
iPhone Application Programming Guide

General



2009-01-06



Apple Inc.
© 2009 Apple 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, mechanical, electronic, photocopying, recording, or otherwise, without prior written permission of Apple Inc., with the following exceptions: Any person is hereby authorized to store documentation on a single computer for personal use only and to print copies of documentation for personal use provided that the documentation contains Apple's copyright notice.

The Apple logo is a trademark of Apple Inc.

Use of the "keyboard" Apple logo (Option-Shift-K) for commercial purposes without the prior written consent of Apple may constitute trademark infringement and unfair competition in violation of federal and state laws.

No licenses, express or implied, are granted with respect to any of the technology described in this document. Apple retains all intellectual property rights associated with the technology described in this document. This document is intended to assist application developers to develop applications only for Apple-labeled computers.

Every effort has been made to ensure that the information in this document is accurate. Apple is not responsible for typographical errors.

Apple Inc.
1 Infinite Loop
Cupertino, CA 95014
408-996-1010

Apple, the Apple logo, Bonjour, Carbon, Cocoa, iPod, iTunes, Mac, Mac OS, Macintosh, Objective-C, Pages, Quartz, Safari, Sand, and Xcode are trademarks of Apple Inc., registered in the United States and other countries.

Finder, iPhone, and Multi-Touch are trademarks of Apple Inc.

NeXT is a trademark of NeXT Software, Inc., registered in the United States and other countries.

Java and all Java-based trademarks are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

OpenGL is a registered trademark of Silicon Graphics, Inc.

Simultaneously published in the United States and Canada.

Even though Apple has reviewed this document, APPLE MAKES NO WARRANTY OR REPRESENTATION, EITHER EXPRESS OR IMPLIED, WITH RESPECT TO THIS DOCUMENT, ITS QUALITY, ACCURACY, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. AS A RESULT, THIS DOCUMENT IS PROVIDED "AS IS," AND YOU, THE READER, ARE ASSUMING THE ENTIRE RISK AS TO ITS QUALITY AND ACCURACY.

IN NO EVENT WILL APPLE BE LIABLE FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES RESULTING FROM ANY DEFECT OR INACCURACY IN THIS DOCUMENT, even if advised of the possibility of such damages.

THE WARRANTY AND REMEDIES SET FORTH ABOVE ARE EXCLUSIVE AND IN LIEU OF ALL OTHERS, ORAL OR WRITTEN, EXPRESS OR IMPLIED. No Apple dealer, agent, or employee is authorized to make any modification, extension, or addition to this warranty.

Some states do not allow the exclusion or limitation of implied warranties or liability for incidental or consequential damages, so the above limitation or exclusion may not apply to you. This warranty gives you specific legal rights, and you may also have other rights which vary from state to state.

Contents

Introduction **Introduction 11**

Who Should Read This Document? 12
Organization of This Document 12
Providing Feedback 12
See Also 13

Chapter 1 **The Core Application 15**

Core Application Architecture 15
 The Application Life Cycle 15
 The Event-Handling Cycle 18
 Fundamental Design Patterns 20
The Application Runtime Environment 21
 Fast Launch, Short Use 21
 The Application Sandbox 21
 The Virtual Memory System 22
 The Automatic Sleep Timer 22
The Application Bundle 23
 The Information Property List 25
 Application Icon and Launch Images 27
 Nib Files 28
Handling Critical Application Tasks 29
 Initialization and Termination 29
 Responding to Interruptions 29
 Observing Low-Memory Warnings 31
Customizing Your Application's Behavior 32
 Launching in Landscape Mode 32
 Communicating with Other Applications 33
 Implementing Custom URL Schemes 34
 Displaying Application Preferences 37
Internationalizing Your Application 37
Tuning for Performance and Responsiveness 39
 Using Memory Efficiently 40
 Floating-Point Math Considerations 41
 Reducing Power Consumption 42
 Tuning Your Code 43

Chapter 2 **Window and Views 45**

What Are Windows and Views? 45
 The Role of UIWindow 45

- The Role of UIView 46
- UIKit View Classes 47
- The Role of View Controllers 50
- View Architecture and Geometry 50
 - The View Interaction Model 50
 - The View Rendering Architecture 52
 - View Coordinate Systems 55
 - The Relationship of the Frame, Bounds, and Center 56
 - Coordinate System Transformations 57
 - Content Modes and Scaling 58
 - Autosizing Behaviors 60
- Creating and Managing the View Hierarchy 61
 - Creating a View Object 63
 - Adding and Removing Subviews 63
 - Converting Coordinates in the View Hierarchy 65
 - Tagging Views 66
- Modifying Views at Runtime 66
 - Animating Views 67
 - Responding to Layout Changes 69
 - Redrawing Your View's Content 69
 - Hiding Views 70
- Creating a Custom View 70
 - Initializing Your Custom View 70
 - Drawing Your View's Content 71
 - Responding to Events 72
 - Cleaning Up After Your View 73

Chapter 3 **Event Handling 75**

- Events and Touches 75
- Event Delivery 77
 - Responder Objects and the Responder Chain 77
 - Regulating Event Delivery 78
- Handling Multi-Touch Events 79
 - The Event-Handling Methods 79
 - Handling Single and Multiple Tap Gestures 80
 - Detecting Swipe Gestures 81
 - Handling a Complex Multi-Touch Sequence 82
 - Event-Handling Techniques 83

Chapter 4 **Graphics and Drawing 85**

- The UIKit Graphics System 85
 - The View Drawing Cycle 85
 - Coordinates and Coordinate Transforms 86
 - Graphics Contexts 87

- Points Versus Pixels 87
- Color and Color Spaces 88
- Supported Image Formats 88
- Drawing Tips 89
 - Deciding When to Use Custom Drawing Code 89
 - Improving Drawing Performance 89
 - Maintaining Image Quality 90
- Drawing with Quartz and UIKit 90
 - Configuring the Graphics Context 91
 - Creating and Drawing Images 92
 - Creating and Drawing Paths 94
 - Creating Patterns, Gradients, and Shadings 94
- Drawing with OpenGL ES 94
 - Setting Up a Rendering Surface 95
 - Best Practices 96
 - Implementation Details 98
 - For More Information 102
- Applying Core Animation Effects 102
 - About Layers 103
 - About Animations 103

Chapter 5 **Text and Web 105**

- About Text and Web Support 105
 - Text Views 105
 - Web View 107
 - Keyboards and Input Methods 108
- Managing the Keyboard 110
 - Receiving Keyboard Notifications 110
 - Displaying the Keyboard 112
 - Dismissing the Keyboard 112
 - Moving Content That Is Located Under the Keyboard 113
- Drawing Text 115

Chapter 6 **Files and Networking 117**

- File and Data Management 117
 - Commonly Used Directories 117
 - Backup and Restore 118
 - Getting Paths to Application Directories 119
 - Reading and Writing File Data 121
 - File Access Guidelines 124
 - Saving State Information 125
 - Case Sensitivity 125
- Networking 125
 - Tips for Efficient Networking 126

Using Wi-Fi 126

Chapter 7 **Multimedia Support 127**

Using Sound in iPhone OS 127
 The Basics: Hardware Codecs, Audio Formats, and Audio Sessions 128
 Playing Audio 131
 Recording Audio 137
 Parsing Streamed Audio 138
 Audio Unit Support in iPhone OS 139
 Best Practices for iPhone Audio 139
 Playing Video Files 141

Chapter 8 **Device Support 143**

Accessing Accelerometer Events 143
 Choosing an Appropriate Update Interval 144
 Isolating the Gravity Component from Acceleration Data 145
 Isolating Instantaneous Motion from Acceleration Data 145
 Getting the Current Device Orientation 146
 Getting the User's Current Location 146
 Taking Pictures with the Camera 148
 Picking a Photo from the Photo Library 150

Chapter 9 **Application Preferences 151**

Guidelines for Preferences 151
 The Preferences Interface 152
 The Settings Bundle 153
 The Settings Page File Format 154
 Hierarchical Preferences 155
 Localized Resources 156
 Adding and Modifying the Settings Bundle 156
 Adding the Settings Bundle 156
 Preparing the Settings Page for Editing 157
 Configuring a Settings Page: A Tutorial 157
 Creating Additional Settings Page Files 161
 Accessing Your Preferences 161
 Debugging Preferences for Simulated Applications 162

Document Revision History 163

Figures, Tables, and Listings

Chapter 1 **The Core Application 15**

Figure 1-1	Application life cycle 16
Figure 1-2	The event and drawing cycle 18
Figure 1-3	Processing events in the main run loop 19
Figure 1-4	The Properties pane of a target's Info window 25
Figure 1-5	The information property list editor 26
Figure 1-6	The flow of events during an interruption 30
Figure 1-7	Defining a custom URL scheme in the <code>Info.plist</code> file 35
Figure 1-8	The Language preference view 38
Table 1-1	Design patterns used by iPhone applications 20
Table 1-2	A typical application bundle 23
Table 1-3	Important keys in the <code>Info.plist</code> file 26
Table 1-4	Responsibilities of the application delegate 29
Table 1-5	Keys and values of the <code>CFBundleURLTypes</code> property 34
Table 1-6	Tips for reducing your application's memory footprint 40
Table 1-7	Tips for allocating memory 41
Listing 1-1	The <code>main</code> function of an iPhone application 16
Listing 1-2	Handling a URL request based on a custom scheme 35
Listing 1-3	The contents of a language-localized subdirectory 38

Chapter 2 **Window and Views 45**

Figure 2-1	View class hierarchy 48
Figure 2-2	UIKit interactions with your view objects 51
Figure 2-3	View coordinate system 55
Figure 2-4	Relationship between a view's frame and bounds 56
Figure 2-5	Altering a view's bounds 57
Figure 2-6	View scaled using the scale-to-fill content mode 58
Figure 2-7	Content mode comparisons 59
Figure 2-8	View autosizing mask constants 61
Figure 2-9	Layered views in the Clock application 62
Figure 2-10	View hierarchy for the Clock application 62
Figure 2-11	Converting values in a rotated view 66
Table 2-1	Autosizing mask constants 60
Table 2-2	Animatable properties 67
Listing 2-1	Creating a window with views 64
Listing 2-2	Initializing a view subclass 71
Listing 2-3	A drawing method 72
Listing 2-4	Implementing the <code>dealloc</code> method 73

Chapter 3 Event Handling 75

- Figure 3-1 A multi-touch sequence and touch phases 76
- Figure 3-2 Relationship of a `UIEvent` object and its `UITouch` objects 76
- Listing 3-1 Detecting a double-tap gesture 80
- Listing 3-2 Tracking a swipe gesture in a view 81
- Listing 3-3 Handling a complex multi-touch sequence 82

Chapter 4 Graphics and Drawing 85

- Table 4-1 Supported image formats 88
- Table 4-2 Tips for improving drawing performance 89
- Table 4-3 Core graphics functions for modifying graphics state 91
- Table 4-4 Usage scenarios for images 93

Chapter 5 Text and Web 105

- Figure 5-1 Text classes in the `UICatalog` application 106
- Figure 5-2 A web view 108
- Figure 5-3 Several different keyboard types 109
- Figure 5-4 Several different keyboards and input methods 110
- Figure 5-5 Relative keyboard sizes in portrait and landscape modes 111
- Figure 5-6 Adjusting content to accommodate the keyboard 113
- Listing 5-1 Handling the keyboard notifications 114

Chapter 6 Files and Networking 117

- Table 6-1 Directories of an iPhone application 117
- Table 6-2 Commonly used search path constants 120
- Listing 6-1 Getting a file-system path to the application's `Documents/` directory 120
- Listing 6-2 Converting a property-list object to an `NSData` object and writing it to storage 122
- Listing 6-3 Reading a property-list object from the application's `Documents` directory 122
- Listing 6-4 Writing data to the application's `Documents` directory 123
- Listing 6-5 Reading data from the application's `Documents` directory 124

Chapter 7 Multimedia Support 127

- Figure 7-1 Media player interface with transport controls 141
- Table 7-1 Features provided by the audio session interface 129
- Table 7-2 Handling audio interruptions 131
- Table 7-3 Supported audio units 139
- Table 7-4 Audio tips 139
- Listing 7-1 Initializing an audio session 130
- Listing 7-2 Setting an audio session category 130

Listing 7-3	Creating a sound ID object	132
Listing 7-4	Playing a system sound	132
Listing 7-5	Triggering vibration	133
Listing 7-6	Configuring an AVAudioPlayer object	133
Listing 7-7	Implementing an AVAudioPlayer delegate method	134
Listing 7-8	Controlling an AVAudioPlayer object	134
Listing 7-9	Creating an audio queue object	135
Listing 7-10	Setting the playback level directly	136
Listing 7-11	The <code>AudioQueueLevelMeterState</code> structure	137
Listing 7-12	Playing full-screen movies	141

Chapter 8 **Device Support 143**

Table 8-1	Common update intervals for acceleration events	144
Listing 8-1	Configuring the accelerometer	143
Listing 8-2	Receiving an accelerometer event	144
Listing 8-3	Isolating the effects of gravity from accelerometer data	145
Listing 8-4	Getting the instantaneous portion of movement from accelerometer data	145
Listing 8-5	Initiating and processing location updates	147
Listing 8-6	Displaying the interface for taking pictures	149
Listing 8-7	Delegate methods for the image picker	149

Chapter 9 **Application Preferences 151**

Figure 9-1	Organizing preferences using child panes	155
Figure 9-2	Formatted contents of the <code>Root.plist</code> file	157
Figure 9-3	A root Settings page	158
Table 9-1	Preference element types	152
Table 9-2	Contents of the <code>Settings.bundle</code> directory	153
Table 9-3	Root-level keys of a preferences Settings Page file	154
Listing 9-1	Accessing preference values in an application	162

Introduction

Note: This document was previously titled *iPhone OS Programming Guide*.

The iPhone SDK provides the tools and resources needed to create native iPhone applications that appear as icons on the user's Home screen. Unlike a web application, which runs in Safari, a native application runs directly as a standalone executable on an iPhone OS–based device. Native applications have access to all the features that make the iPhone interesting, such as the accelerometers, location service, and Multi-Touch interface. They can also save data to the local file system and even communicate with other installed applications through custom URL schemes.



In iPhone OS, you develop native applications using the UIKit framework. This framework provides fundamental infrastructure and default behavior that makes it possible to create a functional application in a matter of minutes. Even though the UIKit framework (and other frameworks on the system) provide a significant amount of default behavior, they also provide hooks that you can use to customize and extend that behavior.

Who Should Read This Document?

This document is intended for both new and experienced iPhone OS developers who are creating native iPhone applications. Its purpose is to orient you to the architecture of an iPhone application and to show you the key customization points in the UIKit and other key system frameworks. Along the way, this document also provides guidance to help you make appropriate design choices. It also points out additional documents that may offer advice or further discussion of a given subject.

Although many of the frameworks described in this document are also present in Mac OS X, this document does not assume any familiarity with Mac OS X or its technologies.

Organization of This Document

This document has the following chapters:

- [“The Core Application”](#) (page 15) contains key information about the basic structure of every iPhone application, including some of the critical tasks every application should be prepared to handle.
- [“Window and Views”](#) (page 45) describes the iPhone windowing model and shows you how you use views to organize your user interface.
- [“Event Handling”](#) (page 75) describes the iPhone event model and shows you how to handle Multi-Touch events.
- [“Graphics and Drawing”](#) (page 85) describes the graphics architecture of iPhone OS and shows you how to draw shapes and images and incorporate animations into your content.
- [“Text and Web”](#) (page 105) describes the text support in iPhone OS, including examples of how you manage the system keyboard.
- [“Files and Networking”](#) (page 117) provides guidelines for working with files and network connections.
- [“Multimedia Support”](#) (page 127) shows you how to use the audio and video technologies available in iPhone OS.
- [“Device Support”](#) (page 143) shows you how to integrate features such as location tracking, the accelerometers, and the built-in camera into your application.
- [“Application Preferences”](#) (page 151) shows you how to configure your application preferences and display them in the Settings application.

Providing Feedback

If you have feedback about the documentation, you can provide it using the built-in feedback form at the bottom of every page.

If you encounter bugs in Apple software or documentation, you are encouraged to report them to Apple. You can also file enhancement requests to indicate features you would like to see in future revisions of a product or document. To file bugs or enhancement requests, go to the Bug Reporting page of the ADC website, which is at the following URL:

<http://developer.apple.com/bugreporter/>

You must have a valid ADC login name and password to file bugs. You can obtain a login name for free by following the instructions found on the Bug Reporting page.

See Also

The following documents provide important information that all developers should read prior to developing applications for iPhone OS:

- *iPhone Development Guide* provides important information about the iPhone development process from the tools perspective. This document covers the configuration of devices and the use of Xcode (and other tools) for building, running, and testing your software.
- *Cocoa Fundamentals Guide* provides fundamental information about the design patterns and practices used to develop iPhone applications.
- *iPhone Human Interface Guidelines* provides guidance and important information about how to design your iPhone application's user interface.

The following reference and conceptual documents provide additional information about key iPhone topics:

- *UIKit Framework Reference* and *Foundation Framework Reference* provide reference information for the classes discussed in this document.
- *View Controller Programming Guide for iPhone OS* provides information on the use of view controllers in creating interfaces for iPhone applications.
- *Table View Programming Guide for iPhone OS* provides information about working with table views, which are used frequently in iPhone applications.
- *The Objective-C 2.0 Programming Language* introduces Objective-C and the Objective-C runtime system, which is the basis of much of the dynamic behavior and extensibility of iPhone OS.

The Core Application

Every iPhone application is built using the UIKit framework and therefore has essentially the same core architecture. UIKit provides the key objects needed to run the application and to coordinate the handling of user input and the display of content on the screen. Where applications deviate from one another is in how they configure these default objects and also where they incorporate custom objects to augment their application's user interface and behavior.

Although customizations to your application's user interface and basic behavior occur down within your application's custom code, there are many customizations that you must make at the highest levels of the application. Because these application-level customizations affect the way your application interacts with the system and other applications installed on a device, it is important to understand when you need to act and when the default behavior is sufficient. This chapter provides an overview of the core application architecture and the high-level customization points to help you make determinations about when to customize and when to use the default behavior.

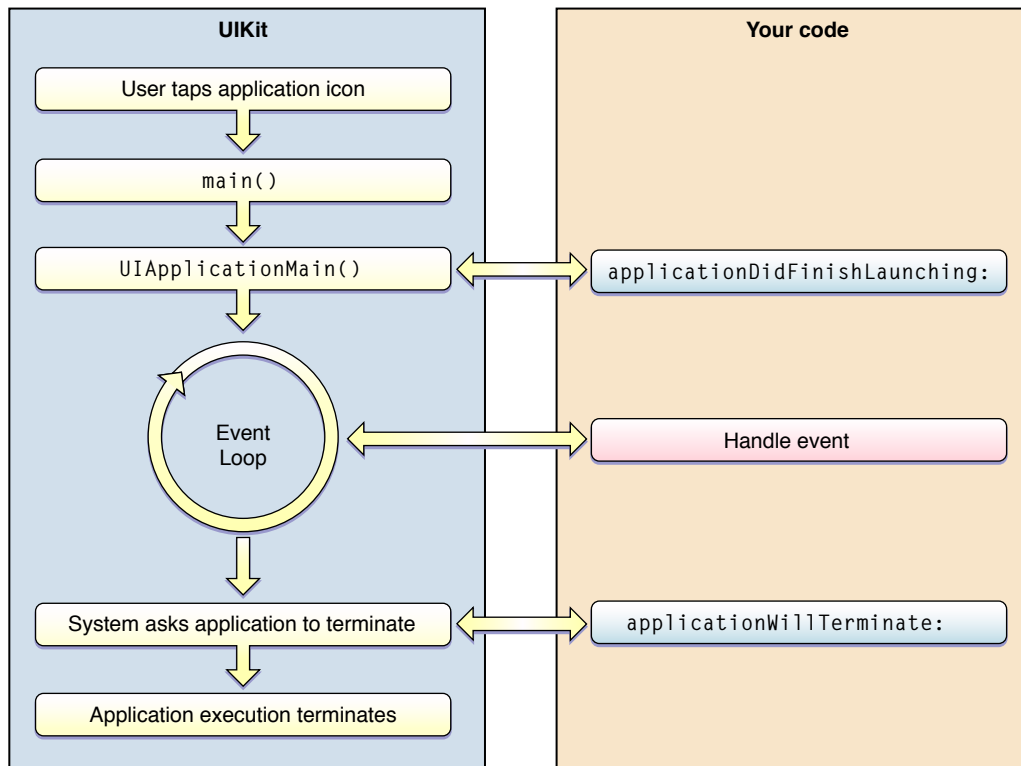
Core Application Architecture

From the time your application is launched by the user, to the time it exits, the UIKit framework manages the majority of the application's key infrastructure. An iPhone application receives events continuously from the system and must respond to those events. Receiving the events is the job of the `UIApplication` object but responding to the events is the responsibility of your custom code. In order to understand where you need to respond to events, though, it helps to understand a little about the overall life cycle and event cycles of an iPhone application. The following sections describe these cycles and also provide a summary of some of the key design patterns used throughout the development of iPhone applications.

The Application Life Cycle

The application life cycle constitutes the sequence of events that occurs between the launch and termination of your application. In iPhone OS, the user launches your application by tapping its icon on the Home screen. Shortly after the tap occurs, the system displays some transitional graphics and proceeds to launch your application by calling its `main` function. From this point on, the bulk of the initialization work is handed over to UIKit, which loads the application's user interface and readies its event loop. During the event loop, UIKit coordinates the delivery of events to your custom objects and responds to commands issued by your application. When the user performs an action that would cause your application to quit, UIKit notifies your application and begins the termination process.

Figure 1-1 depicts the simplified life cycle of an iPhone application. This diagram shows the sequence of events that occur from the time the application starts up to the time it quits. At initialization and termination, UIKit sends specific messages to the application delegate object to let it know what is happening. During the event loop, UIKit dispatches events to your application's custom event handlers. Handling initialization and termination events is discussed later in [“Initialization and Termination”](#) (page 29), and the event handling process is introduced in [“The Event-Handling Cycle”](#) (page 18) and covered in more detail in later chapters.

Figure 1-1 Application life cycle

The Main Function

In an iPhone application, the `main` function is used only minimally. Most of the actual work needed to run the application is handled by the `UIApplicationMain` function instead. As a result, when you start a new application project in Xcode, every project template provides an implementation of the standard `main` function like the one in “Handling Critical Application Tasks.” The `main` routine does only three things: it creates an autorelease pool, it calls `UIApplicationMain`, and it releases the autorelease pool. With few exceptions, you should never have to change the implementation of this function.

Listing 1-1 The `main` function of an iPhone application

```
#import <UIKit/UIKit.h>

int main(int argc, char *argv[])
{
    NSAutoreleasePool * pool = [[NSAutoreleasePool alloc] init];
    int retVal = UIApplicationMain(argc, argv, nil, nil);
    [pool release];
    return retVal;
}
```


Note: An autorelease pool is used in memory management. It is a Cocoa mechanism used to defer the release of objects created during a functional block of code. For more information about autorelease pools, see *Memory Management Programming Guide for Cocoa*. For specific memory-management guidelines related to autorelease pools in iPhone applications, see [“Allocating Memory Wisely”](#) (page 41).

The `UIApplicationMain` function at the heart of the preceding listing takes four parameters and uses them to initialize the application. Although you should never have to change the default values passed into this function, it is worth explaining their purpose in terms of starting the application. In addition to the `argc` and `argv` parameters passed into `main`, this function takes two string parameters that identify the principal class (that is, the class of the application object) and the class of the application delegate. If the value of the principal class string is `nil`, UIKit uses the `UIApplication` class by default. If the value of the application delegate’s class is `nil`, UIKit assumes that the application delegate is one of the objects loaded from your application’s main nib file (which is the case for applications built using the Xcode templates). Setting either of these parameters to a non-`nil` value causes the `UIApplicationMain` function to create an instance of the corresponding class during application startup and use it for the indicated purpose. Thus, if your application uses a custom subclass of `UIApplication` (which is not recommended, but certainly possible), you would specify your custom class name in the third parameter.

The Application Delegate

Monitoring the high-level behavior of your application is the responsibility of the application delegate object, which is a custom object that you provide. Delegation is a mechanism used to avoid subclassing complex UIKit objects, such as the default `UIApplication` object. Instead of subclassing and overriding methods, you use the complex object unmodified and put your custom code inside the delegate object. As interesting events occur, the complex object sends messages to your delegate object. You can use these “hooks” to execute your custom code and implement the behavior you need.

Important: The delegate design pattern is intended to save you time and effort when creating applications and is therefore a very important pattern to understand. For an overview of the key design patterns used by iPhone applications, see [“Fundamental Design Patterns”](#) (page 20). For a more detailed description of delegation and other UIKit design patterns, see *Cocoa Fundamentals Guide*.

The application delegate object is responsible for handling several critical system messages and *must* be present in every iPhone application. The object can be an instance of any class you like, as long as it adopts the `UIApplicationDelegate` protocol. The methods of this protocol define the hooks into the application life cycle and are your way of implementing custom behavior. Although you are not required to implement all of the methods, every application delegate should implement the methods described in [“Handling Critical Application Tasks”](#) (page 29).

For additional information about the methods of the `UIApplicationDelegate` protocol, see *UIApplicationDelegate Protocol Reference*.

The Main Nib File

Another task that occurs at initialization time is the loading of the application’s main nib file. If the application’s information property list (`Info.plist`) file contains the `NSMainNibFile` key, the `UIApplication` object loads the nib file specified by that key as part of its initialization process. The main nib file is the only nib file that is loaded for you automatically; however, you can load additional nib files later as needed.

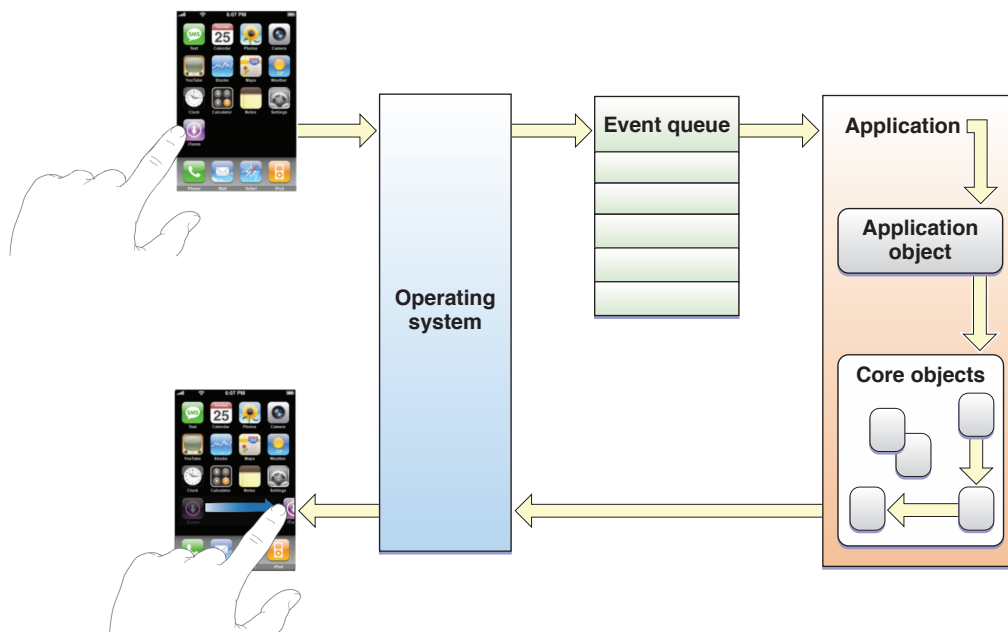
Nib files are disk-based resource files that store a snapshot of one or more objects. The main nib file of an iPhone application typically contains a window object, the application delegate object, and perhaps one or more other key objects for managing the window. Loading a nib file reconstitutes the objects in the nib file, converting each object from its on-disk representation to an actual in-memory version that can be manipulated by your application. Objects loaded from nib files are no different than the objects you create programmatically. For user interfaces, however, it is often more convenient to create the objects associated with your user interface graphically (using the Interface Builder application) and store them in nib files rather than create them programmatically.

For more information about nib files and their use in iPhone applications, see [“Nib Files”](#) (page 28). For additional information about how to specify your application’s main nib file, see [“The Information Property List”](#) (page 25).

The Event-Handling Cycle

After the `UIApplicationMain` function initializes the application, it starts the infrastructure needed to manage the application’s event and drawing cycle, which is depicted in Figure 1-2. As the user interacts with a device, iPhone OS detects touch events and places them in the application’s event queue. The event-handling infrastructure of the `UIApplication` object takes each event off the top of this queue and delivers it to the object that best suited to handle it. For example, a touch event occurring in a button would be delivered to the corresponding button object. Events can also be delivered to controller objects and other objects indirectly responsible for handling touch events in the application.

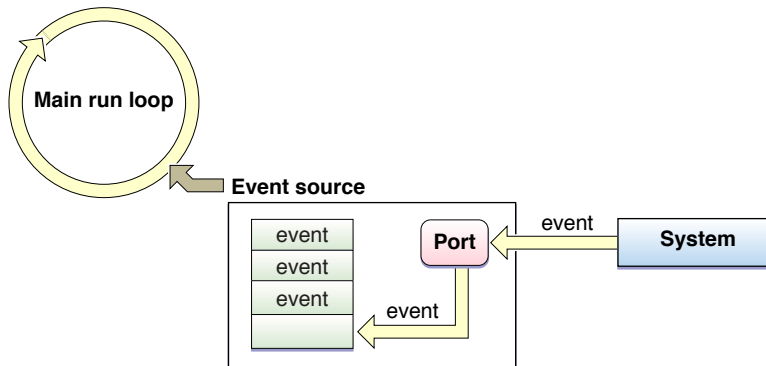
Figure 1-2 The event and drawing cycle



In the iPhone OS Multi-Touch event model, touch data is encapsulated in a single event object (`UIEvent`). To track individual touches, the event object contains touch objects (`UITouch`), one for each finger that is touching the screen. As the user places fingers on the screen, moves them around, and finally removes them from the screen, the system reports the changes for each finger in the corresponding touch object.

When it launches an application, the system creates both a process and a single thread for that application. This initial thread becomes the application's main thread and is where the `UIApplication` object sets up the **main run loop** and configures the application's event-handling code. Figure 1-3 shows the relationship of the event-handling code to the main run loop. Touch events sent by the system are queued until they can be processed by the application's main run loop.

Figure 1-3 Processing events in the main run loop



Note: A run loop monitors sources of input for a given thread of execution. When an input source has data to process, the run loop wakes up the thread and dispatches control to the handler for that input source. When the handler finishes, control passes back to the run loop, which processes the next event or puts the thread to sleep if there is nothing more to do. You can install your own input sources, including ports and timers, on a run loop using the `NSRunLoop` class of the Foundation framework. For more on `NSRunLoop` and run loops in general, see *Threading Programming Guide*.

The `UIApplication` object configures the main run loop with an input source that processes touch events by dispatching them to the appropriate responder objects. A responder object is an object that inherits from the `UIResponder` class and that implements one or more methods for processing the different phases of a touch event. Responder objects in an application include instances of `UIApplication`, `UIWindow`, `UIView`, and all `UIView` subclasses. The application typically dispatches events to the `UIWindow` object representing the application's main window. The window object, in turn, forwards the event to its **first responder**, which is typically the view object (`UIView`) on which the touch took place.

In addition to defining the methods you use to handle events, the `UIResponder` class also defines the programmatic structure of the **responder chain**, which is the Cocoa mechanism for cooperative event handling. The responder chain is a linked series of responder objects in an application, which usually starts at the first responder. If the first responder object cannot handle the event, it passes it to the next responder in the chain. The message continues traveling up the chain—toward higher-level responder objects such as the window, the application, and the application's delegate—until the event is handled. If the event isn't handled, it is discarded.

The responder object that handles the event tends to set in motion a series of programmatic actions that result in the application redrawing all or a portion of its user interface (as well as other possible outcomes, such as the playing of a sound). For example, a control object (that is, a subclass of `UIControl`) handles an event by sending an action message to another object, typically the controller that manages the current set of active views. While processing the action message, the controller might change the user interface or adjust the position of views in ways that require some of those views to redraw themselves. When this happens, the view and graphics infrastructure takes over and processes the required redraw events in the most efficient manner possible.

For more information about events, responders, and how you handle events in your own custom objects, see “[Event Handling](#)” (page 75). For information about how windows and views fit into the event-handling scheme, see “[The View Interaction Model](#)” (page 50). For additional information about the graphics infrastructure and how views are updated, see “[The View Drawing Cycle](#)” (page 85).

Fundamental Design Patterns

The design of the UIKit framework incorporates many of the design patterns found in Cocoa applications in Mac OS X. Understanding these design patterns is crucial to creating iPhone applications, so it is worth taking a few moments to learn about them. The following sections provide a brief overview of these design patterns.

Table 1-1 Design patterns used by iPhone applications

Design pattern	Description
Model-View-Controller	The Model-View-Controller (MVC) design pattern is a way of dividing your code into independent functional areas. The model portion defines your application’s underlying data engine and is responsible for maintaining the integrity of that data. The view portion defines the user interface for your application and has no explicit knowledge of the origin of data displayed in that interface. The controller portion acts as a bridge between the model and view and facilitates updates between them.
Delegation	The delegation design pattern is a way of modifying complex objects without subclassing them. Instead of subclassing, you use the complex object as is and put any custom code for modifying the behavior of that object inside a separate object, which is referred to as the delegate object. At pre-defined times, the complex object then calls the methods of the delegate object to give it a chance to run its custom code.
Target-action	Controls use the target-action design pattern to notify your application of user interactions. When the user interacts with a control in a predefined way (such as by tapping a button), the control sends a message (the action) to an object you specify (the target). Upon receiving the action message, the target object can then respond in an appropriate manner (such as by updating application state in response to the button push).
Managed memory model	The Objective-C language uses a reference-counted scheme for determining when to release objects from memory. When an object is first created, it is given a reference count of 1. Other objects can then use the <code>retain</code> , <code>release</code> , or <code>autorelease</code> methods of the object to increase and decrease that reference count appropriately. When an object’s reference count reaches 0, the Objective-C runtime calls the object’s cleanup routines and then deallocates it.

For a more thorough discussion of these design patterns, see *Cocoa Fundamentals Guide*.

The Application Runtime Environment

The runtime environment of iPhone OS is designed for the fast and secure execution of programs. The following sections describe the key aspects of this runtime environment and provide guidance on how best to operate within it.

Fast Launch, Short Use

The strength of iPhone OS–based devices is their immediacy. A typical user pulls a device out of a pocket or bag and uses it for a few seconds, or maybe a few minutes, before putting it away again. The user might be taking a phone call, looking up a contact, changing the current song, or getting some piece of information during that time.

In iPhone OS, only one foreground application runs at a time. This means that every time the user taps your application's icon on the Home screen, your application must launch and initialize itself quickly to minimize the delay. If your application takes a long time to launch, the user may be less inclined to use it.

In addition to launching quickly, your application must be prepared to exit quickly too. Whenever the user leaves the context of your application, whether by pressing the Home button or by using a feature that opens content in another application, iPhone OS tells your application to quit. At that time, you need to save any unsaved changes to disk and exit as quickly as possible. If your application takes more than 5 seconds to quit, the system may terminate it outright.

Even though your application does not run in the background when the user switches to another application, you are encouraged to make it appear as if that is the case. When your application quits, you should save out information about your application's current state in addition to any unsaved data. At launch time, you should look for this state information and use it to restore your application to the state it was in when it was last used. Doing so provides a more consistent user experience by putting the user right back where they were when they last used your application. Saving the user's place in this way also saves time by potentially eliminating the need to navigate back through multiple screens' worth of information each time an application is launched.

The Application Sandbox

For security reasons, iPhone OS restricts an application (including its preferences and data) to a unique location in the file system. This restriction is part of the security feature known as the application's "sandbox." The **sandbox** is a set of fine-grained controls limiting an application's access to files, preferences, network resources, hardware, and so on. In iPhone OS, an application and its data reside in a secure location that no other application can access. When an application is installed, the system computes a unique opaque identifier for the application. Using a root application directory and this identifier, the system constructs a path to the application's home directory. Thus an application's home directory could be depicted as having the following structure:

/ ApplicationRoot / ApplicationID /

During the installation process, the system creates the application's home directory and several key subdirectories, configures the application sandbox, and copies the application bundle to the home directory. The use of a unique location for each application and its data simplifies backup-and-restore operations and uninstallation. For more information about the application-specific directories created for each application, see [“Commonly Used Directories”](#) (page 117).

Important: The sandbox limits the damage an attacker can cause to other applications and to the system, but it cannot prevent attacks from happening. In other words, the sandbox does not protect your application from direct attacks by malicious entities. For example, if there is an exploitable buffer overflow in your input-handling code and you fail to validate user input, an attacker might still be able to crash your program or use it to execute the attacker's code.

The Virtual Memory System

To manage program memory, iPhone OS uses essentially the same virtual memory system found in Mac OS X. In iPhone OS, each program still has its own virtual address space, but (unlike Mac OS X) its usable virtual memory is constrained by the amount of physical memory available. This is because iPhone OS does not write volatile pages to disk when memory gets full. Instead, the virtual memory system frees up nonvolatile memory, as needed, to make sure the running application has the space it needs. It does this by removing memory pages that are not being used and that contain read-only contents, such as code pages. Such pages can always be loaded back into memory later if they are needed again.

If memory continues to be constrained, the system may also send notifications to the running applications, asking them to free up additional memory. All applications should respond to this notification and do their part to help relieve the memory pressure. For information on how to handle such notifications in your application, see [“Observing Low-Memory Warnings”](#) (page 31).

The Automatic Sleep Timer

One way iPhone OS attempts to save power is through the automatic sleep timer. If the system does not detect touch events for an extended period of time, it dims the screen initially and eventually turns it off altogether. Although most developers should leave this timer on, game developers and developers whose applications do not use touch inputs can disable this timer to prevent the screen from dimming while their application is running. To disable the timer, set the `idleTimerDisabled` property of the shared `UIApplication` object to YES.

Because it results in greater power consumption, disabling the sleep timer should be avoided at all costs. The only applications that should consider using it are mapping applications, games, or applications that do not rely on touch inputs but do need to display visual content on the device's screen. Audio applications do not need to disable the timer because audio content continues to play even after the screen dims. If you do disable the timer, be sure to reenable it as soon as possible to give the system the option to conserve more power. For additional tips on how to save power in your application, see [“Reducing Power Consumption”](#) (page 42).

The Application Bundle

When you build your iPhone application, Xcode packages it as a bundle. A **bundle** is a directory in the file system that groups related resources together in one place. An iPhone application bundle contains the application executable and any resources used by the application (for instance, the application icon, other images, and localized content). Table 1-2 lists the contents of a typical iPhone application bundle, which for demonstration purposes here is called `MyApp`). This example is for illustrative purposes only. Some of the files listed in this table may not appear in your own application bundles.

