

UNIX[®]
Relational
Database
Management

Application
Development
in the UNIX[®]
Environment

ROD MANIS
EVAN SCHAFFER
ROBERT JØRGENSEN

UNIX[®]
Relational Database Management

**Application Development
in the
UNIX Environment**

**Rod Manis
Evan Schaffer
Robert Jørgensen**



PRENTICE HALL
Englewood Cliffs, New Jersey 07632

Library of Congress Cataloging-in-Publication Number 87-36103

Cover design: Lundgren Graphics, Inc.
Manufacturing buyer: Mary Ann Gloriande

© 1988 by Rod Manis

This book was typeset by the authors. UNIX programs formatted the text and printed the final copy. We acknowledge the special assistance of Barbara Wright, who is responsible for the description of ve in Chapter 5.

Prentice Hall Software Series
Brian W. Kernighan, Advisor

The authors and publisher of this book have used their best efforts in preparing this book. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness. The author and publisher make no warranty of any kind, expressed or implied, with regard to these programs or the documentation contained in this book. The author and publisher shall not be liable in any event for incidental or consequential damages in connection with, or arising out of, the furnishing, performance, or use of these programs.

The publisher offers discount on this book when ordered in bulk quantities. For more information, write:

Special Sales/College Marketing
College Technical and Reference Division
Prentice Hall
Englewood Cliffs, New Jersey 07632

/rdb and /act are trademarks of Rod Manis
/menu shell is a trademark of Schmidt Associates
Chiquita Banana is a trademark of United Brands
Ronald McDonald is a trademark of McDonalds
UNIX is a registered trademark of AT&T Bell Laboratories
ve is a trademark of Robinson Schaffer Wright

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 0-13-938622-X

PRENTICE-HALL INTERNATIONAL (UK) LIMITED, *London*
PRENTICE-HALL OF AUSTRALIA PTY. LIMITED, *Sydney*
PRENTICE-HALL CANADA INC., *Toronto*
PRENTICE-HALL HISPANOAMERICANA, S.A., *Mexico*
PRENTICE-HALL OF INDIA PRIVATE LIMITED, *New Delhi*
PRENTICE-HALL OF JAPAN, INC., *Tokyo*
SIMON & SCHUSTER ASIA PTE. LTD., *Singapore*
EDITORIA PRENTICE-HALL DO BRASIL, LTDA., *Rio de Janeiro*

UNIX[®]
Relational Database Management

YES—I'd like \rdb for my UNIX system

Cut out this form and fill in all necessary information. Then enclose this form in an envelope and mail to:

Robinson Schaffer Wright
711 California Street
Santa Cruz, CA 95060
(408) 429-6229

Please send me pricing information for \rdb for my UNIX system. I understand that \rdb runs on *all* UNIX based computers, micro to mainframe, and most UNIX derivatives (XENIX, Microport System V/AT, 4.3 BSD, AIX, ULTRIX, etc.). Prices for \rdb vary by size of system from \$795 to \$2,495.

My computer is _____, it runs the following version of UNIX _____, and supports _____ users.

Name _____
Firm/Department _____
City _____ State _____ Zip _____
Daytime telephone number(_____) _____

NOTE: Prices subject to change without notice.

YES—I'd like /rdb and/or MKS Toolkit for DOS for my PC-DOS or MS-DOS system

The MKS Toolkit provides over 100 UNIX like commands and is required to take advantage of /rdb in a DOS environment (version 2.0 and later releases).

Cut out this form and fill in all necessary information. Then enclose this form with your check or money order *only* in an envelope and mail to:

Please send the items checked below. **PAYMENT ENCLOSED** (check or money order *only*). Robinson Schaffer Wright will pay all shipping and handling charges.

_____ /rdb and MKS Toolkit for DOS—\$275
_____ /rdb for DOS only. I have MKS Toolkit—\$139
_____ MKS Toolkit for DOS only (I don't need /rdb yet)—\$139
California residents please add 7% sales tax

Robinson Schaffer Wright
711 California Street
Santa Cruz, CA 95060
(408) 429-6229

Name _____
Firm/Department _____
City _____ State _____ Zip _____
Daytime telephone number (_____) _____

NOTE: Prices subject to change without notice.

