

DESIGN OF DATABASE STRUCTURES

TOBY J. TEOREY • JAMES P. FRY

General
information
requirements

Processing
requirements

Database
management
system
characteristics

• Step 1
Requirements
formulation
and analysis

Requirements
specifications

Step 2
Conceptual
design

Information
structure

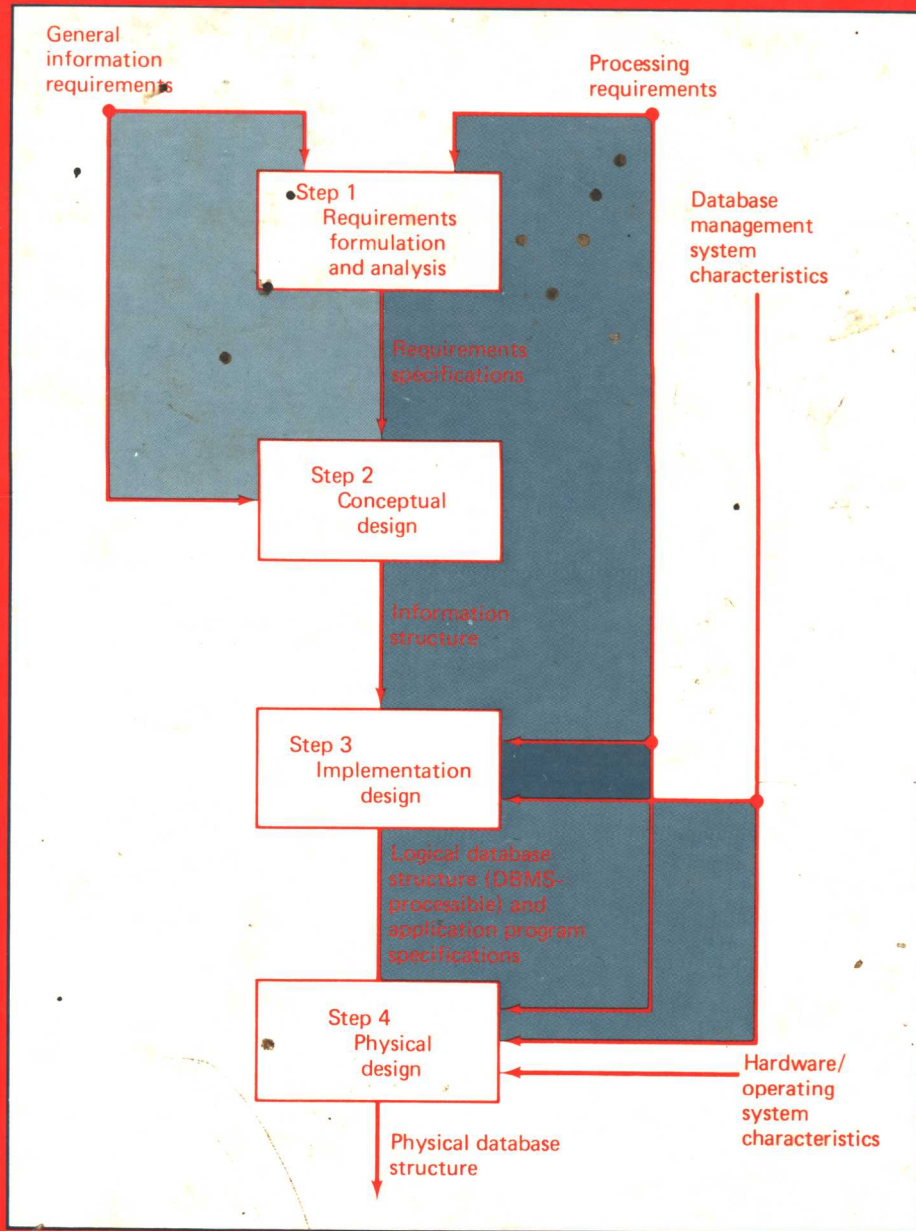
Step 3
Implementation
design

Logical database
structure (DBMS-
processable) and
application program
specifications

Step 4
Physical
design

Hardware/
operating
system
characteristics

Physical database
structure



087860

DESIGN OF DATABASE STRUCTURES


23.879
LT 314

TOBY J. TEOREY
JAMES P. FRY

The University of Michigan



4990087860



69.92
492P
1982

PRENTICE-HALL, INC., Englewood Cliffs, N.J. 07632

Library of Congress Cataloging in Publication Data

Teorey, Toby J.

