

THE EXPERT'S VOICE® IN MOBILE PROGRAMMING

Introducing SQLite for Mobile Developers

Enabling Database Functionality
for Android and iPhone

—
Jesse Feiler

Apress®

Introducing SQLite for Mobile Developers



Jesse Feiler

Apress®

Introducing SQLite for Mobile Developers

Copyright © 2015 by Jesse Feiler

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

ISBN-13 (pbk): 978-1-4842-1765-8

ISBN-13 (electronic): 978-1-4842-1766-5

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Managing Director: Welmoed Spahr

Lead Editor: Jeffrey Pepper

Technical Reviewers: Aaron Crabtree and Cliff Wootton

Editorial Board: Steve Anglin, Pramila Balan, Louise Corrigan, Jonathan Gennick,
Robert Hutchinson, Celestin Suresh John, Michelle Lowman, James Markham,
Susan McDermott, Matthew Moodie, Jeffrey Pepper, Douglas Pundick,
Ben Renow-Clarke, Gwenan Spearing

Coordinating Editor: Mark Powers

Copy Editor: Lori Jacobs

Compositor: SPi Global

Indexer: SPi Global

Artist: SPi Global

Distributed to the book trade worldwide by Springer Science+Business Media New York, 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springer.com. Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc. is a Delaware corporation.

For information on translations, please e-mail rights@apress.com, or visit www.apress.com.

Apress and friends of ED books may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Special Bulk Sales–eBook Licensing web page at www.apress.com/bulk-sales.

Any source code or other supplementary materials referenced by the author in this text is available to readers at www.apress.com/9781484217658. For detailed information about how to locate your book's source code, go to www.apress.com/source-code/. Readers can also access source code at SpringerLink in the Supplementary Material section for each chapter.

Contents at a Glance

About the Author	xi
About the Technical Reviewers	xiii
Acknowledgments	xv
Introduction	xvii
■ Chapter 1: Getting Up to Speed with Databases and SQLite.....	1
■ Chapter 2: Understanding What SQLite Is.....	9
■ Chapter 3: Using SQLite Basics: Storing and Retrieving Data	15
■ Chapter 4: Working with the Relational Model and SQLite.....	29
■ Chapter 5: Using SQLite Features—What You Can Do with SELECT Statements.....	39
■ Chapter 6: Using SQLite with PHP	45
■ Chapter 7: Using SQLite with Android/Java.....	55
■ Chapter 8: Using SQLite with Core Data (iOS and OS X)	61
■ Chapter 9: Using SQLite/Core Data with Swift (iOS and OS X).....	75
■ Chapter 10: Using SQLite/Core Data with Objective-C (iOS and Mac)	97
■ Chapter 11: Using the Simple Database with a PHP Web Site....	121
■ Chapter 12: Using the Simple Database with a Core Data/iOS App	135
Index.....	145

Contents

About the Author	xi
About the Technical Reviewers	xiii
Acknowledgments	xv
Introduction	xvii
■ Chapter 1: Getting Up to Speed with Databases and SQLite.....	1
Moving Beyond <i>Big</i>	1
Databases Are Structured and Organized.....	2
Databases Are Smart.....	2
Relational Databases and SQL to the Rescue	4
Looking Inside a Relational Table and Query.....	5
Basic Query Structure.....	6
Looking at Other Query Choices	7
■ Chapter 2: Understanding What SQLite Is.....	9
Putting a Database in Perspective	9
Defining SQLite.....	10
SQLite Is Designed for a Single User	11
SQLite Is Self-Contained.....	12
SQLite Supports Transactions and Is ACID-Compliant	13

■ **Chapter 3: Using SQLite Basics: Storing and Retrieving Data 15**

- Using `sqlite3`..... 16
 - Run `sqlite3` and Let It Create a New Database..... 16
 - Create and Name a New `sqlite3` Database 17
 - Delete the Database 17
 - Run `sqlite3` and Open an Existing Database 17
- Experimenting with SQLite Syntax 18
- Exploring Your `sqlite3` Database with a Graphical SQLite Editor 19
- Creating a Table..... 21
 - Using a Graphical SQLite Editor 21
 - Creating Table Columns 22
 - Using SQLite3 24
- Inserting Data into a Table..... 24
 - Using a Graphical User Interface 24
 - Using SQLite3 26
- Retrieving Data..... 26
 - Using a Graphical User Interface 26
 - Using `sqlite3`..... 27
- Deleting Data..... 27
- Summary 27