Table 1-2 A typical application bundle

File	Description
<code>MyApp</code>	The executable file containing your application's code. The name of this file is the same as your application name minus the <code>.app</code> extension. This file is required.
<code>Settings.bundle</code>	The settings bundle is a file package that you use to add application preferences to the Settings application. This bundle contains property lists and other resource files to configure and display your preferences. See “Displaying Application Preferences” (page 37) for more information.
<code>Icon.png</code>	The 57 x 57 pixel icon used to represent your application on the device home screen. This icon should not contain any glossy effects. The system adds those effects for you automatically. This file is required. For information about this image file, see “Application Icon and Launch Images” (page 27).
<code>Icon-Settings.png</code>	The 29 x 29 pixel icon used to represent your application in the Settings application. If your application includes a settings bundle, this icon is displayed next to your application name in the Settings application. If you do not specify this icon file, the <code>Icon.png</code> file is scaled and used instead. For information about this image file, see “Displaying Application Preferences” (page 37).
<code>MainWindow.nib</code>	The application's main nib file contains the default interface objects to load at application launch time. Typically, this nib file contains the application's main window object and an instance of the application delegate object. Other interface objects are then either loaded from additional nib files or created programmatically by the application. (The name of the main nib file can be changed by assigning a different value to the <code>NSMainNibFile</code> key in the <code>Info.plist</code> file. See “The Information Property List” (page 25) for further information.)
<code>Default.png</code>	The 480 x 320 pixel image to display when your application is launched. The system uses this file as a temporary background until your application loads its window and user interface. For information about this image file, see “Application Icon and Launch Images” (page 27).

File	Description
iTunesArtwork	The 512 x 512 icon for an application that is distributed using ad-hoc distribution. This icon would normally be provided by the App Store; because applications distributed in an ad-hoc manner do not go through the App Store, however, it must be present in the application bundle instead. iTunes uses this icon to represent your application. (The file you specify for this property should be the same one you would have submitted to the App Store (typically a JPEG or PNG file), were you to distribute your application that way. The filename must be the one shown at left and must not include a filename extension.)
Info.plist	Also known as the information property list, this file is a property list defining key values for the application, such as bundle ID, version number, and display name. See “The Information Property List” (page 25) for further information. This file is required.
sun.png (or other resource files)	Nonlocalized resources are placed at the top level of the bundle directory (sun.png represents a nonlocalized image file in the example). The application uses nonlocalized resources regardless of the language setting chosen by the user.
en.lproj fr.lproj es.lproj other language-specific project directories	Localized resources are placed in subdirectories with an ISO 639-1 language abbreviation for a name plus an .lproj suffix. (For example, the en.lproj, fr.lproj, and es.lproj directories contain resources localized for English, French, and Spanish.) For more information, see “Internationalizing Your Application” (page 37).

An iPhone application should be internationalized and have a *language.lproj* folder for each language it supports. In addition to providing localized versions of your application’s custom resources, you can also localize your application icon (Icon.png), default image (Default.png), and Settings icon (Icon-Settings.png) by placing files with the same name in your language-specific project directories. Even if you provide localized versions, however, you should always include a default version of these files at the top-level of your application bundle. The default version is used in situations where a specific localization is not available.

You use the methods of the `NSBundle` class or the functions associated with the `CFBundleRef` opaque type to obtain paths to localized and nonlocalized image and sound resources stored in the application bundle. For example, to get a path to the image file `sun.png` (shown in [“Responding to Interruptions”](#) (page 29)) and create an image file from it would require two lines of Objective-C code:

```
NSString* imagePath = [[NSBundle mainBundle] pathForResource:@"sun"
ofType:@"png"];
UIImage* sunImage = [[UIImage alloc] initWithContentsOfFile:imagePath];
```

Calling the `mainBundle` class method returns an object representing the application bundle. For information on loading resources, see *Resource Programming Guide*.

The information property list is a file named `Info.plist` that is included with every iPhone application project created by Xcode. It is a property list whose key-value pairs specify essential runtime-configuration information for the application. The elements of the information property list are organized in a hierarchy in which each node is an entity such as an array, dictionary, string, or other scalar type.

Figure 1-4 The Properties pane of a target's Info window

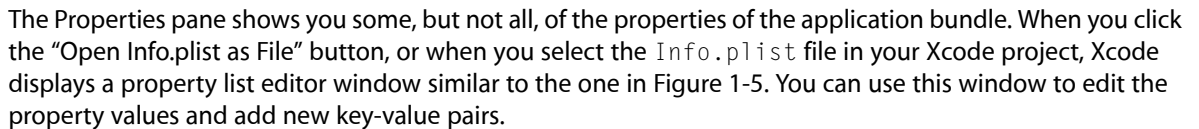
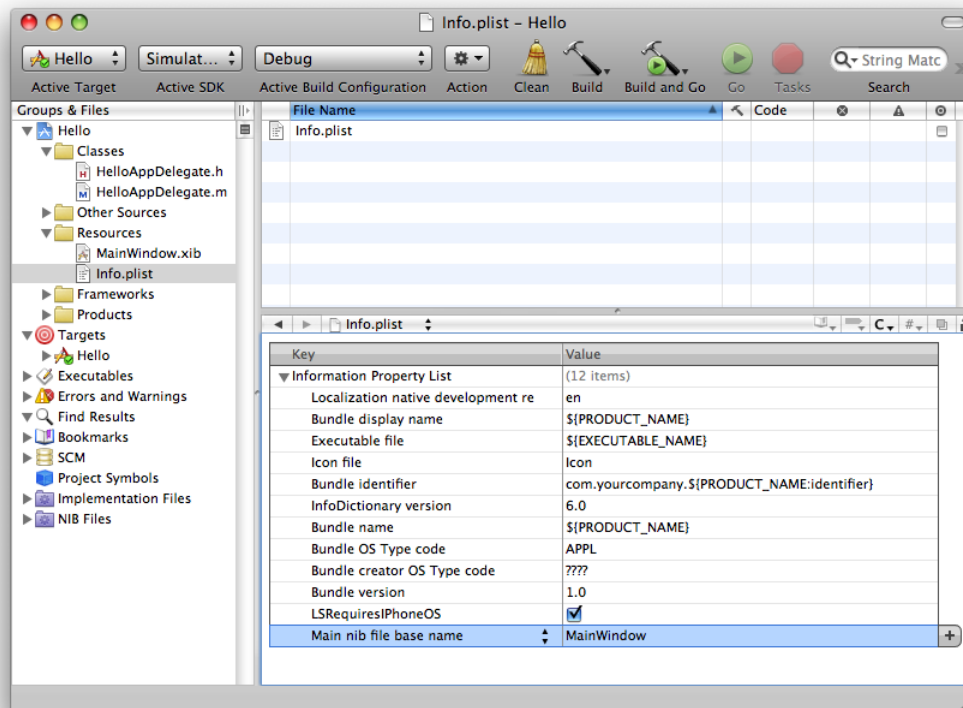


Figure 1-5 The information property list editor

Xcode automatically sets the value of some of these properties, but others you need to set explicitly. Table 1-3 lists some of the important keys you might want to include in your application's `Info.plist` file. For a complete list of properties you can include in this file, see *Runtime Configuration Guidelines*.

Table 1-3 Important keys in the `Info.plist` file

Key	Value
<code>CFBundleDisplayName</code>	The name to display underneath the application icon. This value should be localized for all supported languages.
<code>CFBundleIdentifier</code>	An identifier string that specifies the application type of the bundle. This string should be a uniform type identifier (UTI). For example, if your company name is Ajax and the application is named Hello, the bundle identifier would be <code>com.Ajax.Hello</code> . The bundle identifier is used in validating the application signature.
<code>CFBundleURLTypes</code>	An array of URL types that the application can handle. Each URL type is a dictionary that defines the schemes (for example, such as <code>http</code> or <code>mailto</code>) that the application can handle. This property allows applications to register custom URL schemes.

Key	Value
<code>CFBundleVersion</code>	A string that specifies the build version number of the bundle. This value is a monotonically increased string, comprised of one or more period-separated integers. This value cannot be localized.
<code>LSRequiresiPhoneOS</code>	A Boolean value that indicates whether the bundle can run on iPhone OS only. Xcode adds this key automatically and sets its value to true. You should not change the value of this key.
<code>NSMainNibFile</code>	A string that identifies the name of the application's main nib file. If you want to use a nib file other than the default one created for your project, associate the name of that nib file with this key. The name of the nib file should not include the <code>.nib</code> filename extension.
<code>UIStatusBarStyle</code>	A string that identifies the style of the status bar as the application launches. This value is based on the <code>UIStatusBarStyle</code> constants declared in <code>UIApplication.h</code> header file. The default style is <code>UIBarStyleDefault</code> . The application can change this initial status-bar style when it finishes launching.
<code>UIStatusBarHidden</code>	A Boolean value that determines whether the status bar is initially hidden when the application launches. Set it to true to hide the status bar. The default value is <code>false</code> .
<code>UIInterfaceOrientation</code>	A string that identifies the initial orientation of the application's user interface. This value is based on the <code>UIInterfaceOrientation</code> constants declared in the <code>UIApplication.h</code> header file. The default style is <code>UIInterfaceOrientationPortrait</code> . For more information on launching your application in landscape mode, see “Launching in Landscape Mode” (page 32).
<code>UIPrerenderedIcon</code>	A Boolean value that indicates whether the application icon already includes gloss and bevel effects. This property is <code>false</code> by default. Set it to true if you do not want the system to add these effects to your artwork.
<code>UIRequiresPersistentWiFi</code>	A Boolean value that notifies the system that the application uses the Wi-Fi network for communication. Applications that use Wi-Fi for any period of time must set this key to true; otherwise, after 30 minutes, the device shuts down Wi-Fi connections to save power. Setting this flag also lets the system know that it should display the network selection dialog when Wi-Fi is available but not currently being used. The default value is <code>false</code> .

Properties with string values that are displayed in the user interface should be localized. Specifically, the string value in `Info.plist` should be a key to a localized string in the `InfoPlist.strings` file of a language-localized subdirectory. For more information, see [“Internationalizing Your Application”](#) (page 37).

Application Icon and Launch Images

The file for the icon displayed in the user's Home screen has the default name of `Icon.png` (although the `CFBundleIconFile` property in the `Info.plist` file lets you rename it). It should be a PNG image file located in the top level of the application bundle. The application icon should be a 57 x 57 pixel image

without any shine or round beveling effects. Typically, the system applies these effects to the icon before displaying it. You can override that behavior, however, by including the `UIPrerenderedIcon` key in your application's `Info.plist` file; for more information, see [Table 1-3](#) (page 26).

Note: If you are distributing your application to local users using ad-hoc distribution (instead of going through the App Store), your bundle should also include a 512 x 512 pixel version of your application icon in a file called `iTunesArtwork`. This file provides the icon that iTunes displays when distributing your application.

The file for the application's launch image is named `Default.png`. This image should closely resemble the application's initial user interface; the system displays the launch image before an application is ready to display its user interface, giving users the impression of a quick launch. The launch image should also be a PNG image file, located in the top level of the application bundle. If the application is launched through a URL, the system looks for a launch image named `Default-scheme.png`, where *scheme* is the scheme of the URL. If that file is not present, it chooses `Default.png` instead.

To add an image file to a project in Xcode, choose **Add to Project** from the **Project** menu, locate the file in the browser, and click **Add**.

Note: In addition to the icons and launch image at the top level of your bundle, you can also include localized versions of those images in your application's language-specific project subdirectories. For more information on localizing resources in your application, see [“Internationalizing Your Application”](#) (page 37).

Nib Files

A nib file is a data file that stores a set of “freeze-dried” objects that the application plans to use later. Applications use nib files most often to store the windows and views that make up their user interface. When you load a nib file into your application, the nib-loading code turns the contents into real objects that your application can manipulate. In this way, nib files eliminate the need to create these same objects programmatically from your code.

Interface Builder is the visual design environment that you use to create nib files. You assemble nib files using a combination of standard objects (such as the windows and views provided by the UIKit framework) and custom objects from your Xcode projects. Creating view hierarchies within Interface Builder is a simple matter of dragging and dropping views in place. You can also configure the attributes of each object using the inspector window and create connections between objects to define their runtime relationships. All of the changes you make are subsequently saved to disk as part of your nib file.

At runtime, you load nib files into your application when you need the objects they contain. Typically, you load a nib file when your user interface changes and you need to display some new views on the screen. If your application uses view controllers, the view controller handles the nib loading process for you automatically but you can also load nib files yourself using the methods of the `NSBundle` class.

For information on how to design your application's user interface, see *iPhone Human Interface Guidelines*. For information on how to create nib files, see *Interface Builder User Guide*.

Handling Critical Application Tasks

This section describes the handful of tasks that every iPhone application should perform. These tasks are part of the overall life cycle of the application and are therefore part of the key ways your application integrates with iPhone OS. In the worst case, failure to handle some of these tasks could even lead to the termination of your application by the operating system.

Initialization and Termination

During initialization and termination, the `UIApplication` class sends appropriate messages to the application delegate to let it perform any necessary tasks. Although your application is not required to respond to these messages, nearly all iPhone applications should handle them. Initialization time is where you prepare your application's user interface and put the application into its initial running state. Similarly, termination is the time when you should be writing unsaved data and key application state to disk.

Because one iPhone application must quit before another can be launched, the time it takes to execute your initialization and termination code should be as small as possible. Initialization time is not the time to start loading large data structures that you do not intend to use right away. Your goal during startup should be to present your application's user interface as quickly as possible, preferably in the state it was in when your application last quit. If your application requires additional time at launch to load data from the network or do other tasks that might be slow, you should get your interface up and running first and then launch the slow task on a background thread. Doing so gives you the opportunity to display a progress indicator or other feedback to the user to indicate that your application is loading the necessary data or doing something important.

Table 1-4 lists the methods of the `UIApplicationDelegate` protocol that you implement in your application delegate to handle initialization and termination chores. This table also lists some of the key chores you should perform in each method.

Table 1-4 Responsibilities of the application delegate

Delegate method	Description
<code>applicationDidFinishLaunching:</code>	Use this method to restore the application to the state it was in during the previous session. You can also use this method to perform any custom initialization to your application data structures and user interface.
<code>applicationWillTerminate:</code>	Use this method to save any unsaved data or key application state to disk. You can also use this method to perform additional cleanup operations, such as deleting temporary files.

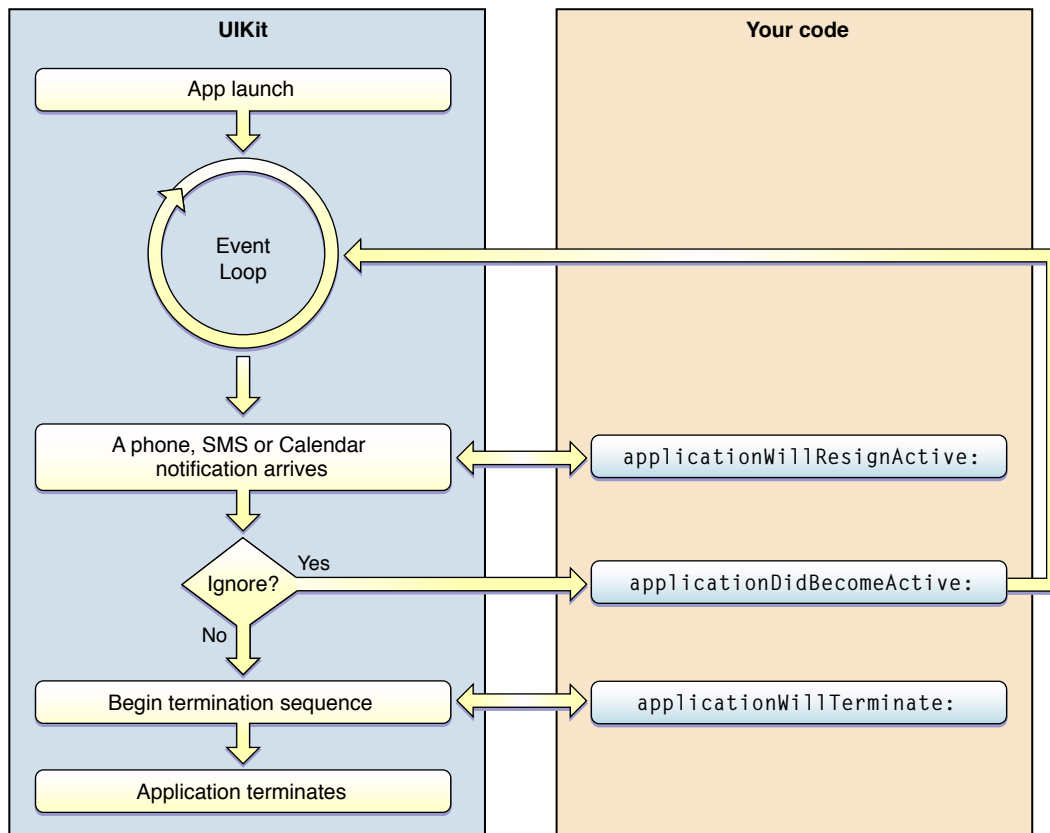
Responding to Interruptions

Besides the Home button, which terminates your application, the system can interrupt your application temporarily to let the user respond to important events. For example, an application can be interrupted by an incoming phone call, an SMS message, a calendar alert, or by the user pressing the Sleep button on a device. Whereas a press of the Home button terminates your application, these interruptions may be only

temporary. If the user ignores the interruption, your application continues running as before. If the user decides to take a call or reply to an SMS message, however, the system does proceed to terminate your application.

Figure 1-6 shows the sequence of events that occurs during the arrival of a phone call, SMS message, or calendar alert. The steps that immediately follow describe the key points in the sequence in more detail, including some of the things your application should do in response to each event. This sequence does not reflect what happens when the user pushes the Sleep/Wake button; that sequence is described after the steps below.

Figure 1-6 The flow of events during an interruption



1. The system detects an incoming phone call or SMS message, or a calendar event occurs.
2. The system calls your application delegate's `applicationWillResignActive:` method. The system also disables the delivery of touch events to your application.

Interruptions amount to a temporary loss of control by your application. If such a loss of control might affect your application's behavior or cause a negative user experience, you should take appropriate steps in your delegate method to prevent that from happening. For example, if your application is a game, you should pause the game. You should also disable timers, throttle back your OpenGL frame rates (if using OpenGL), and generally put your application into a sleep state. While in the resigned state, your application continues to run but should not do any significant work.

3. The system displays an alert panel with information about the event. The user can choose to ignore the event or respond to it.

4. If the user ignores the event, the system calls your application delegate's `applicationDidBecomeActive:` method and resumes the delivery of touch events to your application.

You can use this delegate method to reenable timers, throttle up your OpenGL frame rates, and generally wake up your application from its sleep state. For games that are in a paused state, you should consider leaving the game in that state until the user is ready to resume play. For example, you might display an alert panel with controls to resume play.

5. If the user responds to the event, instead of ignoring it, the system calls your application delegate's `applicationWillTerminate:` method. Your application should terminate as usual, saving any needed contextual information to return the user to the same place in your application upon your next launch.

After terminating your application, the system proceeds to launch the application responsible for the interruption.

Depending on what the user does while responding to an interruption, the system may launch your application again when that interruption ends. For example, if the user takes a phone call and then hangs up, the system relaunches your application. If, while on the call, the user goes back to the Home screen or launches another application, the system does not relaunch your application.

Important: When the user takes a call and then relaunches your application while on the call, the height of the status bar grows to reflect the fact that the user is on a call. Similarly, when the user ends the call, the status bar height shrinks back to its regular size. Your application should be prepared for these changes in the status bar height and adjust its content area accordingly. View controllers handle this behavior for you automatically. If you lay out your user interface programmatically, however, you need to take the status bar height into account when laying out your views and implement the `layoutSubviews` method to handle dynamic layout changes.

If the user presses the Sleep/Wake button on a device while running your application, the system calls your application delegate's `applicationWillResignActive:` method, stops the delivery of touch events, and then puts the device to sleep. When the user wakes the device later, the system calls your application delegate's `applicationDidBecomeActive:` method and begins delivering events to the application again. As you do with other interruptions, you should use these methods to put your application into a sleep state (or pause the game) and wake it up again. While in the sleep state, your application should use as little power as possible.

Observing Low-Memory Warnings

When the system dispatches a low-memory notification to your application, heed the warning. iPhone OS notifies the frontmost application whenever the amount of free memory dips below a safe threshold. If your application receives this notification, it must free up as much memory as it can by releasing objects it does not need or clearing out memory caches that it can recreate easily later.

UIKit provides several ways to receive low-memory notifications, including the following:

- Implement the `applicationDidReceiveMemoryWarning:` method of your application delegate.
- Override the `didReceiveMemoryWarning` method in your custom `UIViewController` subclass.
- Register to receive the `UIApplicationDidReceiveMemoryWarningNotification` notification.

Upon receiving any of these notifications, your handler method should respond by immediately freeing up any unneeded memory. View controllers should purge any views that are currently offscreen, and your application delegate should release any data structures it can or notify other application objects to release memory they own.

If your custom objects have known purgeable resources, you can have those objects register for the `UIApplicationDidReceiveMemoryWarningNotification` notification and release their purgeable resources directly. Have these objects register if you have a few objects that manage most of your purgeable resources and it is appropriate to purge all of those resources. If you have many purgeable objects or want to coordinate the release of only a subset of those objects, however, you might want to use your application delegate to release the desired objects.

Important: Like the system applications, your applications should always handle low-memory warnings, even if they do not receive those warnings during your testing. System applications consume small amounts of memory while processing requests. When a low-memory condition is detected, the system delivers low-memory warnings to all running programs (including your application) and may terminate some background applications (if necessary) to ease memory pressure. If not enough memory is released—perhaps because your application is leaking or still consuming too much memory—the system may still terminate your application.

Customizing Your Application's Behavior

There are several ways to customize your basic application behavior to provide the user experience you want. The following sections describe some of the customizations that you must make at the application level.

Launching in Landscape Mode

Applications in iPhone OS normally launch in portrait mode to match the orientation of the Home screen. If you have an application that runs in both portrait and landscape modes, your application should always launch in portrait mode initially and then let its view controllers rotate the interface as needed based on the device's orientation. If your application runs in landscape mode only, however, you must perform the following steps to make it launch in a landscape orientation initially.

- In your application's `Info.plist` file, add the `UIInterfaceOrientation` key and set its value to the landscape mode. You can set the value of this key to `UIInterfaceOrientationLandscapeLeft` or `UIInterfaceOrientationLandscapeRight`.
- Lay out your views in landscape mode and make sure that their autosizing options are set correctly.
- Override your view controller's `shouldAutorotateToInterfaceOrientation:` method and return YES only for the desired landscape orientation and NO for portrait orientations.

Important: The preceding steps assume your application uses view controllers to manage its view hierarchy. View controllers provide a significant amount of infrastructure for handling orientation changes as well as other complex view-related events. If your application is not using view controllers—as may be the case with games and other OpenGL ES–based applications—you are responsible for rotating the drawing surface (or adjusting your drawing commands) as needed to present your content in landscape mode.

The `UIInterfaceOrientation` key provides a hint to iPhone OS that it should set the orientation of the application status bar (if one is displayed) and your the views in any view controllers at launch time. In iPhone OS 2.1 and later, view controllers respect this key and set their view’s initial orientation to match. Using this key is also equivalent to calling the `setStatusBarOrientation:animated:` method of `UIApplication` early in the execution of your `applicationDidFinishLaunching:` method.

Note: To launch a view controller–based application in landscape mode in versions of iPhone OS prior to v2.1, you need to apply a 90 degree rotation to the transform of the application’s root view in addition to all the preceding steps. Prior to iPhone OS 2.1, view controllers did not automatically rotate their views based on the value of the `UIInterfaceOrientation` key. This step is not necessary in iPhone OS 2.1 and later, however.

Communicating with Other Applications

If an application handles URLs of a known type, you can use that URL scheme to communicate with the application. In most cases, however, URLs are used simply to launch another application and have it display some information that is relevant to your own application. For example, if your application manages address information, you could send a URL containing a given address to the Maps application to show that location. This level of communication creates a much more integrated environment for the user and alleviates the need for your application to implement features that exist elsewhere on the device.

Apple provides built-in support for the `http`, `mailto`, `tel`, and `sms` URL schemes. It also supports `http`–based URLs targeted at the Maps, YouTube, and iPod applications. Applications can register their own custom URL schemes as well. To communicate with an application, create an `NSURL` object with some properly formatted content and pass it to the `openURL:` method of the shared `UIApplication` object. The `openURL:` method launches the application that has registered to receive URLs of that type and passes it the URL. When the user subsequently quits that application, the system often relaunches your application but may not always do so. The decision to relaunch an application is made based on the user’s actions in the handler application and whether returning to your application would make sense from the user’s perspective.

The following code fragment illustrates how one application can request the services of another application (“todolist” in this example is a hypothetical custom scheme registered by an application):

```
NSURL *myURL = [NSURL
URLWithString:@"todolist://www.acme.com?Quarterly%20Report#200806231300"];
[[UIApplication sharedApplication] openURL:myURL];
```

Important: If your URL type includes a scheme that is identical to one defined by Apple, the Apple-provided application is launched instead of your application. If multiple third-party applications register to handle the same URL scheme, it is undefined as to which of the applications is picked to handle URLs of that type.

If your application defines its own custom URL scheme, it should implement a handler for that scheme as described in [“Implementing Custom URL Schemes”](#) (page 34). For more information about the system-supported URL schemes, including information about how to format the URLs, see *Apple URL Scheme Reference*.

Implementing Custom URL Schemes

You can register URL types for your application that include custom URL schemes. A custom URL scheme is a mechanism through which third-party applications can interact with each other and with the system. Through a custom URL scheme, an application can make its services available to other applications.

Registering Custom URL Schemes

To register a URL type for your application, you must specify the subproperties of the `CFBundleURLTypes` property, which was introduced in [“The Information Property List”](#) (page 25). The `CFBundleURLTypes` property is an array of dictionaries in the application’s `Info.plist` file, with each dictionary defining a URL type the application supports. Table 1-5 describes the keys and values of a `CFBundleURLTypes` dictionary.

Table 1-5 Keys and values of the `CFBundleURLTypes` property

Key	Value
<code>CFBundleURLName</code>	<p>A string that is the abstract name for the URL type. To ensure uniqueness, it is recommended that you specify a reverse-DNS style of identifier, for example, <code>com.acme.myscheme</code>.</p> <p>The URL-type name provided here is used as a key to a localized string in the <code>InfoPlist.strings</code> file in a language-localized bundle subdirectory. The localized string is the human-readable name of the URL type in a given language.</p>
<code>CFBundleURLSchemes</code>	An array of URL schemes for URLs belonging to this URL type. Each scheme is a string. URLs belonging to a given URL type are characterized by their scheme components.

Figure 1-7 shows the `Info.plist` file of an application being edited using the built-in Xcode editor. In this figure, the URL types entry in the left column is equivalent to the `CFBundleURLTypes` key you would add directly to an `Info.plist` file. Similarly, the “URL identifier” and “URL Schemes” entries are equivalent to the `CFBundleURLName` and `CFBundleURLSchemes` keys.

Figure 1-7 Defining a custom URL scheme in the `Info.plist` file

Key	Value
▼ Information Property List	(12 items)
Localization native develo	en
Bundle display name	\$(PRODUCT_NAME)
Executable file	\$(EXECUTABLE_NAME)
Icon file	
Bundle identifier	com.acme.\$(PRODUCT_NAME)
InfoDictionary version	6.0
Bundle name	\$(PRODUCT_NAME)
Bundle OS Type code	APPL
Bundle creator OS Type co	????
Bundle version	1.0
Main nib file base name	MainWindow
▼ URL types	(1 item)
▼ Item 1	(2 items)
URL identifier	com.acme.ToDoList
▼ URL Schemes	(1 item)
Item 1	todolist

After you have registered a URL type with a custom scheme by defining the `CFBundleURLTypes` property, you can test the scheme in the following way:

1. Build, install, and run your application.
2. Go to the Home screen and launch Safari. (In the iPhone simulator you can go to the Home screen by selecting Hardware > Home from the menu.)
3. In the address bar of Safari, type a URL that uses your custom scheme.
4. Verify that your application launches and that the application delegate is sent a `application:handleOpenURL: message`.

Handling URL Requests

The delegate of an application handles URL requests routed to the application by implementing the `application:handleOpenURL: method`. You especially need the delegate to implement this method if you have registered custom URL schemes for your application.

A URL request based on a custom scheme assumes a kind of protocol understood by those applications requesting the services of your application. The URL contains information of some kind that the scheme-registering application is expected to process or respond to in some way. Objects of the `NSURL` class, which are passed into the `application:handleOpenURL: method`, represent URLs in the Cocoa Touch framework. `NSURL` conforms to the RFC 1808 specification; it includes methods that return the various parts of a URL as defined by RFC 1808, including user, password, query, fragment, and parameter string. The “protocol” for your custom scheme can use these URL parts for conveying various kinds of information.

In the implementation of `application:handleOpenURL:` shown in Listing 1-2, the passed-in URL object conveys application-specific information in its query and fragment parts. The delegate extracts this information—in this case, the name of a to-do task and the date the task is due—and with it creates a model object of the application.

Listing 1-2 Handling a URL request based on a custom scheme

```
- (BOOL)application:(UIApplication *)application handleOpenURL:(NSURL *)url {
```

```

        if ([[url scheme] isEqualToString:@"todolist"]) {
            ToDoItem *item = [[ToDoItem alloc] init];
            NSString *taskName = [url query];
            if (!taskName || ![self isValidTaskString:taskName]) { // must have a
task name
                [item release];
                return NO;
            }
            taskName = [taskName
stringByReplacingPercentEscapesUsingEncoding:NSUTF8StringEncoding];

            item.toDoTask = taskName;
            NSString *dateString = [url fragment];
            if (!dateString || [dateString isEqualToString:@"today"]) {
                item.dateDue = [NSDate date];
            } else {
                if (![self isValidDateString:dateString]) {
                    [item release];
                    return NO;
                }
                // format: yyyyymmddhhmm (24-hour clock)
                NSString *curStr = [dateString substringWithRange:NSMakeRange(0,
4)];

                NSInteger yeardigit = [curStr integerValue];
                curStr = [dateString substringWithRange:NSMakeRange(4, 2)];
                NSInteger monthdigit = [curStr integerValue];
                curStr = [dateString substringWithRange:NSMakeRange(6, 2)];
                NSInteger daydigit = [curStr integerValue];
                curStr = [dateString substringWithRange:NSMakeRange(8, 2)];
                NSInteger hourdigit = [curStr integerValue];
                curStr = [dateString substringWithRange:NSMakeRange(10, 2)];
                NSInteger minutedigit = [curStr integerValue];

                NSDateComponents *dateComps = [[NSDateComponents alloc] init];
                [dateComps setYear:yeardigit];
                [dateComps setMonth:monthdigit];
                [dateComps setDay:daydigit];
                [dateComps setHour:hourdigit];
                [dateComps setMinute:minutedigit];
                NSCalendar *calendar = [NSCalendar currentCalendar];
                NSDate *itemDate = [calendar dateFromComponents:dateComps];
                if (!itemDate) {
                    [dateComps release];
                    return NO;
                }
                item.dateDue = itemDate;
                [dateComps release];
            }

            [(NSMutableArray *)self.list addObject:item];
            [item release];
            return YES;
        }
        return NO;
    }
}

```

Be sure to validate the input you get from URLs passed to your application; see *Validating Input* in *Secure Coding Guide* to find out how to avoid problems related to URL handling. To learn about URL schemes defined by Apple, see *Apple URL Scheme Reference*.

Displaying Application Preferences

If your application uses preferences to control various aspects of its behavior, how you expose those preferences to the user depends on how integral they are to your program.

- Preferences that are integral to using the application (and simple enough to implement directly) should be presented directly by your application using a custom interface.
- Preferences that are not integral, or that require a relatively complex interface, should be into the system's Settings application.

When determining whether a set of preferences is integral, think about the intended usage pattern. If you expect the user to make changes to preferences somewhat frequently, or if those preferences play a relatively important role in how the application behaves, they are probably integral. For example, the settings in a game are usually integral to playing the game and something the user might want to change quickly. Because the Settings application is a separate application, however, you would use it only for preferences that you do not expect the user to access frequently.

If you choose to implement preferences inside your application, it is up to you to define the interface and write the code to manage those preference. If you choose to use the Settings application, however, your application must provide a Settings bundle to manage them.

A **settings bundle** is a custom resource you include in the top level of your application's bundle directory. An opaque directory with the name `Settings.bundle`, the settings bundle contains specially formatted data files (and supporting resources) that tell the Settings application how to display your preferences. These files also tell the Settings application where to store the resulting values in the preferences database, which your application then accesses using the `NSUserDefaults` or `CFPreferences` APIs.

If you implement your preferences using a settings bundle, you should also provide a custom icon for your preferences. The Settings application looks for an image file with the name `Icon-Settings.png` at the top level of your application bundle and displays that image next to your application name. The image file should be a 29 x 29 pixel PNG image file. If you do not provide this file at the top level of your application bundle, the Settings application uses a scaled version of your application icon (`Icon.png`) by default.

For more information about creating a Settings bundle for your application, see [“Application Preferences”](#) (page 151).

Internationalizing Your Application

Ideally, the text, images, and other content that iPhone applications display to users should be localized for multiple languages. The text that an alert dialog displays, for example, should be in the preferred language of the user. The process of preparing a project for content localized for particular languages is called **internationalization**. Project components that are candidates for localization include:

- Code-generated text, including locale-specific aspects of date, time, and number formatting
- Static text—for example, an HTML file loaded into a web view for displaying application help

- Icons (including your application icon) and other images when those images either contain text or have some culture-specific meaning
- Sound files containing spoken language
- Nib files

Using the Settings application, users select the language they want to see in applications' user interfaces from the Language preferences view (see Figure 1-8). They get to this view from the International group of settings, accessed from General settings.

Figure 1-8 The Language preference view



The chosen language is associated with a subdirectory of the bundle. The subdirectory name has two components: an ISO 639-1 language code and a `.lproj` suffix. The language code may also be modified with a particular region by appending (after an underscore) an ISO 3166-1 region designator. For example, to designate resources localized to English as spoken in the United States, the bundle would be named `en_US.lproj`. By convention, these language-localized subdirectories are called `lproj` folders.

Note: You may use ISO 639-2 language codes instead of those defined by ISO 639-1. See “Language and Locale Designations” in *Internationalization Programming Topics* for descriptions of language and region codes.

An `lproj` folder contains all the localized content for a given language and, possibly, a given region. You use the facilities of the `NSBundle` class or the `CFBundle` opaque type to locate (in one of the application’s `lproj` folders) resources localized for the currently selected language. Listing 1-3 gives an example of such a directory containing content localized for English (`en`).

Listing 1-3 The contents of a language-localized subdirectory

```
en.lproj/
```

```
InfoPlist.strings
Localizable.strings
sign.png
```

This subdirectory example has the following items:

- The `InfoPlist.strings` file contains strings assigned as localized values of certain properties in the project's `Info.plist` file (such as `CFBundleDisplayName`). For example, the `CFBundleDisplayName` key for an application named *Battleship* in the English version would have this entry in `InfoPlist.strings` in the `fr.lproj` subdirectory:

```
CFBundleDisplayName = "Cuirassé";
```

- The `Localizable.strings` file contains localized versions of strings generated by application code.
- The `sign.png` file in this example is a file containing a localized image.

To internationalize strings in your code for localization, use the `NSLocalizedString` macro in place of the string. This macro has the following declaration:

```
NSString *NSLocalizedString(NSString *key, NSString *comment);
```

The first parameter is a unique key to a localized string in a `Localizable.strings` file in a given `lproj` folder. The second parameter is a comment that indicates how the string is used and therefore provides additional context to the translator. For example, suppose you are setting the content of a label (`UILabel` object) in your user interface. The following code would internationalize the label's text:

```
label.text = NSLocalizedString(@"City", @"Label for City text field");
```

You can then create a `Localizable.strings` file for a given language and add it to the proper `lproj` folder. For the above key, this file would have an entry similar to the following:

```
"City" = "Ville";
```

Note: Alternatively, you can insert `NSLocalizedString` calls in your code where appropriate and then run the `genstrings` command-line tool. This tool generates a `Localizable.strings` template that includes the key and comment for each string requiring translation. For further information about `genstrings`, see the `genstrings(1)` man page.

To find out more about internationalization, see *Internationalization Programming Topics*.

Tuning for Performance and Responsiveness

At each step in the development of your application, you should consider the implications of your design choices on the overall performance of your application. The operating environment for iPhone applications is more constrained because of the mobile nature of iPhone and iPod touch devices. The following sections describe the factors you should consider throughout the development process.

Using Memory Efficiently

Because the iPhone OS virtual memory model does not include disk swap space, applications are somewhat more limited in the amount of memory they have available for use. Using large amounts of memory can seriously degrade system performance and potentially cause the system to terminate your application. When you design, therefore, make it a high priority to reduce the amount of memory used by your application.

There is a direct correlation between the amount of free memory available and the relative performance of your application. Less free memory means that the system is more likely to have trouble fulfilling future memory requests. If that happens, the system can always remove code pages and other nonvolatile resources from memory. However, removing those resources may only be a temporary fix, especially when those resources are needed again a short time later. Instead, minimize your memory use in the first place, and clean up the memory you do use in a timely manner.

The following sections provide more guidance on how to use memory efficiently and how to respond when there is only a small amount of available memory.

Reducing Your Application's Memory Footprint

Table 1-6 lists some tips on how to reduce your application's overall memory footprint. Starting off with a low footprint gives you more room for the data you need to manipulate.

Table 1-6 Tips for reducing your application's memory footprint

Tip	Actions to take
Eliminate memory leaks.	Because memory is a critical resource in iPhone OS, your application should not have any memory leaks. Allowing leaks to exist means your application may not have the memory it needs later. You can use the Instruments application to track down leaks in your code, both in the simulator and on actual devices. For more information on using Instruments, see <i>Instruments User Guide</i> .
Make resource files as small as possible.	Files reside on the disk but must be loaded into memory before they can be used. Property list files and images are two resource types where you can save space with some very simple actions. To reduce the space used by property list files, write those files out in a binary format using the <code>NSPropertyListSerialization</code> class. For images, compress all image files to make them as small as possible. (To compress PNG images—the preferred image format for iPhone applications—use the <code>pngcrush</code> tool.)
Use SQLite for large data sets.	If your application manipulates large amounts of structured data, store it in a SQLite database instead of in a flat file. SQLite provides efficient ways to manage large data sets without requiring the entire set to be in memory all at once.
Load resources lazily.	You should never load a resource file until it is actually needed. Prefetching resource files may seem like a way to save time, but this practice actually slows down your application right away. In addition, if you end up not using the resource, loading it simply wastes memory.
Build your program using Thumb.	Adding the <code>-mthumb</code> compiler flag can reduce the size of your code by up to 35%. Be sure to turn this option off for floating-point-intensive code modules, however, because the use of Thumb on these modules can cause performance to degrade.

