Principles Of Transaction Processing Second Edition The Morgan Kaufmann Series In Data Management Systems

Spring Persistence with Hibernate
Management of the Logical Database and its Underlying Physical Structure
Model Rules of Professional Conduct
Principles of Database Management
Transaction Processing: Past, Present, and Future
What You Need to Know about Data Mining and Data-Analytic Thinking
ODMG 3.0
Information Visualization in Data Mining and Knowledge Discovery
Understanding Relational Language Components
The Object Data Standard
Database Modeling with Microsoft® Visio for Enterprise Architects
Advanced SQL:1999
Designing Data-Intensive Applications
Production-Ready Microservices
Java Data Mining: Strategy, Standard, and Practice
Building Standardized Systems Across an Engineering Organization
Managing Reference Data in Enterprise Databases
Joe Celko's Data and Databases
Information Modeling and Relational Databases
30-Second Coding
Joe Celko's Thinking in Sets: Auxiliary, Temporal, and Virtual Tables in SQL
The 50 Essential Principles that Instruct Technology, Each Explained in Half a Minute
How Google Runs Production Systems
Practical Techniques for Data Preparation
Codeless Data Structures and Algorithms
A Guide to SQLJ, JDBC, and Related Technologies
Transaction Processing
An Application-oriented Approach
Data Science for Business
Databases and Transaction Processing
Principles, Programming, and Performance, Second Edition
Component Database Systems
Understanding Object-Relational and Other Advanced Features
Transactional Information Systems
Data Mining

Transaction Processing
The Big Ideas Behind Reliable, Scalable, and Maintainable Systems
Database Design and Implementation
Advanced Database Systems

MOHAMMED ENGLISH

Spring Persistence with Hibernate Morgan Kaufmann
DW 2.0: The Architecture for the Next Generation of Data Warehousing is the first book on the new generation of data warehouse architecture, DW 2.0, by the father of the data warehouse. The book describes the future of data warehousing that is technologically possible today, at both an architectural level and technology level. The perspective of the book is from the top down: looking at the overall architecture and then delving into the issues underlying the components. This allows people who are building or using a data warehouse to see what lies ahead and determine what new technology to buy, how to plan extensions to the data warehouse, what can be salvaged from the current system, and how to justify the expense at the most practical level. This book gives experienced data warehouse professionals everything they need in order to implement the new generation DW 2.0. It is designed for professionals in the IT organization, including data architects, DBAs, systems design and development professionals, as well as data warehouse and knowledge management professionals. * First book on the new generation of data warehouse architecture, DW 2.0. * Written by the "father of the data warehouse", Bill Inmon, a columnist and newsletter editor of The Bill Inmon Channel on the Business Intelligence Network. * Long overdue comprehensive coverage of the implementation of technology and tools that enable the new generation of the DW: metadata, temporal data, ETL, unstructured data, and data quality control.
Management of the Logical Database and its Underlying Physical Structure Morgan Kaufmann
This is the first guide to transaction processing systems that is accessible to the broad audience that needs to know how transaction processing systems work. It includes a wide variety of real-world examples to illustrate key principles and explains all major open and vendor-specific transaction processing standards.
Model Rules of Professional Conduct Morgan Kaufmann
Component Database Systems is a collection of invited chapters by the researchers making the most influential contributions in the database industry's trend toward componentization This book represents the sometimes-divergent, sometimes-convergent approaches taken by leading database vendors as they seek to establish commercially viable componentization strategies. Together, these contributions form the first book devoted entirely to the technical and architectural design of component-based database systems. In addition to detailing the current state of their research, the authors also take up many of the issues affecting the likely future directions of component databases. If you have a stake in the evolution of any of today's leading database systems, this book will make fascinating reading. It will also help prepare you for the technology that is likely to become widely available over the next several years. * Is comprised of contributions from the field's most highly respected researchers, including key figures at IBM, Oracle, Informix, Microsoft, and POET. * Represents the entire spectrum of approaches taken by leading software companies working on DBMS componentization strategies. * Covers component-focused architectures, methods for hooking components into an overall system, and support for component development. * Examines the component technologies that are most valuable to Web-based and multimedia databases. * Presents a thorough classification and overview of component database systems.
Principles of Database Management Elsevier
This guide documents SQL: 1999Us advanced features in the same practical, "programmercentric" way that the first volume documented the language's basic features. This is no mere representation of the standard, but rather authoritative guidance on making an application conform to it, both formally and effectively.
Transaction Processing: Past, Present, and Future Morgan Kaufmann
Whether you are a software developer, systems architect, data analyst, or business analyst, if you want to take advantage of data mining in the development of advanced analytic applications, Java Data Mining, JDM, the new standard now implemented in core DBMS and data mining/analysis software, is a key solution component. This book is the essential guide to the usage of the JDM standard interface, written by contributors to the JDM standard. Data mining introduction - an overview of data mining and the problems it can address across industries; JDM's place in strategic solutions to data mining-related problems JDM essentials - concepts, design approach and design issues, with detailed code examples in Java; a Web Services interface to enable JDM functionality in an SOA environment; and illustration of JDM XML Schema for JDM objects JDM in practice - the use of JDM from vendor implementations and approaches to customer applications, integration, and usage; impact of data mining on IT infrastructure; a how-to guide for building applications that use the JDM API Free, downloadable KJDM source code referenced in the book available here
"O'Reilly Media, Inc."
This book is for database designers and database administrators using Visio, which is the database component of Microsoft's Visual Studio .NET for Enterprise Architects suite, also included in MSDN subscriptions. This is the only guide to this product that tells DBAs how to get their job done. Although primarily focused on tool features, the book also provides an introduction to data modeling, and includes practical advice on managing database projects. The principal author was the program manager of VEA's database modeling solutions. · Explains how to model databases with Microsoft® Visio for Enterprise Architects (VEA), focusing on tool features. · Provides a platform-independent introduction to data modeling using both Object Role Modeling (ORM) and Entity Relationship Modeling (ERM), and includes practical advice on managing database projects. · Additional ORM models, course notes, and add-ins available online.

