

Database Design and Implementation

Second Edition

This textbook examines database systems from the viewpoint of a software developer. This perspective makes it possible to investigate why database systems are the way they are. It is of course important to be able to write queries, but it is equally important to know how they are processed. We e.g. don't want to just use JDBC; we also want to know why the API contains the classes and methods that it does. We need a sense of how hard it is to write a disk cache or logging facility. And what exactly is a database driver, anyway? The first two chapters provide a brief overview of database systems and their use. Chapter 1 discusses the purpose and features of a database system and introduces the Derby and SimpleDB systems. Chapter 2 explains how to write a database application using Java. It presents the basics of JDBC, which is the fundamental API for Java programs that interact with a database. In turn, Chapters 3-11 examine the internals of a typical database engine. Each chapter covers a different database component, starting with the lowest level of abstraction (the disk and file manager) and ending with the highest (the JDBC client interface); further, the respective chapter explains the main issues concerning the component, and considers possible design decisions. As a result, the reader can see exactly what services each component provides and how it interacts with the other components in the system. By the end of this part, s/he will have witnessed the gradual development of a simple but completely functional system. The remaining four chapters then focus on efficient query processing, and focus on the sophisticated techniques and algorithms that can replace the simple design choices described earlier. Topics include indexing, sorting, intelligent buffer usage, and query optimization. This text is intended for upper-level undergraduate or beginning graduate courses in Computer Science. It assumes that the reader is comfortable with basic Java programming; advanced Java concepts (such as RMI and JDBC) are fully explained in the text. The respective chapters are complemented by "end-of-chapter readings" that discuss interesting ideas and research directions that went unmentioned in the text, and provide references to relevant web pages, research articles, reference manuals, and books. Conceptual and programming exercises are also included at the end of each chapter. Students can apply their conceptual knowledge by examining the SimpleDB (a simple but fully functional database system created by the author and provided online) code and modifying it.

This textbook examines database systems from the viewpoint of a software developer. This perspective makes it possible to investigate why database systems are the way they are. It is of course important to be able to write queries, but it is equally important to know how they are processed. We e.g. don't want to just use JDBC; we also want to know why the API contains the classes and methods that it does. We need a sense of how hard it is to write a disk cache or logging facility. And what exactly is a database driver, anyway? The first two chapters provide a brief overview of database systems and their use. Chapter 1 discusses the purpose and features of a database system and introduces the Derby and SimpleDB systems. Chapter 2 explains how to write a database application using Java. It presents the basics of JDBC, which is the fundamental API for Java programs that interact with a database. In turn, Chapters 3-11 examine the internals of a typical database engine. Each chapter covers a different database component, starting with the lowest level of abstraction (the disk and file manager) and ending with the highest (the JDBC client interface); further, the respective chapter explains the main issues concerning the component, and considers possible design decisions. As a result, the reader can see exactly what services each component provides and how it interacts with the other components in the system. By the end of this part, s/he will have witnessed the gradual development of a simple but completely functional system. The remaining four chapters then focus on efficient query processing, and focus on the sophisticated techniques and algorithms that can replace the simple design choices described earlier. Topics include indexing, sorting, intelligent buffer usage, and query optimization. This text



64,19 €

59,99 € (zzgl. MwSt.)

Lieferfrist: bis zu 10 Tage

Artikelnummer: 9783030338350

Medium: Buch

ISBN: 978-3-030-33835-0

Verlag: Springer

Erscheinungstermin: 28.02.2020

Sprache(n): Englisch

Auflage: 1. Auflage 2020

Serie: Data-Centric Systems and Applications

Produktform: Kartoniert

Gewicht: 709 g

Seiten: 458

Format (B x H): 155 x 235 mm