Design of database structures.

(Prentice-Hall software series)

Bibliography: p. 470

Includes index.

1. Database management. 2. System design. I. Fry.

James P. II. Title. III. Series

QA76.9.D3T43

001.64

81-23375

ISBN 0-13-200097-0

AACR2

Editorial/production supervision

and interior design by *Linda Mihatov Paskiet*

Cover designer: *Frederick Charles, Ltd.*

Manufacturing buyer: *Gordon Osbourne*

Prentice-Hall Software Series

Brian W. Kernighan, advisor

© 1982 by PRENTICE-HALL, INC.,
Englewood Cliffs, New Jersey 07632

All rights reserved. No part of this book
may be reproduced in any form or
by any means without permission in writing
from the publisher.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN 0-13-200097-0

Prentice-Hall International, Inc., *London*

Prentice-Hall of Australia Pty. Limited, *Sydney*

Prentice-Hall of Canada, Ltd., *Toronto*

Prentice-Hall of India Private Limited, *New Delhi*

Prentice-Hall of Japan, Inc., *Tokyo*

Prentice-Hall of Southeast Asia Pte. Ltd., *Singapore*

Whitehall Books Limited, *Wellington, New Zealand*

PREFACE

The purpose of this book is to establish a consistent framework for multilevel database design; to define a workable methodology; and to describe a set of general principles, tools, and techniques for database design at each level. The methodology uses the top-down (and iterative) design path, providing for evaluation at many points, allowing for redesign when necessary, and using the same basic terminology throughout. The goal of the design process itself is to formulate a database structure that accurately represents the real-world environment it serves, and one that can be efficiently implemented on an existing hardware/software system. System independence, however, is maintained as long as possible in the design process (i.e., the conceptual design step).

The methodology is illustrated throughout the text using detailed examples from statements of database system requirements through specification of logical and physical database structures that satisfy those requirements. The approach is applicable to network, relational, and hierarchical database systems.

Many steps in the database design process can be assisted with proper use of computer modeling techniques and other tools such as requirements analysis software. These tools range from documentation mechanisms to sophisticated design methodologies. Techniques such as simulation, graph theory, or mathematical optimization are commonly used, although simple expected-value estimates of performance have been found practical at various stages of database design. Categorized references are provided for the most appropriate tools at each design step.

Although emphasis is placed on the design of the database structure, its relationship to efficient program execution is also explored. Constraints for issues such as integrity, recovery, and security are discussed in terms of trade-offs with efficiency goals.

This book is for the professional database analyst, designer, database administrator, and application programmer. It could also be used for a one-semester senior or first-year graduate course in database design. The material is presented tutorially, and

is supplemented with exercises and a bibliography, glossary of terms, index, and summary of variables. The level of mathematics is college algebra with a knowledge of some basic calculus. The reader is assumed to be familiar with data structures, searching and sorting techniques, and characteristics of the major database management systems (DBMSs). A course based on this text would naturally follow an introductory course on database management systems. Both theory and practice are presented, and their interaction is strongly emphasized.

The order of the chapters follows the database system life cycle steps (Chapter 1) very closely and presents the material in the order that a normal design process would follow. Part I (Chapters 1 to 3) is common to an understanding of all chapters, but thereafter each of the design phases is a complete entity and may be studied in isolation. For instance, a short course in physical design would follow Chapters 1 to 3 and 9 to 16. Chapter 18 presents an overview of the database design issues in distributed systems for the purpose of defining problems for future research.

Requirements analysis and conceptual design are illustrated with a common example (Chapters 3 and 6). Implementation design has a separate complete example (Chapter 8), and the physical design discussion contains smaller individual examples to demonstrate specific points. The examples serve to show that database design can be performed manually, but they also point out where computer-based tools can be effectively implemented.