What You Need to Know about Data Mining and Data-Analytic Thinking Elsevier
This book offers a thorough grounding in machine learning concepts combined with practical advice on applying machine learning tools and techniques in real-world data mining situations. Clearly written and effectively illustrated, this book is ideal for anyone involved at any level in the work of extracting usable knowledge from large collections of data. Complementing the book's instruction is fully functional machine learning software.
ODMG 3.0 Cambridge University Press
Written by renowned data science experts Foster Provost and Tom Fawcett, Data Science for Business introduces the fundamental principles of data science, and walks you through the "data-analytic thinking" necessary for extracting useful knowledge and business value from the data you collect. This guide also helps you understand the many data-mining techniques in use today. Based on an MBA course Provost has taught at New York University over the past ten years, Data Science for Business provides examples of real-world business problems to illustrate these principles. You'll not only learn how to improve communication between business stakeholders and data scientists, but also how participate intelligently in your company's data science projects. You'll also discover how to think data-analytically, and fully appreciate how data science methods can support business decision-making. Understand how data science fits in your organization—and how you can use it for competitive advantage Treat data as a business asset that requires careful investment if you're to gain real value Approach business problems data-analytically, using the data-mining process to gather good data in the most appropriate way Learn general concepts for actually extracting knowledge from data Apply data science principles when interviewing data science job candidates
Information Visualization in Data Mining and Knowledge Discovery Apress
The database field has experienced a rapid and incessant growth since the development of relational databases. The progress in database systems and applications has produced a diverse landscape of specialized technology areas that have often become the exclusive domain of research specialists. Examples include active databases, temporal databases, object-oriented databases, deductive databases, imprecise reasoning and queries, and multimedia information systems. This book provides a systematic introduction to and an in-depth treatment of these advanced database areas. It supplies practitioners and researchers with authoritative coverage of recent technological advances that are shaping the future of commercial database systems and intelligent information systems. Advanced Database Systems was written by a team of six leading specialists who have made significant contributions to the development of the technology areas covered in the book. Benefiting from the authors' long experience teaching graduate and professional courses, this book is designed to provide a gradual introduction to advanced research topics and includes many examples and exercises to support its use for individual study, desk reference, and graduate classroom teaching.
Understanding Relational Language Components Springer Science & Business Media
This is a great book! This is the book I wish I had written. --Jim Gray, Microsoft Research, recipient of 1998 A.M. Turing Award for seminal contributions to database and transaction processing researchDatabases and Transaction Processing provides a complete and clear explanation of the conceptual and engineering principles underlying the design and implementation of database and transaction processing applications. Rather than focusing on how to implement the database management system itself, this text focuses on how to build database applications. To provide a solid foundation for these principles, the book thoroughly covers the theory underlying relational databases and relational query languages.To illustrate both database and transaction processing concepts, a case study is carried throughout the book. The technical aspects of each chapter applied to the case study and the software engineering concepts required to implement the case study are discussed.In addition to the more traditional material -- relational databases, SQL, and the ACID properties of transactions -- the book provides in-depth coverage of the most current topics in database and transaction processing tec
The Object Data Standard "O'Reilly Media, Inc."
The key to client/server computing. Transaction processing techniques are deeply ingrained in the fields of databases and operating systems and are used to monitor, control and update information in modern computer systems. This book will show you how large, distributed, heterogeneous computer systems can be made to work reliably. Using transactions as a unifying conceptual framework, the authors show how to build high-performance distributed systems and high-availability applications with finite budgets and risk. The authors provide detailed explanations of why various problems occur as well as practical, usable techniques for their solution. Throughout the book, examples and techniques are drawn from the most successful commercial and research systems. Extensive use of compilable C code fragments demonstrates the many transaction processing algorithms presented in the book. The book will be valuable to anyone interested in implementing distributed systems or client/server architectures.
Database Modeling with Microsoft® Visio for Enterprise Architects Morgan Kaufmann
The text and images in this book are in grayscale. A hardback color version is available. Search for ISBN 9781680922929. Principles of Accounting is designed to meet the scope and sequence requirements of a two-semester accounting course that covers the fundamentals of financial and managerial accounting. This book is specifically designed to appeal to both accounting and non-accounting majors, exposing students to the core concepts of accounting in familiar ways to build a strong foundation that can be applied across business fields. Each chapter opens with a relatable real-life scenario for today's college student. Thoughtfully designed examples are presented throughout each chapter, allowing students to build on emerging accounting knowledge. Concepts are further reinforced through applicable connections to more detailed business processes. Students are immersed in the "why" as well as the "how" aspects of accounting in order to reinforce concepts and promote comprehension over rote memorization.
Advanced SQL:1999 Morgan Kaufmann
Data is at the center of many challenges in system design today. Difficult issues need to be figured out, such as scalability, consistency, reliability, efficiency, and maintainability. In addition, we have an overwhelming variety of tools, including relational databases, NoSQL datastores, stream or batch processors, and message brokers. What are the right choices for your application? How do you make sense of all these buzzwords? In this practical and comprehensive guide, author Martin Kleppmann helps you

