Principles Of Database Query Processing For Advanced Applications The Morgan Kaufmann Series In Data Management Systems

Linked Data ManagementProceedingsProceedings of the 1986 International Conference on Parallel ProcessingProbabilistic DatabasesDistributed Sensor NetworksInstance-level Integration, Query Processing and Optimization in Federated Database SystemsPrinciples of Multimedia Database SystemsMoving Objects DatabasesProceedings of the ACM Symposium on Principles of Database SystemsAdaptive Query ProcessingDatabasePrinciples of Data IntegrationSpatial DatabasesPrinciples of Database Systems with Internet and Java ApplicationsPrinciples of Database ManagementDeep Web Query Interface Understanding and IntegrationPricing Policies and Query Processing in the Mariposa Agoric Distributed Database Management SystemPrinciples of Distributed Database SystemsHigh-Performance Parallel Database Processing and Grid DatabasesProceedings of the Second ACM SIGACT-SIGMOD Symposium on Principles of Database SystemsReadings in Database SystemsQuery Processing on Probabilistic DataAdvanced Principles for Improving Database Design, Systems Modeling, and Software DevelopmentAdvanced Data ManagementPrinciples of Visual Information RetrievalQuery Processing in

Database SystemsAdvanced Metasearch Engine TechnologyProceedings of the ACM SIGACT-SIGMOD Symposium on Principles of Database SystemsPrinciples of Database Query Processing for Advanced ApplicationsPrinciples of Transaction ProcessingArchitecture of a Database SystemAdvances in DatabasesEncyclopedia of Database SystemsDatabase SystemsProceedings of the Fourth ACM SIGACT-SIGMOD Symposium on Principles of Database Systems, March 25-27, 1985, Portland, OregonTwenty-fifth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database SystemsProceedingsData Management and Query Processing in Semantic Web DatabasesComplex Query Processing in Multiprocessor Database MachinesConstraints and Databases

Linked Data Management

A thorough presentation of query processing techniques in a broad range of database systems for advanced applications. Provides the most effective query processing techniques and ways to optimize the information retrieval process. Intended for database systems designers creating advanced applications.

Proceedings

The vision of researchers to create smart environments through the deployment of thousands of sensors, each with a short range wireless

communications channel and capable of detecting ambient conditions such as temperature, movement, sound, light, or the presence of certain objects is becoming a reality. With the emergence of high-speed networks an

Proceedings of the 1986 International Conference on Parallel Processing

Probabilistic Databases

Distributed Sensor Networks

Probabilistic databases are databases where the value of some attributes or the presence of some records are uncertain and known only with some probability. Applications in many areas such as information extraction, RFID and scientific data management, data cleaning, data integration, and financial risk assessment produce large volumes of uncertain data, which are best modeled and processed by a probabilistic database. This book presents the state of the art in representation formalisms and query processing techniques for probabilistic data. It starts by discussing the basic principles for representing large probabilistic databases, by decomposing them into tuple-independent tables, block-independent-disjoint tables, or U-databases. Then it discusses two classes of techniques for query evaluation on probabilistic databases. In extensional query evaluation, the entire

probabilistic inference can be pushed into the database engine and, therefore, processed as effectively as the evaluation of standard SQL queries. The relational queries that can be evaluated this way are called safe queries. In intensional query evaluation, the probabilistic inference is performed over a propositional formula called lineage expression: every relational query can be evaluated this way, but the data complexity dramatically depends on the query being evaluated, and can be #P-hard. The book also discusses some advanced topics in probabilistic data management such as top-k query processing, sequential probabilistic databases, indexing and materialized views, and Monte Carlo databases. Table of Contents: Overview / Data and Query Model / The Query Evaluation Problem / Extensional Query Evaluation / Intensional Query Evaluation / Advanced Techniques

Instance-level Integration, Query Processing and Optimization in Federated Database Systems

Principles of Transaction Processing is a comprehensive guide to developing applications, designing systems, and evaluating engineering products. The book provides detailed discussions of the internal workings of transaction processing systems, and it discusses how these systems work and how best to utilize them. It covers the architecture of Web Application Servers and transactional communication paradigms. The book is divided into 11 chapters, which cover the following:

Overview of transaction processing application and system structure Software abstractions found in transaction processing systems Architecture of multitier applications and the functions of transactional middleware and database servers Queued transaction processing and its internals, with IBM's Websphere MQ and Oracle's Stream AQ as examples Business process management and its mechanisms Description of the two-phase locking function, B-tree locking and multigranularity locking used in SQL database systems and nested transaction locking System recovery and its failures Two-phase commit protocol Comparison between the tradeoffs of replicating servers versus replication resources Transactional middleware products and standards Future trends, such as cloud computing platforms, composing scalable systems using distributed computing components, the use of flash storage to replace disks and data streams from sensor devices as a source of transaction requests. The text meets the needs of systems professionals, such as IT application programmers who construct TP applications, application analysts, and product developers. The book will also be invaluable to students and novices in application programming. Complete revision of the classic "non mathematical" transaction processing reference for systems professionals. Updated to focus on the needs of transaction processing via the Internet-- the main focus of business data processing investments, via web application servers, SOA, and important new TP standards. Retains the practical, non-mathematical, but thorough conceptual basis of the first edition.

Principles of Multimedia Database Systems

The authors explore and explain current techniques for handling the specialised data that describes geographical phenomena in a study that will be of great value to computer scientists and geographers working with spatial databases.

Moving Objects Databases

The proceedings of the April 1991 Conference in Kobe, Japan comprise papers and panels on topics in objected-oriented database systems, distributed database systems, design and human interfaces, data engineering techniques, AI and knowledge-based systems, access methods and file structures, parallel query processing, deductive and extensive databases, distributed database control, heterogeneous or federated or multidatabase systems, query languages and processing, genomic databases, performance evaluation, applications and application systems, DE technology in Japan, database management, multimedia database systems, object-oriented environments, cooperating knowledge-based systems, and AI and databases. No index. Annotation copyrighted by Book News, Inc., Portland, OR.

Proceedings of the ACM Symposium on Principles of Database Systems

Adaptive Query Processing

The Semantic Web, which is intended to establish a machine-understandable Web, is currently changing from being an emerging trend to a technology used in complex real-world applications. A number of standards and techniques have been developed by the World Wide Web Consortium (W3C), e.g., the Resource Description Framework (RDF), which provides a general method for conceptual descriptions for Web resources, and SPARQL, an RDF querying language. Recent examples of large RDF data with billions of facts include the UniProt comprehensive catalog of protein sequence, function and annotation data, the RDF data extracted from Wikipedia, and Princeton University's WordNet. Clearly, querying performance has become a key issue for Semantic Web applications. In his book, Groppe details various aspects of high-performance Semantic Web data management and query processing. His presentation fills the gap between Semantic Web and database books, which either fail to take into account the performance issues of large-scale data management or fail to exploit the special properties of Semantic Web data models and queries. After a general introduction to the relevant Semantic Web standards, he presents specialized indexing and sorting algorithms, adapted approaches for logical and physical query optimization, optimization possibilities when using the parallel database technologies of today's multicore processors, and visual and embedded query languages. Groppe primarily targets researchers, students, and developers of large-scale Semantic Web applications. On the complementary book webpage readers will

find additional material, such as an online demonstration of a query engine, and exercises, and their solutions, that challenge their comprehension of the topics presented.

Database

Linked Data Management presents techniques for querying and managing Linked Data that is available on today's Web. The book shows how the abundance of Linked Data can serve as fertile ground for research and commercial applications. The text focuses on aspects of managing large-scale collections of Linked Data. It offers a detailed introduction to Linked Data and related standards, including the main principles distinguishing Linked Data from standard database technology. Chapters also describe how to generate links between datasets and explain the overall architecture of data integration systems based on Linked Data. A large part of the text is devoted to query processing in different setups. After presenting methods to publish relational data as Linked Data and efficient centralized processing, the book explores lookup-based, distributed, and parallel solutions. It then addresses advanced topics, such as reasoning, and discusses work related to read-write Linked Data for system interoperation. Despite the publication of many papers since Tim Berners-Lee developed the Linked Data principles in 2006, the field lacks a comprehensive, unified overview of the state of the art. Suitable for both researchers and practitioners, this book provides a thorough, consolidated account

of the new data publishing and data integration paradigm. While the book covers query processing extensively, the Linked Data abstraction furnishes more than a mechanism for collecting, integrating, and querying data from the open Web—the Linked Data technology stack also allows for controlled, sophisticated applications deployed in an enterprise environment.