Ann Arbor

Toby J. Teorey

James P. Fry

ACKNOWLEDGMENTS

We would like to express our gratitude for the various examples and manuscript critiques made by Larry Brown, Janis Bubenko, Bob Curtice, Jeffrey Hoffer, Alan Merten, Shamkant Navathe, Donna Rund, Mario Schkolnick, Dennis Severance, Diane Smith, John Smith, Dick Volz, and the several reviewers. Much of the design framework evolved from a course taught at the University of Michigan Engineering Summer Conference; it was also heavily influenced by the work at the 1978 NYU Symposium on Database Design and the 1978 Database Design Workshop in New Orleans. We acknowledge the contribution to our approach from our colleagues Sakti Ghosh, David Jefferson, Paul Jones, Bob Taylor, Vincent Lum, and Bing Yao.

At the University of Michigan we were assisted with software development, design tool evaluation, and manuscript revisions by the many able staff members of the Information Systems Research Group, including Ed Birss, David Chen, Rick Cobb, K. Sundar Das, Mark Deppe, Don DeSmith, Alberto Garcia, Eric Kintzer, Chris Merrill, Don Novak, Lew Oberlander, Don Swartwout, and Mike Wilens.

We thank Connie Allen, Pam Downie, Carol Dunn, and Izena Goulding for their help in manuscript preparation.

We appreciate the use of the library and technical facilities provided by the Department of Management Information Systems at the University of Arizona under Jay Nunamaker.

Exercises B5, B6, and B7 were contributed by Dennis Severance of the University of Michigan. Exercise B8 was suggested by William Weiler of Blue Cross/Blue Shield of Massachusetts, and Design Problem 3 in Appendix A was designed by Paul Helman and Marilyn Mantei of the University of Michigan. Design Problem 2 was revised from Chen [1978].

Finally, we wish to acknowledge the support and encouragement of our families, including Eunice L. and Thomas F. Teorey, Bessie May and Robert W. Teorey, and Anne J. and Palmer E. Fry.

CONTENTS

Preface

xiii

I

INTRODUCTION

1	Database Systems	3
1-1	Data and Database Management	3
1-2	Levels of Data Representation	5
1-3	The Database Design Problem	12
1-4	Database System Life Cycle	14
1-4-1	Analysis and design phase	15
1-4-2	Database implementation and operation phase	16
1-5	Summary	17
2	The Database Design Process	18
2-1	The Concept of a Design Methodology	18
2-2	Overview of Database Design: the Basic Steps	25
2-3	Design Issues	29
2-3-1	Integrity, consistency, and recovery	31
2-3-2	Security	32
2-3-3	Efficiency	33
3	Requirements Formulation and Analysis	35
3-1	Introduction to Requirements Analysis	36
3-1-1	Need for corporate requirements analysis	37
3-1-2	Requirements formulation and analysis steps	37
3-2	Defining Scope	38

- 3-3 Collecting Information About Data Usage 39
 - 3-3-1 *Operational functions of the business* 41
 - 3-3-2 *Conducting interviews* 42
 - 3-3-3 *Management interviews for control and planning functions* 44
- 3-4 Requirements Information Transformation 45
 - 3-4-1 *Identifying data elements* 45
 - 3-4-2 *Identifying operational tasks* 46
 - 3-4-3 *Identifying control and planning tasks* 50
 - 3-4-4 *Identifying current and future operating policies* 50
- 3-5 Tools for Formulating and Documenting Requirements 51

II

CONCEPTUAL DESIGN

- 4 Conceptual Data Modeling 57**
 - 4-1 Framework for Conceptual Design 57
 - 4-2 Object Representation 61
 - 4-3 Entity Modeling 66
 - 4-4 Methodologies for Conceptual Design 72
 - 4-4-1 *Entity analysis* 72
 - 4-4-2 *Attribute synthesis* 74
- 5 Entity Formulation and Analysis 76**
 - 5-1 Introduction 76
 - 5-1-1 *Design objectives and scope* 76
 - 5-1-2 *Definition of design perspectives* 77
 - 5-2 Modeling of Design Perspectives 78
 - 5-2-1 *Model constructs* 78
 - 5-2-2 *Formulation of design perspectives* 79
 - 5-3 Consolidation of User Views 82
 - 5-3-1 *Concepts and principles for consolidation* 83
 - 5-3-2 *Consolidation types* 85
 - 5-3-3 *Consolidation process* 93
- 6 Attribute Synthesis: An Example of Conceptual Design 97**
 - 6-1 The Basic Model 97
 - 6-2 Analyzing Information to Identify the Components of the Information (Conceptual) Structure 99
 - 6-2-1 *Identifying entities and attributes* 99
 - 6-2-2 *Identifying relationships* 112
 - 6-3 Expressing Entities, Attributes, and Relationships Through the Graphical Notations of the Entity-Relationship Information Structure 119