Allocating Memory Wisely

iPhone applications use a managed memory model, whereby you must explicitly retain and release objects. Table 1-7 lists tips for allocating memory inside your program.

Table 1-7 Tips for allocating memory

Tip	Actions to take
Reduce your use of autoreleased objects.	Objects released using the <code>autorelease</code> method stay in memory until you explicitly drain the autorelease pool or until the next time around your event loop. Whenever possible, avoid using the <code>autorelease</code> method when you can instead use the <code>release</code> method to reclaim the memory occupied by the object immediately. If you must create moderate numbers of autoreleased objects, create a local autorelease pool and drain it periodically to reclaim the memory for those objects before the next event loop.
Impose size limits on resources.	Avoid loading large resource files when a smaller one will do. Instead of using a high-resolution image, use one that is appropriately sized for iPhone OS–based devices. If you must use large resource files, find ways to load only the portion of the file that you need at any given time. For example, rather than load the entire file into memory, use the <code>mmap</code> and <code>munmap</code> functions to map portions of the file into and out of memory. For more information about mapping files into memory, see <i>File-System Performance Guidelines</i> .
Avoid unbounded problem sets.	Unbounded problem sets might require an arbitrarily large amount of data to compute. If the set requires more memory than is available, your application may be unable to complete the calculations. Your applications should avoid such sets whenever possible and work on problems with known memory limits.

For detailed information on how to allocate memory in iPhone applications, and for more information on autorelease pools, see Cocoa Objects in *Cocoa Fundamentals Guide*.

Floating-Point Math Considerations

The processors found in iPhone–OS based devices are capable of performing floating-point calculations in hardware. If you have an existing program that performs calculations using a software-based fixed-point math library, you should consider modifying your code to use floating-point math instead. Hardware-based floating-point computations are typically much faster than their software-based fixed-point equivalents.

Important: Of course, if your code does use floating-point math extensively, remember to compile that code without the `-mthumb` compiler option. The Thumb option can reduce the size of code modules but it can also cause degrade the performance of floating-point code.

Reducing Power Consumption

Power consumption on mobile devices is always an issue. The power management system in iPhone OS conserves power by shutting down any hardware features that are not currently being used. In addition to avoiding CPU-intensive operations or operations that involve high graphics frame rates, you can help improve battery life by minimizing the use of the following features:

- Wi-Fi radios
- The baseband cell radios
- The Core Location framework
- The accelerometers
- The disk

The goal of minimizing your use of these features is not to restrict you unnecessarily as you create your application. But as you design it, remember that everything you do on a device has implications for the user's overall battery life. The more code your application executes, the longer the CPU consumes power. The more data you transmit to the network, the more power must be used to run the radios. In fact, accessing the network is the most power-hungry operation you can perform and should be minimized by following these guidelines:

- Connect to external network servers only when needed, and do not poll.
- When you must connect to the network, transmit the smallest amount of data possible.
- Transmit data in bursts rather than spreading out transmission packets over time. The system turns off the Wi-Fi and cell radios when it detects a lack of activity. When it transmits data over a longer period of time, your application uses much more power than when it transmits the same amount of data in a shorter amount of time.
- If you use the Core Location framework to gather location data, disable location updates as soon as you have the data you need. Core Location uses the available GPS, cell, and Wi-Fi networks to determine the user's location. Although Core Location works hard to minimize the use of these radios, turning off the feature altogether is the best way to save power.

Optimizing your application's performance is another good way to minimize the power you use. An application that does work lazily and uses the least code possible is going to be much more power efficient than an unoptimized application. A good design can also help minimize the time the user spends inside your application, which is also a good way to conserve battery life.

Tuning Your Code

iPhone OS comes with several applications for tuning the performance of your application. Most of these tools run on Mac OS X and are suitable for tuning some aspects of your code while it runs in the simulator. For example, you can use the simulator to eliminate memory leaks and make sure your overall memory usage is as low as possible. You can also remove any computational hotspots in your code that might be caused by an inefficient algorithm or a previously unknown bottleneck.

After you have tuned your code in the simulator, you should then use the Instruments application to further tune your code on a device. Running your code on an actual device is the only way to tune your code fully. Because the simulator runs in Mac OS X, it has the advantage of a faster CPU and more usable memory, so its performance is generally much better than the performance on an actual device. And using Instruments to trace your code on an actual device may point out additional performance bottlenecks that need tuning.

For more information on using Instruments, see *Instruments User Guide*.

Window and Views

Windows and views are the visual components you use to construct the interface of your iPhone application. Windows provide the background platform for displaying content but views do most of the work of drawing that content and responding to user interactions. Although this chapter covers the concepts associated with both windows and views, it focuses more on views because of their importance to the system.

Because views play such a vital role in iPhone applications, there is no way to cover every aspect of them in a single chapter. This chapter focuses on the basic properties of windows and views, their relationships to each other, and how you create and manipulate them in your application. This chapter does not cover how views respond to touch events or draw custom content. For more information about those subjects, see [“Event Handling”](#) (page 75) and [“Graphics and Drawing”](#) (page 85) respectively.

What Are Windows and Views?

Like Mac OS X, iPhone OS uses windows and views to present graphical content on the screen. Although there are many similarities between the window and view objects on both platforms, the roles played by both windows and views differ slightly on each platform.

The Role of UIWindow

In contrast with Mac OS X applications, iPhone applications typically have only one **window**, represented by an instance of the `UIWindow` class. Your application creates this window at launch time (or loads it from a nib file), adds one or more views to it, and displays it. After that, you rarely need to refer to the window object again.

In iPhone OS, a window object has no visual adornments such as a close box or title bar and cannot be closed or manipulated directly by the user. All manipulations to a window occur through its programmatic interfaces. The application also uses the window to facilitate the delivery of events to your application. For example, the window object keeps track of its current first responder object and dispatches events to it when asked to do so by the `UIApplication` object.

One thing that experienced Mac OS X developers may find unusual about the `UIWindow` class is its inheritance. In Mac OS X, the parent class of `NSWindow` is `NSResponder`. In iPhone OS, the parent class of `UIWindow` is `UIView`. Thus, in iPhone OS, a window is also a view object. Despite its parentage, you typically treat windows in iPhone OS the same as you would in Mac OS X. That is, you typically do not manipulate the view-related properties of a `UIWindow` object directly.

When creating your application window, you should always set its initial frame size to fill the entire screen. If you load your window from a nib file, Interface Builder does not permit you to create a window smaller than the screen size. If you create your window programmatically, however, you must specifically pass in the desired frame rectangle at creation time. There is no reason to pass in any rectangle other than the screen rectangle, which you can get from the `UIScreen` object as shown here:

```
UIWindow* aWindow = [[[UIWindow alloc] initWithFrame:[UIScreen mainScreen]  
bounds]] autorelease];
```

Although iPhone OS supports layering windows on top of each other, your application should never create more than one window. The system itself uses additional windows to display the system status bar, important alerts, and other types of messages on top of your application's windows. If you want to display alerts on top of your content, use the alert views provided by UIKit rather than creating additional windows.

The Role of UIView

A **view**, an instance of the `UIView` class, defines a rectangular area on the screen. In iPhone applications, views play a key role in both presenting your interface and responding to interactions with that interface. Each view object has the responsibility of rendering content within its rectangular area and for responding to touch events in that area. This dual behavior means that views are the primary mechanism for interacting with the user in your application. In a Model-View-Controller application, view objects obviously are the View portion of the application.

In addition to displaying its own contents and handling events, a view may also manage one or more subviews. A **subview** is simply a view object embedded inside the frame of the original view object, which is referred to as the parent view or **superview**. Views arranged in this manner form what is known as a **view hierarchy** and may contain any number of views. Views can also be nested at arbitrarily deep levels by adding subviews to subviews. The organization of views inside the view hierarchy controls what appears on screen, as each subview is displayed on top of its parent view. The organization also controls how the views react to events and changes. Each parent view is responsible for managing its direct subviews, by adjusting their position and size as needed and even responding to events that its subviews do not handle.

Because view objects are the main way your application interacts with the user, they have a number of responsibilities. Here are just a few:


- Drawing and animation
 - Views draw content in their rectangular area.
 - Some view properties can be animated to new values.
- Layout and subview management
 - Views manage a list of subviews.
 - Views define their own resizing behaviors in relation to their parent view.
 - Views can manually change the size and position of their subviews as needed.
 - Views can convert points in their coordinate system to the coordinate systems of other views or the window.
- Event handling
 - Views receive touch events.
 - Views participate in the responder chain.

In iPhone applications, views work closely with view controllers to manage several aspects of the views behavior. View controllers handle the loading and unloading of views, interface rotations caused by the user physically rotating the device, and interactions with the high-level navigation objects used to construct complex user interfaces. For more information, see [“The Role of View Controllers”](#) (page 50).

Most of this chapter is dedicated to explaining these responsibilities and showing you how to tie your own custom code into the existing `UIView` behaviors.

UIKit View Classes

The `UIView` class defines the basic properties of a view but not its visual representation. Instead, UIKit uses subclasses to define the specific appearance and behavior for standard system elements such as text fields, buttons, and toolbars. Figure 2-1 shows the class hierarchy diagram for all of the views in UIKit. With the exception of the `UIView` and `UIControl` classes, most of the views in this hierarchy are designed to be used as-is or in conjunction with a delegate object.

Figure 2-1 View class hierarchy

This view hierarchy can be broken down into the following broad categories:

- Containers

Container views enhance the function of other views or provide additional visual separation of the content. For example, the `UIScrollView` class is used to display views whose contents are too large to fit onscreen all at once. The `UITableView` class is a subclass of `UIScrollView` that manages lists of data. Because table rows are selectable, tables are commonly used for hierarchical navigation too—for example, to drill down into a hierarchy of objects.

A `UIToolbar` object is a special type of container that visually groups one or more button-like items. A toolbar typically appears along the bottom of the screen. The Safari, Mail, and Photos applications all use toolbars to display buttons representing frequently used commands. Toolbars can be shown all the time or only as needed by the application.

■ Controls

Controls are used to create most of a typical application's user interface. A control is a special type of view that inherits from the `UIControl` superclass. Controls typically display a specific value and handle all of the user interactions required to modify that value. Controls also use standard system paradigms, such as target-action and delegation, to notify your application when user interactions occur. Controls include buttons, text fields, sliders, and switches.

■ Display views

Although controls and many other types of views provide interactive behavior, some views simply display information. The UIKit classes that exhibit this behavior include `UIImageView`, `UILabel`, `UIProgressView`, and `UIActivityIndicatorView`.

■ Text and web views

Text and web views provide a more sophisticated way to display multiline text content in your application. The `UITextView` class supports the display and editing of multiple lines of text in a scrollable area. The `UIWebView` class provides a way to display HTML content, which lets you incorporate graphics and advanced text-formatting options and lay out your content in custom ways.

■ Alert views and action sheets

Alert views and action sheets are used to get the user's attention immediately. They present a message to the user, along with one or more optional buttons that the user can use to respond to the message. Alert views and action sheets are similar in function but look and behave differently. For example, the `UIAlertView` class displays a blue alert box that pops up on the screen and the `UIActionSheet` class displays a box that slides in from the bottom of the screen.

■ Navigation views

Tab bars and navigation bars work in conjunction with view controllers to provide tools for navigating from one screen of your user interface to another. You typically do not create `UITabBar` and `UINavigationController` items directly but configure them through the appropriate controller interface or using Interface Builder instead.

■ The window

A window provides a surface for drawing content and is the root container for all other views. There is typically only one window per application. For more information, see [“The Role of UIWindow”](#) (page 45).

In addition to views, UIKit provides view controllers to manage those objects. For more information, see [“The Role of View Controllers”](#) (page 50).

The Role of View Controllers

Applications running in iPhone OS have many options for organizing their content and presenting it to the user. An application that contains a lot of content might divide that content up into multiple screens' worth of information. At runtime, each screen would then be backed by a set of view objects responsible for displaying the data for that particular screen. The views for a single screen would themselves be backed by a view controller object, whose job is to manage the data displayed by those views and coordinate updates with the rest of the application.

The `UIViewController` class is responsible for creating the set of views it manages and for flushing them from memory during low-memory situations. View controllers also provide automatic responses for some standard system behaviors. For example, in response to a change in the device's orientation, the view controller can resize its managed views to fit the new orientation, if that orientation is supported. You can also use view controllers to display new views modally on top of the current view.

In addition to the base `UIViewController` class, UIKit includes more advanced subclasses for handling some of the sophisticated interface behaviors common to the platform. In particular, navigation controllers manage the display of multiple hierarchical screens worth of content. Tab bar controllers let the user switch between different sets of screens, each of which represents a different operating mode for the application.

For information on how to use view controllers to manage the views in your user interface, see *View Controller Programming Guide for iPhone OS*.

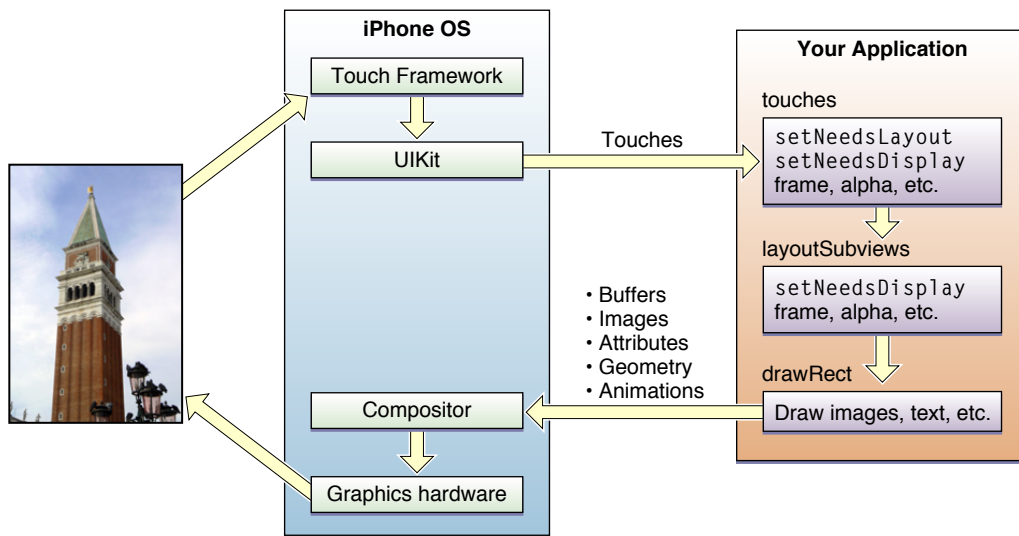
View Architecture and Geometry

Because views are focal objects in iPhone applications, it is important to understand a little about how views interact with other parts of the system. The standard view classes in UIKit provide a considerable amount of behavior to your application for free. They also provide well-defined integration points where you can customize that behavior and do what you need to do for your application.

The following sections explain the standard behavior of views and call out the places where you can integrate your custom code. For information about the integration points of specific classes, see the reference document for that class. You can get a list of all the class reference documents in *UIKit Framework Reference*.

The View Interaction Model

Any time a user interacts with your user interface, or your own code programmatically changes something, a complex sequence of events takes place inside of UIKit to handle that interaction. At specific points during that sequence, UIKit calls out to your view classes and gives them a chance to respond on behalf of your application. Understanding these callout points is important to understanding where your views fit into the system. Figure 2-2 shows the basic sequence of events that starts with the user touching the screen and ends with the graphics system updating the screen content in response. Programmatic events follow the same basic steps without the initial user interaction.

Figure 2-2 UIKit interactions with your view objects

The following steps break the event sequence in [Figure 2-2](#) (page 51) down even further and explain what happens at each stage and how your application might want to react in response.

1. The user touches the screen.
2. The hardware reports the touch event to the UIKit framework.
3. The UIKit framework packages the touch into a `UIEvent` object and dispatches it to the appropriate view. (For a detailed explanation of how UIKit delivers events to your views, see [“Event Delivery”](#) (page 77).)
4. The event-handling methods of your view might respond to the event by doing any of the following:
 - Adjust the properties (frame, bounds, alpha, and so on) of the view or its subviews.
 - Mark the view (or its subviews) as needing a change in its layout.
 - Mark the view (or its subviews) as needing to be redrawn.
 - Notify a controller about changes to some piece of data.

Of course, it is up to the view to decide which of these things must be done and call the appropriate methods to do it.

5. If a view is marked as requiring layout, UIKit calls the view’s `layoutSubviews` method.

You can override this method in your custom views and use it to adjust the position and size of any subviews. For example, a view that provides a large scrollable area would need to use several subviews as “tiles” rather than create one large view, which is not likely to fit in memory anyway. In its implementation of this method, the view would hide any subviews that are now offscreen or reposition them and use them to draw newly exposed content. As part of this process, the view can also mark the new tiles as needing to be redrawn.

6. If any part of the view is marked as needing to be redrawn, UIKit calls the view’s `drawRect:` method.

UIKit calls this method for only those views that need it. Each view's implementation of this method should redraw the specified area as quickly as possible. Each view should draw only its own contents and not the contents of any subviews. Views should not attempt to make any further changes to their properties or layout at this point.

7. Any updated views are composited with the rest of visible content and sent to the graphics hardware for display.
8. The graphics hardware transfers the rendered content to the screen.

Note: The preceding update model applies primarily to applications that use native views and drawing techniques. If your application draws its content using OpenGL ES, you would typically configure a single full-screen view and then draw directly to your OpenGL graphics context. Your view would still handle touch events, but it would not need to lay out subviews or implement a `drawRect:` method. For more information about using OpenGL ES, see [“Drawing with OpenGL ES”](#) (page 94).

Given the preceding set of steps, the primary integration points for your own custom views are as follows:

1. These event-handling methods:
 - `touchesBegan:withEvent:`
 - `touchesMoved:withEvent:`
 - `touchesEnded:withEvent:`
 - `touchesCancelled:withEvent:`
2. The `layoutSubviews` method
3. The `drawRect:` method

These are the methods that most custom views implement to get the behavior they want; you may not need to override all of them. For example, if you are implementing a view whose size never changes, you might not need to override the `layoutSubviews` method. Similarly, if you are implementing a view that displays simple content, such as text and images, you can often avoid drawing altogether by simply embedding `UIImageView` and `UILabel` objects as subviews.

It is also important to remember that these are the primary integration points but not the only ones. Several methods of the `UIView` class are designed to be override points for subclasses. You should look at the method descriptions in *UIView Class Reference* to see which methods might be appropriate for you to override in your custom implementations.

The View Rendering Architecture

Although you use views to represent content onscreen, the `UIView` class itself actually relies heavily on another object for much of its basic behavior. Each view object in UIKit is backed by a Core Animation layer object, which is an instance of the `CALayer` class. This layer class provides the fundamental support for the layout and rendering of a view's contents and for compositing and animating that content.

In contrast with Mac OS X (in which Core Animation support is optional) iPhone OS integrates Core Animation into the heart of the view rendering implementation. Although Core Animation has a central role, UIKit streamlines the programming experience by providing a transparent layer on top of Core Animation. This transparent layer eliminates the need to access Core Animation layers directly most of the time, instead letting you access similar behaviors using the methods and properties of the `UIView` class. Where Core Animation becomes important, however, is when the `UIView` class does not provide everything you need. At that point, you can dive down into the Core Animation layers and do some pretty sophisticated rendering for your application.

The following sections provide an introduction to Core Animation and describe some of the features it provides to you for free through the `UIView` class. For more detailed information about how to use Core Animation for advanced rendering, see *Core Animation Programming Guide*.

Core Animation Basics

Core Animation takes advantage of hardware acceleration and an optimized architecture to implement fast rendering and real-time animations. The first time a view's `drawRect:` method is called, the layer captures the results into a bitmap. Subsequent redraw calls use this cached bitmap whenever possible to avoid calling the `drawRect:` method, which can be expensive. This process allows Core Animation to optimize its compositing operations and deliver the desired performance.

Core Animation stores the layers associated with your view objects in a hierarchy referred to as the **layer tree**. Like views, each layer in the layer tree has a single parent and can have any number of embedded sublayers. By default, objects in the layer tree are organized exactly like the views in your view hierarchy. You can add layers, however, without adding a corresponding view. You might do this to implement special visual effects for which a view is not required.

Layer objects are actually the driving force behind the rendering and layout system in iPhone OS, and most view properties are actually thin wrappers for properties on the underlying layer object. When you change the property of a layer in the layer tree (directly using the `CALayer` object), the changed value is reflected immediately in the layer object. If the change triggers a corresponding animation, however, that change may not be reflected onscreen immediately; instead, it must be animated onto the screen over time. To manage these sorts of animations, Core Animation maintains two additional sets of layer objects in what are referred to as the **presentation tree** and the **render tree**.

The presentation tree reflects the state of the layers as they are currently presented to the user. When you animate the changing of a layer value, the presentation layer reflects the old value until the animation commences. As the animation progresses, Core Animation updates the value in the presentation-tree layer based on the current frame of the animation. The render tree then works together with the presentation tree to render the changes on the screen. Because the render tree runs in a separate process or thread, the work it does does not impact your application's main run loop. While both the layer tree and the presentation tree are public, the render tree is a private API.

The placement of layer objects behind your views has many important implications for the performance of your drawing code. The upside to using layers is that most geometry changes to your views do not require redrawing. For example, changing the position and size of a view does not require the system to redraw the contents of a view; it can simply reuse the cached bitmap created by the layer. Animating this cached content is significantly more efficient than trying to redraw that content every time.

The downside to using layers is that the additional cached data can add memory pressure to your application. If your application creates too many views or creates very large views, you could run out of memory quickly. You should not be afraid to use views in your application, but do not create new view objects if you have existing views that can be reused. In other words, pursue approaches that minimize the number of views you keep in memory at the same time.

For a more detailed overview of Core Animation, the object trees, and how you create animations, see *Core Animation Programming Guide*.

Changing the Layer of a View

Because views are required to have an associated layer object in iPhone OS, the `UIView` class creates this layer automatically at initialization time. You can access the layer that is created through the `layer` property of the view, but you cannot change the layer object after the view is created.

If you want a view to use a different type of layer, you must override the view's `layerClass` class method and return the class object for the layer you want it to use. The most common reason to return a different layer class is to implement an OpenGL-based application. To use OpenGL drawing commands, the layer for the underlying view must be an instance of the `CAEAGLLayer` class. This type of layer interacts with the OpenGL rendering calls to present the desired content on the screen.

Important: You should never modify the delegate property of a view's layer; that property stores a pointer to the view and should be considered private. Similarly, because a view can operate as the delegate for only one layer, you must not assign it as the delegate for any other layer objects. Doing so will cause your application to crash.

Animation Support

One of the benefits of having a layer object behind every view in iPhone OS is that you can animate content more easily. Remember that animation is not necessarily about creating visual eye candy. Animations provide the user with a context for any changes that occur in your application's user interface. For example, when you use a transition to move from one screen to another, you are indicating to users that the screens are related. The system provides automatic support for many of the most commonly used animations, but you can also create animations for other parts of your interface.

Many properties of the `UIView` class are considered to be animatable. An **animatable** property is one for which there is semiautomatic support for animating from one value to another. You must still tell UIKit that you want to perform the animation, but Core Animation assumes full responsibility for running the animation once it has begun. Among the properties you can animate on a `UIView` object are the following:

- frame
- bounds
- center
- transform
- alpha


Even though other view properties are not directly animatable, you can create explicit animations for some of them. Explicit animations require you to do more of the work in managing the animation and the rendered contents, but they still use the underlying Core Animation infrastructure to obtain good performance.

For more information about creating animations using the `UIView` class, see [“Animating Views”](#) (page 67). For more information about creating explicit animations, see *Core Animation Programming Guide*.

View Coordinate Systems

Coordinates in UIKit are based on a coordinate system whose origin is in the top-left corner and whose coordinate axes extend down and to the right from that point. Coordinate values are represented using floating-point numbers, which allow for precise layout and positioning of content and allow for resolution independence. [Figure 2-3](#) (page 55) shows this coordinate system relative to the screen, but this coordinate system is also used by the `UIWindow` and `UIView` classes. This particular orientation was chosen to make it easier to lay out controls and content in user interfaces, even though it differs from the default coordinate systems in use by Quartz and Mac OS X.

Figure 2-3 View coordinate system



As you write your interface code, be aware of the coordinate system currently in effect. Every window and view object maintains its own local coordinate system. All drawing in a view occurs relative to the view's local coordinate system. The frame rectangle for each view, however, is specified using the coordinate system of its parent view, and coordinates delivered as part of an event object are specified relative to the coordinate system of the enclosing window. For convenience, the `UIWindow` and `UIView` classes each provide methods to convert back and forth between the coordinate systems of different objects.

Although the coordinate system used by Quartz does not use the top-left corner as the origin point, for many Quartz calls this is not a problem. Before invoking your view's `drawRect:` method, UIKit automatically configures the drawing environment to use a top-left origin. Quartz calls made within this environment draw correctly in your view. The only time you need to consider these different coordinate systems is when you set up the drawing environment yourself using Quartz.

For more information about coordinate systems, Quartz, and drawing in general, see [“Graphics and Drawing”](#) (page 85).


The Relationship of the Frame, Bounds, and Center

A view object tracks its size and location using its `frame`, `bounds`, and `center` properties. The `frame` property contains a rectangle, the **frame rectangle**, that specifies the view's location and size relative to its parent view's coordinate system. The `bounds` property contains a rectangle, the **bounds rectangle**, that defines the view's position and size relative to its own local coordinate system. And although the origin of the bounds rectangle is typically set to (0, 0), it need not be. The `center` property contains the **center point** of the frame rectangle.

You use the `frame`, `bounds`, and `center` properties for different purposes in your code. Because the bounds rectangle represents the view's local coordinate system, you use it most often during drawing or event-handling code when you need to know where in your view something happened. The center point represents the known center point of your view and is always the best way to manipulate the position of your view. The frame rectangle is a convenience value that is computed using the `bounds` and `center` point and is valid only when the view's transform is set to the identity transform.

Figure 2-4 shows the relationship between the frame and bounds rectangles. The complete image on the right is drawn in the view starting at (0, 0). Because the size of the bounds does not match the full size of the image, however, only part of the image outside the bounds rectangle is clipped automatically. When the view is composited with its parent view, the position of the view inside its parent is determined by the origin of the view's frame rectangle, which in this case is (5, 5). As a result, the view's contents appear shifted down and to the right from the parent view's origin.

Figure 2-4 Relationship between a view's frame and bounds




When there is no transform applied to the view, the location and size of the view are determined by these three interrelated properties. The `frame` property of a view is set when a view object is created programmatically using the `initWithFrame:` method. That method also initializes the `bounds` rectangle to originate at (0.0, 0.0) and have the same size as the view's frame. The `center` property is then set to the center point of the frame.

Although you can set the values of these properties independently, setting the value for one changes the others in the following ways:

- When you set the `frame` property, the size of the `bounds` property is set to match the size of the `frame` property. The `center` property is also adjusted to match the center point of the new frame.
- When you set the `center` property, the origin of the `frame` changes accordingly.
- When you set the size of the `bounds` rectangle, the size of the `frame` rectangle changes to match.

You can change the `bounds` origin without changing the other two properties. When you do, the view displays the portion of the underlying image that you have identified. In [Figure 2-4](#) (page 56), the original `bounds` origin is set to (0.0, 0.0). In [Figure 2-5](#), that origin is moved to (8.0, 24.0). As a result, a different portion of the underlying image is displayed by the view. Because the frame rectangle did not change, however, the new content is displayed in the same location inside the parent view as before.

Figure 2-5 Altering a view's bounds



Note: By default, a view's frame is not clipped to its parent view's frame. If you want to force a view to clip its subviews, set the view's `clipsToBounds` property to YES.

Coordinate System Transformations

Although coordinate system transformations are commonly used in a view's `drawRect:` method to facilitate drawing, in iPhone OS, you can also use them to implement visual effects for your view. For example, the `UIView` class includes a `transform` property that lets you apply different types of translation, scaling, and zooming effects to the entire view. By default, the value of this property is the identity transform, which causes no changes to the view. To add transformations, get the `CGAffineTransform` structure stored in this property, use the corresponding Core Graphics functions to apply the transformations, and then assign the modified transform structure back to the view's `transform` property.

Note: When applying transforms to a view, all transformations are performed relative to the center point of the view.

Translating a view shifts all subviews along with the drawing of the view's content. Because coordinate systems of subviews inherit and build on these alterations, scaling also affects the drawing of the subviews. For more information about how to control the scaling of view content, see [“Content Modes and Scaling”](#) (page 58).


Important: If the `transform` property is not the identity transform, the value of the `frame` property is undefined and must be ignored. After setting the transform, use the `bounds` and `center` properties to get the position and size of the view.

For information about using transforms in conjunction with your `drawRect:` method, see [“Coordinates and Coordinate Transforms”](#) (page 86). For information about the functions you use to modify the `CGAffineTransform` structure, see *CGAffineTransform Reference*.


Content Modes and Scaling

When you change the bounds of a view or apply a scaling factor to the `transform` property of a view, the frame rectangle is changed by a commensurate amount. Depending on the content mode associated with the view, the view’s content may also be scaled or repositioned to account for the changes. The view’s `contentMode` property determines the effect that bounds changes and scaling operations have on the view. By default, the value of this property is set to `UIViewContentModeScaleToFill`, which always causes the view’s contents to be scaled to fit the new frame size. For example, Figure 2-6 shows what happens when the horizontal scaling factor of the view is doubled.

Figure 2-6 View scaled using the scale-to-fill content mode



Scaling of your view’s content occurs because the first time a view is shown, its rendered contents are cached in the underlying layer. Rather than force the view to redraw itself every time its bounds change or a scaling factor is applied, UIKit uses the view’s content mode to determine how to display the cached content. Figure 2-7 compares the results of changing the bounds of a view or applying a scaling factor to it using several different content modes.

Figure 2-7 Content mode comparisons

Although applying a scaling factor always scales the view's contents, there are content modes that do not scale the view's contents when the bounds of the view change. Several `UIViewContentMode` constants (such as `UIViewContentModeTop` and `UIViewContentModeBottomRight`) display the current content in different corners or along different edges of the view. There is also a mode for displaying the content centered inside the view. Changing the bounds rectangle with one of these content modes in place simply moves the existing contents to the appropriate location inside the new bounds rectangle.

Do consider using content modes when you want to implement resizable controls in your application; by doing so you can avoid both control distortion and the writing of custom drawing code. Buttons and segmented controls are particularly suitable for content mode–based drawing. They typically use several images to create the appearance of the control. In addition to having two fixed-size end cap images, a button that can grow horizontally uses a stretchable center image that is only 1 pixel wide. By displaying each image in its own image view and setting the content mode of the stretchable middle image to `UIViewContentModeScaleToFill`, the button can grow in size without distorting the appearance of the end caps. More importantly, the images associated with each image view can be cached by Core Animation and animated without any custom drawing code, which results in much better performance.

Although content modes are good to avoid redrawing the contents of your view, you can also use the `UIViewContentModeRedraw` content mode when you specifically want control over the appearance of your view during scaling and resizing operations. Setting your view's content mode to this value forces Core Animation to invalidate your view's contents and call your view's `drawRect:` method rather than scale or resize them automatically.

Autoresizing Behaviors


When you change the frame rectangle of a view, the position and size of embedded subviews often needs to change to match the new size of the original view. If the `autoresizesSubviews` property of a view is set to YES, its subviews are automatically resized according to the values in the `autoresizingMask` property. (By default, the `autoresizesSubviews` property is NO, so you need to set this value to YES to use this feature.) Often, simply configuring the autoresizing mask for a view provides the appropriate behavior for an application. Otherwise, it is the application's responsibility to reposition and resize the subviews by overriding the `layoutSubviews` method.

To set a view's autoresizing behaviors, combine the desired autoresizing constants using a bitwise OR operator and assign the resulting value to the view's `autoresizingMask` property. Table 2-1 lists the autoresizing constants and describes how each one affects the size and placement of a given view. For example, to keep a view pinned to the lower-left corner of its superview, add the `UIViewAutoresizingFlexibleRightMargin` and `UIViewAutoresizingFlexibleTopMargin` constants and assign them to the `autoresizingMask` property. When more than one aspect along an axis is made flexible, the resize amount is distributed evenly among them.

Table 2-1 Autoresizing mask constants

Autoresizing mask	Description
<code>UIViewAutoresizingNone</code>	If set, the view doesn't autoresize.
<code>UIViewAutoresizingFlexibleHeight</code>	If set, the view's height changes proportionally to the change in the superview's height. Otherwise, the view's height does not change relative to the superview's height.
<code>UIViewAutoresizingFlexibleWidth</code>	If set, the view's width changes proportionally to the change in the superview's width. Otherwise, the view's width does not change relative to the superview's width.
<code>UIViewAutoresizingFlexibleLeftMargin</code>	If set, the view's left edge is repositioned proportionally to the change in the superview's width. Otherwise, the view's left edge remains in the same position relative to the superview's left edge.
<code>UIViewAutoresizingFlexibleRightMargin</code>	If set, the view's right edge is repositioned proportionally to the change in the superview's width. Otherwise, the view's right edge remains in the same position relative to the superview.
<code>UIViewAutoresizingFlexibleBottomMargin</code>	If set, the view's bottom edge is repositioned proportionally to the change in the superview's height. Otherwise, the view's bottom edge remains in the same position relative to the superview.
<code>UIViewAutoresizingFlexibleTopMargin</code>	If set, the view's top edge is repositioned proportionally to the change in the superview's height. Otherwise, the view's top edge remains in the same position relative to the superview.

Figure 2-8 provides a graphical representation of the position of the constant values. When one of these constants is omitted, the view's layout is fixed in that aspect; when a constant is included in the mask, the view's layout is flexible in that aspect.

Figure 2-8 View autoresizing mask constants

If you are using Interface Builder to configure your views, you can set the autoresizing behavior for each view by using the Autosizing controls in the Size inspector. Although the flexible width and height constants from the preceding figure have the same behavior as the Interface Builder springs located in the same position have, the behavior of the margin constants is effectively reversed. In other words, to apply the flexible right margin autoresizing behavior to a view in Interface Builder, you must leave the space on that side of the Autosizing control empty, not place a strut there. Fortunately, Interface Builder provides an animation to show you how changes to the autoresizing behaviors affect your view.

If the `autoresizesSubviews` property of a view is set to `NO`, any autoresizing behaviors set on the immediate subviews of that view are ignored. Similarly, if a subview's autoresizing mask is set to `UIViewAutoresizingNone`, the subview does not change size and so its immediate subviews are never resized either.


Note: For autoresizing to work correctly, the view's `transform` property must be set to the identity transform. The behavior is undefined if it is not.

Although autoresizing behaviors may be suitable for some layout needs, if you want more control over the layout of your views, you should override the `layoutSubviews` method in the appropriate view classes. For more information about managing the layout of your views, see [“Responding to Layout Changes”](#) (page 69).


Creating and Managing the View Hierarchy

Managing the view hierarchy of your user interface is a crucial part of developing your application's user interface. How you organize your views defines not only the way your application appears visually but also how your application responds to changes. The parent-child relationships in the view hierarchy help define the chain of objects that is responsible for handling touch events in your application. When the user rotates the device, parent-child relationships also help define how each view's size and position are altered by changes to the user interface orientation.

Figure 2-9 shows a simple example of how the layering of views creates a desired visual effect. In the case of the Clock application, tab-bar and navigation-bar views are mixed together with a custom view to implement the overall interface.

Figure 2-9 Layered views in the Clock application

If you look at the object relationships for the views in the Clock application, you see that they look something like the relationships shown in “Changing the Layer of a View.” The window object acts as the root view for the application’s tab bar, navigation bar, and custom view.

Figure 2-10 View hierarchy for the Clock application

There are several ways to build view hierarchies in iPhone applications, including graphically in Interface Builder and programmatically in your code. The following sections show you how to assemble your view hierarchies and, having done that, how to find views in the hierarchy and convert between different view coordinate systems.

Creating a View Object

The simplest way to create views is to use Interface Builder and load them from the resulting nib file. From Interface Builder's graphical environment, you can drag new views out of the library and drop them onto a window or another view and build your view hierarchies quickly. Because Interface Builder uses live view objects, when you build your interface graphically you see exactly how it will appear when you load it at runtime. And there is no need to write tedious code to allocate and initialize each view in your view hierarchy.

If you prefer not to use Interface Builder and nib files to create your views, you can create them programmatically. To create a new view object, allocate memory for the view object and send that object an `initWithFrame:` message to initialize it. For example, to create a new instance of the `UIView` class, which you could use as a container for other views, you would use the following code:

```
CGRect viewRect = CGRectMake(0, 0, 100, 100);
UIView* myView = [[UIView alloc] initWithFrame:viewRect];
```

Note: Although all system objects support the `initWithFrame:` message, some may have a preferred initialization method that you should use instead. For information about any custom initialization methods, see the reference documentation for the class.

The frame rectangle that you specify when you initialize the view represents the position and size of the view relative to its intended parent view. You must add views to a window or to another view to make them appear on the screen. When you do, UIKit uses the frame rectangle you specify to place the view inside its parent. For information on how to add views to your view hierarchy, see [“Adding and Removing Subviews”](#) (page 63).

Adding and Removing Subviews

Interface Builder is the most convenient way to build view hierarchies because it lets you see exactly how those views will appear at runtime. It then saves the view objects and their hierarchical relationships in a nib file, which the system uses at runtime to recreate the objects and relationships in your application. When a nib file is loaded, the system automatically calls the `UIView` methods needed to recreate the view hierarchy.

If you prefer not to use Interface Builder and nib files to create your view hierarchies, you can create them programmatically instead. A view that has required subviews should create them in its own `initWithFrame:` method to ensure that they are present and initialized with the view. Subviews that are part of your application design (and not required for the operation of your view) should be created outside of your view's initialization code. In iPhone applications, the two most common places to create views and subviews programmatically are the `applicationDidFinishLaunching:` method of your application delegate and the `loadView` method of your view controllers.

To manipulate views in the view hierarchy, you use the following methods:

- To add a subview to a parent, call the `addSubview:` method of the parent view. This method adds the subview to the end of the parent's list of subviews.
- To insert a subview in the middle of the parent's list of subviews, call any of the `insertSubview:...` methods of the parent view.
- To reorder existing subviews inside their parent, call the `bringSubviewToFront:`, `sendSubviewToBack:`, or `exchangeSubviewAtIndex:withSubviewAtIndex:` methods of the parent view. Using these methods is faster than removing the subviews and reinserting them.

- To remove a subview from its parent, call the `removeFromSuperview` method of the subview (not the parent view).

When adding subviews, the current frame rectangle of the subview is used as the initial position of that view inside its parent. You can change that position at any time by changing the `frame` property of the subview. Subviews whose frame lies outside of their parent's visible bounds are not clipped by default. To enable clipping, you must set the `clipsToBounds` property of the parent view to `YES`.

