
Windows Phone Programming in C#

Rob Miles

Windows Phone Version 7.5

Contents

1	Windows Phone	7
1.1	The Windows Phone Platform.....	7
	A Windows Phone as a Computer	7
	The Windows Phone Hardware	7
	The Windows Phone Processor	8
	The Windows Phone Operating System	8
	Graphical Display	9
	Touch input.....	9
	Location Sensors	10
	Accelerometer	10
	Compass	10
	Gyroscope.....	11
	Sensor Integration and Simulation	11
	Camera	11
	Hardware buttons	11
	Memory and Storage	12
	Network Connectivity	12
	Platform Challenges	13
	The Good News.....	13
1.2	The Windows Phone Ecosystem	13
	The Zune Media Management Software	14
	Windows Live and Xbox Live.....	14
	Bing Maps	14
	Windows Notification Service.....	14
	Windows Phone and Windows Azure	14
	Using the Ecosystem	15
1.3	Windows Phone program execution.....	15
	Application Switching on Windows Phone	16
	Background Processing	16
	Windows Phone and Managed Code.....	16
1.4	Windows Phone application development.....	18
	The Windows Phone Emulator.....	19
	Accessing Windows Phone Facilities	19
	Windows Phone Connectivity	19
	Silverlight and XNA Development	19
	Combining Silverlight and XNA	20
	Data Storage on Windows Phone	20
	Development Tools	20
	Windows Marketplace.....	21
	What We Have Learned.....	21
2	Introduction to Silverlight	23
2.1	Program Design with Silverlight	23
	Development Tools	23
	The Metro Design Style.....	23
	Silverlight Elements and Objects	24
	The Toolbox and Design Surface	26
	Managing Element Names in Visual Studio.....	28
	Properties in Silverlight Elements	29

	Using Properties	30
	Page Design with Silverlight	32
2.2	Understanding XAML	32
	Extensible Markup Languages	33
	XAML and pages	35
2.3	Creating a Silverlight Application	35
	Building the Application	36
	Calculating the Result.....	37
	Events and Programs	38
	Events in Silverlight	38
	Managing Event Properties	40
	Events and XAML.....	41
	What We Have Learned.....	41
3	Visual Studio Solution Management	42
3.1	Getting Started with Projects and Solutions	42
	A Simple Visual Studio Project.....	44
	Adding Resources to a Project.....	47
	Assembly files and Executable files	49
	Visual Studio Solutions	52
	Windows Phone Solutions.....	54
	Creating a Windows Phone solution.....	56
	Running Windows Phone applications	57
3.2	Debugging Programs	59
	Using the Windows Phone emulator	59
	Visual Studio Debugging	60
	Controlling Program Execution.....	62
3.3	Performance Tuning	64
	Using Performance Analysis Tool.....	64
	What We Have Learned.....	65
4	Constructing a program with Silverlight	66
4.1	Improving the User Experience	66
	Manipulating Element Properties	66
	Editing the XAML for Silverlight Elements.....	69
	Displaying a MessageBox	72
	Adding and Using Assets	74
	Adding Images as Items of Content	75
	Adding Images as Resources	77
	Content vs Resources	79
4.2	Data manipulation and display	80
	The TextChanged Event	80
	Data Binding.....	81
	Data Binding using the Data Context	86
4.3	Managing Application Page Layout	87
	Landscape and Portrait Programs	87
	Using Containers to Layout displays	90
4.4	Displaying Lists of Data	91
	Creating a Customer List.....	91
	Making Sample Data	92
	Using the StackPanel to display a List	93
	Using the ListBox to display lists of items	95
	Selecting items in a ListBox	98
4.5	Pages and Navigation	99
	Adding a New Page	99
	Navigation between Pages.....	99
	Passing Data between Pages.....	101
	Sharing Objects between Pages	103
4.6	Using ViewModel classes	105

	Design with a ViewModel Class	105
	Creating a ViewModel Class	106
	ViewModels and Testing.....	108
	Page Navigation using the GoBack method	108
	Observable Collections.....	108
	Saving Data	110
	Observable Collections and Efficiency	110
	What We Have Learned.....	110
5	Isolated Storage on Windows Phone	112
5.1	Storing Data on Windows Phone.....	112
	Using the Isolated Storage File System	112
	Using the Isolated Storage Settings Storage.....	115
	The Isolated Storage Explorer	117
	What We Have Learned.....	117
6	Using Databases on Windows Phone	118
6.1	An Overview of Database Storage	118
	Databases and Queries.....	119
	Connecting to a Database	119
	Using LINQ to connect Databases to Objects	120
6.2	Creating Data Relationships with LINQ.....	129
	LINQ Associations	130
	LINQ Queries and Joining.....	135
	Joining two tables with a query	136
	Deleting Items from a Database	137
	What We Have Learned.....	137
7	Networking with Windows Phone	139
7.1	Networking Overview	139
	Starting with the Signal	139
	Building Up to Packets	141
	Addressing Packets.....	141
	Routing	142
	Networks and Media.....	143
	Networks and Protocols.....	144
	Finding Addresses on the Internet using DNS	144
	Networks and Ports.....	145
	Connections and Datagrams	146
7.2	Creating a User Datagram Protocol (UDP) Connection	146
	The Socket class	147
	Sockets and the SocketEventArgs class	148
	Handling the response from the Socket connection.....	149
	The Complete Method	153
	Sending a Message	153
	Receiving a Message	154
	Sending a datagram to the Echo service	156
	Setting up an Echo Server	156
7.3	Creating a Transmission Control Protocol (TCP) Connection	157
	Reading a Web Page.....	157
7.4	Connecting to a Data Source	163
	Using the WebClient class.....	163
7.5	Using LINQ to Read from an XML Stream	165
7.6	Using Network Services	171
	What We Have Learned.....	177
8	XNA on Windows Phone	178

8.1	XNA in context	178
	2D and 3D Games	178
	XNA and Silverlight.....	178
8.2	Making an XNA program.....	179
	How an XNA Game Runs	180
	Game Content.....	181
	Creating XNA Sprites.....	182
	Drawing Objects.....	182
	Updating Gameplay.....	184
	Adding Paddles to Make a Bat and Ball Game	186
	Controlling Paddles Using the Touch Screen	188
	Displaying Text	189
8.3	Player interaction in games.....	190
	Getting Readings from the Accelerometer Class.....	191
	Using the Accelerometer Values	191
	Threads and Contention.....	193
8.4	Adding sound to a game	194
	Creating Sounds	194
	Sound Samples and Content	195
	Using the SoundEffectInstance Class	196
8.5	Playing Sound in a Silverlight Program	197
	Loading the XNA Namespace	197
	Adding the Sound Resources.....	197
	Loading a SoundEffect in a Silverlight program	198
	Playing the sound	198
8.6	Managing screen dimensions and orientation.....	199
	Selecting a Screen Size.....	199
	Using the Full Screen	200
	Disabling the Screen Timeout	200
8.7	Combining XNA and Silverlight	200
	Creating an XNA Game Environment in Silverlight.....	201
	XNA and Silverlight together.....	202
	What We Have Learned.....	204

9 Creating Windows Phone Applications 206

9.1	The Windows Phone Icons and Splash Screens	206
	Splash Screens	208
	Silverlight Splash Screens	208
	XNA Splash Screens	208
9.2	Fast Application Switching	209
	Task Navigation in Windows Phone	209
	Understanding Fast Application Switching	210
	The Windows Phone Application LifeCycle.....	211
	Fast Application Switching in an application	213
	Fast Application Switching and Development	218
	Fast Application Switching and Design	219
9.3	Launchers and Choosers.....	219
	Using a Launcher.....	220
	Using a Chooser	222
9.4	Background Processing	223
	Background and Periodic Scheduled Tasks.....	224
	Adding a Background Task to Captains Log.....	225
	File Transfer Tasks	231
	Scheduled Notifications.....	231
	Audio Playback Agent.....	232
	What We Have Learned.....	232

10 Windows Phone Marketplace 234

10.1	Preparing an Application for Sale	234
------	---	-----

	Performance Analysis.....	234
	Creating a XAP File for Application Distribution.....	235
	Creating Application Tiles and Artwork	237
	Testing Your Application	237
	Program Obfuscation.....	238
10.2	Distributing Windows Phone Applications and Games.....	239
	Obtaining Windows Phone Applications.....	239
	Creating Windows Phone Applications and Games	239
10.3	Making your Application Stand Out	243
	Design to Sell	243
	Target Different Localisations	244
	Use App Connect.....	244
	Give your Program Away.....	244
	Release Upgrades/Episodes.....	244
	Change Categories.....	244
	Encourage Good Feedback.....	244
10.4	What To Do Next	245
	Register as a Developer	245
	Get the Toolkit.....	245
	Publish Something.....	245
	Make Suggestions.....	245
	Resources	245
Program Ideas		246
	Application Ideas.....	246
	Game Ideas	246
	Fun with Windows Phone	247
What We Have Learned.....		247

Welcome to the wonderful world of Windows Phone development. If you have half as much fun reading this book as I've had writing it, then I've had twice as much fun as you. Which doesn't seem fair really.

Anyhow, I hope you find the content useful and enjoyable.

Rob Miles, October 2011
www.robmiles.com

1 Windows Phone

In this chapter you are going to find out about the Windows Phone platform as a device for running programs. You will learn the key features of the platform itself, how programs are written and also how you can sell your programs via the Windows Marketplace.

1.1 The Windows Phone Platform

In this section we are going to take a look at the actual hardware that makes up a Windows Phone. This is particularly important as we need to put the abilities of the phone into context and identify the effect of the physical limitations imposed by platform that it uses.

A Windows Phone as a Computer

Pretty much everything these days is a computer. Mobile phones are no exception. When you get to the level of the Windows Phone device it is reasonable to think of it as a computer that can make telephone calls rather than a phone that can run programs.

The Windows Phone device has many properties in common with a “proper” computer. It has a powerful processor, local storage, fast 3D graphics and plenty of memory. It also has its own operating system which looks after the device and controls the programs that run on it. If you have used a PC you are used to the Windows operating system which starts running when you turn the computer and even turns the computer off for you when you have finished.

The Windows Phone 7 series is a complete break with previous versions of Windows Mobile devices. You could write programs and run them on earlier versions but you did not use the Silverlight or XNA environments to do this. The number 7 in the name of the product reflects the fact that this is also the 7th incarnation of the Windows Mobile platform. It does not mean that the device shares its underpinnings with desktop PCs running Windows 7. However, as we shall see, it is perfectly possible to take a program you have created for Windows Phone and run it on the Windows desktop, and vice versa.

If you are familiar with computer specifications, then the specifications below are pretty impressive for a portable device. If you are not familiar, then just bear in mind that nobody in the world had a computer like this a few years ago, and now you can carry one around in your pocket.

The Windows Phone Hardware

Before we start programming we can take a look at the hardware we will be working with. This is not a text about computer hardware, but it is worth putting some of the phone hardware into context. All Windows Phones must have a particular minimum specification, so these are the very least you can expect to find on a device.

It is quite likely that different phone manufacturers will add their particular “take” on the platform, so you will find devices with more memory, faster processors, hardware keyboards and larger screens.

Note that a hardware keyboard is not guaranteed to be present on the device (it might be just a touchscreen based phone) and that if you are an XNA game developer you will be wondering where the gamepad has gone. There are some changes to the hardware that you will have to get used to when writing for this platform, but there are also some very interesting input options (particular for game development) where you

can use the accelerometer and the touch screen to good effect. We will look at these later in Chapter 8.

The Windows Phone Processor

The Central Processing Unit (CPU) of a computer is the place where all the work gets done. Whenever a program is running the CPU is in charge of fetching data from memory, changing the data and then putting it back (which is really all computers do). The most popular speed measure in a computer is the *clock speed*. A CPU has a clock that ticks when it is running. At each clock tick the processor will do one part of an operation, perhaps fetch an instruction from memory, perform a calculation and so forth.

The faster the clock speed the faster the computer. Modern desktop computers have clocks that tick at around 3 GHz (that is around 3 thousand million times a second). This is actually incredibly fast. It means that a single clock tick will last a nanosecond. A nanosecond is the time that light takes to travel around 30 cm. If you were wondering why we don't have big computers any more, it is because the time it takes signals to travel around a circuit is a serious limiting factor in performance. Making a computer smaller actually makes it go faster.

A Windows Phone has a clock that ticks at around 1GHz. You might think that this means a Windows Phone will run around a third the speed of a PC, but this turns out not to be the case. This is because of a number of things:

Firstly, clock speed is not directly comparable between processors. The processor in the Windows PC might take five clock ticks to do something that the Windows Phone processor needs ten ticks to perform. The Windows PC processor might be able to do things in hardware (for example floating point arithmetic) which the Windows Phone processor might need to call a software subroutine to perform, which will be much slower. You can regard clock speed as a bit like engine size in cars. A car with a bigger engine might go faster than one with a smaller one, but lots of other factors (weight of car, gearbox, tires) are important too.

Secondly, a Windows PC may well have multiple processors. This doesn't mean a Windows PC can go faster, any more than two motorcycles can go faster than one, but it does mean they can process more data in a given time (two motorcycles can carry twice as many people as one). At some point we will get multiple-processor phones (and the Windows Phone operating system can support this), but at the moment they all have a single processor unit.

Finally, a Windows PC has unlimited mains power. It can run the CPU at full speed all the time if it needs to. The only real problem with doing this is that the processor must be kept cool so that it doesn't melt. The faster a processor runs the more power it consumes. If the phone ran the processor at full speed all the time the battery life would be very short. The phone operating system will speed up and slow down the processor depending on what it needs to do at any given instant. Although the phone has a fast processor this speed is only actually used when the phone has something to do which requires very fast response.

The result of these considerations is that when you are writing a Windows Phone program you cannot regard processing power as an unlimited resource. Desktop PC programmers do not see processor speed as much of an issue but Windows Phone programmers have to remember that poor design can have consequences, both in terms of the experience of the user and the battery life of the phone. The good news for us is that worrying about these things will cause us to turn into better programmers.

The Windows Phone Operating System

The operating system in a Windows Phone is called Windows CE (CE stands for "Compact Edition"). It was specially designed to run on portable computer systems and is very good at getting performance and good battery life out of a device. As we shall

see later this puts some constraints on your programs, however the good news is that as far as we are concerned the underlying operating system is pretty much irrelevant. Our program will run on the Windows Phone in pretty much the same way as they do on the full sized Windows PC.

Graphical Display

The Windows Phone has a high resolution display made up of a very large number of pixels. This provides good quality graphics and also allows lots of text to be displayed on the screen. The more pixels you have on your screen the higher the quality of image that you can display. However, the more pixels you have the greater the amount of memory that you need to store an image, and the more work the computer has to do to change the picture on the screen. This is particularly significant on a mobile device, where more work for the hardware translates to greater power consumption and lower battery life. The display resolution is a compromise between battery life, cost to manufacture and brightness of the display (the smaller the pixels the less light each can give out).

The current versions of Windows Phone have a screen resolution of at least 800x480 pixels. This can be used in both landscape (800 wide and 480 high) and portrait (480 wide by 800 high) modes. The phone contains an accelerometer that detects how the phone is being held. The Windows Phone operating system can then adjust the display to match the orientation. Our programs can decide when they run what orientations they can support. If we design our programs to work in both landscape and portrait mode they can be sent messages to allow them to adjust their display when the user changes the orientation of the device.

One problem faced by phone developers is the multitude of different screen sizes that are available. A program would usually need to be customised for each different sized screen. The Windows Phone screen hardware includes a feature that allows it to scale the screen of an application to fit whatever screen size the device supports. A game can specify that it must have a particular sized screen (say 320x240) and then the display hardware will scale that size to fit whatever physical size is fitted on the device being used. This is very useful and makes it possible to create games that will work on any device including ones with screen sizes that have not been made yet.

The Windows Phone Graphical Processor Unit

In the very first computers all the work was performed by the computer processor itself. This work included putting images on the display. Hardware engineers soon discovered that they could get faster moving images by creating custom devices to drive the screen. A Graphical Processor Unit (GPU) is given commands by the main processor and takes away all the work involved in drawing the screen. More advanced graphical processors have 3D support and are able to do the floating point and matrix arithmetic needed for three dimensions. They also contain pixel shaders which can be programmed to perform image processing on each dot of the screen at high speed as it is drawn, adding things such as lighting effects and blur.

Until quite recently only desktop PC systems and video game consoles had graphical processors, but they are now appearing in mobile phones. The Windows Phone platform contains a graphics processing chip which is used to provide 3D animation effects for the phone display and can also be used from within the XNA game development environment to create fast moving 3D games.

Touch input

Older portable devices used *resistive* touch input screens. When the user touches a resistive touch screen the plastic surface bends and makes a connection with the layer below it. A simple circuit then measures the electrical resistance to the point of contact and uses this to determine where on the screen the touch took place. Resistive touch screens are cheap to make and work very well with a stylus. However the way they

work makes it very difficult to detect multiple simultaneous touches on the screen. It is also difficult to make a resistive screen out of a very hard material, for example glass, as the screen must bend to make the contact that detects the input point.

A capacitive touch screen works in a different way. An array of conductors underneath the screen surface detects the change in capacitance caused by the presence of a finger on the surface. The touch screen hardware then works out where on the screen that the touch input took place. Capacitive touch screens are more expensive to make as they require extra hardware to process the input signals but the sensor circuits can be printed on the back of a glass screen to make a much tougher display. A capacitive touch screen is not quite as precise as a resistive screen which uses a stylus, but you can build a capacitive screen that can detect multiple inputs on different parts of the display.

All Windows Phone devices have touch input which is capable of tracking at least four input points. This means that if you create a Windows Phone piano program it will be able to detect at least four notes being pressed at the same time.

The move to multi-touch input is an important step in the evolution of mobile devices. The user can control software by using multi-touch *gestures* such as “pinch”. The Windows Phone operating system provides built in support for gesture recognition. Your programs can receive events when the user performs particular types of gesture.

Location Sensors

The Windows Phone device is location aware. It contains a Global Positioning System (GPS) device that receives satellite signals to determine the position of the phone to within a few feet. Since the GPS system only works properly when the phone has a clear view of the sky the phone will also use other techniques to determine position, including the location of the nearest cell phone tower and/or the location of the WIFI connection in use. This is called “assisted” GPS.

The Windows Phone operating system exposes methods our programs can call to determine the physical position of the device, along with the level of confidence and resolution of the result supplied. Our programs can also make use of mapping and search facilities which can be driven by location information. We will be exploring this later. The development environment also provides a really useful GPS emulator, so that you can test how your programs use position information without ever leaving your desk.

Accelerometer

The accelerometer is a hardware device that measures acceleration, so no surprises there. You could use it to compare the performance of sports cars if you wish. You can also use the accelerometer to detect when the phone is being shaken or waved around but the thing I like best is that you can also use it to detect how the phone is being held. This is because the accelerometer detects the acceleration due to gravity on the earth. This gives a phone a constant acceleration value which always points downwards. Programs can get the orientation of the phone in three axes, X, Y and Z.

This means we can write programs that respond when the user tips the phone in a particular direction, providing a very useful control mechanism for games. It is very easy for programs to obtain accelerometer readings, as we shall see later.

Compass

The phone also contains an electronic compass so that a program can determine which way the phone is facing. This can be useful in applications that are helping the user find their way around and in “augmented reality” where the phone overlays a computer drawn information on an image of the surroundings.

Gyroscope

A mechanical gyroscope is a device that always points in a particular direction. The Windows Phone contains an electronic version which allows it to detect when the phone is twisted and rotated. Programs can use the accelerometer to get an idea of how the phone is being held but a gyroscope gives much more precise information and can also give information about the rate of turn.

Sensor Integration and Simulation

It is great to have lots of sensors in the phone but there is a downside which is that the poor programmers (that's you and me) have create software that makes good use of all the different signals they produce. It can also be difficult to test how programs respond to a particular set of movements and actions. Fortunately the Windows Phone provides a single software object which integrates the information from the various sensors and provides a simple way a program can find out about the orientation and movement of the phone. This will make a best effort, depending on the sensors in the actual device. It is also possible to create particular actions and then have these replayed by the Windows Phone emulator when you are testing a program.

Camera

All mobile devices have cameras these days, and the Windows Phone is no exception. A phone camera will have at least 5 mega pixels or more. This means that it will take pictures that contain 5 million dots. This equates to a reasonable resolution digital camera (or the best anyone could get around five years ago). A 5 megapixel picture can be printed on a 7"x5" print with extremely good quality.

A Windows Phone application can take control of the camera and use it to take individual pictures or streams of video. It can also have access to the video stream itself, so that it can add things to the display to produce "augmented reality" applications or detect specific items such as barcodes or faces in photographs.

Pictures that are taken by the user can be stored as part of the media content on the phone. Our programs can also open these images and work with them.

Hardware buttons

All Windows Phone systems share a common user interface. As part of this design there are a number of physical buttons which are fitted to every Windows Phone that will always perform the same function, irrespective of the make or model of the phone.

- Start:** The start button is pressed to start a new activity. Pressing the start button will always take you to the program start screen, where you can select a new program and begin to run it. When you press the Start button this causes the currently running application to be paused (we will discuss this more a bit later). However, the Windows Phone operating system "remembers" which application was stopped so that you can return to it later by pressing the Back button.
- Back:** The back button is pressed to move back to the previous menu in a particular application. It is also used to stop one application and return to the one that was previously running. The back button makes the phone very easy to work with. You can start a new application (for example send an email message in the middle of browsing the web) and then once the message has been sent you press Back to return to the browser. Within the mail application the Back button will move you in and out of the various input screens and menus used to create and send the message. Once the message has been sent you can press Back at the top menu of the email application and return to the start menu and either run another program or press Back to return to browsing.

The Back button is also used to activate task switching on the phone. If you press and hold it down (the long press) the phone will present a display of all the active applications which you can move between.

- Lock: Pressing the lock button will always lock the phone and turn off the display to save the phone battery. When the user presses the lock or start button again the phone will display the lock screen. A phone can be configured to automatically lock after a number of minutes of inactivity. It is possible to for an application to ask to be allowed to continue running “under” the lock screen. This is useful for applications such as navigation, where the program needs to remain active, but may have implications for battery life.
- Search: Pressing the search button will start a new search. Precisely what happens when search is pressed depends on what the user is doing at the time. If the user presses search during a browsing session they will see a web search menu. If they press Search when the “People” application is active they can search for contacts.
- Camera: If the user presses the camera button this will stop the currently executing program and start the camera application to take a new picture.

The way these buttons will be used has implications for the programs that we write. A program must be able to cope with the possibility that it will be removed from memory at any time, for example if the users decides to take a photograph while playing a game the game program will be stopped and may be removed from memory. Once the picture has been taken the user should be able to resume the game just it was left. The user should not notice that the game was stopped.

Programs that are about to be stopped are given a warning message and the Windows Phone operating system provides a number of ways that a program can store state information. We will explore how to do this later in the text.

Not all Windows Phone devices will have a physical keyboard for entering text but all devices will be able to use the touch screen for text entry.

Memory and Storage

Memory is one of the things that computer owners have been known to brag about. Apparently the more memory a computer has the “better” it is. Memory actually comes in two flavors. There is the space in the computer itself where programs run and then there is the “mass storage” space that is used to store programs and data on the device.

On a desktop computer these are expressed as the amount of RAM (Random Access Memory) and the amount of hard disk space. A modern desktop computer will probably have around 2 gigabytes (two thousand megabytes) of RAM and around 500 gigabytes of hard disk storage. A megabyte is a million bytes (1,000,000). A gigabyte is a thousand million bytes (1,000,000,000). As a rough guide, a compressed music track uses around six megabytes, a high quality picture around three megabytes and an hour of good quality video will occupy around a gigabyte.

The minimum specification Windows Phone has at least 256 megabytes of RAM and 8 gigabytes of storage. This means that a base specification Windows Phone will have an eighth the amount of memory and around a fiftieth of the amount of storage of a desktop machine. The Windows Phone operating system has been optimized to work in small amounts of memory. However, some of the things it does to make sure that the device doesn’t run out of resources have an impact on how we write programs which we need to be aware of.

Network Connectivity

A mobile phone is actually the most connected device you can get. The one device has a range of networking facilities:

WiFi: All Windows Phones support wireless networking. WiFi provides a high speed networking connection, but only works if you are quite near to a network access point. Fortunately these are appearing all over the place now, with many fast food outlets and even city centers being WiFi enabled.

3G: The next thing down from WiFi in performance is the 3G (Third Generation) mobile phone network. The speed can approach WiFi, but is much more variable. 3G access may also be restricted in terms of the amount of data a mobile device is allowed to transfer, and there may be charges to use this connectivity.

GPRS: In many areas 3G network coverage is not available. The phone will then fall back to the GPRS mobile network, which is much slower.

The networking abilities are exposed as TCP/IP connections (we will explore what this means later in the text). Unfortunately, as you will have undoubtedly experienced, network coverage is not universal, and so software running on the phone must be able to keep going even when there is no data connection. Ideally the software should also cope with different connectivity speeds.

The Windows Phone also provides support for Bluetooth networking. This is used in the context of connecting external devices such as headsets and car audio systems and is not something our programs are able to access in the present version of Windows Phone.

Platform Challenges

The Windows Phone hardware is very powerful for a mobile device, but it is still constrained by the way that it must be portable and battery power is limited. To make matters worse, users who are used to working with high speed devices with rich graphical user interfaces expect their phone to give them the same experience.

As software developers our job is to provide as much of that experience as possible, within the constraints of the environment provided by the phone. When we write programs for this device we need to make sure that they work in an efficient way that makes the most of the available platform. I actually see this as a good thing, I am never happy when programmers are allowed to get away with writing inefficient code just because the underlying hardware is very powerful.

Whenever we do something in we should consider the implications when running on a phone platform. I hope you will see that one effect of this is to make us all into much better programmers.

The Good News

The last few sections read like a list of the bad things, limitations and compromises that surround mobile development. While you have to remember all these issues, it is also the case that writing for mobile devices is actually great fun. The range of features that the device provides and the fact that it is portable make it possible for you to create genuinely new applications that have never been done before.

Also, the tools that we will use to write the program and the built in features provided with the tools make it possible to create really nice looking applications very easily.

1.2 The Windows Phone Ecosystem

The Windows Phone is not designed as a device that stands by itself. It is actually part of an *ecosystem* which contains a number of other software systems that work around it to provide the user experience.

The Zune Media Management Software

A Windows Phone device is connected to a Windows PC by means of the Zune software. This provides a way of managing media and transferring it to and from the phone device. The Zune software is also the means by which the firmware of the Windows Phone can be updated as new versions come along. The Zune software also provides the link between the Visual Studio development environment and the phone itself.

The programs that we write can make use of the media loaded onto the phone via the Zune software. It is very easy to write programs that load pictures or playback music and videos stored on the phone. It is also possible to write music playback programs that run in the background, behind other programs that the user may be running on the phone itself.

Windows Live and Xbox Live

The owner of a Windows Phone can register the phone to their Windows Live account. If their Windows Live account is linked to their Xbox Live gamertag their gamer profile is imported onto their device and so they can participate in handheld gaming using the same identity as they use on their console.

A Windows Phone program can use the Windows Live identity of the owner and XNA gamertag information.

Bing Maps

We know that a phone contains hardware that allows it to be aware of its location. The Microsoft Bing maps service provides maps which cover most of the earth. A Windows Phone program can use its network connection to contact the Bing maps service and download maps and aerial views. There is a Silverlight mapping control that makes it very easy for a program to add location displays.

Windows Notification Service

Although a phone has a wide range of networking technologies available there will still be occasions when the phone is unable to connect to a network. Windows Phone provides a notification service which allows programs to get notifications from the network even when the program is not active. The user is given a notification on the phone that they can use to start the application. Notifications are stored and managed by a notification service which will buffer them if the phone is not connected to the network when the notification is raised.

For example, you might have a sales system that needs to tell a customer when an item they have reserved is in stock. The ordering application on the phone can register with the notification service and then your stock management server can use the notification server to tell the customer when the item becomes available. This system can also be used in gaming, where one player wants to send a challenge to another.

Windows Phone and Windows Azure

Windows Azure is a set of “cloud” services that are made available by Microsoft. Rather than using your own computer to provide data storage and processing for internet users, you can instead rent space and processors in the cloud. If your networked application becomes very popular you don’t have to buy more computers and connect them up, instead you just scale up your use of the cloud resource.

Server Applications

Cloud services are also very useful if you have to do a large amount of processing once, or very infrequently. Rather than having a large number of computers to satisfy

the peak demands, instead you can just offload the work into the cloud. You can write the server applications in C# and test them on the cloud simulator before uploading them to do the work. Client applications can make use of these services via a variety of different network mechanisms.

Database Storage

While a Windows Phone can have a database inside the phone, it is often useful to use a remote database which is accessed via the network. This makes it easier to have really huge amounts of storage available, which you might want to update centrally. It is possible to create databases in the cloud and use them from a Windows Phone application.

BLOB Storage

Blob stands for *Binary Large Object*. It is perhaps my favourite name in this entire book. A blob can be anything, the text of a book, a movie or output from a medical scan. Windows Azure provides a storage facility where an application can store and retrieve blobs.

Authentication Services

Whenever a user wishes to get access to a resource there is always a problem of authentication. The resource provider will want to be sure that the user is who they say they are. This is a particular problem with portable (and stealable) devices such as phones. Writing high quality authentication software is very difficult, and not something you should undertake lightly. Windows Azure provides a way of taking the authentication service into the cloud.

The Windows Azure Toolkit

The Windows Phone Azure toolkit provides project templates, class libraries and sample code that can get you started writing cloud applications using the phone. You can find it here:

watoolkitwp7.codeplex.com

Using the Ecosystem

It is important to remember that a phone is not just a phone these days. The connected nature of the application means that it will function in conjunction with remote services which are available over the network connection. Other parts of the phone make use of this connected ability too. The phone has a Facebook client built in and the camera application and picture storage elements can upload your pictures to Windows Live Skydrive or Facebook. There are methods that make it easy to post updates to live feeds and even share favourites.

You can also use the network features of the phone to create your own client and server applications. You will need to write two programs to do this, one will run on the Windows Phone and the other can run on any computer connected to the internet, or even on a system in the “cloud”.

1.3 Windows Phone program execution

The Windows phone provides a platform to run programs, some of which can be ones that we have written. It is worth spending some time considering how the programs are made to run on the phone and some of the implications for the way that we write them.

Application Switching on Windows Phone

Windows phone was designed on the basis that the user is king (or queen). All of the design decisions have been made to make sure that whatever the user is doing at that time, the phone is giving as much support to that activity as possible.

This means that when a user is running your application the phone is not running any other. If there are any other applications active on the phone they are “standing in the wings” waiting for their turn on stage.

Just like in a theatre show, where it is important that the performers can get on and off the stage quickly between scenes, it is important that it is easy for a user to switch between active applications. We have already seen that the Back button provides a quick way for a user to move from one program to another, when we create our applications we have to make sure the user has a good experience when they do this.

One implication for us developers is that we must live with the possibility that at any point our program might get “kicked off the stage” to make way for another. Then, if the user ever calls the application back from the wings, it should make sure that the show continues at exactly the same place.

The Windows phone operating system provides “Fast Application Switching” where an application is kept in memory while it waits in the wings for the call back on stage. In this situation the program has all its memory and resources intact and just needs to start running again. However, just like on a stage, where the space in the wings is limited, sometimes an application will be removed from memory (or sent back to its dressing room) if there is no longer room for it. The important point to remember though is that from the user’s point of view the experience they get when they return to the application should be *exactly the same*. Just like it would take a performer longer to get from their dressing room back to the stage, an application that needs to be reloaded into memory will take a bit longer to return, but the experience the user gets when the application comes back should be exactly the same. A program is given a message when it is about to be moved “off stage” and it has access to storage that it can use to retain state information. We will see how to do this in Chapter 9.

Background Processing

The limitations imposed by the processor and battery life considerations make it hard for a device as small as a phone to run multiple processes. Windows PC users are familiar with a desktop cluttered with many different programs running simultaneously, but on a phone this is not practical, not least because the display is not large enough to show multiple applications.

The Windows Phone allows applications to contain “background tasks”. A background task can take over when the main application is not able to run. They are allowed to run in carefully managed circumstances, so that if the phone has several background tasks active the user will not notice the phone running more slowly. Background tasks are designed for use in particular scenarios including background music playback, file transfer, regular updates and occasional bulk data processing. Background tasks are invisible to the phone user, although a user can see them running and control which ones are active. A background task can alert the user to an event by sending a notification or by updating the display on a “Live Tile” on the screen. The user can then start the application to find out more. We will see to create background tasks in Chapter 9.

Windows Phone and Managed Code

In the early days a computer ran a program just by, well, running a program. The file containing the program instructions was loaded into memory and then the computer just obeyed each one in turn. This worked OK, but this simple approach is not without its problems.

The first problem is that if you have a different kind of computer system you have to have a different file of instructions. Each computer manufacturer made hardware that understood a particular set of binary instructions, and it was not possible to take a program created for one hardware type and then run it on another.

The second problem is that if the instructions are stupid/dangerous the computer will obey them anyway. Stupid instructions in the file that is loaded into memory may cause the actual hardware in the computer to get stuck (a bit like asking me to do something in Chinese). Dangerous instructions could cause the program to damage other data in the computer (a bit like asking me to delete all the files on my computer).

The Microsoft Intermediate Language (MSIL)

Microsoft .NET addresses these problems by using an intermediate language to describe what a program wants to do. When a C# program is compiled the compiler will produce a file containing instructions in this intermediate language. When the program actually runs these instructions are compiled again, this time into the low level instructions that are understood by the target hardware in use. During the compilation process the instructions are checked to make sure they don't do anything stupid and when the program runs it is closely monitored to make sure that it doesn't do anything dangerous.

Microsoft .NET programs are made up of individual components called *assemblies*. An assembly contains MSIL code along with any data resources that the code needs, for example images, sounds and so on. An assembly can either be an executable (something you can run) or a library (something that provides resources for an application). A Windows Phone can run C# executable assemblies and libraries produced from any .NET compatible compiler. This means that you could write some of the library code in your application in another language, for example Visual Basic, C++ or F#. It also means that if you have libraries already written in these languages you can use them in phone applications.

The idea behind .NET was to provide a single framework for executing code which was independent of any particular computer hardware or programming language. The standards for .NET specify what the intermediate language looks like; the form of the data types stored in the system and also includes the designs of the languages C# and Visual Basic .NET.

Just in Time Compilation

When a program actually gets to run something has to convert the MSIL (which is independent of any particular computer hardware) into machine code instructions that the computer processor can actually execute. This compilation process is called *Just In Time* compilation because the actual machine code for the target device is compiled from the intermediate instructions just before it gets to run. The way in which the program runs within a monitored environment is called *managed code*. The compilation happens in the instant before the program runs, i.e. the user of the phone selects a program from the start page, the MSIL is loaded from storage and then compiled at that time.

The downside of this approach is that rather than just run a file of instructions the computer now has to do more work to get things going. It must load the intermediate language, compile it into machine code and then keep an eye on what the program itself does as it runs. Fortunately modern processors (including the one inside the Windows Phone) can achieve this without slowing things down.

The upside is that the same intermediate code file can be loaded and executed on any machine that has the .NET system on it. You can (and I have) run exactly the same compiled program on a mobile device, Windows PC and Xbox 360, even though the machines have completely different operating systems and underlying hardware.

Another advantage is that this loading process can be expanded to ensure that programs are legitimate. .NET provides mechanisms by which programs can be "signed" using cryptographic techniques that make it very difficult for naughty people to modify the

program or create “fake” versions of your code. If you become a marketplace developer you will be assigned your own cryptographic key that will be used to sign any applications that you create.

Finally, the use of an intermediate language means that we can use a wide range of programming languages. Although the programming tools for Windows Phone are focused on Visual Basic and C# it is possible to use compiled code from any programming language which has a .NET compiler. If you have existing programs in C++ or even F# you can use intermediate code libraries from these programs in your Windows Phone solutions. Remember though that you will need to use C# for the Silverlight or XNA “front end” for these programs.

Managed Code

Programs that you create on Windows Phone run within the “managed” environment provided by the operating system. You will write your C# programs and then they will run inside a safe area in the phone provided by the operating system. Your program will not be allowed direct access to the hardware. This is good news because it stops people from writing programs that stop the proper operation of the phone itself. As far as the developer is concerned this actually makes no difference. When a program calls a method to put a picture on the screen the underlying systems will cause the picture to appear.

Developer Implications

As far as a developer is concerned there is an element of good news/bad news about all this. The good news is that you only need to learn one programming language (C#) to develop on a range of platforms. Programs that you write will be isolated from potential damage by other programs on a system and you can be sure that software you sell can’t be easily tampered with.

The bad news is that all this comes at the expense of extra work for the system that runs your application. Starting up a program is not just a case of loading it into memory and obeying the instructions it contains. The program must be checked to ensure that it has not been tampered with and then “just in time” compiled into the instructions that will really do the work. The result of all this is that the user might have to wait a little while between selecting the application and getting the first screen displayed.

Fortunately the power of Windows Phone means that these delays are not usually a problem but it does mean that if we have very large programs we might want to break them into smaller chunks, so that not everything is loaded at once. But then again, as sensible developers we would probably want to do this kind of thing anyway.

1.4 Windows Phone application development

You write Windows Phone applications in exactly the same way as you write other applications for the Windows desktop. You use the Visual Studio IDE (Integrated Development Environment). You can debug a program running in a Windows Phone device just as easily as you can debug a program on your PC desktop. You can also create solutions that share components across the desktop, Windows Phone and even Xbox platforms.

You can take all your Windows Desktop development skills in Silverlight and your console skills in XNA and use them on the phone. If you learn how to use the Windows Phone you are also learning how to write code for the desktop (Silverlight) or console (XNA). This is great news for you as it means that you can write (and even sell) programs for Windows Phone without having to learn a lot of new stuff. If you have previously written programs for desktop computers then the move to Windows Phone development will be a lot less painful than you might expect.

The Windows Phone Emulator

The Windows Phone development environment is supplied with an emulator which gives you a Windows Phone you can play with on your PC desktop. If you have a PC system that supports multi-touch input you can even use this with the emulator to test the use of multi-touch gestures with your Windows Phone programs.

While the emulator is feature complete, in that it behaves exactly like a real phone would in response to the requests from your software, it does not mimic the performance of the actual phone hardware. Programs running on the emulator are using the power of your PC, which may well be much greater than the processor in the phone. This means that although you can test the functionality of your programs using the emulator you only really get a feel for how fast the program runs, and what the user experience feels like, when you run your program on a real device.

The emulator will allow you to test motion controlled programs by allowing you to tip and turn a “virtual” phone on your PC screen. The inputs to the program will reflect the orientation that you see on your screen. You can also pre-record a set of movements to allow you to perform repeatable tests using particular movements or gestures.

There is also provision for location emulation. You can select on a map where you want the emulated phone to “be”. You can also input paths for the phone to follow and then replay these, so that you can create journeys to test location aware applications.

Accessing Windows Phone Facilities

The Windows Phone platform provides a library of software resources that let your programs make use of the features provided by the device itself. Your programs can make use of the camera in the phone, place phone calls and even send SMS messages. They can also make use of the GPS resources of the phone to allow you to create location aware applications. The facilities are divided into *Launchers* which let your program transfer to another application and *Choosers* which use the system to select an item and then return control to your program. We will see how to use these later in this text.

Windows Phone Connectivity

As you might expect, programs on a Windows Phone are extremely well connected. Applications on the phone can make use of the TCP/IP protocol to connect to servers on the internet. Programs can use Web Services and also set up REST based sessions with hosts. If you are not sure what any of this means, don't worry. We will be exploring the use of network services later in the text.

Silverlight and XNA Development

There are essentially two flavours to Windows Phone development. If you are creating an application (for example a word processor, email client or cheese calculator) then you can use Silverlight. This provides a whole range of features just for creating such programs. If you are creating a game then you can use XNA. XNA provides all the facilities for creating 2D and 3D games with high performance graphics.

You are not forced to work this way. You could create your cheese calculator in XNA or (perhaps more sensibly) you can create simple games (for example word puzzles) very successfully in Silverlight.

You select the type of application you are creating when you make a new project with Visual Studio.

Combining Silverlight and XNA

A single application can contain both Silverlight and XNA behaviours. You can think of this as a Silverlight program with a page that contains an XNA game. This makes it really easy to create the menus and user interface for the game in Silverlight, and put the gameplay into XNA. You can even add Silverlight controls on top of a running XNA game. We will explore how to do this in Chapter 8.

Data Storage on Windows Phone

Useful applications need to store data. Some of this data will be simple things, for example configuration data such as the size of the display. Other items will be larger; perhaps high scores or the place the player has reached in the game. Then there may be a need for even more information that must be held in a structured way, for example a database of customers or products. The Windows Phone provides appropriate storage at all these levels.

Isolated Storage

Isolated storage is so named because the storage space for one application is isolated from all the others. When a program is installed on the Windows Phone it is given access to a storage space that is completely isolated from all others. There is storage can provide simple name/value pairs (for example `ColourScheme:Purple`) for things like program settings. Applications can also create complete file hierarchies within their patch of isolated storage, depending on their needs.

A program called the *Isolated Storage Explorer* can be used to view the folders and files in isolated storage when you are developing an application so that you can make sure they contain the correct data.

Local Database

A database is a collection of data which is managed by a program that is also, rather confusingly, called a database. Databases are used to create programs that store and manipulate huge amounts of data and make it very easy for our applications to provide answers to questions like “Tell me which of our customers from New York bought a pair of shoes in November”. The Windows Phone has a database built-in. This makes it very easy to create an application that stores structured data.

Each Windows Phone application also has access to a single SQL database where they can store very large amounts of structured data. An application can create queries using the LINQ (Language INtegrated Query) libraries to extract and modify data in the database. The actual database file is held in the Isolated Storage for the application. We will see how to use databases and LINQ in Chapter 6.

Development Tools

The tools that you need to get started are free. You can download a copy of the Windows Phone SDK and be started writing Windows Phone applications literally within minutes. Developers who have more advanced, paid for, copies of Visual Studio 2010 can use all the extra features of their versions in mobile device development by adding the Windows Phone SDK plugin to their systems. You can download a copy of the Windows Phone SDK from create.msdn.com

Performance Profiling

Applications for a Windows PC are not usually constrained by the cpu performance. The sheer power of the processors in in desktops and laptops means that quite inefficient code will still be useable. However, for the Windows Phone it is important to make our programs as speedy and efficient as possible, both to give the best user experience and also to make sure that our programs don't use too much battery power.

The Windows Phone SDK includes a performance profiler that we can use to view the demands our applications place on the processor, and also find out where programs are spending most of their time. This makes it possible to determine which methods in a program should be optimised to get the best improvements.

Windows Marketplace

Windows Marketplace is where you can sell programs that you have created. It is moderated and managed by Microsoft to ensure that applications meet certain minimum standards of behaviour. If you want to sell an application you must be a registered developer and submit your program to an approvals process.

It costs \$99 to register as a developer, but students can register for free via the Dreamspark programme. Developers can register their Windows Phone devices as “development” devices. Visual Studio can deploy applications to a developer device and you can even step through programs running in the device itself.

You can distribute free applications as well as paid ones. A registered developer can submit up to 100 free applications for approval in any year. If you want to distribute more than 100 free applications you must pay a \$20 fee for each additional free application. You can also create applications that have a “demo” mode which can be converted into a “full” version when the user purchases them. Your program can tell when it runs whether it is operating in full or demo mode.

When you submit an application to be sold in the marketplace it will go through an approvals process to make sure that it doesn’t do anything naughty and that it works in a way that users expect. If the approvals process fails you will be given a diagnostic report. Before you send an application into the system you should read the certification guidelines of the marketplace.

Testing Tool

The Windows Phone SDK contains a testing tool which can be used to pre-test an application before it is submitted to the marketplace. This tool automatically synchronises with the marketplace approval process, so if the process changes the tool will be updated to match. Some of the testing is automatic, but the tool also describes each of the usage tests that will be performed by the testing team. This greatly improves the chances of us making an application that will be approved first time.

Private Betas

One great way to find out the quality of an application or game is to give it to a bunch of people and see what they say. We can create “invitation only” releases of a product so that up to 100 people can be sent a download link for our new program. This allows the testers up to 90 days of use of the program, during which time they can give us their feedback.

What We Have Learned

1. Windows Phone is a powerful computing platform.
2. All Windows Phone devices have a core specification. This includes a particular size of display, capacitive touch input that can track at least four points, Global Positioning System support, 3D graphics acceleration, high resolution camera and ample memory for program and data storage.
3. The Windows Phone device is connected to a Windows PC by means of the Zune PC software, which provides a media management system for the PC and also allows media to be synchronised to the phone.
4. Windows Phone systems can make use of network based services to receive notifications, determine their position and perform searches.

