

An Introduction to Database Systems

Third Edition

C. J. DATE
IBM Corporation

◆◆ **ADDISON-WESLEY PUBLISHING COMPANY**
Reading, Massachusetts • Menlo Park, California
London • Amsterdam • Don Mills, Ontario • Sydney

GOG 032264

Sponsoring Editor: William B. Gruener
Production Editor: Martha K. Morong
Designer: Herbert E. Caswell
Illustrator: Robert Gallison
Cover Design: Richard Hannus

This book is in the
Addison-Wesley Systems Programming Series
Consulting editors: IBM Editorial Board

Library of Congress Cataloging in Publication Data

Date, C. J.

An introduction to database systems.

(The systems programming series)
Includes bibliographies and index.

1. Data base management. I. Title.

QA76.9.D3D37 1981 001.64 80-17603
ISBN 0-201-14471-9

Reprinted with corrections, February 1982

Copyright © 1981, 1977, 1975 by Addison-Wesley Publishing Company, Inc. Philippines
copyright 1981, 1977, 1975 by Addison-Wesley Publishing Company, Inc.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system,
or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording,
or otherwise, without the prior written permission of the publisher. Printed in the United
States of America. Published simultaneously in Canada. Library of Congress Catalog Card
No. 80-17603.

ISBN 0-201-14471-9
BCDEFGHIJ-HA-898765432

GOG 032265

1

Basic Concepts

1.1 WHAT IS A DATABASE SYSTEM?

Database technology has been described as "one of the most rapidly growing areas of computer and information science" [1.14]. As a field, it is still comparatively young; manufacturers and vendors did not begin to offer database management system products until well into the 1960s (although it is true that certain earlier software packages did include some of the functions now associated with such systems [1.13, 1.15]). Despite its youth, however, the field has quickly become one of considerable importance, both practical and theoretical. The total amount of data now committed to databases can be measured, conservatively, in the billions of bytes; the financial investment involved is represented by a correspondingly enormous figure; and it is no exaggeration to say that many thousands of organizations have become critically dependent on the continued and successful operation of a database system.

So what exactly is a database system? Basically, it is nothing more than a computer-based recordkeeping system: that is, a system whose overall purpose is to record and maintain information.¹ The information concerned can be anything that is deemed to be of significance to the organization the

1. The terms "data" and "information" are treated as synonymous in this book. Some writers distinguish between the two, using "data" to refer to the values physically recorded in the database and "information" to refer to the *meaning* of those values as understood by some user. The distinction is clearly important—so important that it seems preferable to make it explicit, where relevant, instead of relying on a somewhat arbitrary differentiation between two essentially similar terms.

4 Basic Concepts




Fig. 1.1 Simplified picture of a database system.

system is serving—anything, in other words, that may be necessary to the decision-making processes involved in the management of that organization. Figure 1.1 shows a greatly simplified view of a database system.

Figure 1.1 is intended to show that a database system involves four major components: data, hardware, software, and users. We consider each of these briefly below. Later in the chapter we shall discuss each one in rather more detail.

Data

The data stored in the system is partitioned into one or more *databases*. For tutorial purposes it is usually convenient to assume that there is just one database, containing the totality of all stored data in the system, and we will generally make this simplifying assumption, since it does not substantially invalidate any of the subsequent discussion. There are good reasons why such a restriction should not be enforced in practice, however, as will be seen later.

A database, then, is a repository for stored data. In general, it is both *integrated* and *shared*.

1.1

By "integration of files among those files base might contain department, salary, rollment of empl out the process of information, redun be discovered by

By "shared" may be shared an users may have ac purposes). Such s is integrated; in the department in the personnel dep consequence of t given user will n database; moreov ways. In other w in a variety of di of the database, t tailed level. This

The term "sh just described, b different users to same piece of da form of sharing i

Hardware

The hardware co etc.—on which tl control units, cha large to be held in book does not co tem, for the follo their own right; s culiar to database oughly investigat


By "integrated" we mean that the database may be thought of as a unification of several otherwise distinct data files, with any redundancy among those files partially or wholly eliminated. For example, a given database might contain both EMPLOYEE records, giving name, address, department, salary, etc., and ENROLLMENT records, representing the enrollment of employees in training courses. Suppose that in order to carry out the process of course administration, it is necessary to know the department for each enrolled student. There is clearly no need to include this information, redundantly, in ENROLLMENT records, because it can always be discovered by referring to the corresponding EMPLOYEE records.