Listing 2-1 shows a sample `applicationDidFinishLaunching:` method of an application delegate object. In this example, the application delegate creates its entire user interface programmatically at launch time. The interface consists of two generic `UIView` objects, which display primary colors. Each view is then embedded inside a window, which is also a subclass of `UIView` and can therefore act as a parent view. Because parents retain their subviews, this method releases the newly created views to prevent them from being overretained.

Listing 2-1 Creating a window with views

```
- (void)applicationDidFinishLaunching:(UIApplication *)application {
    // Create the window object and assign it to the
    // window instance variable of the application delegate.
    UIWindow *window = [[UIWindow alloc] initWithFrame:[UIScreen mainScreen] bounds];
    window.backgroundColor = [UIColor whiteColor];

    // Create a simple red square
    CGRect redFrame = CGRectMake(10, 10, 100, 100);
    UIView *redView = [[UIView alloc] initWithFrame:redFrame];
    redView.backgroundColor = [UIColor redColor];

    // Create a simple blue square
    CGRect blueFrame = CGRectMake(10, 150, 100, 100);
    UIView *blueView = [[UIView alloc] initWithFrame:blueFrame];
    blueView.backgroundColor = [UIColor blueColor];

    // Add the square views to the window
    [window addSubview:redView];
    [window addSubview:blueView];

    // Once added to the window, release the views to avoid the
    // extra retain count on each of them.
    [redView release];
    [blueView release];

    // Show the window.
    [window makeKeyAndVisible];
}
```


Important: When you're considering memory management, think of the subviews as any other collection object. Specifically, when you insert a view as a subview using `addSubview:`, that subview is retained by its superview. Inversely, when you remove the subview from its superview using the `removeFromSuperview` method, the subview is autoreleased. Releasing views after adding them to your view hierarchy prevents them being overretained, which could cause memory leaks.

For more information about Cocoa memory management conventions, see *Memory Management Programming Guide for Cocoa*.

When you add a subview to a parent view, UIKit sends several messages to both the parent and child to let them know what is happening. You can override methods such as `willMoveToSuperview:`, `willMoveToWindow:`, `willRemoveSubview:`, `didAddSubview:`, `didMoveToSuperview:`, and `didMoveToWindow` in your custom views to process changes before and after they occur and to update the state information in your view accordingly.

After you create a view hierarchy, you can use the `superview` property of a view to get its parent or the `subviews` property to get its children. You can also use the `isDescendantOfView:` method to determine whether a view is in the view hierarchy of a parent view. Because the root view in a view hierarchy has no parent, its `superview` property is set to `nil`. For views currently onscreen, the window object is typically the root view of the hierarchy.

You can use the `window` property of a view to get a pointer to the window that currently contains the view (if any). This property is set to `nil` if the view is not currently attached to a window.

Converting Coordinates in the View Hierarchy

At various times, particularly when handling events, an application may need to convert coordinate values from one frame of reference to another. For example, touch events usually report the touch location using the coordinate system of the window, but view objects need that information in the local coordinate system of the view, which may be different. The `UIView` class defines the following methods for converting coordinates to and from the view's local coordinate system:

```
convertPoint:fromView:
convertRect:fromView:
convertPoint:toView:
convertRect:toView:
```


The `convert...:fromView:` methods convert coordinates to the view's local coordinate system, while the `convert...:toView:` methods convert coordinates from the view's local coordinate system to the coordinate system of the specified view. If you specify `nil` as the reference view for any of the methods, the conversions are made to and from the coordinate system of the window that contains the view.

In addition to the `UIView` conversion methods, the `UIWindow` class also defines several conversion methods. These methods are similar to the `UIView` versions except that instead of converting to and from a view's local coordinate system, these methods convert to and from the window's coordinate system.

```
convertPoint:fromWindow:
convertRect:fromWindow:
convertPoint:toWindow:
convertRect:toWindow:
```

Coordinate conversions are straightforward when neither view is rotated or when dealing only with points. When converting rectangles or sizes between views with different rotations, the geometric structure must be altered in a reasonable way so that the resulting coordinates are correct. When converting a rectangle, the `UIView` class assumes that you want to guarantee coverage of the original screen area. To this end, the converted rectangle is enlarged so that when located in the appropriate view, it completely covers the original rectangle. Figure 2-11 shows the conversion of a rectangle in the `rotatedView` object's coordinate system to that of its superview, `outerView`.

Figure 2-11 Converting values in a rotated view



When converting size information, `UIView` simply treats it as a delta offset from (0.0, 0.0) that you need to convert from one view to another. Though the offset distance remains the same, the balance along the two axes shifts according to the rotation. When converting sizes, UIKit always returns sizes that consist of positive numbers.

Tagging Views

The `UIView` class contains a `tag` property that you can use to tag individual view objects with an integer value. You can use tags to uniquely identify views inside your view hierarchy and to perform searches for those views at runtime. (Tag-based searches are faster than iterating the view hierarchy yourself.) The default value for the `tag` property is 0.

To search for a tagged view, use the `viewWithTag:` method of `UIView`. This method searches the receiver's subviews using a depth-first search, starting with the receiver itself.

Modifying Views at Runtime

As applications receive input from the user, they adjust their user interface in response to that input. An application might rearrange the views in its interface, refresh existing views that contain changed data, or load an entirely new set of views. When deciding which techniques to use, consider your interface and what you are trying to achieve. How you initiate these techniques, however, is the same for all applications. The following sections describe these techniques and how you use them to update your user interface at runtime.

Note: For background information about how UIKit moves events and messages between itself and your custom code, see [“The View Interaction Model”](#) (page 50) before proceeding.

Animating Views

Animations provide fluid visual transitions between different states of your user interface. In iPhone OS, animations are used extensively to reposition views, change their size, and even change their alpha value to make them fade in or out. Because this support is crucial for making easy-to-use applications, UIKit simplifies the process of creating animations by integrating support for them directly into the `UIView` class.

The `UIView` class defines several properties that are inherently **animatable**—that is, the view provides built-in support for animating changes in the property from their current value to a new value. Although the work needed to perform the animation is handled for you automatically by the `UIView` class, you must still let the view know that you want the animation to happen. You do this by wrapping changes to the given property in an animation block.

An **animation block** starts with a call to the `beginAnimations:context:` class method of `UIView` and ends with a call to the `commitAnimations` class method. Between these calls, you configure the animation parameters and change the properties you want to animate. As soon as you call the `commitAnimations` method, UIKit performs the animations, animating changes from their current values to the new values you just set. Animation blocks can be nested, but nested animations do not start until the outermost animation block is committed.

Table 2-2 lists the animatable properties of the `UIView` class.

Table 2-2 Animatable properties

Property	Description
<code>frame</code>	The view’s frame rectangle, in superview coordinates.
<code>bounds</code>	The view’s bounding rectangle, in view coordinates.
<code>center</code>	The center of the frame, in superview coordinates.
<code>transform</code>	The transform applied to the view, relative to the center of its bounds.
<code>alpha</code>	The view’s alpha value, which determines the view’s level of transparency.

Configuring Animation Parameters

In addition to changing property values inside an animation block, you can configure additional parameters that determine how you want the animation to proceed. You do this by calling the following class methods of `UIView`:

- Use the `setAnimationStartDate:` method to set the start date of the animations after the `commitAnimations` method returns. The default behavior is to schedule the animation for immediate execution on the animation thread.
- Use the `setAnimationDelay:` method to set a delay between the time the `commitAnimations` method returns and the animations actually begin.

- Use the `setAnimationDuration:` method to set the number of seconds over which the animations occur.
- Use the `setAnimationCurve:` method to set the relative speed of the animations over their course. For example, the animations can gradually speed up at the beginning, gradually slow down near the end, or remain the same speed throughout.
- Use the `setAnimationRepeatCount:` method to set the number of times the animations repeat.
- Use the `setAnimationRepeatAutoreverses:` method to specify whether the animations reverse automatically when they reach their target value. Combined with the `setAnimationRepeatCount:` method, you can use this method to toggle each property between its initial and final values smoothly over a period of time.

The `commitAnimations` class method returns immediately and before the animations begin. UIKit performs animations in a separate thread and away from your application's main event loop. The `commitAnimations` method posts its animations to this separate thread where they are queued up until they are ready to execute. By default, Core Animation finishes the currently running animation block before starting animations currently on the queue. You can override this behavior and start your animation immediately, however, by passing YES to the `setAnimationBeginsFromCurrentState:` class method within your animation block. This causes the current in-flight animation to stop and the new animation to begin from the current state.

By default, all animatable property changes within an animation block are animated. If you want to prevent some changes made within the block from being animated, use the `setAnimationsEnabled:` method to disable animations temporarily, make your changes, and then reenable them. Any changes made after a `setAnimationsEnabled:` call with the value NO are not animated until a matching call with the value YES occurs or you commit the animation block. Use the `areAnimationsEnabled` method to determine whether animations are currently enabled.

Configuring an Animation Delegate

You can assign a delegate to an animation block and use that delegate to receive messages when the animations begin and end. You might do this to perform additional tasks immediately before and after the animation. You set the delegate using the `setAnimationDelegate:` class method of `UIView`, and use the `setAnimationWillStartSelector:` and `setAnimationDidStopSelector:` methods to specify the selectors that will receive the messages. The signatures of the corresponding methods are as follows:

```
- (void)animationWillStart:(NSString *)animationID context:(void *)context;
- (void)animationDidStop:(NSString *)animationID finished:(NSNumber *)finished
  context:(void *)context;
```

The *animationID* and *context* parameters for both methods are the same parameters that were passed to the `beginAnimations:context:` method at the beginning of the animation block:

- *animationID* - an application-supplied string used to identify animations in an animation block.
- *context* - another application-supplied object you can use to pass additional information to the delegate.

The `setAnimationDidStopSelector:` selector method has an additional argument—a Boolean value that is YES if the animation ran to completion and was not canceled or stopped prematurely by another animation.

Responding to Layout Changes

Whenever the layout of your views changes, UIKit applies each view's autosizing behaviors and then calls its `layoutSubviews` method to give it a chance to adjust the geometry of its contained subviews further. Layout changes can occur when any of the following happens:

- The size of a view's bounds rectangle changes.
- The content offset value—that is, the origin of the visible content region—of a scroll view changes.
- The transform associated with the view changes.
- The set of Core Animation sublayers associated with the view's layer changes.
- Your application forces layout to occur by calling the `setNeedsLayout` or `layoutIfNeeded` methods of the view.
- Your application forces layout by calling the `setNeedsLayout` method of the view's underlying layer object.

A view's autosizing behaviors handle the initial job of positioning any subviews. Applying these behaviors guarantees that your views are close to their intended size. For information about how autosizing behaviors affect the size and position of your views, see [“Autosizing Behaviors”](#) (page 60).

Sometimes, you might want to adjust the layout of subviews manually using `layoutSubviews`, rather than rely exclusively on autosizing behaviors. For example, if you are implementing a custom control that is built from several subview elements, by adjusting the subviews manually you can precisely configure the appearance for your control over a range of sizes. Alternatively, a view representing a large scrollable content area could display that content by tiling a set of subviews. During scrolling, views going off one edge of the screen would be recycled and repositioned at the incoming screen edge along with any new content.

Note: You can also use the `layoutSubviews` method to adjust the size and position of custom `CALayer` objects attached as sublayers to your view's layer. Managing custom layer hierarchies behind your view lets you perform advanced animations directly using Core Animation. For more information about using Core Animation to manage layer hierarchies, see *Core Animation Programming Guide*.

When writing your layout code, be sure to test your code in each of your application's supported interface orientations. Applications that support both landscape and portrait orientations should verify that layout is handled properly in each orientation. Similarly, your application should be prepared to deal with other system changes, such as the height of the status bar changing. This occurs when a user uses your application while on an active phone call and then hangs up. At hang-up time, the managing view controller may resize its view to account for the shrinking status bar size. Such a change would then filter down to the rest of the views in your application.

Redrawing Your View's Content

Occasionally, changes to your application's data model require that you also change the corresponding user interface. To make those changes, you mark the corresponding views as dirty and in need of an update (using either the `setNeedsDisplay` or `setNeedsDisplayInRect:` methods). Marking views as dirty, as opposed to simply creating a graphics context and drawing, gives the system a chance to process drawing operations more efficiently. For example, if you mark several regions of the same view as dirty during a given cycle, the

system coalesces the dirty regions into a single call to the view's `drawRect:` method. As a result, only one graphics context is created to draw all of the affected regions. This practice is much more efficient than creating several graphics contexts in quick succession.

Views that implement a `drawRect:` method should always check the rectangle passed to the method and use it to limit the scope of their drawing operations. Because drawing is a relatively expensive operation, limiting drawing in this way is a good way to improve performance.

By default, geometry changes to a view do not automatically cause the view to be redrawn. Instead, most geometry changes are handled automatically by Core Animation. Specifically, when you change the `frame`, `bounds`, `center`, or `transform` properties of the view, Core Animation applies the geometry changes to the cached bitmap associated with the view's layer. In many cases, this approach is perfectly acceptable, but if you find the results undesirable, you can force UIKit to redraw your view instead. To prevent Core Animation from applying geometry changes implicitly, set your view's `contentMode` property to `UIViewContentModeRedraw`. For more information about content modes, see [“Content Modes and Scaling”](#) (page 58).

Hiding Views

You can hide or show a view by changing the value in the view's `hidden` property. Setting this property to `YES` hides the view; setting it to `NO` shows it. Hiding a view also hides any embedded subviews as if their own `hidden` property were set.

When you hide a view, it remains in the view hierarchy, but its contents are not drawn and it does not receive touch events. Because it remains in the view hierarchy, a hidden view continues to participate in autoresizing and other layout operations. If you hide a view that is currently the first responder, the view does not automatically resign its first responder status. Events targeted at the first responder are still delivered to the hidden view. For more information about the responder chain, see [“Responder Objects and the Responder Chain”](#) (page 77).

Creating a Custom View

The `UIView` class provides the underlying support for displaying content on the screen and for handling touch events, but its instances draw nothing but a background color using an alpha value and its subviews. If your application needs to display custom content or handle touch events in a specific manner, you must create a custom subclass of `UIView`.

The following sections describe some of the key methods and behaviors you might implement in your custom view objects. For additional subclassing information, see *UIView Class Reference*.

Initializing Your Custom View

Every new view object you define should include a custom `initWithFrame:` method. This method is responsible for initializing the class at creation time and putting your view object into a known state. You use this method when creating instances of your view programmatically in your code.

Listing 2-2 shows a skeletal implementation of a standard `initWithFrame:` method. This method calls the inherited implementation of the method first and then initializes the instance variables and state information of the class before returning the initialized object. Calling the inherited implementation is traditionally performed first so that if there is a problem, you can simply abort your own initialization code and return `nil`.

Listing 2-2 Initializing a view subclass

```
- (id)initWithFrame:(CGRect)aRect {
    self = [super initWithFrame:aRect];
    if (self) {
        // setup the initial properties of the view
        ...
    }
    return self;
}
```

If you plan to load instances of your custom view class from a nib file, you should be aware that in iPhone OS, the nib-loading code does not use the `initWithFrame:` method to instantiate new view objects. Instead, it uses the `initWithCoder:` method that is defined as part of the `NSCoding` protocol.

Even if your view adopts the `NSCoding` protocol, Interface Builder does not know about your view's custom properties and therefore does not encode those properties into the nib file. As a result, your own `initWithCoder:` method does not have the information it needs to properly initialize the class when it is loaded from a nib file. To solve this problem, you can implement the `awakeFromNib` method in your class and use it to initialize your class specifically when it is loaded from a nib file.

Drawing Your View's Content

As you make changes to your view's content, you notify the system that parts of that view need to be redrawn using the `setNeedsDisplay` or `setNeedsDisplayInRect:` methods. When the application returns to its run loop, it coalesces any drawing requests and computes the specific parts of your interface that need to be updated. It then begins traversing your view hierarchy and sending `drawRect:` messages to the views that require updates. The traversal starts with the root view of your hierarchy and proceeds down through the subviews, processing them from back to front. Views that display custom content inside their visible bounds must implement the `drawRect:` method to render that content.

Before calling your view's `drawRect:` method, UIKit configures the drawing environment for your view. It creates a graphics context and adjusts its coordinate system and clipping region to match the coordinate system and bounds of your view. Thus, by the time your `drawRect:` method is called, you can simply begin drawing using UIKit classes and functions, Quartz functions, or a combination of them all. If you need to access the current graphics context, you can get a pointer to it using the `UIGraphicsGetCurrentContext` function.

Important: The current graphics context is valid only for the duration of one call to your view's `drawRect:` method. UIKit may create a different graphics context for each subsequent call to this method, so you should not try to cache the object and use it later.

Listing 2-3 shows a simple implementation of a `drawRect:` method that draws a 10-pixel-wide red border around the view. Because UIKit drawing operations use Quartz for their underlying implementations, you can mix drawing calls as shown here and still get the results you expect.

Listing 2-3 A drawing method

```

- (void)drawRect:(CGRect)rect {
    CGContextRef context = UIGraphicsGetCurrentContext();
    CGRect myFrame = self.bounds;

    CGContextSetLineWidth(context, 10);

    [[UIColor redColor] set];
    UIRectFrame(myFrame);
}

```

If you know that your view's drawing code always covers the entire surface of the view with opaque content, you can improve the overall efficiency of your drawing code by setting the `opaque` property of your view to `YES`. When you mark a view as opaque, UIKit avoids drawing content that is located immediately behind your view. This not only reduces the amount of time spent drawing but also minimizes the work that must be done to composite that content together. You should set this property to `YES` only if you know your view provides opaque content. If your view cannot guarantee that its contents are always opaque, you should set the property to `NO`.

Another way to improve drawing performance, especially during scrolling, is to set the `clearsContextBeforeDrawing` property of your view to `NO`. When this property is set to `YES`, UIKit automatically fills the area to be updated by your `drawRect:` method with transparent black before calling your method. Setting this property to `NO` eliminates the overhead for that fill operation but puts the burden on your application to completely redraw the portions of your view inside the update rectangle passed to your `drawRect:` method. Such an optimization is usually a good tradeoff during scrolling, however.

Responding to Events

The `UIView` class is a subclass of `UIResponder` and is therefore capable of receiving touch events corresponding to user interactions with the view's contents. Touch events start at the view in which the touch occurred and are passed up the responder chain until they are handled. Because views are themselves responders, they participate in the responder chain and therefore can receive touch events dispatched to them from any of their associated subviews.

Views that handle touch events typically implement all of the following methods, which are described in more detail in [“Event Handling”](#) (page 75).

```

touchesBegan:withEvent:
touchesMoved:withEvent:
touchesEnded:withEvent:
touchesCancelled:withEvent:

```

Remember that, by default, views respond to only one touch at a time. If the user puts a second finger down, the system ignores the touch event and does not report it to your view. If you plan to track multifinger gestures from your view's event-handler methods, you need to reenableView multi-touch events by setting the `multipleTouchEnabled` property of your view to `YES`.

Some views, such as labels and images, disable event handling altogether initially. You can control whether a view handles events at all by changing the value of the view's `userInteractionEnabled` property. You might temporarily set this property to `NO` to prevent the user from manipulating the contents of your view while a long operation is pending. To prevent events from reaching any of your views, you can also use the

`beginIgnoringInteractionEvents` and `endIgnoringInteractionEvents` methods of the `UIApplication` object. These methods affect the delivery of events for the entire application, not just for a single view.

As it handles touch events, UIKit uses the `hitTest:withEvent:` and `pointInside:withEvent:` methods of `UIView` to determine whether a touch event occurred in a given view. Although you rarely need to override these methods, you could do so to implement custom touch behaviors for your view. For example, you could override these methods to prevent subviews from handling touch events.

Cleaning Up After Your View

If your view class allocates any memory, stores references to any custom objects, or holds resources that must be released when the view is released, you must implement a `dealloc` method. The system calls the `dealloc` method when your view's retain count reaches zero and your view is about to be deallocated itself. Your implementation of this method should release the objects and resources it holds and then call the inherited implementation, as shown in Listing 2-4.

Listing 2-4 Implementing the `dealloc` method

```
- (void)dealloc {  
    // Release a retained UIColor object  
    [color release];  
  
    // Call the inherited implementation  
    [super dealloc];  
}
```


Event Handling

Events in iPhone OS are based on a Multi-Touch model. Instead of using a mouse and a keyboard, users touch the screen of the device to manipulate objects, enter data, and otherwise convey their intentions. iPhone OS recognizes one or more fingers touching the screen as part of a **Multi-Touch sequence**. This sequence begins when the first finger touches down on the screen and ends when the last finger is lifted from the screen. iPhone OS tracks fingers touching the screen throughout a multi-touch sequence and records the characteristics of each of them, including the location of the finger on the screen and the time the touch occurred. Applications often recognize certain combinations of touches as gestures and respond to them in ways that are intuitive to users, such as zooming in on content in response to a pinching gesture and scrolling through content in response to a flicking gesture.

Note: A finger on the screen affords a much different level of precision than a mouse pointer. When a user touches the screen, the area of contact is actually elliptical and tends to be offset below the point where the user thinks he or she touched. This “contact patch” also varies in size and shape based on which finger is touching the screen, the size of the finger, the pressure of the finger on the screen, the orientation of the finger, and other factors. The underlying Multi-Touch system analyzes all of this information for you and computes a single touch point.

Many classes in UIKit handle multi-touch events in ways that are distinctive to objects of the class. This is especially true of subclasses of `UIControl`, such as `UIButton` and `UISlider`. Objects of these subclasses—known as control objects—are receptive to certain types of gestures, such as a tap or a drag in a certain direction; when properly configured, they send an action message to a target object when that gesture occurs. Other UIKit classes handle gestures in other contexts; for example, `UIScrollView` provides scrolling behavior for table views, text views, and other views with large content areas.

Some applications may not need to handle events directly; instead, they can rely on the classes of UIKit for that behavior. However, if you create a custom subclass of `UIView`—a common pattern in iPhone OS development—and if you want that view to respond to certain touch events, you need to implement the code required to handle those events. Moreover, if you want a UIKit object to respond to events differently, you have to create a subclass of that framework class and override the appropriate event-handling methods.

Events and Touches


In iPhone OS, a **touch** is the presence or movement of a finger on the screen that is part of a unique **multi-touch sequence**. For example, a pinch-close gesture has two touches: two fingers on the screen moving toward each other from opposite directions. There are simple single-finger gestures, such as a tap, or a double-tap, or a flick (where the user quickly swipes a finger across the screen). An application might recognize even more complicated gestures; for example, an application might have a custom control in the shape of a dial that users “turn” with multiple fingers to fine-tune some variable.

An **event** is an object that the system continually sends to an application as fingers touch the screen and move across its surface. The event provides a snapshot of all touches during a multi-touch sequence, most importantly the touches that are new or have changed for a particular view. A multi-touch sequence begins

when a finger first touches the screen. Other fingers may subsequently touch the screen, and all fingers may move across the screen. The sequence ends when the last of these fingers is lifted from the screen. An application receives event objects during each phase of any touch.

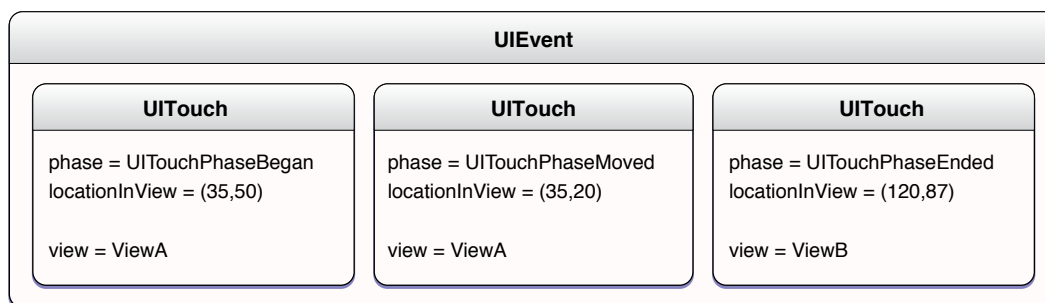
Touches have both temporal and spatial aspects. The temporal aspect, called a phase, indicates when a touch has just begun, whether it is moving or stationary, and when it ends—that is, when the finger is lifted from the screen (see Figure 3-1). A touch also has the current location in a view or window and the previous location (if any). When a finger touches the screen, the touch is associated with a window and a view and maintains that association throughout the life of the event. If multiple touches arrive at once, they are treated together only if they are associated with the same view. Likewise, if two touches arrive in quick succession, they are treated as a multiple tap only if they are associated with the same view.

Figure 3-1 A multi-touch sequence and touch phases



In iPhone OS, a `UITouch` object represents a touch, and a `UIEvent` object represents an event. An event object contains all touch objects for the current multi-touch sequence and can provide touch objects specific to a view or window (see Figure 3-2). A touch object is persistent for a given finger during a sequence, and UIKit mutates it as it tracks the finger throughout it. The touch attributes that change are the phase of the touch, its location in a view, its previous location, and its timestamp. Event-handling code evaluates these attributes to determine how to respond to the event.

Figure 3-2 Relationship of a `UIEvent` object and its `UITouch` objects



The system can cancel a multi-touch sequence at any time and an event-handling application must be prepared to respond appropriately. Cancellations can occur as a result of overriding system events, such as an incoming phone call.

Event Delivery

The delivery of an event to an object for handling occurs along a specific path. As described in “[Core Application Architecture](#)” (page 15), when users touch the screen of a device, iPhone OS recognizes the set of touches and packages them in a `UIEvent` object that it places in the current application’s event queue. The event object encapsulates the touches for a given moment of a multi-touch sequence. The singleton `UIApplication` object that is managing the application takes an event from the top of the queue and dispatches it for handling. Typically, it sends the event to the application’s key window—the window currently the focus for user events—and the `UIWindow` object representing that window sends the event to the first responder for handling. (The first responder is described in “Responder Objects and the Responder Chain.”)

An application uses hit-testing to find the first responder for an event; it recursively calls `hitTest:withEvent:` on the views in the view hierarchy (going down the hierarchy) to determine the subview in which the touch took place. The touch is associated with that view for its lifetime, even if it subsequently moves outside the view. “[Event-Handling Techniques](#)” (page 83) discusses some of the programmatic implications of hit-testing.

The `UIApplication` object and each `UIWindow` object dispatches events in the `sendEvent:` method. (Both classes declare an identically named method). Because these methods are funnel points for events coming into an application, you can subclass `UIApplication` or `UIWindow` and override the `sendEvent:` method to monitor events or perform special event handling. However, most applications have no need to do this.

Responder Objects and the Responder Chain

A **responder object** is an object that can respond to events and handle them. `UIResponder` is the base class for all responder objects. It defines the programmatic interface not only for event handling but for common responder behavior. `UIApplication`, `UIView`, and all `UIKit` classes that descend from `UIView` (including `UIWindow`) inherit directly or indirectly from `UIResponder`.

The **first responder** is the responder object in the application (usually a `UIView` object) that is the current recipient of touches. A `UIWindow` object sends the first responder an event in a message, giving it the first shot at handling the event. If the first responder doesn’t handle the event, it passes the event (via message) to the next responder in the responder chain to see if it can handle it.

The **responder chain** is a linked series of responder objects. It allows responder objects to delegate responsibility for handling an event to other, higher-level objects. An event proceeds up the responder chain as the application looks for an object capable of handling the event. The responder chain consists of a series of “next responders” in the following sequence:

1. The first responder passes the event to its view controller (if it has one) and then on to its superview.
2. Each subsequent view in the hierarchy similarly passes to its view controller first (if it has one) and then to its superview.
3. The topmost enclosing view passes the event to the `UIWindow` object.
4. The `UIWindow` object passes the event to the singleton `UIApplication` object.

If the application finds no responder object to handle the event, it discards the event.

Any responder object in the responder chain may implement a `UIResponder` event-handling method and thus receive an event message. But a responder may decline to handle a particular event or may handle it only partially. In that case, it can forward the event message to the next responder in a message similar to the following one:

```
- (void)touchesBegan:(NSSet *)touches withEvent:(UIEvent *)event {
    UITouch* touch = [touches anyObject];
    NSUInteger numTaps = [touch tapCount];
    if (numTaps < 2) {
        [self.nextResponder touchesBegan:touches withEvent:event];
    } else {
        [self handleDoubleTap:touch];
    }
}
```

Note: If a responder object forwards event-handling messages to the next responder for the initial phase of multi-touch sequence (in `touchesBegan:withEvent:`), it should forward all other event-handling messages for that sequence.

Action messages also make use of the responder chain. When users manipulate a `UIControl` object such as button or page control, the control object (if properly configured) sends an action message to a target object. But if `nil` is specified as the target, the application initially routes the message as it does an event message: to the first responder. If the first responder doesn't handle the action message, it sends it to its next responder, and so on up the responder chain.

Regulating Event Delivery

UIKit gives applications programmatic means to simplify event handling or to turn off the stream of events completely. The following list summarizes these approaches:

- **Turning off delivery of events.** By default, a view receives touch events, but you can set its `userInteractionEnabled` property to `NO` to turn off delivery of events. A view also does not receive events if it's hidden or if it's transparent.
- **Turning off delivery of events for a period.** An application can call the `UIApplication` method `beginIgnoringInteractionEvents` and later call the `endIgnoringInteractionEvents` method. The first method stops the application from receiving touch event messages entirely; the second method is called to resume the receipt of such messages. You sometimes want to turn off event delivery when your code is performing animations.
- **Turning on delivery of multiple touches.** By default, a view ignores all but the first touch during a multi-touch sequence. If you want the view to handle multiple touches you must enable multiple touches for the view. This can be done programmatically by setting the `multipleTouchEnabled` property of your view to `YES`, or in Interface Builder using the inspector for the related view.
- **Restricting event delivery to a single view.** By default, a view's `exclusiveTouch` property is set to `NO`. If you set the property to `YES`, you mark the view so that, if it is tracking touches, it is the only view in the window that is tracking touches. Other views in the window cannot receive those touches. However, a view that is marked "exclusive touch" does not receive touches that are associated with other views in the same window. If a finger contacts an exclusive-touch view, then that touch is delivered only if that view is the only view tracking a finger in that window. If a finger touches a non-exclusive view, then that touch is delivered only if there is not another finger tracking in an exclusive-touch view.

- **Restricting event delivery to subviews.** A custom `UIView` class can override `hitTest:withEvent:` to restrict the delivery of multi-touch events to its subviews. See “[Event-Handling Techniques](#)” (page 83) for a discussion of this technique.

Handling Multi-Touch Events

To handle multi-touch events, your custom `UIView` subclass (or, less frequently, your custom `UIApplication` or `UIWindow` subclass), must implement at least one of the `UIResponder` methods for event handling. The following sections describe these methods, discuss approaches for handling common gestures, show an example of a responder object that handles a complex sequence of multi-touch events, and suggest some techniques for event handling.

The Event-Handling Methods

During a multi-touch sequence, the application dispatches a series of event messages. To receive and handle these messages, the class of a responder object must implement at least one of the following methods declared by `UIResponder`:

- `(void)touchesBegan:(NSSet *)touches withEvent:(UIEvent *)event;`
- `(void)touchesMoved:(NSSet *)touches withEvent:(UIEvent *)event;`
- `(void)touchesEnded:(NSSet *)touches withEvent:(UIEvent *)event;`
- `(void)touchesCancelled:(NSSet *)touches withEvent:(UIEvent *)event;`

The application sends these messages when there are new or changed touches for a given touch phase:

- It sends the `touchesBegan:withEvent:` message when one or more fingers touch down on the screen.
- It sends the `touchesMoved:withEvent:` message when one or more fingers move.
- It sends the `touchesEnded:withEvent:` message when one or more fingers lift up from the screen.
- It sends the `touchesCancelled:withEvent:` message when the touch sequence is cancelled by a system event, such as an incoming phone call.

Each of these methods is associated with a touch phase (for example, `UITouchPhaseBegan`), which for any `UITouch` object you can find out by evaluating its `phase` property.

Each message that invokes an event-handling method passes in two parameters. The first is a set of `UITouch` objects that represent new or changed touches for the given phase. The second parameter is a `UIEvent` object representing this particular event. From the event object you can get all touch objects for the event (`allTouches`) or a subset of those touch objects filtered for specific views or windows. Some of these touch objects represent touches that have not changed since the previous event message or that have changed but are in different phases.

A responder object frequently handles an event for a given phase by getting one or more of the `UITouch` objects in the passed-in set and then evaluating their properties or getting their locations. (If any of the touch objects will do, it can send the `NSSet` object an `anyObject` message.) One important method is `locationInView:`, which, if passed a parameter of `self`, yields the location of the touch in the responder object’s coordinate system (assuming the responder is a `UIView` object and the view passed as a parameter

is not nil). A parallel method tells you the previous location of the touch (`previousLocationInView:`). Properties of the `UITouch` instance tell you how many taps have been made (`tapCount`), when the touch was created or last mutated (`timestamp`), and what phase it is in (`phase`).

A responder class does not have to implement all three of the event methods listed above. For example, if it is looking for only fingers when they're lifted from the screen, it need only implement `touchesEnded:withEvent:`.

If a responder creates persistent objects while handling events during a multi-touch sequence, it should implement `touchesCancelled:withEvent:` to dispose of those objects when the system cancels the sequence. Cancellation often occurs when an external event—for example, an incoming phone call—disrupts the current application's event processing. Note that a responder object should also dispose of those same objects when it receives the last `touchesEnded:withEvent:` message for a multi-touch sequence. (See [“Event-Handling Techniques”](#) (page 83) to find out how to determine the last touch-up in a sequence.)

Handling Single and Multiple Tap Gestures

A very common gesture in iPhone applications is the tap: the user taps an object with his or her finger. A responder object can respond to a single tap in one way, a double-tap in another, and possibly a triple-tap in yet another way. To determine the number of times the user tapped a responder object, you get the value of the `tapCount` property of a `UITouch` object.

The best places to find this value are the methods `touchesBegan:withEvent:` and `touchesEnded:withEvent:`. In many cases, the latter method is preferred because it corresponds to the touch phase in which the user lifts a finger from a tap. By looking for the tap count in the touch-up phase (`UITouchPhaseEnded`), you ensure that the finger is really tapping and not, for instance, touching down and then dragging.

Listing 3-1 shows the way to determine whether a double-tap occurred in one of your views.

Listing 3-1 Detecting a double-tap gesture

```
- (void) touchesEnded:(NSSet*)touches withEvent:(UIEvent*)event
{
    UITouch      *touch = [touches anyObject];

    if([touch tapCount] == 2) {
        // Process a double-tap gesture
    }
}
```

A complication arises when a responder object wants to handle a single-tap *and* a double-tap gesture in different ways. For example, a single tap might select the object and a double tap might display a view for editing the item that was double-tapped. How is the responder object to know that a single tap is not the first part of a double tap? Here is how a responder object could handle this situation using the event-handling methods just described:

1. In `touchesEnded:withEvent:`, when the tap count is one, the responder object sends itself a `performSelector:withObject:afterDelay:` message. The selector identifies another method implemented by the responder to handle the single-tap gesture; the object for the second parameter is the related `UITouch` object; the delay is some reasonable interval between a single- and a double-tap gesture.

2. In `touchesBegan:withEvent:`, if the tap count is two, the responder object cancels the pending `delayed-perform` invocation by sending itself a `cancelPreviousPerformRequestsWithTarget:` message. If the tap count is not two, the method identified by the selector in the previous step for single-tap gestures is invoked after the delay.
3. In `touchesEnded:withEvent:`, if the tap count is two, the responder performs the actions necessary for handling double-tap gestures.

Detecting Swipe Gestures

Horizontal and vertical swipes are a simple type of gesture that you can track easily from your own code and use to perform actions. To detect a swipe gesture, you have to track the movement of the user's finger along the desired axis of motion, but it is up to you to determine what constitutes a swipe. In other words, you need to determine whether the user's finger moved far enough, if it moved in a straight enough line, and if it went fast enough. You do that by storing the initial touch location and comparing it to the location reported by subsequent touch-moved events.

Listing 3-2 shows some basic tracking methods you could use to detect horizontal swipes in a view. In this example, the view stores the initial location of the touch in a `startTouchPosition` member variable. As the user's finger moves, the code compares the current touch location to the starting location to determine whether it is a swipe. If the touch moves too far vertically, it is not considered to be a swipe and is processed differently. If it continues along its horizontal trajectory, however, the code continues processing the event as if it were a swipe. The processing routines could then trigger an action once the swipe had progressed far enough horizontally to be considered a complete gesture. To detect swipe gestures in the vertical direction, you would use similar code but would swap the x and y components.

Listing 3-2 Tracking a swipe gesture in a view

```
#define HORIZ_SWIPE_DRAG_MIN 12
#define VERT_SWIPE_DRAG_MAX 4

- (void)touchesBegan:(NSSet *)touches withEvent:(UIEvent *)event
{
    UITouch *touch = [touches anyObject];
    startTouchPosition = [touch locationInView:self];
}

- (void)touchesMoved:(NSSet *)touches withEvent:(UIEvent *)event
{
    UITouch *touch = [touches anyObject];
    CGPoint currentTouchPosition = [touch locationInView:self];

    // If the swipe tracks correctly.
    if (fabsf(startTouchPosition.x - currentTouchPosition.x) >=
        HORIZ_SWIPE_DRAG_MIN &&
        fabsf(startTouchPosition.y - currentTouchPosition.y) <=
        VERT_SWIPE_DRAG_MAX)
    {
        // It appears to be a swipe.
        if (startTouchPosition.x < currentTouchPosition.x)
            [self myProcessRightSwipe:touches withEvent:event];
        else
            [self myProcessLeftSwipe:touches withEvent:event];
    }
}
```

```

    }
    else
    {
        // Process a non-swipe event.
    }
}

```

Handling a Complex Multi-Touch Sequence

Taps and swipes are simple gestures. Handling a multi-touch sequence that is more complicated—in effect, interpreting an application-specific gesture—depends on what the application is trying to accomplish. You may have to track all touches through all phases, recording the touch attributes that have changed and altering internal state appropriately.

The best way to convey how you might handle a complex multi-touch sequence is through an example. Listing 3-3 shows how a custom `UIView` object responds to touches by animating the movement of a “Welcome” placard around the screen as a finger moves it and changing the language of the welcome when the user makes a double-tap gesture. (The code in this example comes from the *MoveMe* sample code project, which you can examine to get a better understanding of the event-handling context.)