5. When developing programs for Windows Phone the Zune software is used to transfer programs into the phone for testing. The Zune software is also used to upgrade the firmware in the phone.
6. The Windows Phone operating system supports full multi-tasking, but to for the best performance only one user application can be active at one time. However, applications can create their own “background agents” that can perform specific tasks when the application is not running. Windows Phone also provides “Fast Application Switching” that retains applications in memory so that they can be resumed very quickly.
7. The Windows Phone runs programs that have been compiled into Microsoft Intermediate Language (MSIL). This MSIL is compiled inside the phone just before the program runs. The programs themselves run in a “managed” environment which stops them from interfering with the operating of the phone itself.
8. When developing software for Windows Phone you can create Silverlight and XNA applications. These are written in C# using Visual Studio 2010. Programmers can use a Windows Phone emulator that runs on Windows PC and provides a simulation of the Windows Phone environment.
9. Programs have access to all the phone facilities and can place calls, send SMS messages etc.
10. The Windows Phone SDK can be used to create Windows Phone applications. It is a free download from create.msdn.com However, to deploy applications to a phone device you must be a registered Windows Phone Developer. This costs \$99 per year but registration is free to students via the Microsoft Dreamspark initiative. A registered developer can upload their applications to Windows Phone Marketplace for sale.

2 Introduction to Silverlight

A user interface is a posh name for the thing that people actually see when they use your program. It comprises the buttons, text fields, labels and pictures that the user actually works with to get their job done. Part of the job of the programmer is to create this “front end” and then put the appropriate behaviours behind the screen display to allow the user to drive the program and get what they want from it. In this section we are going to find out about Silverlight, and how to use it to create a user interface for our programs.

2.1 Program Design with Silverlight

It turns out that most programmers are not that good at designing attractive user interfaces (although I’m sure that you are). In real life a company will employ graphic designers who will create artistic looking front ends. The role of the programmer will then be to put the code behind these displays to get the required job done. Silverlight recognises this process by enforcing a very strong separation between the screen display design and the code that is controlled by it. This makes it easy for a programmer to create an initial user interface which is subsequently changed by a designer into a much more attractive one. It is also possible for a programmer to take a complete user interface design and then fit the behaviours behind each of the display components.

Development Tools

The user interface designer uses the Expression Blend tool and the programmer uses the Visual Studio tool to create the code. The Silverlight system provides a really easy way to combine user interface designs with code. A good way to work is for the programmer to use a “placeholder” design to create the program behaviours and then incorporate the final designs later in the production process.

The Windows Phone SDK (Software Development Kit) includes versions of Visual Studio and Expression Blend that can be used in this way.

The Metro Design Style

From a design point of view a nice thing about Windows Phone is that it brings with it a whole set of design guidelines which are referred to as “Metro”. This sets out how controls are supposed to look and establishes a set of criteria that your applications should meet if they are to be “good looking” Windows Phone applications. This style regime is carried over into the built in components that are provided for you to use in your programs. The happy consequence of this is that if you use the Buttons, TextFields and other components that are supplied with Visual Studio you will be automatically adhering to the style guidelines.

This is a great help those of us who are not very good at design because it means that we are only ever going to use things that look right. Of course you can completely override the properties of the supplied components if you really want purple text on an orange background but I would not really advise this.

There is actually a Metro style document you can read if you really want to know how to make “proper looking” programs. This is well worth a look if you want to publish your programs in the Marketplace, where people will have certain expectations of how things should look and work. You can find the style guide here:

msdn.microsoft.com/en-us/library/hh202915.aspx

For the purpose of this course we are going to use the Silverlight design tools in Visual Studio. These do not give all the graphical richness that you can achieve with Expression Blend, but for programmers they are more than adequate. This should make sure that our applications adhere to the Metro guidelines and are clean and simple to use.

Silverlight Elements and Objects

From a programming point of view each of the elements on the screen of a display is actually a software object. We have come across these before. An object is a lump of behaviours and data. If we were creating an application to look after bank accounts we could create a class to hold account information:

```
public class Account
{
    private decimal balance ;
    private string name ;

    public string GetName ()
    {
        return name;
    }

    public bool SetName (string newName){
    {
        // Final version will validate the name
        name = newName;
        return true;
    }

    // Other get and set methods here
}
}
```

This class holds the amount of money the account holder has (in the data member called `balance`) and the name of the account holder (in the data member called `name`). If I want to make a new `Account` instance I can use the `new` keyword:

```
Account rob = new Account();
rob.SetName("Rob");
```

This makes a new `Account` instance which is referred to by the reference `rob`. It then sets the name member of this `Account` to "Rob". We are used to creating objects which match things we want to store. When we make games we could invent a `Sprite` class to hold the picture of the sprite on the screen, the position of the sprite and other information. Objects are a great way to represent things we want to work with. It turns out that objects are also great for representing other things too, such as items on a display. If you think about it, a box displaying text on a screen will have properties such as the position on the screen, the colour of the text, the text itself and so on.

Consider the following:



This is a very simple Windows Phone program that I've called an "Adding Machine". You can use to perform very simple sums. You just enter two numbers into the two text boxes at the top and then press the equals button. It then rewards you with the sum of these two numbers. At the moment it is showing us that 0 plus 0 is 0. Each individual item on the screen is called a *UIElement* or User Interface element. I'm going to call these things elements from now on. There are seven of them on the screen above:


1. The small title "Adding Machine". This is known in the Windows Phone style guidelines as the 'Application Title'.
2. The larger title "Add". This is known as the 'Page Title'.
3. The top textbox, where I can type a number.
4. A text item holding the character +.
5. The bottom textbox, where I can type another number.
6. A button, which we press to perform the sum.
7. A result textbox, which changes to show the result when the button is pressed.

Each of these items has a particular position on the screen, particular size of text and lots of other properties too. We can change the colour of the text in a text box, whether it is aligned to the left, right or centre of the enclosing box and lots of other things too.

There are three different types of element on the screen:

1. **TextBox** – allows the user to enter text into the program.
2. **TextBlock** – a block of text that just conveys information.
3. **Button** – something we can press to cause events in our program.

If you think about it, you can break the properties of each of these elements down into two kinds, those that all the elements need to have, for example position on the screen, and those that are specific to that type of element. For example only a **TextBox** needs to record the position of the cursor where text is being entered. From a software design point of view this is a really good use for a class hierarchy.



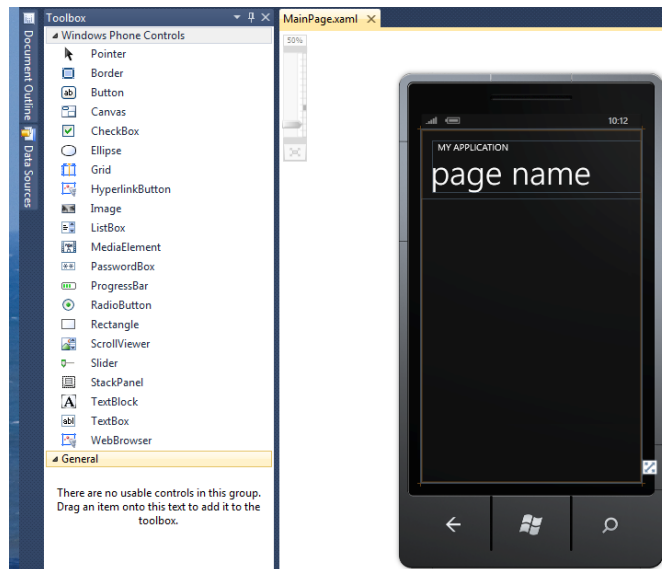
Above you can see part of the hierarchy that the Silverlight designers built. The top class is called `FrameworkElement`. It contains all the information that is common to all controls on the screen. Each of the other classes is a child of this class. Children pick up all the behaviours and properties of their parent class and then add some of their own. Actually the design is a bit more complex than shown above, the `FrameworkElement` class is a child of a class called `UIElement`, but it shows the principles behind the controls.

Creating a class hierarchy like this has a lot of advantages. If we want a custom kind of textbox we can extend the `TextBox` class and add the properties and behaviours that we need. As far as the Silverlight system is concerned it can treat all the controls the same and then ask each control to draw itself in a manner appropriate to that component.

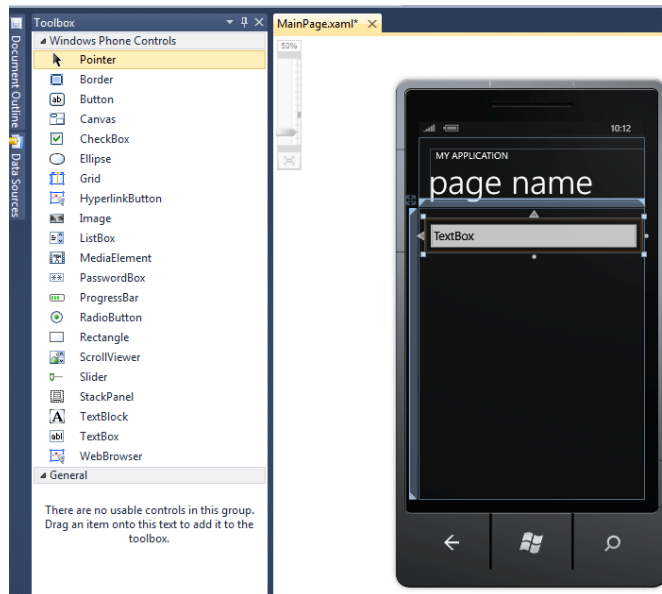
So, remember that when we are adjusting things on a display page and creating and manipulating controls we are really just changing the properties of objects, just as we would change the name and balance values of a bank account object. When we design a Silverlight user interface we set the data inside the display elements to position them on the display. Next we are going to find out how to do this.

The Toolbox and Design Surface

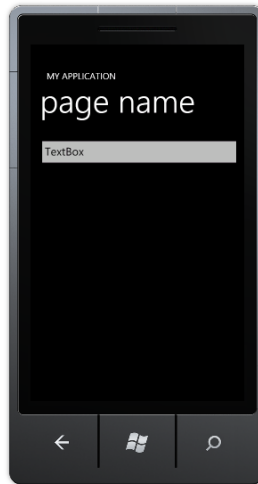
We could start by investigating how the Adding Machine above was created. It turns out to be really easy (we will take a really detailed look at Visual Studio in the next section). When we create a brand new Silverlight project we get an empty page and we can open up a `ToolBox` which contains all the controls that we might want to add to the page:



We can assemble the user interface by simply dragging the controls out of the toolbox onto the design surface.



Above shows the effect of dragging a `TextBox` out of the Toolbox area and dropping it onto the Windows Phone page. If we run the program now we can see our textbox on the screen:



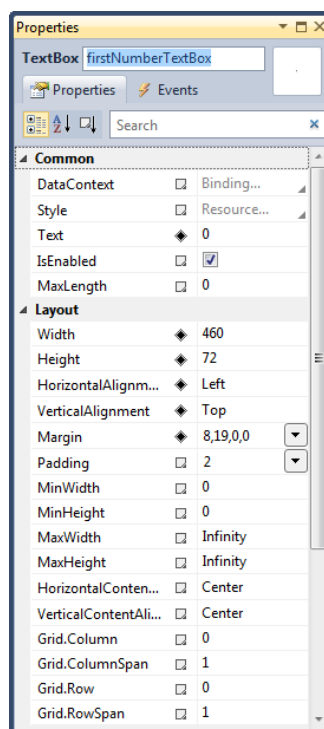
This is the Windows Phone emulator running our program, which is showing the textbox on the screen.

If you have ever done any Windows Forms development (perhaps using earlier versions of .NET) then this will all be very familiar to you. If you haven't then it is a quick and easy way to build up a user interface. Because all the controls are already styled in the Metro style, just like the Windows Phone user interface, you are automatically building an application that looks like a "real" one. The designer has lots of nice features that make it very easy to align your controls too, which is nice.

Managing Element Names in Visual Studio

Once we have put some elements on the screen we need to set the names of them to sensible values. The designer will give each element a meaningless name like "TextBox1", which I like to change to something which has meaning in the context of the system being created.

We can change the properties of things on the screen by clicking on the item in the designer and then seeking out the Properties pane in Visual Studio for that item.



Above you can see the properties window for the top `TextBox` on the page. The name of an element is given at the very top of the window.

I have changed the name of this textbox to `firstNumberTextBox`. You'll never guess what the second text box is called. Note that the name of a property in this context is actually going to set the name of the variable declared inside the adding machine program. In other words, as a result of what I've done above there will now be the following statement in my program somewhere:

```
TextBox firstNumberTextBox;
```

Visual Studio looks after the declaration of the actual C# variables that represent the display elements that we create and so we don't need to actually worry about where the above statement is. We just have to remember that this is how the program works.

Properties in Silverlight Elements

Once we have given our `TextBox` variable a proper name we can move on to give it all the properties that are required for this application. We can also change lots of properties for the textbox including the width of the textbox, the margin (which sets the position) and so on. The values that you see in the properties windows above are ones that reflect the current position and size of the item on the screen. If I drag the item around the margin values will change. If I change the values in the window above I will see the item move in the design surface. Again, if you are a Windows Forms kind of person you will see nothing strange here. But you must remember that all we are doing is changing the content of an object. The content of the properties window will change depending on what item you select

Silverlight properties and C# properties

When we talk about the "properties" of Silverlight elements on the page (for example the text displayed in a `TextBox`) we are actually talking about property values in the class that implements the `TextBox`. In other words, when a program contains a statement such as:

```
resultTextBlock.Text = "0";
```

- this will cause a `Set` method to run inside the `resultTextBlock` object which sets the text on the `TextBlock` to the appropriate value. At this point it is worth refreshing our understanding of properties in C# so that can get a good understanding of what is happening here.

C# classes and properties

C# classes can contain member data (the name of the person owning a particular bank account) and behaviours (methods called `GetName` and `SetName` that let a program find out this name and change it). Properties provide a way of representing data inside a class that is easier to work with than using `Get` and `Set` methods. They are used extensively in managing Silverlight elements, and therefore worth a refresher at this point.

A property is a member of a class that holds a value. We have seen that we can use a member variable to do this kind of thing, but we need to make the member value public. For example, the bank may ask us to keep track of staff members. One of the items that they may want to hold is the age of the bank account holder. We could do it like this:

```
public class Account
{
    public int Age;
}
```

The class contains a public member. I can get hold of this member in the usual way:

```
Account s = new Account ();
s.Age = 21;
```

I can access a public member of a class directly; just by giving the name of the member. The problem is that we have already decided that this is a bad way to manage our objects. There is nothing to stop things like:

```
s.Age = -100210232;
```

This is very naughty, but because the Age member is **public** we cannot stop it.

Creating Get and Set methods

To get control and do useful things we can create get and set methods which are **public**. These provide access to the member in a managed way. We then make the Age member **private** and nobody can tamper with it:

```
public class Account
{
    private int age;
    public int GetAge()
    {
        return this.age;
    }
    public void SetAge( int inAge )
    {
        if ( (inAge > 0) && (inAge < 120) )
        {
            this.age = inAge;
        }
    }
}
```

We now have complete control over our property, but we have had to write lots of extra code. Programmers who want to work with the age value now have to call methods:

```
Account s = new Account ();
s.SetAge(21);
Console.WriteLine ( "Age is : " + s.GetAge() );
```

Using Properties

Properties are a way of making the management of data like this slightly easier. An age property for the StaffMember class would be created as follows:

```
public class Account
{
    private int ageValue;

    public int Age
    {
        set
        {
            if ( (value > 0) && (value < 120) )
                ageValue = value;
        }
        get
        {
            return ageValue;
        }
    }
}
```

The age value has now been created as a property. Note how there are get and set parts to the property. These equate directly to the bodies of the get and set methods that I wrote earlier. The really nice thing about properties is that they are used just as the class member was:

```
Account s = new Account ();
s.Age = 21;
Console.WriteLine ( "Age is : " + s.Age );
```

When the `Age` property is given a value the `set` code is run. The keyword `value` means “the thing that is being assigned”. When the `Age` property is being read the `get` code is run. This gives us all the advantages of the methods, but they are much easier to use and create.

Data validation in properties

If we set an age value to an invalid one (for example we try to set the age to 150) the set behaviour above will perform validation and reject this value (nobody over 120 is allowed to have an account with our bank), leaving the age as it was before. A program that is using our `Age` property has no way of knowing whether an assignment to the property has failed unless it actually checks that the assignment worked.

```
Account s = new Account ();
int newAge = 150;
s.Age = newAge;
if (s.Age != newAge )
    Console.WriteLine ( "Age was not set" );
```

The code above sets the age to 150, which is not allowed. However, the code then tests to see if the value has been set. In other words, it is up to the users of your properties to make sure that values have been set correctly. In the case of a `Set` method the method itself could return false to indicate that the set has failed, here the user of the property has to do a bit more work.

Multiple ways of reading a property

Properties let us do other clever things too:

```
public int AgeInMonths
{
    get
    {
        return this.ageValue*12;
    }
}
```

This is a new property, called `AgeInMonths`. It can only be read, since it does not provide a set behaviour. However, it returns the age in months, based on the same source value as was used by the other property. This means that you can provide several different ways of recovering the same value. You can also provide read-only properties by leaving out the set behaviour. Write only properties are also possible if you leave out the get.

Properties and notifications

You may be asking the question “Why do we use properties in Silverlight elements?” It makes sense to use them in a bank account where I want to be able to protect the data inside my objects but in a Silverlight program, where I can put any text I like inside a `TextBlock`, there seems to be no point in having validation of the incoming value. In fact running this code will slow down the setting process. So, by making the `Text` value a public string we could make the program smaller and faster, which has to be a good thing. Right?

Well, sort of. Except that when we change the text on a `TextBlock` we would like the text on the Silverlight page in front of the user to change as well. This is how our

adding machine will display the result. If a program just changed the value of a data member there would be no way the Silverlight system could know that the message on the screen needs to be updated.

However, if the `Text` member is a property when a program updates it the set behaviour in the `TextBlock` will get to run. The code in the set behaviour can update the stored value of the text field and it can also trigger an update of the display to make the new value visible. Properties provide a means by which an object can get control when a value inside the object is being changed, and this is extremely important. The simple statement:

```
resultTextBlock.Text = "0";
```

- may cause many hundreds of C# statements to run as the new value is stored in the `TextBlock` and an update of the display is triggered to make the screen change.

Page Design with Silverlight

We can complete the page of our Adding machine by dragging more elements onto the screen and setting their properties. Remember that the first thing I do after dragging an item onto the page is set the name of that element. One of the first symptoms of a badly written program (for me) is a page that contains lots of elements called “Button1” and “Button2”.

From the sequence above I hope you can see that designing a page looks quite easy. You just have to drag the elements onto the page and then set them up by using the properties of each one. If you read up on the Silverlight language you will find that you can give elements graphical properties that can make them transparent, add images to their backgrounds and even animate them around the screen. At this point we have moved well beyond programming and entered the realm of graphic design. And I wish you the best of luck.

2.2 Understanding XAML

In the above section we saw that Silverlight design can be broken down into some element manipulation using the Visual Studio designer and some property setting on the elements that you create. Once you know how to do these things, you can create every user interface you need.

It is perfectly true that you can work this way, and you could even think about making every program you ever need without touching XAML, although I think you would be missing out on some of the most powerful features of Silverlight.

At the moment everything seems quite sensible and we are all quite happy (at least I hope so). So this would seem a good point to confuse everybody and start to talk about XAML. The letters stand for “Extensible Application Markup Language” which I’m sure makes everything clear. Or not. XAML is the language used by Silverlight to describe what a page should look like.

Why have XAML?

You might be wondering why we have XAML. If you have used Windows Forms before you might think that you managed perfectly well without this new language getting in the way. XAML provides a well-defined interface between the look of an application (the properties of the display elements that the user sees) and the behaviour of the application (what happens behind the display to make the application work).

The split is a good idea because, as we have mentioned before, there is no guarantee that a good programmer will be a good graphic designer. If you want to make good looking applications you really have to have both designers and programmers working on your solution. This leads to a problem, in that ideally you want to separate the work of the two, and you don’t want the programmer unable to do anything until the graphical designer has finished creating all the menu screens.

XAML solves both these problems. As soon as the requirements of the user interface have been set out (how many text boxes, how many buttons, what is there on each page etc) then the designer can work on how each should look, while the programmer can focus on what they do. The XAML files contain the description of the user interface components and the designer can work on the appearance and position of these elements to their heart's content while the programmer gets on with making the code that goes behind them. There is no need for the programmer to give the designer their code files since they are separate from the design, and vice versa.

XAML file content

A XAML file for a given page on the display will contain constructions that describe all the elements on it. Each description will contain a set of named properties. The line of XAML that describes our first `TextBox` is as given below:

```
<TextBox Height="72" HorizontalAlignment="Left" Margin="8,19,0,0"
Name="firstNumberTextBox" Text="" VerticalAlignment="Top"
Width="460" TextAlignment="Center" />
```

If you compare the information in the Properties pane with the values in the XAML above you will find that they all line up exactly. The name of the `TextBox` is `firstNumberTextBox`, the width is `460` and so on.

If you were wondering how the page and the properties pane kept themselves synchronised it is because they both used the XAML description file for the page. When I move things around the page the editor updates the XAML with the new positions. The properties window reads this file and updates its values accordingly.

XAML is described as a *declarative* language. This means it just tells us about stuff. It is designed to be understandable by humans, which is why all the properties above have sensible names. If you want to, you can edit the text contents of the XAML files within Visual Studio and change both the appearance in the designer and also the property values.

XAML turns out to be very useful. Once you get the hang of the information used to describe components it turns out to be much quicker to add things to a page and move them about by just editing the text in the XAML file, rather than dragging things around or moving between property values. I find it particularly useful when I want a large number of similar elements on the screen. Visual Studio is aware of the syntax used to describe each type of element and will provide Intellisense support help as you go along.

When an application is built the XAML file is converted into a file containing low level instructions that create the actual display components when the program runs. This means that declaring a control inside the XAML file for a page will mean that the control will exist as an object our program can use.

XAML looks a lot like XML, an eXtensible Markup language you may have heard of. The way that items and elements are expressed is based on XML syntax and it is quite easy to understand.

Extensible Markup Languages

At this point you may be wondering what an eXtensible Markup Language actually is. Well, it is a markup language that is extensible. I'm sure that helped. What we mean by this is that you can use the rules of the language to create constructions that describe anything. English is a lot like this. We have letters and punctuation which are the symbols of English text. We also have rules (called grammar) that set out how to make up words and sentences and we have different kinds of words. We have nouns that describe things and verbs that describe actions. When something new comes along we invent a whole new set of word to describe them. Someone had to come up with the word "computer" when the computer was invented, along with phrases like "boot up", "crash" and "too slow".

XML based languages are extensible in that we can invent new words and phrases that fit within the rules of the language and use these new constructions to describe anything we like. They are called *markup* languages because they are often used to describe the arrangement of items on a page. The word markup was originally used in printing when you wanted to say things like “Print the name Rob Miles in very large font”. The most famous markup language is probably HTML, HyperText Markup Language, which is used by the World Wide Web to describe the format of web pages.

Programmers frequently invent their own data storage formats using XML. As an example, a snippet of XML that describes a set of high scores might look as follows:

```
<?xml version="1.0" encoding="us-ascii" ?>
<HighScoreRecords count="2">
  <HighScore game="Breakout">
    <playername>Rob Miles</playername>
    <score>1500</score>
  </HighScore>
  <HighScore game="Space Invaders">
    <playername>Rob Miles</playername>
    <score>4500</score>
  </HighScore>
</HighScoreRecords>
```

This is a tiny XML file that describes some high score records for a video game system. The `HighScoreRecords` element contains two `HighScore` items, one for `Breakout` and one for `Space Invaders`. The two high score items are contained within the `HighScoreRecords` item. Each of the items has a property which gives the name of the game and also contains two further elements, the name of the player and the score that was achieved. This is quite easy for us to understand. From the text above it is not hard to work out the high score for the `Space Invaders` game.

The line at the very top of the file tells whatever wants to read the file the version of the XML standard it is based on and the encoding of the characters in the file itself. XAML takes the rules of an extensible markup language and uses them to create a language that describes components on a page of a display.

```
<TextBox Height="72" HorizontalAlignment="Left" Margin="8,19,0,0"
Name="firstNumberTextBox" Text="0" VerticalAlignment="Top"
Width="460" TextAlignment="Center" />
```

If we now take a look at the description of a `TextBox` I think we can see that the designers of XAML have just created field names that match their requirements.

XML Schema

The XML standard also contains descriptions how to create a *schema* which describes a particular document format. For example, in the above the schema for the high score information would say that a `HighScore` must contain a `PlayerName` and a `Score` property. It might also say things like the `HighScore` can contain a `Date` value (the date when the high score was achieved) but that this is not required for every `HighScore` value.

This system of standard format and schema means that it is very easy for developers to create data formats for particular purposes. This is helped by the fact that there are a lot of design tools to help create documents and schemas. The .NET framework even provides a way by which a program can save an object as a formatted XML document. In fact the Visual Studio solution file is actually stored as an XML document.

As far as we are concerned, it is worth remembering that XML is useful for this kind of thing, but at the moment I want to keep the focus on XAML itself.

XAML and pages

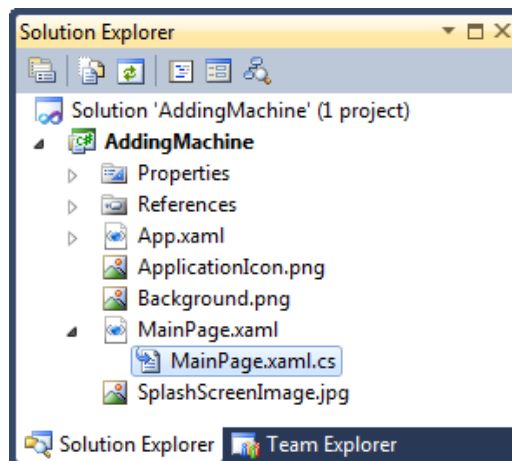
A file of XAML can describe a complete page of the Windows Phone display. When you create a brand new Windows Phone project you get a page that just contains a few elements. As you put more onto the page the file grows as each description is added. Some elements work as *containers* which means that they can hold other components. They are very useful when you want to lay things out, for example there is a Grid element which can hold a set of other elements in a grid arrangement. The XAML code can also contain the descriptions of animations and transitions that can be applied to items on the page to make even more impressive user interfaces.

One thing you need to bear in mind at this point is that XAML has no concept of hierarchy. It is not possible to create an XAML element which is derived from another. We do this in software so that we can create a new object which is similar to an existing one. We know that Silverlight elements are represented in software terms by an object which is part of a class hierarchy. However, when they are written in a XAML file they are all expressed as items which are at the same level.

We are not going to spend too much time on the layout aspects of XAML; suffice it to say that you can create incredibly impressive front ends for your programs using this tool. There is also a special design tool called “Expression Blend” for use by graphical designers. However, it is important to remember that at the end of the day the program is working in terms of objects that expose properties and methods that we can use to work on the data inside them. In the case of our user interface elements, if we change the properties of the element the display the user sees will change to reflect this. The objects represent the items on the screen of our program.

2.3 Creating a Silverlight Application

Now that we know that items on the screen are in fact the graphical realisation of software objects, the next thing we need to know is how to get control of these objects and make them do useful things for us in our application. To do this we will need to add some C# program code that will perform the calculation that the adding machine needs.



Whenever Visual Studio makes a XAML file that describes a page on the Windows Phone display it also makes a program file to go with it. This is where we can put code that will make our application work. If you actually take a look in the file `MainPage.xaml.cs` above you will find that it doesn't actually contain much code:

```

using System;
using System.Collections.Generic;
using System.Linq;
using System.Net;
using System.Windows;
using System.Windows.Controls;
using System.Windows.Documents;
using System.Windows.Input;
using System.Windows.Media;
using System.Windows.Media.Animation;
using System.Windows.Shapes;
using Microsoft.Phone.Controls;

namespace AddingMachine
{
    public partial class MainPage : PhoneApplicationPage
    {
        // Constructor
        public MainPage()
        {
            InitializeComponent();
        }
    }
}

```

Most of the file is `using` statements which allow our program to make direct use of classes without having to give the fully formed name of each of them. For example, rather than saying `System.Windows.Controls.Button` we can say `Button`, because the file contains the line `using System.Windows.Controls`.

The only method in the program is the constructor of the `MainPage` class. As we know, the constructor of a class is called when an instance of the class is created.

All the constructor does is call the method `InitializeComponent`. If you go take a look inside this method you will find the code that actually creates the instances of the display elements. This code is automatically created for you by Visual Studio, based on the XAML that describes your page. It is important that you leave this call as it is and don't change the content of the method itself as this will most likely break your program.

Note that at this point we are deep in the “engine room” of Silverlight. I'm only really telling you about this so that you can understand that actually there is no magic here. If the only C# programs you have seen so far start with a call of the `Main` method then it is important that you understand that there is nothing particularly special about a Silverlight one. There is a `Main` method at the base of a Silverlight application; it starts the process of building the components and putting them on the screen for the user to interact with.

The nice thing as far as we are concerned is that we don't need to worry about how these objects are created and displayed, we can just use the high level tools or the easy to understand XAML to design and build our display.

Building the Application

We can now see how to create the user interface for our number adding program. If we add all the components and then start the application it even looks as though it might do something for us. We can type in the numbers that we want to add, and even press the equals button if we like:



If we touch inside a `TextBox` item the keyboard appears and we can type numbers into the box. If we then touch anywhere else on the screen the keyboard moves out of the way. We seem to have got a lot of behaviour for very little effort, which is nice. However, we need to add some business logic of our own now to get the program to work out the answer and display it.

Calculating the Result

At the moment our program looks good, but doesn't actually do anything. We need to create some code which will perform the required calculation and display the result. Something like this.

```
private void calculateResult()
{
    float v1 = float.Parse(firstNumberTextBox.Text);
    float v2 = float.Parse(secondNumberTextBox.Text);

    float result = v1 + v2;

    resultTextBlock.Text = result.ToString();
}
```

The `TextBox` objects expose a property called `Text`. This can be read from or written to. Setting a value into the `Text` property will change the text in the textbox. Reading the `Text` property allows our program to read what has been typed into the textbox.

The text is given as a string, which must be converted into a numeric value if our program is to do any sums. You may have seen the `Parse` method before. This takes a string and returns the number that the string describes. Each of the numeric types (`int`, `float`, `double` etc) has a `Parse` behaviour which will take a string and return the numeric value that it describes. The adding machine that we are creating can work with floating point numbers, so the method parses the text in each of the input textboxes and then calculates a result by adding them together.

Finally the method takes the number that was calculated, converts it into a string and then sets the text of the `resultTextBlock` to this string. `ToString` is the reverse of `Parse`, the “anti-parse” if you like. It provides the text that describes the contents of an object. In the case of the `float` type, this is the text that describes that value.

Now we have our code that works out the answer, we just have to find a way of getting it to run when the user presses the equals button.

Events and Programs

If you have done any kind of form based programming you will know all about events. If you haven't then don't worry, we now have a nice example of a situation where we need them, and they are not that frightening anyway. In the olden days, before graphical user interfaces and mice, a program would generally start at the beginning, run for a while and then finish.

But now we have complicated and rich user interfaces with lots of buttons and other elements for the user to interact with. The word processor I am using at the moment exposes hundreds of different features via buttons on the screen and menu items. In this situation it would be very hard to write a program which checked every user element in turn to see if the user has attempted to use that. Instead the system waits for these display elements to raise an event when they want attention. Teachers do this all the time. They don't go round the class asking each child in turn if they know the answer. Instead they ask the children to raise their hands. The "hand raising" is treated as an event which the teacher will respond to.

Using events like this makes software design a lot easier. Our program does not have to check every item on the screen to see if the user has done anything with it, instead it just binds to the events that it is interested in.

To make events work a programming language needs a way of expressing a reference to a method in an object. C# provides the *delegate* type which does exactly that. You can create a delegate type that can refer to a particular kind of method and then create instances of that delegate which refer to a method in an object. Delegates are very powerful, but can be a bit tricky to get your head around. Fortunately we don't have to worry about how delegates work just at the moment; because we can get Silverlight and Visual Studio to do all the hard work for us.

Events in Silverlight

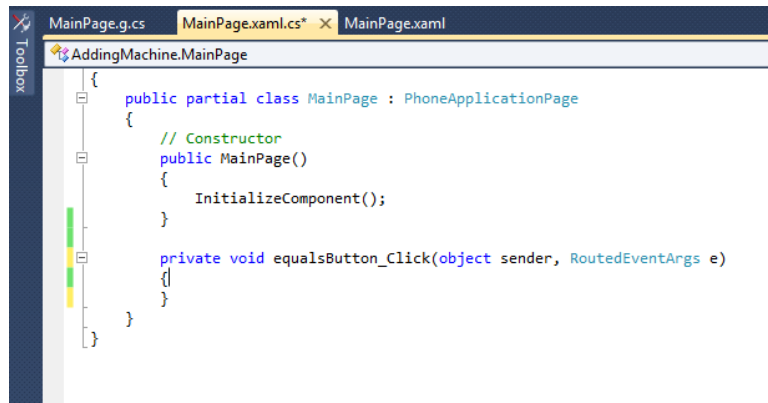
In C# an event is delivered to an object by means of a call to a method in that object. In this respect you can regard an event and a message as the same thing. From the point of view of our adding machine we would like to have a particular method called when the "equals" button is pressed by the user. This is the only event that we are interested in. When the event fires we want it to call the `calculateResult` method above.

We can use the editor in Visual Studio to bind this event for us. This is actually so easy as to appear magical. To connect a method to the click event of a particular button we just have to double click on that button on the design surface.



The picture above shows where we should double click in Visual Studio. When we double click on the button Visual Studio sets the properties of the button in XAML to connect the button to a method in our page. It also creates an empty version of that method and takes us straight to this method, so that we can start adding the code that must run when the button is clicked.

If we just single click on the button this has the effect of selecting it, so that we can move it around the display and change its size. If you want to connect an event handler it must be a double click. If we do this we will see a display like the one shown below.



Now, all we have to do is add the code to make this work.

```

private void calculateResult()
{
    float v1 = float.Parse(firstNumberTextBox.Text);
    float v2 = float.Parse(secondNumberTextBox.Text);

    float result = v1 + v2;

    resultTextBlock.Text = result.ToString();
}

private void equalsButton_Click(object sender, RoutedEventArgs e)
{
    calculateResult();
}

```

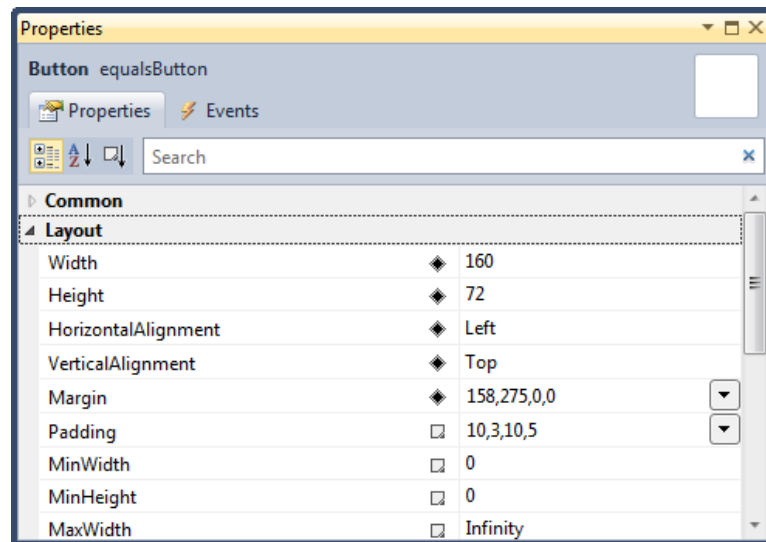
The event handler that is created is automatically given a sensible name by Visual Studio. It takes the name of the element and adds the text “_Click” on the end. This underlines the importance of giving your elements sensible names when you create them, because this name makes it very easy for anyone to understand what is happening.



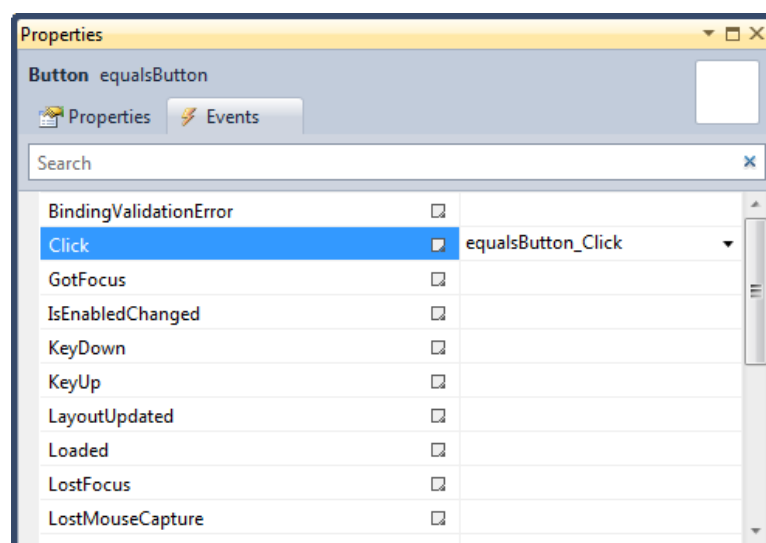
The program works fine, as you can see above. Of course it is not perfect. If a user types in text rather than a numeric value the program will fail with an exception, but we will deal with this later.

Managing Event Properties

At the moment the process of connecting event handlers to events looks a bit magic. And if there is one thing I'm sure of, it is that computers are not magical. Ever. So now we should really find out how this event process actually works. We can start by taking a look at the properties of the equals button in our application.



If we single click on the button in the editor and then take a look at the property information in Visual Studio we see something very like the display above. The information here sets out the position of the button on the screen and a bunch of other things too. The property window has two tab panels though. If we look at the top we will find one with the name Events, and a rather scary looking lightning bolt icon. If we click on the Events tab the display changes to display the events that this element can generate, and which are presently connected to code:



The Properties window is now showing us all the events that a button can generate. At the moment only one event has a method connected to it, and that is the Click event. This is linked to the `equalsButton_Click` method. If we wanted to disconnect the Click event from the method we could do this by just deleting the name of the handler from this window. We could cause problems for ourselves by replacing the content of that part of the form with something stupid like "ILikeCheese". If we do this the

program will not work (or even build) any more because the Click event is now connected to something that does not exist.

Events and XAML

Up until now we have worked on the principle that the properties display in Visual Studio is driven by the XAML file which describes the components on the screen. We can prove that this is the case by taking a look at the Button description in the XAML file for the MainPage:

```
Button Content="equals" Height="72" HorizontalAlignment="Left"
Margin="158,275,0,0" Name="equalsButton" VerticalAlignment="Top"
Width="160" Click="equalsButton_Click" />
```

As you can see; the description of Button now contains an item that connects the Click event with the name of a method.

One important point to note here is that if the XAML file says that a button is connected to a method with a particular name the program will fail to run correctly if that method is not there.

What We Have Learned

1. Silverlight provides a way of designing graphical user interfaces.
2. A Silverlight user interface is made up of elements which are things such as text boxes and buttons.
3. The Visual Studio integrated development environment provides an editor which can be used to add Silverlight user interface elements onto the pages of an application.
4. From a software point of view each of the user interface elements is represented by a particular type of object in a class hierarchy which is part of Silverlight.
5. Designers can modify the properties of the elements using the design tools in Visual Studio, either by changing them on the design surface or by modifying their properties.
6. The actual properties of Silverlight elements in a project are held in text files in the XAML format. These files are updated by the design tools and used when the program is built to create the software objects that are used in the solution.
7. XAML (eXtensible Application Markup Language) is an XML based language that specifies all the properties of the design elements on a page. It provides a separation between the appearance and properties of the display elements and the program code that sits behind them.
8. XML (eXtensible Markup Language) is a way of creating customised languages that can be used to describe things.
9. Elements can generate events which may be bound to methods inside a C# program. The method name is given in the XML description for the element.
10. The methods that are bound to the events can provide the business logic for a user application.

3 Visual Studio Solution Management

When you write programs for Windows Phone you will be using Visual Studio. In this section we will take a look at the process of creating and managing Windows Phone projects. We will also find out how to run and debug Windows Phone programs using the Windows Phone emulator program. This allows you to test your programs without needing to have a device.

However, this is not just a look at how to use Visual Studio. That would be too easy. We are also going to take a peek underneath the hood and find out how Visual Studio manages the content that makes up our solutions. This will stand us in good stead when we come to create Windows Phone applications.

3.1 Getting Started with Projects and Solutions

The first thing we are going to look at is just how Visual Studio manages the programs that you create.

Creating programs

We know that when we create a program we write a set of instructions in a high level language (which if we are lucky is C#). If you would rather work with Visual Basic .NET than C# you will be pleased to discover that you can create Windows Phone programs in that language too. This text is written for C# developers, but all the fundamental principles and the application programmer interface (API) calls can be mapped directly into the Visual Basic Language.

The *compiler* is a program that takes the high level instructions which the programmer created and converts them into lower level instructions for execution on the computer. We also know now that in Microsoft .NET the compiler produces output in an *intermediate language* (MSIL) which is then “Just In Time” compiled into hardware instructions when the program actually runs.

The bare minimum you need to create a working program is therefore a C# source file (a text file with the language extension .CS) and a compiler. The compiler takes the source file and converts it into an executable file (a binary file with the language extension .EXE) which can then be used by the .NET runtime system for whatever platform you want to execute the program on. The use of MSIL means that I can take exactly the same .EXE file and run it on completely different hardware platforms if I wish.

The .NET SDK

If you just want to create programs for the Windows PC and you don't want to download all of Visual Studio you can actually just get the compiler and runtimes for .NET from here:


<http://msdn.microsoft.com/en-us/netframework/aa569263.aspx>

When you install the framework you get a command line compiler which can take C# source files and convert them into executable ones using console commands such as:

```
csc MyProg.cs
```

This command will compile the file `MyProg.cs` and create a file called `MyProg.exe`.

The command prompt compiler is also available if you have installed Visual Studio.



```
Visual Studio Command Prompt (2010)
C:\code>type HelloWorld.cs
using System;

class Demo
{
    public static void Main ()
    {
        Console.WriteLine("Hello world");
    }
}
C:\code>esc HelloWorld.cs
Microsoft (R) Visual C# 2.0.50727.4842 Compiler version 4.0.30319.1
Copyright (C) Microsoft Corporation. All rights reserved.

C:\code>HelloWorld
Hello world
C:\code>
```

It is possible to use these simple tools to create very large programs, but it is also rather hard work. You have to create large, complicated commands for the compiler and you need to make sure that whenever you make a new version of your system that you compile every source file that you need to. Furthermore you have to manage and incorporate all the media that your system needs and when it comes to debugging your program all you are going to get is text messages when the program fails.

The Windows Phone SDK

You don't need to download the .NET SDK if you want to create Windows Phone applications using Visual Studio. For that you just need to go to:

create.msdn.com

This site contains links to the latest versions of the Windows Phone SDK and you can also register here to sell your programs in the Windows Phone Marketplace. Note that if you are a student you can get free access to the marketplace via the DreamSpark program. Dreamspark can also give you access to free copies of the Professional versions of Visual Studio.

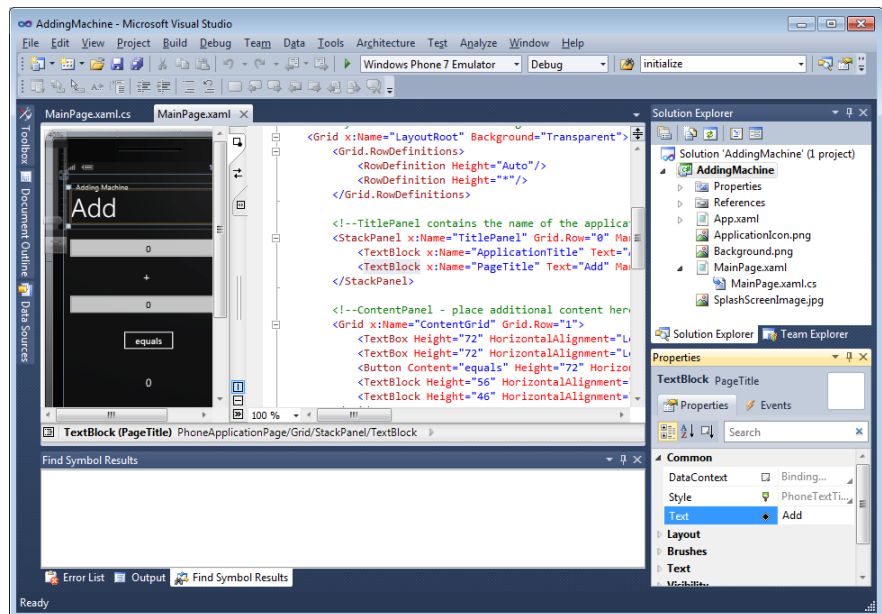
www.dreamspark.com

The Windows Phone SDK is all installed at once. Once the installer has finished you will have the Interactive Development Environment, the Windows Phone emulator and other tools to help you create Windows Phone applications. The tools can also be used to create desktop applications and XNA games for Windows PC and Xbox 360. If you have another version of Visual Studio 2010 present when you install the Windows Phone SDK it will add the new project types and the tools alongside that version.

Visual Studio

Visual Studio provides a single environment in which you can edit, build and run your programs. It manages projects and solutions which let you organise your systems and control what goes into them. Finally, just to make it even more wonderful, it provides a full debugging solution that lets you single step through your program and even execute individual statements and change the values in variables as your program runs. Skill with Visual Studio and an understanding of how it works is very marketable. Remember that this tool is used to create PC applications as well as those for the Windows Phone.