Principles of Data Integration

Among the search tools currently on the Web, search engines are the most well known thanks to the popularity of major search engines such as Google and Yahoo!. While extremely successful, these major search engines do have serious limitations. This book introduces large-scale metasearch engine technology, which has the potential to overcome the limitations of the major search engines. Essentially, a metasearch engine is a search system that supports unified access to multiple existing search engines by passing the queries it receives to its component search engines and aggregating the returned results into a single ranked list. A large-scale metasearch engine has thousands or more component search engines. While metasearch engines were initially motivated by their ability to combine the search coverage of multiple search engines, there are also other benefits such as the potential to obtain better and fresher results and to reach the Deep Web. The following major components of large-scale metasearch engines will be discussed in detail in this book: search engine selection, search engine incorporation, and result

merging. Highly scalable and automated solutions for these components are emphasized. The authors make a strong case for the viability of the large-scale metasearch engine technology as a competitive technology for Web search. Table of Contents: Introduction / Metasearch Engine Architecture / Search Engine Selection / Search Engine Incorporation / Result Merging / Summary and Future Research

Spatial Databases

Principles of Database Systems with Internet and Java Applications

Principles of Database Management

How do you approach answering queries when your data is stored in multiple databases that were designed independently by different people? This is first comprehensive book on data integration and is written by three of the most respected experts in the field. This book provides an extensive introduction to the theory and concepts underlying today's data integration techniques, with detailed, instruction for their application using concrete examples throughout to explain the concepts. Data integration is the problem of answering queries that span multiple data sources (e.g., databases, web pages). Data integration problems surface in multiple contexts, including enterprise information integration, query processing on the Web, coordination between

government agencies and collaboration between scientists. In some cases, data integration is the key bottleneck to making progress in a field. The authors provide a working knowledge of data integration concepts and techniques, giving you the tools you need to develop a complete and concise package of algorithms and applications. Offers a range of data integration solutions enabling you to focus on what is most relevant to the problem at hand Enables you to build your own algorithms and implement your own data integration applications

Deep Web Query Interface Understanding and Integration

The latest techniques and principles of parallel and grid database processing The growth in grid databases, coupled with the utility of parallel query processing, presents an important opportunity to understand and utilize high-performance parallel database processing within a major database management system (DBMS). This important new book provides readers with a fundamental understanding of parallelism in data-intensive applications, and demonstrates how to develop faster capabilities to support them. It presents a balanced treatment of the theoretical and practical aspects of high-performance databases to demonstrate how parallel query is executed in a DBMS, including concepts, algorithms, analytical models, and grid transactions. High-Performance Parallel Database Processing and Grid Databases serves as a valuable resource for researchers working in parallel databases

and for practitioners interested in building a high-performance database. It is also a much-needed, self-contained textbook for database courses at the advanced undergraduate and graduate levels.

Pricing Policies and Query Processing in the Mariposa Agoric Distributed Database Management System

Principles of Distributed Database Systems

Adaptive Query Processing surveys the fundamental issues, techniques, costs, and benefits of adaptive query processing. It begins with a broad overview of the field, identifying the dimensions of adaptive techniques. It then looks at the spectrum of approaches available to adapt query execution at runtime - primarily in a non-streaming context. The emphasis is on simplifying and abstracting the key concepts of each technique, rather than reproducing the full details available in the papers. The authors identify the strengths and limitations of the different techniques, demonstrate when they are most useful, and suggest possible avenues of future research. Adaptive Query Processing serves as a valuable reference for students of databases, providing a thorough survey of the area. Database researchers will benefit from a more complete point of view, including a number of approaches which they may not have focused on within the scope of their own research.

High-Performance Parallel Database Processing and Grid Databases

Advanced data management has always been at the core of efficient database and information systems. Recent trends like big data and cloud computing have aggravated the need for sophisticated and flexible data storage and processing solutions. This book provides a comprehensive coverage of the principles of data management developed in the last decades with a focus on data structures and query languages. It treats a wealth of different data models and surveys the foundations of structuring, processing, storing and querying data according these models. Starting off with the topic of database design, it further discusses weaknesses of the relational data model, and then proceeds to convey the basics of graph data, tree-structured XML data, key-value pairs and nested, semi-structured JSON data, columnar and record-oriented data as well as object-oriented data. The final chapters round the book off with an analysis of fragmentation, replication and consistency strategies for data management in distributed databases as well as recommendations for handling polyglot persistence in multi-model databases and multi-database architectures. While primarily geared towards students of Master-level courses in Computer Science and related areas, this book may also be of benefit to practitioners looking for a reference book on data modeling and query processing. It provides both theoretical depth and a concise treatment of open source technologies currently on the market.