Listing 3-3 Handling a complex multi-touch sequence

```

- (void)touchesBegan:(NSSet *)touches withEvent:(UIEvent *)event
{
    UITouch *touch = [[event allTouches] anyObject];
    // Only move the placard view if the touch was in the placard view
    if ([touch view] != placardView) {
        // On double tap outside placard view, update placard's display string
        if ([touch tapCount] == 2) {
            [placardView setupNextDisplayString];
        }
        return;
    }
    // "Pulse" the placard view by scaling up then down
    // Use UIView's built-in animation
    [UIView beginAnimations:nil context:NULL];
    [UIView setAnimationDuration:0.5];
    CGAffineTransform transform = CGAffineTransformMakeScale(1.2, 1.2);
    placardView.transform = transform;
    [UIView commitAnimations];

    [UIView beginAnimations:nil context:NULL];
    [UIView setAnimationDuration:0.5];
    transform = CGAffineTransformMakeScale(1.1, 1.1);
    placardView.transform = transform;
    [UIView commitAnimations];

    // Move the placardView to under the touch
    [UIView beginAnimations:nil context:NULL];
    [UIView setAnimationDuration:0.25];
    placardView.center = [self convertPoint:[touch locationInView:self]
fromView:placardView];
    [UIView commitAnimations];
}

```

```

- (void)touchesMoved:(NSSet *)touches withEvent:(UIEvent *)event
{
    UITouch *touch = [[event allTouches] anyObject];

    // If the touch was in the placardView, move the placardView to its location
    if ([touch view] == placardView) {
        CGPoint location = [touch locationInView:self];
        location = [self convertPoint:location fromView:placardView];
        placardView.center = location;
        return;
    }
}

- (void)touchesEnded:(NSSet *)touches withEvent:(UIEvent *)event
{
    UITouch *touch = [[event allTouches] anyObject];

    // If the touch was in the placardView, bounce it back to the center
    if ([touch view] == placardView) {
        // Disable user interaction so subsequent touches don't interfere with
        animation
        self.userInteractionEnabled = NO;
        [self animatePlacardViewToCenter];
        return;
    }
}

```

Note: Custom views that redraw themselves in response to events they handle generally should only set drawing state in the event-handling methods and perform all of the drawing in the `drawRect:` method. To learn more about drawing view content, see [“Graphics and Drawing”](#) (page 85).

Event-Handling Techniques

Here are some event-handling techniques you can use in your code.

■ Tracking the mutations of `UITouch` objects

In your event-handling code you can store relevant bits of touch state for later comparison with the mutated `UITouch` instance. As an example, say you want to compare the final location of each touch with its original location. In the `touchesBegan:withEvent:` method, you can obtain the original location of each touch from the `locationInView:` method and store those in a `CFDictionary` object using the addresses of the `UITouch` objects as keys. Then, in the `touchesEnded:withEvent:` method you can use the address of each passed-in `UITouch` object to obtain the object's original location and compare that with its current location. (You should use a `CFDictionary` object rather than an `NSDictionary` object; the latter copies its keys, but the `UITouch` class does not adopt the `NSCopying` protocol, which is required for object copying.)

■ Hit-testing for a touch on a subview or layer

A custom view can use the `hitTest:withEvent:` method of `UIView` or the `hitTest:` method of `CALayer` to find the subview or layer that is receiving a touch, and handle the event appropriately. The following example detects when an “Info” image in a layer of the custom view is tapped.

```

- (void)touchesEnded:(NSSet*)touches withEvent:(UIEvent*)event {
    CGPoint location = [[touches anyObject] locationInView:self];

```

```

    CALayer *hitLayer = [[self layer] hitTest:[self convertPoint:location
fromView:nil]];

    if (hitLayer == infoImage) {
        [self displayInfo];
    }
}

```

If you have a custom view with subviews, you need to determine whether you want to handle touches at the subview level or the superview level. If the subviews do not handle touches by implementing `touchesBegan:withEvent:`, `touchesEnded:withEvent:`, or `touchesMoved:withEvent:`, then these messages propagate up the responder chain to the superview. However, because multiple taps and multiple touches are associated with the subviews where they first occurred, the superview won't receive these touches. To ensure reception of all kinds of touches, the superview should override `hitTest:withEvent:` to return itself rather than any of its subviews.

■ Determining when the last finger in a multi-touch sequence has lifted

When you want to know when the last finger in a multi-touch sequence is lifted from a view, compare the number of `UITouch` objects in the passed-in set with the number of touches for the view maintained by the passed-in `UIEvent` object. For example:

```

- (void)touchesEnded:(NSSet*)touches withEvent:(UIEvent*)event {
    if ([touches count] == [[event touchesForView:self] count]) {
        // last finger has lifted....
    }
}

```

Graphics and Drawing

High-quality graphics are an important part of your application's user interface. Providing high-quality graphics not only makes your application look good, but it also makes your application look like a natural extension to the rest of the system. iPhone OS provides two primary paths for creating high-quality graphics in your system: OpenGL or native rendering using Quartz, Core Animation, and UIKit.

The OpenGL frameworks are geared primarily toward game development or applications that require high frame rates. OpenGL is a C-based interface used to create 2D and 3D content on desktop computers. iPhone OS supports OpenGL drawing through the OpenGL ES framework, which is based on the OpenGL ES v1.1 specification and is designed specifically for use on embedded hardware systems. This version differs in many ways from desktop versions of OpenGL, so you should be sure to follow the advice for using it later in this chapter.

For developers who want a more object-oriented drawing approach, iPhone OS provides Quartz, Core Animation, and the graphics support in UIKit. Quartz is the main drawing interface, providing support for path-based drawing, anti-aliased rendering, gradient fill patterns, images, colors, coordinate-space transformations, and PDF document creation, display, and parsing. UIKit provides Objective-C wrappers for Quartz images and color manipulations. Core Animation provides the underlying support for animating changes in many UIKit view properties and can also be used to implement custom animations.

This chapter provides an overview of the drawing process for iPhone applications, along with specific drawing techniques for each of the supported drawing technologies. This chapter also provides tips and guidance on how to optimize your drawing code for the iPhone OS platform.

The UIKit Graphics System

In iPhone OS, all drawing—regardless of whether it involves OpenGL, Quartz, UIKit, or Core Animation—occurs within the confines of a `UIView` object. Views define the portion of the screen in which drawing occurs. If you use system-provided views, this drawing is handled for you automatically. If you define custom views, however, you must provide the drawing code yourself. For applications that draw using OpenGL, once you set up your rendering surface, you use the drawing model specified by OpenGL.

For Quartz, Core Animation, and UIKit, you use the drawing concepts described in the following sections.

The View Drawing Cycle

The basic drawing model for `UIView` objects involves updating content on demand. The `UIView` class makes the update process easier and more efficient, however, by gathering the update requests you make and delivering them to your drawing code at the most appropriate time.

Whenever a portion of your view needs to be redrawn, the `UIView` object's built-in drawing code calls its `drawRect:` method. It passes this method a rectangle indicating the portion of your view that needs to be redrawn. You override this method in your custom view subclasses and use it to draw the contents of your

view. The first time your view is drawn, the rectangle passed to the `drawRect:` method contains your view's entire visible area. During subsequent calls, however, this rectangle represents only the portion of the view that actually needs to be redrawn. There are several actions that can trigger a view update:

- Moving or removing another view that was partially obscuring your view
- Making a previously hidden view visible again by setting its `hidden` property to `NO`
- Scrolling a view off the screen and then back on
- Explicitly calling the `setNeedsDisplay` or `setNeedsDisplayInRect:` method of your view

After calling your `drawRect:` method, the view marks itself as updated and waits for new actions to arrive and trigger another update cycle. If your view displays static content, then all you need to do is respond to changes in your view's visibility caused by scrolling and the presence of other views. If you update your view's content periodically, however, you must determine when to call the `setNeedsDisplay` or `setNeedsDisplayInRect:` method to trigger an update. For example, if you were updating content several times a second, you might want to set up a timer to update your view. You might also update your view in response to user interactions or the creation of new content in your view.

Coordinates and Coordinate Transforms

As described in “[View Coordinate Systems](#)” (page 55), the origin of a window or view is located in its top-left corner, and positive coordinate values extend down and to the right of this origin. When you write your drawing code, you use this coordinate system to specify the location of individual points for the content you draw.

If you need to make changes to the default coordinate system, you do so by modifying the current transformation matrix. The **current transformation matrix (CTM)** is a mathematical matrix that maps points in your view's coordinate system to points on the device's screen. When your view's `drawRect:` method is first called, the CTM is configured so that the origin of the coordinate system matches the your view's origin and its positive axes extend down and to the right. However, you can change the CTM by adding scaling, rotation, and translation factors to it and thereby change the size, orientation, and position of the default coordinate system relative to the underlying view or window.

Modifying the CTM is the standard technique used to draw content in your view because it involves much less work. If you want to draw a 10 x 10 square starting at the point (20, 20) in the current drawing system, you could create a path that moves to (20, 20) and then draws the needed set of lines to complete the square. If you decide later that you want to move that square to the point (10, 10), however, you would have to recreate the path with the new starting point. In fact, you would have to recreate the path every time you changed the origin. Creating paths is a relatively expensive operation, but creating a square whose origin is at (0, 0) and modifying the CTM to match the desired drawing origin is cheap by comparison.

In the Core Graphics framework, there are two ways to modify the CTM. You can modify the CTM directly using the CTM manipulation functions defined in *CGContext Reference*. You can also create a `CGAffineTransform` structure, apply any transformations you want, and then concatenate that transform onto the CTM. Using an affine transform lets you group transformations and then apply them to the CTM all at once. You can also evaluate and invert affine transforms and use them to modify point, size, and rectangle values in your code. For more information on using affine transforms, see *Quartz 2D Programming Guide* and *CGAffineTransform Reference*.

Graphics Contexts

Before calling your custom `drawRect:` method, the view object automatically configures its drawing environment so that your code can start drawing immediately. As part of this configuration, the `UIView` object creates a graphics context (a `CGContextRef` opaque type) for the current drawing environment. This graphics context contains the information the drawing system needs to perform any subsequent drawing commands. It defines basic drawing attributes such as the colors to use when drawing, the clipping area, line width and style information, font information, compositing options, and several others.

You can create custom graphics context objects in situations where you want to draw somewhere other than your view. In Quartz, you primarily do this when you want to capture a series of drawing commands and use them to create an image or a PDF file. To create the context, you use the `CGBitmapContextCreate` or `CGPDFContextCreate` function. Once you have the context, you can pass it to the drawing functions needed to create your content.

When creating custom contexts, the coordinate system for those contexts is different than the native coordinate system used by iPhone OS. Instead of the origin being in the upper-left corner of the drawing surface, it is in the lower-left corner and the axes point up and to the right. The coordinates you specify in your drawing commands must take this into consideration or the resulting image or PDF file may appear wrong when rendered.

Important: Because you use a lower-left origin when drawing into a bitmap or PDF context, you must compensate for that coordinate system when rendering the resulting content into a view. In other words, if you create an image and draw it using the `CGContextDrawImage` function, the image will appear upside down by default. To correct for this, you must invert the y axis of the CTM (by multiplying it by -1) and shift the origin from the lower-left corner to the upper-left corner of the view.

If you use a `UIImage` object to wrap a `CGImageRef` you create, you do not need to modify the CTM. The `UIImage` object automatically compensates for the inverted coordinate system of the `CGImageRef` type.

For more information about graphics contexts, modifying the graphics state information, and using graphics contexts to create custom content, see *Quartz 2D Programming Guide*. For a list of functions used in conjunction with graphics contexts, see *CGContext Reference*, *CGBitmapContext Reference*, and *CGPDFContext Reference*.

Points Versus Pixels

The Quartz drawing system uses a vector-based drawing model. Compared to a raster-based drawing model, in which drawing commands operate on individual pixels, drawing commands in Quartz are specified using a fixed-scale drawing space, known as the **user coordinate space**. iPhone OS then maps the coordinates in this drawing space onto the actual pixels of the device. The advantage of this model is that graphics drawn using vector commands continue to look good when scaled up or down using an affine transform.

In order to maintain the precision inherent in a vector-based drawing system, drawing coordinates are specified using floating-point values instead of integers. The use of floating-point values for coordinates makes it possible for you to specify the location of your program's content very precisely. For the most part, you do not have to worry about how those values are eventually mapped to the device's screen.

The user coordinate space is the environment that you use for all of your drawing commands. The units of this space are measured in points. The **device coordinate space** refers to the native coordinate space of the device, which is measured in pixels. By default, one point in user coordinate space is equal to one pixel in device space, which results in 1 point equaling 1/160th of an inch. You should not assume that this 1-to-1 ratio will always be the case, however.

Color and Color Spaces

iPhone OS supports the full range of color spaces available in Quartz; however, most applications should need only the RGB color space. Because iPhone OS is designed to run on embedded hardware and display graphics on a screen, the RGB color space is the most appropriate one to use.

The `UIColor` object provides convenience methods for specifying color values using RGB, HSB, and grayscale values. When creating colors in this way, you never need to specify the color space. It is determined for you automatically by the `UIColor` object.

You can also use the `CGContextSetRGBStrokeColor` and `CGContextSetRGBFillColor` functions in the Core Graphics framework to create and set colors. Although the Core Graphics framework includes support for creating colors using other color spaces, and for creating custom color spaces, using those colors in your drawing code is not recommended. Your drawing code should always use RGB colors.

Supported Image Formats

Table 4-1 lists the image formats supported directly by iPhone OS. Of these formats, the PNG format is the one most recommended for use in your applications.

Table 4-1 Supported image formats

Format	Filename extensions
Portable Network Graphic (PNG)	.png
Tagged Image File Format (TIFF)	.tiff, .tif
Joint Photographic Experts Group (JPEG)	.jpeg, .jpg
Graphic Interchange Format (GIF)	.gif
Windows Bitmap Format (DIB)	.bmp, .BMPf
Windows Icon Format	.ico
Windows Cursor	.cur
XWindow bitmap	.xbm

Drawing Tips

The following sections provide tips on how to write quality drawing code while ensuring that your application looks appealing to end users.

Deciding When to Use Custom Drawing Code

Depending on the type of application you are creating, it may be possible to use little or no custom drawing code. Although immersive applications typically make extensive use of custom drawing code, utility and productivity applications can often use standard views and controls to display their content.

The use of custom drawing code should be limited to situations where the content you display needs to change dynamically. For example, a drawing application would need to use custom drawing code to track the user's drawing commands and a game would be updating the screen constantly to reflect the changing game environment. In those situations, you would need to choose an appropriate drawing technology and create a custom view class to handle events and update the display appropriately.

On the other hand, if the bulk of your application's interface is fixed, you can render the interface in advance to one or more image files and display those images at runtime using `UIImageView` objects. You can layer image views with other content as needed to build your interface. For example, you could use `UILabel` objects to display configurable text and include buttons or other controls to provide interactivity.

Improving Drawing Performance

Drawing is a relatively expensive operation on any platform, and optimizing your drawing code should always be an important step in your development process. Table 4-2 lists several tips for ensuring that your drawing code is as optimal as possible. In addition to these tips, you should always use the available performance tools to test your code and remove hotspots and redundancies.

Table 4-2 Tips for improving drawing performance

Tip	Action
Draw minimally	During each update cycle, you should update only the portions of your view that actually changed. If you are using the <code>drawRect:</code> method of <code>UIView</code> to do your drawing, use the update rectangle passed to that method to limit the scope of your drawing. For OpenGL drawing, you must track updates yourself.
Mark opaque views as such	Compositing a view whose contents are opaque requires much less effort than compositing one that is partially transparent. To make a view opaque, the contents of the view must not contain any transparency and the <code>opaque</code> property of the view must be set to <code>YES</code> .
Remove alpha channels from opaque PNG files	If every pixel of a PNG image is opaque, removing the alpha channel avoids the need to blend the layers containing that image. This simplifies compositing of the image considerably and improves drawing performance.

Tip	Action
Reuse table cells and views during scrolling	Creating new views during scrolling should be avoided at all costs. Taking the time to create new views reduces the amount of time available for updating the screen, which leads to uneven scrolling behavior.
Avoid clearing the previous content during scrolling	By default, UIKit clears a view's current context buffer prior to calling its <code>drawRect:</code> method to update that same area. If you are responding to scrolling events in your view, clearing this region repeatedly during scrolling updates can be expensive. To disable the behavior, you can change the value in the <code>clearsContextBeforeDrawing</code> property to <code>NO</code> .
Minimize graphics state changes while drawing	Changing the graphics state requires effort by the window server. If you need to draw content that uses similar state information, try to draw that content together to reduce the number of state changes needed.

Maintaining Image Quality

Providing high-quality images for your user interface should be a priority in your design. Images provide a reasonably efficient way to display complicated graphics and should be used wherever they are appropriate. When creating images for your application, keep the following guidelines in mind:

- **Use the PNG format for images.** The PNG format provides high-quality image content and is the preferred image format for iPhone OS. In addition, iPhone OS includes an optimized drawing path for PNG images that is typically more efficient than other formats.
- **Create images so that they do not need resizing.** If you plan to use an image at a particular size, be sure to create the corresponding image resource at that size. Do not create a larger image and scale it down to fit, because scaling requires additional CPU cycles and requires interpolation. If you need to present an image at variable sizes, include multiple versions of the image at different sizes and scale down from an image that is relatively close to the target size.

Drawing with Quartz and UIKit

Quartz is the general name for the native window server and drawing technology in iPhone OS. The Core Graphics framework is at the heart of Quartz, and is the primary interface you use for drawing content. This framework provides data types and functions for manipulating the following:

- Graphics contexts
- Paths
- Images and bitmaps
- Transparency layers
- Colors, pattern colors, and color spaces
- Gradients and shadings
- Fonts

- PDF content

UIKit builds on the basic features of Quartz by providing a focused set of classes for graphics-related operations. The UIKit graphics classes are not intended as a comprehensive set of drawing tools—Core Graphics already provides that. Instead, they provide drawing support for other UIKit classes. UIKit support includes the following classes and functions:

- `UIImage`, which implements an immutable class for displaying images
- `UIColor`, which provides basic support for device colors
- `UIFont`, which provides font information for classes that need it
- `UIScreen`, which provides basic information about the screen
- Functions for generating a JPEG or PNG representation of a `UIImage` object
- Functions for drawing rectangles and clipping the drawing area
- Functions for changing and getting the current graphics context

For information about the classes and methods that comprise UIKit, see *UIKit Framework Reference*. For more information about the opaque types and functions that comprise the Core Graphics framework, see *Core Graphics Framework Reference*.

Configuring the Graphics Context

By the time your `drawRect:` method is called, your view's built-in drawing code has already created and configured a default graphics context for you. You can retrieve a pointer to this graphics context by calling the `UIGraphicsGetCurrentContext` function. This function returns a reference to a `CGContextRef` type, which you pass to Core Graphics functions to modify the current graphics state. Table 4-3 lists the main functions you use to set different aspects of the graphics state. For a complete list of functions, see *CGContext Reference*. This table also lists UIKit alternatives where they exist.

Table 4-3 Core graphics functions for modifying graphics state

Graphics state	Core Graphics functions	UIKit alternatives
Current transformation matrix (CTM)	<code>CGContextRotateCTM</code> <code>CGContextScaleCTM</code> <code>CGContextTranslateCTM</code> <code>CGContextConcatCTM</code>	None
Clipping area	<code>CGContextClipToRect</code>	None
Line: Width, join, cap, dash, miter limit	<code>CGContextSetLineWidth</code> <code>CGContextSetLineJoin</code> <code>CGContextSetLineCap</code> <code>CGContextSetLineDash</code> <code>CGContextSetMiterLimit</code>	None

Graphics state	Core Graphics functions	UIKit alternatives
Accuracy of curve estimation (flatness)	<code>CGContextSetFlatness</code>	None
Anti-aliasing setting	<code>CGContextSetAllowsAntialiasing</code>	None
Color: Fill and stroke settings	<code>CGContextSetRGBFillColor</code> <code>CGContextSetRGBStrokeColor</code>	<code>UIColor</code> class
Alpha value (transparency)	<code>CGContextSetAlpha</code>	None
Rendering intent	<code>CGContextSetRenderingIntent</code>	None
Color space: Fill and stroke settings	<code>CGContextSetFillColorSpace</code> <code>CGContextSetStrokeColorSpace</code>	None
Text: Font, font size, character spacing, text drawing mode	<code>CGContextSetFont</code> <code>CGContextSetFontSize</code> <code>CGContextSetCharacterSpacing</code>	<code>UIFont</code> class
Blend mode	<code>CGContextSetBlendMode</code>	The <code>UIImage</code> class and various drawing functions let you specify which blend mode to use.

The graphics context contains a stack of saved graphics states. When Quartz creates a graphics context, the stack is empty. Using the `CGContextSaveGState` function pushes a copy of the current graphics state onto the stack. Thereafter, modifications you make to the graphics state affect subsequent drawing operations but do not affect the copy stored on the stack. When you are done making modifications, you can return to the previous graphics state by popping the saved state off the top of the stack using the `CGContextRestoreGState` function. Pushing and popping graphics states in this manner is a fast way to return to a previous state and eliminates the need to undo each state change individually. It is also the only way to restore some aspects of the state, such as the clipping path, back to their original settings.

For general information about graphics contexts and using them to configure the drawing environment, see Graphics Contexts in *Quartz 2D Programming Guide*.

Creating and Drawing Images

iPhone OS provides support for loading and displaying images using both the UIKit and Core Graphics frameworks. How you determine which classes and functions to use to draw images depends on how you intend to use them. Whenever possible, though, it is recommended that you use the classes of UIKit for representing images in your code. Table 4-4 lists some of the usage scenarios and the recommended options for handling them.

Table 4-4 Usage scenarios for images

Scenario	Recommended usage
Display an image as the content of a view	Use a <code>UIImageView</code> class to load and display the image. This option assumes that your view's only content is an image. You can still layer other views on top of the image view to draw additional controls or content.
Display an image as an adornment for part of a view	Load and draw the image using the <code>UIImage</code> class.
Save some bitmap data into an image object	Use the <code>UIGraphicsBeginImageContext</code> function to create a new image-based graphics context. After creating this context, you can draw your image contents into it and then use the <code>UIGraphicsGetImageFromCurrentImageContext</code> function to generate an image based on what you drew. (If desired, you can even continue drawing and generate additional images.) When you are done creating images, use the <code>UIGraphicsEndImageContext</code> function to close the graphic context. If you prefer using Core Graphics, you can use the <code>CGBitmapContextCreate</code> function to create a bitmap graphics context and draw your image contents into it. When you finish drawing, use the <code>CGBitmapContextCreateImage</code> function to create a <code>CGImageRef</code> from the bitmap context. You can draw the Core Graphics image directly or use this it to initialize a <code>UIImage</code> object.
Save an image as a JPEG or PNG file	Create a <code>UIImage</code> object from the original image data. Call the <code>UIImageJPEGRepresentation</code> or <code>UIImagePNGRepresentation</code> function to get an <code>NSData</code> object, and use that object's methods to save the data to a file.

The following example shows how to load an image from your application's bundle. You can subsequently use this image object to initialize a `UIImageView` object, or you can store it and draw it explicitly in your view's `drawRect:` method.

```
NSString* imagePath = [[NSBundle mainBundle] pathForResource:@"myImage"
ofType:@"png"];
UIImage* myImageObj = [[UIImage alloc] initWithContentsOfFile:imagePath];
```

To draw an image explicitly in your view's `drawRect:` method, you can use any of the drawing methods available in `UIImage`. These methods let you specify where in your view you want to draw the image and therefore do not require you to create and apply a separate transform prior to drawing. Assuming you stored the previously loaded image in a member variable called `anImage`, the following example draws that image at the point (10, 10) in the view.

```
- (void)drawRect:(CGRect)rect
{
    // Draw the image
    [anImage drawAtPoint:CGPointMake(10, 10)];
}
```

Important: If you use the `CGContextDrawImage` function to draw bitmap images directly, the image data is inverted along the y axis by default. This is because Quartz images assume a coordinate system with a lower-left corner origin and positive coordinate axes extending up and to the right from that point. Although you can apply a transform before drawing, the simpler (and recommended) way to draw Quartz images is to wrap them in a `UIImage` object, which compensates for this difference in coordinate spaces automatically. For more information on creating and drawing images using Core Graphics, see *Quartz 2D Programming Guide*.

Creating and Drawing Paths

A path is a description of a 2D geometric scene that uses a sequence of lines and Bézier curves to represent that scene. UIKit includes the `UIRectFrame` and `UIRectFill` functions (among others) for creating drawing simple paths such as rectangles in your views. Core Graphics also includes convenience functions for creating simple paths such as rectangles and ellipses. For more complex paths, you must create the path yourself using the functions of the Core Graphics framework.

To create a path, you use the `CGContextBeginPath` function to configure the graphics context to receive path commands. After calling that function, you use other path-related functions to set the path's starting point, draw lines and curves, add rectangles and ellipses, and so on. When you are done specifying the path geometry, you can paint the path directly or create a `CGPathRef` or `CGMutablePathRef` data type to store a reference to that path for later use.

When you want to draw a path in your view, you can stroke it, fill it, or do both. Stroking a path with a function such as `CGContextStrokePath` creates a line centered on the path using the current stroke color. Filling the path with the `CGContextFillPath` function uses the current fill color or fill pattern to fill the area enclosed by the path's line segments.

For more information on how to draw paths, including information about how you specify the points for complex path elements, see Paths in *Quartz 2D Programming Guide*. For information on the functions you use to create paths, see *CGContext Reference* and *CGPath Reference*.

Creating Patterns, Gradients, and Shadings

The Core Graphics framework includes additional functions for creating patterns, gradients, and shadings. You use these types to create non monochrome colors and use them to fill the paths you create. Patterns are created from repeating images or content. Gradients and shadings provide different ways to create smooth transitions from color to color.

The details for creating and using patterns, gradients, and shadings are all covered in *Quartz 2D Programming Guide*.

Drawing with OpenGL ES

The **Open Graphics Library (OpenGL)** is a cross-platform C-based interface used to create 2D and 3D content on desktop systems. It is typically used by games developers or anyone needing to perform drawing with high frame rates. You use OpenGL functions to specify primitive structures such as points, lines, and polygons

and the textures and special effects to apply to those structures to enhance their appearance. The functions you call send graphics commands to the underlying hardware, where they are then rendered. Because rendering is done mostly in hardware, OpenGL drawing is usually very fast.

OpenGL for Embedded Systems is a pared-down version of OpenGL that is designed for mobile devices and takes advantage of modern graphics hardware. If you want to create OpenGL content for iPhone OS–based devices—that is, iPhone or iPod Touch—you'll use OpenGL ES. The OpenGL ES framework (`OpenGLES.framework`) provided with iPhone OS conforms to the [OpenGL ES v1.1 specification](#). You can find out more about OpenGL ES by reading [Polygons In Your Pocket: Introducing OpenGL ES](#).

This section is designed to get you started writing OpenGL ES applications for iPhone OS–based devices. [“Setting Up a Rendering Surface”](#) (page 95) provides step-by-step instructions for creating a surface that you can draw to using OpenGL ES. But before you start writing the OpenGL ES portion of your application, you'll want to read [“Implementation Details”](#) (page 98) to learn about the capabilities of iPhone OS–based devices and the specifics of the OpenGL ES implementation in iPhone OS. [“Best Practices”](#) (page 96) provides coding guidelines that can help your application perform optimally.

Setting Up a Rendering Surface

The setup for drawing with OpenGL ES is straightforward and requires the same types of tasks you'd perform to draw to surfaces on a Macintosh computer. The primary difference is that you'll use the EAGL API to set up the window surface instead of an API such as CGL or AGL. The EAGL API provides the interface between the OpenGL ES renderer and the windows and views of an iPhone application. (See *OpenGL ES Framework Reference*.)

When you set up your Xcode project, make sure you link to `OpenGLES.framework`. Then, set up a surface for rendering by following these steps:

1. Subclass `UIView` and set up the view appropriately for your iPhone application.
2. Override the `layerClass` method of the `UIView` class so that it returns a `CAEAGLLayer` object rather than a `CALayer` object.

```
+ (Class) layerClass
{
    return [CAEAGLLayer class];
}
```

3. Get the layer associated with the view by calling the `layer` method of `UIView`.

```
myEAGLLayer = (CAEAGLLayer*)self.layer;
```

4. Set the layer properties.

For optimal performance, it's recommended that you mark the layer as opaque by setting the `opaque` property provided by the `CALayer` class. See [“Best Practices”](#) (page 96).

5. Optionally configure the surface properties of the rendering surface by assigning a new dictionary of values to the `drawableProperties` property of the `CAEAGLLayer` object.

The EAGL framework lets you specify custom color formats and whether or not the native surface retains its contents after presenting them. You identify these properties in the dictionary using the `kEAGLDrawablePropertyColorFormat` and `kEAGLDrawablePropertyRetainedBacking` keys. For a list of values for these keys, see the `EAGLDrawable` protocol.

6. Create a new `EAGLContext` object to manage the drawing context. You typically create and initialize this object as follows:

```
EAGLContext* myContext = [[EAGLContext alloc]
initWithAPI:kEAGLRenderingAPIOpenGLES1];
```

If you want to share objects (texture objects, vertex buffer objects, and so forth) between multiple contexts, use the `initWithAPI:sharegroup:` initialization method instead. For each context that should share a given set of objects, pass the same `EAGLSharegroup` object in the `sharegroup` parameter.

7. Make your `EAGLContext` object the context for the current thread using the `setCurrentContext:` class method.

You can have one current context per thread.

8. Create and bind a new render buffer to the `GL_RENDERBUFFER_OES` target. (Typically you would do this in a two-step process using the `glGenRenderbufferOES` function to allocate an unused name and the `glBindRenderbufferOES` functions create and bind the render buffer to that name.)
9. Attach the newly created render buffer target to your view's layer object using the `renderbufferStorage:fromDrawable:` method of your `EAGLContext` object. (The layer provides the underlying storage for the render buffer.) For example, given the context created previously, you would use the following code to bind the buffer to the view's layer (obtained previously and stored in the `myEAGLLayer` variable):

```
[myContext renderbufferStorage:GL_RENDERBUFFER_OES fromDrawable:myEAGLLayer];
```

The width, height, and format of the render buffer storage are derived from the bounds and properties of the `CAEAGLLayer` object at the moment you call the `renderbufferStorage:fromDrawable:` method. If you change the layer's bounds later, Core Animation scales the content by default. To avoid scaling, you must recreate the renderbuffer storage by calling `renderbufferStorage:fromDrawable:`.

10. Configure your frame buffer as usual and bind your render buffer to the attach points of your frame buffer.

Best Practices

To create great OpenGL ES applications that run optimally, whenever possible follow the guidelines discussed in this section. You'll also want to read ["Implementation Details"](#) (page 98) to see how your code can take advantage of the specific features of the GPU.

General Guidelines

When developing OpenGL ES applications for iPhone OS–based devices:

- Use an EAGL surface that is the same size as size of the screen. Read the `bounds` property of the `UIScreen` class to get the screen size.
- Do not apply any Core Animation transforms to the `CAEAGLLayer` object that contains the EAGL window surface.
- Set the `opaque` property of the `CALayer` class to mark the `CAEAGLLayer` object as opaque.

- Do not place any other Core Animation layers or UIKit views above the `CAEAGLLayer` object that you're rendering to.
- If your application needs to present landscape content, avoid transforming the layer. Instead, set the OpenGL ES state to rotate everything by changing the Model/View/Projection transform, and swapping the width and height arguments to the `glViewport` and `glScissor` functions.
- Limit interactions between OpenGL ES and UIKit or Core Animation rendering. For example, avoid rendering with OpenGL ES while you render notifications, messages, or any other user interface controls provided by UIKit or Core Animation.
- Economize memory usage. iPhone OS–based devices use a shared memory system. Memory used by your application's graphics is not available for the system. For example, after loading GL textures, free your copy of the pixel data if you no longer need to use it. (See “Memory” (page 100))
- To implement transparency, use alpha blending instead of alpha testing. Alpha testing is considerably more expensive than alpha blending.
- Disable unused or unnecessary features. For instance, do not enable lighting or blending if you do not need them.
- Minimize scissor state changes. They are considerably more expensive on OpenGL ES for iPhone OS than they are on OpenGL for Mac OS X.
- Minimize the number of times you call the `EAGLContext` class method `setCurrentContext:` during a rendering frame. Changing the current surface on OpenGL ES for iPhone OS is considerably more expensive than it is on OpenGL for Mac OS X. Rather than using multiple contexts, consider using multiple frame buffer objects instead.
- Operations that depend on completing previous rendering commands (such as `glTexSubImage`, `glCopyTexImage`, `glCopyTexSubImage`, `glReadPixels`, `glFlush`, `glFinish`, or `setCurrentContext:`) can be very expensive if you perform them in the middle of a frame. If you need these operations, perform them at the beginning or end of a frame.

CPU Usage

Respecifying OpenGL ES state can cause the CPU to perform unnecessary work. Use techniques to reduce CPU usage. For example, use a texture atlas. This allows you to perform additional draw calls without changing the texture binding. (Or even better, collapse multiple draw calls into one.) If you create a texture atlas, you also need to sort state calls to avoid rebinding the texture. Otherwise, you won't see a performance benefit.

Vertex Data

When creating geometry:

- Reduce overall geometry by doing such things as using indexed triangle strips and providing only as much detail as the user can see.
- To achieve fine detail in your drawing without increasing the vertex count, consider using DOT3 lighting or textures.
- Memory bandwidth is limited, so use the smallest acceptable type for your data. Specify vertex colors using 4 unsigned byte values. Specify texture coordinates with 2 or 4 unsigned byte or short values instead of floating-point values, if you can.

Textures

To get the best performance with textures:

- Use textures with the smallest size per pixel that you can afford. If possible, use 565 textures instead of 8888.
- Use compressed textures stored in the PVRTC format. See the specification for the extension `GL_IMG_texture_compression_pvrtc`.
- Use mipmapping with the `LINEAR_MIPMAP_NEAREST` option.
- Create and load all textures prior to rendering. Don't upload or modify textures during a frame. Specifically, avoid calling the functions `glTexSubImage` or `glCopyTexSubImage` in the middle of a frame.
- Use multitexturing rather than applying textures over multiple passes.

Drawing Order

Drawing order is important for hardware that uses tile based deferred rendering. (See “[Rendering Path](#)” (page 100).)

- Don't waste CPU time sorting objects front to back. The tile based deferred rendering model used by the GPU makes sorting unnecessary.
- Draw opaque objects first; draw alpha blended objects last.

Lighting

Simplify lighting as much as possible.

- Use the fewest lights possible and the simplest lighting type for your application. For example, consider using directional lights instead of spot lighting, which incurs a higher performance cost.
- If you can, precompute lighting. Static lighting gives better performance than dynamic. You can compute the lighting ahead of time and store the result in textures or color arrays that you can look up later.

Debugging and Tuning

The Instruments application includes an OpenGL ES instrument that you can use to gather information about the runtime behavior of your OpenGL code. In addition, you can set a breakpoint on the `opengl_error_break` symbol in GDB to see when OpenGL errors are generated.

Implementation Details

Understanding the features of the hardware and the specifics of the implementation of OpenGL ES can help you tailor your code to get the best performance. Use the information in this section, along with “[Best Practices](#)” (page 96), as you design your OpenGL ES application.

OpenGL ES Implementation

The OpenGL ES implementation in iPhone OS differs from other implementations of OpenGL ES in the following ways:

- The maximum texture size is 1024 x 1024.
- 2D texture targets are supported; other texture targets are not.
- Stencil buffers aren't available.

Hardware Capabilities

When developing any OpenGL application, it's important to check for the functionality that your application uses and have a contingency in place if the hardware doesn't support the particular feature or extension that you want to use. This is true for Macintosh computers and it is even more important when writing OpenGL ES applications for iPhone OS-based devices. It's essential that you understand the capabilities of the specific hardware that you are writing for. Keep in mind that, unlike OpenGL on Macintosh computers, there is no software rendering fallback option for iPhone OS-based devices.

The graphics hardware for iPhone OS-based devices has the following limitations:

- The texture magnification and magnification filters (within a texture level) must match. For example:
 - ❑ **Supported:** `GL_TEXTURE_MAG_FILTER = GL_LINEAR, GL_TEXTURE_MIN_FILTER = GL_LINEAR_MIPMAP_LINEAR`
 - ❑ **Supported:** `GL_TEXTURE_MAG_FILTER = GL_NEAREST, GL_TEXTURE_MIN_FILTER = GL_LINEAR_MIPMAP_NEAREST`
 - ❑ **Not Supported:** `GL_TEXTURE_MAG_FILTER = GL_NEAREST, GL_TEXTURE_MIN_FILTER = GL_LINEAR_MIPMAP_LINEAR`

There are a few, rarely used texture environment operations that aren't available:

- If the value of `GL_COMBINE_RGB` is `GL_MODULATE`, only one of the two operands may read from an `GL_ALPHA` source.
- If the value of `GL_COMBINE_RGB` is `GL_INTERPOLATE`, `GL_DOT3_RGB`, or `GL_DOT3_RGBA`, then several combinations of `GL_CONSTANT` and `GL_PRIMARY_COLOR` sources and `GL_ALPHA` operands do not work properly.
- If the value of `GL_COMBINE_RGB` or `GL_COMBINE_ALPHA` is `GL_SUBTRACT`, then `GL_SCALE_RGB` or `GL_SCALE_ALPHA` must be 1.0.
- If the value of `GL_COMBINE_ALPHA` is `GL_INTERPOLATE` or `GL_MODULATE`, only one of the two sources can be `GL_CONSTANT`.
- The value of `GL_TEXTURE_ENV_COLOR` must be the same for all texture units.

Supported Extensions

These are the OpenGL ES extensions that you can use when developing OpenGL ES applications for iPhone OS-based devices:

- `GL_OES_blend_subtract`
- `GL_OES_compressed_paletted_texture`
- `GL_OES_depth24`
- `GL_OES_draw_texture`
- `GL_OES_framebuffer_object`
- `GL_OES_mapbuffer`
- `GL_OES_matrix_palette`
- `GL_OES_point_size_array`
- `GL_OES_point_sprite`
- `GL_OES_read_format`
- `GL_OES_rgb8_rgba8`
- `GL_OES_texture_mirrored_repeat`
- `GL_EXT_texture_filter_anisotropic`
- `GL_EXT_texture_lod_bias`
- `GL_IMG_read_format`
- `GL_IMG_texture_compression_pvrtc`
- `GL_IMG_texture_format_BGRA8888`

Memory

OpenGL ES applications should use no more than 24 MB of memory for both textures and surfaces. This 24 MB is not dedicated graphics memory but comes from the main system memory. Because main memory is shared with other iPhone applications and the system, your application should use as little of it as possible. “[Best Practices](#)” (page 96) provides several guidelines for ways to use memory economically. In particular, see “[General Guidelines](#)” (page 96) and “[Vertex Data](#)” (page 97).

Rendering Path

The GPU in the iPhone and iPod touch is a [PowerVR MBX Lite](#). This GPU uses a technique known as **Tile Based Deferred Rendering** (TBDR). When you submit OpenGL ES commands for rendering, TBDR behaves very differently from a streaming renderer. A streaming renderer simply executes rendering commands in order, one after another. In contrast, a TBDR defers any rendering until it accumulates a large number of rendering commands, and then operates on this command list as a single scene. The framebuffer is divided up into a number of tiles, and the scene is drawn once for each tile, each time drawing only the content that is actually visible within that tile. The TBDR approach has several advantages and disadvantages compared to streaming renderers. Understanding these differences will help you write better performing software.