Once you have installed the Windows Phone SDK you will find Microsoft Visual Studio 2010 in your programs menu. Simply start that to begin writing programs.



The first thing that hits you with Visual Studio is the sheer number of buttons and controls that you have to deal with. If we start opening menus it gets even worse. However, the good news is that you don't have to press all the buttons to get started. We are just going to take a look at the basics of the program. You can pick up more advanced features as we go along.

The first aspect of Visual Studio we are going to take a look at is that of projects and solutions. You will use these to organise any systems that you create. Visual Studio is not able to run any program code which is not in a project or solution. The good news is that to get started you can use one of the templates that provide a starting point for particular project types. We will be using the Windows Phone templates for now. There are many other templates for different kinds of programs.

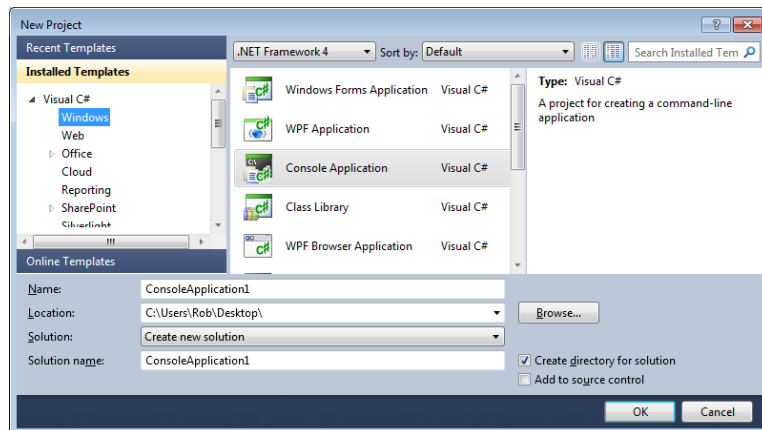
Visual Studio makes a distinction between a *project* and a *solution* when organising your work. It is important that you understand the difference between these two. Let's start by looking at a project.

A Simple Visual Studio Project

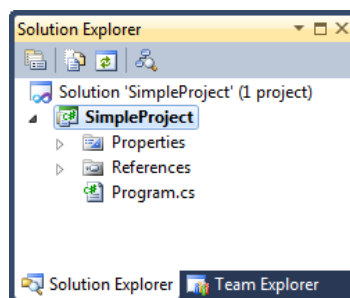
A project is a container for a set of program files and resources which generate a particular *assembly* file. So, what does that mean? Well, the simplest possible project would be one containing a single program file. You can create one of these by asking Visual Studio to create a Console Application. We can do this by using the **File>New>Project** command to open up the New Project dialogue:

The New Project dialogue shows us all the possible project templates that are supplied with Visual Studio. A template is a pre-set collection of files, folders and resources which is used for a particular type of application. We can add to the template once it has been created, and it is even possible to create customised templates of our own. The projects that we saw earlier were created using templates for creating Windows Phone applications. Note that if you install just the Windows Phone SDK you will not see these templates, to get access to these you need to install the Visual Studio Express program from

www.microsoft.com/express



If we select the template for a Windows Console Application as shown above Visual Studio will create the simplest possible project for us. This is what it would look like in Visual Studio Solution explorer.



This project contains a single program file, called `Program.cs`. The program source in this file is very simple too:

```
using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;

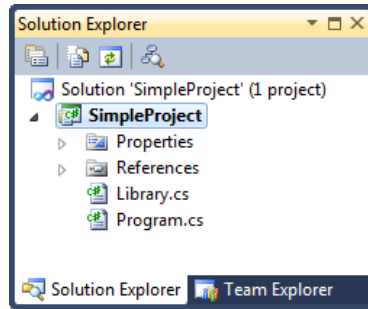
namespace SimpleProject
{
    class Program
    {
        static void Main(string[] args)
        {
        }
    }
}
```

If we ask Visual Studio to run the program it will compile the `Program` class and then run the application by calling the `Main` method. This is how C# programs work. The output of the build process on this program would be a file called `Program.exe` which can be run as a program on a Windows PC.

The program runs in the command window, which will pop up briefly onto the screen when the `Main` method runs. We will look at running and debugging programs a little later in this section.

Making a program from several source files

It is often sensible to create a program from a number of class files. These can be separate parts written by another programmer or you might use a library of C# code provided by somebody else. You can do this in Visual Studio by adding a new class to your project. The solution explorer in Visual Studio will show you the new file in the project:



The file `Library.cs` is now part of the project and when Visual Studio builds the project it knows to compile the `Library.cs` source file along with the `Program.cs` one. However, the output of this build will be a single executable file called `Program.exe`. Of course only one of these source files can have a `Main` method, otherwise the compiler will become confused about where the program is supposed to start running. The `Library` class is presently empty:

```
using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;

namespace SimpleProject
{
    class Library
    {
    }
}
```

We can create instances of the class in our `Main` method:

```
namespace SimpleProject
{
    class Program
    {
        static void Main(string[] args)
        {
            Library l = new Library();
        }
    }
}
```

Namespaces and Projects

The `Library` class and the `Program` class are in the `SimpleProject` namespace. From a design point of view it is often sensible to have different parts of a system in separate namespaces. This is particularly important if the programs are being built by different people. If the programmer writing the library code decides that the name `Save` is appropriate for a particular method you don't want that clashing with a `Save` method that you want to use.

Remember that namespaces are a *logical* way of grouping items together. They have no bearing on where the actual code is physically located. We can put the library class into a separate namespace simply by changing the name:


```

using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;

namespace SimpleLibrary
{
    class Library
    {
    }
}

```

The Library class is now in the SimpleLibrary namespace. If we try to build the project we get some errors:



This is because the Main method in the Program class is trying to create an instance of the Library class. This class is no longer in the SimpleProject namespace, and so we get the errors. Fortunately the Description field of the error list tries to help by telling us that we might be missing a using directive. This is the case. We can fix the problem by adding a using directory to tell the compiler to look in the SimpleLibrary namespace if it needs to find a class:

```

using SimpleLibrary;

namespace SimpleProject
{
    class Program
    {
        static void Main(string[] args)
        {
            Library l = new Library();
        }
    }
}

```

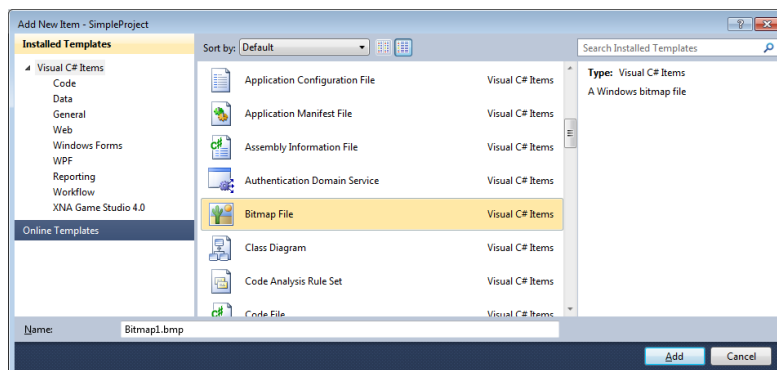
Another way to solve the problem is to use fully qualified names to access the resources in the library:

```
SimpleLibrary.Library l = new SimpleLibrary.Library();
```


The thing to remember about all this is that we are simply creating names for items in our project; we are not specifying where in the system they are stored. When Visual Studio builds SimpleProject it will look in all the source files in the project to find elements that it needs.

Adding Resources to a Project

If your program needs to use resources (perhaps an image for splash screen picture) you can add them to a project too. We can add a bitmap image in just the same way as we added the new class.



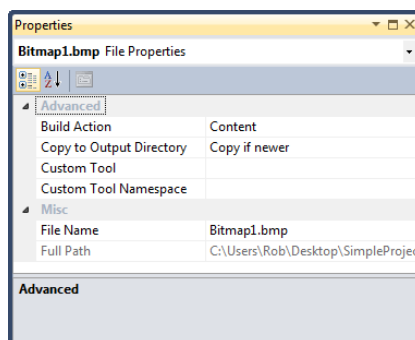
When we have added the bitmap resource it appears in the Solution Explorer for the project.



This solution now contains a bitmap file called `Bitmap1.bmp`. We can now select how Visual Studio manages this content item by modifying the properties of `Bitmap1` in the solution.

Content Properties

Like everything else that Visual Studio manages, the resources added to a project also have properties. We can work with the properties just as we work with the properties of any other item that we add to a project, we use the properties pane.



The settings above ask Visual Studio to copy the bitmap into the same directory as the program when it builds the project. The program can then open this file and use it when it runs:

```
Bitmap b = new Bitmap("Bitmap1.bmp");
```

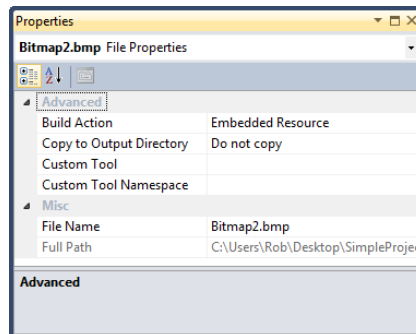
The above code would load the bitmap resource into a bitmap called `b`.

This is a very good way to add items to a project. Note that the properties above are not the default ones which are set when a new item of content is added. By default (i.e. unless you specify otherwise) the content is not copied into the output directory.

Embedding Content in the Application itself

The above settings ask that the bitmap be stored in the same directory as the program. This works well, but assumes that when we install the program the image files are copied as well as the program binary.

If we want tighter coupling between a program and its resources we can request that this resource be embedded in the program file when the program is built. When the program is built the output executable program will contain the bitmap. We do this by changing the properties of the content item.



The Build Action property of the `Bitmap2.bmp` file is now set to `Embedded Resource`. Our image is now stored in the actual program file. Getting hold of this is slightly more involved, but guaranteed to work wherever our program is deployed, because there are no extra files:

```
Assembly assembly;
Stream imageStream;

assembly = Assembly.GetExecutingAssembly();
imageStream =
    assembly.GetManifestResourceStream("SimpleProject.Bitmap2.bmp");

Bitmap b = new Bitmap(imageStream);
```

The above code will load a bitmap from an embedded resource and store it in a bitmap called `b`. Note that when we get a resource from the assembly we have to give the namespace of the project along with the resource name.

One of the design decisions we will have to make later in this text is whether or not to incorporate things like pictures in the program executable file, or as resources which are stored in the same folder as the program itself.

Assembly files and Executable files

At this point you might be starting to wonder about the output file. It seemed simple enough when we were just compiling a simple C# program, but now it contains a picture as well. This is because the output of a Visual Studio build is not just a program; it is in fact an *assembly*. An assembly is a container which can hold a whole range of things, including program code. When a program runs the .NET runtime system will open the assembly and go and find the class in the assembly with the `Main` method in it. It then starts that method running. The project file tells Visual Studio what to put in the assembly.

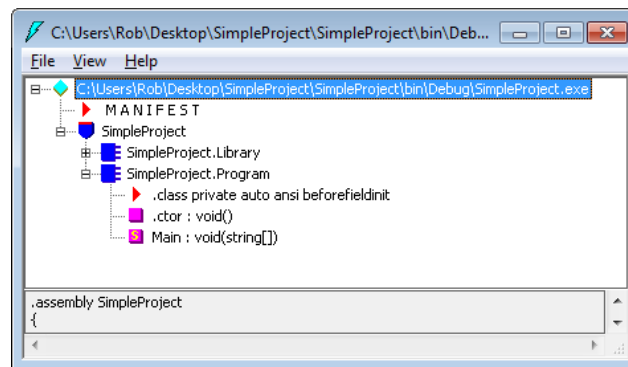
If you are very old, and remember when a file that ends in `.exe` just contained program instructions, then this is a bit confusing. The good news is that you don't need to worry; Visual Studio does a very good job of managing what goes into a program.

Using *ildasm* to look inside assembly files

If you've ever wanted to look inside an assembly file and see what makes it tick you can use the program `ildasm`. The program is hidden away in the Visual Studio installation. The best way to find it is to open your Program Files(x86) and use the `Serach` function to look for it.

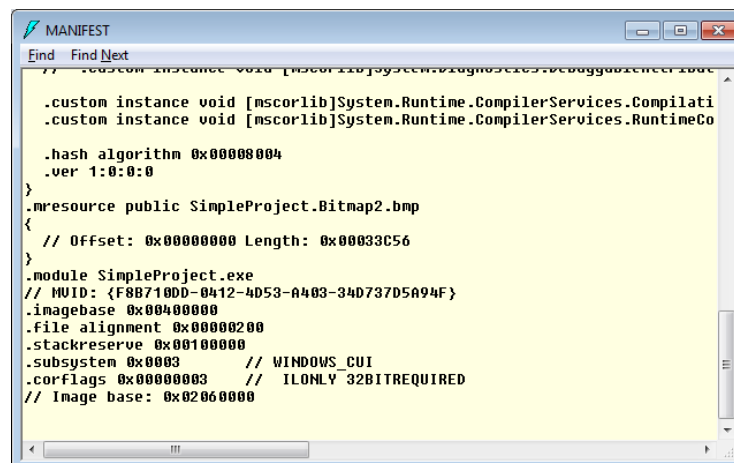
The `ildasm` program lets you can open any assembly and take a look inside. This also lets you see the MSIL (Microsoft Intermediate Language) instructions that the compiler produced.

Once the program is running you can find your way to an assembly and open it:



This is the ildasm display for the simple project I made above. The assembly contains the compiled classes that make up the program content. You can see that the `SimpleProject.Program` class contains two methods, the `ctor` method (the constructor for the class) and the `Main` method. You can open these and take a look at the code inside them.

The MANIFEST part of the assembly contains a “list of contents” for the assembly.

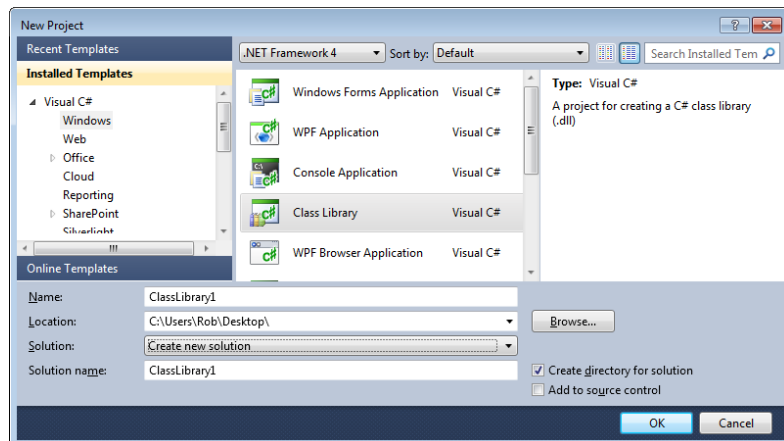


This is the manifest for an assembly that contains an embedded resource. You can see the name of the resource and the size of it.

The ildasm program can show us this information because a compiled class contains *metadata* which describes the methods and members and can even include information for Intellisense to use in Visual Studio. If you have ever wondered exactly what a program does and how it gets to run you can have a lot of fun going through the views that ildasm provides. However you might not have the same idea about what constitutes fun as I do.

Library Assemblies

There are actually two kinds of assembly; ones that contain a `Main` method somewhere (the executable ones) and libraries. When you create a new project you can create a library project if you like:



This will create a brand new project which will only contain library classes, i.e. none of the classes in the assembly will contain a `Main` method. If any of them do you will get an error when you try to compile the project. A library assembly compiles to a file with the language extension `.dll`. This stands for *Dynamic Link Library*. The library word sounds sensible enough; the dynamic link part means that the library classes are loaded dynamically when the program runs.

When the program refers to a method in a class in the library that class will be loaded into memory and the Just In Time compiler will convert the method into machine code so that it can run. If the class is never used, it is never loaded. This is a nice way of saving memory, at the expense of a little more effort when the program runs.

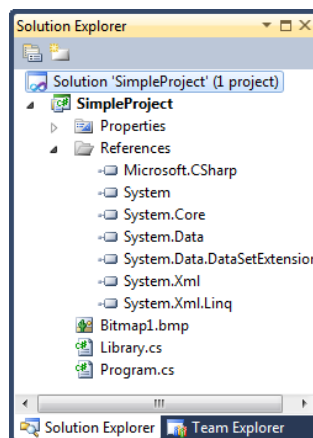
We can give other programmers our `dll` files and they can use the classes inside them without needing the source code of the program. Even better, the metadata which the libraries contain makes it easy for the programmers to find their way around them.

System Libraries and references

A program running under `.NET` makes use of system resources as it runs. When the program makes wants to send a message to the console it will call a system method to do this:

```
System.Console.WriteLine("Hello World");
```

These methods are held in the system libraries, which are actually `dll` files that were created by Microsoft and distributed as part of the `.NET` system. An assembly file contains references to the library files that it works with. This list of references is maintained in Visual Studio via the Solution Explorer pane.




The References section of a project gives a list of all the resources used by an application. When Visual Studio creates a new project of a particular type it will add references to the resources that it things that kind of project needs. Sometimes you will have to manage this list yourself, for example a Windows Phone project does not

usually contain a reference to the `Microsoft.Devices.Sensors` component. If you want to use the accelerometer in a Windows Phone application you have to add this reference yourself. We will see how to do this a bit later in the course.

Note that a reference is just that, a reference. If we include a reference to a resource this does not mean that the resource becomes part of our project; rather that the assembly is able to “reach out” to that resource when it runs. If the resource isn’t available at run time the program will fail. A given resource has a version number attached to it, so that if a new version of a library is added to a system any older programs will still be able to use the previous version.

The build process

We should now be happy with the idea that a project contains all the content required to create a single assembly. When a program is built Visual Studio will gather together all the elements that are needed to make the assembly and then use them to create the file. The output assembly is stored in a folder created when the project was set up.



Here you can see the folder that contains the output from the simple project that we created right at the start. The application file at the top is the executable assembly that runs on the computer. You could give someone this file to run and, virus scanners permitting, they would be able to execute it.

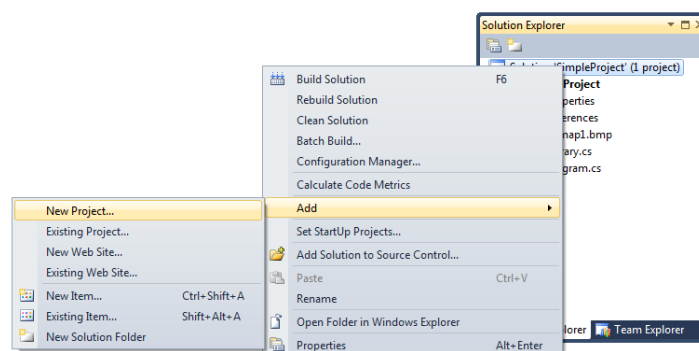
The screenshot above is for the debug version of the program. When Visual Studio makes the debug version it also writes out a database of debug information that provides information about the source when we re debugging.

Visual Studio Solutions

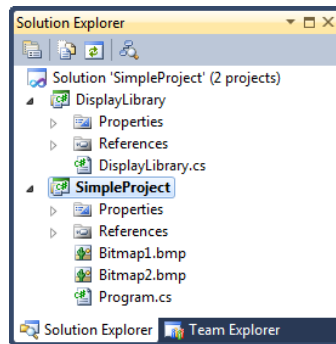
If you only ever want to produce a single executable assembly a project is all you need. However, if you want to create a number of separate assembly files and use this in a single application Visual Studio lets you do this too. A solution file is a container that can hold project files. When you create a new project, as we did with `SimpleProject` earlier, we actually get a solution created for us which encloses the project. The fact that we have been viewing all this content with the Solution Explorer could be considered a clue here...

Creating a multi-project solution


We can add new projects to a solution simply by right clicking the solution and selecting `New Project` from the menu that appears.



We can then pick up any of the project templates and make a new project which is then held as part of the solution:



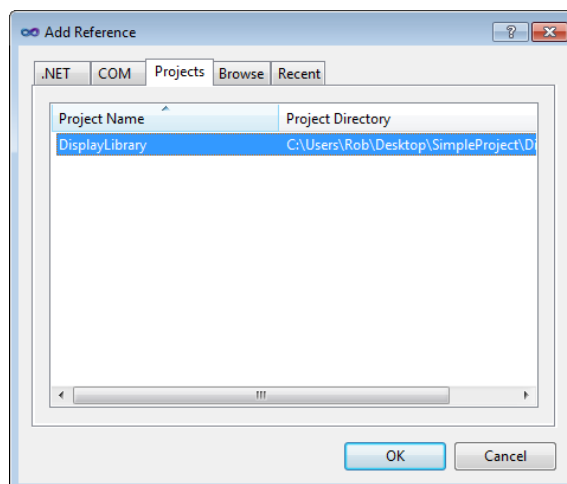
In the above I made a new library project called `DisplayLibrary`. This will compile to produce a dll file which will be stored in the binary folder for the new project.



This library can be used by any projects that add a reference to the dll file.

Linking Projects

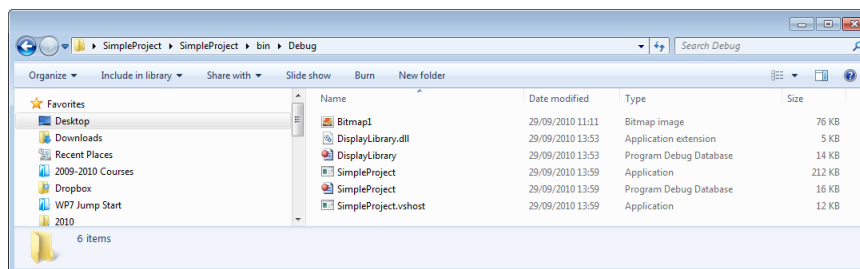
At the moment the projects in our solution are not connected in any way. If programs in `SimpleProject` want to use resources from `ResourceLibrary` they have to add them as a resource reference:



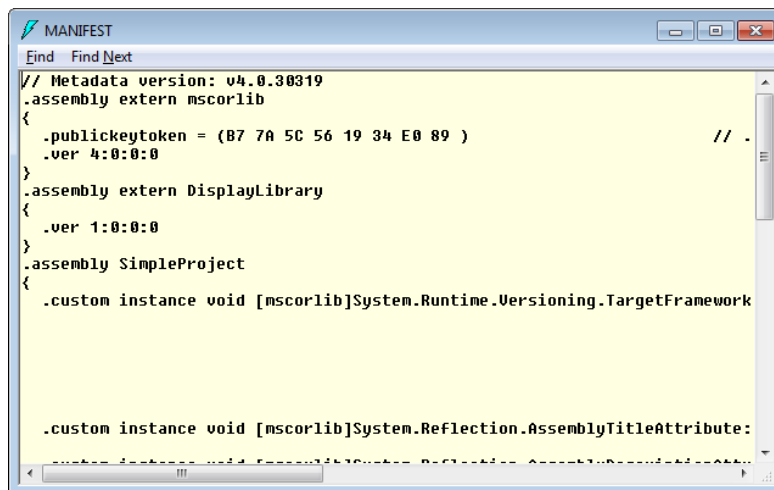
In the dialog above I'm adding a reference to `DisplayLibrary` to the project in `SimpleProject`. To make my programming easier I can add a using directive into the program source file.

```
using DisplayLibrary;
```

This all helps to remind us that the using directive doesn't actually connect our program to any resources that it wants to use. We only need the using directive to reduce the amount of typing that we have to do when writing programs. When we build the `SimpleProject` we find that the binary directory for this program now contains copies of the `DisplayLibrary` dll files:



This is a version of the `SimpleProject` program that uses the `DisplayLibrary` to load and display images. The external resources used by an assembly are listed in the manifest for that assembly.



Here you can see the manifest for a version of the `SimpleProject` program that uses a library called `DisplayLibrary`. Note that the version of the library is added to the information so that if a new version of the library is released this assembly will continue to use the old one.

Multi-project solutions

You can put many projects of different types into a single solution. For example you could create an XNA game for Windows Phone, Xbox 360 and Windows PC using a single solution file. Each of the devices could have their own project file, the game engine could be held in a fourth project that was then included by the others.

If you are writing a Silverlight application that connects to a web site that provides services you can put both the Silverlight application project and the web server project inside the same solution.

The projects in your solution and the tasks performed by each of them should be designed at the start of your development.

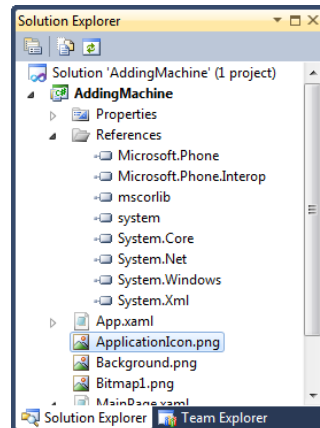
If your solution contains several projects that produce executable outputs you can set one as the “Startup Project” which will be given control when the project is executed.

Windows Phone Solutions

There are a number of kinds of Windows Phone solutions that can be created using Visual Studio. These are Silverlight application, XNA applications and a type that combines Silverlight and XNA. The combined project is very interesting. This makes it possible to create a program that uses Silverlight for the user menus and XNA for the gameplay. There are also Windows Phone Projects that can be added to solutions that can add extra elements such as background tasks. We will find out how these are created later in Chapter 9.

Silverlight Windows Phone Solution

The adding machine that we created in the previous chapter was created as a Silverlight project. If you take a look at Solution for this project you get something like this:

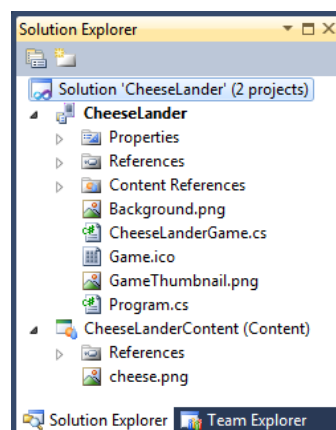


A Windows Phone Silverlight solution contains the `MainPage.xaml` file that describes the appearance of the main screen. It also contains the images that will be used to display the application icon and splash screen when the program runs. If you add new pages to your Silverlight program these will be added to this project as an xaml file and also the code behind file that we saw earlier. Note that the project also includes references to all the system libraries that are used by the phone.

XNA Windows Phone Solution

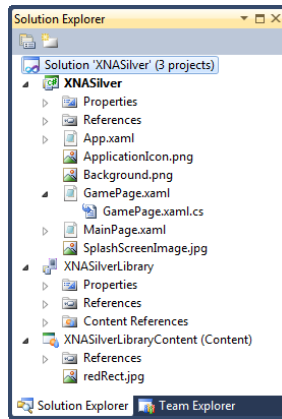
The XNA framework was developed for writing games. It provides a complete game engine into which you can put your game update and draw behaviours. It also provides comprehensive content management so that images, textures and sounds can be easily incorporated into a game.

When you make an XNA solution using the template supplied by Visual Studio you actually get two projects created. One of these contains the executable code. The other will hold all the content the game uses. The idea is that if we want to make a game that runs on multiple platforms we just have to add new projects to the solution. These can share the single resource project.



Above you can see a solution for a partially completed `CheeseLander` game. The game uses a single item of content, the images of the cheese to be landed. If we add resources such as images to the project these will be stored in the Content project and could be used by any game project in this solution.

Silverlight and XNA Solution

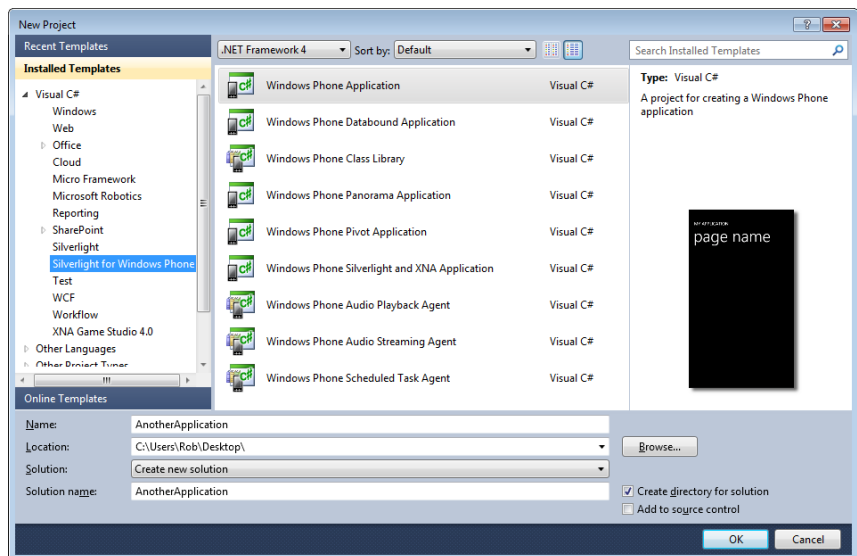


A solution built for XNA and Silverlight contains three projects. The Silverlight project actually contains the application, with the XNA game running on one page of the Silverlight application. However, the solution needs a library of XNA components and also the XNA content project. Don't worry if this seems a little complicated, the Visual Studio project template does all the hard work, we just have to fill in program code.

Creating a Windows Phone solution

We have seen how to create a very simple application which just contains a Main method. The process of creating a Windows Phone application is exactly the same. We choose a template and then Visual Studio creates the projects and source files that match that template.

The starting point is File>New>Project from within Visual Studio. This produces the New Project dialogue:



On the left of the dialog box we can see all the projects available. This screenshot was taken from my “working” machine, as you can see I’ve installed lots of extra toys, including the Microsoft Robotics framework and the Micro Framework. You might not have those present, but you should have “Silverlight for Windows Phone” available, along with the ones for “XNA Game Studio 4.0”. If you haven’t this means you are using a machine that does not have the Windows Phone SDK installed on it. The SDK is a free download from create.msdn.com.

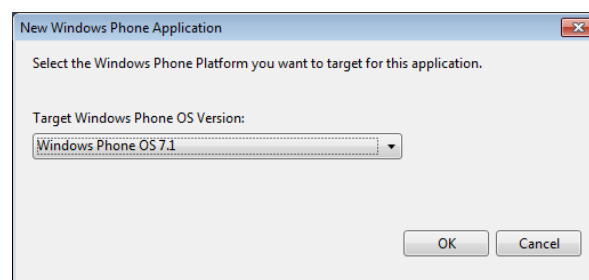
Above in the middle window you can also see the range of Silverlight projects available from the “Silverlight for Windows Phone” group. On the right hand side is a preview of what such an application would look like.

We can also use a template to create a project that is added to an existing Windows Phone application. We do this when we have a special task to perform. For example we might add the “Windows Phone Scheduled Task Agent” project to a solution. You would never create a new solution that just contains this project, but you might add a scheduled task to your application so that it can perform actions at regular intervals when the user is not actually using your application.

To start with we are going to create a simple “Windows Phone Application” using the top template. There are other templates you might want to explore later. Note that there is nothing particularly special about a template; it just creates certain combination of projects and source files. You could actually create the same effect by adding projects and source files and adjusting the settings, but the templates make it much easier.

We need to type in a name for the application. In the screenshot above I’ve used the rather daft name `AnotherApplication`. Note that if you have the “Create Directory for Solution” box checked (and you should) Visual Studio will create a folder with the application name and place all the solution files and project folders into it.

When we click OK Visual Studio asks us what kind of application we want to create



The first version of Windows Phone was called, rather confusingly, Windows Phone 7. There is nothing new in manufacturers playing around with version numbers. This has been going on ever since Version 3.0 of Fire was introduced way back in cave man days. The latest version of Windows Phone is going to be called Windows Phone 7.5. But it runs on an operating and tools version which is numbered 7.1.

This is only confusing if you think of version numbers as anything other than “things you use to tell one version from another”. It is unfortunate that the actual phone will not have the same number as the operating system that it runs on, but this has been happening for a very long time, and is just a quirk of the way that computer software is created and sold. The good news is that in general the bigger the version number, the more things the device can do. As far as we are concerned the choice is between:

Windows Phone OS 7.0 – the original version of the phone which has a restricted feature set. However this feature set is available on all phones.

Windows Phone OS 7.1 – an upgraded version which can run lots of new cool things, but is only available on new and upgraded phones

Many of the topics in this course are based on the new elements in Windows Phone 7.5 (running on Windows Phone OS 7.1) and so we should ensure that `Windows Phone OS 7.1` is selected above and click OK.

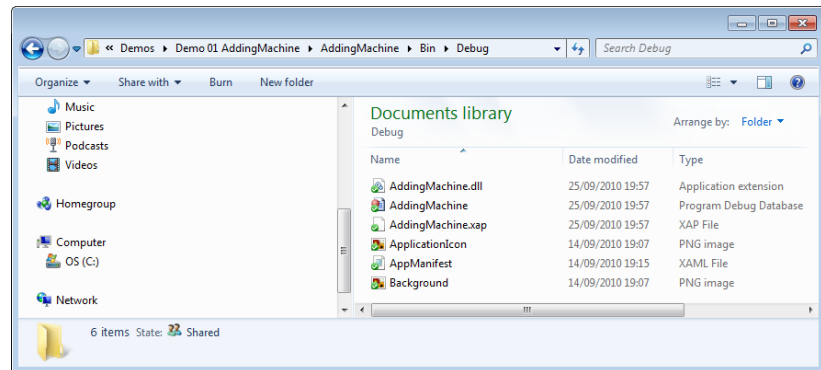
Visual Studio will now set to work building an empty Windows Phone Silverlight application. The application will run, and you can even load it into a Windows Phone, but it won’t be very interesting. It is simply a placeholder, into which you can add the display elements and behaviours that make the application into a useful thing to have.

Running Windows Phone applications

Starting the Compile and Run process in Visual Studio is as simple as pressing the F5 function key, or selecting `Debug>Start Debugging` from the menu. When this happens Visual Studio will assemble all the resources required to make the program work and begin running it. We will find out more about this process when we debug a program in the next section.

We know that when you compile and build a Windows PC application the result will be a file with a language extension of `.exe` which contains a method named `Main` which will be called when the program starts running.

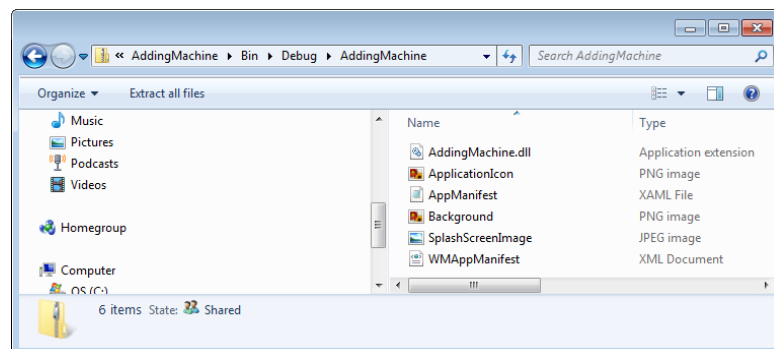
If you compile and build a Windows Phone application things are a little more complicated. The program will not be running inside the PC, instead it must be transferred into the Windows Phone device, or the emulator, before being started. Furthermore, any resources that the program needs (for example content that is not part of the program assembly) must be made available in the target device. This problem is solved by the use of a container file that holds the entire application and all required resources. Below you can see the output directory for the `AddingMachine` program. It contains all the resources along with the `AddingMachine.dll` file that contains the program code.



There is also a file that has the language extension `.XAP`. This is the `XAP` file. This is the container that holds the entire application.

The XAP file

The `XAP` file is an archive file (stored in the same format as `Zip` files) which contains all the files that make up the application. It also contains a manifest file which describes the contents of the `XAP`.



Above you can see the contents of the `XAP` file for the `AddingMachine` application. The `XAP` file is the file that is transferred into the Windows Phone device and made to run. When you press the run button in Visual Studio this file is created and then sent into the device or emulator. The target device opens the file, reads the manifest and then unpacks the content. Finally the application is executed. If your program is made up of multiple libraries and contains images and sound files these will all placed in the `XAP` file to be picked up in the target.

When you submit an application to the Windows Phone Marketplace you actually send the `XAP` file, this is the file that is loaded into the customer's phone if they buy your program.

3.2 Debugging Programs

One of the great things about the Windows Phone development environment is the ease with which you can debug programs. We can get a really good idea of what is going on inside a program by single stepping through the code and taking a look at the values in variables. You can debug programs on either the real device or the emulator.

To start a program we press F5 function key or select **Debug>Start Debugging** from the Visual Studio menus. We could also press the green triangular start button that you can see on the screenshot below. But before we do that, we want to decide whether we are running the program on a real device or the emulator.

Using the Windows Phone emulator

The Windows Phone emulator gives you a chance to try your programs and discover what they look like when they run. The emulator works in just the same way as a real phone. Effectively it is a Windows PC implementation of the Windows Phone platform, running exactly the same source code as the actual phone. You can use the emulator to prove that your program works correctly.

Deploying to the emulator

You can deploy programs to the emulator or a real device. To select the destination when you run a program you use the combo-box to the right of the run button as shown below.



It is possible to select a Windows Phone device even when one is not connected. In this case you will get an error when you try to run the program and Visual Studio detects that a device is not attached.



When the program starts it runs in exactly the same way as it would on the Windows Phone device. You can use the Windows PC mouse to click on items. If you have a multi-touch monitor on your Windows 7 PC you can use multiple touches on the emulator screen to access the multi-touch features of the user interface.

To stop the program you can click the back button (the left pointing arrow) at the base of the emulator or press the stop button in Visual Studio.

Emulator features


Note that the emulator is not provided with all the built-in programs provided with a real device. You can browse the internet using Internet Explorer and you can change some machine settings but lots of things, for example the Calculator, the Calendar and

Zune behaviour are all absent. There are some emulations provided to allow you to test your programs however, if you write a program that works with contacts from the address book (something we will do later) you will find that some contacts have been set up inside the phone for you to work with.

Emulator performance

The emulator cannot run at exactly the same speed as the Windows Phone hardware. The actual display speed of an application is restricted on the emulator so that the display will match the real device to a certain extent, but it is important to run your programs on a real device to get an accurate impression of what the user experience will be.

Programs in the emulator



When Visual Studio deploys an application to the Windows Phone emulator it installs it on the emulated device. The application will remain there until the emulator is reset. Above you can see the start menu for the emulator after the `AddingMachine` application had been run on it. We can run this application again just by clicking the icon.

Visual Studio Debugging

The debugging features are integrated into the Visual Studio environment. You can run programs from within Visual Studio. You can add *breakpoints* into your program. When a program reaches a statement nominated as a breakpoint Visual Studio will pause the program and allow you to look at the contents of variables inside the code.

You can add breakpoints to programs running inside the Windows Phone device. You can even set and remove breakpoints as the program itself is running.


Debugging is an important part of the development process. If you are going to become a successful programmer you are going to have to get used to that feeling in the pit of your stomach that you get when the program doesn't do what you expected. Fortunately there are some even more powerful tools and techniques that you can use to debug your programs.

Debugging and Programming

However, one important thing to remember here is that I don't consider it acceptable to debug your programs into life. Every line of code you write should be put there on the basis that you know exactly what the line is going to do, and why it is there. I use debugging to fix an implementation of a solution I have fully thought through. My programs tend to go wrong because I have miss-typed something or because I don't have the correct understanding of the behaviour of a library I'm using. If you don't know how to solve the problem it is very unlikely that throwing a few lines together might just solve the problem, any more than throwing a bucket of components at a wall will get you a new Xbox.


Adding a breakpoint at a statement


Breakpoints are added in just the same way as for any program you are working with in Visual Studio. You click in the left margin of the line you want the program to break at.




The line is highlighted as shown above. You can clear a breakpoint by repeating the process. When the program reaches the breakpoint it will stop running and you will be able to take a look at the values of the variables.

Running the Program

We run the program by clicking the Run button in the Visual Studio menus. The button looks like this: . When you press the button the program is deployed into the target device and starts running. You can use the function key F5 to instead of pressing the button if you prefer. When the program runs, if it reaches this line the execution will pause. Of course, to get to this statement in the adding machine we will have to enter some numbers into the textboxes and then press the enter button.




The yellow highlight shows the line reached. We can view the contents the variables in the program at this point just by resting the mouse pointer on the variable we want to see:




The variable result has not been set to the result of the calculation yet because the line that is highlighted has not been performed. To get the line performed we need to step through the statement.

Single stepping

We can step through a single line of the program by clicking the Single Step button in the Visual Studio menu system. The button looks like this: .

Each time you click this button the program will obey a single statement. Clicking the button once will move the program onto the next statement. The function key F11 will have the same effect.



```
private void calculateResult()
{
    float v1 = float.Parse(firstNumberTextBox.Text);
    float v2 = float.Parse(secondNumberTextBox.Text);

    float result = v1 + v2;

    resultTextBox.Text = result.ToString();
}
```

There are actually three versions of the Single Step button.



This means obey a single statement. If the single statement is a call of a method this version of single step will step into the method. This is called the “Step into” version of single step. Use F11 to perform this.



This means obey a single statement. If the single statement is a call of a method this version of single step will step over the method call. This is useful if you don’t care what happens inside the method. This is called the “Step over” version of single step. Use F10 to perform this.




This means run to the end of the current method and exit from it. Then break. This is very useful if you step into a method by mistake, or you have seen all you need from the method you are stepping through. Use SHIFT+F11 to perform this.


Controlling Program Execution

Visual Studio maintains contact with the running program in the target device. You can issue commands to stop, pause and resume execution and these will directly affect the program.

Starting or Resuming Execution


If you want to start the program you press the  (run) button. You also use this button to resume execution after a breakpoint. You can also press the F5 function key.

Pausing program execution

Pausing execution is not usually that useful, in that the program might not be at a place in the code where it is doing anything interesting, but you can sometimes use it to find out what a program is doing when it appears to have got “stuck”. To pause a running program you click the  button. The program stops at the statement that it was running when you pressed pause.

If you want to pause a running program you should remember that you can set breakpoints in the program code even when the program itself is running. We could have put our breakpoint in the `calculateResult` method while the code was running and then pressed the equals button in the emulator to send the program that way.

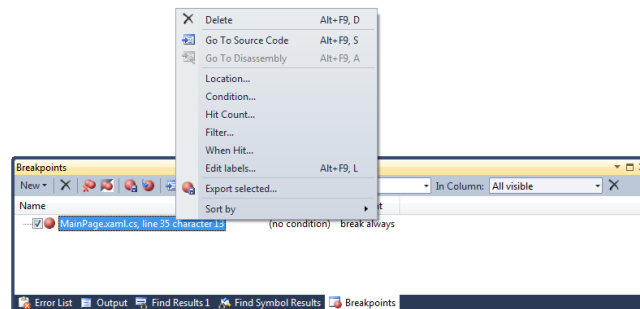
Stopping the program

The  button can be used to stop the program from running. When we click this button Visual Studio will stop program after the next statement it obeys. We can use it if the program appears to have got stuck. We need to be a bit careful with this button though.

When a Windows Phone program is stopped properly it is sent a sequence of messages that tell it the end is nigh and ask the application to do any tidying up that is required. We will discover how this works later in the course. If the program is stopped using the button above these messages are not sent, and so a program interrupted in this way may lose or corrupt data that it was using when the stop event was fired.

Managing Breakpoints

Visual Studio also has a breakpoint window that you can use to see the breakpoints you have created. You can display this window by going to the **Debug>Windows** menu in Visual Studio and selecting the **Breakpoints** option.




You can use this window to set breakpoint properties. You can request that a breakpoint only fire after a given number of hits, or only fire when a particular condition fires. This is terribly useful. You might have a program that fails after the fiftieth time round a loop. It is very nice to be able to add a counter or a test so that the program skips past the breakpoints you are not interested in.

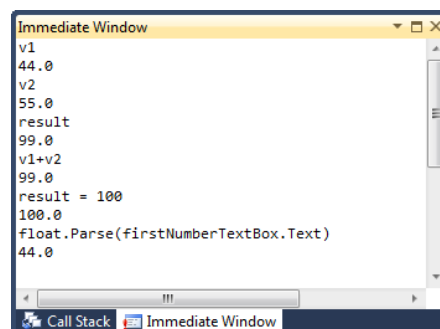
To open the properties of a breakpoint, either in the code itself or in the Breakpoints window, right click the breakpoint and then select **Breakpoint** from the context menu that appears.

Using the immediate window

The Immediate Window can be displayed by going to the **Debug>Windows** menu and selecting it. It is where you can view and change the contents of variables. You can also use it to evaluate expressions. As an example of the immediate window in action, we could set a breakpoint just before the `calculateResult` method returns:



When I run the program and press the equals button the program will hit the breakpoint. I can now move to the Immediate Window and play with the variables in this method.



If I just type the name of a variable the window displays the value in that variable. I can also type in expressions (e.g. `v1+v2` as shown above) and these are evaluated and displayed. I can also assign a value (e.g. `result = 100`) and call methods on objects to work out results.

The Immediate windows uses Intellisense to help you write the statements you want to run. It even works when you are debugging a program stored inside the Windows Phone device. It is very powerful, but should be used with care.