- 6-4 Interpreting the Information So That It May Be Verified by All Users 123

III

IMPLEMENTATION DESIGN

7 Implementation Design Concepts 133

- 7-1 Implementation Design Components 133
- 7-2 Implementation Design Steps 135
- 7-3 Logical Database Structure Performance 138
 - 7-3-1 Logical record access approach 139
 - 7-3-2 Relationship between logical and physical performance 143
- 7-4 Alternative Implementation Design Tools 146

8 An Example Schema Design Problem 148

- 8-1 Requirements Specifications 148
- 8-2 Conceptual Design 150
- 8-3 Implementation Design 151
- 8-4 Summary of the Major Design Issues 160

IV

PHYSICAL DESIGN

9 Physical Database Design Principles: Basic Concepts 167

- 9-1 Introduction to Physical Design 167
 - 9-1-1 Physical design steps 168
 - 9-1-2 Physical design environment 170
 - 9-1-3 Performance measures 171
- 9-2 Performance Computation 177
 - 9-2-1 I/O service time: the basic model 178
 - 9-2-2 I/O service time for disk storage 180
 - 9-2-3 Dedicated and shared computing environments 181
- 9-3 Secondary Storage Space 186
 - 9-3-1 File storage 186
 - 9-3-2 Database storage: multiple-record types 190
- 9-4 Summary 193

10 Record Structure Design 194

- 10-1 Data Item Encoding and Compression 194
 - 10-1-1 Data item representation 195
 - 10-1-2 Compression techniques 196

10-2	Record partitioning 201	
10-2-1	Record Segmentation Algorithms 201	
10-2-2	The bond energy algorithm 205	
10-2-3	A heuristic partitioning model 211	
11	Record Clustering	212
11-1	Clustering in Hierarchical Databases 213	
11-2	Clustering in Network Database Structures 219	
12	Primary Access Methods: Sequential Processing	225
12-1	Introduction 225	
12-1-1	Access method terminology 226	
12-1-2	Classification of access methods 226	
12-2	Physical Sequential Processing: GET MANY, GET ALL 228	
12-2-1	Physical sequential retrieval 228	
12-2-2	Physical sequential update 231	
12-2-3	Trade-offs in physical sequential organizations 233	
12-3	Linked Sequential Processing 237	
12-3-1	Linked sequential retrieval 237	
12-3-2	Linked sequential update 238	
12-3-3	Trade-offs in linked sequential organizations 240	
12-3-4	Storage space for sequential organizations 243	
12-4	Report Generation: Total Cost 243	
13	Primary Access Methods: Random Processing	248
13-1	Direct Access 248	
13-2	Identifier Hashing (Random Access) 251	
13-2-1	Hashing functions 253	
13-2-2	Overflow techniques 255	
13-2-3	Performance characteristics 267	
13-2-4	Scatter tables 269	
13-3	Full Index (Indexed Random) 271	
13-4	Indexed Sequential 274	
13-4-1	Retrieval from an indexed sequential file 276	
13-4-2	Update to an indexed sequential file 287	
13-4-3	Storage space for indexed sequential organizations 289	
13-4-4	Trade-off analysis in indexed sequential organizations 289	
14	Primary Access Methods: Search Trees and Random Processing	293
14-1	Binary Search Trees 293	
14-1-1	Retrieval performance 296	
14-1-2	Update performance 298	
14-1-3	Storage space 304	