Proceedings of the Second ACM SIGACT-SIGMOD Symposium on Principles of Database Systems

Readings in Database Systems

Abstract: "High performance multiprocessor database machines have been made feasible with the advent of cheap, powerful microprocessors and large main memories. However, exploiting these platforms to support high speed complex query processing has lagged behind the hardware technology. The thrust of this dissertation has concentrated on developing strategies for efficiently processing join queries consisting of on the order of 10 joins in a parallel database machine with hundreds of processors. Although the algorithms were developed with a shared-nothing architecture in mind, the algorithms can be applied to shared-memory systems with little modification.

Query Processing on Probabilistic Data

Constraints and Databases contains seven contributions on the rapidly evolving research area of constraints and databases. This collection of original research articles has been compiled as a tribute to Paris C. Kanellakis, one of the pioneers in the field. Constraints have long been used for maintaining the integrity of databases. More recently, constraint databases have emerged where databases store and manipulate data in the form of constraints. The

generality of constraint databases makes them highly attractive for many applications. Constraints provide a uniform mechanism for describing heterogenous data, and advanced constraint solving methods can be used for efficient manipulation of constraint data. The articles included in this book cover the range of topics involving constraints and databases; join algorithms, evaluation methods, applications (e.g. data mining) and implementations of constraint databases, as well as more traditional topics such as integrity constraint maintenance. Constraints and Databases is an edited volume of original research comprising invited contributions by leading researchers.

Advanced Principles for Improving Database Design, Systems Modeling, and Software Development

First uniform treatment of moving objects databases, the technology that supports GPS and RFID data analysis.

Advanced Data Management

Principles of Visual Information Retrieval

This volume contains the proceedings of the eleventh British National Conference on Databases, held at Keele University, England. A dominant themein the volume is the provision of the means to enhance the capabilities of databases to handle information that has a rich semantic structure. A major research

question is how to achieve such a semantic scale-up without sacrificing performance. There are currently two main paradigms within which it is possible to propose answers to this question, deduction-oriented and object-oriented. Both paradigms are well represented in this collection, with the balance in the direction of the deductive approach, which is followed by both the invited papers, by Michael Freeston from the European Computer-Industry Research Centre in Munich and Carlo Zaniolo from the University of California at Los Angeles. In addition, the volume contains 13 full papers selected from a total of36 submissions.

Query Processing in Database Systems

Until recently, databases contained easily indexed numbers and text. Today, in the age of powerful, graphically based computers, and the world wide web, databases are likely to contain a much greater variety of data forms, including images, sound, video clips, and even handwritten documents. When multimedia databases are the norm, traditional methods of working with databases no longer apply. How do you query a video library, or an image database containing x-rays, or sounds in an audio database? Principles of Multimedia Database Systems explains how to work with these new multimedia data forms. It is the first comprehensive treatment of the skills and techniques required to build, maintain, and query multimedia databases. This book presents the mix of techniques necessary for working with multimedia databases, including synthetic solutions

for the design and deployment of multimedia database systems. Because rapid technological developments are constantly changing the landscape of multimedia databases, the book teaches basic theoretical principles applicable to any database. * Covers the major issues of multimedia database design, with a strong focus on distributed multimedia databases. * Discusses important topics including how to organize the vast data types, storage and retrieval, and creation and delivery of multimedia presentations. * Organized around the lively scenario of a crime-fighting database that evolves as new concepts are introduced. * Includes numerous exercises and suggestions for programming projects. * Additional materials on the web include updates, on-line supplements, and links to downloadable software.