By "shared" we mean that individual pieces of data in the database may be shared among several different users, in the sense that each of those users may have access to the same piece of data (and may use it for different purposes). Such sharing is really a consequence of the fact that the database is integrated; in the EMPLOYEE/ENROLLMENT example cited above, the department information in EMPLOYEE records is shared by users in the personnel department and users in the education department. Another consequence of the same fact (that the database is integrated) is that any given user will normally be concerned only with some subset of the total database; moreover, different users' subsets will overlap in many different ways. In other words, a given database will be perceived by different users in a variety of different ways. (Even when two users share the same subset of the database, their views of that subset may differ considerably at a detailed level. This topic is discussed more fully in Section 1.4.)

The term "shared" is frequently extended to cover, not only sharing as just described, but also *concurrent* sharing: that is, the ability for several different users to be actually accessing the database—possibly even the same piece of data—at the same time. (A database system supporting this form of sharing is sometimes referred to as a multiuser system.)

Hardware

The hardware consists of the secondary storage volumes—disks, drums, etc.—on which the database resides, together with the associated devices, control units, channels, and so forth. (We assume that the database is too large to be held in its entirety within the computer's primary storage.) This book does not concern itself very greatly with hardware aspects of the system, for the following reasons: First, these aspects form a major topic in their own right; second, the problems encountered in this area are not peculiar to database systems; and third, those problems have been very thoroughly investigated and documented elsewhere.



ary to the organiza- stem.

ives four sider each ch one in

ases. For just one id we will stantially sons why as will be

it is both

6 Basic Concepts

Software

Between the physical database itself (i.e., the data as actually stored) and the users of the system is a layer of software, usually called the database management system or DBMS. All requests from users for access to the database are handled by the DBMS. One general function provided by the DBMS is thus the shielding of database users from hardware-level detail (in much the same way that programming-language systems for languages such as COBOL shield programming users from hardware-level detail). In other words, the DBMS provides a view of the database that is elevated somewhat above the hardware level, and supports user operations (such as "get the EMPLOYEE record for employee Smith") that are expressed in terms of that higher-level view. We shall discuss this function, and other functions of the DBMS, in considerably more detail later.

Users

We consider three broad classes of user. First, there is the *application programmer*, responsible for writing application programs that use the database, typically in a language such as COBOL or PL/I. These application programs operate on the data in all the usual ways: retrieving information, creating new information, deleting or changing existing information. (All these functions are performed by issuing the appropriate request to the DBMS.) The programs themselves may be conventional batch applications, or they may be "on-line" programs that are designed to support an end-user (see below) interacting with the system from an on-line terminal.

The second class of user, then, is the *end-user*, accessing the database from a terminal. An end-user may employ a *query language* provided as an integral part of the system, or (as mentioned above) he or she may invoke a user-written application program that accepts commands from the terminal and in turn issues requests to the DBMS on the end-user's behalf. Either way the user may again, in general, perform all the functions of retrieval, creation, deletion, and modification, although it is probably true to say that retrieval is the most common function for this class of user.

The third class of user is the *database administrator*, or DBA (not shown in Fig. 1.1). Discussion of the DBA function is deferred to Section 1.5.

This completes our preliminary description of the major aspects of a database system. The rest of this chapter will consider these topics in somewhat more detail.

1.2

1.2 OPERAT

In one of the ea in a database a put data, and Engles's origin

- A databas cation syst

This defin convenient gen scientific, techn

- Manufactu
- Bank,
- Hospital,
- University,
- Government

Any enterp tion. This is its listed above wo

- Product da
- Account da
- Patient dat
- Student da
- Planning d

As already men data, work que data" refers to (typically on ce change to be ma base. Similarly, from the system report contains self part of the

As an illust case of a manuf will wish to reta

GOG 032269