<<数据库系统实现>>

图书基本信息

书名：<<数据库系统实现>>

13位ISBN编号：9787111288602

10位ISBN编号：7111288602

出版时间：2010-1

出版时间：机械工业出版社

作者：（美）加西亚－莫利纳　等著

页数：1181

<<数据库系统实现>>

前言

This book covers the core of the material taught in the database sequence at Stanford. The introductory course, CS145, uses the first twelve chapters, and is designed for all students -those who want to use database systems as well as those who want to get involved in database implementation. The second course, CS245 on database implementation, covers most of the rest of the book. However, some material is covered in more detail in special topics courses. These include CS346 （implementation project）, which concentrates on query optimization as in Chapters 15 and 16. Also, CS345A, on data mining and Web mining, covers the material in the last two chapters.What's New in the Second Edition After a brief introduction in Chapter 1, we cover relational modeling in Chapters 2-4. Chapter 4 is devoted to high-level modeling. There, in addition to the E/R model, we now cover UML （Unified Modeling Language）. We also have moved to Chapter 4 a shorter version of the material on ODL, treating it as a design language for relational database schemas.The material on functional and multivalued dependencies has been mod- ified and remains in Chapter 3. We have changed our viewpoint, so that a functional dependency is assumed to have a set of attributes on the right. We have also given explicitly certain algorithms, including the "chase," that allow us to manipulate dependencies. We have augmented our discussion of third normal form to include the 3NF synthesis algorithm and to make clear what the tradeoff between 3NF and BCNF is.Chapter 5 contains the coverage of relational algebra from the previous edition, and is joined by （part of） the treatment of Dataiog from the old Chap- ter 10. The discussion of recursion in Datalog is either moved to the book's Web site or combined with the treatment of recursive SQL in Chapter 10 of this edition.Chapters 6-10 are devoted to aspects of SQL programming, and they repre- sent a reorganization and augmentation of the earlier book's Chapters 6, 7, 8, and parts of 10. The material on views and indexes has been moved to its own chapter.

<<数据库系统实现>>

内容概要

本书是关于数据库系统实现方面内容最为全面的著作之一，是美国斯坦福大学计算机科学专业数据库系列课程第二门课程的指定教材。
书中从数据库实现者的角度对数据库系统实现原理进行了深入阐述。
并具体讨论了数据库管理系统的三个主要成分——存储管理器、查询处理器和事务管理器的实现技术。
斯坦福大学计算机科学专业数据库系列课程第一门课程的内容包括数据库设计和数据库编程。
本书的后两位作者Jeffrey D.UIIman和Jennifer Widom为该课程编写的教材《数据库系统基础教程》（A First Course in Database Systems）第3版的中文翻译版和英文影印版已由机械工业出版社出版。
本书内容深入且全面，技术实用且先进，叙述深入浅出，是一本难得的高层次的教材，适合作为高等院校计算机专业研究生的教材或本科生的教学参考书，也适合作为从事相关研究或开发工作的专业技术人员的高级参考资料。

作者简介

加西亚－莫利纳（Hector Garcia-Molina）斯坦福大学计算机科学与电子工程系的Leonard Bosack和Sandra Lerner教授。
他在数据库系统、分布式系统和数字图书馆领域中发表了大量论文。
研究兴趣包括分布式计算系统、数据库系统和数字图书馆。
他是ACM会士、美国艺术与科学院会士和

<<数据库系统实现>>

书籍目录

章节摘录

插图：13.2 DisksThe use of secondary storage is one of the important characteristics of a DBMS, and secondary storage is almost exclusively based on magnetic disks. Thus, to motivate many of the ideas used in DBMS implementation, we must examine the operation of disks in detail.13.2.1 Mechanics of DisksThe two principal moving pieces of a disk drive are shown in Fig. 13.2; they are a disk assembly and a head assembly. The disk assembly consists of one or more circular platters that rotate around a central spindle. The upper and lower surfaces of the platters are covered with a thin layer of magnetic material, on which bits are stored. O's and l's are represented by different patterns in the magnetic material. A common diameter for disk platters is 3.5 inches, although disks with diameters from an inch to several feet have been built.The disk is organized into tracks, which are concentric circles on a single platter. The tracks that are at a fixed radius from the center, among all the surfaces, form one cylinder. Tracks occupy most of a surface, except for the region closest to the spindle, as can be seen in the top view of Fig. 13.3. The density of data is much greater along a track than radially. In 2008, a typical disk has about 100,000 tracks per inch but stores about a million bits per inch along the tracks.Tracks are organized into sectors, which are segments of the circle separated by gaps that are not magnetized to represent either O's or l's.1 The sector is an indivisible unit, as far as reading and writing the disk is concerned. It is also indivisible as far as errors are concerned.

<<数据库系统实现>>

编辑推荐

《数据库系统实现(英文版第2版)》：经典原版书库

<<数据库系统实现>>

版权说明