■ **Chapter 4: Working with the Relational Model and SQLite..... 29**

- Building the Users Table..... 30
- Building the Scores Table 31
- Relating the Tables 32
 - Using Aliases to Identify Multiple Tables in a SELECT Statement..... 32
 - Using the `rowid` Primary Key..... 33
 - Changing a Name in One Table..... 34

Using a Foreign Key.....	34
Joining the Tables	37
Summary.....	38
■ Chapter 5: Using SQLite Features—What You Can Do with SELECT Statements.....	39
Looking at the Test Data.....	40
Ordering Data Makes It Easier to Use.....	40
Grouping Data Can Consolidate It.....	41
Using Variables in Queries.....	42
Summary.....	43
■ Chapter 6: Using SQLite with PHP	45
Putting PHP and SQLite Together: The Basics	46
Verifying PHP in Your Environment	46
Preparing the SQLite Database.....	47
Connecting to Your SQLite Database	50
1. Create a New PDO Object	51
2. Create and Prepare the Query	52
3. Execute the Query	52
4. Fetch the Results.....	52
5. Use the Results.....	52
Summary.....	54
■ Chapter 7: Using SQLite with Android/Java.....	55
Integrating SQLite with Any Operating System, Framework, or Language	55
Using Android and SQLite.....	57
Using the Static Values.....	57
Extend SQLiteOpenHelper.....	58
Summary.....	60

- **Chapter 8: Using SQLite with Core Data (iOS and OS X) 61**
 - Introducing the Core Data Framework 62
 - Using the Core Data Model Editor 63
 - Using Entities..... 66
 - Working with Attributes 68
 - Managing Relationships 69
 - Summary 73
- **Chapter 9: Using SQLite/Core Data with Swift (iOS and OS X)..... 75**
 - Looking at the Core Data Stack..... 75
 - Fetching Data to the Core Data Stack 76
 - Structuring a Core Data App..... 76
 - Passing a Managed Object Context to a View Controller in iOS..... 77
 - Setting Up the Core Data Stack in AppDelegate for iOS 78
 - Setting Up the Core Data Stack in AppDelegate for OS X 81
 - Creating a Fetch Request in iOS..... 84
 - Saving the Managed Object Context 85
 - Saving in iOS 85
 - Saving in OS X 86
 - Working with NSManagedObjectContext 87
 - Creating a New NSManagedObjectContext Instance 88
 - Working with a Subclass of NSManagedObjectContext 90
 - Summary 95
- **Chapter 10: Using SQLite/Core Data with Objective-C (iOS and Mac) 97**
 - Looking at the Core Data Stack..... 98
 - Fetching Data to the Core Data Stack 99
 - Objective-C Highlights..... 99

Using Quoted Strings	99
Objective-C Is a Messaging Language	99
Using Brackets in Objective-C	99
Chaining Messages	100
Ending Statements with a Semicolon.....	100
Separating Headers and Bodies in Objective-C.....	100
Looking at Method Declarations.....	101
Handling nil in Objective-C.....	101
Structuring a Core Data App with Objective-C	102
Passing a Managed Object Context to a View Controller in iOS.....	102
Setting up the Core Data Stack in AppDelegate for iOS.....	103
Setting Up the Core Data Stack in AppDelegate for OS X	107
Creating a Fetch Request in iOS.....	110
Saving the Managed Object Context	111
Saving in iOS	111
Saving in OS X	112
Summary.....	120
■ Chapter 11: Using the Simple Database with a PHP Web Site....	121
Reviewing the Database.....	121
Previewing the Web Site	124
Implementing the PHP Web Site.....	127
Looking at the Basic PHP/SQLite Structure	128
Building the Drop-Down Selection List (phpsql1.php).....	130
Showing the Selected Data (phpsql2.php)	131
Adding New Data (phpsql3.php).....	133
Using Try/Catch Blocks with PHP and PDO.....	134
Summary.....	134

- Chapter 12: Using the Simple Database with a Core Data/iOS App 135**
- The Story Continues... 135
- Adjusting the Data Model and Template for Core Data..... 136
 - Getting Rid of Keys and Revising the Data Model 137
 - Changing timeStamp to name 137
 - Create a New Database on Your Device or Simulator 138
- Add the Score Table and Interface to the App 138
 - Making Sure You Can Add New Users with + in the Master View Controller 138
 - Working with the Detail View..... 139
- Working with the Detail View for Score..... 139
 - Use NSManagedObject Subclasses 141
 - Use a Table View Controller for DetailViewController 141
 - Modify DetailViewController Code for DetailViewController 142
 - Modify MasterViewController to Pass the User to DetailViewController..... 143
- Summary 144
- Index..... 145**

