehouse section of Chapter 16. Dana Florescu, Alon Levy, Ioana Manolescu and Anthony Tomasic provided input for research prototypes in the section on Web and databases in Chapter 16. The material in push-based technologies section was reviewed (a number of times) by Stan Zdonik and Mike Franklin. Both of them also provided significant feedback about the characterization of data delivery alternatives. We are indebted to all of them, as well as to those who helped with the original edition of the book and whom we cite in the Preface to the First Edition. Many other colleagues have asked questions and provided suggestions over the years; unfortunately, we have not kept their names. Our thanks to everyone who has provided input. We look forward to receiving more suggestions on the second edition.

We have had very good luck with our editors at Prentice Hall. Our current editor, Alan Apt, and our development editor, Sondra Chavez, have been tremendously helpful in both pushing us forward and providing the necessary institutional support. Our production editors, Ed DeFelippis and Irwin Zucker, have managed the production process so that the production of earlier chapters could proceed in parallel with our writing of the later chapters. This allowed the revised edition to be ready within one year. Stephen Lee, as our copy editor, made the entire text significantly more readable. Anne Nield helped us in many ways editing chapters, correcting the text and keeping us organized. Paul Iglinski wrote a number of scripts that helped immensely with cleaning up the bibliography. We thank them all.

M. Tamer Ozsu
Patrick Valduriez

Users Review

From reader reviews:

David Crockett:

Why don't make it to become your habit? Right now, try to ready your time to do the important take action, like looking for your favorite guide and reading a guide. Beside you can solve your short lived problem; you can add your knowledge by the publication entitled Principles of Distributed Database Systems (2nd Edition). Try to stumble through book Principles of Distributed Database Systems (2nd Edition) as your close friend. It means that it can be your friend when you feel alone and beside that course make you

smarter than in the past. Yeah, it is very fortunate for you personally. The book makes you far more confidence because you can know everything by the book. So, let us make new experience in addition to knowledge with this book.

Patricia Gross:

As people who live in typically the modest era should be change about what going on or facts even knowledge to make these people keep up with the era that is certainly always change and move forward. Some of you maybe will update themselves by examining books. It is a good choice for you personally but the problems coming to you actually is you don't know what kind you should start with. This Principles of Distributed Database Systems (2nd Edition) is our recommendation to make you keep up with the world. Why, since this book serves what you want and want in this era.

William Reynolds:

The guide with title Principles of Distributed Database Systems (2nd Edition) has a lot of information that you can learn it. You can get a lot of profit after read this book. This kind of book exist new expertise the information that exist in this book represented the condition of the world at this point. That is important to you to be aware of how the improvement of the world. This kind of book will bring you inside new era of the syndication. You can read the e-book in your smart phone, so you can read the idea anywhere you want.

Shawn Hernandez:

Playing with family within a park, coming to see the marine world or hanging out with pals is thing that usually you may have done when you have spare time, after that why you don't try factor that really opposite from that. 1 activity that make you not feeling tired but still relaxing, trilling like on roller coaster you have been ride on and with addition info. Even you love Principles of Distributed Database Systems (2nd Edition), it is possible to enjoy both. It is excellent combination right, you still would like to miss it? What kind of hangout type is it? Oh can happen its mind hangout folks. What? Still don't buy it, oh come on its identified as reading friends.

Download and Read Online Principles of Distributed Database Systems (2nd Edition) By M. Tamer Ozsu, Patrick Valduriez #VFDLT0XAO41

Read Principles of Distributed Database Systems (2nd Edition) By M. Tamer Ozsu, Patrick Valduriez for online ebook

Principles of Distributed Database Systems (2nd Edition) By M. Tamer Ozsu, Patrick Valduriez Free PDF d0wnl0ad, audio books, books to read, good books to read, cheap books, good books, online books, books online, book reviews epub, read books online, books to read online, online library, greatbooks to read, PDF best books to read, top books to read Principles of Distributed Database Systems (2nd Edition) By M. Tamer Ozsu, Patrick Valduriez books to read online.

Online Principles of Distributed Database Systems (2nd Edition) By M. Tamer Ozsu, Patrick Valduriez ebook PDF download

Principles of Distributed Database Systems (2nd Edition) By M. Tamer Ozsu, Patrick Valduriez Doc

Principles of Distributed Database Systems (2nd Edition) By M. Tamer Ozsu, Patrick Valduriez Mobipocket

Principles of Distributed Database Systems (2nd Edition) By M. Tamer Ozsu, Patrick Valduriez EPub

VFDLT0XAO41: Principles of Distributed Database Systems (2nd Edition) By M. Tamer Ozsu, Patrick Valduriez