The most significant advantage of TBDR is that it can make much more efficient use of available bandwidth to memory. Constraining rendering to only one tile allows the GPU to more effectively cache the framebuffer, making depth testing blending much more efficient. Otherwise, the memory bandwidth consumed by these framebuffer operations often becomes a significant performance bottleneck.

When using deferred rendering, some operations become more expensive. For example, if you call the function `glTexSubImage` in the middle of a frame, the accumulated command list may include commands from both before and after the call to `glTexSubImage`. This command list needs to reference both the old

and new version of the texture image at the same time, forcing the entire texture to be duplicated even if only a small portion of the texture is updated. Duplication can make functions such as `glTexSubImage` significantly more expensive on a deferred renderer than a streaming renderer.

The PowerVR GPU relies on more than just TBDR to optimize performance; it performs hidden surface removal before fragment processing. If the GPU determines that a pixel won't be visible, it discards the pixel without performing texture sampling or fragment color calculations. Removing hidden pixels can significantly improve performance for scenes that have obscured content. To gain the most benefit from this feature, you should try to draw as much of the scene with opaque content as possible and minimize use of blending and alpha testing.

For more information on exactly how these features are implemented and how your application can best take advantage of them, see [PowerVR Technology Overview](#) and [PowerVR MBX 3D Application Development Recommendations](#).

Simulator Capabilities

The iPhone simulator includes a complete and conformant implementation of OpenGL ES 1.1 that you can use for your application development. This implementation differs in a few ways from the implementation found in iPhone OS–based devices. In particular, the simulator does not have the same limitations regarding texture magnification filters or texture environment operations that are described in “[Hardware Capabilities](#)” (page 99). In addition, the simulator supports antialiased lines while iPhone OS–based devices do not.

Important: It is important to understand that the rendering performance of OpenGL ES in the simulator has no relation to the performance of OpenGL ES on an actual device. The simulator provides an optimized software rasterizer that takes advantage of the vector processing capabilities of your Macintosh computer. As a result, your OpenGL ES code may run faster or slower in Mac OS X (depending on your computer and what you are drawing) than on an actual device. Therefore, you should always profile and optimize your drawing code on a real device and not assume that the simulator reflects real-world performance.

The following sections provide additional details about the OpenGL ES support available in the iPhone simulator.

Supported Extensions

The iPhone simulator supports all of the OpenGL ES 1.1 core functionality and most of the extensions supported by iPhone OS–based devices. The following extensions are not supported by the simulator, however:

- [GL_OES_draw_texture](#)
- [GL_OES_matrix_palette](#)
- [GL_EXT_texture_filter_anisotropic](#)
- [GL_IMG_texture_compression_pvrtc](#)

For a list of the extensions supported by the hardware, see “[Supported Extensions](#)” (page 99).

Memory

On a device, OpenGL ES applications can use no more than 24 MB of memory for both textures and surfaces. The simulator does not enforce this limit. As a result, your code can allocate as much memory as your computer's rendering hardware supports. Be sure to keep track of the size of your assets during development.

Rendering Path

In contrast to the Tile Based Deferred Rendering technique used in devices, the simulator's software rasterizer uses a traditional streaming model for OpenGL ES commands. Objects are transformed and rendered immediately as you specify them. Consequently, the performance of some operations can differ significantly from that on actual devices.

As with any two different implementations of OpenGL ES, there may be small differences between the pixels rendered by the simulator and those rendered by the device. For example, OpenGL ES allows some calculations, such as color interpolation and texture mipmap filtering, to be approximated. In general, the two implementations will produce similar results, but do not rely on them to be bit-for-bit identical.

For More Information

You may want to consult these resources as you develop OpenGL ES applications for iPhone OS–based devices:

- [OpenGL ES 1.X Specification](#) is the official definition of this technology provided by the [Khronos Group](#). You'll also find other useful information on this website.
- [PowerVR MBX OpenGL ES 1.x SDK page](#) provides information about the specific OpenGL ES implementation supported by the PowerVR MBX graphics hardware.
- [OpenGL ES 1.1 Reference Pages](#) provides a complete reference to OpenGL ES specification, indexed alphabetically as well as by theme.
- *OpenGL ES Framework Reference* describes the functions and constants that provide the interface between OpenGL ES and the iPhone user interface.

Applying Core Animation Effects

Core Animation is an Objective-C framework that provides infrastructure for creating fluid, real-time animations quickly and easily. Core Animation is not a drawing technology itself, in the sense that it does not provide primitive routines for creating shapes, images, or other types of content. Instead, it is a technology for manipulating and displaying content that you created using other technologies.

Most applications can benefit from using Core Animation in some form in iPhone OS. Animations provide feedback to the user about what is happening. For example, when the user navigates through the Settings application, screens slide in and out of view based on whether the user is navigating further down the preferences hierarchy or back up to the root node. This kind of feedback is important and provides contextual information for the user. It also enhances the visual style of an application.

In most cases, you may be able to reap the benefits of Core Animation with very little effort. For example, several properties of the `UIView` class (including the view's frame, center, color, and opacity—among others) can be configured to trigger animations when their values change. You have to do some work to let UIKit know that you want these animations performed, but the animations themselves are created and run automatically for you. For information about how to trigger the built-in view animations, see [“Animating Views”](#) (page 67).

When you go beyond the basic animations, you must interact more directly with Core Animation classes and methods. The following sections provide information about Core Animation and show you how to work with its classes and methods to create typical animations in iPhone OS. For additional information about Core Animation and how to use it, see *Core Animation Programming Guide*.

About Layers

The key technology in Core Animation is the layer object. Layers are lightweight objects that are similar in nature to views, but that are actually model objects that encapsulate geometry, timing, and visual properties for the content you want to display. The content itself is provided in one of three ways:

- You can assign a `CGImageRef` to the `contents` property of the layer object.
- You can assign a delegate to the layer and let the delegate handle the drawing.
- You can subclass `CALayer` and override one of the display methods.

When you manipulate a layer object's properties, what you are actually manipulating is the model-level data that determines how the associated content should be displayed. The actual rendering of that content is handled separately from your code and is heavily optimized to ensure it is fast. All you must do is set the layer content, configure the animation properties, and then let Core Animation take over.

For more information about layers and how they are used, see *Core Animation Programming Guide*.

About Animations

When it comes to animating layers, Core Animation uses separate animation objects to control the timing and behavior of the animation. The `CAAnimation` class and its subclasses provide different types of animation behaviors that you can use in your code. You can create simple animations that migrate a property from one value to another, or you can create complex keyframe animations that track the animation through the set of values and timing functions you provide.

Core Animation also lets you group multiple animations together into a single unit, called a transaction. The `CATransaction` object manages the group of animations as a unit. You can also use the methods of this class to set the duration of the animation.

For examples of how to create custom animations, see *Animation Types and Timing Programming Guide*.

Text and Web

The text system in iPhone OS was designed with the modest needs of mobile users in mind. The text system was designed to handle the single and multi-line text input that are commonly used in email and SMS programs. The text system also supports Unicode and has several different input methods, making it easy to display and read text in many different languages.

About Text and Web Support

The text system in iPhone OS provides a tremendous amount of power while still being very simple to use. The UIKit framework includes several high-level classes for managing the display and input of text. This framework also includes a more advanced class for displaying HTML and JavaScript-based content.

The following sections describe the basic support for text and web content in iPhone OS. For more information about each of the classes listed in this section, see *UIKit Framework Reference*.

Text Views

The UIKit framework provides three primary classes for displaying text content:

- `UILabel` displays static text strings
- `UITextField` displays a single line of editable text
- `UITextView` displays multiple lines of editable text

These classes support the display of arbitrarily large amounts of text, although labels and text fields are typically used for relatively small amounts of text. To make the displayed text easier to read on the smaller screens of iPhone OS–based devices, however, these classes do not support the kinds of advanced formatting you might find in desktop operating systems like Mac OS X. All three classes still allow you to specify the font information, including size and styling options, that you might otherwise want, but the font information you specify is applied to all of the text associated with the object.

Figure 5-1 shows examples of the available text classes as they appear on screen. These examples were taken from the *UICatalog* sample application, which demonstrates many of the views and controls available in UIKit. The image on the left shows several different styles of text fields while the image on the right shows a single text view. The callouts displayed on the gray background are themselves `UILabel` objects embedded inside the table cells used to display the different views. There is also a `UILabel` object with the text “Left View” at the bottom of the screen on the left.

Figure 5-1 Text classes in the UICatalog application



When working with editable text views, you should always provide a delegate object to manage the editing session. Text views send several different notifications to the delegate to let them know when editing begins, when it ends, and to give them a chance to override some editing actions. For example, the delegate can decide if the current text contains a valid value and prevent the editing session from ending if it does not. When editing does finally end, you also use the delegate to get the resulting text value and update your application's data model.

Because there are slight differences in their intended usage, the delegate methods for each text view are slightly different. A delegate that supports the `UITextField` class implements the methods of the `UITextFieldDelegate` protocol. Similarly, a delegate that supports the `UITextView` class implements the methods of the `UITextViewDelegate` protocol. In both cases, you are not required to implement any of the protocol methods but if you do not, the text field is not going to be of much use to you. For more information on the methods in these two protocols, see *UITextFieldDelegate Protocol Reference* and *UITextViewDelegate Protocol Reference*.

Web View

The `UIWebView` class lets you integrate what is essentially a miniature web browser into your application's user interface. The `UIWebView` class makes full use of the same web technologies used to implement Safari in iPhone OS, including full support for HTML, CSS, and JavaScript content. The class also supports many of the built-in gestures that users are familiar with in Safari. For example, you can double-click and pinch to zoom in and out of the page and you can scroll around the page by dragging your finger.

In addition to displaying content, you can also use a web view object to gather input from the user through the use of web forms. Like the other text classes in UIKit, if you have an editable text field on a form in your web page, tapping that field brings up a keyboard so that the user can enter text. Because it is an integral part of the web experience, the web view itself manages the displaying and dismissing of the keyboard for you.

Figure 5-2 shows an example of a `UIWebView` object from the the *UICatalog* sample application, which demonstrates many of the views and controls available in UIKit. Because it just displays HTML content, if you want the user to be able to navigate pages much like they would in a web browser, you need to add controls to do so. For example, the web view in the figure occupies the space below the text field containing the target URL and does not contain the text field itself.

Figure 5-2 A web view



A web view provides information about when pages are loaded, and whether there were any load errors, through its associated delegate object. A web delegate is an object that implements one or more methods of the `UIWebViewDelegate` protocol. Your implementations of the delegate methods can respond to failures or perform other tasks related to the loading of a web page. For more information about the methods of the `UIWebViewDelegate` protocol, see *UIWebViewDelegate Protocol Reference*.

Keyboards and Input Methods

Whenever the user taps in an object capable of accepting text input, the object asks the system to display an appropriate keyboard. Depending on the needs of your program and the user's preferred language, the system might display one of several different keyboards. Although your application cannot control the user's preferred language (and thus the keyboard's input method), it can control attributes of the keyboard that indicate its intended use, such as the configuration of any special keys and its behaviors.

You configure the attributes of the keyboard directly through the text objects of your application. The `UITextField` and `UITextView` classes both conform to the `UITextInputTraits` protocol, which defines the properties for configuring the keyboard. Setting these properties programmatically or in the Interface Builder inspector window causes the system to display the keyboard of the designated type.


Note: Although the `UIWebView` class does not support the `UITextInputTraits` protocol directly, you can configure some keyboard attributes for text input elements. In particular, you can include `autocorrect` and `autocapitalization` attributes in the definition of an input element to specify the keyboard's behaviors, as shown in the following example.

```
<input type="text" size="30" autocorrect="off" autocapitalization="on">
```

You cannot specify the keyboard type in input elements. The web view displays a custom keyboard that is based on the default keyboard but includes some additional controls for navigating between form elements.

The default keyboard configuration is designed for general text input. Figure 5-3 displays the default keyboard along with several other keyboard configurations. The default keyboard displays an alphabetical keyboard initially but the user can toggle it and display numbers and punctuation as well. Most of the other keyboards offer similar features as the default keyboard but provide additional buttons that are specially suited to particular tasks. However, the phone and numerical keyboards offer a dramatically different layout that is tailored towards numerical input.

Figure 5-3 Several different keyboard types



To facilitate the language preferences of different users, iPhone OS also supports different input methods and keyboard layouts for different languages, some of which are shown in Figure 5-4. The input method and layout for the keyboard is determined by the user's language preferences.

Figure 5-4 Several different keyboards and input methods

Managing the Keyboard

Although many UIKit objects display the keyboard automatically in response to user interactions, your application still has some responsibilities for configuring and managing the keyboard. The following sections describe those responsibilities.

Receiving Keyboard Notifications


When the keyboard is shown or hidden, iPhone OS sends out the following notifications to any registered observers:

- `UIKeyboardWillShowNotification`
- `UIKeyboardDidShowNotification`
- `UIKeyboardWillHideNotification`
- `UIKeyboardDidHideNotification`

The system sends keyboard notifications when the keyboard first appears, when it disappears, any time the owner of the keyboard changes, or any time your application's orientation changes. In each situation, the system sends only the appropriate subset of messages. For example, if the owner of the keyboard changes, the system sends a `UIKeyboardWillHideNotification` message, but not a `UIKeyboardDidHideNotification` message, to the current owner because the change never causes the keyboard to be hidden. The delivery of the `UIKeyboardWillHideNotification` is simply a way to alert the current owner that it is about to lose the keyboard focus. Changing the orientation of the keyboard does send both will and did hide notifications, however, because the keyboards for each orientation are different and thus the original must be hidden before the new one is displayed.

Each keyboard notification includes information about the size and position of the keyboard on the screen. You should always use the information in these notifications as opposed to assuming the keyboard is a particular size or in a particular location. The size of the keyboard is not guaranteed to be the same from one input method to another and may also change between different releases of iPhone OS. In addition, even for a single language and system release, the keyboard dimensions can vary depending on the orientation of your application. For example, Figure 5-5 shows the relative sizes of the URL keyboard in both the portrait and landscape modes. Using the information inside the keyboard notifications ensures that you always have the correct size and position information.

Figure 5-5 Relative keyboard sizes in portrait and landscape modes



Note: The rectangle contained in the `UIKeyboardBoundsUserInfoKey` of the info dictionary should be used only for the size information it contains. Because the keyboard is animated into position, the actual bounding rectangle of the keyboard changes over time. The starting and ending positions of the keyboard are therefore stored in the info dictionary under the `UIKeyboardCenterBeginUserInfoKey` and `UIKeyboardCenterEndUserInfoKey` keys.

One reason to use keyboard notifications is so that you can reposition content that is obscured by the keyboard when it is visible. For information on how to handle this scenario, see [“Moving Content That Is Located Under the Keyboard”](#) (page 113).

Displaying the Keyboard

When the user taps a view, the system automatically designates that view as the first responder. When this happens to a view that contains editable text, the view initiates an editing session for that text. At the beginning of that editing session, the view asks the system to display the keyboard, if it is not already visible. If the keyboard is already visible, the change in first responder causes text input from the keyboard to be redirected to the newly tapped view.

Because the keyboard is displayed automatically when a view becomes the first responder, you often do not need to do anything to display it. However, you can programmatically display the keyboard for an editable text view by calling that view's `becomeFirstResponder` method. Calling this method makes the target view the first responder and begins the editing process just as if the user had tapped on the view.

If your application manages several text-based views on a single screen, it is a good idea to track which view is currently the first responder so that you can dismiss the keyboard later.

Dismissing the Keyboard

Although it typically displays the keyboard automatically, the system does not dismiss the keyboard automatically. Instead, it is your application's responsibility to dismiss the keyboard at the appropriate time. Typically, you would do this in response to a user action. For example, you might dismiss the keyboard when the user taps the Return or Done button on the keyboard or taps some other button in your application's interface. Depending on how you configured the keyboard, you might need to add some additional controls to your user interface to facilitate the keyboard's dismissal.

To dismiss the keyboard, you call the `resignFirstResponder` method of the text-based view that is currently the first responder. When a text view resigns its first responder status, it ends its current editing session, notifies its delegate of that fact, and dismisses the keyboard. In other words, if you have a variable called `myTextField` that points to the `UITextField` object that is currently the first responder, dismissing the keyboard is as simple as doing the following:

```
[myTextField resignFirstResponder];
```

Everything from that point on is handled for you automatically by the text object.

Moving Content That Is Located Under the Keyboard

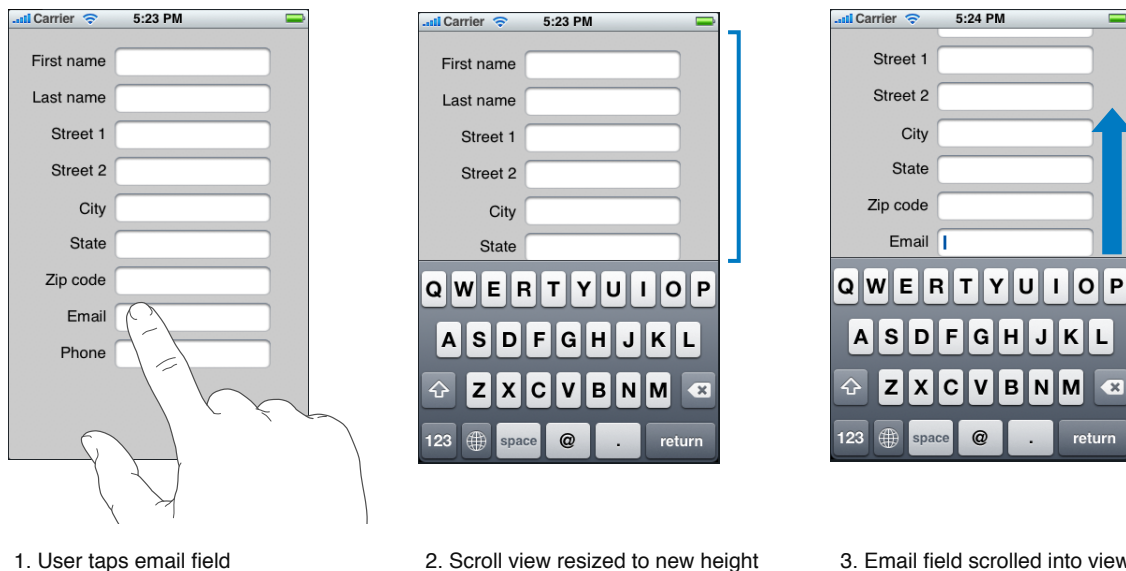
When asked to display the keyboard, the system slides it in from the bottom of the screen and positions it over your application's content. Because it is placed on top of your content, it is possible for the keyboard to be placed on top of the text object that the user wanted to edit. When this happens, you must adjust your content so that the target object remains visible.

Adjusting your content typically involves temporarily resizing one or more views and positioning them so that the text object remains visible. The simplest way to manage text objects with the keyboard is to embed them inside a `UIScrollView` object (or one of its subclasses like `UITableView`). When the keyboard is displayed, all you have to do is resize the scroll view and scroll the desired text object into position. Thus, in response to a `UIKeyboardDidShowNotification`, your handler method would do the following:

1. Get the size of the keyboard.
2. Subtract the keyboard height from the height of your scroll view.
3. Scroll the target text field into view.

Figure 5-6 illustrates the preceding steps for a simple application that embeds several text fields inside a `UIScrollView` object. When the keyboard appears, the notification handler method resizes the scroll view and then uses the `scrollRectToVisible:animated:` method of `UIScrollView` to scroll the tapped text field (in this case the email field) into view.

Figure 5-6 Adjusting content to accommodate the keyboard



Note: When setting up your own scroll views, be sure to configure the autosizing rules for any content views appropriately. In the preceding figure, the text fields are actually embedded inside a generic `UIView` object, which is itself embedded in the `UIScrollView` object. If the generic view's `UIViewAutoresizingFlexibleWidth` and `UIViewAutoresizingFlexibleHeight` autosizing options are set, changing the scroll view's frame size also changes the frame of the generic view, which could yield undesirable results. Disabling these options for that view ensures that the view retains its size and its contents are scrolled correctly.

Listing 5-1 shows the code for registering to receive keyboard notifications and shows the handler methods for those notifications. This code is implemented by the view controller that manages the scroll view, and the `scrollView` variable is an outlet that points to the scroll view object. Each handler method gets the keyboard size from the info dictionary of the notification and adjusts the scroll view height by the corresponding amount. In addition, the `keyboardWasShown:` method scrolls the rectangle of the active text field, which is stored in the `activeField` variable and set when the `textFieldDidBeginEditing:` delegate method is called. (In this example, the view controller also acts as the delegate for each of the text fields.)

Listing 5-1 Handling the keyboard notifications

```
- (void)registerForKeyboardNotifications
{
    [[NSNotificationCenter defaultCenter] addObserver:self
        selector:@selector(keyboardWasShown:)
        name:UIKeyboardDidShowNotification object:nil];

    [[NSNotificationCenter defaultCenter] addObserver:self
        selector:@selector(keyboardWasHidden:)
        name:UIKeyboardDidHideNotification object:nil];
}

// Called when the UIKeyboardDidShowNotification is sent.
- (void)keyboardWasShown:(NSNotification*)aNotification
{
    if (keyboardShown)
        return;

    NSDictionary* info = [aNotification userInfo];

    // Get the size of the keyboard.
    NSValue* aValue = [info objectForKey:UIKeyboardBoundsUserInfoKey];
    CGSize keyboardSize = [aValue CGRectValue].size;

    // Resize the scroll view (which is the root view of the window)
    CGRect viewFrame = [scrollView frame];
    viewFrame.size.height -= keyboardSize.height;
    scrollView.frame = viewFrame;

    // Scroll the active text field into view.
    CGRect textFieldRect = [activeField frame];
    [scrollView scrollRectToVisible:textFieldRect animated:YES];

    keyboardShown = YES;
}
```

```
// Called when the UIKeyboardDidHideNotification is sent
- (void)keyboardWasHidden:(NSNotification*)aNotification
{
    NSDictionary* info = [aNotification userInfo];

    // Get the size of the keyboard.
    NSValue* aValue = [info objectForKey:UIKeyboardBoundsUserInfoKey];
    CGSize keyboardSize = [aValue CGRectValue].size;

    // Reset the height of the scroll view to its original value
    CGRect viewFrame = [scrollView frame];
    viewFrame.size.height += keyboardSize.height;
    scrollView.frame = viewFrame;

    keyboardShown = NO;
}
```

The `keyboardShown` variable from the preceding listing is a Boolean value used to track whether the keyboard is already visible. If your interface has multiple text fields, the user can tap between them to edit the values in each one. When that happens, however, the keyboard does not disappear but the system does still generate `UIKeyboardDidShowNotification` notifications each time editing begins in a new text field. By tracking whether the keyboard was actually hidden, this code prevents the scroll view from being reduced in size more than once.

Drawing Text

In addition to the UIKit classes for displaying and editing text, iPhone OS also includes several ways to draw text directly on the screen. The easiest and most efficient way to draw simple strings is using the UIKit additions to the `NSString` class. These extensions include methods for drawing strings using a variety of attributes wherever you want them on the screen. There are also methods for computing the size of a rendered string before you actually draw it, which can help you lay out your application content more precisely.

Important: Because of the performance implications, you should avoid drawing text directly whenever possible. Static text can be drawn much more efficiently using one or more `UILabel` objects than it can be using a custom drawing routine. Similarly, the `UITextField` class includes different styles that make it easier to integrate editable text areas into your content.

When you need to draw custom text strings in your interface, use the methods of `NSString` to do so. UIKit includes extensions to the basic `NSString` class that allow you to draw strings in your views. These methods allow you to adjust the position of the rendered text precisely and blend it with the rest of your view's content. The methods of this class also let you compute the bounding rectangle for your text in advance based on the desired font and style attributes. For information, see *NSString UIKit Additions Reference*.

If you need more control over the fonts you plan to use during drawing, you can also use the functions of the Core Graphics framework to do your drawing. The Core Graphics framework provides methods for the precise drawing and placement of glyphs and text. For more information about these functions and their use, see *Quartz 2D Programming Guide* and *Core Graphics Framework Reference*.

Files and Networking

An application running in iPhone OS has access to the local file system and network, which you can access using assorted Core OS and Core Services frameworks. Being able to read and write files in the local file system lets you save user data and application state for subsequent uses. Accessing the network gives you the ability to communicate with network servers to perform remote operations and send and retrieve data.

File and Data Management

Files in iPhone OS share space with the user's media and personal files on the flash-based memory. For security purposes, your application is placed in its own directory and is limited to reading and writing files in that directory only. The following sections describe the structure of an application's local file system and several techniques for reading and writing files.

Commonly Used Directories

For security purposes, an application has only a few locations in which it can write its data and preferences. When an application is installed on a device, a home directory is created for the application. Table 6-1 lists some of the important subdirectories inside the home directory that you might need to access. This table describes the intended usage and access restrictions for each directory and whether the directory's contents are backed up by iTunes. For more information about the backup and restore process, see [“Backup and Restore”](#) (page 118). For more information about the application home directory itself, see [“The Application Sandbox”](#) (page 21).

Table 6-1 Directories of an iPhone application

Directory	Description
<code><Application_Home>/AppName.app</code>	<p>This is the bundle directory containing the application itself. Because an application must be signed, you must not make changes to the contents of this directory at runtime. Doing so may prevent your application from launching later.</p> <p>In iPhone OS 2.1 and later, the contents of this directory are not backed up by iTunes. However, iTunes does perform an initial sync of any applications purchased from the App Store.</p>

Directory	Description
<code><Application_Home>/Documents/</code>	<p>This is the directory you should use to write any application-specific data files. Use this directory to store user data or other information that should be backed up regularly. For information about how to get the path of this directory, see “Getting Paths to Application Directories” (page 119).</p> <p>The contents of this directory are backed up by iTunes.</p>
<code><Application_Home>/Library/Preferences</code>	<p>This directory contains application-specific preference files. You should not create preferences files directly but should instead use the <code>NSUserDefaults</code> class or <code>CFPreferences</code> API to get and set application preferences; see “Adding the Settings Bundle” (page 156).</p> <p>The contents of this directory are backed up by iTunes.</p>
<code><Application_Home>/Library/Caches</code>	<p>Use this directory to write any application-specific support files that you want to persist between launches of the application. Your application is generally responsible for adding and removing these files. However, iTunes removes these files during a full restore of the device so you should be able to recreate them as needed. To access this directory, use the interfaces described in “Getting Paths to Application Directories” (page 119) to get the path to the directory.</p> <p>In iPhone OS 2.2 and later, the contents of this directory are not backed up by iTunes.</p>
<code><Application_Home>/tmp/</code>	<p>Use this directory to write temporary files that you do not need to persist between launches of your application. Your application should remove files from this directory when it determines they are no longer needed. (The system may also purge lingering files from this directory when your application is not running.) For information about how to get the path of this directory, see “Getting Paths to Application Directories” (page 119).</p> <p>In iPhone OS 2.1 and later, the contents of this directory are not backed up by iTunes.</p>

Backup and Restore

You do not have to prepare your application in any way for backup and restore operations. In iPhone OS 2.2 and later, when a device is plugged into a computer and synced, iTunes performs an incremental backup of all files, except for those in the following directories:

- `<Application_Home>/AppName.app`
- `<Application_Home>/Library/Caches`
- `<Application_Home>/tmp`

Although iTunes does backup the application bundle itself, it does not do this during every sync operation. Applications purchased from the App Store on the device are backed up when that device is next synced with iTunes. Applications are not backed up during subsequent sync operations though unless the application bundle itself has changed (because the application was updated, for example).

To prevent the syncing process from taking a long time, you should be selective about where you place files inside your application's home directory. The `<Application_Home>/Documents` directory should be used to store user data files or files that cannot be easily recreated by your application. Files used to store temporary data should be placed inside the `Application_Home/tmp` directory and deleted by your application when they are no longer needed. If your application creates data files that can be used during subsequent launches, it should place those files in the `Application_Home/Library/Caches` directory.

Note: If your application creates large data files, or files that change frequently, you should consider storing them in the `Application_Home/Library/Caches` directory and not in the `<Application_Home>/Documents` directory. Backing up large data files can slow down the backup process significantly. The same is true for files that change regularly (and therefore must be backed up frequently). Placing these files in the `Caches` directory prevents them from being backed up (in iPhone OS 2.2 and later) during every sync operation.

For additional guidance about how you should use the directories in your application, see [Table 6-1](#) (page 117).

Getting Paths to Application Directories

At various levels of the system, there are programmatic ways to obtain file-system paths to the directory locations of an application's sandbox. However, the preferred way to retrieve these paths is with the Cocoa programming interfaces. The `NSHomeDirectory` function (in the Foundation framework) returns the path to the top-level home directory—that is, the directory containing the application, `Documents`, `Library`, and `tmp` directories. In addition to that function, you can also use the `NSSearchPathForDirectoriesInDomains` and `NSTemporaryDirectory` functions to get exact paths to your `Documents`, `Caches`, and `tmp` directories.

Both the `NSHomeDirectory` and `NSTemporaryDirectory` functions return the properly formatted path information in an `NSString` object. You can use the path-related methods of `NSString` to modify the path information or create new path strings. For example, upon retrieving the temporary directory path, you could append a file name and use the resulting string to create a file with the given name in the temporary directory.

Note: If you are using frameworks with ANSI C programmatic interfaces—including those that take paths—recall that `NSString` objects are “toll-free bridged” with their Core Foundation counterparts. This means that you can cast a `NSString` object (such as the return by one of the above functions) to a `CFStringRef` type, as shown in the following example:

```
CFStringRef homeDir = (CFStringRef)NSHomeDirectory();
```

For more information on toll-free bridging, see *Carbon-Cocoa Integration Guide*.

The `NSSearchPathForDirectoriesInDomains` function of the Foundation framework lets you obtain the full path to several application-related directories. To use this function in iPhone OS, specify an appropriate search path constant for the first parameter and `NSUserDomainMask` for the second parameter. Table 6-2 lists several of the most commonly used constants and the directories they return.

Table 6-2 Commonly used search path constants

Constant	Directory
<code>NSDocumentDirectory</code>	<code><Application_Home>/Documents</code>
<code>NSCachesDirectory</code>	<code><Application_Home>/Library/Caches</code>
<code>NSApplicationSupportDirectory</code>	<code><Application_Home>/Library/Application Support</code>

Because the `NSSearchPathForDirectoriesInDomains` function was designed originally for Mac OS X, where multiple such directories could exist, it returns an array of paths rather than a single path. In iPhone OS, the resulting array should contain the single path to the desired directory. Listing 6-1 shows a typical use of this function.

Listing 6-1 Getting a file-system path to the application's `Documents/` directory

```
NSArray *paths = NSSearchPathForDirectoriesInDomains(NSDocumentDirectory,
NSUserDomainMask, YES);
NSString *documentsDirectory = [paths objectAtIndex:0];
```

You can call `NSSearchPathForDirectoriesInDomains` using a domain-mask parameter other than `NSUserDomainMask` or a directory constant other than those in Table 6-2 (page 120), but the application will be unable to write to any of the returned directory locations. For example, if you specify `NSApplicationDirectory` as the directory parameter and `NSSystemDomainMask` as the domain-mask parameter, you get the path `/Applications` returned (on the device), but your application cannot write any files to this location.

Another consideration to keep in mind is the difference in directory locations between platforms. The paths returned by `NSSearchPathForDirectoriesInDomains`, `NSHomeDirectory`, `NSTemporaryDirectory`, and similar functions differ depending on whether you're running your application on the device or on the Simulator. For example, take the function call shown in Listing 6-1 (page 120). On the device, the returned path (`documentsDirectory`) is similar to the following:

```
/var/mobile/Applications/30B51836-D2DD-43AA-BCB4-9D4DADFED6A2/Documents
```

However, on the Simulator, the returned path takes the following form:

```
/Volumes/Stuff/Users/johnDoe/Library/Application Support/iPhone
Simulator/User/Applications/118086A0-FAAF-4CD4-9A0F-CD5E8D287270/Documents
```

To read and write user preferences, use the `NSUserDefaults` class or the `CFPreferences` API. These interfaces eliminate the need for you to construct a path to the `Library/Preferences/` directory and read and write preference files directly. For more information on using these interfaces, see [“Adding the Settings Bundle”](#) (page 156).

If your application contains sound, image, or other resources in the application bundle, you should use the `NSBundle` class or `CFBundle` opaque type to load those resources. Bundles have an inherent knowledge of where resources live inside the application. In addition, bundles are aware of the user's language preferences and are able to choose localized resources over default resources automatically. For more information on bundles, see [“The Application Bundle”](#) (page 23).

Reading and Writing File Data

The iPhone OS provides several ways to read, write, and manage files.

■ Foundation framework:

- ❑ If you can represent your application's data as a property list, convert the property list to an `NSData` object using the `NSPropertyListSerialization` API. You can then write the data object to disk using the methods of the `NSData` class.
- ❑ If your application's model objects adopt the `NSCoding` protocol, you can archive a graph of these model objects using the `NSKeyedArchiver` class, and especially its `archivedDataWithRootObject:` method.
- ❑ The `NSFileHandle` class in Foundation framework provides random access to the contents of a file.
- ❑ The `NSFileManager` class in Foundation framework provides methods to create and manipulate files in the file system.

■ Core OS calls:

- ❑ Calls such as `fopen`, `fread`, and `fwrite` also let you read and write file data either sequentially or via random access.
- ❑ The `mmap` and `munmap` calls provide an efficient way to load large files into memory and access their contents.

Note: The preceding list of Core OS calls is just a sample of the more commonly used calls. For a more complete list of the available functions, see the list of functions in section 3 of *iPhone OS Manual Pages*.

The following sections show examples of how to use some of the higher-level techniques for reading and writing files. For additional information about the file-related classes of the Foundation framework, see *Foundation Framework Reference*.

Reading and Writing Property List Data

A property list is a form of data representation that encapsulates several Foundation (and Core Foundation) data types, including dictionaries, arrays, strings, dates, binary data, and numerical and Boolean values. Property lists are commonly used to store structured configuration data. For example, the `Info.plist` file found in every Cocoa and iPhone applications is a property list that stores configuration information about the application itself. You can use property lists yourself to store additional information, such as the state of your application when it quits.

In code, you typically construct a property list starting with either a dictionary or array as a container object. You then add other property-list objects, including (possibly) other dictionaries and arrays. The keys of dictionaries must be string objects. The values for those keys are instances of `NSDictionary`, `NSArray`, `NSString`, `NSDate`, `NSData`, and `NSNumber`.

For an application whose data can be represented by a property-list object (such as an `NSDictionary` object), you could write that property list to disk using a method such as the one shown in Listing 6-2. This method serializes the property list object into an `NSData` object, then calls the `writeApplicationData:toFile:` method (the implementation for which is shown in Listing 6-4 (page 123)) to write that data to disk.

Listing 6-2 Converting a property-list object to an `NSData` object and writing it to storage

```
- (BOOL)writeApplicationPlist:(id)plist toFile:(NSString *)fileName {
    NSString *error;
    NSData *pData = [NSPropertyListSerialization dataFromPropertyList:plist
format:NSPropertyListBinaryFormat_v1_0 errorDescription:&error];
    if (!pData) {
        NSLog(@"%@", error);
        return NO;
    }
    return ([self writeApplicationData:pData toFile:(NSString *)fileName]);
}
```

When writing property list files in iPhone OS, it is important to store your files in binary format. You do this by specifying the `NSPropertyListBinaryFormat_v1_0` key in the *format* parameter of the `dataFromPropertyList:format:errorDescription:` method. The binary property-list format is much more compact than the other format options, which are text based. This compactness not only minimizes the amount of space taken up on the user's device, it also improves the time it takes to read and write the property list.

Listing 6-3 shows the corresponding code for loading a property-list file from disk and reconstituting the objects in that property list.

Listing 6-3 Reading a property-list object from the application's Documents directory

```
- (id)applicationPlistFromFile:(NSString *)fileName {
    NSData *retData;
    NSString *error;
    id retPlist;
    NSPropertyListFormat format;

    retData = [self applicationDataFromFile:fileName];
    if (!retData) {
        NSLog(@"Data file not returned.");
        return nil;
    }
    retPlist = [NSPropertyListSerialization propertyListFromData:retData
mutabilityOption:NSPropertyListImmutable format:&format errorDescription:&error];
    if (!retPlist){
        NSLog(@"Plist not returned, error: %@", error);
    }
    return retPlist;
}
```

For more on property lists and the `NSPropertyListSerialization` class, see *Property List Programming Guide*.

Using Archivers to Read and Write Data

An archiver converts an arbitrary collection of objects into a stream of bytes. Although this may sound similar to the process employed by the `NSPropertyListSerialization` class, there is an important difference. A property-list serializer can convert only a limited set of (mostly scalar) data types. Archivers can convert arbitrary Objective-C objects, scalar types, arrays, structures, strings, and more.

The key to the archiving process is in the target objects themselves. The objects manipulated by an archiver must conform to the `NSCoding` protocol, which defines the interface for reading and writing the object's state. When an archiver encodes a set of objects, it sends an `encodeWithCoder:` message to each one, which the object then uses to write out its critical state information to the corresponding archive. The unarchiving process reverses the flow of information. During unarchiving, each object receives an `initWithCoder:` message, which it uses to initialize itself with the state information currently in the archive. Upon completion of the unarchiving process, the stream of bytes is reconstituted into a new set of objects that have the same state as the ones written to the archive previously.

The Foundation framework supports two kinds of archivers—sequential and keyed. Keyed archivers are more flexible and are recommended for use in your application. The following example shows how to archive a graph of objects using a keyed archiver. The `representation` method of the `_myDataSource` object returns a single object (possibly an array or dictionary) that points to all of the objects to be included in the archive. The data object is then written to a file whose path is specified by the `myFilePath` variable.

```
NSData *data = [NSKeyedArchiver archivedDataWithRootObject:[_myDataSource
representation]];
[data writeToFile:myFilePath atomically:YES];
```

Note: You could also send a `archiveRootObject:toFile:` message to the `NSKeyedArchiver` object to create the archive and write it to storage in one step.

To load the contents of an archive from disk, you simply reverse the process. After loading the data from disk, you use the `NSKeyedUnarchiver` class and its `unarchiveObjectWithData:` class method to get back the model-object graph. For example, to unarchive the data from the previous example, you could use the following code:

```
NSData* data = [NSData dataWithContentsOfFile:myFilePath];
id rootObject = [NSKeyedUnarchiver unarchiveObjectWithData:data];
```

For more information on how to use archivers and how to make your objects support the `NSCoding` protocol, see *Archives and Serializations Programming Guide for Cocoa*.

Writing Data to Your Documents Directory

After you have an `NSData` object encapsulating the application data (either as an archive or a serialized property list), you can call the method shown in Listing 6-4 to write that data to the application Documents directory.