About the Author



Jesse Feiler is a developer, consultant, and author specializing in database technologies and location-based apps. He has worked with databases and data management on computers from mainframes to iPhone, iPad, and Apple TV using data management tools from DB2 (IBM) and DMSII (Burroughs) to Enterprise Objects Framework and Core Data, MySQL, Oracle, and, of course, SQLite.

In the early days of the web, he built the page caching mechanism for the Prodigy web browser for Mac using a relational database library similar in some ways to SQLite.

He is the creator of Minutes Machine the meeting management app, as well as Saranac River Trail app a guide to the Trail that includes location-based updates as well as social media tools. His apps are available in the App Store and are published by Champlain Arts Corp (champlainarts.com). As a consultant, he has worked with small businesses and nonprofits on projects such as production control, publishing, marketing and project management usually involving FileMaker and other databases.

His books include:

iOS Programming with Swift for Dummies (Wiley, 2015)

Swift for Dummies (Wiley, 2015)

iOS App Development for Dummies (Wiley, 2014)

iWork for Dummies (Wiley, 2012)

Videos include:

Mixed Language App Development with Objective-C and Swift (O'Reilly, 2015)

iOS Developer's Guide to Views and View Controller (O'Reilly, 2015)

Learning Objective-C Programming (O'Reilly, 2015)

He is heard regularly on WAMC Public Radio for the Northeast's The Roundtable, and is the founder of Friends of Saranac River Trail. A native of Washington DC, he has lived in New York City and currently lives in Plattsburgh NY.

He can be reached at northcountryconsulting.com (consulting) and champlainarts.com (app development).

About the Technical Reviewers



Trained in computer engineering, multimedia, and graphic design, **Aaron Crabtree** has spent the last 15 years working in the industry. Experience has taught Aaron that, in most cases, programmatically creating each piece of an application is necessary for greater flexibility and control. Drop him a line on Twitter: @aaron_crabtree



Cliff Wootton is a former Interactive TV systems architect at BBC News. The “News Loops” service developed there was nominated for a BAFTA and won a Royal Television Society Award for Technical Innovation. An invited speaker on pre-processing for video compression at the Apple WWDC conference. Taught post graduate MA students about real-world computing, multimedia, video compression, metadata and researching the deployment of next generation interactive TV systems based on open standards.

Currently working on R&D projects investigating new Interactive TV technologies, involved with MPEG standards working groups, writing more books on the topic and speaking at conferences when not lecturing on Multimedia at the University of the Arts in London.

Acknowledgments

What I love most about working with databases is putting the pieces of the puzzle together – what a friend of mine used to call “making order out of chaos.” When it comes to thinking about SQLite, there are many pieces to put together – everything from the first work on relational databases back in the 1960s to the idea of lightweight code libraries that can fit appropriately onto heavy-duty computers, as well as small battery-powered mobile phones and autonomous satellites designed to explore other worlds far from the relative safety of their home planet.

This book has been made possible, first and foremost, by D. Richard Hipp who began (and continues) the SQLite project starting in 2000. The many people who have added and tested code are too many to thank, but you know who you are.

Closer to home, Jeff Pepper at Apress has been great to work with (again), as has Mark Powers. Together with Carole Jelen at Waterside Productions, they have all made it possible to provide this basic introduction to SQLite.

Introduction

I have worked with databases of many types on many platforms. Despite the fact that almost all of them have been based on the relational model. I've had some excursions into ancient database structures such as IMS (a hierarchical database) and, more recently into unstructured data, but in every case it has been a matter of a basic principle: finding the simplest way to organize data so that it makes sense. (As I often say to clients, "The data doesn't lie." When you lay data out this way, that way, and every which way, I've found that gradually its inherent structure becomes clear - if there is one. If there isn't an inherent structure, what may become clear is some structure no one has ever thought of for the data. And sometimes, the data stubbornly refuses to reveal a structure that we can understand.

This may sound very highfalutin and esoteric, but I do think that our job as database designers ultimately becomes a job of finding patterns. If we cannot find a logically inherent and inevitable pattern, the best choice is to find the best pattern that is usable for the task at hand.