This book is dedicated to the people who have made
/rdb
possible

Thomas Arnow
Robert Berkley
Andy Chang
Michael J. Defazio
Steve Eberhart
Keith Eisenstark
Raymond Eisenstark
David Fiedler
Tom Johnson
Ray Jones
Brian Kernighan
Paul S. Kleppner
Mark Krieger
Gig Graham
William Hamilton
Joel Harrison
Beth Harsaghy
Chris Lynch
John Mashey
Marc Meyer
Phoebe Miller
Eli Nielsen
Fred Pack
Durk Pearson
Ron Posner
Pat Rafter
William Robinson
Pat Sarma
David Schmidt
Tom Slezak
William P. Spencer
Shawn Steel
Ed Taylor
Roy Takai
Zhang Yuchao
Richard Vento
Peter Vizel
Jerry Walker
Colin Watanabee
Barbara Wright
Myron Zimmerman

Preface

This book shows you how to develop software applications in a UNIX environment, in a small fraction of the time and effort needed with other systems. The UNIX environment is so different that it makes possible a new approach to computing.

There are many books available on writing C programs, administering a UNIX system, and specific software like spreadsheets and word processors. But this book tells you how to put them all together. System integration is the hardest part of software development. Therefore, we have filled this book with examples, solutions, source code, tricks, hints, and useable theory to help you.

Each of the authors has over 20 years of experience in computing in which we have had to put together an enormous number of applications on many different systems for a wide range of users. We have tried to put as much of what we have learned into this book as possible.

We emphasize database management because almost all applications require data to be entered, manipulated, searched, or reported. By learning simple database operations, you can approach any application with a powerful tool kit.

Model Business System

We have even included, as an example to get you started, a model business system with general ledger, tax, accounts receivable and payable, inventory, payroll, and operations. Included is a tutorial, manual pages, and source code of the UNIX shell scripts. They show you how to use the system, and also how to manage a business with a computer.

New Approach to Relational Database Management

Most books on relational database theory take a very mathematical approach and use a difficult language. They talk of relations, tuples, attributes, degrees, cardinality, and normalization. We've translated these terms into familiar words like tables, rows, and columns. We also show a new approach to normalizing tables which anyone can understand. We have tried to make relational database theory accessible to everyone.

The UNIX System

The UNIX system, of course, is among the greatest software tools available. It is portable, multi-user, and multi-tasking. It runs on hundreds of computers, from the IBM PC/XT to the Cray 2. It has been adopted as the standard operating system by the U.S. federal government (POSNIX), and many corporations and universities. But the most important features of the UNIX system are the UNIX shell, the software tools, and the *pipe and filter* approach.

To get the full power of the UNIX system, one must relearn and rethink how to do software development. Most software people are not aware of what they can do with the UNIX system. They use the same approach that they learned on other systems. It is as if they had a Porsche, and hitched a horse to it, not knowing that they could start its engine and roar away. In this book we will teach you how to drive, but you may have to give up your attachment to old fashioned horse and buggy styles of programming.

UNIX Shell is the Best Fourth Generation Programming Language

It is the UNIX shell that makes it possible to do applications in a small fraction of the code and time it takes in third generation languages. In the shell you process whole files at a time, instead of only a line at a time. And, a line of code in the UNIX shell is one or more programs, which do more than pages of instructions in a 3GL. Applications can be developed in hours and days, rather than months and years with traditional systems. Most of the other 4GLs available today look more like COBOL or RPG, the most tedious of the third generation languages.

Flat File Database

There are two approaches to database management. Traditionally, a large database management program had to be written to run on the old mainframe and mini computer operating systems which provided only a minimum of assistance. Therefore, all of the functions needed had to be coded. Each system had its own language and data formats. They seldom talked to each other. It was hard to pass data from one to the other. It took a lot of work to learn and code them. Then when you moved to another system, you had to learn a new language and figure out how to transfer and convert your data and rewrite your programs.