navigate this diverse landscape by examining the pros and cons of various technologies for processing and storing data. Software keeps changing, but the fundamental principles remain the same. With this book, software engineers and architects will learn how to apply those ideas in practice, and how to make full use of data in modern applications. Peer under the hood of the systems you already use, and learn how to use and operate them more effectively Make informed decisions by identifying the strengths and weaknesses of different tools Navigate the trade-offs around consistency, scalability, fault tolerance, and complexity Understand the distributed systems research upon which modern databases are built Peek behind the scenes of major online services, and learn from their architectures

Designing Data-Intensive Applications Morgan Kaufmann

This text surveys research from the fields of data mining and information visualisation and presents a case for techniques by which information visualisation can be used to uncover real knowledge hidden away in large databases.

Production-Ready Microservices Elsevier

Principles of Transaction ProcessingMorgan Kaufmann

Java Data Mining: Strategy, Standard, and Practice IBM Redbooks

Transactions are a concept related to the logical database as seen from the perspective of database application programmers: a transaction is a sequence of database actions that is to be executed as an atomic unit of work. The processing of transactions on databases is a well- established area with many of its foundations having already been laid in the late 1970s and early 1980s. The unique feature of this textbook is that it bridges the gap between the theory of transactions on the logical database and the implementation of the related actions on the underlying physical database. The authors relate the logical database, which is composed of a dynamically changing set of data items with unique keys, and the underlying physical database with a set of fixed-size data and index pages on disk. Their treatment of transaction processing builds on the “do-redo-undo” recovery paradigm, and all methods and algorithms presented are carefully designed to be compatible with this paradigm as well as with write-ahead logging, steal-and-no-force buffering, and fine-grained concurrency control. Chapters 1 to 6 address the basics needed to fully appreciate transaction processing on a centralized database system within the context of our transaction model, covering topics like ACID properties, database integrity, buffering, rollbacks, isolation, and the interplay of logical locks and physical latches. Chapters 7 and 8 present advanced features including deadlock-free algorithms for reading, inserting and deleting tuples, while the remaining chapters cover additional advanced topics extending on the preceding foundational chapters, including multi-granular locking, bulk actions, versioning, distributed updates, and write-intensive transactions. This book is primarily intended as a text for advanced undergraduate or graduate courses on database management in general or transaction processing in particular.

Building Standardized Systems Across an Engineering Organization 30 Second

A key task that any aspiring data-driven organization needs to learn is data wrangling, the process of converting raw data into something truly useful. This practical guide provides business analysts with an overview of various data wrangling techniques and tools, and puts the practice of data wrangling into context by asking, "What are you trying to do and why?" Wrangling data consumes roughly 50-80% of an analyst’s time before any kind of analysis is possible. Written by key executives at Trifacta, this book walks you through the wrangling process by exploring several factors—time, granularity, scope, and structure—that you need to consider as you begin to work with data. You’ll learn a shared language and a comprehensive understanding of data wrangling, with an emphasis on recent agile analytic processes used by many of today’s data-driven organizations. Appreciate the importance—and the satisfaction—of wrangling data the right way. Understand what kind of data is available Choose which data to use and at what level of detail Meaningfully combine multiple sources of data Decide how to distill the results to a size and shape that can drive downstream analysis

Managing Reference Data in Enterprise Databases "O'Reilly Media, Inc."