Design for debug

Single stepping through code is a great way to find out what it is doing. When you write code it useful to design it so that you can easily step through it and find out what it is doing. The code in the `calculateResult` method could easily be as a single statement.

```
private void calculateResult()
{
    resultTextBlock.Text =
        (float.Parse(firstNumberTextBox.Text) +
         float.Parse(secondNumberTextBox.Text)).ToString();
}
```

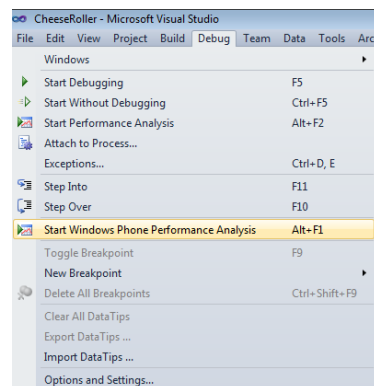
However this version is impossible to step through. I am a great believer in spreading my code over a number of statements and even using extra temporary variables when it runs. This makes debugging easier and doesn't even cost your program more memory since in the "smaller" version of the code above the compiler will still have to create some internal temporary variables to perform the calculation. The only difference with the "more efficient" version above is that you can't take a look at them.

3.3 Performance Tuning


Performance tuning is something you can do during development to make sure that your program gives the user a good experience and doesn't flatten the battery too quickly. The Windows Phone SDK includes a profiling tool that you can use to find out what demands your program is placing on the phone, and where the program is spending most of its time. You can also find out how much memory the program is using, and how often the Garbage Collector is being used to remove discarded objects from memory.

Using Performance Analysis Tool

The Performance Tuning tool is started from within Visual Studio:



The profiling can take place on the emulator, or on the Windows Phone device itself. There are a number of different levels of profiling detail available and the profiler produces graphs that you can view once the program has finished running.



The analyser will also display warnings if it detects that your program is making too many demands on the processor.

We will do some profiling of an application Chapter 10, as part of preparing a product for the marketplace.

What We Have Learned

1. Visual Studio is the development environment that is used to create Windows Phone applications. You can obtain a free version of this SDK.
2. Microsoft .NET applications are organised into *assemblies*. An assembly is either an executable program (.exe) or a library (.dll).
3. Visual Studio organises the elements required to produce an assembly into a project. A project brings together the source files and other program resources such as images and sounds required to produce the assembly. When the assembly is created a manifest is added to the project which describes the contents.
4. You can use the ildasm program which is supplied with Visual Studio to view the contents of assemblies.
5. A Visual Studio solution brings together a number of projects which can describe the creation of different elements of an application. This can include the production of completely different components, for example the server and client parts of a large installation.
6. A Windows Phone application can contain a number of assemblies, including library resources. An assembly will contain versioned references to all the resources that it uses.
7. When a Windows Phone application is prepared for deployment to the target device it is packaged into a XAP file which contains a manifest file, program assemblies and any resources that the program uses.
8. The Windows Phone emulator runs on the Windows PC and provides an emulation of the functionality of the device, but does not emulate the performance.
9. Visual Studio provides a means by which a programmer can insert breakpoints into program. These are statements where program execution will stop and the programmer can view the contents of the variables in the program. Breakpoints and single stepping of code can be used on programs executing in the Windows Phone device as well as the emulator.
10. The Windows Phone Performance tools allow a developer to find and target those parts of their program that could best be optimised for speed. It also helps to detect “memory leaks” and allows a developer to see how many “garbage” objects it is creating as it runs.

4 Constructing a program with Silverlight

In chapter 2 we looked at some of the elements provided with Silverlight. Now we are going to build on our skills and delve a bit more deeply into what Silverlight can do for us. By the end of this section you should be able to create useable, multi-page Silverlight applications that will work on Windows Phone and allow the user to manipulate data. This is a large chapter, with a lot of demonstrations. But don't worry.

4.1 Improving the User Experience

At the moment we are able to build user interfaces out of three Silverlight elements:

- `TextBlock` to display messages and labels
- `TextBox` to receive input from the user
- `Button` to receive events

Now we are going to build on these three items to create slightly more interesting places for the user to work. At the same time we are going to learn a bit more about how the Silverlight elements work together to build user interfaces.

Manipulating Element Properties

We have seen that Silverlight display elements are actually implemented as classes in the code, and that we can change how they appear by modifying the properties in instances of these elements. The properties of a Silverlight element can be set using the Properties pane in Visual Studio or by editing the XAML that describes each element. However, these settings are made when the program is built. We can also set the properties of the elements when the program runs. We have already done this in our `AddingMachine` application, the result is displayed by changing the properties of the `ResultTextBlock`.

```
resultTextBlock.Text = result.ToString();
```

Programs can manipulate any of the properties of a display element. This includes their position on screen, making it possible for us to create moving sprites in games.

We are going to take a look at some other properties that are available. We are going to use further property manipulation to improve the user experience of our `AddingMachine` program. However, we are just going to scratch the surface of what is actually possible. I strongly advise you to spend some time looking at just what you can do with elements. There are some very powerful features available.

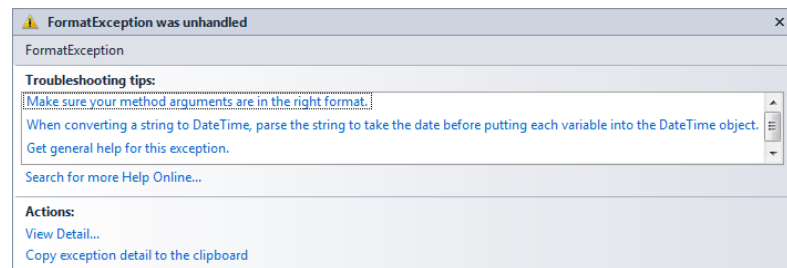
Adding Error Handling

We are going to start by improving the error handling of our `AddingMachine`. In my experience I spend at least as much time writing code to deal with invalid inputs for a system as I do writing programs that deal with the working parts.

At the moment the error handling in our `AddingMachine` program is non-existent. If a user enters text into a `TextBox` the program will try to use the `Parse` method to convert this into a number, and this will fail.



Above you can see trouble brewing for our adding machine program. The user will claim that they just wanted to work out the answer to “two plus two” but the result will not be a successful one:



This is because this version of the program uses a very simple method to convert the text in the `TextBox` into a number:

```
float v1 = float.Parse(firstNumberTextBox.Text);
```

If `firstNumberTextBox` contains text that cannot be converted into a number (for example “two”) it will throw an exception. We can improve on this by using a different version of the `Parse` method:

```
float v1 = 0;
if (!int.TryParse(firstNumberTextBox.Text, out v1))
{
    // Invalid text in textbox
}
```

The code above uses the `TryParse` method which returns `false` if the parse fails. It will obey the statements in the conditional clause if the text in `firstNumberTextBox` contains invalid text. It would be nice to turn the text red in the `TextBox` to indicate there is a problem.

Changing the colour of text

A good way to show an error status is to change the colour of the text. This is a very good way to highlight a problem. Silverlight uses a system of brushes to draw items on the screen. This is a very powerful mechanism. It means that we can draw text using images or textures and get some very impressive effects. In the case of our error display we don’t want to do anything quite so interesting, we just want to draw the text

with a solid red brush. Silverlight has a set of such colours built in, and so to turn the text in a TextBox red we just have to set the foreground brush to a solid red brush:

```
float v1 = 0;

if (!float.TryParse(firstNumberTextBox.Text, out v1))
{
    firstNumberTextBox.Foreground =
        new SolidColorBrush(Colors.Red);
    return;
}
```

This version of the code works OK. If a user enters an invalid number (i.e. one which contains text) the colour of the foreground for that textbox is set to a solid red brush and the method ends without calculating a result. However, this is not a very good solution. The program should really check more than just the first value when it runs. If the user has typed invalid text into both boxes they would much rather find out in a single message, rather than having to fix one error and then be told later that there was also a second fault item.

There is a more serious fault too, in that the code above will turn the text red when the user enters an invalid item, but it will not turn it back to the original colour when the value is corrected. We must add an else behaviour to put the colour back to the proper brush when the value that is entered is correct. We need to store the “proper” brush so that we can use it for this. A complete version of calculateResult looks like this:

```
private SolidColorBrush errorBrush =
    new SolidColorBrush(Colors.Red);
private Brush correctBrush = null;

private void calculateResult()
{
    bool errorFound = false;

    if (correctBrush == null)
        correctBrush = firstNumberTextBox.Foreground;

    float v1 = 0;

    if (!float.TryParse(firstNumberTextBox.Text, out v1)) {
        firstNumberTextBox.Foreground = errorBrush;
        errorFound = true;
    }
    else
        firstNumberTextBox.Foreground = correctBrush;

    // Repeat for v2

    if (errorFound)
        resultTextBlock.Text = "Invalid inputs";
    else
    {
        float result = v1 + v2;
        resultTextBlock.Text = result.ToString();
    }
}
```

This code turns the text red if a value is invalid. It also maintains a copy of the correct brush, so that if a value is correct it can be set to the correct colour.

This code works well and provides a reasonable user experience. It also illustrates that when you write a program with a user interface you actually end up writing a lot more code than just the statements that are needed to calculate the result. The code to deal with input errors can greatly increase the amount of code that you have to write.

Editing the XAML for Silverlight Elements

XAML can be described as a ‘declarative’ language. This means that it doesn’t actually describe any behaviour (i.e. it does not tell the computer how to do something). Instead it just tells the computer above stuff. In C# terms we tell the compiler about things we want to use (for example variables) by *declaring* them. The XAML language is only ever used to tell Silverlight about things, hence it is called a declarative language. We will have to get used to putting some aspects of a program into the declaration of the Silverlight elements it uses. One such situation is in the configuration of the input mode for the `TextBox` elements that are used to enter the numbers in our `AddingMachine` program.

Configuring a `TextBox` to use the numeric keyboard

At the moment the `AddingMachine` uses the standard keyboard behaviour on the Windows Phone which is to display the text keyboard when the user selects that control. It would make sense to provide the user with a keyboard set to enter numbers, rather than the text one. There are a number of different keyboard configurations for the keyboard that a Windows Phone can display. Below you can see the standard “text” one.



To enter digits the user must press the ‘&123’ key to switch the keyboard into numeric mode.



It would be nice if each `TextBox` in the adding machine could be set for numeric entry when that `TextBox` is opened. We could do this by modifying the properties of the `TextBox` when the program starts running but the best place to put such settings is in the XAML description of the control itself. At the moment the XAML that describes the `secondNumberTextBox` is as follows:

```
<TextBox Height="72" HorizontalAlignment="Left" Margin="8,175,0,0" Name="secondNumberTextBox"
Text="0" VerticalAlignment="Top" Width="460" TextAlignment="Center" />
```

This tells Silverlight where on the screen to put the element, the size of the element and the name it has. There is also some text alignment information.

Each of the values enclosed in `<TextBox .. />` is an *attribute* of the textbox. Attributes are simple items that are not structured and can be expressed as a name – value pair. Above you can see that the `Height` of the element, the `Name` of the element and lots of the other properties are expressed in this way.

To tell the `TextBox` what form of text should be entered we can add extra information to it. This is called the `InputScope` of the `TextBox`.

This information can be given as a *property element* which is part of a `TextBox` value. To express this we have to use a slightly more complex version of `TextBox`:

```
<TextBox Height="72" HorizontalAlignment="Left" Margin="8,19,0,0" Name="firstNumberTextBox"
Text="0" VerticalAlignment="Top" Width="460" TextAlignment="Center">
  <TextBox.InputScope>
    <InputScope>
      <InputScopeName NameValue="Digits" />
    </InputScope>
  </TextBox.InputScope>
</TextBox>
```

Good thing I said “slightly” more complex, eh? If you want a XAML element to contain properties you have to use a different form of the element description. This is a bit confusing in the Silverlight context, so we might want to take a look at something a bit easier to understand.

Element with no properties

We can use the XML format to create a language that will hold information about people. We could start by just storing the name of a person:

```
<Person name="Rob"/>
```

This element is called “Person” and has a single attribute called “Name” which is set to “Rob”. The XAML compiler sees the `/>` at the end of the element definition and knows that it has reached the end of the description of this element.

Element with properties

However, we might want to store more complicated information about a person, such as where they live. Their address will be made up of lots of different elements, and so we should make this address a property of the person.

```
<Person name="Rob">
  <Address addressType="HomeAddress">
    <FirstLine text="18 Pussycat Mews"/>
    <Town text="Seldom"/>
    <County text="Wilts"/>
    <PostCode text="NE1 410S"/>
  </Address>
</Person>
```

This element contains one property, which is the address of the person. The address itself is another element which contains a single attribute (the type of the address) and then four properties. This illustrates an important point with XML based languages. A property is an XML element like any other, which means that it can have attributes and contain properties of its own. Silverlight uses this to good effect with the `Grid` element, which can hold a collection of Silverlight elements (including `Grids`) which are to be laid out on the page.

The XAML compiler can tell that an element contains properties because the first line of the element description ends with `>` rather than ending `/>`. The XAML compiler will regard everything it sees after the first line as a property of the element, until it sees the `</ElementName>` which rounds things off.

If you find this confusing at first, welcome to the club. The best way to think about this is to consider what is being stored. In the case of the `Person` above the name is just a single item which can be expressed in a simple attribute. However the address will contain multiple properties of its own and so is best made into a property. There is also the possibility that a `Person` will have several address values, for example a `HomeAddress` and a `WorkAddress` and so the above design makes it possible to add new address types easily.

The type of the address has been made an attribute of the address, which I think is the most appropriate. Whether to make a data item an attribute or a property is something

you have to think about when you use XML (eXtensible Markup Language) to describe things.


If we go back to our `TextBox` we should find that things make a bit more sense now.

```
<TextBox Height="72" HorizontalAlignment="Left" Margin="8,19,0,0" Name="firstNumberTextBox"
Text="0" VerticalAlignment="Top" Width="460" TextAlignment="Center">
  <TextBox.InputScope>
    <InputScope>
      <InputScopeName NameValue="Digits" />
    </InputScope>
  </TextBox.InputScope>
</TextBox>
```

Simple information about the `TextBox` is given as attributes, whereas the `InputScope` value is given as a property. The `InputScope` itself contains a set of properties, each of which describes a particular input scope. This means that one day we might be able to ask for an input scope that includes `Digits` and `Text`, but not punctuation. However, we don't want anything as complex as that. We just want to ask for `Digit` input, and so that is the `InputScope` that we set up. If we use the above description for the two textboxes in the adding machine we get a vastly improved user interface where the user is given the numeric keyboard when they move to those input elements.

This is an example of a situation where just adding text to the XAML is the quickest and easiest way to configure the input behaviour. We could have used the Visual Studio property pane, or even written some complex code to build a collection of `InputScopeName` values and assign them to the `TextBox`, but this solution is much neater, and also allowed us to refine our understanding of XAML.

```
<TextBox.InputScope>
  <InputScope>
    <InputScopeName NameValue="" />
  </InputScope>
</TextBox.InputScope>
</TextBox>
```



Visual Studio provides a feature called "Intellisense" where the editor tries to give you help by figuring out what you are trying to type. Because Visual Studio knows all the possible values for an `InputScope` item it suggests them when you start to enter that expression. Above you can see what happens when Intellisense detects that we are entering an `InputScope` value.

XAML Attributes and Properties – the truth

It turns out that from the point of view of XAML attributes and properties are actually interchangeable. The software that uses the XML description can get information from either parts of an XML element. Rather than the somewhat complicated format above we can actually write:

```
<TextBox InputScope="Number" Height="72" HorizontalAlignment="Left" Margin="144,44,0,0"
Name="startHourTextBox" Text="00" VerticalAlignment="Top" Width="104"
TextAlignment="Center" />
```

This sets the input scope of the `TextBox` to numbers. It doesn't work if we want to set multiple scope values at the same time, but it turns out that it is only useful to set one value anyway. The reverse of this is also true.


```
<TextBox HorizontalAlignment="Left" Margin="8,175,0,0" Name="secondNumberTextBox" Text="0"
VerticalAlignment="Top" Width="460" TextAlignment="Center">
    <TextBox.Height>
        72
    </TextBox.Height>
</TextBox>
```

In the XAML above the height of the `TextBox` has been brought out of the attributes and turned into a property. The property name must contain the name of the element type that it is part of. We can't just say `Height`, we have to put `TextBox.Height`. This is not a restriction of XML (we didn't have to put this in our `Person` example above) but is a requirement for XAML.

From the point of view of an XAML writer it doesn't really matter which of these forms we use. For simple name/value pairs we can use attributes, but for more complicated structures we may decide to use properties.

C# Property Setting

If you were wondering what the C# to do the same setting would look like, this is it:

```
// Make a new input scope
InputScope digitScope = new InputScope();

// Make a new input scope name
InputScopeName digits = new InputScopeName();
// Set the new name to Digits
digits.NameValue = InputScopeNameValue.Digits;

// Add the name to the new scope
digitScope.Names.Add(digits);

// Set the scope of the textbox to the new scope
firstNumberTextBox.InputScope = digitScope;
```

If you read carefully through this code you will see that it is actually building up each of the elements in the XAML above.

The solution in *Demo 01 AddingMachine with Error Checking* implements a version of the `AddingMachine` that performs validation and turns invalid entries red. It also sets the input scope of the `TextBoxes` to numeric.

Displaying a MessageBox

Silverlight provides a `MessageBox` object which you can use to inform the user of errors:

```
MessageBox.Show("Invalid Input");
```

This will display a message at the top of the phone screen that the user must clear before they can continue. An alert sound is also played.



If you want to send larger messages to the user of the phone you can do this by displaying several lines in the message:

```
MessageBox.Show("Invalid Input" +  
    System.Environment.NewLine +  
    "Please re-enter");
```



The element `System.Environment.NewLine` will provide a newline character in a manner appropriate to the target platform. A `MessageBox` is useful if you want to be sure that the user has acknowledged a message before continuing. However, you need to remember that when it is displayed it will stop the program at that point.


The solution in *Demo 02 AddingMachine with MessageBox* displays a message box when the user enters an invalid value.

Message Box with selection

We can ask the user to make a choice in a `MessageBox`:

```
if (MessageBox.Show("Do you really want to do this?",  
                    "Scary Thing", MessageBoxButton.OKCancel)  
    == MessageBoxResult.OK)  
{  
    // do scary thing here  
}  
else  
{  
    // do something else  
}
```

This version displays a message with a cancel option as shown below. The if condition is evaluated according to the response from the user.



Adding and Using Assets

In the previous section we saw how Visual Studio manages assets such as images and sounds and adds these to the XAP file for transfer into the target device. Now we are going to see how our applications can add assets of our own and then use it.

A background screen for the Adding Machine



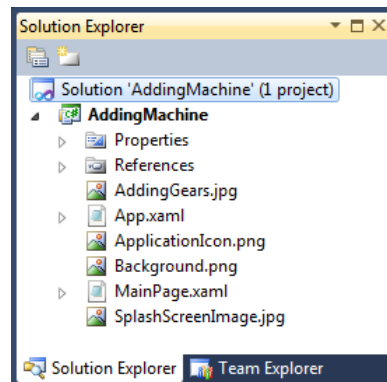
We have been asked by the son of the boss of our customer (who seems to think he is a graphic designer) to add the above image as a background to the adding machine application. This is actually a photograph of some of the gears of Charles Babbage's "Difference Engine" and he thinks that this addition will make the program much better. We're not convinced, but as he is the son of the boss we have to go along with it.

Assets for Windows Phone

When you add assets such as this it is important to remember that the target device does not have a lot of memory or a particularly high resolution screen. If the above picture was from a high resolution camera it might contain a lot more detail than would be required for a phone application. Remember that the highest resolution for a Windows Phone is 800x480 pixels. There is no point in using a background image with a higher resolution than this.

Adding Images as Items of Content


We can add an image as an item of content to the `AddingMachine` project. The boss's son has given us a jpeg image called `AddingGears.jpg` to use. It turns out that the quickest way to add a content item to a project in Visual Studio is just to drag the item out of the folder and drop it onto the project:



This makes a copy of the resource file in the project directory at the appropriate point.

Links to assets

If we want to share an asset amongst a number of different projects (perhaps you have a company logo that is to be used by lots of different programs) then we can add a link to a resource, but to do this we have to add the item using the `Add Existing Item` dialog and then select "Add as Link":



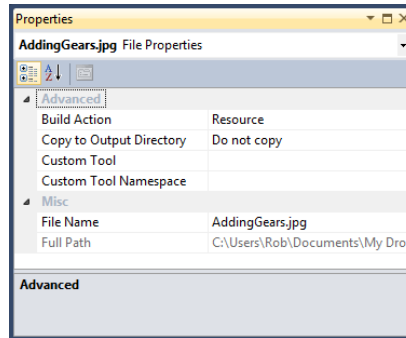
If we change the logo file the new version will be picked up by projects that use it next time they are rebuilt.

The "Build Action" of an Asset

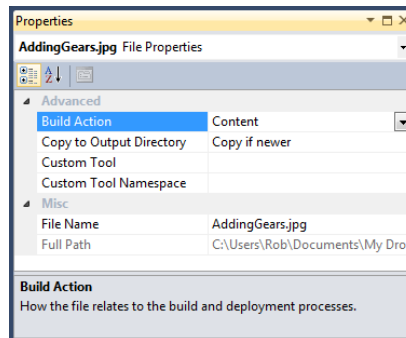
When Visual Studio builds a project there are a number of ways that assets can be added to the completed program. We are going to consider two; adding the asset as an item of content or adding the asset as a resource inside the program assembly. The decision of which to use can have a bearing on the way that programs load and run, so it is worth knowing a bit about how this works.

Adding an asset as Content

When an asset is added to a project it is initially added with a build action of "Resource".



If we want to use the asset as content we need to change its properties so that it is an item of content. We do this by changing the Build Action entry for the item of content.




Content assets are copied by Visual Studio into the application directory when the program is built. The “Copy if newer” setting means that Visual Studio will copy a new version of the image into the binary folder if the current image in the project is replaced by a newer file.

Using Image content in a program

An asset of type content is simply made available to the application as a file in the same folder as the program itself. To display an image held as content in our program we just need to add the line of XAML that describes the image file that we want to use:

```
<Image Height="611" HorizontalAlignment="Left" Margin="8,0,0,0" Name="gearBackgroundImage" Stretch="Fill" VerticalAlignment="Top" Width="472" Source="AddingGears.jpg" />
```

The **Image** element displays an image on the page. It has a **Source** attribute that identifies the file from which the image is to be loaded. If we enter a filename as the source the image will be loaded from the file. When the program runs the file will be loaded from the file and then drawn on the screen. Silverlight draws the elements on the screen in the order they are defined in the XAML. If this line is put at the top of the **Grid** containing the elements the image will be drawn in the background:



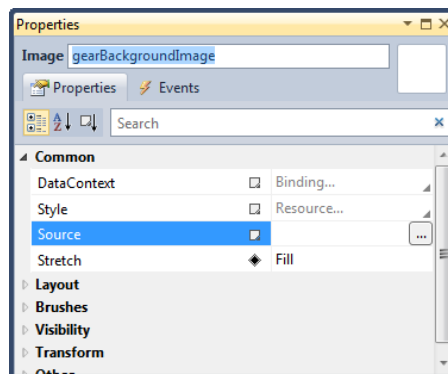
This is the result of the program. From an artistic point of view we are not convinced, by the choice of picture but the process itself works.

The solution in *Demo 03 AddingMachine with Background Content* displays a background image for the application which is loaded as an item of content.

Adding Images as Resources


There is another way to add an image to a project. This involves creating a new **Image** item from the **ToolBox** in the Visual Studio editor and then selecting the image resource to use for this item by using the Visual Studio the **Property** pane.

The first step in this process is to drag an image from the Visual Studio Toolbox onto the page. Then open up the **Property** pane for that resource:

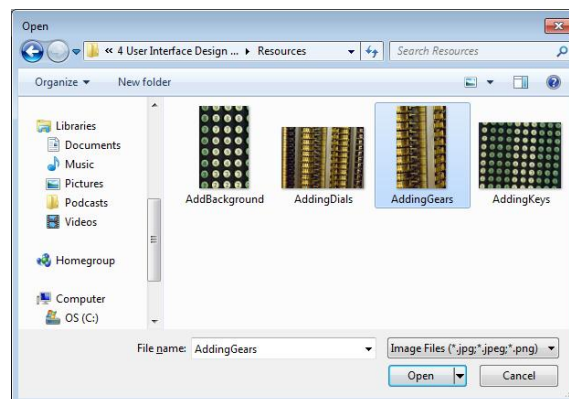


I've changed the name of the image to `gearBackgroundImage` and the next thing to do is to set the source of the image by clicking on the "... " button to the right of the source item. This opens the **Choose Image** dialogue.

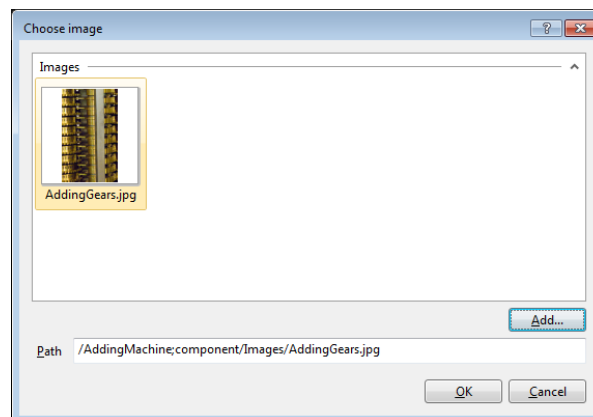
Constructing a program with Silverlight



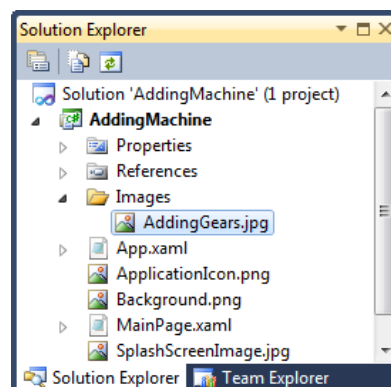
This dialog shows all the images which are presently stored in the application. As we have not added any the dialogue is empty. If we click on Add we can open a file browser to look for suitable pictures on the development computer.



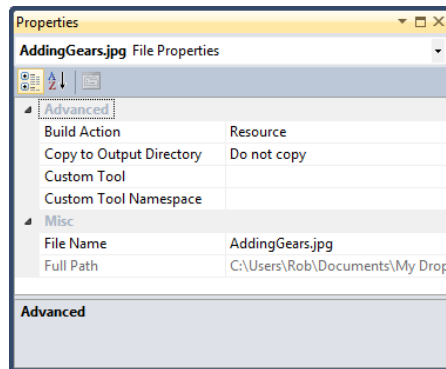
This adds the specified image to those available to the project:



It also creates the resource item in the project for this item:



Any other images that we add to the project will be added in the folder that Visual Studio created for us. The properties of these images will be set as Resources, not content.



The XAML that describes the image contains an image source that refers to the resource in the application:

```
<Image Height="611" HorizontalAlignment="Left" Name="gearBackgrounds" Stretch="Fill"
VerticalAlignment="Top" Width="472"
Source="/AddingMachine;component/Images/AddingGears.jpg" />
```

When the program runs the display is exactly the same, but there is a crucial difference in how the resource is located and used.

The solution in *Demo 04 AddingMachine with Background Resource* displays a background image for the application which is loaded as a resource.

Content vs Resources

When creating Windows Phone applications it is important that we understand the difference between content and resources in an application.

An item of content is simply a file that is stored in the application directory. Our programs can open and use content files when they need them. Adding content has no effect on the size of the program itself, although of course when a program is running and loads content it will use more memory.

A resource item is stored in the program assembly. Adding the gears background image to the assembly above made the program around 317K bytes larger. When a Windows Phone device loads a program it looks through the entire program assembly and does some checks to make sure that it is legal code. These checks take longer if the program is bigger. If your program takes too long to start up when it is loaded it will be killed by the Windows Phone operating system. This means that you should be careful about putting too many things in the program assembly.

Adding our image as an item of content made the program itself smaller, meaning that it starts running more quickly. However, content items take slightly longer to load because they have to be fetched from the filestore when they are used. Adding our image as a resource makes it faster to load when the program tries to fetch it, but makes the program larger and will make the program itself slower to start.

If this is confusing, here are some hints:

- If you have lots of large resources, and you don't need to load all of them into the program at once, use content. An example would be where you allow the user to select from a number of different background images for their screen.
- If you have lots of small resources which are used all the time throughout the program, use resource. An example would be if you had icons that represent particular program setting or options.
- If you find your program is taking too long to get started you should move some of the resources out of the program assembly and add a loading screen

of your own to show them being brought into memory when the program starts up. You can also speed up the program startup process by breaking your program into smaller assemblies which are linked together.

4.2 Data manipulation and display

Our adding machine is now quite useable, but the customer is still not completely happy. He doesn't like the way that users have to press the Calculate button each time they want a new answer. What he wants is for the program to detect when the text in the number boxes changes and then update the total automatically. Fortunately for us this turns out to be quite easy to do. We need to revisit events to find out more.

The TextChanged Event

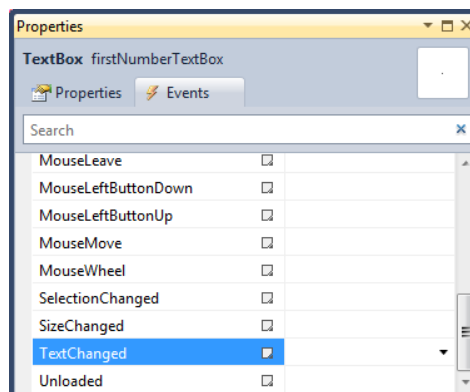
We first saw events when we considered the Button element. A Button element generates an event when it is activated. From a program point of view an event is a call to a method. We can get Visual Studio to "hook up" a button to an event for us just by double clicking the button in the designer. At the moment the Calculate button has a click behaviour that calculates and displays the result. The method to be called is identified in the Button description:

```
<Button Content="equals" Height="72" HorizontalAlignment="Left" Margin="158,275,0,0"
Name="equalsButton" VerticalAlignment="Top" Width="160" Click="equalsButton_Click" />
```

When the button is clicked this causes the nominated event handler to run inside the page class. In the case of our AddingMachine this will call the calculateResult method.

```
private void equalsButton_Click(object sender, RoutedEventArgs e)
{
    calculateResult();
}
```

We can find out what events a TextBox can create by selecting a TextBox in the Visual Studio editor and then selecting Events in the properties pane.



A TextBox can produce a lot of events, including some which seem somewhat meaningless on a Windows Phone. Visual Studio highlights the event that is most useful from a TextBox. This is the TextChanged event. This is fired each time the text in the TextBox is changed. If we double click on TextChanged in the Properties pane Visual Studio will create a method and hook it up to this event:

```
private void firstNumberTextBox_TextChanged(object sender,
                                           TextChangedEventArgs e)
{
}
}
```

We just need to add a call of calculateResult into this method:

```
private void firstNumberTextBox_TextChanged(object sender,
                                           TextChangedEventArgs e)
{
    calculateResult();
}
```

We can do the same thing for the second text box, so that if either of the boxes is changed the sum is updated. This means we can now remove the Calculate button.

Double Changed Events

If we run this program we find that it works quite well. However, we do notice a problem with data validation. If the user types text instead of a number our program will spot this and display a message box:



However, this message is displayed twice for each invalid character that is entered. It is as if the `TextChanged` event is firing twice for each character. It turns out that this is the case and it is a known issue with Windows Phone applications. We can get around it by checking for successive changed events with the same data:

```
string oldFirstNumber = "";
private void firstNumberTextBox_TextChanged(object sender,
                                           TextChangedEventArgs e)
{
    if (firstNumberTextBox.Text == oldFirstNumber) return;
    oldFirstNumber = firstNumberTextBox.Text;

    calculateResult();
}
```

This method only calls `calculateResult` if the text in the `TextBox` has really changed.

The solution in *Demo 05 AddingMachine with no button* implements an adding machine which uses `TextChanged` events and does not have a calculate button.

Data Binding

Data binding is wonderful. Perhaps there are even statues out there to it. Perhaps we should build one, but then again perhaps not. Data binding takes the place of the “glue” code that we are writing to connect display elements to the methods that work on them. It allows us to take a property of a display element and connect it directly to an object. All we have to do is create an object that behaves in a particular way and we

can get values in the object transferred into the Silverlight element with no effort on our part.

Data binding can work in either direction. We can bind the text a `TextBox` to our application so that when a user changes the text in the textbox our application is informed of this. We can also bind a `TextBlock` to an object so that when a property in the object is changed the content of the `TextBlock` updates to match. This is “two-way” binding, where a program can both change an element and respond to changes in it. We can do “one way” binding so that a program can display output just by changing properties in an object. Finally we can bind to any of the properties of an object so that a program can move a Silverlight element around the screen by changing X and Y properties in a game object.

Creating an object to bind to

We are going to begin by creating a version of the adding machine which uses data binding. The starting point is a class that will do all the work. This will expose data properties that will be connected to element properties in the display object. There are three things that our object must expose:

- The text in the top `TextBox`
- The text in the bottom `TextBox`
- The text in the result `TextBlock`

We could build an `AdderClass` to express all this:

```
public class AdderClass
{
    private int topValue;
    public int TopValue
    {
        get
        {
            return topValue;
        }
        set
        {
            topValue = value;
        }
    }

    private int bottomValue;
    public int BottomValue
    {
        get
        {
            return bottomValue;
        }
        set
        {
            bottomValue = value;
        }
    }

    public int AnswerValue
    {
        get
        {
            return topValue + bottomValue;
        }
    }
}
```

This class has three properties. Two of them have get and set behaviours so that a user of this class can find out the values of the top and bottom items (the things we are adding together) and set them to new values. The third property is just the answer value (the result of the sum). This is never set to a value; the user of the `AdderClass` will only ever read this property to get the latest answer.

This class does all the “thinking” required to make our adding machine work. A program could make an instance of this, set top and bottom values and then read back the answer. If we wanted to change to a “subtracting” machine we would just have to replace this class with one that worked differently. This would be a good piece of design whether we were using Silverlight or not.

Adding Notification Behaviour

In order for a class to be able to bind to Silverlight objects it must implement the `INotifyPropertyChanged` interface:

```
public interface INotifyPropertyChanged
{
    // Summary:
    //     Occurs when a property value changes.
    event PropertyChangedEventHandler PropertyChanged;
}
```

For a class to implement this interface it must contain the event delegate which will be used by `AdderClass` to tell anything that is interested when a property has changed.

```
public event PropertyChangedEventHandler PropertyChanged;
```

The `PropertyChangedEventHandler` type is used by Silverlight elements to manage event messages. It is described in the `System.ComponentModel` namespace. If a Silverlight component wants to bind to any of the properties in the class it can add delegates to this event.

The `AdderClass` will use this delegate to tell the outside world when one of the properties in the class has changed. This will trigger a display update. Silverlight objects that are connected to our properties bind to this delegate so that they can be notified when required. The final version of the class looks like this:

```
namespace AddingMachine
{
    public class AdderClass : INotifyPropertyChanged
    {
        public event PropertyChangedEventHandler PropertyChanged;

        private int topValue;

        public int TopValue
        {
            get
            {
                return topValue;
            }
            set
            {
                topValue = value;

                if (PropertyChanged != null)
                {
                    PropertyChanged(this,
                        new PropertyChangedEventArgs("AnswerValue"));
                }
            }
        }
    }
}
```

```

        // repeat for bottom value

        public int AnswerValue
        {
            get
            {
                return topValue + bottomValue;
            }
        }
    }
}

```

The `AdderClass` now contains a member called `PropertyChanged`. Whenever our class wants to change a property value it must test the value of `PropertyChanged`. If it is not null this means that something has bound to the delegate. You can regard the `PropertyChanged` event as a mailing list. Process can subscribe to the list and be informed of interesting events that occur in the `AdderClass`. If nobody is subscribed to the mailing list (i.e. the `PropertyChanged` value is null) then there is no point in calling a method to describe a change. However, if it is not null this means that we have something out there that wants to know about changes inside our `AdderClass`.

The format of the method call that delivers our event is as follows:

```
PropertyChanged(this, new PropertyChangedEventArgs("AnswerValue"));
```

The first parameter to the call is a reference to the object that is generating the event. The second is an argument value loaded with the name of the property that has changed. The Silverlight element will use reflection (i.e. it will look at the public properties exposed by `AdderClass`) to know what properties are available. If it gets a notification that `AnswerValue` has changed it will then update any display element properties that are bound to the property in this class.

Note that this is just a string of text though, if we write “Answervalue” by mistake this will prevent the update message from being picked up correctly.

When the program fires the changed event it will cause the result display to fetch an updated value of `AnswerValue` for display to the user. When the `get` behaviour in `AnswerValue` is called it computes the total and returns it.

```

public int AnswerValue
{
    get
    {
        return topValue + bottomValue;
    }
}

```

If you are finding this confusing, don't worry. Just remember why we are doing this. We are going to connect this object to the user interface. The user interface will tell our `AdderClass` when the values of the two inputs are changed by the user. When the input changes the `AdderClass` instance needs to have a way of telling the user interface that there is now a new total to be displayed.

Adding the Namespace containing our class to MainPage

We now have our object that we want to data bind. The next thing we need to do is link our code to the XAML that describes our user interface. To do this we must add the namespace containing the `AdderClass` to the XAML.

```
xmlns:local="clr-namespace:AddingMachine"
```

If you look at the final version of `AdderClass` above you will find that it is in the `AddingMachine` namespace. By adding this line at the top of the XAML file for `MainPage.xaml` we tell the system that any classes in the `AddingMachine` namespace should be made available for use in our application XAML.

In the same way that a C# program must have all the using directives at the top of the source file a XAML file must have all the namespaces that it uses at the top as well. If you look in the file at the top you will find lots of namespaces being added. These are the classes for the Silverlight elements that make up the user interface.

Once we have added a namespace we now have to create a name for the resources that can be found in this namespace:

```
<phone:PhoneApplicationPage.Resources>
  <local:AdderClass x:Key="AdderClass" />
</phone:PhoneApplicationPage.Resources>
```

If this seems a lot of work then remember that you only have to do this once, and that the class that we add can contain many properties that the user interface can interact with.

Adding the Class to an element on the page

Now that the namespace is available we can connect a particular class from this namespace to the elements in our page. The element we are going to add the class to will be the Grid that holds all our display elements. Adding a class to an element automatically makes it available to any elements inside it, and so the TextBox and TextBlock elements can all use this class.

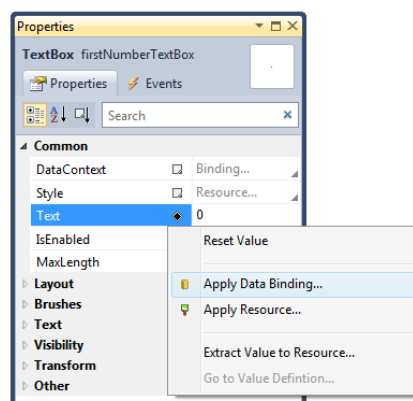
```
<Grid x:Name="LayoutRoot" Background="Transparent" DataContext="{StaticResource AdderClass}">
```

We get an element to use a particular class by setting a value for the DataContext for that component. For our application this will be a static resource, in that it will be part of the program. This might seem a bit over complicated, but it does mean that you can create a test data context that you can use instead of a “real” one.

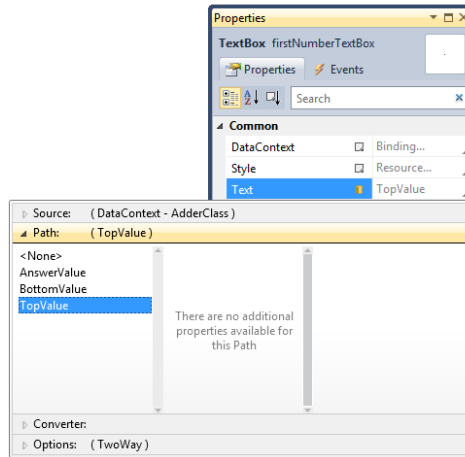
Connecting the Silverlight element to the property

We now have the Adder class available as a data context for elements on a page. This sets up a link between the Silverlight page and an object that is going to provide business behaviours. The next thing we need to do is connect properties of each element to properties that our business object exposes.

We can do this from the Properties pane for each element. Let’s start by connecting the text in the firstNumberTextBox with the TopValue property in our AdderClass. If we click on the TextBox in the Visual Studio editor we can then open the Property pane for the item. We can then go and find the Text property. At the moment it is just set to a string, we want to apply a Data Binding to our adding machine object:

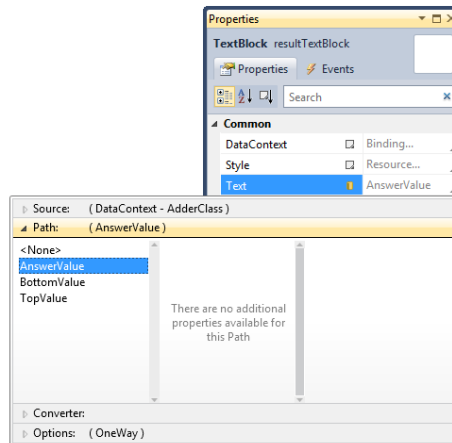


If we click “Apply Data Binding” the next thing we need to do is find the property that is going to be bound. Behind the scenes Visual Studio has gone off and opened the AdderClass and used found out what properties it exposes. It can then display a menu of those that are available:



This is where the magic starts to happen. Simply by selecting the required item from the list we can connect the Silverlight element property with the object. Note that the TwoWay option is selected so that when the user changes the value in the TextBox the object is informed.

When we bind to the TextBlock we are only given the option of OneWay binding, as it is not possible to send data from a TextBlock into a program.



Once we have made these bindings our program will just work. The interesting thing about this is that if you look in the MainPage.xaml.cs file for the program code that makes the application work you will find that it is completely empty. All the work is now being done by our tiny little class, which is connected to the display items. Changes to the content of the TextBox elements used to hold the input values cause events inside the Adder class. This class then changes the content of the AnswerValue which is bound to the ResultTextBlock on the display.

The solution in *Demo 06 AddingMachine with data binding* implements an adding machine which uses data binding to connect the inputs to the AdderClass.

Data Binding using the Data Context

You might be looking at data binding and thinking “This looks OK, and I understand how it works, but you do seem to have to do a lot of work to add the resources and everything”. The good news is that there is a way to short circuit this process which makes the whole thing really easy. It means adding a statement to our code but I reckon we can live with this.

Setting the Data Binding in XAML

We have seen how the elements in a control can be bound to object properties. We can just bind an object property to an element property and Silverlight will provide the glue to tie these together.

```
<TextBox Height="72" HorizontalAlignment="Left" Margin="8,19,0,0" Name="firstNumberTextBox"
Text="{Binding TopValue, Mode=TwoWay}" VerticalAlignment="Top" Width="460"
TextAlignment="Center" >
```

The XAML above shows how this works. I've highlighted the data binding for the first number textbox in the XAML above. This gives the name of the property that the text will bind to and the mode of the binding. We have set the binding to `TwoWay` because we want program changes (when the program has calculated a result) to be bound to the display. Sometimes we only want to use a display element to show a value (for example the result of the calculation). In that situation we can use `OneWay` binding.

Setting the DataContext

Now all we have to do is create an `AdderClass` instance and set the `DataContext` of the element that contains the `firstNumberTextbox` to this instance. We can do this in the constructor for the main page.

```
// Constructor
public MainPage()
{
    InitializeComponent();

    AdderClass adder = new AdderClass();
    ContentGrid.DataContext = adder;
}
```

The `DataContext` for a Silverlight element identifies the object that contains all the properties that have been bound to by that control and any that the control contains.

The `ContentGrid` on the page contains our `TextBox` and `TextBlock` elements and now every binding that they contain will be mapped onto the `AdderClass` instance that we created.

In other words we have replaced all the hard work with property panes and resources with a small amount of XAML and a single statement.


The solution in *Demo 06 AddingMachine with data binding* implements an adding machine which uses data binding to connect the inputs to the `AdderClass`.

4.3 Managing Application Page Layout

Now we are going to take a break from data binding for a while and consider some other issues of application layout, particularly with the Windows Phone in mind. We are going to start with phone orientation and then go off to consider some very interesting and useful ways you can get a display to design itself.

Landscape and Portrait Programs

You might have noticed that there are two ways that users can hold the device, landscape (device on its side) and Portrait (device held upright). The phone itself is able to detect the way it is being held and some built in applications respond correctly when the phone is tipped. Unfortunately our adding machine doesn't handle this very well:



Now, there is nothing wrong with only working in one orientation, the Start Menu in Windows Phone only works if the phone is held in the portrait orientation. But remember that if you decide to support both ways of holding the phone you will have to spend twice as much time designing the user interface and then add the code to respond when the user tilts the phone. The bottom line (for me at least) is that if your customer says “You know, it would be great if the program worked when the phone was held on its side” you don’t immediately say “I can do that”. Instead you should say “That will cost extra”. Because it will. However, if you have been offered big bucks to add this feature, it is useful to know the best way to do this.

Landscape and Portrait Selection in MainPage.XAML

A page can tell Silverlight the orientations that it supports. This is expressed as a property of the PhoneApplicationPage:

```
SupportedOrientations="Portrait" Orientation="Portrait"
```

The settings that you see above are the ones that Visual Studio puts in the page when it is created. This means that the page only work in Portrait mode and its initial setting is Portrait, which makes sense. We can modify this if we want to:

```
SupportedOrientations="PortraitOrLandscape" Orientation="Portrait"
```

This is a bit of an improvement, that when the program runs we can now tip the screen and see this:



The technical term for this is “Close, but no cigar”. Our program now redraws the screen when the orientation changes, but because all the components are positioned absolutely on a grid, some of them are in the wrong place, and the answer and the calculate button don’t fit on the screen any more.