UNIX solved most of these problems by providing a huge set of tools and facilities and the ability to link them together with pipes and shell programs. Therefore, the burden on a database manager is significantly reduced. To take advantage of the UNIX environment, *flat file* database management systems were developed, including */rdb* (pronounced slash-r-d-b) from Robinson Schaffer Wright, *Prelude* from VenturCom, *Unity* from AT&T, among others. These database managers work with UNIX and keep their data in flat ASCII UNIX files. You can mix UNIX and database commands in the same shell scripts and pipe commands. Therefore, you do not lose the power of UNIX by going into the database program. You use your UNIX knowledge, and/or learn UNIX skills once. Then you can move to most computers without having to learn *yet another system and language*.

***/rdb* Relational Database Management System for UNIX**

We use */rdb* for the examples in this book because it is so simple to learn and use, and does not get in the way of your understanding the principles of UNIX database management and applications development. It makes a good teaching aid. Examples with */rdb* show you how to do basic tasks, without having to spend time learning complex systems. So even if you choose other software packages, you have learned the principles in the easiest way. You can then focus on the more cumbersome syntax of other software packages. */rdb* is also one of the least expensive of the database managers and is available on any UNIX computer.

At the end of this book is a manual of the */rdb* commands including examples of

their use, and the source code of all of the UNIX shell programs. They will give you many ideas about how to handle specific problems in more detail than is appropriate for the tutorial part of this book.

UNIX and DOS

There is a joke going around that in the future there will be only two operating systems, *UNIX running DOS* and *DOS running UNIX*. It has already happened. The latest UNIX operating systems run DOS applications. And you can get the MKS Toolkit of over 100 UNIX commands, including the latest UNIX korn shell, vi, awk, grep, etc., to run on your DOS. `\rdb` (pronounced backslash-r-d-b), is the DOS version of `/rdb`. Therefore, most of the programs described in this book will also run on DOS (MS-DOS, PC-DOS, OS/2) with the MKS Toolkit and `\rdb`. See the tear out card in this book, which offers a discount on these packages.

Since UNIX and DOS run on most of the world's computers, you can develop an application on one computer and run it on many others. You can also network computers together and run the same programs on each, transferring data in a standard format. It is an old dream come true. But there are lots of details to make it all work, which is what this book is about.

Macintosh

We see no conflict between the UNIX/DOS approach and the Macintosh approach. In an integrated computing environment, we'll be able to 'point and click' on, for example, a shell script, or any arbitrary collection of commands. The concept of relations as flat files and relational operators as programs is the common ground.

Required Knowledge

You should know how to log on to UNIX, to use its commands, to use a text editor like vi. You should be aware of the method of using UNIX programs, often piped together, to get things done without having to write C or other language programs. You do not have to know how to program, although it helps.

There are many introductory books on UNIX and a few books on advanced UNIX and shell programming: *The UNIXTM Shell Programming Language* by Rod Manis and Marc H. Meyer [Manis 1986]. For more advanced UNIX, you should read Stephen Prata's *Advanced UNIXTM - A Programmer's Guide* [Prata 1985], and Brian Kernighan and Rob Pike's book *The UNIXTM Programming Environment* [Kernighan 1984], and the new Kernighan book on `awk`. See the bibliography for more complete citations.

This book starts where those leave off, and covers the details of putting together real applications with databases. One or more of these books should be read because they teach the UNIX and shell programming basics that we build upon. We assume that basic knowledge so that we can focus on application development. But even if you are new to UNIX, you should try to get started. If you get stuck, you can look up ideas in the other books and in your system's UNIX documentation.

To help overcome the old beliefs, we show how to develop and use software applications in the UNIX environment, covering all of the important tasks of programming and setting up computer systems to do real work for users. A beginner will learn the

basics of database design and fourth generation programming. Experienced programmers will learn how to work in this new environment. Many UNIX gurus understand in principle that UNIX is a far better environment, but need to know more of the details of how to handle certain problems.

Computer Revolution

This book aims at nothing less than a revolution in computing! It challenges some of the most encrusted assumptions and proposes a simple, elegant approach which solves many of the major problems in computing today.