We can't produce these data models without knowing something (a lot, in fact) about the ways to organize data, and, today that comes down to the relational model and the tools of SQL. This process is iterative (woe betide the design team that adopts a partial database model too early), and it relies on flexible tools as well as imagination and constant questioning of everyone working on a project.

SQLite has become a critical component of many data modeling projects because it is so lightweight and flexible. You can put up a SQLite database with basic functionality very quickly. That's where this book comes in. If you have the data and the mission to come up with something that uses the data - website, mobile app, or standalone software product - SQLite can be just what you want.

SQLite runs on many operating systems. With SQLite, you're working with a database that you can move around from one platform to another and one development environment to another. Your prototyping, experimentation, and development can proceed without committing yourself to anything besides the relational model and SQLite. You may move on to other relational tools, and your little SQLite app may turn into a production app for a multitude of devices or for a single many-multi-threaded corporate data center.

This flexibility and very low barrier to entry are two of the features that many people including me find attractive in SQLite.

I hope that you'll find this book useful. There's more information on SQLite at sqlite.org, and there's more information (and downloadable examples from the book) on my website at northcountryconsulting.com.

—Jesse Feiler

CHAPTER 1



Getting Up to Speed with Databases and SQLite

Ask people to tell you words they associate with *database* and you'll probably get *big* among the responses. Databases handle large amounts of data, and everyone knows that. How big? It depends whether you count the number of items in a database or the size of the items. The Library of Congress has 160 million items on 858 miles of shelving as of this writing. The catalog, which is based on a database, is available online at www.loc.gov/about/fascinating-facts/. There's no question among people who are interested in databases that the Library of Congress catalog is far from the largest database in the world. (In all likelihood, the largest databases are not visible to the public because they contain classified corporate and governmental information.)

This chapter provides an introduction or reminder about databases today and how they are used. If you've used databases in the past, much has changed, and if you haven't used them much in the past, this chapter will give you a quick overview.

Moving Beyond *Big*

Beyond *big*, you should start thinking about databases as being *structured* and *organized*. In fact, as a mobile developer, the structured and organized aspects of databases are much more important than their size. From the earliest days of the Web in the early 1990s, web browsers have used databases to store and organize their data. What data, you may wonder? How do you think your browser is able to store your passwords for the web sites you visit? How do you think browsers store web pages in a cache so that they can be retrieved without a new network access when possible? And what about your preferences for default font sizes? Or the download folder for files you can change when you feel like it, and even, on many browsers, restrictions on certain users or types of users? These are just simple examples, but the principles apply even to the largest (and smallest) databases today.

Databases Are Structured and Organized

All those examples can be implemented with a database because all those items need to be structured and organized. Some of the items can turn into large amounts of data (particularly a browser history if you don't choose an option to clear out the old information occasionally—another option that can be stored in a preference database), but some of the preferences listed are very small amounts of data: even with a fully modified file name for your download folder, 128 characters is enough to store that data element. Furthermore, most browsers store the name of your preferred download folder, but they don't store the history of download folders (last week's preferred download folder, the folder you used last month, etc.).

In thinking about databases, “big” is often relevant, but very frequently it's misleading. It is structure and organization that matter most.

And they matter a great deal to a mobile developer.

Precisely because mobile devices need to function in a world of constrained resources (e.g., limited storage space and battery power), structuring and organizing the data that is stored is particularly important. Built into a data management system are various optimizations that matter to databases large and small. For example, many databases can store character or string data. Such data is typically quite variable. Even if a database designer specifies that a text field can contain 50, 500, or 50,000 characters, behind the scenes trailing blanks are often discarded. The user (and even the developer) may never know this, but it makes the entire database function more efficiently. (Features like this function behind the scenes and they can be accessed and turned on or off by the database designer in many cases.)

Databases Are Smart

Being able to optimize storage to work with the presence or absence of characters in a text field is a powerful behind-the-scenes tool, but databases have intelligence built into them.

In most database engines today, the designer can specify many attributes of a field such as the following:

- *Name.* A name that is used internally is often quite different from the name that appears in the user interface.
- *Type.* You can specify that a certain database field must be an integer or a string or any other type that the database supports.
- *Optionality.* Some fields are optional, and you can specify that when you set up a database. (Many people have a car with a license number; many other people do not own cars.)
- *Default values.* A database can be set to provide a default value—either a simple value or a calculated value based on other data in the database.