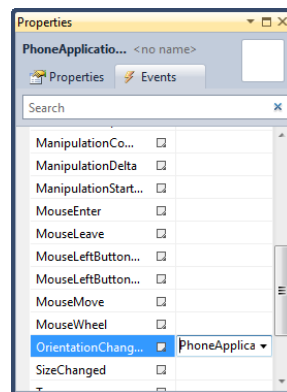
```
<TextBox Height="72" HorizontalAlignment="Left" Margin="8,19,0,0"
  Name="firstNumberTextBox" Text="0" VerticalAlignment="Top" Width="460"
  TextAlignment="Center" />
<Button Content="equals" Height="72" HorizontalAlignment="Left"
  Margin="158,275,0,0" Name="equalsButton" VerticalAlignment="Top" Width="160"
  Click="equalsButton_Click" />
```

These two components have Margin values that are used to precisely position them on the screen. Unfortunately, some of the positions are off the bottom of the screen. Silverlight will not complain about this, but the components won’t be displayed either.

There are two ways we can fix this. One way is to provide a completely different layout for the landscape screen and switch between the two layouts when the orientation changes.

Orientation Changed Event

We have already seen how Silverlight elements can generate events. We attached an event to the Click event generated by the Button that the user presses when they want the result of a calculation. We can also attach event handlers to other events too. The PhoneApplicationPage can generate a whole set of events, including one which fires when the orientation of the page changes.



Above you can see how we can add a method to this event. When the orientation of the phone changes this method will be called to adjust the display. We just have to add some code that does this:

```
private void PhoneApplicationPage_OrientationChanged(
    object sender, OrientationChangedEventArgs e)
{
    if (e.Orientation == PageOrientation.PortraitUp)
    {
        setPortrait();
    }
    else
    {
        setLandscape();
    }
}
```

This code uses the parameter to the method to detect the new orientation and call the appropriate method. The methods modify the margins of the controls to lay them out properly:

```
private void setLandscape()
{
    firstNumberTextBox.Margin = new Thickness(8, 19, 0, 0);
    firstNumberTextBox.Width = 207;
    secondNumberTextBox.Margin = new Thickness(252, 19, 0, 0);
    secondNumberTextBox.Width = 207;
    plusTextBlock.Margin = new Thickness(221, 35, 0, 0);
    resultTextBlock.Margin = new Thickness(538, 35, 0, 0);
}
```

This is the setLandscape method. The setPortrait method is very similar. Controls are positioned by setting a Margin around them. The margin of a control is given by a Thickness value which has margin size and a border size. We don't want to actually draw any borders for these controls, and so the border size values are set to 0. If I add these methods to my new application it now works fine in either mode:



There is only one tiny difference, which is that when a “proper” Windows Phone application is rotated it often shows a really nice “spinning” animation as it changes. You can do this kind of thing if you like, but you need to use a tool like Expression Blend to create the animation.

The solution in *Demo 07 AddingMachine with auto orientation* implements an adding machine which automatically handles orientation changes.

Using Containers to Layout displays

We now know how to make applications that adjust the properties of the controls under program control. However, we still have to come up with the values to make the layouts work. Fortunately Silverlight can help us with this as well; in fact it can even automatically lay out an entire display for us. This can save a lot of work. The key to understanding how this works is to take a look at how Silverlight containers can be used to hold other items. There are a number of different container types in Silverlight. We are going to use the [StackPanel](#) container.

A [StackPanel](#) contains a stack of Silverlight elements. It lays them out in the order that they are given. It can be vertical (down the display) or horizontal (across the display). Rather than giving each of our display elements a margin that positions them on the screen, we can instead get a [StackPanel](#) to do the layout for us:

```
<StackPanel>
  <TextBox InputScope="Digits" Height="72" HorizontalAlignment="Center"
    Name="firstNumberTextBox" VerticalAlignment="Top" Width="460"
    TextAlignment="Center" />
  <TextBlock Height="56" HorizontalAlignment="Center" Name="plusTextBlock"
    Text="+" VerticalAlignment="Top" FontSize="32" Width="25" />
  <TextBox InputScope="Digits" Height="72" HorizontalAlignment="Center"
    Name="secondNumberTextBox" VerticalAlignment="Top" Width="460"
    TextAlignment="Center"/>
  <TextBlock Height="46" HorizontalAlignment="Center" Name="resultTextBlock"
    VerticalAlignment="Top" FontSize="30" Width="160" TextAlignment="Center" />
</StackPanel>
```

The great thing about this is that we don't have to decide where on the screen each item needs to go, they are just stacked on top of each other. I've changed the alignment of each of the elements to “Center” so that if the [StackPanel](#) gives an item a particular area of the screen it will centralise itself in that area. When we run the program with the layout you can see below:



Each of the items is displayed in a stack. One very useful side effect of this approach is that the `StackPanel` will re-arrange the elements if the screen dimensions change. This means that if the orientation of the phone is changed it automatically adjusts:



Note that because the items are centered they now draw themselves in the middle of the landscape screen.

We can add a `StackPanel` to a display design by dragging it from the Visual Studio Toolbox onto the design surface or by typing the XAML text by hand. If we want to make a horizontal `StackPanel` we just have to add an additional attribute:

```
<StackPanel Orientation="Horizontal">
```

We can put one `StackPanel` inside another to make complex displays, as we shall see next.

There are other container types available. One type, which is used in the Windows Phone pages themselves, is the `Grid`. This allows you to create a rectangular grid of elements, each of which contains one or more display items positioned within it. In fact the Windows Phone program display that we have been using up to now has been one element in the grid that makes up the entire display area.

The solution in *Demo 08 AddingMachine with StackPanel* implements an adding machine which uses a `StackPanel` to automatically handle orientation changes.

4.4 Displaying Lists of Data

Displaying single items is all very well, but we would like to be able to display lists of things. Which brings us to another problem; we need to have some things to list.

Creating a Customer List

We are going to make a little customer management program. It will be very simple, just holding the customer name, address and ID number, but we will expand it later, even going as far as putting the whole thing into an SQL database. Our first version of the customer class looks like this:

```

public class Customer
{
    public string Name { get; set; }
    public string Address { get; set; }
    public int ID { get; set; }

    public Customer(string inName, string inAddress, int inID)
    {
        Name = inName;
        Address = inAddress;
        ID = inID;
    }
}

```

I have implemented the data fields as properties, we will see why this is a good idea later. The `Customer` class just contains a single method, the constructor, which creates an instance of the class from three input values.

Now we need to store a list of customers. The best way to store lists of things is to use the `List` collection class provided by C#. We can create a `Customers` class that contains the name of the customer list as a string and the list of customers:

```

public class Customers
{
    public string Name { get; set; }

    public List<Customer> CustomerList;

    public Customers(string inName)
    {
        Name = inName;
        CustomerList = new List<Customer>();
    }
}

```

The `Customers` class only contains a single method which is the constructor for the class. This sets the name of the class and then creates an empty list to store them in. If I wanted to create a customer list for my mail order customers I could write it like this:

```
Customers mailCustomers = new Customers("Mail Order Customers");
```

We now have two classes, one which can hold a single customer and one that can hold a large number of them.

Making Sample Data

The next thing we want is a whole bunch of customers. We could type these in by hand, and spend many happy hours inventing interesting names, but this would be silly. A better way of doing this is to create a method to do the hard work for us. We can put the method in the `Customers` class so that we can ask the class to give us some test data:

```

public static Customers MakeTestCustomers()
{
    string [] firstNames = new string [] { "Rob", "Jim", "Joe",
                                           "Nigel", "Sally", "Tim" };
    string[] lastsNames = new string[] { "Smith", "Jones",
                                         "Bloggs", "Miles", "Wilkinson", "Brown"
    };
}

```

```

};
Customers result = new Customers("Test Customers");
int id = 0;

foreach (string lastName in lastsNames)
{
    foreach (string firstname in firstNames)
    {
        //Construct a customer name
        string name = firstname + " " + lastName;
        //Add the new customer to the list
        result.CustomerList.Add(new Customer(name,
            name + "'s House", id));
        // Increase the ID for the next customer
        id++;
    }
}
return result;
}

```

This code will create 36 customers by combining all the first and last names that are given. If we want many hundreds of customers we just have to add more name strings. Each customer has a unique ID number. The `MakeTestCustomers` method has been made static, so that we don't need to have a customers list before we can create the test data. The customer list that is created is called "Test Customers". We can get our list very simply:

```
customerList = Customers.MakeTestCustomers();
```

Whenever I make something that is supposed to hold a large number of items the next thing I do is create a quick way of getting a bunch of them. If you are worried about this affecting the size of the program you can use conditional compilation so that the `MakeTestCustomers` method is only included in the class if a particular debug symbol is defined when the program is built.

I hope that this illustrates a very important principle. As soon as you make something you need to create a way of testing that thing and proving it works well. If we were selling our customer management system to a shop owner she would not be impressed if we showed it working with only a couple of customers that we had entered by hand. She would want to see the program managing many thousands of entries. It is not that hard to create code that does this, as you see above, and it makes the testing of the program much easier, and more convincing.

Using the StackPanel to display a List

The first thing we are going to do with our customer list is just display all of the customer names on the phone screen. We can put each name into a `TextBlock`. We've created `TextBlock` elements before; they are how our adding machine displays the answer to the user. What we want is something that will lay out the textboxes for us automatically, even if we don't know just how many there will be when the program runs. We know how to do this as well. We can use a `StackPanel` to lay out a large number of elements automatically. This is what the XAML for the display looks like:

```

<StackPanel HorizontalAlignment="Left"
    Margin="0,0,0,0" Name="customersStackPanel"
    VerticalAlignment="Top"/>

```

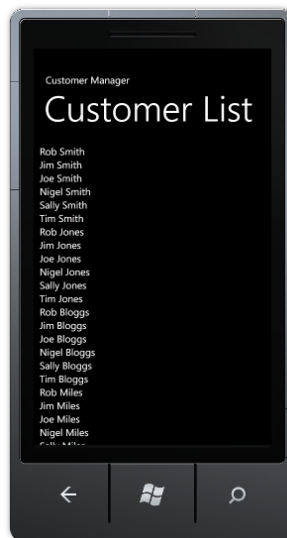
You may be wondering where all the `TextBlock` values are? The answer is that we can create them in the program:

```

customers = Customers.MakeTestCustomers();
foreach (Customer c in customers.CustomerList)
{
    TextBlock customerBlock = new TextBlock();
    customerBlock.Text = c.Name;
    customersStackPanel.Children.Add(customerBlock);
}

```

This code creates a set of test customers and then works through them. For each customer it creates a new `TextBlock` instance called `customerBlock`. It then sets the `Text` property of `customerBlock` to the name of the customer. Finally it adds `customerBlock` to the children of `customersStackPanel`. This illustrates a very important point. We have seen that you can create display elements by describing them in the XAML for a page. However, we can also create display elements by constructing them as objects in our programs. The above program will work well for any number of customers.



Above you can see how the `StackPanel` has displayed a stack of customer names. However, we have a problem. There are some customers below Nigel Miles, but we can't see them because they have "fallen off" the screen. The height of a `StackPanel` will grow to accommodate the items that are placed in it, but they might not all fit on the display.

Silverlight will not produce an error if we try to draw something that won't fit on the screen, but the display won't look right. We need to find a way of allowing the user to scroll the display up and down to see the rest of the list. We can do this with a `ScrollViewer`. This provides a scrollable window which lets the user pan around a displayed item.

```

<ScrollViewer>
    <StackPanel HorizontalAlignment="Left" Margin="9,6,0,0"
        Name="customersStackPanel" VerticalAlignment="Top"/>
</ScrollViewer>

```



Above you can see the scrolled view of the list.

The solution in *Demo 09 CustomerManager List* implements the display of the customer list as we have seen above. If you run this you will see that scroll bars are automatically displayed, and that the screen can be flicked and squeezed in a most pleasing way. All this behaviour is provided for you by the `ScrollViewer` control.

Using the `ListBox` to display lists of items

We have seen that it is perfectly OK for an application to create Silverlight elements when it runs. Display elements can be described inside the XAML file, in which case they are made automatically when our program starts; or a program can create new element instances and add them to controls on the screen as it runs. We have just used this ability to display a list of customers but our program had to do all the hard work. It had to create all the display objects and then add them to the `StackPanel` where they were to be displayed.

Last time we had this problem we found that Silverlight allowed us to use `Data Binding` to make life a lot easier. `Data binding` gets Silverlight to do all the hard work for us. We just have to bind the content of a Silverlight object to the data we want to display and it just works. Up until now we have just connected individual elements to single values (for example the result from the `AddingMachine` program). Now we are going to expand this to bind the contents of a customer object to a list.

However, this is a bit more complicated than the binding we did last time. The good news is that is actually very easy, and really, really powerful once you get the hang of it. So, let's do that.

Data Binding Single Items

In the adding machine program we used data binding to read numbers from the user (the content of the `TextBox` elements were bound to the value which needed to be added together) and we also used data binding to show the result of the calculation (the `AnswerValue` was bound to the `resultTextBlock`). We did all this in the XAML that described the screen display.

```
<TextBlock Height="46" HorizontalAlignment="Left" Margin="158,208,0,0" Name="resultTextBlock"
Text="{Binding Path=AnswerValue}" VerticalAlignment="Top" FontSize="30" Width="160"
TextAlignment="Center" />
```

Above you can see the binding that we used in the adding machine to display the result of a calculation. The text in the `resultTextBlock` display element was bound to the `AnswerValue` property in the `Adder` class. When the class worked out a new answer, it was automatically displayed by data binding.

Creating a Data Template

We'd like to bind a `Customer` to a display element, but we can't just bind one directly because the value of a customer includes a bunch of different components. A `Customer` value contains a name, an address and a customer number. What we'd really like to do is design a template which sets out how the customer is to be displayed. This would let us choose which parts of a customer we want to display on the screen (perhaps we want to omit the customer number) and how they are formatted. Perhaps something like this:

```
<DataTemplate>
  <StackPanel>
    <TextBlock Text="{Binding Name}"/>
    <TextBlock Text="{Binding Address}"/>
  </StackPanel>
</DataTemplate>
```

This is a data template. It is a lump of Silverlight that describes how some data is to be displayed. We are telling Silverlight that we want to display the data in the form of a `StackPanel` that contains two items, the name and the address. It should result in a display a bit like this, with the two text strings stacked above each other:

```
Rob Miles
Rob Miles's House
```

Note that we are only displaying the name and address of a `Customer`. There is no data binding for the Customer ID value and so it will not be shown.

Using a DataTemplate in a list

We use a data template when we to tell Silverlight how to display properties of an object. We are going to use a template to display each of the customers in the list. We do this by placing this `DataTemplate` inside the `ItemTemplate` for the `ListBox` that will display our customers:

```
<ListBox Name="customerList">
  <ListBox.ItemTemplate>
    <DataTemplate>
      <StackPanel>
        <TextBlock Text="{Binding Name}"/>
        <TextBlock Text="{Binding Address}"/>
      </StackPanel>
    </DataTemplate>
  </ListBox.ItemTemplate>
</ListBox>
```

OK. Time to step back little. The Silverlight `ListBox` display element called `customerList` is going to display a list of our customers. To do this it needs to know how to lay out the display of a single `Customer`. A `ListBox` uses an `ItemTemplate` element to get the design of a list item, and this case it will be using a `DataTemplate` to do this.

So, once we have set up this Silverlight the `ListBox` just needs to be told the collection of data it is to display:

```
customers = Customers.MakeTestCustomers();
customerList.ItemsSource = customers.CustomerList;
```

This code creates a list of test customers and then sets the `ItemsSource` of the `customerList` to the list of customers. The `ListBox` will display each object in the collection as an element in the list.



Our program doesn't have to do anything other than just set the `ItemsSource` property of the `ListBox`. Everything else is done automatically. A `ListBox` even has scroll bars provided automatically, so if the list is too large for the screen the user can scroll up and down it.

From this example you should see how easy it is to display collections of data. You just have to design the template and then use data binding to get your data displayed.

Adding style to a display

The display above works OK, but it is not very interesting. It is also hard to see which are the customer names and which are the customer addresses. We could decide on different font sizes and styles for the name and address but we can also make use of some built in styles provided as resources for Windows Phone applications.

These styles have the advantage that they will work in different colour schemes and light and dark backgrounds. We add the style information to the `TextBox` items in the `DataTemplate`:

```
<DataTemplate>
  <StackPanel>
    <TextBlock Text="{Binding Name}"
              Style="{StaticResource PhoneTextExtraLargeStyle}"/>
    <TextBlock Text="{Binding Address}"
              Style="{StaticResource PhoneTextSubtleStyle}"/>
  </StackPanel>
</DataTemplate>
```

We are using `ExtraLargeStyle` for the customer name and `SubtleStyle` for the address. This has the result of making the display look much better.



Selecting items in a ListBox

We can now display a list of items really easily. The next thing we want to do is make a proper data editing application. You will have seen this kind of application before. When the user selects one of the items in the list we want to open up a screen that allows this item to be edited:



The screens above show how this will work. The user scrolls up and down to find the customer they want to work with. They then select the required customer. The application then displays an edit screen for that customer. We can use this pattern any time we have a list of things that the user can pick one from.

It is very easy to get a `ListBox` to produce an event when an item is selected:

```
<ListBox Name="customerList"
         SelectionChanged="customerList_SelectionChanged">
```

This is where the `customerList` is declared in the XAML for the main page. The event handler must then be created inside the `MainPage.xaml.cs` source file.

```
private void customerList_SelectionChanged(object sender,
                                           SelectionChangedEventArgs e)
{
    // when we get here the user has selected a customer
    Customer selectedCustomer =
        customerList.SelectedItem as Customer;

    MessageBox.Show(selectedCustomer.Name + " is selected");
}
```

The `customerList` provides a property called `SelectedItem` which is a reference to the object that has been selected. The method must convert this to a reference to a customer and then the program can display a message box containing this name.

The solution in *Demo 10 CustomerManager Binding List* implements the display of the customer list as we have seen above. The Silverlight for `MainPage.xaml` contains the `ListDisplay` element and the template, and the program uses data binding to display the result. If we select an item in the list the program displays a `MessageBox` with the name of the item selected.

This pattern is very useful if you want to display a list of anything and have the user choose which one they want to work on.

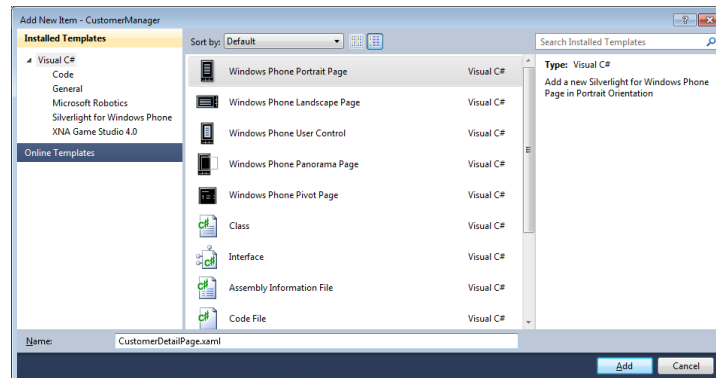
The next thing we want to do for our customer manager application is take the user to a new page in our application that allows them to interact with the settings for this customer. To do this we need to know how to perform page navigation in Silverlight.

4.5 Pages and Navigation

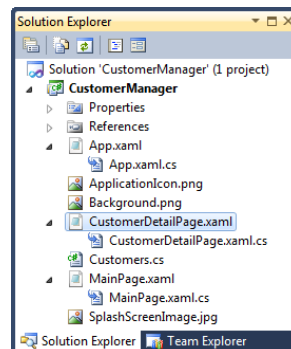
At the moment all the programs we have written have run on a single Silverlight page. However, many applications run over a number of pages. The Customer List application that we are creating will need a different page that will be used to display the information about the customer we are working on.

Adding a New Page

It is easy to add a new page to an application. If we select **Project > Add New Item** from within Visual Studio we are presented with a menu of possible items that we might like to add.



If we select **Windows Phone Portrait Page** we get a new page added. Before adding a page it is sensible to set the name to something more meaningful than “Page1”. We can change the name of the item later if we like by renaming it. An application can contain as many pages as we need.



Here we can see that the project now contains a **CustomerDetailPage**. New pages themselves are edited in exactly the same way as **MainPage**. When the page is open we can drag elements onto it from the Visual Studio ToolBox, we can manipulate their properties and we can edit the XAML directly if we wish.

Navigation between Pages

The Silverlight model for moving between pages has more in common with the internet than Windows Forms. If you are used to writing Windows Forms applications you will be familiar with methods such as “**Show**” and “**ShowDialog**” which are used to allow one form to cause the display of another.

Silverlight doesn’t work like that. When you think about moving around a Silverlight application you should think in terms of navigating from one page to another. Each page has an address expressed as a **uri** (Uniform Resource Indicator). The **NavigationService** object provides methods that perform the navigation for us.

```
// Navigate to the detail page
NavigationService.Navigate(new Uri("/CustomerDetailPage.xaml",
    UriKind.RelativeOrAbsolute));
```

This code will cause the application to switch to a page called `CustomerDetailPage`. Note that the `Navigate` method is given a uri value which contains the name of the page and also the kind of Uri being used.

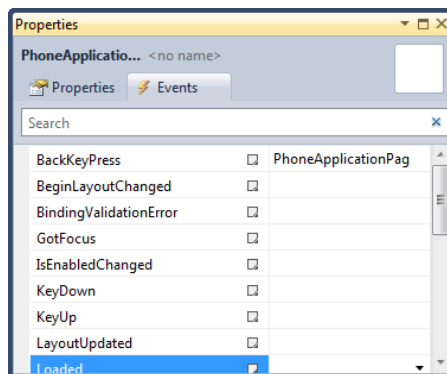
When we create the uri we also need to select what kind of uri it is, we are setting this uri to be `RelativeOrAbsolute`, even though this particular Uri is relative to the present one. We do this to ensure maximum flexibility when loading resources.

There is an issue here that you need to be aware of though. The Uri value is given as a string of text in your program. If you miss-type the string (for example put `"/CuustomerDeytailsPagerino.whamlxaml"`) the program will compile just fine, but then when it runs the navigation will fail with a run time error.

Using the Back button

Users will expect the Back button on Windows Phone to return to the previous page. This is exactly what it does. If the user has navigated to `CustomerDetailPage` as shown above, and then presses the Back button they will be returned to the page they came from. If the user presses Back when they are at the `MainPage` the application is ended.

Sometimes we want to override this behaviour. We might not want our user to leave the page if they have not saved the data they have entered. We can do this by adding an event handler to the `BackKeyPress` event:



The event handler method can cancel the back behaviour if required.

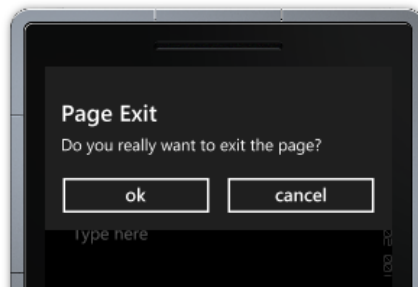
```
private void PhoneApplicationPage_BackKeyPress(object sender,
    System.ComponentModel.CancelEventArgs e)
{
    e.Cancel = true;
}
```

If the method above is bound to the `BackKeyPress` event the user will not be able to use the Back key to move away from the page.

Sometimes you want the user to confirm that they want to move away from the page. You can use the `MessageBox` class to do this:

```
private void PhoneApplicationPage_BackKeyPress(object sender,
    System.ComponentModel.CancelEventArgs e)
{
    if (MessageBox.Show("Do you really want to exit the page?",
        "Page Exit", MessageBoxButton.OKCancel)
        != MessageBoxResult.OK)
    {
        e.Cancel = true;
    }
}
```

This code displays a message box and asks the user if they really want to leave the page. If they don't want to leave the back button is cancelled. This is what the screen looks like when this code runs.



You can use this form of `MessageBox` any time you want to put up a quick question for the user to answer.

Passing Data between Pages

Each Silverlight page is distinct and separate from any other. They do not share data. Sometimes you want to pass data from one page to another. If the data is a simple string of text the easiest way to do this is by adding it to the uri that is passed into the destination page. This technique is based on the way that web pages send queries to a web server.

As an example, we could consider our Customer Manager. When the user selects a customer from the list we would like to pass the name and address of the customer which has been selected into the edit page.

```
private void customerList_SelectionChanged(object sender,
                                           SelectionChangedEventArgs e)
{
    // Get the selected customer from the list
    Customer selectedCustomer = customerList.SelectedItem as Customer;

    // Build a navigation string with the customer information in it
    NavigationService.Navigate(new Uri("/CustomerDetailPage.xaml?" +
        "name=" + selectedCustomer.Name + "&" +
        "address=" + selectedCustomer.Address,
        UriKind.Relative));
}
```

This is the method that runs when the user selects a customer. The first thing it does is get the selected customer. Then it creates a `Uri` which contains the name and the address of the selected item. This is then used to navigate to the customer details page.

`"/CustomerDetailPage.xaml?name=Sally Smith&address=Sally Smith's House"`

This is the format of the `Uri` that would be created for Sally Smith. We can see that the query is the part after the `?` character in the `Uri`. This is then followed by two “name-value” pairs, separated by the `&` character. The first item is the name part of the query, and the second the address.

Now we need to understand how a receiving page can get hold of this information and display it. To do this we have to consider events that are raised when navigating Silverlight pages.

Using Page Navigation Events

There are two really important events in the life of a Silverlight page. One is fired when the page is navigated to (`OnNavigatedTo`) and the other is fired when the page is left (`OnNavigatingFrom`). To get control when these events fire we need to override these methods in our page.

When we create a new page the class that controls it is derived from a parent class called `PhoneApplicationPage`:

```
public partial class CustomerDetailPage: PhoneApplicationPage
{
    // CustomerDetailPage methods go here
    // we override the ones that we want to provide
    // our own behaviours for
}
```

This is the declaration of `CustomerDetailPage`. We add new methods to the page and we can override existing methods and add new behaviours to them. We want our application to get control when the user navigates to the page. This is the point where we want to load information from the uri and display it in a control on the page. Our replacement `OnNavigatedTo` method looks like this:

```
protected override void OnNavigatedTo(
    System.Windows.Navigation.NavigationEventArgs e)
{
    string name, address;
    if (NavigationContext.QueryString.TryGetValue("name",
                                                out name))
        nameTextBlock.Text = name;
    if (NavigationContext.QueryString.TryGetValue("address",
                                                out address))
        addressTextBlock.Text = address;
}
```

The method uses the `QueryString` in the `NavigationContext` object and tries to get the data values out of that string. It asks for each value by name and, if a value is recovered, displays it in a text block. Each time the user returns to this page the method runs and puts the latest content onto the page. We can use the `OnNavigatedTo` event any time we want to set up a page when the user moves to it. We can pass as many named items into the destination page as we like.

One thing to remember is that we must make sure that the names match up in both pages. The compiler will not detect mistakes such as spelling the word `"address"` as `"addresss"` in the above code. If we did this the program would run correctly, but the address information will not be displayed properly.

We can use this code to create a two page application that lets the user select a customer and see the details page for this customer. Unfortunately there is a tiny problem with page navigation if we do this. There is one behaviour that will be confusing:

1. Select an item from the list
2. Navigate to the customer details page.
3. Press the back button to return to the list
4. Select the same item on the list.

If the user does this they will notice that they are not returned to the same customer. They have to click a different customer and then return to the original one to select it. The reason for this is that the program is bound to the `SelectionChanged` event. As far as it is concerned, the selection has not changed when the user re-selects the same one. We can fix this by clearing the selection whenever this page is navigated to:

```
protected override void OnNavigatedTo(
    System.Windows.Navigation.NavigationEventArgs e)
{
    customerList.SelectedItem = null;
}
```

Now, if the user selects the same item again it will count as a valid selection event. Of course, setting the selection to null is actually a selection event and so if our program tries to use this property it will fail with a null exception. This is easy to fix; the event

handler just has to make sure that an element in the list is selected before it tries to use it.

```
private void customerList_SelectionChanged(object sender,
                                           SelectionChangedEventArgs e)
{
    // Abandon if nothing selected
    if (customerList.SelectedItem == null) return;

    // Rest of method here
}
```

The above version of the method returns if there is no selected item.

The solution in *Demo 10 CustomerManager Query String* is a two page Silverlight application. You can select a customer and see this customer displayed in the customer details page.

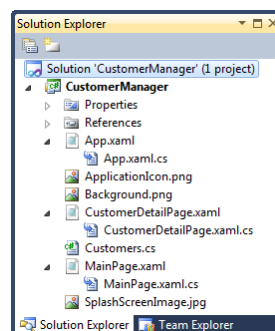
Sharing Objects between Pages

Passing strings of text between pages is useful, but you often need to share larger items which contain structured data. Often your user will be working with a “thing” of some kind; perhaps a document or a game, and the different pages will provide different views of that object. In our Customer Manager application we would really like to pass a Customer instance into the display page, so that it can work on the content of that customer. At the moment the display page is just being given the values of the items in the selected customer which are being passed as strings.

What you really want is an object which is common to all the pages and can be accessed by any of them. You could put your shared data into the object and they could all make use of it. Turns out that this is very easy to do, because all Windows Phone Silverlight programs have such a thing as part of the way they work. It is called the `App.xaml` page.

The App.xaml page

When we create a Silverlight application we get a `MainPage.xaml` and an `App.xaml` page.



We can regard `App.xaml` as the container for the phone application. It provides the context in which the pages are displayed. It doesn't ever get drawn on the screen, but it is the starting point for the application when it is loaded. The `App.xaml.cs` file contains the methods that start an application running. It also contains handler methods for application lifecycle events which we will be using later.

We can edit the `App.xaml.cs` file and add our own code and data members. If we want to store any global data for use by all the pages in our application we can declare it here.


```
public partial class App : Application
{
    // To be used from all pages in the application
    public Customer ActiveCustomer;
}
```

Here we have added an `ActiveCustomer` value that we want to be able to use in all the pages. It will hold a reference to the customer that is currently being edited by the user.

Any Silverlight page can obtain a reference to the `App` which it is part of:

```
protected override void OnNavigatedTo(
    System.Windows.Navigation.NavigationEventArgs e)
{
    base.OnNavigatedTo(e);

    // Get the parent application that contains the active customer
    App thisApp = Application.Current as App;

    // Set the data context for the display grid to the active
    // customer
    customerDisplayGrid.DataContext = thisApp.ActiveCustomer;
}
```

The `Current` member of the `Application` class is a reference that can refer to objects of type `Application`. When the application starts it is set to refer to the currently executing application instance.

The `App.xaml.cs` class defines a child class `App` that extends the `Application` parent. To get hold of objects defined in the `App` class we have to make a reference of type `App` and then make this refer to the current application. Above you can see how the `as` operator is used to convert the reference to one with the correct type. This is the code that runs inside the `customerDetailsPage` when this page is navigated to. It gets the `ActiveCustomer` out of the application and sets the data context on the `customerDisplayGrid` to refer to the selected customer. We have already used data binding to allow us to directly connect all the customers to the customer selection page; here we are using it to display the values on a page.

Data Design

Since there may be other pages in the application that want to use the customer list it makes sense to put this list into the `App.xaml.cs` class as well. The test data is then set up when the application starts running. In the finished application this is the point where the data would be loaded from storage.

The solution in *Demo 11 CustomerManager Shared Object* is a two page Silverlight application that uses a shared variable to allow a user to select one of the items from the list and view this on the details page. It also uses a Silverlight Grid control to layout the customer display. This is very useful if you want to align objects on a page.

4.6 Using ViewModel classes

At the moment our program is very good for displaying the customer, but it doesn't actually allow us to edit their details. The user can change the contents of the textboxes on the customer display page but the changes will not be reflected in the data in the program.



What we really want to happen is that when the user updates the Address field, perhaps to "Nigel Jones's New House", this is automatically put back into the customer that is presently selected. We can do this with Data Binding; it is how we got the numbers out of the text boxes in the adding machine that we created earlier.

The first thing we do is tell Silverlight that the binding of the text in the `TextBox` is two-way. This means that changes to the data will be displayed on the screen, and that changes in the `TextBox` will update the property in the class that is bound to it.

```
<TextBox Name="nameTextBox" Text="{Binding Name,Mode=TwoWay}"/>
```

Next we could add code to the `Customer` class to manage the binding events. This code would make sure that when we change the name the customer the text on the screen would change, and if the content of the `TextBox` changes the name property of the customer changes automatically. We did this in the Adding Machine with the `Adder.cs` class, and it seemed to work fine.

But first I think we need to stop and consider what we are doing.

Design with a ViewModel Class

We could add the data binding behaviours to the `Customer` class and from a programming point of view this would work fine. However, there are some reasons why it would not be a good idea from a design point of view.

Division of Purpose

There is a rule in software design that a class should usually do one thing; and one thing only. The job of the `Customer` class should be to hold customer data, not take part in an editing process. When we store or transfer customer data around we do not want to be worrying about methods and behaviours which are concerned with editing. If we transfer our customer onto another device for storage (and we will look at this later) any property binding behaviour will just be a waste of space, since it will not be used in this context.

Providing Validation and Data Conversion

Ideally something should make sure that what is entered into the form conforms to our business rules. These are things like "the name must contain only letters and spaces". We could put these rules into the `Customer` class, but it would be much easier if we could directly connect them to the data input process, since this is the point at which they are going to be used. It might also be useful to be able to perform conversion of some input values, for example dates and times, from the format that they are entered and displayed to the ones used for storage.

Providing an Undo Behaviour

This is a user interface issue, but it is also very important. If you look back at the edit screen you will notice that there is a big **Save** button. The idea is that the user will press **Save** when they have completed editing, and this is the point that the changes they have made will be committed to the data. However, there is also a large **Cancel** button which would close the screen down and effectively cancel any changes.

If we are using direct data binding this would be a problem, as our data object would have already been changed to reflect what the user had entered. This means that we would have to make a copy of the data on entry to the form and then put it back in the object. Since the reason we are using data binding is to stop us having to write this kind of code, it seems a pity to have to start writing it all over again, just because the user wants to be able to change their mind.

Creating a ViewModel Class

We can address these problems by creating a *ViewModel* class. This has the job of representing a view of the data that we wish to present. It takes care of the data presentation and binding to the form components and can also perform any data validation as required. We create the *ViewModel* class to represent the activities we are going to perform in a particular user interface. We can think of it as a kind of “glue” that links the user interface and the actual data objects. It contains the properties that we wish to work with in the user interface. I’ve written one below called *CustomerView*.

```
public class CustomerView : INotifyPropertyChanged
{
    private string name;

    public string Name
    {
        get {
            return name;
        }
        set {
            name = value;

            if (PropertyChanged != null)
            {
                PropertyChanged(this,
                    new PropertyChangedEventArgs("name"));
            }
        }
    }

    // Address information here

    private int id;

    public int ID
    {
        get
        {
            return id;
        }
    }

    public event PropertyChangedEventHandler PropertyChanged;
}
```

```

public void Load(Customer cust)
{
    Name = cust.Name;
    Address = cust.Address;
    id = cust.ID;
}

public void Save(Customer cust)
{
    cust.Name = Name;
    cust.Address = Address;
}
}

```

This looks a lot like the `Adder.cs` class that we created earlier. It creates events when properties are changed by our program, and the Silverlight controls can set new values in the class when they are changed by the user. Note that the ID property does not have a set behaviour. This is because the user is not allowed to change the ID of a customer. This is set when the customer is created and is unique for that customer.

Moving Data in and out of the Viewmodel

The `CustomerView` class also contains `Load` and `Save` methods. These are how the program sets the contents of the view class to the data being worked on. As an example, consider what happens when we start editing a customer:

```

// Viewmodel for the details page
CustomerView view = new CustomerView();

protected override void OnNavigatedTo(
    System.Windows.Navigation.NavigationEventArgs e)
{
    // Get the parent application that contains the active customer
    App thisApp = Application.Current as App;

    // Load the active customer into the viewmodel
    view.Load(thisApp.ActiveCustomer);

    // Set the data context for the display to the viewmodel
    customerDisplayGrid.DataContext = view;
}

```

When the program moves to the customer details page the `OnNavigatedTo` method is called by Silverlight to set up the page. This method finds the currently active customer (this is the one that will have been selected) and then loads this into the `CustomerView` being used on this page.

When the user presses the `Save` button the program must take the content of the viewmodel and put it back into the customer being edited:

```

private void saveButton_Click(object sender, RoutedEventArgs e)
{
    // Get the parent application that contains the active customer
    App thisApp = Application.Current as App;

    // Copy the data from the viewmodel into the active customer
    view.Save(thisApp.ActiveCustomer);

    // Go back to the previous page
    NavigationService.GoBack();
}

```

ViewModels and Testing

They ViewModel class looks nice, but you might be wondering why we bother with it. It looks as if we can get a very similar behaviour just by putting the members of the Customer class directly into the TextBox elements on the page and then reading them back if the user presses the Save button.

While this is true, it overlooks one important aspect of program development, and that is testing. If the customer data editor interacts directly with the display elements it is very hard to automatically test whether it works or not. Test code would have to enter things into textboxes and look at the content of other textboxes to see if the application was doing what it was supposed to do. This becomes even more difficult if the user interface has to perform validation and data translation.

However, with a ViewModel class doing this work it is much easier to test. A test environment can simply connect to the same elements in the ViewModel as the user interface and inject test messages into it. There is no need to test by entering text onto the screen because a test environment can simply simulate the data binding.

Page Navigation using the GoBack method

The `NavigationService` provides a `GoBack` method that takes the user back to the previous page. This is more efficient than using a URI to navigate because it does not recreate the page. If you navigate to a page using the `Navigate` method provided by the `NavigationService` class this will create a new page each time you visit it. This can slow your application down and also lead to the creation and destruction of lots of Silverlight elements that will need to be recovered by the Garbage Collector.

Furthermore the user will be a bit annoyed if, having scrolled down the list to select the item they want to work on, they are then taken back to the top of the list after an edit. I hate it when a badly written program does this kind of thing to me, and I'm sure you do to.

However, using the `GoBack` method to take us back to the original page does cause us problems. The difficulty is that the display on the original page will not reflect the changes that the user has made to the data. Consider the following sequence:

1. User selects customer Rob Miles
2. User changes name to "The Great Rob Miles" (and quite right too)
3. User Presses the Save button to perform the changes
4. Upon return to the selection screen the name is still "Rob Miles".
5. Confusion results.

The problem is that the `ListBox` is unaware that the content of a customer record has changed. When the customer list was displayed it made a copy of the data onto the screen and it does not know that this needs to be updated. The actual data has been updated, but the view of it does not reflect this. To solve the problem we have to go even deeper into Silverlight, and find out how to use an `ObservableCollection` to solve the problem

Observable Collections

We have seen that the best way to connect a data object to a Silverlight page is to put a ViewModel class in between the two. The ViewModel can be customised for the particular view required, and the data object can focus on just storing the data.

This is also true when considering collections of data. At the moment we are binding a list of customers directly to the `ListBox`. This is fine for just viewing the content of a list, but as we have seen, if any elements of the list change there we do not have a notification method that can be used to tell Silverlight that the display on the screen needs to be updated. What we need is another ViewModel class that can represent the list of customers and provide a way that the `ListBox` can be notified of changes to elements in it.

This is what the `ObservableCollection` was created for. It can be used to hold a collection of items and provides notification support so that if the content of the collection changes the `ListBox` can be bound to change events. We can create an `ObservableCollection` of customers very easily:

```
ObservableCollection<Customer> observableCustomers;

private void PhoneApplicationPage_Loaded(object sender,
                                         RoutedEventArgs e)
{
    // Get the parent application that contains the customers list
    App thisApp = Application.Current as App;

    // Make an observable collection of these
    observableCustomers =
        new ObservableCollection<Customer>(
            thisApp.ActiveCustomerList.CustomerList);

    // Set the list to display the observable collection
    customerList.ItemsSource = observableCustomers;
}
```

This is the code that runs when the main page of our Customer Manager starts. It creates a new list called `observableCustomers` from the customer list. It then sets the source of the `customerList` to this newly created list. Now, if the program makes changes to the list the display will update to reflect this.

Unfortunately the program still doesn't reflect changes we want. If the user updates a name in a customer this will not be reflected by the text on the display. This is because the `ObservableCollection` responds to the changes in the contents of the list, not changes to the data in an element in the list.

Actually, what you get is worse than just the display not updating correctly. What will happen is that if the `ListBox` ever needs to redraw that element it will go back to the original data and display it. This means that if the user scrolls up and down the list for a while there is a good chance that the correct data will appear eventually, as the `ListBox` will only create display objects that are visible, and may discard ones that have been moved off the screen. This can lead to some very confusing effects where the program seems to catch up after a while.

To get the view to update we have to do something to the content of the list to force the update to take place. One way to force an update is to remove something from the list and put it back again. Changes to the content of an observable collection generate events that are picked up by the `ListBox` that is displaying the list.

```
protected override void OnNavigatedTo(
    System.Windows.Navigation.NavigationEventArgs e)
{
    // Get the parent application that contains the customer list
    App thisApp = Application.Current as App;

    if (thisApp.ActiveCustomer != null)
    {
        // Find the customer in the list
        int pos = observableCustomers.IndexOf(
            thisApp.ActiveCustomer);

        // Remove it
        observableCustomers.RemoveAt(pos);
        // Put it back again - to force a change
        observableCustomers.Insert(pos, thisApp.ActiveCustomer);
    }
}
```

This is the code that runs when the user navigates back to the main page. If there is an active customer (i.e. one has been selected for edit) the method removes this customer

from the list and then puts it back in the same place. This forces the redraw of the item and the user will now see changes reflected immediately. Even better, because we are just updating the one item, the display will update really quickly. This is particularly important if we have a list which contains many thousands of customers.

Saving Data

When we used a `ViewModel` class to edit a customer we had `Load` and `Save` methods that put the customer into the `CustomerView` class and then recovered the updated values from it. In the case of the `ObservableCollection` we have to do something similar. Fortunately there are a number of ways that we can ask the `ObservableCollection` to give back the data it is displaying.

```
thisApp.ActiveCustomerList.CustomerList =  
    observableCustomers.ToList<Customer>();
```

The `ToList` method extracts the list of customers and returns it.

Observable Collections and Efficiency

You might be thinking that adding an `ObservableCollection` makes the program run more slowly and use up more memory. This is actually not the case. The collection is another set of references to the same customer objects that are in the customer store. This means that the only extra memory that is required is for the observable collection itself, which is not large.

If your program only needs to view the content of a list then you do not need to use the `ObservableCollection`. It is only when the contents of the list need to change because of updates to the data in the list that you need to add this.

The solution in *Demo 12 Complete CustomerManager* is a two page Silverlight application that allows the user to select and edit items in the list of customers. Changes to an item are updated in the list view and the user is returned to the original position at the end of each edit.

What We Have Learned

1. Programs can manipulate the properties of Silverlight elements to display information to the user of a program. This includes the position and colour of items on the display.
2. Some element properties are best set by editing the XAML directly. The XAML information for an element is structured in terms of properties which can be nested inside each other. The `TextBox` has a set of properties that give the initial settings for the keyboard to be used to enter data. These can be used to cause a numeric keyboard to be displayed instead of a text one.
3. Windows Phone can display messageboxes to the user. These can display multi-line messages and can also be used to confirm or cancel user actions.
4. Assets such as images can be added to a Windows Phone application as content or resources. An item of content is copied into the application directory as a separate file and can be used from there by the program. A resource item is embedded into the assembly file for an application. Content items do not slow the loading of assemblies but may be slower to load when the program runs. Resource items are more readily available to a program but they increase the size of the program assembly.
5. Silverlight elements can generate events in response to user actions. One such action is the `TextChanged` event produced by the `TextBox`.
6. Silverlight provides data binding support where the properties of an object in the program can be connected to those of a Silverlight display element. The

binding can be “one way”, where the display element is used to output the value of a program item or “two way” where changes to the item on the page result in an update of the bound property in the class.

7. It is possible to bind a collection of items to a ListBox Silverlight element to show a list of items. A data template is used to express how the individual data values in each item are to be displayed.
8. Silverlight applications can be made up of multiple pages. The navigation between pages is performed by a navigation helper class which is given the uri of the page to be navigated to. Simple items of text can be passed between pages by placing them in a query string attached to the uri.
9. Pages can receive events when they are navigated to and from. The OnNavigatedFrom event provides cancellation option so that a user can be asked to confirm the navigation.
10. Larger data objects can be shared between pages in an application by the use of the App class which is part of a Windows Phone application. Any Silverlight page in an application can obtain a reference to the application object that it is part of.
11. Programmers can create ViewModel classes that are set to the data being edited and bound to the Silverlight controls on a page. These act as “glue” between the actual data in the program and the particular presentational and editing needs of the user interface being developed.
12. The ObservableCollection mechanism allows changes to the content of a collection to be reflected in the list view displaying it.

5 Isolated Storage on Windows Phone

Each application or game has its own particular data storage needs. The storage can be as simple as a name – value pair (for example HighScore=100). On the other hand it you might be storing an entire product database along with all your customers. Or you might be storing text documents or pictures taken with the phone camera.