1. Data should be kept in flat ASCII tables (files), not binary, so that we can always see what we are doing, and do not have to depend upon some special program to decode our data for us.
2. Programming should be done in a fourth generation programming language, the most powerful of which is the UNIX shell, not in lower level languages, except in extreme cases.
3. Programs should be small and should pass data on to other programs. Software prisons, or large programs with self-contained environments, must be avoided because they require learning and they make extracting data difficult.
4. All programs should be integrated with a common interface with both users and data. All of our tools should look more or less the same for quick learning and easy use.
5. We should build software and systems to meet interface standards so that we can share software and stop dreaming that any individual or company can do it all from scratch.
6. We should use networks to tie together computers and software so that simple tables of data can be passed from one to another.
7. Industrial strength artificial intelligence and expert systems will have to be developed in the same environment previously described, and not in toy languages like LISP and PROLOG.
8. The UNIX system embodies the most advanced facilities for implementing these principles.

*Rod Manis
Evan Schaffer
Robert Jørgensen*

Foreword

I stumbled into using UNIX in late 1978, in a Bio-Medical Research environment. The machine was a PDP 11-34 running V6 UNIX, with custom automated microscopy gear used for image processing of mammalian cells. In those days there was the Kernighan and Ritchie C book and the UNIX manuals themselves; nothing more. Two degrees in computer science prepared me to tackle UNIX but I was still perhaps understandably shy about investigating all of the 100+ tools, including such strange ones as **awk**, **lex**, and **yacc**.

I quickly discovered that the use of the UNIX tools freed me from the great bulk of the drudgery of supporting real end users on a computer system. I could string tools together quickly instead of writing small *one shot* programs to do the myriad of data manipulation tasks that characterize real life in the scientific computing world. In the course of my work I acquired a relational database package, but found that it did not mesh well with the UNIX tool-kit philosophy. Like the mammoth database systems I had previously used on large IBM systems, one had to *drop into* the database package, where a whole new command language applied and the UNIX tools could only be reached with extreme difficulty, if at all.

By 1983 I had added an early 16-bit microprocessor-based UNIX timesharing system running V7 UNIX to my computer shop. I had a growing number of users who were doing general purpose data processing and needed a variety of tools to handle their needs. I was fortunate enough to be in attendance at the Usenix conference where Rod Manis first described his **/rdb** database, that existed as a set of UNIX filters. As he presented his paper I realized, along with numerous other members of the audience, that Rod had grasped and implemented the central abstraction I had been toying with but never formalized: a UNIX database should use ASCII files and operate at the shell level as a series of filters. This simple but brilliant insight of Rod's, brought about by his mastery of a UNIX tool (**awk**) that I had ignored to that point, provided my researchers with the single most powerful tool I have been able to acquire to this date. Along with many others who remained after his talk to discuss his ideas, I was one of the first users of the **/rdb** package. The medical research environment is characterized by chronic lack of funding for computing equipment and scarcity of programming expertise. Without mega-funding, I had no choice but to find ways to let end users solve their own problems as much as possible, without making them all *programmers*.

To my mind, the power of this tool is the freedom it gives to the end users. For the first time I could sit down with a novice and create a working sample database using real data in a matter of minutes. Queries could be demonstrated and stored in *user-friendly* shell scripts. Within hours the end users were fully capable of creating their own databases. Their creativity was amazing: One researcher who had had a single Fortran class eight years earlier wrote a two-page shell script using **/rdb** filters to

completely computerize her laboratory data processing and analysis. Other users discovered that the innocent-appearing simple report writer was in reality a powerful meta-tool that allowed them to write shell scripts that wrote shell scripts, yielding in effect a *two-dimensional* program. By word of mouth it was demanded to be bought for at least a dozen other machines at this site. The database for the world's largest melanoma (skin cancer) epidemiology study runs on **/rdb**, as do many other applications.

I have found this abstraction of a database (regular ASCII files, filter programs, normal file access) to serve well for moderate size databases and the frequency of update and query that are common in the scientific research fields. Like most UNIX utilities and indeed UNIX itself, it is always possible to write a single-point solution that would be faster. I have found, however, that the ability to apply the full power of the UNIX toolkit to *database problems* far outweighs any speed penalties that I might be paying on the current generation of super-micros and super-minis. Even in situations where a "TRADITIONAL DATABASE" is required, the rapid prototyping of the **/rdb** approach is often used to provide the early insights into the proper way to tackle problems requiring exceptional speed or size. I find that writing a custom program to manipulate or calculate data is virtually unnecessary, since such problems can nearly always be solved with the **/rdb** database and regular UNIX filters.