ODMG is a widely accepted standard for object database modelling; every year more companies implement it. ODMG 3.0 integrates programming languages with databases and ensures the portability of applications across platforms and DBMS products

Joe Celko's Data and Databases Principles of Transaction Processing

In the era of self-taught developers and programmers, essential topics in the industry are frequently learned without a formal academic foundation. A solid grasp of data structures and algorithms (DSA) is imperative for anyone looking to do professional software development and engineering, but classes in the subject can be dry or spend too much time on theory and unnecessary readings. Regardless of your programming language background, Codeless Data Structures and Algorithms has you covered. In this book, author Armstrong Subero will help you learn DSAs without writing a single line of code. Straightforward explanations and diagrams give you a confident handle on the topic while ensuring you never have to open your code editor, use a compiler, or look at an integrated development environment. Subero introduces you to linear, tree, and hash data structures and gives you important insights behind the most common algorithms that you can directly apply to your own programs. Codeless Data Structures and Algorithms provides you with the knowledge about DSAs that you will need in the professional programming world, without using any complex mathematics or irrelevant information. Whether you are a new developer seeking a basic understanding of the subject or a decision-maker wanting a grasp of algorithms to apply to your projects, this book belongs on your shelf. Quite often, a new, refreshing, and unpretentious approach to a topic is all you need to get inspired. What You'll Learn Understand tree data structures without delving into unnecessary details or going into too much theory Get started learning linear data structures with a basic discussion on computer memory Study an overview of arrays, linked lists, stacks and queues Who This Book Is ForThis book is for beginners, self-taught developers and programmers, and anyone who wants to understand data structures and algorithms but don’t want to wade through unnecessary details about quirks of a programming language or don’t have time to sit and read a massive book on the subject. This book is also useful for non-technical decision-makers who are curious about how algorithms work.

Information Modeling and Relational Databases Springer Science & Business Media

The role of IT is becoming more prominent in people's daily lives and we are becoming increasingly dependent on computers. More and more business transactions are being automated, for

example, ordering a book at an online bookstore or transferring money to a bank account in another part of the world. No matter the type of transaction, we want it to be accurate and we want to have no doubts about its outcome. Transactions are also becoming more complex, driven by new ways of conducting business and new technologies. Smartphones now allow us to conduct transactions anywhere and at anytime. Technology paradigms, such as Web 2.0 and business event processing, enable businesses to increase the dynamics of a transaction through instrumentation that captures events, analyzes the associated data, and proactively interacts with the client in order to improve the customer experience. To adapt to the increasing volume and complexity of transactions requires an ongoing assessment of the current way of supporting transactions with IT. No matter what your business is, you need to ensure that your transactions are properly completed with integrity. Wrong or incomplete results can adversely affect client loyalty, affect company profits, and lead to claims, lawsuits, or fines. Companies need to be able to rely on computer systems that are 100% reliable and guarantee transaction integrity at all times. The IBM® mainframe is such a platform. Clients that have been using an IBM mainframe are conscious of its added value. For this IBM RedguideTM publication, we surveyed a number of companies that use the IBM mainframe and we asked them to tell us its most distinguishing qualities. They answered unanimously "reliability, availability, and scalability." They also do not see an alternative for running their mission-critical business workloads other than the IBM mainframe. When we surveyed our clients, we also asked them about the future. Clearly, major future trends demand significantly smarter, faster, and bigger transaction processing systems than we have today. Some of these trends are the availability of new computing paradigms, continuing growth of the mobile channel, further integration of organizations, massive growth of unstructured and uncertain data, and increasing complexity of IT systems. IBM continues to invest in mainframe technology leadership, which protects years of client investments on this platform. Today, well-known transaction processing (TP) middleware, such as the IBM CICS, IBM IMS, IBM z/TPF, and IBM WebSphere Application Server products, and also solutions for service-oriented architecture (SOA) and business process management (BPM) are available and fully optimized on the IBM mainframe running the mission-critical business workloads of many companies the world over. In 2010, IBM announced the IBM zEnterprise® system introducing a hybrid computing platform that combines the traditional IBM mainframe capabilities and the ability to use IBM blade servers, managed by a single management software. With zEnterprise, you can significantly reduce the complexity of your IT and achieve better service levels, while continuing to benefit from traditional mainframe strengths in transaction processing.