14-2	B-Trees 305	
14-2-1	<i>Retrieval performance</i>	309
14-2-2	<i>Update performance</i>	313
14-2-3	<i>Storage space</i>	321
14-2-4	<i>B*-trees</i>	321
14-2-5	<i>Prefix B*-trees</i>	326
14-3	TRIE Structures 327	
14-3-1	<i>Retrieval performance</i>	330
14-3-2	<i>Update performance</i>	330
14-3-3	<i>Storage space</i>	331
14-3-4	<i>TRIE/B-tree trade-offs</i>	332
15	Secondary Access Methods	334
15-1	Introduction	334
15-2	Multilist File	335
15-2-1	<i>Retrieval performance for query applications</i>	338
15-2-2	<i>Generalized update performance</i>	340
15-2-3	<i>Storage space</i>	341
15-2-4	<i>Cellular multilist</i>	342
15-3	Inverted File	344
15-3-1	<i>Retrieval performance for query applications</i>	346
15-3-2	<i>Generalized update performance</i>	348
15-3-3	<i>Storage space</i>	349
15-3-4	<i>Trade-offs in secondary storage access methods</i>	350
15-3-5	<i>Cellular inverted file</i>	352
15-4	Doubly Chained Tree	355
15-4-1	<i>Retrieval performance for query applications</i>	359
15-4-2	<i>Generalized update performance</i>	361
15-4-3	<i>Storage space</i>	362
15-4-4	<i>Comparison with inverted and multilist files</i>	363
15-5	Inverted Index Maintenance Alternatives	366
16	Secondary Index Selection	370
16-1	Multiple-Key (Combined) Indexing	370
16-2	Categorization of Secondary Access Methods	373
16-3	The Index Selection Problem	377
16-4	Optimal Secondary Index Selection	380

V

SPECIAL DESIGN ISSUES

17	Reorganization	387
17-1	Introduction	387
17-2	Reorganization Strategies	389

- 17-3 The Database Administrator's Role 391
- 17-4 When to Reorganize: a Heuristic Rule 392
- 17-5 Restructuring 397
 - 17-5-1 Network restructuring 399
 - 17-5-2 Technical approaches to restructuring 401

18 Distributed Database Design: An Overview 408

- 18-1 Introduction 408
- 18-2 Distributed DBMS Architecture 409
- 18-3 Design Issues in a Distributed Database Environment 414
 - 18-3-1 Data distribution strategies 415
 - 18-3-2 Network data directory distributions 419
 - 18-3-3 Homogeneous versus heterogeneous environments 420
- 18-4 A Framework for Distributed Database Design 422
 - 18-4-1 Database partitioning 422
 - 18-4-2 Database allocation 425
- 18-5 Differential Files 427

APPENDICES

A	Exercises in Conceptual and Implementation Schema Design	437
B	Exercises in Physical Database Design	445
C	List of Variables	455
	Glossary	460
	References	470
	Index	485

I INTRODUCTION

1 DATABASE SYSTEMS

Designing an integrated database is a difficult, time-consuming, and often unstructured process. It is a complex problem that pervades not only the data-processing function but eventually the entire organization. The quality of the resulting database structure is dependent upon the design methodology, the design techniques used in the steps of the methodology, the validity of the information requirements, and the commitment of the organization's developmental and operational resources to the endeavor.

In this chapter we begin to build the foundations for a practical database design methodology by establishing a standard set of terminology to be used throughout the text. We also review some of the basic concepts of data as used in the organization, as specified in application programs and database management systems, and as stored in computer systems.

1-1 DATA AND DATABASE MANAGEMENT

Data, as defined by Webster's is "a fact; something upon which an inference or an intellectual system of any sort is based" [Webster's, 1974]. The primitive components of data are the characters and numbers upon which a natural language is based, or their coded representations in strings of binary bits. A *data item* is the smallest unit of data that has meaning in the real world; it is the smallest named unit of data. A group of related data items treated as a unit by an application program is known as a *record*, and the collection of records of a single type is referred to as a *file*. Files can be manually stored or mechanized (electronically stored) on a computer. The distinction between logical records, as viewed by the applications programmer, and stored records, as viewed by the storage devices and systems programmer/analyst, is emphasized throughout the text.