A completely unexpected benefit of this tool-kit approach to databases has been the education of users, many of whom have had no prior computer training or exposure. The early confidence gained in putting up their own simple databases often led to users taking the plunge and developing more complex shell scripts using those tools. In numerous cases, their continued interest led to them delving into **awk** and writing custom front or back end programs. A few adventurous souls then began writing C code in the **awk** scripts, as well as learning to use **lex**, which after all is merely a C program-generator cleverly (?) disguised as a lexical-analyzer. The net result is that I have had several novice users bootstrap themselves up to be excellent applications programmers in C in an amazingly short time, all due to the good first experience with the database tools!

The evolution of personal computers into machines with sufficient power to run UNIX well (IBM AT, Mac II, etc.) puts the ability to use the UNIX tool-kit philosophy into the reach of every research lab and small business. A book like this is as much a guide to using UNIX itself and the UNIX philosophy of problem solving as it is to being a guide to a specific database. The authors have quietly effected a revolution in the use of databases as a problem-solving tool that is as startling in its clarity and simplicity as the development of the spreadsheet program.

Tom Slezak
Computer Scientist, UNIX Support
Lawrence Livermore National Laboratory

Contents

Preface		xiii
Foreword		xvii
Chapter 1. Introduction		1
1.1	Overview: Relational Database and UNIX	1
1.2	Operating Systems that Support UNIX Style Environments	1
1.3	Applications Development	1
1.4	Databases	2
1.5	UNIX Environment	2
1.6	Fourth Generation Programming Language	2
1.7	UNIX Shell Programming	2
1.8	Assemble Software Modules	3
1.9	Fast Prototyping	3
Chapter 2. Tables: Where to store information		5
2.1	A World of Tables	5
2.2	Files	5
2.3	Column	6
2.4	Row	6
2.5	Column and Row	7
2.6	Shell Programming and Flat ASCII Files	7
2.7	Lists	7
Chapter 3. Query and Report: How to get data out		9
3.1	column	10
3.2	row	10
3.3	Quotes	11
3.4	compute	12
3.5	justify	13
3.6	total	14
3.7	subtotal	14
3.8	Pipes	15
3.9	Syntax: How to enter commands right	16
3.10	UNIX Shell Scripts	18
3.11	Report Writing: How to get reports from your data	19
3.12	Conclusion	21

Chapter 4.	Data Entry: How to put data in	23
4.1	ve	23
4.2	Text Editor or Word Processor	24
4.3	Enter	24
4.4	Update	25
4.5	Programs and Other Formats	25
4.6	Rules for Table and List File Setup	25
4.7	Tab Problems	26
4.8	Seeing the Tabs	27
4.9	Table Width Problems	27
4.10	Special Characters	28
4.11	Data Validation: How to get data right	28
Chapter 5.	ve Form and Screen Database Editor	31
5.1	File Structure	31
5.2	The Command Line	32
5.3	The Screen File	37
5.4	The Header File	40
5.5	The Audit File	40
5.6	The Validation File	41
5.7	Commands	45
5.8	Multi-User Considerations	56
5.9	/rdb Compatability	56
5.10	Limits	57
Chapter 6.	Database Design: How to set it up	59
6.1	One-to-one Relationships in One Table	59
6.2	One-to-many Relationships in Two Tables	59
6.3	Many-to-many Relationships in Three Tables	60
6.4	Planning	61
6.5	Normalization	63
6.6	Functional Dependency	63
6.7	Keys	63
6.8	Universal Relation	64
6.9	Problems	64
6.10	First Normal Form	64
6.11	Second Normal Form	65
6.12	Third Normal Form	66
6.13	Normalizing Example	66
6.14	Complex Queries with Joins	68
6.15	Pipeline Join	69
6.16	Phone Book	71