Advanced Metasearch Engine Technology

This book is an anthology of the results of research and development in database query processing during the past decade. The relational model of data provided tremendous impetus for research into query processing. Since a relational query does not specify access paths to the stored data, the database management system (DBMS) must provide an intelligent query-processing subsystem which will evaluate a number of potentially efficient strategies for processing the query and select the one that optimizes a given performance measure. The degree of sophistication of this subsystem, often called the optimizer, critically affects the performance of the DBMS. Research into query processing thus started

has taken off in several directions during the past decade. The emergence of research into distributed databases has enormously complicated the tasks of the optimizer. In a distributed environment, the database may be partitioned into horizontal or vertical fragments of relations. Replicas of the fragments may be stored in different sites of a network and even migrate to other sites. The measure of performance of a query in a distributed system must include the communication cost between sites. To minimize communication costs for-queries involving multiple relations across multiple sites, optimizers may also have to consider semi-join techniques.

Proceedings of the ACM SIGACT-SIGMOD Symposium on Principles of Database Systems

Principles of Database Query Processing for Advanced Applications

Architecture of a Database System presents an architectural discussion of DBMS design principles, including process models, parallel architecture, storage system design, transaction system implementation, query processor and optimizer architectures, and typical shared components and utilities.

Principles of Transaction Processing

The latest edition of a popular text and reference on database research, with substantial new material and revision; covers classical literature and recent hot topics. Lessons from database research have been applied in academic fields ranging from bioinformatics to next-generation Internet architecture and in industrial uses including Web-based e-commerce and search engines. The core ideas in the field have become increasingly influential. This text provides both students and professionals with a grounding in database research and a technical context for understanding recent innovations in the field. The readings included treat the most important issues in the database area--the basic material for any DBMS professional. This fourth edition has been substantially updated and revised, with 21 of the 48 papers new to the edition, four of them published for the first time. Many of the sections have been newly organized, and each section includes a new or substantially revised introduction that discusses the context, motivation, and controversies in a particular area, placing it in the broader perspective of database research. Two introductory articles, never before published, provide an organized, current introduction to basic knowledge of the field; one discusses the history of data models and query languages and the other offers an architectural overview of a database system. The remaining articles range from the classical literature on database research to treatments of current hot topics, including a paper on search engine architecture and a paper on application servers, both written expressly for this edition. The result is a collection of papers that are seminal and also accessible to a reader who has a basic familiarity

with database systems.

Architecture of a Database System

This book provides a concise and modern treatment of introductory database topics that enlists Java and the Internet to present core Database Management (DBMS) theory from an applications perspective. It incorporates programming and database applications when presenting the core theory behind DBMS and their applications. Information management is the central theme of Principles of Database Systems with Internet and Java Applications. The book motivates the development of data models and the representation of information in relational database systems. Students learn how to define database content with Entity-Relationship models, and how to represent that content in relational systems. They become thoroughly familiar with the SQL language, and learn exactly what is required to build quality information-rich applications. Students also learn how the World Wide Web and Java can work together to publish and collect information in the widest possible context. This book covers the basic material of information management in detail. Topics covered include analyzing information requirements, conceptual data modeling, translation of conceptual models to relational needs, normalization of relational schemas, SQL, and database application programming. Additional topics include object-oriented modeling and object databases, database performance and optimization, constraints and triggers, transactions, and file structures. The

interaction between applications and databases is discussed and illustrated in the context of Web sites. The JDBC classes of Java provide a database- and platform-independent method of creating database applications, and all of these classes are thoroughly discussed with abundant examples. After learning the fundamentals of HTML and CGI programming, students create their own Web sites using Java programs to service CGI requests and generate HTML responses. Further topics include the use of Java servlets to replace CGI programs and the use of Java I/O classes for the development of file structures. The Java language provides the foundation for all programming examples because of its portable approach to database access through the JDBC classes. Students do not need extensive experience with Java before using this book, only knowledge of an object-oriented language.

Advances in Databases

This third edition of a classic textbook can be used to teach at the senior undergraduate and graduate levels. The material concentrates on fundamental theories as well as techniques and algorithms. The advent of the Internet and the World Wide Web, and, more recently, the emergence of cloud computing and streaming data applications, has forced a renewal of interest in distributed and parallel data management, while, at the same time, requiring a rethinking of some of the traditional techniques. This book covers the breadth and depth of this re-emerging field. The coverage consists of two parts.