In this section we are going to explore the ways in which C# applications on the device can connect to and use data services provided by the phone.

5.1 Storing Data on Windows Phone

A proper application will need to store data which it can use each time it runs. A Silverlight application will need to hold settings and information the user is working on and an XNA game will want to store information about the progress of a player through the game along with high score and player settings. Windows Phone programs can use “isolated storage” to hold this kind of information. The storage is called “isolated” because it is not possible for one application to access data stored by another. This is different from a Windows PC where any program can access any file on the system.

The Isolated Storage area can hold very large amounts of data, up to the limit of the storage available in the phone itself. A Windows Phone will have at least 8G of built in storage which is shared between media (stored music, pictures and videos) and all the applications on the device.

When an application is removed from the phone all the isolated storage assigned to it is deleted. When an application is upgraded via Marketplace (i.e. we release a new version of our program or game) the isolated storage retains the original contents. If the data there must be updated for the new version our application must do this.

Using the Isolated Storage File System

You can use the isolated storage in the same way as you use a file system. The only difference is the way that a program makes a connection to the storage itself. However, once a program has a connection to the file system it can use it as it would any other, creating streams to transfer data to files and even using folders to create a structured filestore. A program can store very large amounts of data in this way.

As an example we could create a simple Silverlight application which gives the user a simple jotter where they store simple messages. This version of the program uses a single file but could easily be extended to provide a note storage system.



The user can type in jottings which are saved when the Save button is pressed. If the Load button is pressed the jottings are loaded from storage.

```
private void saveButton_Click(object sender, RoutedEventArgs e)
{
    saveText("jot.txt", jotTextBox.Text);
}
```

This is the code for the Save button. It takes the text out of the textbox and passes it into a call of the saveText method along with the name of the file to store it in.

```
private void saveText(string filename, string text)
{
    using (IsolatedStorageFile isf =
            IsolatedStorageFile.
                GetUserStoreForApplication())
    {
        using (IsolatedStorageFileStream rawStream =
                isf.CreateFile(filename))
        {
            StreamWriter writer = new StreamWriter(rawStream);
            writer.Write(text);
            writer.Close();
        }
    }
}
```

The saveText method creates a stream connected to the specified file in isolated storage and then writes the text out to it. If you have used streams in C# programs already you will be pleased to find that these work in exactly the same way. In this case the method creates a StreamWriter and then just sends the text out to that stream.

The read button uses a loadText method to do the reverse of this operation.

```
private void loadButton_Click(object sender, RoutedEventArgs e)
{
    string text;

    if ( loadText("jot.txt", out text ) )
    {
        jotTextBox.Text = text;
    }
    else
    {
        jotTextBox.Text = "Type your jottings here....";
    }
}
```

The `loadText` method tries to open a file with the name it has been given. It then tries to read a string from that file. If any part of the read operation fails the `loadText` method returns `false` and `jotTextBox` is set to a starting message.

If the file can be read correctly the `loadText` method returns `true` and sets the output parameter to the contents of the file. This is then used to set the text in `jotTextBox`.

Note that we are using an `out` parameter to allow the `loadText` method to return the string that was read as well as whether it worked or not.

```
private bool loadText(string filename, out string result)
{
    result = "";
    using (IsolatedStorageFile isf =
        IsolatedStorageFile.GetUserStoreForApplication())
    {
        if (isf.FileExists(filename))
        {
            try
            {
                using (IsolatedStorageFileStream rawStream =
                    isf.OpenFile(filename,
                        System.IO.FileMode.Open))
                {
                    StreamReader reader =
                        new StreamReader(rawStream);
                    result = reader.ReadToEnd();
                    reader.Close();
                }
            }
            catch
            {
                return false;
            }
        }
        else
        {
            return false;
        }
    }
    return true;
}
```

The `loadText` method will return `false` if the input file cannot be found or the read process fails. Note that it uses the `ReadToEnd` method to read from the file so that the program can read multi-line jottings.

The solution in *Demo 1 Jotpad* contains a Windows Phone Silverlight jotter pad that stores text in a file and loads it back. If you stop the program and run it again from the Application list in the emulator you will notice that the stored data is recovered when you press Load. Note that you could expand the save and load methods to store even more data simply by adding extra write and read statements and further parameters. Alternatively you could make the load and save methods work on an object rather than a number of individual items.

Using the Isolated Storage Settings Storage

Quite often the only thing a program needs to store is a simple settings value. In this case it seems a lot of effort to create files and streams just to store a simple value. To make it easy to store settings a Windows Phone program can use the Isolated Storage settings store. This works like a *Dictionary*, where you can store any number of name/value pairs in the isolated storage area.

An introduction to dictionaries

Dictionaries are a very useful collection mechanism provided as part of the .NET framework. A dictionary is made using two types. One is the type of the key, and the other is the type of the item being stored. For example, we might create a class called *Person* which holds data about, surprise surprise, a person:

```
class Person
{
    public string Name;
    public string Address;
    public string Phone;
}
```

This *Person* class above only contains three fields, but it could hold many more. We would like to be able to create a system where we can give the name of a *Person* and the system will then find the *Person* value with that name. A *Dictionary* will do this for us very easily:

```
Dictionary<string, Person> Personnel =
    new Dictionary<string, Person>();
```

When we create a new *Dictionary* type we give it the type of the key (in this case a *string* because we are entering the name of the person) and the type of the item being stored (in this case a *Person*). We can now add things to the dictionary:

```
Person p1 = new Person { Name = "Rob", Address = "His House",
    Phone = "1234"
};
```

```
Personnel.Add(p1.Name, p1);
```

This code makes a new *Person* instance and then stores it in the *Personnel* dictionary. Note that the code takes the name out of the person that has been created and uses that as the key. The program can now use a string as an indexer to find items in the dictionary:

```
Person findPerson = Personnel["Rob"];
```

The *findPerson* reference is set to refer to the *Person* instance with the name Rob. The dictionary does all the hard work of finding that particular entry. This makes it very easy to store large numbers of items and locate them later.

The dictionary will refuse to store an item with the same key as one already present. In other words if I try to add a second person with the name "Rob" the Add operation will throw an exception. A program will also get an exception if it asks the dictionary for a record that doesn't exist:

```
Person findPerson = Personnel["Jim"];
```

Fortunately there is also a mechanism for determining whether or not a given key is present in the dictionary:

```
if (Personnel.ContainsKey("Jim"))
{
    // If we get here the dictionary contains Jim
}
```

Dictionaries are very useful for storing name-value pairs. They remove the need to write code to search through collections to find things. We can use multiple dictionaries too, for example we could add a second dictionary to our Personnel system that would allow us to find a person from their phone number.

Dictionaries and Isolated Storage

The `IsolatedStorageSettings` class provides us with a dictionary based storage system for settings that we can use to store setting values. The settings dictionary stores a collection of objects using a string as the key.

Saving text in the isolated storage settings

We could convert our jotter to use the settings class as follows:

```
private void saveText(string filename, string text)
{
    IsolatedStorageSettings isolatedStore =
        IsolatedStorageSettings.ApplicationSettings;
    isolatedStore.Remove(filename);
    isolatedStore.Add(filename, text);
    isolatedStore.Save();
}
```

This version of the `saveText` method uses the filename as the key. It removes an existing key with the given filename and then adds the supplied text as a new entry. The `Remove` method can be called to remove an item from a dictionary. It is given the key to the item to be removed. If the key is not present in the dictionary the `Remove` method returns false, but we can ignore this. Once we have made our changes to the store we need to call the `Save` method to persist these changes.

Loading text from the isolated storage settings

The `loadText` method will fetch the value from the settings dictionary:

```
private bool loadText(string filename, out string result)
{
    IsolatedStorageSettings isolatedStore =
        IsolatedStorageSettings.ApplicationSettings;
    result = "";
    try
    {
        result = (string)isolatedStore[filename];
    }
    catch
    {
        return false;
    }
    return true;
}
```

The `loadText` method fetches the requested item from the settings storage. It is made slightly more complicated by the fact that, unlike the `Dictionary` class, the `IsolatedStorageSettings` class does not provide a `ContainsKey` method that we can use to see if a given item is present. The above method simply catches the exception that is thrown when an item cannot be found and returns false to indicate that the item is not present. Note that it must cast to string the value that is loaded from the

settings dictionary. This is because the dictionary holds objects (so that we can put any type into it).

We now have two ways to persist data on a Windows Phone device. We can store large amounts of data in a filestore that we can create and structure how we like.

Alternatively we can store individual data items by name in a settings dictionary. The isolated storage mechanism can be used by Silverlight programs and XNA games alike.

The solution in *Demo 02 Settings JotPad* contains a Windows Phone Silverlight jotter pad that stores text in the `IsolatedSettings` dictionary.

The Isolated Storage Explorer

When we are developing an application it is useful to be able to see the files that have been stored there. The Windows Phone Developer Toolkit contains a tool that lets you do just that. It is a program that is installed with the Windows Phone developer toolkit. On my system I can find it on the path below:

```
C:\Program Files (x86)\Microsoft SDKs\Windows Phone\v7.1\Tools\IsolatedStorageExplorerTool
```

The Isolated Storage Explorer is given the Product GUID (Globally Unique Identifier) of the application. This tells it which area of Isolated Storage to read. We can find the product GUID in the `WMAppManifest.xml` file for the application:

```
<App xmlns="" ProductID="{3363ac33-4f45-4b21-b932-fa2084b6deb0}"
Title="JotPad" RuntimeType="Silverlight" Version="1.0.0.0"
Genre="apps.normal" Author="JotPad author"
Description="Sample description" Publisher="JotPad">
```

To read the contents of an isolated store we issue the `ISETool` command with appropriate parameters.

```
ISETool ts xd 3363ac33-4f45-4b21-b932-fa2084b6deb0 c:\isoStore
```

This command takes a snapshot (`ts`) of the emulator (`xd`) isolated storage for the Jotpad application above and places it in the folder `c:\isoStore`.

There are lots of other things you can do with the explorer. The full description of the command is as follows:

```
Command Parameters: <ts|rs|dir[:device-folder]> <xd|de> <Product GUID> [<desktop-path>]
```

Where:

```
ts - (take snapshot) to download the contents of isolated store from Device to desktop.
rs - (restore snapshot) to upload the contents of isolated store from desktop to device.
dir- (list directory) to list the contents of the device folder or root if not mentioned
xd - Emulator.
de - Device.
Product GUID - (product ID) located in the WMAppManifest.xml file of the project.
desktop-path - desktop path for download and upload.
```

What We Have Learned

1. A Windows Phone provides a system of isolated storage by which each application on the phone is able to store its own data on the device.
2. Programs can create folders and files in their isolated storage area and use standard input/output constructions based on streams to communicate with them.
3. An alternative method to store data is to make use of a dictionary based mechanism to store name-value pairs such as settings values (Background colour = blue).
4. The Isolated Storage Explorer tool allows programmers to view the contents of the isolated storage for an application under development.

6 Using Databases on Windows Phone

6.1 An Overview of Database Storage

A database is a collection of, well, data, which is organised and managed by a computer program which is often, and rather confusingly, called a database.

Another computer program can ask a database questions and the database will come back with the results. You don't think of a database as part of your program as such, it is the component in your solution that deals with data storage.

We could discover how a database works by taking a quick look at a very simple example. Consider how we would create a database to hold information for an internet shop. We have a large number of customers that will place orders for particular products.

If we hire a database designer they will charge us a huge amount of money and then they will come up with some designs for database tables that will hold the information in our system. Each of the tables will have a set of columns that hold one particular property of an item, and a row across the table will describe a single item. This is the design for the customer table:

Customer ID	Name	Address	Bank Details
123456	Rob	18 Pussycat Mews	Nut East Bank
654322	Jim	10 Motor Drive	Big Fall Bank
111111	Ethel	4 Funny Address	Strange bank

This is a **Customers** table that our sales system will use. It contains the same information that was stored in the Customer Manager, with the addition of a string which gives bank details for that customer. Each row of the table will hold information about one customer. The table could be a lot larger than this; it might also hold the email address of the customer, their birthday and so on.

Product ID	Product Name	Supplier	Price
1001	Windows Phone 7	Microsoft	200
1002	Cheese grater	Cheese Industries	2
1003	Boat hook	John's Dockyard	20

This is the **Products** table. Each row of this table describes a single product that the store has in stock. This idea might be extended so that rather than the name of the supplier the system uses a Supplier ID which identifies a row in the Suppliers table.

Order ID	Customer ID	Product ID	Quantity	Order Date	Status
1	123456	1001	1	21/10/2010	Shipped
2	111111	1002	5	10/10/2010	Shipped
3	654322	1003	4	1/09/2010	On order

This is the **Orders** table. Each row of the table describes a particular order that has been placed. It gives the Product ID of the product that was bought, and the Customer ID of the customer that bought it.

By combining the tables you can work out that Rob has bought a Windows phone, Ethel has bought five Cheese graters and the four Boat hooks for Jim are on order.

This is a bit like detective work or puzzle solving. If you look in the Customer table you can find that the customer ID for Rob is 123456. You can then look through the order table and find that customer 123456 ordered something with an ID of 1001. You can then look through the product table and find that product 1001 is a Windows Phone (which is actually not that surprising).

Databases and Queries

You drive a database by asking it questions, or queries. We could build a query that would find all the orders that Rob has placed. The database system would search through the Orders table and return all the rows that have the Customer ID of 123456. We could then use this table to find out just what products had been purchased by searching through the Products table for each of the Product IDs in the orders that have been found. This would return another bunch of results that identify all the things I have bought. If I want to ask the database to give me all the orders from Rob I could create a query like this:

```
SELECT * FROM Orders WHERE CustomerID = "123456"
```

This would return a “mini-table” that just contained rows with the Customer ID of “123456”, i.e. all the orders placed by Rob. The commands I’m using above are in a language called SQL or Structured Query Language. This was specifically invented for asking questions of databases.

Companies like Amazon do this all the time. It is how they manage their enormous stocks and huge number of customers. It is also how they create their “Amazon has recommendations for you” part of the site, but we will let that slide for now.

Connecting to a Database

Unfortunately object oriented programs do not work in terms of tables, rows and queries. They work in terms of objects that contain data properties and methods. When a program connects to a database we need a way of managing this transition from tables to objects.

Databases and Classes

We have already seen that we can represent data in classes. The **Customer** class in the Customer Manager application we have been working on holds information about each customer:


```

public class Customer
{
    public int CustomerID { get; set; }
    public string Name { get; set; }
    public string Address { get; set; }
    public string BankDetails { get; set; }

    public Customer(int inID, string inName, string inBank,
                    string inAddress,)
    {
        Name = inName;
        Address = inAddress;
        BankDetails = inBank;
        ID = inID;
    }
}

```

To hold a large number of customers we created a class that holds a list of them.

```
List<Customer> Customers = new List<Customer>();
```

Whenever I add a new customer I add it to the `Customers` list and to find customers I use the `foreach` loop to work through the customers and locate the one I want. As an example, to find all the orders made by a customer I could do something like this:

```

public List<Order> FindCustomerOrders(int CustomerID)
{
    List<Order> result = new List<Order>();

    foreach ( Order order in Orders )
    {
        if (order.CustomerID == CustomerID)
        {
            result.Add(order);
        }
    }
    return result;
}

```

The method searches through all the orders and builds a new list containing only ones which have the required customer ID. It is a bit more work than the SQL query, but it has the same effect.

Using LINQ to connect Databases to Objects

People like using databases because it is easier to create a query than write the code to perform the search. However, they also like using objects because they are a great way to structure programs. To use databases and object oriented programs together the programmer would usually have to write lots of “glue” code that transferred data from SQL tables into objects and then back again when they are stored. This is a big problem in large developments where they may have many data tables and classes based on information in the tables.

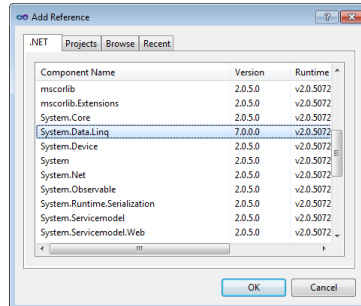
Language Integrated Query, or LINQ, was the result of an effort to remove the need for all this glue. It is called “Language Integrated” because they actually had to change the design of the C# language to be able to make the feature work.

Windows Phone supports the creation of databases for our programs to use and all the database interaction is performed using LINQ. So now we are going to see how we can take our class design and use it to create a database. We are going to use SQL to store the customers in our Customer Manager. Then we are going to add support for product orders that they place. This will let our user hold a large number of customers, orders and products in a database on the phone.

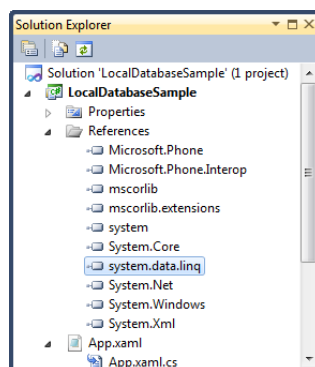
The next thing we are going to do is start to design the database tables, but before we can do that we need to make LINQ work with this project.

The LINQ Libraries

Before a program can use the LINQ code we have to add a reference to the `System.Data.Linq` libraries to the project. Because not all programs need database support, this reference is normally left out when a new project is created. We can open the Add Reference dialog by selecting Add Reference from the Project menu in Visual Studio. This gives us a list of libraries from which we can select the one we want.



When OK is clicked the reference is added to the list of resources that the program requires.



The next thing we should do is add some “using” directives to make it easier to access the classes in this library.

```
using System.Linq;
using System.Data.Linq;
using System.Data.Linq.Mapping;
using System.ComponentModel;
using System.Collections.ObjectModel;
```

We can now use all the classes from these libraries without having to use the fully qualified form of each name.

Creating a LINQ Table

Next we want to create some table designs that our database is going to manage. There will be three tables in our database.

- Customers
- Orders
- Products

We could start by considering how we are going to store information about a customer. From the original table we created a `Customer` class that describes the data a customer must contain. Now we want to use this class to create a design for a database table that will hold all our customers. We can do this by taking our original `Customer` class and modifying it so that LINQ can use the class design to create a table in the database from it. This is our `Customer` class design.

```
public class Customer
{
    public string Name {get; set;}
    public string Address { get; set; }
    string BankDetails { get; set; }
    public int CustomerID { get; set; }
}
```

This class holds three strings and an integer, which are all members of the class. They are implemented as properties, with `get` and `set` behaviours. The start of our `Customer` class for LINQ will look rather similar:

```
[Table]
public class Customer : INotifyPropertyChanged,
                       INotifyPropertyChanging
{
    // Customer table design goes here
}
```

The `[Table]` item is an *attribute*. Attributes are used to flag classes with information which can be picked up by programs that look at the *metadata* in the compiled code. Metadata, as I am sure you are all aware, is “data about data”. In this case the data being added is the attribute, and the data it is about is the `Customer` class.

When a C# class is compiled the compiler adds lots of metadata to the output that is produced, including the precise version of the code, details of all the methods it contains and so on. Other programs can read this metadata, along with the class information. This how `ildasm` produced the output we saw in chapter 3 when we inspected the contents of the files produced by Visual Studio.

The code above adds the `[Table]` attribute to the `Customer` class, which LINQ interprets as meaning “This class can be used as the basis of a data table”. There is actually nothing much inside the `[Table]` attribute itself, it just serves as a tag.

The code above also states that `Customer` class implements the `INotifyPropertyChanged` and `INotifyPropertyChanging` interfaces. A class implements an interface when it contains all the methods that are specified in the interface. The two interfaces contain methods that will be used to tell LINQ when the content of the data in the class are being changed.

We have seen this situation before, where data bound to a Silverlight display element needed to communicate changes so that the display updated automatically when the data was changed. This situation is very similar, except that a change to the data is going to trigger updates in the database.

Each interface contains a single event delegate. This is the delegate for `INotifyPropertyChanged`.

```
public event PropertyChangedEventHandler PropertyChanged;
```

The LINQ infrastructure will bind to the `PropertyChanged` event so that it can be told when a property value has changed. There is also a delegate to be used to tell LINQ when a value is changing. This event is fired before the change is executed:

```
public event PropertyChangingEventHandler PropertyChanging;
```

Our job is to make sure that the `Customer` class fires these events when data changes. In other words, if a programmer does something like this:

```
activeCustomer.Name = "Trevor";
```

- the code that updates the `Name` property should also tell LINQ that the data has changed. This means that the `Name` property could look like this:

```

private string nameValue;

public string Name
{
    get
    {
        return nameValue;
    }
    set
    {
        if (PropertyChanging != null)
        {
            PropertyChanging(this,
                new PropertyChangingEventArgs("Name"));
        }
        nameValue = value;
        if (PropertyChanged != null)
        {
            PropertyChanged(this,
                new PropertyChangedEventArgs("Name"));
        }
    }
}

```

The `Get` behaviour for the property just returns the value of the data member that holds the name.

The `Set` behaviour tests to see if anything is attached to the event handlers for the events and calls one before the property is changed, and the other after. The delegate calls are supplied with two parameters. The first is `this`, which is a reference to the currently active `Customer` instance (the one whose name is changing). The second is an event argument that contains the name of the property that is being changed. LINQ can use these to work out what has changed and perform an appropriate database update. As we saw when using Silverlight, if the name of the property is specified incorrectly (we say “name” or “nayme”) the update will not work correctly.

The change generates “before” and “after” events so that LINQ can manage the data changes more efficiently. If this seems like a lot of extra work, remember that the result of our efforts will be that when we update objects from our database we don’t have to worry about storing the effects of our changes. This will all happen automatically.

If our objects contain a lot of data fields it makes sense to simplify the changed event management by writing some methods that do the notification management for us:

```

private void NotifyPropertyChanged(string propertyName)
{
    if (PropertyChanged != null)
    {
        PropertyChanged(this,
            new PropertyChangedEventArgs(propertyName));
    }
}

private void NotifyPropertyChanging(string propertyName)
{
    if (PropertyChanging != null)
    {
        PropertyChanging(this,
            new PropertyChangingEventArgs(propertyName));
    }
}

```

We can then simplify the name property to this behaviour:

```

private string nameValue;

public string Name
{
    get
    {
        return nameValue;
    }
    set
    {
        NotifyPropertyChanging("Name");
        nameValue = value;
        NotifyPropertyChanged("Name");
    }
}

```

If this `NotifyPropertyChanged` stuff seems rather tiresome, remember why we are doing it. We are doing it so that we can write code such as:

```
activeCustomer.Name = "Fred Blogs";
```

- and have the database entry for the currently active customer automatically update. This is really useful. If we make use of `TwoWay` data binding the result will be that once we have connected our data objects to the display elements they will update themselves automatically, which is very useful indeed.

There is one final thing we need to add to the `Name` property to allow it to be used by LINQ. We have to mark the property with the `[Column]` attribute, so that LINQ knows that it is to be used in the database:

```

[Column]
public string Name
{
    // Name property behaviour goes here.
}

```

The address and bank details items can be added to the `Customer` class in exactly the same way. Each of them will map onto another column in the database.

Creating a Primary Key

The `CustomerID` item is slightly different however. In our original `Customer Manager` it was an integer value that identified each customer. If the customer gets married and changes their name the customer ID is there to make sure we can still find them in the application. In database terms we are using the `Customer ID` as the *primary key* for this table. The primary key is used to uniquely identify particular customers in this table. We might have two customers with the name "John Smith", but we will **never** have two customers with the same ID. The ID works in exactly the same way as your Social Security number or bank account number.

When we create the ID column we have to tell LINQ that this is the primary key for the table. We can also ask LINQ to make sure that all values in this column are unique, and we can even ask that the values for this element be auto-generated. We do this by adding information to the `Column` attribute for the `CustomerID` property:

```
[Column(IsPrimaryKey = true, IsDbGenerated = true)]
```

```
public int CustomerID { get; set; }
```

When we add a new customer entry to the database we don't have to worry about the ID, as it will be generated for us. We can of course view the ID value, so that we can tell the next "John Smith" what their unique id is.

Note that we are **never** allowed to change the ID for a customer. This is a unique key which will be automatically created by the database for each customer in turn. In database terms this property is called the *primary key* for the customer record. The

primary key of a customer will be used in our database to implement a relationship between the Customers table and the Orders table. Each row of the Orders database will contain a customer ID value that identifies the customer who placed that order.

When LINQ reads the metadata for this property it will use these settings to decide what kind of database column to create.

Creating a LINQ Data Context

Now that we have the design for a row in the table we can create the table itself. In our previous customer manager application we used a List to hold the customers. In LINQ we create a `DataContext` which will manage the connection to the database we are about to create. This will contain all the tables in the application, which at the moment just means the customers:

```
public class SalesDB : DataContext
{
    public Table<Customer> Customers;

    public Customers(string connection) : base(connection)
    {
    }
}
```

The `SalesDB` class extends the `DataContext` class, which is provided by LINQ as the basis of a database design. The constructor for this class simply calls the constructor for the parent class. The connection string gives location of the database we are going to use. One of the great things about database connection strings is that they can be used to connect to databases over a network, so that a program can work with a database on a distant server. We are not going to do this; we are going to connect the database to a file held in isolated storage on the phone.

The `DataContext` serves as a connection to a database. You can think of it as a bit like a stream that connects a program to a file. It exposes methods that we can use to manage the database contents.

The connection string describes how the program will connect to the database. In some applications this connection can be to a database server on a distant machine. In this case the database requests are sent to the remote system and the program works with the responses. In the case of the database on a Windows Phone this is a path to a file held in Isolated Storage.

Creating Sample Data

When we created our original Customer Manager program we created a method to create sample data. We are going to do exactly the same thing with our database version. This will populate the database with known data which we can then work with. Note that the method is now called `MakeTestDB`, because we will add further test setup for the other tables when we add them later. We can start with exactly the same sample information:

```

public static void MakeTestDB(string connection)
{
    string[] firstNames = new string[] { "Rob", "Jim", "Joe",
                                         "Nigel", "Sally", "Tim" };
    string[] lastsNames = new string[] { "Smith", "Jones",
                                         "Bloggs", "Miles", "Wilkinson", "Brown" };

    SalesDB newDB = new SalesDB(connection);

    if (newDB.DatabaseExists())
    {
        newDB.DeleteDatabase();
    }

    newDB.CreateDatabase();

    foreach (string lastName in lastsNames)
    {
        foreach (string firstname in firstNames)
        {
            //Construct some customer details
            string name = firstname + " " + lastName;
            string address = name + "'s address";
            string bank = name + "'s bank";
            Customer newCustomer = new Customer();
            newCustomer.Name = name;
            newCustomer.Address = address;
            newCustomer.BankDetails = bank;
            newDB.CustomerTable.InsertOnSubmit(newCustomer);
        }
    }

    newDB.SubmitChanges();
}

```

The first thing the method does is make a connection to the database:

```
Customers newDB = new Customers(connection);
```

The variable `newDB` now represents a connection to the database with the given connection string. The method can use this connection to issue commands to manage the database.

The next thing that the method does is check to see if the database already exists. If it does, the database is deleted.

```

if (newDB.DatabaseExists())
{
    newDB.DeleteDatabase();
}

```

An important principle of testing is that a test should always start from exactly the same place each time. If there is already a database in a file in isolated storage it will need to be cleared before the test data is added.

```
newDB.CreateDatabase();
```

The next thing the method does is create a new database which contains a set of empty tables.

Once we have our new database we can create a set of test customer values and add each one to the customers table.

```

foreach (string lastName in lastsNames)
{
    foreach (string firstname in firstNames)
    {

```

```

        //Construct some customer details
        string name = firstname + " " + lastName;
        string address = name + "'s address";
        string bank = name + "'s bank";
        Customer newCustomer = new Customer();
        newCustomer.Name = name;
        newCustomer.Address = address;
        newCustomer.BankDetails = bank;
        newDB.CustomerTable.InsertOnSubmit(newCustomer);
    }
}

```

These two nested loops work through the same set of first and last names we used last time. The major difference from the previous test data generator is that this time a completed customer is added to the database:

```
newDB.CustomerTable.InsertOnSubmit(newCustomer);
```

The `InsertOnSubmit` method is how we add new entries into a table. We are using the `CustomerTable` property of the new database. If the database contained other tables we could add items to them in the same way. Note that this is typesafe; in other words if we tried to add a different type to the `CustomerTable` the program would not compile.

The final, and perhaps the most important, part of the method is the call that actually commits the changes to the database:

```
newDB.SubmitChanges();
```

This is directly analogous to the `Close` method when using file streams. This is the point at which the changes that we have asked for will be committed to the database. Up until the `SubmitChanges` method is called the system may keep some of the changes in memory to speed things up. This makes good sense if the same data item is being updated repeatedly. Rather than changing the database file each time, the system will use copy in memory and work with that. Only when this method is called will the memory copies be written back to the database file.

If you don't submit your changes there is a good chance that this will leave the database in a mess. Which would be bad.

We now have a database which we can use in our application. We can also create a test database with values we can look at. The database context that we are going to use will live in the `App.xaml.cs` program file along with a reference to the currently active customer.

```
public SalesDB ActiveDB;
public Customer ActiveCustomer;
```

When the program starts running the variable `ActiveDB` is set to the database we are using.

```
ActiveDB = new SalesDB("Data Source=isostore:/Sample.sdf");
```

Note that the file path has a particular format and identifies a file in the isolated storage. We can make use of multiple databases in our program if we wish, each will be held in a different file. This database file format is that of a standard SQL database. A PC based database program could read the tables out of this file and it could be managed by an SQL database editor. The reverse is also true, a Windows Phone application could make use of a database that was prepared on a different computer and loaded into the application.

Binding a ListBox to the result of a LINQ Query

Next we want to see how we can use this database in our application. In the simple Customer Manager program we used a `List` to hold all the customers. We found that the `ListBox` display element can take a list and display it. Now we want to take data out of the database and display this.

We get information out of the database by issuing queries. Previously we saw how the SELECT command could be used to fetch information. What we are now going to do is build a LINQ query to get all the customers out of the database.

```
var customers = from Customer customer
                in thisApp.ActiveDB.CustomerTable
                select customer;
```

This line of C# creates a variable called `customers`. The `customer` variable is of type `var`. If you've not seen this type before it can look a bit confusing. What `var` means in this context is "Get the type of this variable from the expression that is being put into it". This doesn't mean that the C# is giving up on strong typing, if we ever try to use the variable `customers` in an invalid way our program will still fail to compile.

The rest of the statement tells LINQ to get all the customer items from the `CustomerTable` member of the `ActiveDB` database context. This is returned as a list of customers that can behave as an `ObservableCollection`. Which is just what we need to give the `ListBox` to display the customers:

```
customerList.ItemsSource = customers;
```

At this point we have a fully working customer database. The only thing that we have to add to the program is code to submit the changes to the database when the user leaves the program. The best way to arrange this is to put the call into the `OnNavigatedFrom` method in `MainPage.xaml.cs`:

```
protected override void OnNavigatedFrom(
    System.Windows.Navigation.NavigationEventArgs e)
{
    App thisApp = Application.Current as App;

    thisApp.ActiveDB.SubmitChanges();
}
```

When the user moves away from this page the method will find the currently active database and submit any outstanding changes to it.

The solution in *Demo 01 SalesManagement* contains a Windows Phone Silverlight application that implements a fully working customer manager. It does persist the data into isolated storage, but because the data storage is created fresh each time you will not see this. Once you have run the program once to create you can comment out the statement in `App.xaml.cs` that creates a new database and see this working for yourself.

Database and Data Binding Magic

I think this is actually quite magical. We have only added a small amount of code to our application and how we have the data elements stored and retrieved for us automatically. Changes that we make in the user interface are automatically being persisted in the database without any effort from us. It is worth just going through exactly what is happening here so that we can completely understand how it works. Consider the following sequence:

1. Users changes "Rob Miles" to "Rob Miles the Wonderful" in the `TextBox` holding the name field on the `CustomerDetailsPage`.
2. Because this `TextBox` is bound to the `Name` property in the `CustomerView` bound to this page it now updates that property in the `CustomerView` to the new name text.
3. User presses the `Save` button on the `CustomerView` page.
4. The `saveButton_Click` method in the `CustomerDetailsPage` runs and calls the `Save` method in the `CustomerView` instance behind this page.
5. The `Save` method copies the page properties from the view class into the active customer record.
6. Changes to the `Name` property in the active customer fires the property changed events in LINQ informing the database that the values are being changed.

7. LINQ flags the customer information as having changed and also raises a change event in the `ObservableCollection` which holds our customer list and is connected to the `ListBox` on `MainPage`.
8. The program returns to the `ListBox` view on `MainPage` and the display of the name “Rob Miles” is updated to “Rob Miles the Wonderful”.
9. When the user exits the application the `OnNavigatedFrom` event fires and calls `SubmitChanges` on the database, storing the changes back into isolated storage.

If you don't believe this works, then try it and find out. You can use the above pattern anywhere you want to automatically save and load data with very little effort.

Adding Filters

By slightly changing the format of the LINQ query to fetch the data from the database we can select records that match a particular criteria:

```
var customers = from Customer customer
                in thisApp.ActiveDB.CustomerTable
                where customer.Name.StartsWith("S")
                select customer;
```

This query selects only the customers whose name starts with the letter S.



This would be the result of such a query. This makes it very easy to extract particular elements from the database. We could easily add a search box to our application if we wanted to search for customers with a particular name.

6.2 Creating Data Relationships with LINQ

We are now able to store a large number of customer items and change and update them by using data binding.

However, to make our sales management system work we will need to create other tables and also link them together. To refresh our memories this is the design for the customer table:

Customer ID	Name	Address	Bank Details
123456	Rob	18 Pussycat Mews	Nut East Bank
654322	Jim	10 Motor Drive	Big Fall Bank
111111	Ethel	4 Funny Address	Strange bank

We have created a class which contains the required rows and can use LINQ to persist this data in a database file in isolated storage. The **Products** table is very similar to the

Customers table, in that it just holds a list of items with certain properties.

Product ID	Product Name	Supplier	Price
1001	Windows Phone 7	Microsoft	200
1002	Cheese grater	Cheese Industries	2
1003	Boat hook	John's Dockyard	20

The Products class is actually very similar to the Customer one; it just contains a set of data items. The **Order** table is the more complex one. It actually contains references to rows in the other two classes.

Order ID	Customer ID	Product ID	Quantity	Order Date	Status
1	123456	1001	1	21/10/2010	Shipped
2	111111	1002	5	10/10/2010	Shipped
3	654322	1003	4	1/09/2010	On order

Each row of the Order table describes a particular order that has been placed. It gives the Product ID of the product that was bought, and the Customer ID of the customer that bought it. In C# this would be easy to create:

```
public class Order
{
    public DateTime OrderDate;
    public int Quantity;
    public Customer OrderCustomer;
    public Product OrderProduct;
}
```

The **Customer** and **Order** references would connect an order to the customer that made the order, and the product that was ordered.

However, when we are using a database we don't have any references as such. Rather than having a reference to a customer and product, instead the Order table holds the ID of the items that are related to each order. We then use the value in the Customer ID column to find the customer who has placed the order. For example, we know that order 1 was placed by Rob Miles because it contains a Customer ID of 123456, and so on.

In the good old days before computers this was how things worked. A company would have a filing cabinet full of customer data, another full of product data and a third full of orders. To find out the details of a particular order a clerk would have to look up the customer and product information from information written on the order details. The relational database provided a way for computers to manage information that contains relationships like these, what we want to do now is take the database relationship information and use it to create object references that make it easy to manipulate in an object oriented program.

LINQ Associations

In LINQ a relationship between one table and another is called an Association. This is implemented in the database classes by an **EntityRef** value.

Linking an Order to the Customer that placed it

From our system requirements the **Order** class needs to hold a reference to the **Customer** who placed the order. This is implemented as a relationship between two tables. The **EntityRef** is a special LINQ class we can use to connect two tables together, using a database to provide the underlying storage.

```

[Table]
public class Order : INotifyPropertyChanged,
                    INotifyPropertyChanging
{
    ...

    private EntityRef<Customer> orderCustomer;

    [Association(IsForeignKey = true, Storage = "orderCustomer")]
    public Customer OrderCustomer
    {
        get
        {
            return orderCustomer.Entity;
        }
        set
        {
            NotifyPropertyChanging("OrderCustomer");
            orderCustomer.Entity = value;
            NotifyPropertyChanged("OrderCustomer");
        }
    }
}

```

The `EntityRef` acts as a kind of glue between a simple reference (which we would like to use) and a lookup in a table (which is what LINQ will have to do when we use this reference). The good news is that as programmers we can write:

```

Customer newCustomer = new Customer();
Order newOrder = new Order();
newOrder.OrderCustomer = newCustomer;

```

This simple assignment will have the effect of connecting the order to the customer, so that the database record for that order holds the id of the customer the order is for. We can use this technique every time we want to implement a relationship between a row in one table and a row in another.

When we define the association property we also tell LINQ two other things. We tell it that the association is a “foreign key” and we also identify the property of the class that is going to store the value. If you have used databases before you will be familiar with the idea of a foreign key. This is a primary key from another database. In this case it is the primary key from the Customer database that can be used to identify the particular customer who placed the order.

The storage item in the association tells LINQ which private property of our database will actually hold the property information.

Databases and Navigability

Navigability in database systems is important. It allows a system to find one piece of information given another. At the moment we can start from an order and directly find the customer for that order. This is because a row in the order table contains the `CustomerID` of the customer for that order. However, we can’t find our way back. Starting with an order we have no way of finding out the customer who placed it.

To allow proper navigability we have to put something in the `Customer` class that will allow us to find the orders that the customer has placed. However, things now get even more complicated because a customer could create lots of orders. This means that the nature of the relationships changes depending on which direction you are travelling:

- A particular order will only ever be associated with a single customer, the customer who created the order. This means that the relationship travelling from order to customer is one to one (one order to one customer).

- A particular customer will be associated with many orders. This means that the relationship travelling from customer to order is one to many (one customer to many orders)

When we design our databases we need to consider the navigability, in particular we need to think about how we are going to navigate using our relationships, and whether they are one to one or one to many.

Linking a Customer to all their orders

A customer may place many orders and so the relationship that we need to implement is a “one to many” relationship. If we were implementing this using C# classes we would use a collection of some kind, perhaps an array or a List. In LINQ such a relationship is provided by the `EntitySet` class. This is a bit like the `EntityRef`, except that it can manage a set of items, rather than just one. We can link a customer to their orders by adding an `EntitySet` to the `Customer` class.

```
[Table]
public class Customer : INotifyPropertyChanged,
                      INotifyPropertyChanging
{
    [Column(IsPrimaryKey = true, IsDbGenerated = true,
           AutoSync = AutoSync.OnInsert)]
    public int CustomerID { get; set; }

    private EntitySet<Order> orders = new EntitySet<Order>();

    [Association(Storage="orders", ThisKey="CustomerID",
                OtherKey="OrderCustomerID")]
    public EntitySet<Order> Orders
    {
        get
        {
            return orders;
        }
        set
        {
            orders = value;
        }
    }
}
```

This looks very like the `EntityRef` that we created earlier, except that it has extra information in the Association to describe the relationship.

The item `ThisKey` gives the name of the property that will be used by the `Order` to find the customer it is associated with. We can use `CustomerID`, the primary key for the `Customer` to do this.

The item `OtherKey` gives the name of the element in the `Order` class that will implement the association in the other direction, i.e. allows an order to find the customer it is related to.

You may be wondering why we have not fired any `NotifyPropertyChanged` events when `Orders` is changed. This is because changing the value of the `EntitySet` is something that we would hardly ever do. Remember that this is the container for the orders, not the orders themselves. We will add orders to the customer by putting things into the `EntitySet`, not by replacing the `EntitySet` itself.

To make the association work properly we have to make some changes to the reference in the `Order` class. These will bind both ends of the association together.

```

[Table]
public class Order : INotifyPropertyChanged, INotifyPropertyChanging
{
    ...

    [Column]
    public int OrderCustomerID;

    private EntityRef<Customer> orderCustomer =
        new EntityRef<Customer>();

    [Association(IsForeignKey = true, Storage = "orderCustomer",
                ThisKey = "OrderCustomerID")]
    public Customer OrderCustomer
    {
        get
        {
            return orderCustomer.Entity;
        }
        set
        {
            NotifyPropertyChanging("OrderCustomer");
            orderCustomer.Entity = value;
            NotifyPropertyChanged("OrderCustomer");
            if (value != null)
                OrderCustomerID = value.CustomerID;
        }
    }
}

```

The Association is defined as holding a foreign key (the primary key of the customer database) and using the value of OrderCustomerID to hold this value.

This version of the reference makes a copy of the CustomerID value into the OrderCustomerID column in the table when a customer is assigned to the order. In other words, if I write something like this:

```

Customer c = new Customer();
Order o = new Order();
o.OrderCustomer = c;

```

When the OrderCustomer property is assigned the set method above will take the CustomerID value and copy it into the OrderCustomerID value to record the customer for this order.

If all this is hurting your head a bit, and I must admit it hurt mine first time round, remember the problem that we are trying to solve. The ThisKey and the OtherKey descriptions in the association tell each side of the relationship how to find each other.

The good news is that once we have implemented this pattern we can easily add new orders to a customer:

```

Customer c = new Customer();
Order o = new Order();
o.OrderCustomer = c;
c.Orders.Add(o);

```

The EntitySet class provides an Add method that lets us add orders to the database. To work through the contents of the entity set we can use any of the constructions we would normally use to process collections:

```

foreach (Order o in c.Orders)
{
    // do something with each order
}

```

This doesn't look that special, but the special thing is that we are writing C# to work with the database.

Orders, Products and Database Design

We now have a link between customers and orders which works properly. A customer can have a very large number of orders, and an order is always connected to the customer who placed it. The next thing we need to do is fill in the order with the products that the customer might have bought. This is the point at which we hit a snag with our original design of our `Order` table:

Order ID	Customer ID	Product ID	Quantity	Order Date	Status
1	123456	1001	1	21/10/2010	Shipped
2	111111	1002	5	10/10/2010	Shipped
3	654322	1003	4	1/09/2010	On order

If you look carefully at the design it turns out that there is no way we can have an order that holds more than one product. A particular order can be for multiple items, but they must all be of the same item.

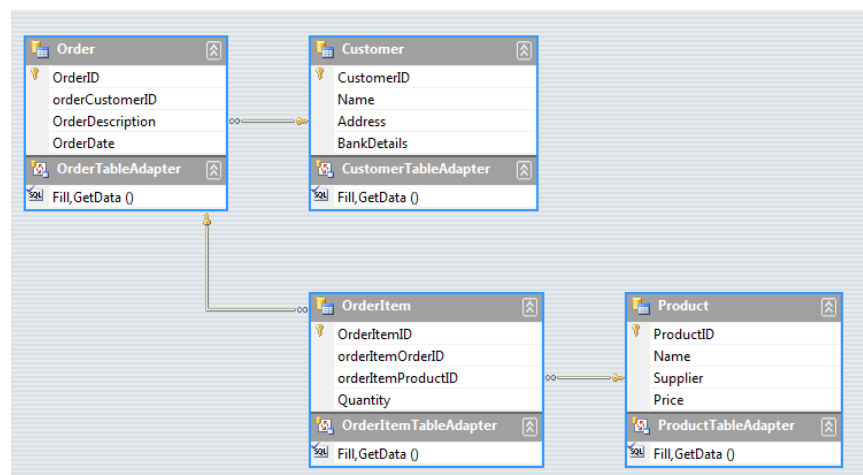
It might be that this is perfectly OK. The person who wants to use your system might be selling items that are very expensive and are only ever ordered one at a time. Alternatively this might be disastrous; the system may be used by a supermarket that wishes to assemble orders which contain a huge number of different products.

To solve the problem we need to add another table, which we will hold all the items in each order. Each `OrderItem` will be linked to a particular order. Each `Order` will then hold a collection of these items.

OrderItem ID	Order ID	Product ID	Quantity
1	56	1001	1
2	56	1002	5
3	12343	1003	4

The table above shows that the order 65 was made up of two products, a Windows Phone and 5 Cheese Graters. (If you are wondering where the products come from, take a look at the Products table way back at the start of this chapter).

Database designers like to use diagrams to show how the various elements fit together. The diagram for the database that we have built looks like this:



The little keys show where a foreign key is being used at one end of the association. The little infinities (∞) show that an association can connect to any number of that item.

If you do a lot of database design you will be used to drawing diagrams like the one above. It was actually created from the database that was formed by the classes that we created.



The solution in *Demo 02 SalesManagement* contains a Windows Phone Silverlight application that adds Order, OrderItem and Product tables to the customer database. You can navigate the lists of customers and products and also find the orders that a customer has created and their contents.

LINQ Queries and Joining

Now that we have our data in a database we can use the very powerful query mechanisms to get information from it. For example, we might want to find all the orders that were placed on a particular date.

```
DateTime searchDate = new DateTime(2011, 8, 19);

var orders = from Order order
              in activeDB.OrderTable
              where order.OrderDate == searchDate
              select order;
```

This query searches the order table of the currently active DB and finds all the orders created on the 19th of August 2011. We don't have to write any of the searching; the result is just returned in the variable called `orders`. You might be wondering what the actual type of `orders` is. Actually it is of type "LINQ query". The statement above doesn't do anything until you try to consume the data that it holds. For example, if we tried to work out the total value of the sales for that day we could write something like this:

```
int totalSales = 0;

foreach (Order order in orders)
{
    foreach (OrderItem item in order.OrderItems)
    {
        totalSales += item.OrderItemProduct.Price;
    }
}
```

The first `foreach` works through `orders` getting each order in turn. It is only when the program starts to actually consume the data that LINQ will leap into action and start producing results.