Listing 6-4 Writing data to the application's Documents directory

```
- (BOOL)writeApplicationData:(NSData *)data toFile:(NSString *)fileName {
    NSArray *paths = NSSearchPathForDirectoriesInDomains(NSDocumentDirectory,
    NSUserDomainMask, YES);
    NSString *documentsDirectory = [paths objectAtIndex:0];
    if (!documentsDirectory) {
```

```

        NSLog(@"Documents directory not found!");
        return NO;
    }
    NSString *appFile = [documentsDirectory
stringByAppendingPathComponent:fileName];
    return ([data writeToFile:appFile atomically:YES]);
}

```

Reading Data from the Documents Directory

To read a file from your application's Documents directory, construct the path for the file name and use the desired method to read the file contents into memory. For relatively small files—that is, files less than a few memory pages in size—you could use the code in Listing 6-5 to obtain a data object for the file contents. This example constructs a full path to the file in the Documents directory, creates a data object from it, and then returns that object.

Listing 6-5 Reading data from the application's Documents directory

```

- (NSData *)applicationDataFromFile:(NSString *)fileName {
    NSArray *paths = NSSearchPathForDirectoriesInDomains(NSDocumentDirectory,
NSUserDomainMask, YES);
    NSString *documentsDirectory = [paths objectAtIndex:0];
    NSString *appFile = [documentsDirectory
stringByAppendingPathComponent:fileName];
    NSData *myData = [[[NSData alloc] initWithContentsOfFile:appFile]
autorelease];
    return myData;
}

```

For files that would require multiple memory pages to hold in memory, you should avoid loading the entire file all at once. This is especially important if you plan to use only part of the file. For larger files, you should consider mapping the file into memory using either the `mmap` function or the `initWithContentsOfFileMappedFile:` method of `NSData`.

Choosing when to map files versus load them directly is up to you. It is relatively safe to load a file entirely into memory if it requires only a few (3-4) memory pages. If your file requires several dozen or a hundred pages, however, you would probably find it more efficient to map the file into memory. As with any such determination, though, you should measure your application's performance and determine how long it takes to load the file and allocate the necessary memory.

File Access Guidelines

When creating files or writing out file data, keep the following guidelines in mind:

- Minimize the amount of data you write to the disk. File operations are relatively slow and involve writing to the Flash disk, which has a limited lifespan. Some specific tips to help you minimize file-related operations include:
 - ❑ Write only the portions of the file that changed, but aggregate changes when you can. Avoid writing out the entire file just to change a few bytes.
 - ❑ When defining your file format, group frequently modified content together so as to minimize the overall number of blocks that need to be written to disk each time.

- ❑ If your data consists of structured content that is randomly accessed, store it in a SQLite database. This is especially important if the amount of data you are manipulating could grow to be more than a few megabytes in size.
- Avoid writing cache files to disk. The only exception to this rule is when your application quits and you need to write state information that can be used to put your application back into the same state when it is next launched.

Saving State Information

When the user presses the Home button, iPhone OS quits your application and returns to the Home screen. Similarly, if your application opens a URI whose scheme is handled by a different application, iPhone OS quits your application and opens the URI in the other application. In other words, any action that would cause your application to suspend or go to the background in Mac OS X causes your application to quit in iPhone OS. Because these actions happen regularly on mobile devices, your application must change the way it manages volatile data and application state.

Unlike most desktop applications, where the user manually chooses when to save files to disk, your application should save changes automatically at key points in your workflow. Exactly when you save data is up to you, but there are two potential options. Either you can save each change immediately as the user makes it, or you can batch changes on the same page together and save them when the page is dismissed, a new page is displayed, or the application quits. Under no circumstances should you let the user navigate to a new page of content without saving the data on the previous page.

When your application is asked to quit, you should save the current state of your application to a temporary cache file or to the preferences database. The next time the user launches your application, use that information to restore your application to its previous state. The state information you save should be as minimal as possible but still let you accurately restore your application to an appropriate point. You do not have to display the exact same screen the user was manipulating previously if doing so would not make sense. For example, if a user edits a contact and then leaves the Phone application, upon returning, the Phone application displays the top-level list of contacts, rather than the editing screen for the contact.

Case Sensitivity

The file system for iPhone OS–based devices is case sensitive. Whenever you work with filenames, you should be sure that the case matches exactly or your code may be unable to open or access the file.

Networking

The networking stack in iPhone OS includes several interfaces over the radio hardware of iPhone and iPod touch devices. The main programming interface is the CFNetwork framework, which builds on top of BSD sockets and opaque types in the Core Foundation framework to communicate with network entities. You can also use the `NSStream` classes in the Foundation framework and the low-level BSD sockets found in the Core OS layer of the system.

The following sections provide iPhone-specific tips for developers who need to incorporate networking features into their applications. For information about how to use the `CFNetwork` framework for network communication, see *CFNetwork Programming Guide* and *CFNetwork Framework Reference*. For information about using the `NSStream` class, see *Foundation Framework Reference*.

Tips for Efficient Networking

When implementing code to receive or transmit across the network, remember that doing so is one of the most power-intensive operations on a device. Minimizing the amount of time spent transmitting or receiving helps improve battery life. To that end, you should consider the following tips when writing your network-related code:

- For protocols you control, define your data formats to be as compact as possible.
- Avoid communicating using chatty protocols.
- Transmit data packets in bursts whenever you can.

The cellular and Wi-Fi radios are designed to power down when there is no activity. Doing so can take several seconds though, depending on the radio. If your application transmits small bursts of data every few seconds, the radios may stay powered up and continue to consume power, even when they are not actually doing anything. Rather than transmit small amounts of data more often, it is better to transmit a larger amount of data once or at relatively large intervals.

When communicating over the network, it is also important to remember that packets can be lost at any time. When writing your networking code, you should be sure to make it as robust as possible when it comes to failure handling. It is perfectly reasonable to implement handlers that respond to changes in network conditions, but do not be surprised if those handlers are not called consistently. For example, the Bonjour networking callbacks may not always be called immediately in response to the disappearance of a network service. The Bonjour system service does immediately invoke browsing callbacks when it receives a notification that a service is going away, but network services can disappear without notification. This might occur if the device providing the network service unexpectedly loses network connectivity or the notification is lost in transit.

Using Wi-Fi

If your application accesses the network using the Wi-Fi radios, you must notify the system of that fact by including the `UIRequiresPersistentWiFi` key in the application's `Info.plist` file. The inclusion of this key lets the system know that it should display the network selection panel if it detects any active Wi-Fi hot spots. It also lets the system know that it should not attempt to shut down the Wi-Fi hardware while your application is running.

To prevent the Wi-Fi hardware from using too much power, iPhone OS has a built-in timer that turns off the hardware completely after 30 minutes if no application has requested its use through the `UIRequiresPersistentWiFi` key. If the user launches an application that includes the key, iPhone OS effectively disables the timer for the duration of the application's life cycle. As soon as that application quits, however, the system reenables the timer.

For more information on the `UIRequiresPersistentWiFi` key and the keys of the `Info.plist` file, see [“The Information Property List”](#) (page 25).

Multimedia Support

iPhone OS supports audio features in your application through the Core Audio and OpenAL frameworks, and provides video playback support using the Media Player framework. Core Audio provides an advanced interface for playing, recording, and manipulating sound and for parsing streamed audio. If you are a game developer and already have code that takes advantage of OpenAL, you can use your code in iPhone OS.

Using Sound in iPhone OS

Core Audio offers a rich set of tools for working with sound in your application. These tools are arranged into four frameworks:

- The Audio Toolbox framework, for recording and playing audio, parsing audio streams, and converting audio formats.
- The Audio Unit framework, for using audio processing plug-ins.
- The AV Foundation framework, for playing audio using a simple Objective-C interface.
- The Core Audio framework (not an umbrella framework), which provides data types used by all Core Audio services.

In addition to Core Audio, you can use the OpenAL framework in iPhone OS for positional audio playback. OpenAL 1.1 support in iPhone OS is built on top of Core Audio.

You select among the various iPhone audio APIs based on the needs of your application. The following list lets you zoom in on the audio technologies to use:

- To play alerts and user-interface sound effects, or to invoke vibration on devices that provide that feature, use System Sound Services. See [“Playing Short Sounds or Invoking Vibration Using System Sound Services”](#) (page 131).
- To play sounds in the fewest lines of code, use the `AVAudioPlayer` class. See [“Playing Sounds Easily with the AVAudioPlayer Class”](#) (page 133).
- To provide full-featured audio playback including stereo positioning, level control, and simultaneous sounds, use OpenAL. See [“Playing Sounds with Positioning Using OpenAL”](#) (page 137).
- To provide lowest latency audio, especially when doing simultaneous input and output (such as for a VOIP application), use the I/O audio unit. See [“Audio Unit Support in iPhone OS”](#) (page 139).
- To play sounds with the highest degree of control, including support for synchronization, use Audio Queue Services. See [“Playing Sounds with Control Using Audio Queue Services”](#) (page 134).
- To record audio, use Audio Queue Services. See [“Recording Audio”](#) (page 137).
- To parse audio streamed from a network connection, use Audio File Stream Services. See [“Parsing Streamed Audio”](#) (page 138).

Be sure to read the next section, “[The Basics: Hardware Codecs, Audio Formats, and Audio Sessions](#)” (page 128), for critical information on how audio works on an iPhone OS-based device. Also read “[Best Practices for iPhone Audio](#)” (page 139), which offers guidelines and lists the audio and file formats to use for best performance and best user experience.

When you’re ready to dig deeper, the iPhone Dev Center contains guides, reference books, sample code, and more. For tips on how to perform common audio tasks, see *Audio & Video Coding How-To’s*. For in-depth explanations of audio development in iPhone OS, see *Core Audio Overview*, *Audio Queue Services Programming Guide*, and *Audio Session Programming Guide*.

The Basics: Hardware Codecs, Audio Formats, and Audio Sessions

To get oriented toward iPhone audio development, it’s helpful to understand a few things about the hardware and software architecture of iPhone OS-based devices.

iPhone Audio Hardware Codecs

iPhone OS applications can use a wide range of audio data formats, as described in the next section. Some of these formats use software-based encoding and decoding. You can simultaneously play multiple sounds in these formats. Moreover, your application and a background application (notably, the iPod application) can simultaneously play sounds in these formats.

Other iPhone OS audio formats employ a hardware codec for playback. These formats are:

- AAC
- ALAC (Apple Lossless)
- MP3

The device can play only a single instance of one of these formats at a time. For example, if you are playing a stereo MP3 sound, you cannot simultaneously play another, nor can you simultaneously play an AAC or an ALAC sound. If the iPod application is playing an MP3 sound in the background, your application is unable to play AAC, ALAC, or MP3 audio.

To play multiple, high-quality sounds, or to play sounds while the iPod is playing in the background, use linear PCM (uncompressed) or IMA4 (compressed) audio.

Audio Playback and Recording Formats

Here are the audio playback formats supported in iPhone OS:

- AAC
- AMR (Adaptive Multi-Rate, a format for speech)
- ALAC (Apple Lossless)
- iLBC (internet Low Bitrate Codec, another format for speech)
- IMA4 (IMA/ADPCM)
- linear PCM (uncompressed)
- μ -law and a-law

- MP3 (MPEG-1 audio layer 3)

Here are the audio recording formats supported in iPhone OS:

- ALAC (Apple Lossless)
- iLBC (internet Low Bitrate Codec, for speech)
- IMA/ADPCM (IMA4)
- linear PCM
- μ -law and a-law

The following list summarizes how iPhone OS supports audio formats for single or multiple playback:

- **Linear PCM and IMA4 (IMA/ADPCM)** You can play multiple linear PCM or IMA4 format sounds simultaneously in iPhone OS without incurring CPU resource problems. The same is true for the AMR and iLBC speech-quality formats, and for the μ -law and a-law compressed formats.
- **AAC, MP3, and ALAC (Apple Lossless)** Playback for AAC, MP3, and ALAC sounds uses efficient hardware-based decoding on iPhone OS–based devices, but these codecs all share a single hardware path. The device can play only a single instance of one of these formats at a time.

The single hardware path for AAC, MP3, and ALAC audio entails implications for "play along" style applications, such as a virtual piano. If the user is playing a sound in one of these three formats in the iPod application, then your application—to play along over that audio—must use one of the software-decoded formats: linear PCM, IMA4, AMR, iLBC, μ -law, or a-law.

Audio Sessions

Core Audio's audio session interface (declared in `AudioToolbox/AudioServices.h`) lets your application define an audio context for itself and to work well within the larger audio context of the device it's running on.

Your context includes such things as whether you want your audio to be silenced by the Ring/Silent switch and whether you want iPod audio to continue playing when your audio starts. The larger audio context includes changes made by users—such as when they plug in headsets. It also includes events such as incoming phone calls.

Each iPhone OS application gets a singleton **audio session object**, usually called an *audio session*. This object comes with some default behavior that you can use to get started in development. Except for certain special cases, however, the default behavior is unsuitable for a shipping application. By configuring and using the audio session, you can express your audio intentions and respond to OS-level audio decisions.

Audio Session Services provides three programmatic features, described in Table 7-1.

Table 7-1 Features provided by the audio session interface

Audio session feature	Description
Categories	A category is a key that identifies a set of audio behaviors for your application. By setting a category, you indicate your audio intentions to iPhone OS, such as whether your audio should continue when the screen locks.

Audio session feature	Description
Interruptions and route changes	Your audio session posts notifications when your audio is interrupted, when an interruption ends, and when the hardware audio route changes. These notifications let you respond to changes in the larger audio environment—such as an interruption due to an incoming phone call—gracefully.
Hardware characteristics	You can query the audio session to discover characteristics of the device your application is running on, such as hardware sample rate, number of hardware channels, and whether audio input is available.

When using the default audio session, audio in your application stops when the Auto-Lock period times out and the screen locks. If you want to ensure that playback continues with the screen locked, you perform the following steps:

1. In your application's initialization code, initialize your audio session as shown in Listing 7-1. Your audio session must be initialized before you can assign a category to it, which you do in the next step.

Listing 7-1 Initializing an audio session

```
AudioSessionInitialize (
    NULL,                // uses the default run loop
    NULL,                // uses the default run loop mode
    interruptionCallback, // your callback function (can be set to NULL)
    userData              // data you want sent to your callback
);
```

(This function uses the *interruptionCallback* parameter to register your interruption listener callback function. You can (at your option) provide that callback to specify how your application behaves when an audio interruption starts or ends. That callback is not involved with setting the category for your audio session, and does not affect screen-locking behavior.)

2. Before starting playback, you set your audio session's category to `kAudioSessionCategory_MediaPlayback`, as shown in Listing 7-2. Among other things, this category indicates that you intend for your audio to continue on screen lock or when the Ring/Silent switch is set to silent.

Listing 7-2 Setting an audio session category

```
UInt32 sessionCategory = kAudioSessionCategory_MediaPlayback;

AudioSessionSetProperty (
    kAudioSessionProperty_AudioCategory,
    sizeof (sessionCategory),
    &sessionCategory
);
```

For detailed information on categories, see *Audio Session Categories* in *Audio Session Programming Guide*.

How you handle the interruption caused by an incoming phone call depends on the audio technology you are using in your application, as shown in Table 7-2.

Table 7-2 Handling audio interruptions

Audio technology	How interruptions work
System Sound Services	System sounds and alerts go silent when an interruption starts. They can automatically be used again if the interruption ends—which happens if a user elects to ignore an incoming phone call. Applications cannot influence the interruption behavior for sounds that use this technology.
Audio Queue Services, OpenAL, I/O audio unit	These technologies provide the greatest flexibility, putting you in control of handling interruptions. You write an interruption listener callback function, as explained in “Responding to Audio Interruptions” in <i>Audio Session Programming Guide</i> .
AVAudioPlayer class	The AVAudioPlayer class provides delegate methods for interruption start and end. You can implement the <code>audioPlayerBeginInterruption:</code> method, as appropriate, to update your user interface. The audio player object takes care of pausing playback and saving state. You can use the <code>audioPlayerEndInterruption:</code> method to resume playback and, as appropriate, update your user interface. The audio player takes care of reactivating your audio session.

Every iPhone OS application—with rare exception—should adopt Audio Session Services. To learn how, read *Audio Session Programming Guide*.

Playing Audio

This section introduces you to playing sounds in iPhone OS using System Sound Services, Audio Queue Services, the AVAudioPlayer class, and OpenAL.

Playing Short Sounds or Invoking Vibration Using System Sound Services

To play user-interface sound effects (such as button clicks) or alert sounds, or to invoke vibration on devices that support it, use System Sound Services. This compact interface is described in *System Sound Services Reference*. You can find sample code in the *SysSound* sample in the [iPhone Dev Center](#).

Note: Sounds played with System Sound Services are not subject to configuration using your audio session. As a result, you cannot keep the behavior of System Sound Services audio in line with other audio behavior in your application. This is the most important reason to avoid using System Sound Services for any audio apart from its intended uses.

The `AudioServicesPlaySystemSound` function lets you very simply play short sound files. The simplicity carries with it a few restrictions. Your sound files must be:

- Shorter than 30 seconds in duration
- In linear PCM or IMA4 (IMA/ADPCM) format
- Packaged in a `.caf`, `.aif`, or `.wav` file

In addition, when you use the `AudioServicesPlaySystemSound` function:

- Sounds play at the current system audio level, with no level control available
- Sounds play immediately
- Looping and stereo positioning are unavailable

The similar `AudioServicesPlayAlertSound` function plays a short sound as an alert. If a user has configured their device to vibrate in Ring Settings, calling this function invokes vibration in addition to playing the sound file.

Note: System-supplied alert sounds and system-supplied user-interface sound effects are not available to your application. For example, using the `kSystemSoundID_UserPreferredAlert` constant as a parameter to the `AudioServicesPlayAlertSound` function will not play anything.

To play a sound with the `AudioServicesPlaySystemSound` or `AudioServicesPlayAlertSound` function, you first create a sound ID object, as shown in Listing 7-3.

Listing 7-3 Creating a sound ID object

```
// Get the main bundle for the app
CFBundleRef mainBundle = CFBundleGetMainBundle ();

// Get the URL to the sound file to play. The file in this case
// is "tap.aiff"
soundFileURLRef = CFBundleCopyResourceURL (
    mainBundle,
    CFSTR ("tap"),
    CFSTR ("aif"),
    NULL
);

// Create a system sound object representing the sound file
AudioServicesCreateSystemSoundID (
    soundFileURLRef,
    &soundFileObject
);
```

You can then play the sound, as shown in Listing 7-4.

Listing 7-4 Playing a system sound

```
- (IBAction) playSystemSound {
    AudioServicesPlaySystemSound (self.soundFileObject);
}
```

In typical use, which includes playing a sound occasionally or repeatedly, retain the sound ID object until your application quits. If you know that you will use a sound only once—for example, in the case of a startup sound—you can destroy the sound ID object immediately after playing the sound, freeing memory.

Applications running on iPhone OS–based devices that support vibration can trigger that feature using System Sound Services. You specify the vibrate option with the `kSystemSoundID_Vibrate` identifier. To trigger it, use the `AudioServicesPlaySystemSound` function, as shown in Listing 7-5.

Listing 7-5 Triggering vibration

```
#import <AudioToolbox/AudioToolbox.h>
#import <UIKit/UIKit.h>
- (void) vibratePhone {
    AudioServicesPlaySystemSound (kSystemSoundID_Vibrate);
}
```

If your application is running on an iPod touch, this code does nothing.

Playing Sounds Easily with the AVAudioPlayer Class

The `AVAudioPlayer` class provides a simple Objective-C interface for playing sounds. If your application does not require stereo positioning or precise synchronization, and if you are not playing audio captured from a network stream, Apple recommends that you use this class for playback.

Using an audio player you can:

- Play sounds of any duration
- Play sounds from files or memory buffers
- Loop sounds
- Play multiple sounds simultaneously
- Control relative playback level for each sound you are playing
- Seek to a particular point in a sound file, which supports such application features as fast forward and rewind
- Obtain audio power data that you can use for audio level metering

The `AVAudioPlayer` class lets you play sound in any audio format available in iPhone OS, as described in [“Audio Playback and Recording Formats”](#) (page 128). For a complete description of this class’s interface, see *AVAudioPlayer Class Reference*.

To configure an audio player for playback, you assign a sound file to it, prepare it to play, and designate a delegate object. The code in Listing 7-6 would typically go into an initialization method of the controller class for your application.

Listing 7-6 Configuring an AVAudioPlayer object

```
NSString *soundFilePath =
    [[NSBundle mainBundle] pathForResource:@"sound"
                                         ofType:@"wav"];

NSURL *fileURL = [[NSURL alloc] initWithURLWithPath: soundFilePath];

AVAudioPlayer *newPlayer =
    [[AVAudioPlayer alloc] initWithContentsOfURL: fileURL
                                             error: nil];

[fileURL release];

self.player = newPlayer;
[newPlayer release];

[self.player prepareToPlay];
```

```
[self.player setDelegate: self];
```

You use a delegate object (which can be your controller object) to handle interruptions and to update the user interface when a sound has finished playing. The delegate methods for the `AVAudioPlayer` class are described in *AVAudioPlayerDelegate Protocol Reference*. Listing 7-7 shows a simple implementation of one delegate method. This code updates the title of a Play/Pause toggle button when a sound has finished playing.

Listing 7-7 Implementing an `AVAudioPlayer` delegate method

```
- (void) audioPlayerDidFinishPlaying: (AVAudioPlayer *) player
      successfully: (BOOL) flag {
    if (flag == YES) {
        [self.button setTitle: @"Play" forState: UIControlStateNormal];
    }
}
```

To play, pause, or stop an `AVAudioPlayer` object, call one of its playback control methods. You can test whether or not playback is in progress by using the `playing` property. Listing 7-8 shows a basic play/pause toggle method that controls playback and updates the title of a `UIButton` object.

Listing 7-8 Controlling an `AVAudioPlayer` object

```
- (IBAction) playOrPause: (id) sender {

    // if already playing, then pause
    if (self.player.playing) {
        [self.button setTitle: @"Play" forState: UIControlStateHighlighted];
        [self.button setTitle: @"Play" forState: UIControlStateNormal];
        [self.player pause];

        // if stopped or paused, start playing
    } else {
        [self.button setTitle: @"Pause" forState: UIControlStateHighlighted];
        [self.button setTitle: @"Pause" forState: UIControlStateNormal];
        [self.player play];
    }
}
```

The `AVAudioPlayer` class uses the Objective-C declared properties feature for managing information about a sound—such as the playback point within the sound’s timeline, and for accessing playback options—such as volume and looping. For example, you set the playback volume for an audio player as shown here:

```
[self.player setVolume: 1.0];    // available range is 0.0 through 1.0
```

For more information on the `AVAudioPlayer` class, see *AVAudioPlayer Class Reference*.

Playing Sounds with Control Using Audio Queue Services

Audio Queue Services adds playback capabilities beyond those available with the `AVAudioPlayer` class. Using Audio Queue Services for playback lets you:

- Precisely schedule when a sound plays, allowing synchronization
- Precisely control volume—on a buffer-by-buffer basis
- Play audio that you have captured from a stream using Audio File Stream Services

Audio Queue Services lets you play sound in any audio format available in iPhone OS, as described in “[Audio Playback and Recording Formats](#)” (page 128). You also use this technology for recording, as explained in “[Recording Audio](#)” (page 137).

For detailed information on using this technology, see *Audio Queue Services Programming Guide* and *Audio Queue Services Reference*. For sample code, see the *SpeakHere* sample in the [iPhone Dev Center](#). (For a Mac OS X implementation, see the AudioQueueTools project available in the Core Audio SDK. When you install Xcode tools in Mac OS X, the AudioQueueTools project is available at `/Developer/Examples/CoreAudio/SimpleSDK/AudioQueueTools`.)

Creating an Audio Queue Object

To create an audio queue object for playback, perform these three steps:

1. Create a data structure to manage information needed by the audio queue, such as the audio format for the data you want to play.
2. Define a callback function for managing audio queue buffers. The callback uses Audio File Services to read the file you want to play. (In iPhone OS 2.1 and later, you can also use Extended Audio File Services to read the file.)
3. Instantiate the playback audio queue using the `AudioQueueNewOutput` function.

Listing 7-9 illustrates these steps using ANSI C. The *SpeakHere* sample project shows these same steps in the context of an Objective-C program.

Listing 7-9 Creating an audio queue object

```
static const int kNumberBuffers = 3;
// Create a data structure to manage information needed by the audio queue
struct myAQStruct {
    AudioFileID                mAudioFile;
    CAStreamBasicDescription    mDataFormat;
    AudioQueueRef              mQueue;
    AudioQueueBufferRef         mBuffers[kNumberBuffers];
    SInt64                     mCurrentPacket;
    UInt32                     mNumPacketsToRead;
    AudioStreamPacketDescription *mPacketDescs;
    bool                       mDone;
};
// Define a playback audio queue callback function
static void AQTestBufferCallback(
    void                *inUserData,
    AudioQueueRef       inAQ,
    AudioQueueBufferRef inCompleteAQBuffer
) {
    myAQStruct *myInfo = (myAQStruct *)inUserData;
    if (myInfo->mDone) return;
    UInt32 numBytes;
    UInt32 nPackets = myInfo->mNumPacketsToRead;

    AudioFileReadPackets (
        myInfo->mAudioFile,
        false,
        &numBytes,
```

```

        myInfo->mPacketDescs,
        myInfo->mCurrentPacket,
        &nPackets,
        inCompleteAQBuffer->mAudioData
    );
    if (nPackets > 0) {
        inCompleteAQBuffer->mAudioDataByteSize = numBytes;
        AudioQueueEnqueueBuffer (
            inAQ,
            inCompleteAQBuffer,
            (myInfo->mPacketDescs ? nPackets : 0),
            myInfo->mPacketDescs
        );
        myInfo->mCurrentPacket += nPackets;
    } else {
        AudioQueueStop (
            myInfo->mQueue,
            false
        );
        myInfo->mDone = true;
    }
}
// Instantiate an audio queue object
AudioQueueNewOutput (
    &myInfo.mDataFormat,
    AQTestBufferCallback,
    &myInfo,
    CFRRunLoopGetCurrent(),
    kCFRunLoopCommonModes,
    0,
    &myInfo.mQueue
);

```

Controlling the Playback Level

Audio queue objects give you two ways to control playback level.

To set playback level directly, use the `AudioQueueSetParameter` function with the `kAudioQueueParam_Volume` parameter, as shown in Listing 7-10. Level change takes effect immediately.

Listing 7-10 Setting the playback level directly

```

Float32 volume = 1;    // linear scale, range from 0.0 through 1.0
AudioQueueSetParameter (
    myAQstruct.audioQueueObject,
    kAudioQueueParam_Volume,
    volume
);

```

You can also set playback level for an audio queue buffer by using the `AudioQueueEnqueueBufferWithParameters` function. This lets you assign audio queue settings that are, in effect, carried by an audio queue buffer as you enqueue it. Such changes take effect when the buffer begins playing.

In both cases, level changes for an audio queue remain in effect until you change them again.

Indicating Playback Level

You can obtain the current playback level from an audio queue object by:

1. Enabling metering for the audio queue object by setting its `kAudioQueueProperty_EnableLevelMetering` property to `true`
2. Querying the audio queue object's `kAudioQueueProperty_CurrentLevelMeter` property

The value of this property is an array of `AudioQueueLevelMeterState` structures, one per channel. Listing 7-11 shows this structure:

Listing 7-11 The `AudioQueueLevelMeterState` structure

```
typedef struct AudioQueueLevelMeterState {  
    Float32    mAveragePower;  
    Float32    mPeakPower;  
}; AudioQueueLevelMeterState;
```

Playing Multiple Sounds Simultaneously

To play multiple sounds simultaneously, create one playback audio queue object for each sound. For each audio queue, schedule the first buffer of audio to start at the same time using the `AudioQueueEnqueueBufferWithParameters` function.

Audio format is critical when you play sounds simultaneously on an iPhone OS–based device. To play simultaneous sounds, you use the linear PCM (uncompressed) audio format or certain compressed audio formats, as described in “[Audio Playback and Recording Formats](#)” (page 128).

Playing Sounds with Positioning Using OpenAL

The open-sourced OpenAL audio API, available in iPhone OS in the OpenAL framework, provides an interface optimized for positioning sounds in a stereo field during playback. Playing, positioning, and moving sounds is simple when you use OpenAL—working the same way as it does on other platforms. OpenAL also lets you mix sounds. OpenAL uses Core Audio’s I/O unit for playback, resulting in the lowest latency.

For all of the reasons, OpenAL is your best choice for playing sound effects in game applications on iPhone OS–based devices. However, OpenAL is also a good choice for general iPhone OS application audio playback needs.

OpenAL 1.1 support in iPhone OS is built on top of Core Audio. For more information, see *OpenAL FAQ for iPhone OS*. For OpenAL documentation, see the OpenAL website at <http://openal.org>. For sample code showing you how to play OpenAL audio, see *oalTouch*.

Recording Audio

Core Audio provides support in iPhone OS for recording audio using Audio Queue Services. This interface does the work of connecting to the audio hardware, managing memory, and employing codecs as needed. You can record audio in any of the formats listed in “[Preferred Audio Formats in iPhone OS](#)” (page 140).

To record audio, your application configures the audio session, instantiates a recording audio queue object, and provides a callback function. The callback stores the audio data in memory for immediate use or writes it to a file for long-term storage.

Recording takes place at a system-defined level in iPhone OS. The system takes input from the audio source that the user has chosen—the built-in microphone or, if connected, the headset microphone or other input source.

Just as with playback, you can obtain the current recording audio level from an audio queue object by querying its `kAudioQueueProperty_CurrentLevelMeter` property, as described in [“Indicating Playback Level”](#) (page 137).

For detailed examples of how to use Audio Queue Services to record audio, see Recording Audio in *Audio Queue Services Programming Guide*. For sample code, see the *SpeakHere* sample in the [iPhone Dev Center](#).

Parsing Streamed Audio

To play streamed audio content, such as from a network connection, use Audio File Stream Services in concert with Audio Queue Services. Audio File Stream Services parses audio packets and metadata from common audio file container formats in a network bitstream. You can also use it to parse packets and metadata from on-disk files.

In iPhone OS, you can parse the same audio file and bitstream formats that you can in Mac OS X, as follows:

- MPEG-1 Audio Layer 3, used for .mp3 files
- MPEG-2 ADTS, used for the .aac audio data format
- AIFC
- AIFF
- CAF
- MPEG-4, used for .m4a, .mp4, and .3gp files
- NeXT
- WAVE

Having retrieved audio packets, you can play back in any of the formats supported in iPhone OS, as listed in [“Audio Playback and Recording Formats”](#) (page 128).

For best performance, network streaming applications should use data from Wi-Fi connections only. iPhone OS lets you determine which networks are reachable and available through its System Configuration framework and its `SCNetworkReachability.h` interfaces. For sample code, see the *Reachability* sample in the [iPhone Dev Center](#).

To connect to a network stream, you can use interfaces from Core Foundation in iPhone OS, such as the `CFHTTPMessage` interface, described in *CFHTTPMessage Reference*. You parse the network packets to recover audio packets using Audio File Stream Services. You then buffer the audio packets and send them to a playback audio queue object.

Audio File Stream Services relies on interfaces from Audio File Services, such as the `AudioFramePacketTranslation` structure and the `AudioFilePacketTableInfo` structure. These are described in *Audio File Services Reference*.

For more information on using streams, refer to *Audio File Stream Services Reference*. For sample code, see the `AudioFileStream` sample project located in the `<Xcode>/Examples/CoreAudio/Services/` directory, where `<Xcode>` is the path to your developer tools directory.

Audio Unit Support in iPhone OS

iPhone OS provides a set of audio plug-ins, known as *audio units*, that you can use in any application. The interfaces in the Audio Unit framework let you open, connect, and use these audio units. You can also define custom audio units and use them inside your application. Because you must statically link custom audio unit code into your application, audio units that you develop cannot be used by other applications in iPhone OS.

Table 7-3 lists the audio units provided in iPhone OS.

Table 7-3 Supported audio units

Audio unit	Description
3D Mixer unit	The 3D Mixer unit, of type <code>kAudioUnitSubType_AU3DMixerEmbedded</code> , lets you mix multiple audio streams, specify stereo output panning, manipulate sample rate, and more.
Multichannel Mixer unit	The Multichannel Mixer unit, of type <code>kAudioUnitSubType_MultiChannelMixer</code> , lets you mix multiple audio streams.
Converter unit	The Converter unit, of type <code>kAudioUnitSubType_AUConverter</code> , lets you convert audio data from one format to another.
I/O unit	The I/O unit, of type <code>kAudioUnitSubType_RemoteIO</code> , lets you connect to audio input and output hardware and supports realtime I/O. For sample code showing you how to use this audio unit, see <i>aurioTouch</i> .
iPod EQ unit	The iPod EQ unit, of type <code>kAudioUnitSubType_AUiPodEQ</code> , provides a simple, preset-based equalizer you can use in your application.

Best Practices for iPhone Audio

Tips for Manipulating Audio

Table 7-4 lists some basic tips to remember when manipulating audio content in iPhone OS.

Table 7-4 Audio tips

Tip	Action
Use compressed audio appropriately	For AAC, MP3, and ALAC (Apple Lossless) audio, decoding takes place in hardware and, while efficient, is limited to one audio stream at a time. If you need to play multiple sounds simultaneously, store those sounds using the IMA4 (compressed) or linear PCM (uncompressed) format.

Tip	Action
Convert to the data format and file format you need	The <code>afconvert</code> tool in Mac OS X lets you convert to a wide range of audio data formats and file types. See “Preferred Audio Formats in iPhone OS” (page 140) and the <code>afconvert</code> man page.
Evaluate audio memory issues	When playing sound with Audio Queue Services, you write a callback that sends short segments of audio data to audio queue buffers. In some cases, loading an entire sound file to memory for playback, which minimizes disk access, is best. In other cases, loading just enough data at a time to keep the buffers full is best. Test and evaluate which strategy is best for your application.
Reduce audio file sizes by limiting sample rates and bit depths	Sample rates and the number of bits per sample have a direct impact on the size of your uncompressed audio. If you need to play many such sounds, consider reducing these values to reduce the memory footprint of the audio data. For example, rather than use 44.1 kHz sampling rate for sound effects, you could use a 32 kHz (or possibly lower) sample rate and still provide reasonable quality.
Pick the appropriate technology	Use Core Audio’s System Sound Services to play alerts and user-interface sound effects. Use OpenAL when you want a convenient, high-level interface for positioning sounds in a stereo field or when you need low latency playback. To parse audio packets from a file or a network stream, use Audio File Stream Services. For simple playback of single or multiple sounds, use the <code>AVAudioPlayer</code> class. For other audio applications, including playback of streamed audio and audio recording, use Audio Queue Services.
Code for low latency	For the lowest possible playback latency, use OpenAL or use the I/O unit directly.

Preferred Audio Formats in iPhone OS

For uncompressed (highest quality) audio, use 16-bit, little endian, linear PCM audio data packaged in a CAF file. You can convert an audio file to this format in Mac OS X using the `afconvert` command-line tool.

```
/usr/bin/afconvert -f caff -d LEI16 {INPUT} {OUTPUT}
```

The `afconvert` tool lets you convert to a wide range of audio data formats and file types. See the `afconvert` man page, and enter `afconvert -h` at a shell prompt, for more information.

For compressed audio when playing one sound at a time, and when you don’t need to play audio simultaneously with the iPod application, use the AAC format packaged in a CAF or m4a file.

For less memory usage when you need to play multiple sounds simultaneously, use IMA4 (IMA/ADPCM) compression. This reduces file size but entails minimal CPU impact during decompression. As with linear PCM data, package IMA4 data in a CAF file.

Playing Video Files

iPhone OS supports the ability to play back video files directly from your application using the Media Player framework (`MediaPlayer.framework`). Video playback is supported in full screen mode only and can be used by game developers who want to play cut scene animations or by other developers who want to play media files. When you start a video from your application, the media player interface takes over, fading the screen to black and then fading in the video content. You can play a video with or without user controls for adjusting playback; enabling some or all of these controls (shown in Figure 7-1) gives the user the ability to change the volume, change the playback point, or start and stop the video. If you disable all of these controls, the video plays until completion.

Figure 7-1 Media player interface with transport controls



To initiate video playback, you must know the URL of the file you want to play. For files your application provides, this would typically be a pointer to a file in your application's bundle; however, it can also be a pointer to a file on a remote server. You use this URL to instantiate a new instance of the `MPMoviePlayerController` class. This class presides over the playback of your video file and manages user interactions, such as user taps in the transport controls (if shown). To initiate playback, simply call the `play` method of the controller.

Listing 7-12 shows a sample method that plays back the video at the specified URL. The `play` method is an asynchronous call that returns control to the caller while the movie plays. The movie controller loads the movie in a full-screen view, and animates the movie into place on top of the application's existing content. When playback is finished, the movie controller sends a notification to the object, which releases the movie controller now that it is no longer needed.

Listing 7-12 Playing full-screen movies

```
-(void)playMovieAtURL:(NSURL*)theURL
{
    MPMoviePlayerController* theMovie = [[MPMoviePlayerController alloc]
initWithContentURL:theURL];

    theMovie.scalingMode = MPMovieScalingModeAspectFill;
    theMovie.movieControlMode = MPMovieControlModeHidden;

    // Register for the playback finished notification.
    [[NSNotificationCenter defaultCenter] addObserver:self
        selector:@selector(myMovieFinishedCallback:)
        name:MPMoviePlayerPlaybackDidFinishNotification
        object:theMovie];
```

```
        // Movie playback is asynchronous, so this method returns immediately.
        [theMovie play];
    }

    // When the movie is done, release the controller.
    -(void)myMovieFinishedCallback:(NSNotification*)aNotification
    {
        MPMoviePlayerController* theMovie = [aNotification object];

        [[NSNotificationCenter defaultCenter] removeObserver:self
            name:MPMoviePlayerPlaybackDidFinishNotification
            object:theMovie];

        // Release the movie instance created in playMovieAtURL:
        [theMovie release];
    }
```

For more information about the classes of the Media Player framework, see *Media Player Framework Reference*.
For a list of supported video formats, see *iPhone OS Technology Overview*.

Device Support

iPhone OS supports a variety of features that make the mobile computing experience compelling for users. Through iPhone OS, your applications can access hardware features, such as the accelerometers and camera, and software features, such as the user's photo library. The following sections describe these features and show you how to integrate them into your own applications.

Accessing Accelerometer Events

An accelerometer measures changes in velocity over time along a given linear path. The iPhone and iPod touch each contain three accelerometers, one along each of the primary axes of the device. This combination of accelerometers lets you detect movement of the device in any direction. You can use this data to track both sudden movements in the device and the device's current orientation relative to gravity.

Every application has a single `UIAccelerometer` object that can be used to receive acceleration data. You get the instance of this class using the `sharedAccelerometer` class method of `UIAccelerometer`. Using this object, you set the desired reporting interval and a custom delegate to receive acceleration events. You can set the reporting interval to be as small as 10 milliseconds, which corresponds to a 100 Hz update rate, although most applications can operate sufficiently with a larger interval. As soon as you assign your delegate object, the accelerometer starts sending it data. Thereafter, your delegate receives data at the requested update interval.