The first part discusses the fundamental principles of distributed data management and includes distribution design, data integration, distributed query processing and optimization, distributed transaction management, and replication. The second part focuses on more advanced topics and includes discussion of parallel database systems, distributed object management, peer-to-peer data management, web data management, data stream systems, and cloud computing. New in this Edition: • New chapters, covering database replication, database integration, multidatabase query processing, peer-to-peer data management, and web data management. • Coverage of emerging topics such as data streams and cloud computing • Extensive revisions and updates based on years of class testing and feedback Ancillary teaching materials are available.

Encyclopedia of Database Systems

Query Processing on Probabilistic Data: A Survey presents the main approaches developed in the literature on probabilistic relational data, reconciling concepts developed in parallel by the Database and Artificial Intelligence communities.

Database Systems

"This book presents cutting-edge research and analysis of the most recent advancements in the fields of database systems and software development"--Provided by publisher.

Proceedings of the Fourth ACM SIGACT-SIGMOD Symposium on Principles of Database Systems, March 25-27, 1985, Portland, Oregon

Twenty-fifth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems

There are millions of searchable data sources on the Web and to a large extent their contents can only be reached through their own query interfaces. There is an enormous interest in making the data in these sources easily accessible. There are primarily two general approaches to achieve this objective. The first is to surface the contents of these sources from the deep Web and add the contents to the index of regular search engines. The second is to integrate the searching capabilities of these sources and support integrated access to them. In this book, we introduce the state-of-the-art techniques for extracting, understanding, and integrating the query interfaces of deep Web data sources. These techniques are critical for producing an integrated query interface for each domain. The interface serves as the mediator for searching all data sources in the concerned domain. While query interface integration is only relevant for the deep Web integration approach, the extraction and understanding of query interfaces are critical for both deep Web exploration approaches. This book aims to provide in-depth and comprehensive

coverage of the key technologies needed to create high quality integrated query interfaces automatically. The following technical issues are discussed in detail in this book: query interface modeling, query interface extraction, query interface clustering, query interface matching, query interface attribute integration, and query interface integration. Table of Contents: Introduction / Query Interface Representation and Extraction / Query Interface Clustering and Categorization / Query Interface Matching / Query Interface Attribute Integration / Query Interface Integration / Summary and Future Research

Proceedings

Data Management and Query Processing in Semantic Web Databases

This comprehensive textbook teaches the fundamentals of database design, modeling, systems, data storage, and the evolving world of data warehousing, governance and more. Written by experienced educators and experts in big data, analytics, data quality, and data integration, it provides an up-to-date approach to database management. This full-color, illustrated text has a balanced theory-practice focus, covering essential topics, from established database technologies to recent trends, like Big Data, NoSQL, and more. Fundamental concepts are supported by real-world examples, query and code walkthroughs, and figures,

making it perfect for introductory courses for advanced undergraduates and graduate students in information systems or computer science. These examples are further supported by an online playground with multiple learning environments, including MySQL; MongoDB; Neo4j Cypher; and tree structure visualization. This combined learning approach connects key concepts throughout the text to the important, practical tools to get started in database management.

Complex Query Processing in Multiprocessor Database Machines

This text introduces the basic concepts and techniques in VIR. In doing so, it develops a foundation for further research and study. Divided into two parts, the first part describes the fundamental principles. A chapter is devoted to each of the main features of VIR, such as colour, texture and shape-based search. There is coverage of search techniques for time-based image sequences or videos, and an overview of how to combine all the basic features described and integrate them into the search process. The second part looks at advanced topics such as multimedia query. This book is essential reading for researchers in VIR, and final-year undergraduate and postgraduate students on courses such as Multimedia Information Retrieval, Multimedia Databases, and others.

Constraints and Databases

Database: Principles Programming Performance provides an introduction to the fundamental principles of database systems. This book focuses on database programming and the relationships between principles, programming, and performance. Organized into 10 chapters, this book begins with an overview of database design principles and presents a comprehensive introduction to the concepts used by a DBA. This text then provides grounding in many abstract concepts of the relational model. Other chapters introduce SQL, describing its capabilities and covering the statements and functions of the programming language. This book provides as well an introduction to Embedded SQL and Dynamic SQL that is sufficiently detailed to enable students to immediately start writing database programs. The final chapter deals with some of the motivations for database systems spanning multiple CPUs, including client-server and distributed transactions. This book is a valuable resource for database administrators, application programmers, specialist users, and end users.