Note that this has implications if you want to use the result of a query more than once. If you use the same query twice you will find that your program will run a bit more slowly, as LINQ will be fetching things from the database twice. If you want to use a result more than once it is best to use the query to product a list of items that you can then work through as many times as you like:

```
List<Order> DateOrders = orders.ToList<Order>();
```

This creates a list called `DateOrders` which contains the result of the query. We can use the list in our program without generating any further transactions on the database. The `ToList` method applies the query and generates a list of results.

Joining two tables with a query

LINQ has another trick up its sleeve. It can create queries that go across more than one table, a process called “joining”. Perhaps we might want to create a complete list of all the orders made by all customers, listed in customer order. We want to find `Order` and `Customer` pairs that match up and list them. Using the `join` keyword LINQ can do this in one action:

```
var allOrders = from Customer c in newDB.CustomerTable
                join Order o in newDB.OrderTable on
                    c.CustomerID equals o.OrderCustomerID
                select new { c.Name, o.OrderDescription };
```

If you haven’t seen one of these you would be forgiven wondering how on earth it compiles. The key is in the name LINQ, in that the `from`, `join` and `select` keywords are all integrated into the C# language to make these kinds of queries work. The first part of query finds all the customers in the customer table. This query is then joined onto a second, which works through all the orders doing a join on the two ID values stored in customer and order.

Essentially what it is saying is “Find me pairs of orders and customers where the orderID in the customer matches the customer id in the order”.

Once it finds a match it then creates a new type that just contains the results that we want, in this case the `Name` from the customer and the `OrderDescription` from the order. Note the magic `var` is here again. In this case we get an SQL query that generates a collection of type containing two strings of the appropriate names.

We could use this query to produce a list of order descriptions by working through it

```
List<String> OrderDescriptions = new List<String>();

foreach (var orderDesc in allOrders)
{
    OrderDescriptions.Add(orderDesc.Name + " order: " +
        orderDesc.OrderDescription);
}
```

As with the previous query, LINQ doesn’t actually do anything until the program starts to work through the results. Note that we need to use a `var` type for the `orderDesc` value that we are fetching from the query. This contains the two items that we added when we created the type in the query.

The result of his query could be bound to a display element to show the list:

```
Rob Miles order: Camera Order
Rob Miles order: Cheese Order
Rob Miles order: Barge Pole Order
```

The `var` keyword

The `var` keyword can seem a little scary if, like me, you are used to declaring variables of a particular type and then just using them. I like the way that the C# compiler fusses about my programs and stops me from doing stupid things like:

```
int i = "hello";
```

An attempt to put a string into an integer is a bad thing to do, and will not compile. However, at first glance the `var` keyword looks like a way to break this lovely safety net. We can create things that don't seem to have a type associated with them, which must be bad. However, you can rest assured that although we don't give a name for a type like this, the C# compiler knows exactly what a `var` type is made from, and will refuse to compile a program that tries to use such a type in an incorrect way.

Deleting Items from a Database

We have seen how to add items to a database. It is also possible to delete items as well. For example, if the customer decided that they didn't want a particular order item we can delete it from the database:

```
ActiveDB.OrderItemTable.DeleteOnSubmit(item);
```

We can also delete collections of items using `DeleteAllOnSubmit`, including a collection identified as a result of a query.

Foreign Key Constraints and Delete

Some of our data items have relationships with others, for example a Customer will contain a number of Orders, and an Order will contain a number of OrderItem values. If the parent object is deleted before the children (i.e. if we try to delete a customer which contains some orders) the database will refuse to perform this, and will throw an exception as a result. If you want to delete a customer you will first have to delete all the order items in each order, and then delete the orders, and then delete the customer.

What We Have Learned

1. Windows Phone applications can use Language Integrated Query (LINQ) to interact with databases can be stored in files in isolated storage in the device. The databases themselves are managed by an SQL database server, but it is not possible to use SQL commands to interact with a database, instead queries are expressed using LINQ.
2. A database is made up of a number of tables. A table is made up of columns (different data items such as name, address, and bank details) and rows (all the information about a particular item – the name, address and bank details of an individual customer).
3. One column in a table acts as the “primary key”, holding a value which can uniquely identify that row (the unique ID of that particular customer). The database can be asked to automatically create primary keys that are unique for each row in a table.
4. Windows Phone applications can create classes which are stored by LINQ as database tables. The attributes `[Table]` and `[Column]` are used within the class definition to specify the role of the components. Additional information can be added to the attributes to identify primary keys.
5. Database classes managed by LINQ can add notification behaviours to the data properties so that LINQ will automatically update the database entries when the properties are changed. If these are used with changed events in Silverlight display elements this can allow automatic database update when items are changed by the user.
6. A database can contain relationships, where a column in one table contains primary key values which identify rows in another. These are called “foreign keys” as they are primary keys, but not for the table which contains them.
7. LINQ implements relationships by using the `EntityRef` class, which contains a reference to a foreign key in another database. This is used for “one to one”

relationships. The EntitySet class is used to allow a member of a table to refer to a large number of rows in another table. This is used for “one to many” relationships. These data elements are modified with the Association attribute that describes how the table relationship is implemented.

8. A database will frequently have a one to one going in one direction (an order is attached to one customer) and a one to many going in the other (a customer can place many orders). For such associations to work correctly is important that the associations are correctly configured and that foreign keys are stored correctly.
9. A LINQ query can be used to fetch structured data from the database. The query is only actually executed when the results of the query are being consumed by the program. LINQ queries can use the join keyword to combine the results of multiple queries.
10. When deleting items from a database all child items must be deleted before the parent object is removed, otherwise there would be elements in the database containing primary key values for items that don't exist.

7 Networking with Windows Phone

A networked device is much more useful than one that is just free standing. Once a system is connected to the internet all kinds of things become possible. Mobile phones are around the most connected devices there are, in that they have a whole variety of connection possibilities. However, they also bring challenges to the application developer in that any, or perhaps all, of these network connections could be removed at any instant, and the user of our application will expect it to handle this automatically.

In this section we will look at the network connection options available to programs that we write. We will also give some consideration to how we can make sure that our applications can respond gracefully to network problems.

All Windows Phones support network connectivity via the mobile telephone network and also by means of their built in WIFI. As far as our programs are concerned the two network technologies are interchangeable. The programs we are going to write will initiate connections to a data service and the underlying phone systems will make the connection the best way that they can.

Of course, if there is no mobile signal or WIFI available these connections will fail. Applications that we write must be able to deal with this situation gracefully. At the very least this means displaying a “Sorry we can’t connect you just now, please try later” message. A more advanced program may store data locally and then send it later. It is up to you how your program will work in this respect.

The Windows phone emulator has the same networking abilities as the real device. It uses the network connectivity of the host PC it is running on. This means that we can use a desktop PC to test any Windows Phone applications we are writing.

However it is also important to test programs on a real device. On a real Windows Phone we can disable network services and test our applications to ensure that they perform properly over lower speed connections and when the network connection is suddenly removed.


7.1 Networking Overview

Before we look at how Windows Phone uses network connections, we need to learn a little bit about networks. This is not a detailed examination of the field, but it should give you enough background to understand how our programs are going to work.

Starting with the Signal

The first computer networks used wires to send their data signals, although more modern networks can use radio or fibre optic cables. Whatever the medium is, the fundamental principle is that you have some hardware that can put data onto the medium in the form of bits and get it off again. A bit is either 0 or 1 (or you can think of a bit as either true or false) and can be signalled by the presence or absence of a voltage, a light from a light-emitting diode (LED), or a radio signal.

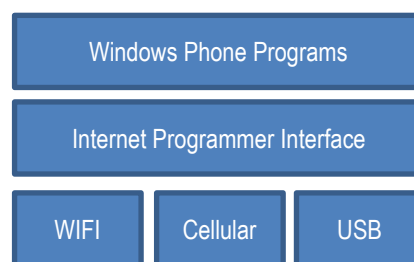
If you imagine signalling your friend in the house across the road by flashing your bedroom light on and off, you have an idea of the starting point of network communications.



Once we have this raw ability to send a signal from one place to another, we can start to transfer useful data.

Signals in Windows Phone

In the case of Windows Phone there are a whole range of signal types that it can use to connect with a network. There is the WIFI connection, the mobile phone cellular connection and even the USB connection to the Zune software on the PC. All of these technologies can be used to pass data between programs on the phone and the Internet.



This is actually made to work as you see above. The Windows Phone programs talk to the Internet Programmer Interface (the thing we are going to learn how to use) and the Internet Programmer Interface talks to the three connection technologies. The good news is that the programs we write don't need to be changed to work over the different types of connection. However, there are things that we need to remember, in that if our phone user goes abroad they may not allow data roaming, so the cellular data connection will be disabled. Even worse for our programs, a user may put their phone into "aircraft" mode and disable all network connections for a while.

A Windows Phone program can determine which connections are available by using methods in the `DeviceNetworkInformation` class.

```
using Microsoft.Phone.Net.NetworkInformation;

...

System.Text.StringBuilder sb = new System.Text.StringBuilder();

sb.Append("Operator: ");
sb.AppendLine(DeviceNetworkInformation.CellularMobileOperator);

sb.Append("Network available: ");
sb.AppendLine(DeviceNetworkInformation.IsNetworkAvailable.ToString());

sb.Append("Cellular enabled: ");
sb.AppendLine(DeviceNetworkInformation.IsCellularDataEnabled.ToString());

sb.Append("Roaming enabled: ");
sb.AppendLine(DeviceNetworkInformation.IsCellularDataRoamingEnabled.ToString());
;

sb.Append("Wi-Fi enabled: ");
sb.AppendLine(DeviceNetworkInformation.IsWiFiEnabled.ToString());
```

The above code assembles a message that describes the status of the phone which can then be displayed. We can also use these values to control the way that our application behaves.



The solution in *Demo 01 NetworkDiagnostics* contains a program that displays the network status for a phone or emulator.

Building Up to Packets

Just flashing your light to your friend willy-nilly does not allow you to send much information. To communicate useful signals you have to agree on a message format. You could say, “If my light is off and I flash it twice, it means it is safe to come round because my sister is out. If I flash it once, it means don’t come, and if I flash it three times, it means come and bring pizza with you.” This is the basis of a thing called a protocol, which is an arrangement by parties on the construction and meaning of messages.

A modern network sends data in 8 bit chunks that are called *octets* by the people who design them. In C# we can hold the value of a single octet in a variable of type `byte`. If we have 8 data bits we can arrange them in 256 different patterns, from 00000000 (all the bits clear) to 11111111 (all the bits set). An octet is the smallest unit of data that is transferred by a network. A give octet might be a number (in the range 0 to 255) a single character (a letter digit or punctuation) or part of a larger item, for example a floating point number or a picture file. On the internet a data packet can be several thousand octets in size.

Addressing Packets

Your bedroom light communication system would be more complicated if you had two friends on your street who want to use your bedroom light network. You would have to agree with them that you would send two sets of flashes for each message. The first one would indicate who the message was for, and the second would be the message itself. Computer networks function in exactly the same way. Every station on a physical network must have a unique address. Packets sent to that address are picked up by the network hardware in that station.

Networks also have what is called a broadcast address. This allows a system to send a message which will be picked up by every system. This is the network equivalent of “Calling all cars . . .” In our communication network, this could be used to warn everyone that your sister has come home and brought her boyfriend, so your house is to be avoided at all costs. In computer networks a broadcast is how a new computer can find out the addresses of important systems on the network. It can send out a “Hi. I’m

new around here!” message and get a response from a system that can then send it configuration information.

Everyone on a particular network can receive and act on a broadcast sent around it. In fact, if it wanted to, a station could listen to all the messages traveling down its part of the wire or WiFi channel. This illustrates a problem with networks. Just as both of your friends can see all the messages from your bedroom light, including ones not meant for them, there is nothing to stop someone from eavesdropping on your network traffic. When you connect to a secure Web site, your computer is encoding all the messages that it sends out so that someone listening other than the intended recipient would not be able to learn anything.


Internet Addressing

We have already seen the networks we are using can send 8 bit sized items. This would mean that if we used a single octet to address stations on the network we could not connect more than 256 systems. To get around this the present version of the Internet (IP V4) uses four 8 bit values which are strung together to allow 32 bits to be used as an address. In theory this allows for over four thousand million different systems on the internet. However, this large number is hard to achieve in practice and so a new version of the Internet (IP V6) is being rolled out which has 128 bit addresses (16 octets in each address). This would allow for a huge increase in the number of possible stations, but since it is a fundamental change to way the network functions it is taking a while to put into practice. At the moment the networking software in the Windows Phone does not support IPV6 addressing.

Routing


Not everything on the internet is connected to the same physical network. Signals sent around the wires in my house do not make as far as the network operated by my next door neighbour. Instead we have to view the Internet as a large number of separate networks which are connected together. This means that to get messages from one physical network to another we have to introduce the idea of routing.

If you had a friend on the next block, she might not be able to see your bedroom light. But she might be able to see the light of your friend across the road. This means that you could ask your friend across the road to receive messages and then send them on for you. Your friend across the road would read the address of the message coming in, and if it was for your friend on the next block, she would transmit it again. Figure 16-2 shows how this works. Your friend uses the window on the left to talk to you and the window on the right to relay messages to your more distant friend.



You can think of your friend in the middle as performing a routing role. She has a connection to both “networks”, the people you can see and also the people that your distant friend can see. She is therefore in a position to take messages from one network and resend them on the other one. Note that this is not like passing a letter from one person to another. Instead your friend is receiving your message and then re-transmitting it to your distant friend.

An address of a given system on the Internet is made up of two elements, the address of the local network containing the system and the address of that system on the network. The important thing to remember about a local network is that the machines in it are actually connected to the same physical media.



The diagram shows how this works. The machines on your home network are physically connected and can talk directly to each other. The Desktop PC can fetch files directly from the Home Server. However, if the Desktop PC needs to contact Microsoft the messages must leave the local network and travel via the Internet. Just before the Desktop PC sends any message out it looks at the network address of the system it is talking to. If the remote system is on the same network it will connect directly the system. However if, like Microsoft, the system is on a different network the Desktop PC will send the message to the router which sends such messages into the internet, which delivers them to the router at the Microsoft.

Packets that you send from your home PC to distant machines are sent to the network at your Internet service provider (ISP) which then re-transmits them to the next system in the path to the destination. Packets might have to be sent over several machines to reach their destination. The Internet constantly changes the routes that packets take. This makes it very reliable and able to manage sudden surges in traffic and failures of parts of the network but it can lead to situations where a packet arrives before another which was sent first. Sometimes packets can get lost (although this is fairly rare), so you can't be sure that one has arrived until you receive an acknowledgement. One thing you should remember is that you do not really "connect" your system to the Internet. Whenever your system is connected, it actually becomes part of the Internet.

You can regard a local-level network as the same as the internal mail that is used in many organizations, including my university at Hull. If I want to send a message to our chemistry department, I just put the address "Chemistry Department" on the envelope and drop it in the internal mail. There is a local protocol (called the internal mail system) that makes sure that the message gets there. However, if I want to send a message to the chemistry department at York University, I must put a longer address on the envelope. When the letter gets to the mailroom the staff notice that the destination is not local, and they route it out to the postal system, which sends it to York. This is the "Internet Protocol" for letters.

The Internet is powered by a local protocol (Transport Control Protocol, or TCP) and an internetwork protocol (Internet Protocol, or IP). Put these together, and you have the familiar TCP/IP name that refers to the combination.

You can also use the TCP/IP protocol to connect machines without linking them to the Internet. If you plug your Windows Phone into your PC it actually creates a "tiny internet" between the two machines so that the Zune software can transfer media to and from your phone.

Anything that you send via the Internet will be transferred using one or more individual packets. Each packet contains the address of the destination and each packet is numbered, so that missing packets can be detected and packets can be put into the right order when they are received (if you want that). If you need to transfer a large file, this will be broken down into a number of packets.

Networks and Media

Note that links between different parts of a network can use different media. A laptop could connect via WIFI to a home network which would be connected via the service

provider to the Internet backbone. The internet backbone is connected to your mobile phone provider which then sets up a network connection from a cell phone tower to your Windows Phone. This is how you can use your Windows Phone to message your friends. Furthermore many devices have multiple network connections. Your phone can talk to the Internet via a cell connection to your mobile phone provider or it can use your home WIFI. The good news as far as our programs are concerned is that the way we make connections is independent of the precise connection.

Networks and Protocols

A protocol is a set of rules that tells you how to behave in a certain situation. There is a protocol that tells you which knife and fork to use in a posh banquet and another that tells you to kiss a maiden on her hand having just rescued her from a dragon. You already have one with your friend, where you have agreed on the meaning of the various messages that you send with your bedroom lights.

In networking terms, a protocol sets out the design of all the messages and how stations in a network should cooperate to move data around. There are essentially two levels at which this must take place. There must be a “local-level” protocol that lets local stations (ones on the same piece of physical media) exchange data, and there must be an “internetwork” protocol that allows messages to be sent from one local network to another. The protocol should also extend to cover how systems on the network are addressed and discovered by other systems.

The Internet is based on a large number of different protocols which all work together to describe how systems can work together to exchanged data.

Finding Addresses on the Internet using DNS

The Internet addresses everything by their 32 bit network/system values, which are usually expressed as four “octet” values. For example the address on the internet of the University of Hull web server is 150.237.176.24. This 32 bit value is what Internet messages use to find their way to the web server system at our university. An address is split into 150.237 which is the address of the university network in the world and 176.24 which is the address on the Hull network of our web server. This form of a network address is called the Internet Protocol or *IP* address of a system.

However, nobody wants to have to remember an IP address like this just because they want to see our university web pages – beautiful though they are. People would much rather use a name like `www.hull.ac.uk` to find the site. To solve this problem a computer on the Internet will connect to a *name server* which will convert hostnames into IP addresses. The system behind this is called the *Domain Name System*. It is a collection of servers who pass naming requests around amongst themselves until they find a system with authority for a particular set of addresses who can match the name with the required address. We can think of a name server as a kind of “directory enquiries” for computers. In the old days if I wanted to know the phone number of the local movie house I would ring up directory enquiries who would tell me this. When a computer wants to know the IP address of a web site a user wants to visit it will ask its Domain Name service to provide the result.

Addresses and Subnets

You might think that if every device on the internet has an IP address it should be possible to uniquely address that specific device. In other words, if you get hold of the IP address of my phone you should be able to send network messages directly to it. After all, every phone in the world has a unique number, doesn't it?

This is not actually the case though. Because of the worldwide shortage of IP addresses many networks are arranged into subnets. In this arrangement systems on the subnet have IP addresses that work on that subnet but are not visible to machines outside the subnet. The subnet router, which has an IP address on the “proper” internet, sends messages on behalf of all the machines on the subnet. When messages come from the


internet the router performs some address translation to ensure the packet is sent to the correct local system. This arrangement works well, and is used by companies and home networks to reduce the number of “worldwide” IP addresses that they need. It is also used by mobile phone carriers. This rather like a company having an internal telephone system with many staff telephones and only a few external phone numbers. Calls to the company phone number are switched to particular phones inside the company.

Unfortunately this internet addressing limitation can make it difficult for two mobile devices to directly exchange data, as their addresses are not meaningful to each other. The only way to solve this problem is to use a third party machine which is on the wired network and has a “proper” IP address. Both phones can originate a connection to this third machine and use it to pass data backwards and forwards.

Of course, if the phones are both connected to the same WIFI network when you are at home it should be possible for them to connect directly as they will both be on the same physical network.

Networks and Ports

We now know how the Internet sends messages from one computer to another. Each computer has an address that identifies the internet the network the computer is connected to and the address of that computer on the network. Now we have to add another layer to the protocol, and consider ports.



If we go back to our original network, where we are communicating with our friend by flashing our bedroom light, we have a problem if we want to talk about different topics at the same time. If we wanted to send messages about what is coming up on TV, and also order food, we would have to improve our protocol so that some messages could be flagged as being about TV shows, and others identify the type of pizza we want.

In the case of a computer system we have exactly the same problem. A given computer server can provide a huge variety of different services to the clients that connect to it. The server might be sending web pages to one user, sharing video with another and hosting a multiplayer game all at the same time. The different clients need a way of locating the service that they want on from the server.


The Internet achieves this by using “ports”. A port is a number which identifies a particular service provided by a computer. Some of the ports are “well known”, for example port number 80 is traditionally the one used for web pages. In other words, when your browser connects to a web page it is using the internet address of the server to find the actual computer and then connecting to port 80 to get the web page from that server.

When a program starts a service it tells the Internet software which port that service is sitting behind. When messages arrive for that port the program is passed these messages. If you think about it, the Internet is really just a way that we can make one program talk to another on a distant computer. The program (perhaps a web server) you want to talk to sits behind a port on a computer connected to the Internet. You run another program (perhaps a web browser) that creates a connection to that port which lets you request web pages and display them on your computer.

Programmers can write programs that use any port number, but many networks use *firewalls* that only allow particular packets addressed to particular well know ports to be passed through. This reduces the chance of systems on the network coming under attack from malicious network programs.

Connections and Datagrams

The Internet provides two ways systems can exchange information: connections and datagrams. A datagram is a single message that is sent from one system to another, in the same way that we could just flash the lights in our inter-house network to deliver a message to someone who may or may not be watching. This is like flashing your bedroom light three times to ask for pizza and then just waiting for someone to turn up with your favourite pepperoni thin and crispy.



If you want to be sure that a message has got through you could agree with your friend that she would flash her light once to indicate that she has received it. Then perhaps you could send another message asking for drinks. When she was leaving to fetch the pizza, she could flash her light twice to indicate that she was “going off the air.” This would be the basis of a *connection* between the two of you.

When two systems are in a connection, they have to perform extra work to manage the connection itself. When one system sends a message that is part of a connection the network either confirms the message was successfully transferred (once the network has received an acknowledgement) or gives an error saying that it could not be delivered.

Connections are used when it is important that the entire message gets through. When your browser is loading a Web page, your computer and the Web server share a connection across the network. This makes sure that all parts of the Web page are delivered and that any pieces that don’t arrive first time are retransmitted. The effort of setting up and managing the connection and requesting retransmission when things fail to turn up means that data in connections will be transferred more slowly. Managing a connection places heavier demands on the systems communicating this way.

The Internet has a protocol, called Transmission Control Protocol (or TCP), which describes how to set up and manage a connection between two network stations.

Datagrams are used for sending data in situations where you don’t need an acknowledgement. If you were flashing your bedroom light every ten seconds to show the score in a football game it wouldn’t matter if anyone receiving your messages missed a message as they would just have to wait for the next one to arrive. If a datagram is lost by the network, there is no point in asking for it again because by the time the replacement arrives, it will be out of date. Datagrams work well when you want to send data as fast as you can and it doesn’t matter if some gets lost on the way. They are used for streaming media, where moving video is being sent and the priority is to get the signal delivered as fast as possible.

The internet has a protocol called User Datagram Protocol (or UDP) that describes how systems can send datagrams between each other.

The Windows Phone networking system can send datagrams or create connections, depending on what you want to do. You need to decide which connection type you should use in a particular situation.

7.2 Creating a User Datagram Protocol (UDP) Connection

We are going to start by creating a program that will use UDP to exchange messages with a remote service. The remote service we are going to use is called “echo”. As the name implies, we send echo a message and the service just sends the same message right back to us. The echo service is very useful for testing a connection. It was one of the first ever Internet services and was given port number 7. Nowadays not all systems

have this service running, but we can start the echo service on our PC and use this to explore how datagrams are sent and received. We will see how to set up this service later in the section, when we actually have a program that we want to test.

We are going to write a method that uses UDP to send a message to a distant server. The user will type in the message that they want to send and identify the host they want to send to and the program will do the rest. The text will be entered into TextBox elements in our Silverlight program. The method will look like this:

```

/// <summary>
/// Send a message to a host using UDP
/// </summary>
/// <param name="message">message to send</param>
/// <param name="hostUrl">Url of host</param>
/// <param name="portNumber">number of port</param>
/// <returns>empty string or error message</returns>
public string SendMessageUDP(string message, string hostUrl,
                             int portNumber)
{
    // Creates a socket and sends the message
    // returns an error message anything fails
    // returns an empty string if the message was sent OK
}

```

We can use the method like this:

```

const int ECHO_PORT = 7;
string sendResponse = SendMessageUDP(MessageTextBox.Text,
                                     HostTextBox.Text, ECHO_PORT);

```

This method takes text from two Silverlight TextBoxes that contain the message text and the name of the host and sends them to the ECHO_PORT on the remote server. The above call will send to port number 7, but it could be used to send a datagram to any port.

To manage the connection between the program and the network the program will create an instance of the `Socket` class.

The Socket class

We are well used to objects standing in for something that has meaning in the real world. We use Silverlight elements to represent user interface components and we use streams to represent connections to files. The `Socket` class is how a program represents a connection to an internet resource. If a program wants to connect to a number of different network services at the same time it must create a number of socket instances.

Note that although a Windows Phone has many different ways of physically connecting to the Internet a program will use the same socket technology to manage the connection. The layer underneath the `Socket` will select the most appropriate medium for the connection. Normally this will mean using WIFI if that is available. If there is no WIFI the connection will be made using the slower and less reliable cellular phone network.

Unlike desktop PCs with their fixed, wired, connections there is always the possibility that a Windows Phone has no network coverage, or the user has selected “flight” mode and turned off the network completely. In these situations it is not acceptable for a program to just crash. Instead it must either work around the missing network or give the user sensible status messages.

The `Socket` class is described in the `System.Net.Sockets` namespace, so we need to include this in any C# program that wants to use this class:

```

using System.Net.Sockets;

```

The network library files are already part of the project, so there is no need to add any extra resources.

The starting point for a network connection is the creation of a socket instance to represent the connection that the program wants to make.

```
Socket hostSocket;
```

When a socket is created the constructor is given the settings to be used for that particular socket type.

```
hostSocket = new Socket(AddressFamily.InterNetwork,
                        SocketType.Dgram, ProtocolType.Udp);
```

This creates a new `Socket` which uses internetnetwork V4 addressing and is *datagram* based using the UDP protocol. This is the basis of a datagram connection. We will use the tcp connection type later on.

Sockets and the SocketAsyncEventArgs class

Now that we have our socket we need to configure it and give it the data that we want to send. To do this we create an instance of the `SocketAsyncEventArgs` class and use this to configure the `hostSocket` value.

The `SocketAsyncEventArgs` class contains the properties for a socket which will be used to create an *asynchronous* connection.

```
// Create the argument object for this request
SocketAsyncEventArgs socketEventArgs = new SocketAsyncEventArgs();
```

The first thing we need to do is set the *endpoint* of the connection to identify the system and port that we want to communicate with.

Setting the EndPoint of a SocketAsyncEventArgs value

An endpoint is the thing that we want to talk to over the internet. It is made up of a host address and the number of the port on that host. In effect this is identifying the program on the distant machine that our program is trying to connect to. The endpoint with the internet address `www.hull.ac.uk` and the port number `80` actually identifies the Web server program on a machine on the University of Hull campus. We want our endpoint to identify an echo program (behind port 7) on a server.

The `SocketAsyncEventArgs` class exposes a property called `RemoteEndPoint` that we set to describe the connection endpoint we want the `Socket` to use.

```
socketEventArgs.RemoteEndPoint =
    new DnsEndPoint(hostUrl, portNumber);
```

The statement above creates a new `DnsEndPoint` value from the `hostUrl` (the internet address of the host) and the `portNumber` (number of the port on that host) that we want to use. The `DnsEndPoint` class is very useful, as it will perform DNS lookups for us. We have already seen that systems on the internet use the Domain Naming System to perform a mapping between the names we want to use (`www.hull.ac.uk`) and the addresses that are used to route messages (`150.237.176.24`).

When a `DnsEndPoint` instance is created it is given the text url of the target system and will use the Doman Name System to find the actual IP address that is required by the socket to make the connection. If this address cannot be found the send will fail and the program will be given an appropriate error code.

Encoding the message we want to send

The next thing we need to do is load a message into the `SocketAsyncEventArgs` that describes what we want the `Socket` to do. This is the actual payload of our network action. When your browser fetches a web page from a server the data payload of the

message is the text of the web request that identifies the page you want to read. In our echo program it will be the text that we want the echo service to send back to us.

```
// Put the message into the buffer
byte[] messageBytes = Encoding.UTF8.GetBytes(message);
socketEventArgs.SetBuffer(messageBytes, 0, messageBytes.Length);
```

The internet works in terms of “octets” which are the 8 bit data values it moves around. This means that we can use an internet connection to transfer any kind of data, whether it is images, sounds, video or text, as long as we can break that data down into 8 bit values. However, it also means that before we can send any kind of data we have to encode it into a sequence of octet values.

The Echo service that we are going to use returns the octet sequence that it receives. This means that it could echo an image, a sound or anything else that we send in the form of data.

To convert text into 8 bit values we can make use of the `Encoding` class which is described in the `System.Text` namespace:

```
using System.Text;
```

This class contains a number of encoding methods that can take a string and convert it into a sequence of 8 bit values in an array of bytes. The UTF8 encoding scheme provides a mapping for UNICODE characters (like the ones manipulated by the Windows Phone) into sequences of bytes.

The `GetBytes` method takes a string and returns the contents of the string encoded as an array of 8 bit byte values that can be transferred over the internet. When we get the response back from the echo server we will have to reverse this translation.

Handing the response from the Socket connection

We have now identified the end point we are sending the message to and also set the message that we want to send. The final thing we need to do is specify the method that will be called when the socket completes the action. We need to do this because the `SocketAsyncEventArgs` class, as its name implies, describes a connection that is managed asynchronously.

There are two ways you can do things, synchronous and asynchronous. When you do something synchronously what we are really saying is “while you wait”. If I have my car serviced “synchronously” it means that I stand in the garage waiting for the “service my car” method to complete, at which point I can take my car and drive home. If I have my car serviced “asynchronously” this means that I go off and do some shopping while the job is being done. At some point the garage will ring me up and deliver a “service complete” event. At which point I go and pick up my car and drive home.

Making asynchronous use of data services is very sensible, particularly on a device such as a Windows Phone. The limitations of the network connection to the phone mean that a network connection may take several seconds to deliver an answer. A program should not be stopped for the time it takes for the server to respond.

We are used to writing programs that respond to events generated by display elements. You make a program respond to events from network requests in exactly the same way. The `SocketAsyncEventArgs` class contains a `Completed` property which generates events when a network action completes. When an action completes we need to run some code that will check to see if the network request has worked. This code will be connected to the `Completed` property.

```

public string SendMessageUDP(string message, string hostUrl,
                             int portNumber)
{
    // String to hold the response message for a caller
    // to SendMessage
    string response = "";

    // Create the argument object for this request
    SocketAsyncEventArgs socketEventArgs =
        new SocketAsyncEventArgs();

    // Make an endpoint that connects to the server
    socketEventArgs.RemoteEndPoint = new DnsEndPoint(hostUrl,
        portNumber, AddressFamily.InterNetwork);

    // Put the message into the buffer
    byte[] messageBytes = Encoding.UTF8.GetBytes(message);
    socketEventArgs.SetBuffer(messageBytes, 0,
        messageBytes.Length);

    // Connect an event handler to run when a request completes
    socketEventArgs.Completed +=
        new EventHandler<SocketAsyncEventArgs>
            (delegate(object s, SocketAsyncEventArgs e)
            {
                if (e.SocketError == SocketError.Success)
                {
                    response = "";
                }
                else
                {
                    response = e.SocketError.ToString();
                }
                // Enable any threads waiting on this thread
                transferDoneFlag.Set();
            });

    ...
}

```

We can bind a method to the `Completed` event which is exposed by a `SocketAsyncEventArgs` value. The code above does this binding, but it works in a way you might not have seen before. It binds a block of code directly to the event itself. This is a good idea because it means that the event handler code can access local variables in the `SendMessageUDP` method, in this case the string variable called `response` that we are using to hold error reports. This can be a bit confusing though. Remember that the statements in the event handler (those shaded in the above code) are the ones that run when the event fires. This code is not run when `SendMessageUDP` is called, only when a socket action has completed.

The `Completed` event is passed a reference to the `SocketAsyncEventArgs` value that we set up to describe the transaction. This exposes a status flag called `SocketError` which our program can test to determine whether or not the message was sent successfully.

```

if (e.SocketError != SocketError.Success)
{
    response = "";
}
else
{
    response = e.SocketError.ToString();
}

```

The code above tests the result flag and if the message did not get through it sets the string `response` to a description of the error that was detected. I'm using an error reporting technique that I call "Silence is Golden". If the method completes with the `response` variable holding an empty string this means the message has been sent successfully. If `response` holds anything it means bad news.

Flags and Threads

The second thing that the event handler does is set a flag to indicate that our program has had a response from the socket:

```

// Enable any threads waiting on this thread
transferDoneFlag.Set();

```

To understand what this means we have to think a bit about how asynchronous operations work. An asynchronous operation happens "in the background". At some point the action completes and calls an event method to indicate that it has finished. This is just like me dropping my car off at the garage for a service. The mechanic will work on my car while I go off to do something else, perhaps relax in a café with a good book. At some point my car will be ready, and the garage will give me a call to let me know.

In .NET asynchronous actions are performed by threads of execution. When the phone is switched on it actually contains many different threads that are looking after different aspects of the phone behaviour. Our program is just one more thread that is active on the device. Programs can start further threads to perform asynchronous activities and sometimes a thread will need to wait for another one to finish (rather like me sitting reading a book and waiting for a call from the garage). The .Net framework provides a means by which threads can communicate by using the `ManualResetEvent` class.

```

static ManualResetEvent transferDoneFlag =
    new ManualResetEvent(false);

```

The program contains a variable called `transferDoneFlag`. This can be used as a flag. This has been made `static` so that only one flag will exist for all instances of the transfer class. This is because we can only create a limited number of `ManualResetEvent` instances on a given system. Note that this could lead to problems if we had multiple socket requests active at the same time, since they would all be using the same flag variable. In this situation we would have to add extra code to manage the different requests.

The `transferDoneFlag` variable has two possible states. It is initially set to `false` (i.e. not set). A thread can wait for another thread to set this flag (change it to `true`) by using the `WaitOne` method:

```
transferDoneFlag.WaitOne();
```

This will stop the present thread until another thread sets the `transferDoneFlag` flag:

```
transferDoneFlag.Set();
```

In "waiting for the car to be serviced" terms we would use the flags as follows:


```
serviceDoneFlag.WaitOne(); // Me drinking coffee until car ready
pickUpCar();              // method to pick up the car
```

...

```
doingCarService(); // Garage doing service on car
serviceDoneFlag.Set(); // Garage calling me to say car is ready
```

If the two threads are arranged like this it means that the `pickUpCar` method in the “me” thread will not run until the `doingCarService` at the garage has completed.

One problem with the above design is that the `WaitOne` method will cause a thread to pause until the flag is set. If the flag is never set the thread will be paused for ever. The same thing would happen if the garage forgot to call me, or I had given them the wrong phone number. With the code above I would be sat waiting for ever, which would not be good. In real life I’d probably sit for an hour or so and if I’d not heard anything from the garage I’d give them a call. This is called a timeout, and we can do just the same thing with the `WaitOne` method. We can give the `WaitOne` method a timeout value in milliseconds.

```
transferDoneFlag.WaitOne(1000); // Wait for 1 second
```

The waiting thread will continue if the flag is set by another thread or after one second. This is actually rather impatient. Network requests can take a while to complete, especially when you add the time taken to look up the address of a system, with this in mind a timeout of several seconds is appropriate.

This approach is especially useful on the phone platform, as there is a much higher chance of a network request failing on a mobile device. When you write code that uses the network you must make sure that you include timeouts and your program does sensible things when a connection is not possible.

You might be thinking that making one thread wait for another is a bad idea for performance reasons. However this is not the case. When a thread is waiting for a flag it is actually suspended and using no system resources. The operating system manages the threads and will “wake up” any threads waiting for flags.

Resetting the Flag

The technique we are using is that the “foreground” thread (the one running in the `SendMessageUDP` method) will start a socket request and then wait for the “background” thread (the one running in the `Completed` event handler) to set the flag indicating it has completed.

If the foreground thread is going to wait for the flag to be set it must first clear the flag.

```
transferDoneFlag.Reset();
```

If we didn’t do this it might mean that our program will not wait when it should, as the flag might already be set when the `WaitOne` method is called.

Sending the message and waiting for a response

Now that we have set up the address of the endpoint of the message, the contents of the message and what to do when the message completes we can send the message to the remote system:

```
// Send the message
hostSocket.SendToAsync(socketEventArgs);

// Wait on the event
// Will continue after the timeout or
// if a response is received
transferDoneFlag.WaitOne(MESSAGE_TIMEOUT_MSECS);
```

The Complete Method

```

static ManualResetEvent transferDoneFlag =
    new ManualResetEvent(false);
const int MESSAGE_TIMEOUT_MSECS = 10000;

public string SendMessageUDP(string message, string hostUrl,
    int portNumber)
{
    string response = "Timed Out";

    SocketAsyncEventArgs socketEventArgs =
        new SocketAsyncEventArgs();

    socketEventArgs.RemoteEndPoint =
        new DnsEndPoint(hostUrl, portNumber,
            AddressFamily.InterNetwork);

    byte[] messageBytes = Encoding.UTF8.GetBytes(message);
    socketEventArgs.SetBuffer(messageBytes, 0,
        messageBytes.Length);

    socketEventArgs.Completed +=
        new EventHandler<SocketAsyncEventArgs>
            (delegate(object s, SocketAsyncEventArgs e)
            {
                if (e.SocketError == SocketError.Success) {
                    response = "";
                }
                else {
                    response = e.SocketError.ToString();
                }
                transferDoneFlag.Set();
            });

    transferDoneFlag.Reset();

    hostSocket.SendToAsync(socketEventArgs);

    transferDoneFlag.WaitOne(MESSAGE_TIMEOUT_MSECS);

    return response;
}

```

This is the complete method that we have been working on over the last few pages. The response string is initially set to “Timed Out”. If the socket does not respond before the timeout, this is the string that is returned from the method. If the socket request does respond and the response indicates success, the method will return an empty string.

Sending a Message

The method can be called by giving it the text to send, the name of the host and the port to use. It will return a response which we can then test in our program:

```

string sendResponse = SendMessageUDP(
    MessageTextBox.Text, HostTextBox.Text, ECHO_PORT);

```

If the address of the host is invalid the message will fail. However, that is pretty much the only way that this method call will fail. Remember that this is a datagram. We have no way of knowing if the response was received by the remote system unless we get something back from it.

Receiving a Message

If we send a message to a host that has the echo service running behind its echo port we would expect to get a response back from the host containing a copy of the message we sent. We therefore need to set up another socket connection to the remote host to read back the response. We could write another method to send this read request:

```
public string ReceiveMessageUDP(int portNumber, out string result)
{
    string response = "Error: Request Timed Out";
    string message = "";

    SocketAsyncEventArgs socketEventArgs =
        new SocketAsyncEventArgs();
    socketEventArgs.RemoteEndPoint =
        new IPEndPoint(IPAddress.Any, portNumber);

    byte [] responseBytes = new Byte[MAX_BUFFER_SIZE];
    socketEventArgs.SetBuffer(responseBytes, 0, MAX_BUFFER_SIZE);

    socketEventArgs.Completed +=
        new EventHandler<SocketAsyncEventArgs>
            (delegate(object s, SocketAsyncEventArgs e)
            {
                if (e.SocketError == SocketError.Success) {
                    response = "";
                    message = Encoding.UTF8.GetString(e.Buffer, e.Offset,
                                                    e.BytesTransferred);

                    message = message.Trim('\0');
                }
                else {
                    response = e.SocketError.ToString();
                }

                transferDoneFlag.Set();
            });

    transferDoneFlag.Reset();

    hostSocket.ReceiveFromAsync(socketEventArgs);

    transferDoneFlag.WaitOne(MESSAGE_TIMEOUT_MSECS);

    result = message;
    return response;
}
```

This method looks very similar to the send message method. It will listen on a particular port for a response from the network. It differs from the send method in a few respects though:

```
socketEventArgs.RemoteEndPoint =
    new IPEndPoint(IPAddress.Any, portNumber);
```

The endpoint for a receive request can be set to an `IPEndPoint` value which can receive on any IP Address. This would mean that any remote system could send a message to the socket on our system and we would receive it. If we only want to receive messages from one particular system we could put the address of that system here.

Since we are requesting some data from a distant system we need to create a buffer to hold the bytes that are sent back to us.

```
byte [] responseBytes = new Byte[MAX_BUFFER_SIZE];
socketEventArgs.SetBuffer(responseBytes, 0, MAX_BUFFER_SIZE);
```

These two statements create a byte array to receive the message and then set the buffer on the `socketEventArgs` value to this array.

We also need to make some changes to the code that runs when the receive request completes. When the response is received the program must fetch the data that the remote system has sent us.

```
socketEventArgs.Completed +=
    new EventHandler<SocketAsyncEventArgs>
        (delegate(object s, SocketAsyncEventArgs e)
{
    if (e.SocketError == SocketError.Success) {
        response = "";
        // Retrieve the data from the buffer
        message = Encoding.UTF8.GetString(e.Buffer, e.Offset,
            e.BytesTransferred);
        message = message.Trim('\0');
    }
    else {
        response = e.SocketError.ToString();
    }
    transferDoneFlag.Set();
});
```

The string `response` is set to the result of the request as before. The response is an empty string if the request succeeded or an error message if something went wrong.

If the receive request succeeds the program must convert the stream of 8 bit values that have been sent into the data our program will work with. The code uses the `Encoding.UTF` class to convert the buffer that has been received into a string that can be displayed. Finally the event handler sets a flag to allow a thread to synchronise on the text received event. Remember that, as with the send request we saw earlier, this code runs when the asynchronous request to the socket has completed.

Finally the method must return the string of text that was received.

```
// Clear the flag this is set when the request is completed
transferDoneFlag.Reset();

// Send off the read request to the socket
hostSocket.ReceiveFromAsync(socketEventArgs);

// Pause this thread until the socket responds or the wait
// times out
transferDoneFlag.WaitOne(MESSAGE_TIMEOUT_MSECS);

// copy the text that was received into the out result
result = message;
// Return status of our request, an empty string if it worked
return response;
```

This sequence of statements resets the thread synchronisation flag and then makes a Socket request to receive some data. It then waits for the socket to complete and, when it does, it sets the result parameter to the message that was received and returns the response.

A call of this method looks like this:

```
string receiveResponse;
string message;

receiveResponse = ReceiveMessageUDP(ECHO_PORT, out message);
```

Sending a datagram to the Echo service

We can now put together some statements that will make a call to an Echo service. We will first call `SendMessageUDP` to send a message to the service and then call `ReceiveMessageUDP` to get the echo back.

```
string sendResponse = SendMessageUDP(MessageTextBox.Text,
                                     HostTextBox.Text, ECHO_PORT);

if (sendResponse.Length == 0) {
    string receiveResponse;
    string message;

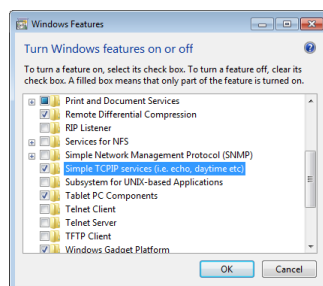
    receiveResponse = ReceiveMessageUDP(ECHO_PORT, out message);
    if (receiveResponse.Length == 0) {
        ResponseTextBox.Text = message;
    }
    else {
        ErrorTextBox.Text = "Receive Error: " + receiveResponse;
    }
}
else {
    ErrorTextBox.Text = "Send Error: " + sendResponse;
}
```

If the `SendMessageUDP` returns an empty string it means that the message was sent successfully and the program can call the `ReceiveMessageUDP` method to get the text back from the Echo server.

The solution in *Demo 02 UDPNetworkClient* contains a program that connects to the echo service on a host. You can type in some text and this will be sent to the echo port on the host that you identify. The program will display any errors that are detected.

Setting up an Echo Server

To test the example program above you need to have a system that provides an echo server. Fortunately a Windows PC can work as an echo server, but this service has to be enabled before you can use it. To turn on the service you first need to open the Windows Control panel and open “Turn Windows Features on or off”. This will open the Windows Features dialog.



Find the “Simple TCPIP services” option and ensure that it is selected. Then click OK. You will now have to reboot your PC to start the services running. Once you have done this you will be able to enter your computer name into the address for a service and it will respond to echo requests.



This shows the echo client working on my machine once I had started the service. My machine is called “martini”. If you happen to know any servers on the internet which provide the echo service you can enter their names in the Host dialog box.

7.3 Creating a Transmission Control Protocol (TCP) Connection

The problem with User Datagram Protocol (UDP) is that systems that use it to transfer data do not know if the recipient received anything. If you use the demo program above to connect to a machine that is on the Internet but not running the echo service your outgoing message will be sent with no reported errors. However, as there is not a service listening behind port 7 at the destination it will not receive anything back in response. This means that the receive request will timeout when the host does not respond with the echoed message.