The remarkable continued growth in the use of computers for a wide variety of

industrial, administrative, and scientific applications has led to the mechanization of extremely large quantities of data. During the late 1950s and early 1960s, businesses, government, and other organizations began to collect and store data on computer-accessible files. As each new need for data or data processing arose, a new file was created to meet that need. Individual groups within organizations developed their own computer applications and collected and maintained any required data in private files. The data files and applications programs were designed for one another, with much of the information being stored implicitly in the relationship of the program to the file.

Organizations eventually developed an awareness of the need for centralized management of data and applications. This awareness developed in several ways. Primarily, higher-level managers discovered that the information required to support their decision making was not easily obtainable. To fulfill a request for information, an application program had to be written to access several private files, with each file using its own format. The manager often canceled the request, either because the information would no longer be useful by the time it could be obtained, or because the information was not worth the cost involved. Second, decision making was hampered by the lack of data integrity. Computer reports conflicted because logically identical data items had different values. This was caused by data duplication among the private files, combined with inconsistent updating policies. There was also duplication of effort in the collection of data and the production of reports. Computer resources, both storage and processing, were being wasted. Finally, technology improved to the point where it became feasible to design, build, and operate large-scale collections of data in a computer environment. In summary, organizations realized that data were a valuable resource and needed to be centrally managed.

The concept of a database has thus emerged fully only in recent years. A *database* can be defined as a computerized collection of stored operational data that serves the needs of multiple users within one or more organizations. A key point is that the database is an integrated resource to be used by all members of the organizations who need information contained in it. Information is no longer implicit in the file-program combination but is stored explicitly in the database, which may include many different record types. The database is not single-program oriented as were private data files, but has an integrated requirements orientation.

Databases support managerial decision making in the following ways:

1. *Speed.* The computer system allows on-line querying for information.
2. *Total availability.* All the information contained within a database is available for use.
3. *Flexibility.* Previously unanswerable questions become answerable, and changes become relatively easy to implement.
4. *Integrity.* Data duplication is reduced and updating policies can be standardized, resulting in consistent data.

Simply having a database does not entirely solve the organization's data-processing and decision-making needs, however. Since the database is an integrated

resource for multiple users within an organization, it should be managed for the organization's benefit and from its viewpoint, not by individual users. Users naturally tend to develop applications for their own purposes, unconcerned about the impact that the new applications will have on other users. Without centralized management, the usefulness of a database declines over time.

Two additional concepts have been developed to solve the problem of controlling and managing the organization's database resource. Initially, software was developed to provide a common interface between all users and the integrated database. A common interface promotes privacy and data integrity. Also, users cannot store information implicitly and must use and modify data in a manner consistent with the organization's viewpoint. The software known as a database management system allows computer control of the data resource. A *database management system* (DBMS) is a generalized tool for manipulating a database; it is made available through special software for the interrogation, maintenance, and analysis of data. Its interfaces generally provide a broad range of languages to aid a wide variety of users. It also provides a convenient framework in which databases can be designed and used. A DBMS can be obtained by an organization through either a software vendor or in-house development.

The second concept is that of the *database administrator* (DBA). The DBA can be thought of as one or more individuals, possibly aided by a staff, who manage the organization's database resource. The DBA should be a dynamic and capable individual, primarily management oriented but also with a technical background. This person must be able to communicate with both upper-level management and data-processing users as well as manage the staff of technical specialists. The staff should include individuals experienced in a variety of areas, such as DBMS software, operating systems, computer hardware, applications programming, and system design. It is important that the staff also include individuals with a knowledge of the organization and its information requirements. The DBA's staff must be able to relate well to groups outside the data-processing department.

The DBA position was created when organizations recognized the need for centralized control of the data resource, data processing, and all other aspects that concern a database. The user community and individual users must be served as fairly as possible while keeping in mind the objectives of the organization as a whole. Thus, the DBA is responsible for determining all users' needs, designing a database, implementing the database, modifying or converting the database as needed, and assisting users with documentation and through education [Canning, 1972; DeBlasis and Johnson, 1977; Kent, 1978; Lyon, 1976a].

1-2 LEVELS OF DATA REPRESENTATION

At least three levels of abstraction are now recognized to exist in the specification of a database structure: the conceptual or enterprise administrator view, the implementation view of the applications programmer or end user, and the physical view of the