Listing 8-1 shows the basic steps for configuring the accelerometer. In this example, the update frequency is 50 Hz, which corresponds to an update interval of 20 milliseconds. The `myDelegateObject` is a custom object that you define; it must support the `UIAccelerometerDelegate` protocol, which defines the method used to receive acceleration data.

Listing 8-1 Configuring the accelerometer

```
#define kAccelerometerFrequency      50 //Hz
-(void)configureAccelerometer
{
    UIAccelerometer* theAccelerometer = [UIAccelerometer sharedAccelerometer];
    theAccelerometer.updateInterval = 1 / kAccelerometerFrequency;

    theAccelerometer.delegate = self;
    // Delegate events begin immediately.
}
```

The shared accelerometer delivers event data at regular intervals to your delegate's `accelerometer:didAccelerate:` method, shown in Listing 8-2. You can use this method to process the accelerometer data however you want. In general it is recommended that you use some sort of filter to isolate the component of the data in which you are interested.

Listing 8-2 Receiving an accelerometer event

```

- (void)accelerometer:(UIAccelerometer *)accelerometer
didAccelerate:(UIAcceleration *)acceleration
{
    UIAccelerationValue x, y, z;
    x = acceleration.x;
    y = acceleration.y;
    z = acceleration.z;

    // Do something with the values.
}

```

To stop the delivery of acceleration events, set the delegate of the shared `UIAccelerometer` object to `nil`. Setting the delegate object to `nil` lets the system know that it can turn off the accelerometer hardware as needed, and thus save battery life.

The acceleration data you receive in your delegate method represents the instantaneous values reported by the accelerometer hardware. Even when a device is completely at rest, the values reported by this hardware can fluctuate slightly. When using these values, you should be sure to account for these fluctuations by averaging out the values over time or by calibrating the data you receive. For example, the Bubble Level sample application provides controls for calibrating the current angle against a known surface. Subsequent readings are then reported relative to the calibrated angle. If your own code requires a similar level of accuracy, you should also include some sort of calibration option in your user interface.

Choosing an Appropriate Update Interval

When configuring the update interval for acceleration events, it is best to choose an interval that minimizes the number of delivered events while still meeting the needs of your application. Few applications need acceleration events delivered 100 times a second. Using a lower frequency prevents your application from running as often and can therefore improve battery life. Table 8-1 lists some typical update frequencies and what you can do with the acceleration data generated at that frequency.

Table 8-1 Common update intervals for acceleration events

Event frequency (Hz)	Usage
10–20	Suitable for use in determining the vector representing the current orientation of the device.
30–60	Suitable for games and other applications that use the accelerometers for real-time user input.
70–100	Suitable for applications that need to detect high-frequency motion. For example, you might use this interval to detect the user hitting the device or shaking it very quickly.

Isolating the Gravity Component from Acceleration Data

If you are using the accelerometer data to detect the current orientation of a device, you need to be able to filter out the portion of the acceleration data that is caused by gravity from the portion that is caused by motion of the device. To do this, you can use a low-pass filter to reduce the influence of sudden changes on the accelerometer data. The resulting filtered values would then reflect the more constant effects of gravity.

Listing 8-3 shows a simplified version of a low-pass filter. This example uses a low-value filtering factor to generate a value that uses 10 percent of the unfiltered acceleration data and 90 percent of the previously filtered value. The previous values are stored in the `accelX`, `accelY`, and `accelZ` member variables of the class. Because acceleration data comes in regularly, these values settle out quickly and respond slowly to sudden but short-lived changes in motion.

Listing 8-3 Isolating the effects of gravity from accelerometer data

```
#define kFilteringFactor 0.1

- (void)accelerometer:(UIAccelerometer *)accelerometer
didAccelerate:(UIAcceleration *)acceleration {
    // Use a basic low-pass filter to keep only the gravity component of each
    axis.
    accelX = (acceleration.x * kFilteringFactor) + (accelX * (1.0 -
kFilteringFactor));
    accelY = (acceleration.y * kFilteringFactor) + (accelY * (1.0 -
kFilteringFactor));
    accelZ = (acceleration.z * kFilteringFactor) + (accelZ * (1.0 -
kFilteringFactor));

    // Use the acceleration data.
}
```

Isolating Instantaneous Motion from Acceleration Data

If you are using accelerometer data to detect just the instant motion of a device, you need to be able to isolate sudden changes in movement from the constant effect of gravity. You can do that with a high-pass filter.

Listing 8-4 shows a simplified high-pass filter computation. The acceleration values from the previous event are stored in the `accelX`, `accelY`, and `accelZ` member variables of the class. This example computes the low-pass filter value and then subtracts it from the current value to obtain just the instantaneous component of motion.

Listing 8-4 Getting the instantaneous portion of movement from accelerometer data

```
#define kFilteringFactor 0.1

- (void)accelerometer:(UIAccelerometer *)accelerometer
didAccelerate:(UIAcceleration *)acceleration {
    // Subtract the low-pass value from the current value to get a simplified
    high-pass filter
    accelX = acceleration.x - ( (acceleration.x * kFilteringFactor) + (accelX
* (1.0 - kFilteringFactor)) );
    accelY = acceleration.y - ( (acceleration.y * kFilteringFactor) + (accelY
* (1.0 - kFilteringFactor)) );
```

```

    accelZ = acceleration.z - ( (acceleration.z * kFilteringFactor) + (accelZ
* (1.0 - kFilteringFactor)) );

    // Use the acceleration data.
}

```

Getting the Current Device Orientation

If you need to know only the general orientation of the device, and not the exact vector of orientation, you should use the methods of the `UIDevice` class to retrieve that information. Using the `UIDevice` interface is simpler and does not require you to calculate the orientation vector yourself.

Before getting the current orientation, you must tell the `UIDevice` class to begin generating device orientation notifications by calling the `beginGeneratingDeviceOrientationNotifications` method. Doing so turns on the accelerometer hardware (which may otherwise be off to conserve power).

Shortly after enabling orientation notifications, you can get the current orientation from the `orientation` property of the shared `UIDevice` object. You can also register to receive `UIDeviceOrientationDidChangeNotification` notifications, which are posted whenever the general orientation changes. The device orientation is reported using the `UIDeviceOrientation` constants, which indicate whether the device is in landscape or portrait mode or whether the device is face up or face down. These constants indicate the physical orientation of the device and need not correspond to the orientation of your application's user interface.

When you no longer need to know the orientation of the device, you should always disable orientation notifications by calling the `endGeneratingDeviceOrientationNotifications` method of `UIDevice`. Doing so gives the system the opportunity to disable the accelerometer hardware if it is not in use elsewhere.

Getting the User's Current Location

The Core Location framework lets you locate the current position of the device and use that information in your application. The framework takes advantage of the device's built-in hardware, triangulating a position fix from available signal information. It then reports the location to your code and occasionally updates that position information as it receives new or improved signals.

If you do use the Core Location framework, be sure to do so sparingly and to configure the location service appropriately. Gathering location data involves powering up the onboard radios and querying the available cell towers, Wi-Fi hotspots, or GPS satellites, which can take several seconds. In addition, requesting more accurate location data may require the radios to remain on for a longer period of time. Leaving this hardware on for extended periods of time can drain the device's battery. Given that position information does not change too often, it is usually sufficient to establish an initial position fix and then acquire updates periodically after that. If you are sure you need regular position updates, you can also configure the service with a minimum threshold distance to minimize the number of position updates your code must process.

To retrieve the user's current location, create an instance of the `CLLocationManager` class and configure it with the desired accuracy and threshold parameters. To begin receiving location notifications, assign a delegate to the object and call the `startUpdatingLocation` method to start the determination of the user's current location. When new location data is available, the location manager notifies its assigned delegate object. If a location update has already been delivered, you can also get the most recent location data directly from the `CLLocationManager` object without waiting for a new event to be delivered.

Listing 8-5 shows implementations of a custom `startUpdates` method and the `locationManager:didUpdateToLocation:fromLocation:` delegate method. The `startUpdates` method creates a new location manager object (if one does not already exist) and uses it to start generating location updates. (In this case, the `locationManager` variable is a member variable declared by the `MyLocationGetter` class, which also conforms to the `CLLocationManagerDelegate` protocol.) The handler method uses the timestamp of the event to determine how recent it is. If it is an old event, the handler ignores it and waits for a more recent one, at which point it disables the location service.

Listing 8-5 Initiating and processing location updates

```
#import <CoreLocation/CoreLocation.h>

@implementation MyLocationGetter
- (void)startUpdates
{
    // Create the location manager if this object does not
    // already have one.
    if (nil == locationManager)
        locationManager = [[CLLocationManager alloc] init];

    locationManager.delegate = self;
    locationManager.desiredAccuracy = kCLLocationAccuracyKilometer;

    // Set a movement threshold for new events
    locationManager.distanceFilter = 500;

    [locationManager startUpdatingLocation];
}

// Delegate method from the CLLocationManagerDelegate protocol.
- (void)locationManager:(CLLocationManager *)manager
    didUpdateToLocation:(CLLocation *)newLocation
    fromLocation:(CLLocation *)oldLocation
{
    // If it's a relatively recent event, turn off updates to save power
    NSDate* eventDate = newLocation.timestamp;
    NSTimeInterval howRecent = [eventDate timeIntervalSinceNow];
    if (abs(howRecent) < 5.0)
    {
        [manager stopUpdatingLocation];

        printf("latitude %+.6f, longitude %+.6f\n",
               newLocation.coordinate.latitude,
               newLocation.coordinate.longitude);
    }
    // else skip the event and process the next one.
}
@end
```

Checking the timestamp of an event is recommended because the location service often returns the last cached location event immediately. It can take several seconds to obtain a rough location fix so the old data simply serves as a way to reflect the last known location. You can also use the accuracy as a means of determining whether you want to accept an event. As it receives more accurate data, the location service may return additional events, with the accuracy values reflecting the improvements accordingly.

Note: The Core Location framework records timestamp values at the beginning of each location query, not when that query returns. Because Core Location uses several different techniques to get a location fix, queries can sometimes come back in a different order than their timestamps might otherwise indicate. As a result, it is normal for new events to sometimes have timestamps that are slightly older than those from previous events. The framework concentrates on improving the accuracy of the location data with each new event it delivers, regardless of the timestamp values.

For more information about the objects and methods of the Core Location framework, see *Core Location Framework Reference*.

Taking Pictures with the Camera

UIKit provides access to a device's camera through the `UIImagePickerController` class. This class displays the standard system interface for taking pictures using the available camera. It also supports optional controls for resizing and cropping the image after the user takes it. This class can also be used to select photos from the user's photo library.

The view representing the camera interface is a modal view that is managed by the `UIImagePickerController` class. You should never access this view directly from your code. To display it, you must call the `presentModalViewController:animated:` method of the currently active view controller, passing a `UIImagePickerController` object as the new view controller. Upon being installed, the picker controller automatically slides the camera interface into position, where it remains active until the user approves the picture or cancels the operation. At that time, the picker controller notifies its delegate of the user's choice.

Interfaces managed by the `UIImagePickerController` class may not be available on all devices. Before displaying the camera interface, you should always make sure that the interface is available by calling the `isSourceTypeAvailable:` class method of the `UIImagePickerController` class. You should always respect the return value of this method. If this method returns `NO`, it means that the current device does not have a camera or that the camera is currently unavailable for some reason. If the method returns `YES`, you display the camera interface by doing the following:

1. Create a new `UIImagePickerController` object.
2. Assign a delegate object to the picker controller.

In most cases, the current view controller acts as the delegate for the picker, but you can use an entirely different object if you prefer. The delegate object must conform to the `UIImagePickerControllerDelegate` protocol.

3. Set the picker type to `UIImagePickerControllerSourceTypeCamera`.
4. Optionally, enable or disable the picture editing controls by assigning an appropriate value to the `allowsImageEditing` property.
5. Call the `presentModalViewController:animated:` method of the current view controller to display the picker.

Listing 8-6 shows the code representing the preceding set of steps. As soon as you call the `presentModalViewController:animated` method, the picker controller takes over, displaying the camera interface and responding to all user interactions until the interface is dismissed. To choose an existing photo from the user's photo library, all you have to do is change the value in the `sourceType` property of the picker to `UIImagePickerControllerSourceTypePhotoLibrary`.

Listing 8-6 Displaying the interface for taking pictures

```
- (BOOL)startCameraPickerFromViewController:(UIViewController*)controller
usingDelegate:(id<UIImagePickerControllerDelegate>)delegateObject
{
    if ( (![UIImagePickerController
isSourceTypeAvailable:UIImagePickerControllerSourceTypeCamera])
        || (delegateObject == nil) || (controller == nil))
        return NO;

    UIImagePickerController* picker = [[UIImagePickerController alloc] init];
    picker.sourceType = UIImagePickerControllerSourceTypeCamera;
    picker.delegate = delegateObject;
    picker.allowsImageEditing = YES;

    // Picker is displayed asynchronously.
    [controller presentModalViewController:picker animated:YES];
    return YES;
}
```

When the user taps the appropriate button to dismiss the camera interface, the `UIImagePickerController` notifies the delegate of the user's choice but does not dismiss the interface. The delegate is responsible for dismissing the picker interface. (Your application is also responsible for releasing the picker when done with it, which you can do in the delegate methods.) It is for this reason that the delegate is actually the view controller object that presented the picker in the first place. Upon receiving the delegate message, the view controller would call its `dismissModalViewControllerAnimated:` method to dismiss the camera interface.

Listing 8-7 shows the delegate methods for dismissing the camera interface displayed in Listing 8-6 (page 149). These methods are implemented by a custom `MyViewController` class, which is a subclass of `UIViewController` and, for this example, is considered to be the same object that displayed the picker in the first place. The `useImage:` method is an empty placeholder for the work you would do in your own version of this class and should be replaced by your own custom code.

Listing 8-7 Delegate methods for the image picker

```
@implementation MyViewController (ImagePickerDelegateMethods)

- (void)imagePickerController:(UIImagePickerController *)picker
    didFinishPickingImage:(UIImage *)image
    editingInfo:(NSDictionary *)editingInfo
{
    [self useImage:image];

    // Remove the picker interface and release the picker object.
    [[picker parentViewController] dismissModalViewControllerAnimated:YES];
    [picker release];
}

- (void)imagePickerControllerDidCancel:(UIImagePickerController *)picker
{
}
```

```
[[picker parentViewController] dismissModalViewControllerAnimated:YES];  
[picker release];  
}  
  
// Implement this method in your code to do something with the image.  
- (void)useImage:(UIImage*)theImage  
{  
}  
}  
@end
```

If image editing is enabled and the user successfully picks an image, the `image` parameter of the `imagePickerController:didFinishPickingImage:editingInfo:` method contains the edited image. You should treat this image as the selected image, but if you want to store the original image, you can get it (along with the crop rectangle) from the dictionary in the `editingInfo` parameter.

Picking a Photo from the Photo Library

UIKit provides access to the user's photo library through the `UIImagePickerController` class. This controller displays a photo picker interface, which provides controls for navigating the user's photo library and selecting an image to return to your application. You also have the option of enabling user editing controls, which let the user pan and crop the returned image. This class can also be used to present a camera interface.

Because the `UIImagePickerController` class is used to display the interface for both the camera and the user's photo library, the steps for using the class are almost identical for both. The only difference is that you assign the `UIImagePickerControllerSourceTypePhotoLibrary` value to the `sourceType` property of the picker object. The steps for displaying the camera picker are discussed in ["Taking Pictures with the Camera"](#) (page 148).

Note: As you do for the camera picker, you should always call the `isSourceTypeAvailable:` class method of the `UIImagePickerController` class and respect the return value of the method. You should never assume that a given device has a photo library. Even if the device has a library, this method could still return `NO` if the library is currently unavailable.

Application Preferences

In traditional desktop applications, preferences are application-specific settings used to configure the behavior or appearance of an application. iPhone OS also supports application preferences, although not as an integral part of your application. Instead of each application displaying a custom user interface for its preferences, all application-level preferences are displayed using the system-supplied Settings application.

In order to integrate your custom application preferences into the Settings application, you must include a specially formatted settings bundle in the top-level directory of your application bundle. This settings bundle provides information about your application preferences to the Settings application, which is then responsible for displaying those preferences and updating the preferences database with any user-supplied values. At runtime, your application retrieves these preferences using the standard retrieval APIs. The sections that follow describe both the format of the settings bundle and the APIs you use to retrieve your preferences values.

Guidelines for Preferences

Adding your application preferences to the Settings application is most appropriate for productivity-style applications and in situations where you have preference values that are typically configured once and then rarely changed. For example, the Mail application uses these preferences to store the user's account information and message-checking settings. Because the Settings application has support for displaying preferences hierarchically, manipulating your preferences from the Settings application is also more appropriate when you have a large number of preferences. Providing the same set of preferences in your application might require too many screens and might cause confusion for the user.

When your application has only a few options or has options that the user might want to change regularly, you should think carefully about whether the Settings application is the right place for them. For instance, utility applications provide custom configuration options on the back of their main view. A special control on the view flips it over to display the options and another control flips the view back. For simple applications, this type of behavior provides immediate access to the application's options and is much more convenient for the user than going to Settings.

For games and other full-screen applications, you can use the Settings application or implement your own custom screens for preferences. Custom screens are often appropriate in games because those screens are treated as part of the game's setup. You can also use the Settings application for your preferences if you think it is more appropriate for your game flow.

Note: You should never spread your preferences across the Settings application and custom application screens. For example, a utility application with preferences on the back side of its main view should not also have configurable preferences in the Settings application. If you have preferences, pick one solution and use it exclusively.

The Preferences Interface

The Settings application implements a hierarchical set of pages for navigating application preferences. The main page of the Settings application displays the system and third-party applications whose preferences can be customized. Selecting a third-party application takes the user to the preferences for that application.

Each application has at least one page of preferences, referred to as the main page. If your application has only a few preferences, the main page may be the only one you need. If the number of preferences gets too large to fit on the main page, however, you can add more pages. These additional pages become child pages of the main page. The user accesses them by tapping on a special type of preference, which links to the new page.

Each preference you display must be of a specific type. The type of the preference defines how the Settings application displays that preference. Most preference types identify a particular type of control that is used to set the preference value. Some types provide a way to organize preferences, however. Table 9-1 lists the different element types supported by the Settings application and how you might use each type to implement your own preference pages.

Table 9-1 Preference element types

Element Type	Description
Text Field	The text field type displays an optional title and an editable text field. You can use this type for preferences that require the user to specify a custom string value. The key for this type is <code>PSTextFieldSpecifier</code> .
Title	The title type displays a read-only string value. You can use this type to display read-only preference values. (If the preference contains cryptic or nonintuitive values, this type lets you map the possible values to custom strings.) The key for this type is <code>PSTitleValueSpecifier</code> .
Toggle Switch	The toggle switch type displays an ON/OFF toggle button. You can use this type to configure a preference that can have only one of two values. Although you typically use this type to represent preferences containing Boolean values, you can also use it with preferences containing non-Boolean values. The key for this type is <code>PSToggleSwitchSpecifier</code> .
Slider	The slider type displays a slider control. You can use this type for a preference that represents a range of values. The value for this type is a real number whose minimum and maximum you specify. The key for this type is <code>PSSliderSpecifier</code> .

Element Type	Description
Multi value	The multi value type lets the user select one value from a list of values. You can use this type for a preference that supports a set of mutually exclusive values. The values can be of any type. The key for this type is <code>PSMultiValueSpecifier</code> .
Group	The group type is a way for you to organize groups of preferences on a single page. The group type does not represent a configurable preference. It simply contains a title string that is displayed immediately before one or more configurable preferences. The key for this type is <code>PSGroupSpecifier</code> .
Child Pane	The child pane type lets the user navigate to a new page of preferences. You use this type to implement hierarchical preferences. For more information on how you configure and use this preference type, see “Hierarchical Preferences” (page 155). The key for this type is <code>PSChildPaneSpecifier</code> .

For detailed information about the format of each preference type, see *Settings Application Schema Reference*. To learn how to create and edit Setting page files, see [“Adding and Modifying the Settings Bundle”](#) (page 156).

The Settings Bundle

In iPhone OS, you specify your application’s preferences through a special settings bundle. This bundle has the name `Settings.bundle` and resides in the top-level directory of your application’s bundle. This bundle contains one or more Settings Page files that provide detailed information about your application’s preferences. It may also include other support files needed to display your preferences, such as images or localized strings. Table 9-2 lists the contents of a typical settings bundle.

Table 9-2 Contents of the `Settings.bundle` directory

Item name	Description
<code>Root.plist</code>	The Settings Page file containing the preferences for the root page. The contents of this file are described in more detail in “The Settings Page File Format” (page 154).
Additional <code>.plist</code> files.	If you build a set of hierarchical preferences using child panes, the contents for each child pane are stored in a separate Settings Page file. You are responsible for naming these files and associating them with the correct child pane.
One or more <code>.lproj</code> directories	These directories store localized string resources for your Settings Page files. Each directory contains a single strings file, whose title is specified in your Settings Page. The strings files provide the localized content to display to the user for each of your preferences.
Additional images	If you use the slider control, you can store the images for your slider in the top-level directory of the bundle.

In addition to the settings bundle, your application bundle can contain a custom icon for your application settings. If a file with the name `Icon-Settings.png` is located in the top of your application's bundle directory, that icon is used to identify your application preferences in the Settings application. If no such image file is present, the Settings application uses your application's icon file (`Icon.png` by default) instead, scaling it as necessary. Your `Icon-Settings.png` file should be a 29 x 29 pixel image.

When the Settings application launches, it checks each custom application for the presence of a settings bundle. For each custom bundle it finds, it loads that bundle and displays the corresponding application's name and icon in the Settings main page. When the user taps the row belonging to your application, Settings loads the `Root.plist` Settings Page file for your settings bundle and uses that file to display your application's main page of preferences.

In addition to loading your bundle's `Root.plist` Settings Page file, the Settings application also loads any language-specific resources for that file, as needed. Each Settings Page file can have an associated `.strings` file containing localized values for any user-visible strings. As it prepares your preferences for display, the Settings application looks for string resources in the user's preferred language and substitutes them into your preferences page prior to display.

The Settings Page File Format

Each Settings Page file in your settings bundle is stored in the iPhone Settings property-list file format, which is a structured file format. The simplest way to edit Settings Page files is using Xcode's built in editor facilities; see [“Preparing the Settings Page for Editing”](#) (page 157). You can also edit property-list files using the Property List Editor application that comes with the Xcode tools.

Note: Xcode automatically converts any XML-based property files in your project to binary format when building your application. This conversion saves space and is done for you automatically at build time.

The root element of each Settings Page file contains the keys listed in Table 9-3. Only one key is actually required, but it is recommended that you include both of them.

Table 9-3 Root-level keys of a preferences Settings Page file

Key	Type	Value
<code>PreferenceSpecifiers</code> (required)	Array	The value for this key is an array of dictionaries, with each dictionary containing the information for a single preference element. For a list of element types, see Table 9-1 (page 152). For a description of the keys associated with each element type, see <i>Settings Application Schema Reference</i> .
<code>StringsTable</code>	String	The name of the strings file associated with this file. A copy of this file (with appropriate localized strings) should be located in each of your bundle's language-specific project directories. If you do not include this key, the strings in this file are not localized. For information on how these strings are used, see “Localized Resources” (page 156).

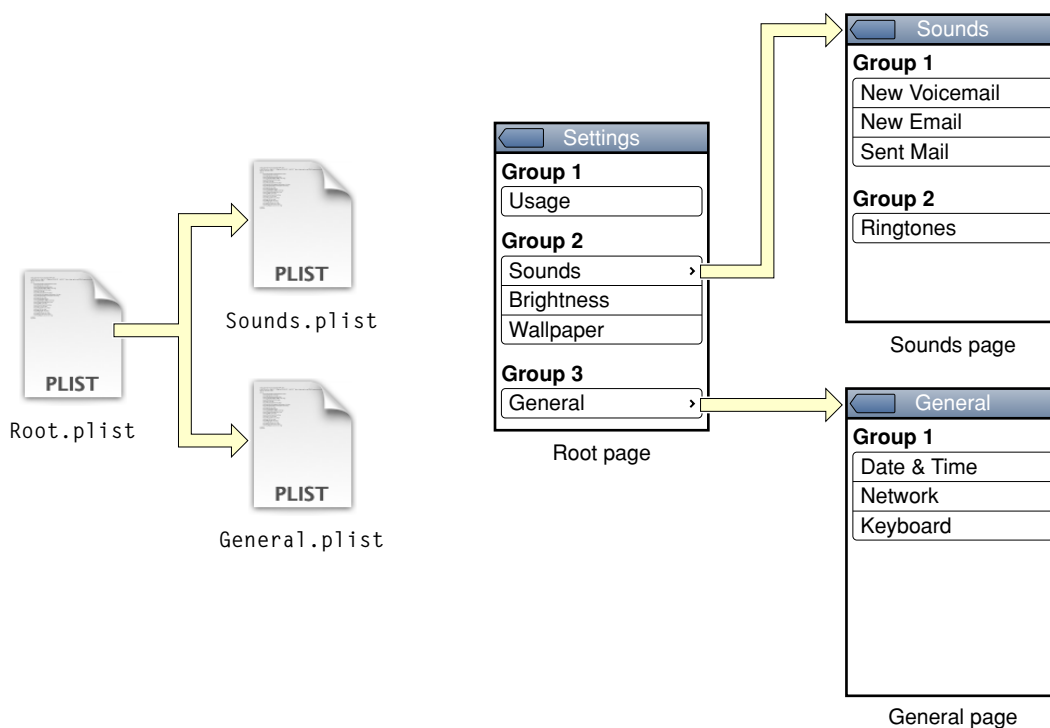
Hierarchical Preferences

If you plan to organize your preferences hierarchically, each page you define must have its own separate `.plist` file. Each `.plist` file contains the set of preferences displayed only on that page. Your application's main preferences page is always stored in the `Root.plist` file. Additional pages can be given any name you like.

To specify a link between a parent page and a child page, you include a child pane element in the parent page. A child pane element creates a row that, when tapped, displays a new page of settings. The `File` key of the child pane element identifies the name of the `.plist` file that defines the contents of the child page. The `Title` key identifies the title of the child page; this title is also used as the text of the row that the user taps to display the child page. The Settings application automatically provides navigation controls on the child page to allow the user to navigate back to the parent page.

Figure 9-1 shows how this hierarchical set of pages works. The left side of the figure shows the `.plist` files, and the right side shows the relationships between the corresponding pages.

Figure 9-1 Organizing preferences using child panes



For more information about child pane elements and their associated keys, see *Settings Application Schema Reference*.

Localized Resources

Because preferences contain user-visible strings, you should provide localized versions of those strings with your settings bundle. Each page of preferences can have an associated `.strings` file for each localization supported by your bundle. When the Settings application encounters a key that supports localization, it checks the appropriately localized `.strings` file for a matching key. If it finds one, it displays the value associated with that key.

When looking for localized resources such as `.strings` files, the Settings application follows the same rules that Mac OS X applications do. It first tries to find a localized version of the resource that matches the user's preferred language setting. If a resource does not exist for the user's preferred language, an appropriate fallback language is selected.

For information about the format of strings files, language-specific project directories, and how language-specific resources are retrieved from bundles, see *Internationalization Programming Topics*.

Adding and Modifying the Settings Bundle

Xcode provides a template for adding a Settings bundle to your current project. The default settings bundle contains a `Root.plist` file and a default language directory for storing any localized resources. You can then expand this bundle to include additional property list files and resources needed by your Settings bundle.

Adding the Settings Bundle

To add a settings bundle to your Xcode project:

1. Choose File > New File.
2. Choose the iPhone OS > Settings > Settings Bundle template.
3. Name the file `Settings.bundle`.

In addition to adding a new Settings bundle to your project, Xcode automatically adds that bundle to the Copy Bundle Resources build phase of your application target. Thus, all you have to do is modify the property list files of your Settings bundle and add any needed resources.

The newly added `Settings.bundle` bundle has the following structure:

```
Settings.bundle/  
  Root.plist  
  en.lproj/  
    Root.strings
```

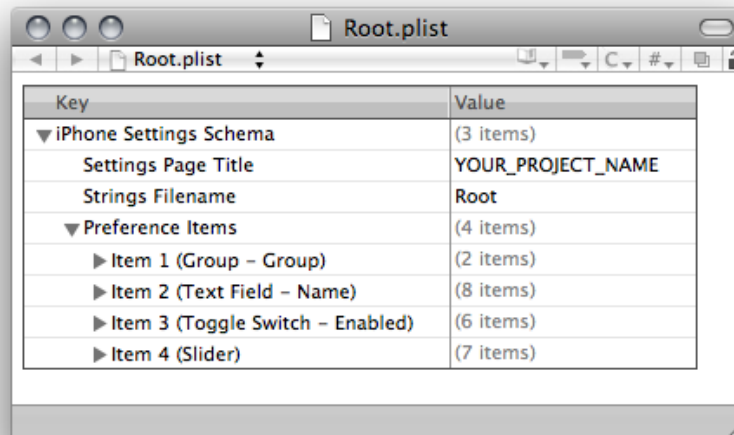
Preparing the Settings Page for Editing

After creating your Settings bundle using the Settings Bundle template, you can format the contents of your schema files to make them easier to edit. The following steps show you how to do this for the `Root.plist` file of your Settings bundle but the steps are the same for any other schema files you create.

1. Display the contents of the `Root.plist` file of your Settings bundle.
 - a. In the Groups & Files list, disclose `Settings.bundle` to view its contents.
 - b. Select the `Root.plist` file. Its contents appear in the Detail view.
2. In the Detail view, select the Root key of the `Root.plist` file.
3. Choose View > Property List Type > iPhone Settings plist.

This command formats the contents of the property list inside the Detail view. Instead of showing the property list key names and values, Xcode substitutes human-readable strings (as shown in Figure 9-2) to make it easier to understand and edit the file's contents.


Figure 9-2 Formatted contents of the `Root.plist` file



Key	Value
▼ iPhone Settings Schema	(3 items)
Settings Page Title	YOUR_PROJECT_NAME
Strings Filename	Root
▼ Preference Items	(4 items)
▶ Item 1 (Group – Group)	(2 items)
▶ Item 2 (Text Field – Name)	(8 items)
▶ Item 3 (Toggle Switch – Enabled)	(6 items)
▶ Item 4 (Slider)	(7 items)

Configuring a Settings Page: A Tutorial

This section contains a tutorial that shows you how to configure a Settings page to display the controls you want. The goal of the tutorial is to create a page like the one in Figure 9-2. If you have not yet created a Settings bundle for your project, you should do so as described in [“Preparing the Settings Page for Editing”](#) (page 157) before proceeding with these steps.

Figure 9-3 A root Settings page

1. Change the value of the Settings Page Title key to the name of your application.
Double-click the `YOUR_PROJECT_NAME` text and change the text to `MyApp`.
2. Disclose the Preference Items key to display the default items that come with the template.
3. Change the title of Item 1 to Sound:
 - Disclose Item 1 of Preference Items.
 - Change the value of the Title key from Group to Sound.
 - Leave the Type key set to Group.
4. Create the first toggle switch for the newly renamed Sound group.
 - Select Item 3 of Preference Items and choose Edit > Cut.
 - Select Item 1 and choose Edit > Paste. (This moves the toggle switch item in front of the text field item.)
 - Disclose the toggle switch item to reveal its configuration keys.
 - Change the value of the Title key to Play Sounds.
 - Change the value of the Identifier key to `play_sounds_preference`. The item should now be configured as shown in the following figure.

▼ Item 2 (Toggle Switch – Play Sounds) ▼ (6 items)	
Type	Toggle Switch
Title	Play Sounds
Identifier	play_sounds_preference
Default Value	<input checked="" type="checkbox"/>
Value for ON	YES
Value for OFF	NO

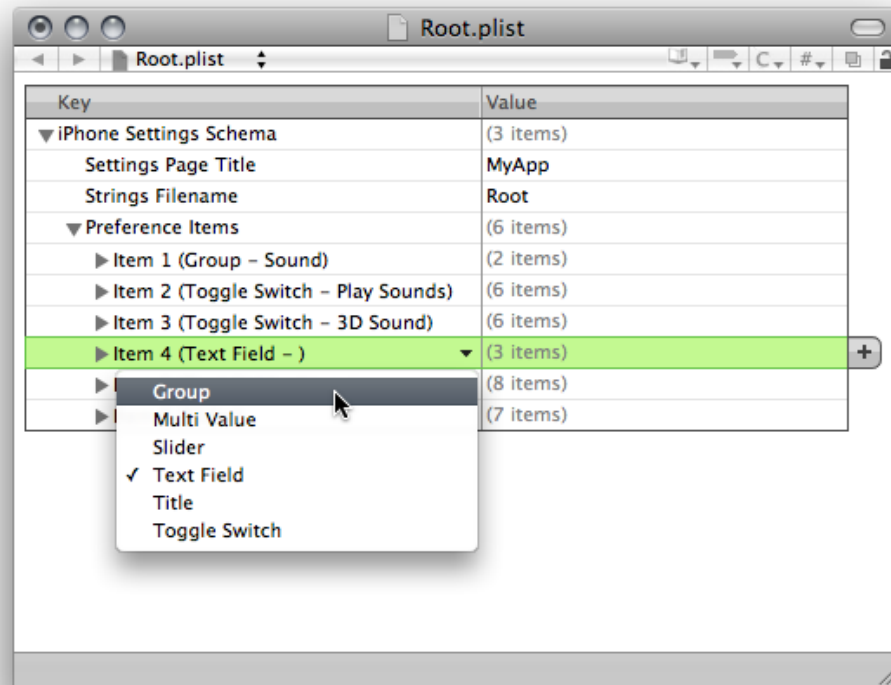
5. Create a second toggle switch for the Sound group.

- Select `Item 2` (the Play Sounds toggle switch).
- Select `Edit > Copy`.
- Select `Edit > Paste`. This places a copy of the toggle switch right after the first one.
- Disclose the new toggle switch item to reveal its configuration keys.
- Change the value of its `Title` key to `3D Sound`.
- Change the value of its `Identifier` key to `3D_sound_preference`.

At this point, you have finished the first group of settings and are ready to create the User Info group.

6. Change `Item 4` into a Group element and name it `User Info`.

- Click `Item 4` in the `Preferences Items`. This displays a drop-down menu with a list of item types.
- From the drop-down menu, select `Group` to change the type of the element.



- Disclose the contents of Item 4.
 - Set the value of the Title key to User Info.
7. Create the Name field.
- Select Item 5 in the Preferences Item.
 - Using the drop-down menu, change its type to Text Field.
 - Set the value of the Title key to User Info.
 - Set value of the Identifier key to user_name.
 - Toggle the disclosure button of the item to hide its contents.
8. Create the Experience Level settings.
- Select Item 5 and click the plus (+) button (or press Return) to create a new item.
 - Click the new item and set its type to Multi Value.
 - Disclose the items contents and set its title to Experience Level, its identifier to experience_preference, and its default value to 0.
 - With the Default Value key selected, click the plus button to add a Titles array.
 - Open the disclosure button for the Titles array and click the items button along the right edge of the table. Clicking this button adds a new subitem to Titles.

- Select the new subitem and click the plus button 2 more times to create 3 total subitems.
 - Set the values of the subitems to `Beginner`, `Expert`, and `Master`.
 - Select the `Titles` key again and click its disclosure button to hide its subitems.
 - Click the plus button to create the `Values` array.
 - Add 3 subitems to the `Values` array and set their values to 0, 1, and 2.
 - Click the disclosure button of `Item 6` to hide its contents.
9. Add the final group to your settings page.
- Create a new item and set its type to `Group` and its title to `Gravity`.
 - Create another new item and set its type to `Slider`, its identifier to `gravity_preference`, its default value to 1, and its maximum value to 2.

Creating Additional Settings Page Files

The Settings Bundle template includes the `Root.plist` file, which defines your application's top Settings page. To define additional Settings pages, you must add additional property list files to your Settings bundle. You can do this either from the Finder or from Xcode.

To add a property list file to your Settings bundle in Xcode, do the following:

1. In the Groups and Files pane, open your Settings bundle and select the `Root.plist` file.
2. Choose `File > New`.
3. Choose `Other > Property List`.
4. Select the new file and choose `View > Property List Type > iPhone Settings plist` to configure it as a settings file.

After adding a new Settings page to your Settings bundle, you can edit the page's contents as described in [“Configuring a Settings Page: A Tutorial”](#) (page 157). To display the settings for your page, you must reference it from a Child Pane element as described in [“Hierarchical Preferences”](#) (page 155).

Accessing Your Preferences

iPhone applications get and set preferences values using either the Foundation and Core Foundation frameworks. In the Foundation framework, you use the `NSUserDefaults` class to get and set preference values. In the Core Foundation framework, you use several preferences-related functions to get and set values.

Listing 9-1 shows a simple example of how to read a preference value from your application. This example uses the `NSUserDefaults` class to read a Boolean value from preferences and assign it to an application-specific instance variable.

Listing 9-1 Accessing preference values in an application

```
- (void)applicationDidFinishLaunching:(UIApplication *)application
{
    NSUserDefaults *defaults = [NSUserDefaults standardUserDefaults];
    [self setMyAppBoolProperty:[defaults boolForKey:MY_BOOL_PREF_KEY]];

    // Finish app initialization...
}
```

For information about the `NSUserDefaults` methods used to read and write preferences, see *NSUserDefaults Class Reference*. For information about the Core Foundation functions used to read and write preferences, see *Preferences Utilities Reference*.

Debugging Preferences for Simulated Applications

When running your application, the iPhone Simulator stores any preferences values for your application in `~/Library/Application Support/iPhone Simulator/User/Applications/<APP_ID>/Library/Preferences`, where `<APP_ID>` is a programmatically generated directory name that iPhone OS uses to identify your application.

Each time you reinstall your application, iPhone OS performs a clean install, which deletes any previous preferences. In other words, building or running your application from Xcode always installs a new version, replacing any old contents. To test preference changes between successive executions, you must run your application directly from the simulator interface and not from Xcode.

Document Revision History

This table describes the changes to *iPhone Application Programming Guide*.

Date	Notes
2009-01-06	Fixed several typos and clarified the creation process for child pages in the Settings application.
2008-11-12	Added guidance about floating-point math considerations
	Updated information related to what is backed up by iTunes.
2008-10-15	Reorganized the contents of the book.
	Moved the high-level iPhone OS information to <i>iPhone OS Technology Overview</i> .
	Moved information about the standard system URL schemes to <i>Apple URL Scheme Reference</i> .
	Moved information about the development tools and how to configure devices to <i>iPhone Development Guide</i> .
	Created the Core Application chapter, which now introduces the application architecture and covers much of the guidance for creating iPhone applications.
	Added a Text and Web chapter to cover the use of text and web classes and the manipulation of the onscreen keyboard.
	Created a separate chapter for Files and Networking and moved existing information into it.
	Changed the title from <i>iPhone OS Programming Guide</i> .
2008-07-08	New document that describes iPhone OS and the development process for iPhone applications.