Chapter 7.	Shell Programming: Let UNIX do the work	73
7.1	Database Programming in UNIX Shell Language	73
7.2	UNIX Utilities	74
7.3	awk - Language to Produce Complex Reports	74
7.4	cat - Display a Table or List File	74
7.5	grep - Find All Rows That Contain a Matching String	74
7.6	od -c - Octal Dump All Bytes as Characters	74
7.7	sed - Stream Editor to Edit File in a Pipe	74
7.8	sh - UNIX Shell Programming Language	75
7.9	spell - Check Spelling in a Table or List File	75
7.10	tail - Display Bottom Rows of a Table or List File	75
7.11	vi - Text Editor to Enter and Update Files	75
7.12	Reading and Writing Database Files	75
7.13	Parsing Rows	76
7.14	Tables to Shell Variables	76
7.15	Lists to Shell Variables	77
7.16	4GL	79
Chapter 8.	Menus: One from column A and ...	83
8.1	Shell Menu	83
8.2	Screen Menu Example	83
8.3	case Actions	84
8.4	termput and tput Commands	85
8.5	clear Command	85
8.6	cursor Command	85
8.7	Help Facility: How to get help	86
Chapter 9.	Screen Form Form and more forms	87
9.1	Building a Screen Form	87
9.2	screen Command	87
9.3	update.inv: Advanced Screen Program	88
Chapter 10.	Fast Access Methods: How to get data fast	91
10.1	When Appropriate	91
10.2	index	91
10.3	search	92
10.4	Multi-Rows, Multi-Columns, and Multi-Keys	93
10.5	Methods	93
10.6	Analysis	97
10.7	Management	97
Chapter 11.	Swiss Army Knife: A program for all reasons	99
11.1	Record Locking: It's Mine	99
11.2	Dates: Conversion and Math	100
11.3	Set Theory Commands	103

Chapter 12.	How to Solve Traditional Database Problems	105
12.1	Multi-User Concurrent Access to Files	105
12.2	Screen Form Entry	105
12.3	Security	106
12.4	Distributed Data	106
12.5	Backup	106
12.6	Checkpoint and Recovery	106
12.7	Validation	106
12.8	Audit Trails and Logging	106
Chapter 13.	Comparing Systems: Who's on First?	109
13.1	Resource Use	109
13.2	C Programming Unnecessary	109
13.3	Speed	110
13.4	Size	111
Chapter 14.	Database Theory: Ways of thinking about data	113
14.1	Can You Say that in English?	113
14.2	Database Models	113
14.3	Hierarchical	114
14.4	Network	114
14.5	Relational	115
14.6	Entity-Relationship	116
14.7	Binary	117
14.8	Semantic Network	117
14.9	Infological	117
14.10	Prolog: Programming in Logic	117
14.11	The Grand Unified Field Theory of Information	117
Chapter 15.	AI and Prolog: Programming in Logic	119
15.1	Prolog Language and Environment	119
15.2	/rdb Interface to Prolog	121
15.3	Problems of Prolog	122
15.4	searchtree: Database Tree Searching	123

Chapter 16.	C Interface: Speak C to the database	125
16.1	Don't	125
16.2	System	125
16.3	Execl	126
16.4	Execl Shell Programs	127
16.5	Fork	128
16.6	One-Way Pipe	129
16.7	Pipe to Standard-in	131
16.8	Two-Way Pipe	132
16.9	Programming Style	135
16.10	Fast Access	135
16.11	tabletostruct	138
16.12	Read Table into Memory: getfile and fsize	140
16.13	/rdb Functions: librdb.a	142
16.14	Colroutines	142
16.15	Display Example	142
16.16	Debugging	145
Chapter 17.	Converting from Other Databases to /rdb	149
17.1	SQL: How to talk SQL to /rdb	149
17.2	dBASE: How to talk dBASE to /rdb	149
17.3	R:base: How to talk R:base to UNIX and /rdb	152
Chapter 18.	Installation: How to set up /rdb	157
18.1	Step 1: tar from Tape, Floppy or Cassette	157
18.2	Step 2: Change Path Name	158
18.3	Step 3: Set RDB Environment Variable	159
18.4	Step 4: Source Installation (If you have source code)	160
18.5	Summary	160
Chapter 19.	/rdb Commands	163
19.1	Features	163
19.2	Most Important Commands	164
19.3	Commands Grouped by Function	164