If we want to create reliable connections we have to use a different protocol. This is the Transmission Control Protocol we saw earlier. This provides a connection between two machines. You can compare Transmission Control Protocol (TCP) with User Datagram Protocol (UDP) by comparing SMS text messages with phone calls. If I send someone a text message I have no idea whether or not they received it. I will only know that they have got the message when I get a response from them. This is the service provided via UDP. However if I call them on the phone we make a connection. I can send and receive messages during the call and be sure that they are at the other end listening. This is the service provided by TCP.

When we use TCP we actually create an on-going connection. Messages are created and transferred in exactly the same way, but each message is part of a call, rather than being an individual event. The call will exist until one of the parties in it decides to terminate the conversation. The response that we receive from sending a message will allow our programs to determine whether the message was delivered successfully or not. At the end of the call one of the parties will close the connection.

Reading a Web Page

The method we are going to create will look like this:

```
private string RequestWebPage(string url, string page,
                             out string pageContent)
```

It is given the web address of the server, the name of the page to fetch and places the text in the third parameter. It returns a response that is an empty string if the request

worked or an error description if the request failed. The following call would get the home page of my blog.

```
string resp;
string page;
resp = RequestWebPage ("www.robmiles.com",
                      "index.html", out page);
```

The World Wide Web uses a protocol called Hyper Text Transfer Protocol or HTTP. This protocol defines the language used by a browser to request particular pages of content from a web server. For example, to get a web page from the server the browser will send the command GET, along with the location of the web page on the server.

Because it is important that all the parts of a web page are delivered to the client the HTTP protocol uses a TCP connection. To fetch a web page a program must do the following:

1. Create a TCP connection to the web server using port 80, which is the one reserved for HTTP requests.
2. Send the GET request to the server on this connection.
3. Read the web page content from the server. This may require several read operations if the page is large.
4. Close the connection.

Setting up the TCP Connection

We can use the same format for the connection methods as we did before.

```
private string ConnectTCP(string host, int port)
{
    string response = "Connect timed out";

    DnsEndPoint hostEntry = new DnsEndPoint(host, port);

    hostSocket = new Socket(AddressFamily.InterNetwork,
                           SocketType.Stream, ProtocolType.Tcp);

    SocketAsyncEventArgs socketEventArgs =
        new SocketAsyncEventArgs();
    socketEventArgs.RemoteEndPoint = hostEntry;

    socketEventArgs.Completed +=
        new EventHandler<SocketAsyncEventArgs>
            (delegate(object s, SocketAsyncEventArgs e)
            {
                if (e.SocketError == SocketError.Success) {
                    response = "";
                }
                else {
                    // Retrieve the result of this request
                    response = e.SocketError.ToString();
                    // Ensure the socket is flagged unavailable
                    hostSocket = null;
                }
                transferDoneFlag.Set();
            });

    transferDoneFlag.Reset();

    hostSocket.ConnectAsync(socketEventArgs);

    transferDoneFlag.WaitOne(MESSAGE_TIMEOUT_MSECS);
    return response;
}
```

This code is very similar to that we used to send a UDP message. We set up a `SocketAsyncEventArgs` value and then use this to configure the call that we make to the socket. However, in this case we call the `ConnectAsync` method:

```
hostSocket.ConnectAsync(socketEventArgs);
```

If this method completes successfully it will return with a TCP connection set up to the remote system. We can then use this to send and receive messages until the connection is closed by either party. However, note that the connection only lasts for the transmission of a single page of content. If a user navigates to a different page this will require the creation of another TCP connection to fetch the new page.

Sending a Message to a TCP Connection

To send a message we simply need to assemble the message text and then use the socket to send it. The socket that we have set up contains properties that we can use to configure the message transfer:

```
private string SendMessageTCP(string message)
{
    if (hostSocket == null) return "Send socket not open";

    string response = "Connect timed out";

    SocketAsyncEventArgs socketEventArgs =
        new SocketAsyncEventArgs();

    // Use the endpoint in the socket to set up the arguments
    socketEventArgs.RemoteEndPoint = hostSocket.RemoteEndPoint;

    byte[] messageBytes = Encoding.UTF8.GetBytes(message);
    socketEventArgs.SetBuffer(messageBytes, 0, messageBytes.Length);

    socketEventArgs.Completed +=
        new EventHandler<SocketAsyncEventArgs>
            (delegate(object s, SocketAsyncEventArgs e)
            {
                if (e.SocketError == SocketError.Success) {
                    response = "";
                }
                else {
                    response = e.SocketError.ToString();
                    CloseTCP();
                }

                transferDoneFlag.Set();
            });

    transferDoneFlag.Reset();

    hostSocket.SendAsync(socketEventArgs);

    transferDoneFlag.WaitOne(MESSAGE_TIMEOUT_MSECS);

    return response;
}
```

This code is almost identical to the send message that we created earlier for UDP connections. This version checks to make sure that the socket has been set up before it uses it:

```
if (hostSocket == null) return "Send socket not open";
```

It also closes the socket if the request returns with an error. The `CloseTCP` method just closes the socket if it exists:


```

void CloseTCP()
{
    if (hostSocket != null) {
        hostSocket.Close();
        hostSocket = null;
    }
}

```

Receiving a Message from a TCP Connection

The final method that we need to read a web page using a TCP connection is one to receive a message. Again, this looks remarkably like the ones that we have used to read UDP messages:

```

private string ReceiveMessageTCP(out string message)
{
    message = "";

    if (hostSocket == null) return "Send socket not open";

    string response = "Receive timed out";
    string result = "";

    SocketAsyncEventArgs socketEventArg =
        new SocketAsyncEventArgs();
    socketEventArg.RemoteEndPoint = hostSocket.RemoteEndPoint;

    // Create a buffer for the response
    byte[] responseBytes = new Byte[MAX_BUFFER_SIZE];
    socketEventArg.SetBuffer(responseBytes, 0, MAX_BUFFER_SIZE);

    socketEventArg.Completed +=
        new EventHandler<SocketAsyncEventArgs>
            (delegate(object s, SocketAsyncEventArgs e)
            {
                if (e.SocketError == SocketError.Success) {
                    response = "";
                    if (e.BytesTransferred > 0) {
                        result = Encoding.UTF8.GetString(e.Buffer,
                            e.Offset, e.BytesTransferred);
                        result = result.Trim('\0');
                    }
                }
                else {
                    response = e.SocketError.ToString();
                }
                transferDoneFlag.Set();
            });

    transferDoneFlag.Reset();

    hostSocket.ReceiveAsync(socketEventArg);

    transferDoneFlag.WaitOne(MESSAGE_TIMEOUT_MSECS);

    message = result;

    return response;
}

```

This version of the method checks to see if any bytes have been transferred from the server and if they have it decodes the bytes and returns them as a string result.

Building the Complete Request

Now that we have the methods that we need to build a request for a web page we can put them together to do the job. This method will send the request and then assemble the web page text from the responses. It knows when the transfer is complete as it will receive an empty buffer from the read method call.

```
private string RequestWebPage(string url, string page,
                             out string pageContent)
{
    pageContent = "";

    // Set up the connection
    string response = ConnectTCP(url, 80);

    if (response != "") {
        return response;
    }

    // Send the page request using the HTTP Get command
    response = SendMessageTCP("GET " + page + " HTTP/1.1\r\nHost: "
                              + url + "\r\nConnection: Close\r\n\r\n");

    if (response != "") {
        CloseTCP();
        return response;
    }

    // Repeatedly ask the server for packets until we
    // get an empty one sent back
    string wholePage = "";
    string fetchText;

    do {
        response = ReceiveMessageTCP(out fetchText);

        if (response != ""){
            CloseTCP();
            return response;
        }

        if (fetchText == "") break;

        wholePage = wholePage + fetchText;
    } while (true);

    pageContent = wholePage;

    CloseTCP();

    return response;
}
```

This method calls each of our methods in turn, building up the text of the complete web page from successive responses. If any of the methods returns an error response the method will close the TCP connection and return.

The method is called when the user presses the Load button in our application:

```
private void LoadButton_Click(object sender, RoutedEventArgs e)
{
    string response;
    string webPageText;

    response = RequestWebPage(UriTextBox.Text, PageTextBox.Text,
        out webPageText);

    StatusMessage.Text = response;
    ResponseTextBlock.Text = webPageText;
}
```

This loads the url and the page name from the textboxes and then calls the method to request the web page. It then displays any response that was received.



This shows the program in operation. We can see the web page content that has been loaded.

The solution in *Demo 03 TCPNetworkClient* contains a program will fetch web pages for you using the methods described above.

Page Requests and Threads

At the moment all the activity for the page loading takes place in the button event handler. This means that the button pressed event will not complete until either the page has been fetched and displayed or something has gone wrong. We can see this in the user interface of the program above; the button seems to take a while to return to the “un-pressed” state when we perform a load. This is not necessarily a problem, but if we wanted the user to be able to do other things while our page was loading this would not be possible.

If we want to allow the user to perform actions in parallel – for example load a web page and use the echo server at the same time we have to use multiple threads to do the work. This makes our life much more difficult as developers, but the user will prefer this approach. If you want to do this kind of thing you need to find out more about threading, which works in exactly the same way on Windows Phone as on any other .NET platform.

Data Chunks and the UTF8 Encoding

If you use the program above you might notice that every now and then it will corrupt a web page. This is because of the way that it works. We are using socket methods that fetch an array of bytes from a web server and then convert that array into a string of

text. The encoding method converts the text that we have typed into one or more 8 bit values that are transferred for us. This would work fine if the web server sent us the entire page in response to a single read request, but this is not how networks send things. Networks transfer large amounts of data by breaking it down into smaller chunks. And the network is completely unaware of the nature of the data that it is sending.

This gives us a problem, in that if a UTF8 character is split over two packets it will become corrupted. As an example, the UTF8 character for the pound currency sign (£) is the pair of bytes 0xC2 and 0xA3 in sequence. If the 0xC2 is transferred as the very last character of a transfer and the 0xA3 sent as the first character of the next one my code above will cause this character to be corrupted.

The HTML protocol itself will normally encode these kinds of characters into values that should not cause problems like this, but if we decide to transfer text other than web pages we need to address this. There are two ways to solve it, we can either keep track of “trailing” bytes that the UTF8 decoder did not understand or we can create an array of byte values from the incoming messages and decode the whole array at once. I’ve written some sample code that does the latter, it assembles a large byte array. This is actually quite a useful thing to have, in that it makes it easy for us to receive other kinds of binary data that may be sent from a web host.

The solution in *Demo 04 Binary TCPNetwork Client* contains a program will fetch web pages for you and will assemble the message in bytes and decode them all at once.

7.4 Connecting to a Data Source

We have seen that we can create a TCP connection and read text from a web page. Reading information from network datasources is something that programs do a lot. However, this is rather tedious to do. Fortunately there is a much easier way of reading web pages. We can use the `WebClient` class.

Using the WebClient class

The `WebClient` class is provided as part of the Windows .NET networking library, so if you learn how to use it on the Windows Phone you can also use it in your Windows desktop programs. We can use the class to perform an HTTP (Hyper Text Transfer Protocol) interaction with a remote web server. There are a number of different ways we can use a `WebClient`, we are going to look at one of the simplest.

We are going to create a `WebClient` variable in the `MainPage.xaml.cs` of our application. When the user presses the `loadButton` our application will load the url from a `TextBox` and then display the text content of that page.



Above you can see the program in action, displaying one of the best pages on the web. Note that this is exactly what a browser does, except that a browser will then take the HTML content received from the server and render it along with any associated graphical content.

The first thing we need to do is declare our `WebClient` as a member of the `MainPage` class. This will be used manage the web requests that we are going to make.

```
WebClient client;
```

At the moment we just have the reference to a `WebClient` object. The next thing we need to do is create an actual instance of a `WebClient`. We can do this in the constructor for the page class:

```
public MainPage()
{
    InitializeComponent();

    client = new WebClient();
    client.DownloadStringCompleted +=
        new DownloadStringCompletedEventHandler(
            client_DownloadStringCompleted);
}
```

The constructor creates a new `WebClient` and then binds a method to the `DownloadStringCompleted` event that the `WebClient` exposes. The method, called `client_DownloadStringCompleted`, will display the web page text in the in the `pageTextBlock`:

```
void client_DownloadStringCompleted(object sender,
                                     DownloadStringCompletedEventArgs e)
{
    if (e.Error == null)
    {
        pageTextBlock.Text = e.Result;
    }
}
```

When this method is called it tests to see if any errors were reported. If there are no errors the method displays the result of the download in the `pageTextBlock` element.

The last thing we need to add to our program is the event handler for the button that will start the download of a page at the url that the user has typed in:

```
private void loadButton_Click(object sender, RoutedEventArgs e)
{
    client.DownloadStringAsync(new Uri(urlTextBox.Text));
}
```

When the load button is clicked the following sequence takes place:

1. The text from the `urlTextBox` is used to create a new `Uri` which describes the web site to be loaded.
2. The `DownloadStringAsync` method is called on the `WebClient`. This starts a new web transaction. This method returns immediately, so the button event handler is now completed.
3. Some-time later the `WebClient` will have fetched the string of content at the url. It fires the `DownloadStringCompleted` event to indicate that it has finished, and to deliver the result.
4. The `client_DownloadStringCompleted` method now runs, which checks for errors and displays the received text in the `TextBlock`.

The solution in *Demo 05 Web Scraper* contains this code. You can type in any url and download the html from that page. The program will work on either the emulator or a real device.

Note that this is a very basic web page reading program (as you might expect for only fifteen or so lines of code). It does not stop the user from pressing the Load button repeatedly and interrupting downloads and it does not provide any form of time out or progress display to the user.

7.5 Using LINQ to Read from an XML Stream

The ability to read from the web can be used for much more than just loading the text part of a web page. We can also interact with many other data services, for example Twitter. This uses a REST (Representational State Transfer) protocol to expose the activities of Twitter users. REST makes use of the structure of the url being used to denote the required action. As an example, if you point a web browser at the following url:

`http://twitter.com/statuses/user_timeline/robmiles.xml`

- you will be rewarded with an XML document that contains my most recent Twitter posts. By replacing robmiles with the name of another Twitter user you can read their feed.



This makes it very easy to make a slightly modified version of our WebClient program that assembles a web request for the Twitter server and uses that to read feed information from it. The only change we have to make is to the code that produces the web request:

```
private void loadButton_Click(object sender, RoutedEventArgs e)
{
    string url = "http://twitter.com/statuses/user_timeline/" +
        nameTextBox.Text + ".xml";

    client.DownloadStringAsync(new Uri(url));
}
```

The solution in *Demo 06 Twitter Reader* contains this code. You can type in any twitter username and download the feed from that user. The program will work on either the emulator or a real device.

Structured Data and LINQ

Twitter returns status reports in the form of an XML formatted document. XML is something we know and love (a bit). What we want to do now is pull all the information out of the XML content we get from Twitter and display this information

properly. As great programmers we could of course write some software to do this. The good news is that we don't have to. We can use LINQ to do this for us.

An XML document is very like a database, in that contains a number of elements each of which has property information attached to it. Along with tools to create database queries and get structured information from them, LINQ also contains tools that can create objects from XML data. We are going to use these to read our Twitter feed.

The Twitter feed as a data source

The feed information that you get from Twitter in response to a status request is a structured document that contains attributes and properties that describe the Twitter content. If you have been paying attention during the discussions of XAML and XML you should find this very familiar.

```
<?xml version="1.0" encoding="UTF-8"?>
<statuses type="array">
<status>
  <created_at>Tue Oct 11 11:57:37 +0000 2011</created_at>
  <id>27131245083</id>
  <text>Had 12,000 hits on www.csharpcourse.com for my C# Yellow Book. No idea why.</text>
  <source>web</source>
  <truncated>>false</truncated>
  <in_reply_to_status_id></in_reply_to_status_id>
  <in_reply_to_user_id></in_reply_to_user_id>
  <favorited>>false</favorited>
  <in_reply_to_screen_name></in_reply_to_screen_name>
  <retweet_count>0</retweet_count>
  <retweeted>>false</retweeted>
  <user>
    <id>2479801</id>
    <name>Rob Miles</name>
    <screen_name>robmiles</screen_name>
    <location>Hull, UK</location>
    <description>A Computer Science lecturer and Microsoft MVP.</description>
    <profile_image_url>http://a3.twimg.com/prof_im/157/me2_normal.jpg</profile_image_url>
    <url>http://www.robmiles.com</url>
    <protected>>false</protected>
    <friends_count>21</friends_count>
    <created_at>Tue Mar 27 12:35:28 +0000 2007</created_at>
    <favourites_count>0</favourites_count>
    <utc_offset>0</utc_offset>
    <time_zone>London</time_zone>
    <profile_background_image_url>http://s.twimg.com/a/1286/images/themes/theme1/bg.png
      </profile_background_image_url>
    <notifications>>false</notifications>
    <geo_enabled>>true</geo_enabled>
    <verified>>false</verified>
    <following>>true</following>
    <statuses_count>731</statuses_count>
    <lang>en</lang>
    <contributors_enabled>>false</contributors_enabled>
    <follow_request_sent>>false</follow_request_sent>
    <listed_count>61</listed_count>
    <show_all_inline_media>>false</show_all_inline_media>
  </user>
</status>
</statuses>
```

I've removed some elements to make it fit nicely on the page, but if you look carefully you can find the text of the tweet itself, along with lots of other stuff. The real document contains a number of status elements. Above I have only shown one.

In the same way as we can ask LINQ to query a database and return a collection of objects it has found we can ask LINQ to create a construction that represents the contents of an XML document:

```
XElement TwitterElement = XElement.Parse(twitterText);
```

The `XElement` class is the part of LINQ that represents an XML element. The `twitterText` string contains the text that we have loaded from the Twitter web site using our `WebClient`. Once this statement has completed we now have an `XElement` that contains all the Twitter status information as a structured document. The next thing we need to do is convert this information into a form that we can display. We are going to pull just the status items that we need out of the `TwitterElement` and use these to create a collection of items that we can display using data binding.

Creating Posts for Display

We know that it is very easy to bind a C# object to display elements on a Silverlight page. So now we have to make some objects that we can use in this way.

The items we are going to create will expose properties that our Silverlight elements will use to display their values. The three things we want to display for each post are the image of the Twitter user making the post, the date of the post and the post text itself. We can make a class that holds this information.

```
public class TwitterPost
{
    public string PostText { get; set; }

    public string DatePosted { get; set; }

    public string UserImage { get; set; }
}
```

These properties are all “one way” in that the `TwitterPost` object does not want to be informed of changes in them. This means that we can just expose them via properties and all will be well. The `UserImage` property is a string because it will actually give the uri of the image on the internet.

The next thing we have to do is to use LINQ to get an object collection from the `XElement` that was built from the Twitter feed we loaded.

```
var postList =
    from tweet in twitterElements.Descendants("status")
    select new TwitterPost
    {
        UserImage = tweet.Element("user").Element("profile_image_url").Value,
        PostText = tweet.Element("text").Value,
        DatePosted = tweet.Element("created_at").Value
    };
```

This code will do exactly that. It is well worth a close look. The `from` keyword requests an iteration through a collection of elements in the document. The elements are identified by their name. The program will work through each status element in turn in the document. Each time round the loop the variable `tweet` will hold the information from the next status element in the `XElement`.

The `select` keyword identifies what is to be returned from each iteration. For each selected tweet we want to return a new `TwitterPost` instance with settings pulled from that `tweet`. The `Element` method returns an element with a particular name. As you can see we can call the `Element` method on other elements. That is how we get the `image_profile_url` out of the `user` element for a tweet.

The variable `postList` is set to the result of all this activity. We can make it `var` type, but actually it is a collection of `TwitterPost` references. It must be a collection, since the `from` keyword is going to generate an iteration through something.

We now have a collection of `TwitterPost` values that we want to put onto the screen.

Laying out the post display using XAML

We have seen how easy it is to put a Silverlight element on a page and then bind object properties to it. However, we have a little extra complication with our program. We want to display lots of posts, not just one value. Ideally we want to create something that shows the layout for a single item and use that to create lots of display elements, each of which is bound to a different `TwitterPost` value. This turns out to be quite easy.



Above you can see what I want an individual post to look like. It will have an image on the side, with the date and post text on the right, one above each other.



From a Silverlight point of view this is three panels, arranged as shown above. We can get Silverlight to do a lot of the work in laying these out for us by using the `StackPanel` container element.

The first `StackPanel` runs horizontally. This contains the user image and then another `StackPanel` that runs vertically and holds the data and time above the post text.

```
<StackPanel Orientation="Horizontal" Height="132">
  <Image Source="{Binding UserImage}" Height="73" Width="73"
    VerticalAlignment="Top" Margin="0,10,8,0"/>
  <StackPanel Width="370">
    <TextBlock Text="{Binding DatePosted}" Foreground="#FFC8AB14" FontSize="22" />
    <TextBlock Text="{Binding PostText}" TextWrapping="Wrap" FontSize="24" />
  </StackPanel>
</StackPanel>
```

Above you can see the XAML that expresses all this. The `StackPanel` class is a lovely way of laying out items automatically. We don't have to do lots of calculations about where things should be placed, instead we just give the items and the `StackPanel` puts them all in the right place. You can also see the bindings that link the items to the properties in the `TwitterPost` class. The `UserImage` property in `TwitterPost` is bound to the `Source` property of the image. If we add the complete XAML above to our main page we have an area that can hold a list of Twitter posts.

Creating a complete layout

```
<ListBox Height="442" HorizontalAlignment="Left" Name="tweetsListBox" VerticalAlignment="Top"
        Width="468">
    <ListBox.ItemTemplate>
        <DataTemplate>
            <StackPanel Orientation="Horizontal" Height="132">
                <Image Source="{Binding UserImage}" Height="73" Width="73"
                    VerticalAlignment="Top" Margin="0,10,8,0"/>
                <StackPanel Width="370">
                    <TextBlock Text="{Binding DatePosted}" Foreground="#FFC8AB14" FontSize="22" />
                    <TextBlock Text="{Binding PostText}" TextWrapping="Wrap" FontSize="24" />
                </StackPanel>
            </StackPanel>
        </DataTemplate>
    </ListBox.ItemTemplate>
</ListBox>
```

This is exactly how we do it. This XAML repays a lot of study. The first thing we notice is that this is a `ListBox` element. We've not seen these before but they are, as you might expect, a way of holding a list of boxes. Within a `ListBox` we can put a template that describes what each item in the list will look like. This includes the data binding that will link element properties to our `TwitterPost` instances. It also includes the layout elements that will describe how the post will look.

Creating Posts for Display

The next thing our Twitter reader must do is get the posts data onto the display. Before we do this, it is worth stepping back and considering what we are about to do. This is worth doing because it will help us to understand how Silverlight binds collections to displays.

We could write some test code that made us a collection of `TwitterPost` items:

```
TwitterPost p1 = new TwitterPost
{
    DatePosted = "Tue Oct 11 11:57:37 +0000 2011",
    UserImage = "http://a3.twimg.com/profile_images/150108107/me2_normal.jpg",
    PostText = "This is a test post from Rob"
};

TwitterPost p2 = new TwitterPost
{
    DatePosted = "Wed Oct 12 14:21:04 +0000 2011",
    UserImage = "http://a3.twimg.com/profile_images/150108107/me2_normal.jpg",
    PostText = "This is another test post from Rob"
};

List<TwitterPost> posts = new List<TwitterPost>();
posts.Add(p1);
posts.Add(p2);
```

This code creates two `TwitterPost` instances and then adds them to a list of posts. This is not particularly interesting code. It becomes interesting when we do this:

```
tweetsListBox.ItemsSource = posts;
```

The magic of Silverlight data binding takes the list of posts and the template above, and does this with it:



This is very nice. It means that you can display any collection of data in a list on the page just by setting them as the source for a `ListBox`. All you have to do is create a data template that binds the properties of each instance to display elements and then assign a list of these instances to the `ItemsSource` property of the `ListBox`.

Since we know that we can assign a collection of `TwitterPost` references to a `ListBox` and get this displayed the next statement is not going to come as much of a surprise.

```
tweetsListBox.ItemsSource = postList;
```

At this point our work is done, and the Twitter posts are displayed.




The `ListBox` provides a lot of useful functionality. If you run this program you will find that you can scroll the list up and down by dragging them with your finger. The scroll action will slow down and the items will even “scrunch” and bounce as they move.

The solution in *Demo 7 Complete Twitter Reader* contains this code. You can type in any twitter username and download the feed from that user and display the feed on the screen. When the program starts it is pre-loaded with the test tweets shown above.

7.6 Using Network Services

We can create Windows Phone applications that interact with web servers. If we want to pass data to and from a server using web protocols we can also use the GET and POST elements of HTTP (Hyper Text Transfer Protocol). In this section we are going to go beyond that and investigate how we can create network based services and then consume them on the Windows Phone device.

WCF (Windows Communication Foundation) services provide a way that two processes can communicate over a network. They take care of all the transfer of messages between two systems and can work over a variety of communication media. The services are exposed by a server and then accessed by a client program. In our case the client program will be running on a Windows Phone.




The diagram shows how this works. Software on the phone interacts with a “proxy object” that contains the methods provided by the service. A call to a method on the proxy object will cause a network message to be created that is sent to the service on the server. When the server receives the message it will then call the method code in a server object. The result to be returned by the method will then be packaged up into another network message which is sent back to the client which then sends the return value back to the calling software. We don’t need to worry about how the messages are constructed and transferred. We just have to create our server methods and calls from the client systems.

The service also exposes a description of the methods that it provides. This is used by the development tool (in our case Visual Studio) to actually create the proxy object in the client program. This means that we can create a client application very easily.

Joke of the Day service

To discover how all this works we are going to create a “Joke of the Day” service. This will return a string containing something suitably rib-tickling on request from the client. The user will be able to select the “strength” of the joke, ranging from 0 to 2 where 0 is mildly amusing and 2 is a guaranteed roll on the floor laughing experience.

The application will be made up of two parts, the server program that exports the service and the client program that uses it. The server will be created as a WCF (Windows Communication Foundation) Service application. The client will be a Windows Phone application that connects to the service and requests a joke.



Services are created as Visual Studio projects. They can run inside Visual Studio in a test environment. After testing the project files would be loaded onto a production server.

Services and Clients

Services and clients look like methods. When we write the server we write a method that does something and returns a result. When we write the client we call a method and something comes back. The communication service underneath the programs takes care of packaging the parameters and the method return value. When we create the service we create an interface that defines the method the service will provide.

```
[ServiceContract]
public interface IJokeOfTheDayService
{
    [OperationContract]
    string GetJoke(int jokeStrength);
}
```

We are only going to create one method in our service, but a service could provide many methods if required. This method only accepts a single parameter and returns a result, but we could create more complex methods if we wish.

Once we have the interface we can now create the method to provide the service that the service describes. The `[ServiceContract]` and `[OperationContract]` attributes on the classes and methods provide the means by which the build process can generate the service descriptions that will be used by clients to discover and bind to these services.

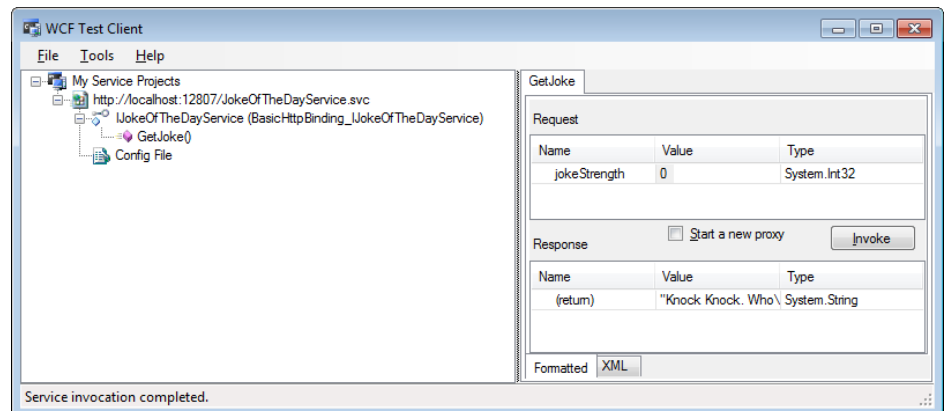
```

public class JokeOfTheDayService : IJokeOfTheDayService
{
    public string GetJoke(int jokeStrength)
    {
        string result = "Invalid strength";
        switch (jokeStrength)
        {
            case 0:
                result =
                "Knock Knock. Who's there? Oh, you've heard it";
                break;
            case 1:
                result =
                "What's green and hairy and goes up and down? A gooseberry in a lift";
                break;
            case 2:
                result =
                "A horse walks into a bar and the barman asks 'Why the long face?";
                break;
        }
        return result;
    }
}

```

As you can see from above, the method is quite simple. As are the jokes. The input parameter is used to select one of three jokes and return them. If there is no match to the input the message “Invalid strength” is returned.

If we create a service like this we can use the WCF Test Client to invoke the methods and view the results.



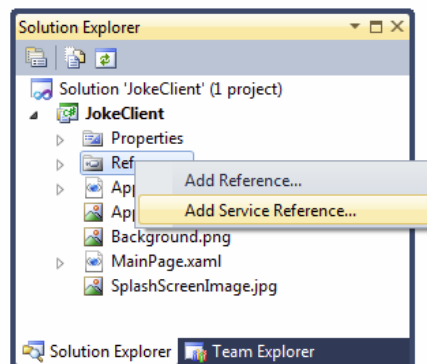
This tool lets us call methods in the service and view the results that they return. You can see the results of a call to the method with a parameter of 0. We can also view the service description in a browser:



This gives a link to the service description, as well some sample code that shows how to use it.

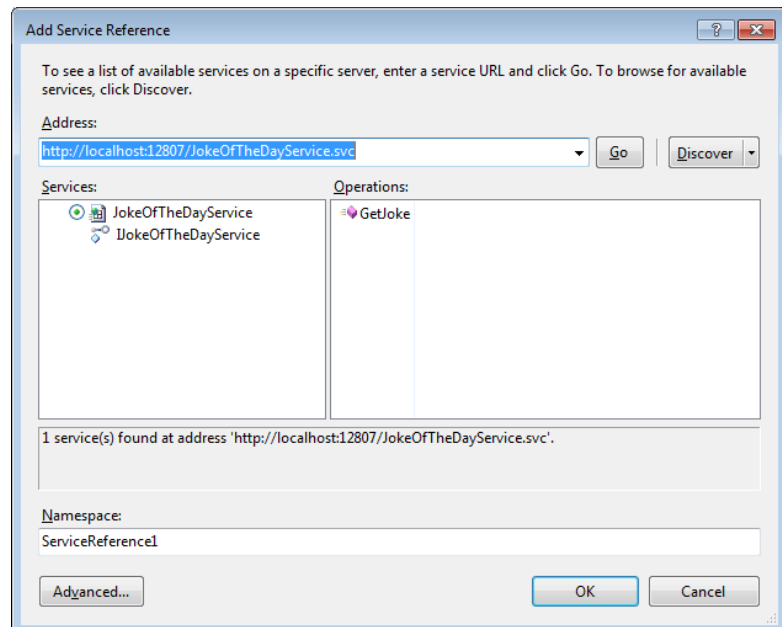
Joke of the Day Client

The client application needs to have a connection to the JokeOfTheDayService. This is added as a resource like any other, by using the Solution Explorer pane in Visual Studio:

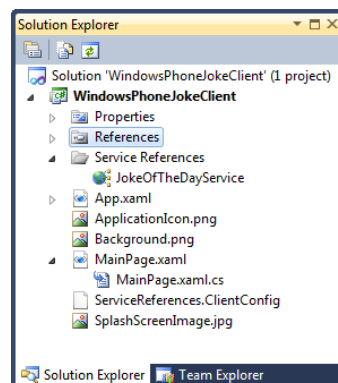


Rather than being a reference to a library assembly we instead start to create a Service reference. At this point Visual Studio needs to find the description of the service that is to be used. The Add Service Reference dialog lets us type in the network address of a service and it will then read the service description provided by the service.

It is very easy to create and test WCF services as these can be run on your Windows PC using a Development Server which is provided as part of Visual Studio. Once we have the service running and have obtained the url of the service on from the development server we just have to type this into the “Add Service Reference” dialog as shown below.



The Add Service Reference dialog will read the service description and show the operations available from that service. Above you can see that the JokeOfTheDay service only provides one method. At the bottom of this dialog you can see the namespace that we want the service to have in our Windows Phone client application. We can change this name to JokeOfTheDayService.



Once the service has been added it now takes its place in the solution. Our application must now create a proxy object that will be used to invoke the methods in the service. The best place to do this is in the constructor of the main page:

```
JokeOfTheDayService.JokeOfTheDayServiceClient jokeService;

// Constructor
public MainPage()
{
    InitializeComponent();

    jokeService = new JokeOfTheDayService.JokeOfTheDayServiceClient();

    jokeService.GetJokeCompleted +=
        new EventHandler<JokeOfTheDayService.GetJokeCompletedEventArgs>
            (jokeService_GetJokeCompleted);
}
```

This code creates a service instance and binds a method to the “completed” event that is fired when a service call returns. You might remember that when we got content from a web server the action was performed asynchronously, in that the program sent off a request to read the service and then another method was called when the data had

returned. Calls to WCF methods are exactly the same. When the program wants a new joke it must call the `GetJoke` method. This will not return with the joke, instead it will start off the process of fetching a joke. At some point in the future, when the joke arrives back from the server, a method will be called to indicate this. All the method needs to do is display the joke on the screen:

```
void jokeService_GetJokeCompleted(object sender,
    JokeOfTheDayService.GetJokeCompletedEventArgs e)
{
    if (!e.Cancelled)
    {
        jokeTextBlock.Text = e.Result;
    }
}
```

This method just checks to see if the call was successful and if it was the method takes the result from the call and puts it on the screen.

The final piece of code that we need to add is the button event handler that will actually call the WCF method to ask the server to do something:

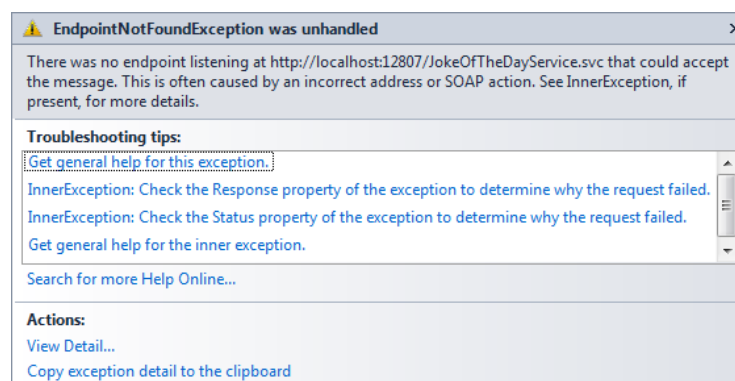
```
private void getJokeButton_Click(object sender, RoutedEventArgs e)
{
    int strength = 0;

    if (int.TryParse(strengthTextBox.Text, out strength))
    {
        jokeService.GetJokeAsync(strength);
    }
}
```

This method gets the strength value and then uses this in a call of the `GetJokeAsync` method. Note that `TryParse` will return false if the text in `strengthTextBox` could not be converted into an integer.

Errors and TimeOuts

If the service cannot be accessed the call to `GetJokeAsync` will throw an exception:



We know that the connectivity of a Windows Phone is not guaranteed. If they do not have a network signal when they run this program then it will throw the above exception. In production code you must make sure that these exceptions never get passed through to the user. This means that you need to enclose all your attempts to use these facilities in `try – catch` constructions and provide appropriate feedback to the user when these events happen.

Using Network Services

For the purpose of the demonstration above I used a network server that runs on the same machine as the client. In reality you will be connecting to different systems on the internet which may be a long way away. This is not a problem. As long as Visual

Studio can locate the service and download the service information it can then build the correct proxy objects to use the service correctly. If you want to test a service you can connect to a local service initially and then re-configure the service uri later to connect to the actual service.

The solution in *Demo 08a JokeService* contains a server which provides the joke service as described above. You can run this service on your machine and use the solution in *Demo 08b WindowsPhoneJokeClient* to connect to this service and read jokes from it. When you run the program you need to make sure that the development server url assigned to the service matches that of the service in the client.

We can now create systems that use Windows Phones as clients to servers. A given Windows Phone client can connect to multiple remote services at the same time if required. This is just an introduction to the things you can do with network services. Later we will look at some of the Windows Phone controls that build on these abilities.

What We Have Learned

1. Computer networks provide a means by which one system can send a message to another. A given system is physically connected to its *local* network. Each local network also contains a *router* component which will send packets off the local network to remote destinations. The routers are connected by data transfer devices which provide *inter-network* connections.
2. The full address of a system on a network contains the address of that station on its physical network along with the address of their physical network on the internet.
3. Data can be transferred between systems on the basis of unacknowledged datagrams (UDP) or as part of a connection set up between two systems (TCP).
4. Individual services on a system can be addressed by the use of ports, some of which have “well known” address values, for example an HTTP web server will typically be located behind port 80 on a host.
5. Windows Phone applications can create and use objects that represent a connection to a network service. The use of these objects is *asynchronous*, in that the results of a request are not returned to that request, but delivered later by a further call made by the network client.
6. A socket provides an abstraction of a connection to a network.
7. Some network calls may return structured data that contains XML describing the content.
8. The LINQ (Language INtegrated Query) library provided for use on the phone can create a representation of structured data which can then be traversed to create structured data for use in a program.
9. LINQ can be used to provide an automatic transfer of data from the rows and tables in a database into objects that can be used in a program.
10. XAML allows the creation of display templates which can be bound to data elements. These templates can include lists, so that a collection of structured data can be bound to a list.
11. Windows Communication Foundation (WCF) provides a means by which servers can be set up that expose methods for clients to call.
12. A client application can read the service description provided by a WCF server and use this to create a *proxy object* that contains the methods provided by the service. When code in the client calls the method in the proxy object a network transaction is performed that transfers the parameters of the call into the server. The server method is then invoked and a further network transaction delivers the result back to the caller.

8 XNA on Windows Phone

After all the hard work of the previous sections, now would seem a good place to have some fun and play some games. You can write games in Silverlight, but that is not really what it was designed for. XNA on the other hand was built from the ground up to be an efficient and powerful tool for game creation. In this section we are going to take a look at XNA and how to use it to create games for the Windows Phone device.

8.1 XNA in context

XNA is for creating games. It provides a complete ecosystem for game creation, including Content Management (how you get your sound effects, maps and textures into a game) and the game build process (how you combine all the elements into the single distributable game element). We are not going to delve too deeply into all these aspects of the system, instead we are going to focus on the XNA Framework, which is a library of C# objects that is used to create the games programs themselves. XNA allows game developers to use C# to create high performance games on a variety of platforms, including Windows PC, Xbox 360 and Windows Phone. By careful design of your program structure it is possible to reuse the base majority of your game code across these three platforms.


2D and 3D Games

Games can be “2D” (flat images that are drawn in a single plane) or “3D” (a visual simulation of a 3D environment). XNA provides support for both kinds of games and the Windows Phone hardware acceleration makes it capable of displaying realistic 3D worlds. For the purpose of this section we are going to focus on 2D, sprite based, games however. We can create a good gameplay experience this way, particularly on a mobile device.

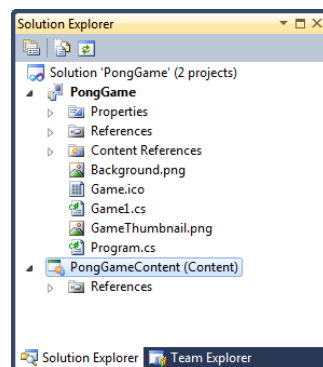
XNA and Silverlight

When we create a Windows Phone project we can create either a Silverlight application or an XNA one. Visual Studio uses the appropriate template and builds the solution with the required elements to make that kind of program. It is not possible to combine these. It would be lovely to think of a system which allowed us to use Silverlight to make the game menus and XNA to provide the gameplay but at the present this is not possible.

8.2 Making an XNA program




The New Project dialog in Visual Studio lets us select Windows Phone Game as the project type. If we want to make an XNA game for Windows PC or Xbox 360 we can do that here as well. One thing worth bearing in mind is that it is possible to use an existing XNA project as the basis of one for another platform, inside the same solution. Visual Studio will copy a project and make a new version which targets the new device. This means that we can take a Windows Phone XNA game and make a version for the Xbox 360 if we wish.



The Solution for an XNA project looks rather like a Silverlight one, but there are some differences. There is a second project in the solution which is specifically for game content. This is managed by the Content Manager which provides a set of input filters for different types of resources, for example PNG images, WAV sound files etc. These resources are stored within the project and deployed as part of the game onto the target platform. When the game runs the Content Manager loads these resources for use in the game. The result of this is that whatever the platform you are targeting the way that content assets are managed is the same as far as you are concerned.

One thing you need to be aware of though is that Windows Phone has a reduced display size compared to the Windows PC and Xbox 360 platforms. The maximum resolution of a Windows Phone game is 800x480 pixels. To make your XNA games load more quickly and take up less space it is worth resizing your game assets for the Windows Phone platform.

If you run the new game project as created above you will be rewarded by a blue screen. This is not the harbinger of doom that it used to be, but simply means that the initial behaviour of an XNA game is to draw the screen blue.



How an XNA Game Runs

If you look at a Silverlight application you will find that a lot of the time the program never seems to be doing anything. Actions are only carried out in response to events, for example when a user presses a button, or a network transaction completes, or a new display page is loaded.

XNA programs are quite different. An XNA program is continuously active, updating the game world and drawing the display. Rather than waiting for an input from a device, an XNA game will check devices that may have input and use these to update the game model and then redraw the screen as quickly as possible. In a game context this makes a lot of sense. In a driving game your car will continue moving whether you steer it or not (although this might not end well). The design of XNA recognises that the game world will update all the time around the player, rather than the player initiating actions.

When an XNA game runs it actually does three things:

1. Load all the content required by the game. This includes all the sounds, textures and models that are needed to create the game environment.
2. Repeatedly run the Game Engine
 - Update the game world. Read controller inputs, update the state and position of game elements.
 - Draw the game world. Take the information in the game world and use this to render a display of some kind, which may be 2D or 3D depending on the game type.

These three behaviours are mapped onto three methods in the `Game` class that is created by Visual Studio as part of a new game project.

```
partial class Game1 : Microsoft.Xna.Framework.Game
{
    protected override void LoadContent
        (bool loadAllContent)
    {
    }
    protected override void Update(GameTime gameTime)
    {
    }
    protected override void Draw(GameTime gameTime)
    {
    }
}
```

When we write a game we just have to fill in the content of these methods. Note that we never actually call these methods ourselves, they are called by the XNA framework when the game runs. `LoadContent` is called when the game starts. The `Draw` method is called as frequently as possible and `Update` is called thirty times a second. Note that this is not the same as XNA games on the Windows PC or Xbox 360. In order to reduce power consumption a Windows Phone game only updates 30 times a second, rather than the 60 times a second of the desktop or console version of the game. If you


want to write a game which will work on all three platforms you will have to bear this in mind.

Before we can start filling in the `Update` and `Draw` methods in our game we need to have something to display on the screen. We are going to start by drawing a white ball. Later we will start to make it bounce around the screen and investigate how we can use this to create a simple bat and ball game.

Game Content

We use the word *content* to refer to all the assets that make a game interesting. This includes all the images on the textures in a game, the sound effects and 3D models. The XNA framework Content Management system can take an item of content from its original source file (perhaps a PNG image) all the way into the memory of the game running on the target device. This is often called a *content pipeline* in that raw resources go into one end and they are then processed appropriately and finally end up in the game itself.

The content management is integrated into Visual Studio. Items of content are managed in the same way as files of program code. We can add them to a project and then browse them and manage their properties.



Above you can see an item of content which has been added to a Content project. The item is a PNG (Portable Network Graphics) image file that contains a White Dot which we could use as a ball in our game. This will have the asset name “WhiteDot” within our XNA game. Once we have our asset as part of the content of the game we can use it in games.

We have seen something that looks a bit like this in the past when we added resources to Silverlight games. Those resources have actually been part of Visual Studio solutions as well. However, it is important that you remember that the Content Manager is specifically for XNA games and is how you should manage XNA game content.

Loading Content

The `LoadContent` method is called to load the content into the game. (the clue is in the name). It is called once when the game starts running.

```
protected override void LoadContent()
{
    // Create a new SpriteBatch, which can be used to
    // draw textures.
    spriteBatch = new SpriteBatch(GraphicsDevice);
}
```

The first thing that `LoadContent` does is actually nothing to do with content at all. It creates a new `SpriteBatch` instance which will be used to draw items on the screen. We will see what the `SpriteBatch` is used for later.

We can add a line to this method that loads image into our game:

```
ballTexture = Content.Load<Texture2D>("WhiteDot");
```

The `Load` method provided by the content manager is given the type of resource to be loaded (in this case `Texture2D`). The appropriate loader method is called which loads the texture into the game.

```
Texture2D ballTexture;
```

XNA provides a type called `Texture2D` which can hold a single two-dimensional texture in a game. Later in the program we will draw this on the screen. A game may contain many textures. At the moment we are just going to use one.

Creating XNA Sprites

In computer gaming terms a sprite is an image that can be positioned and drawn on the display. In XNA terms we can make a sprite from a texture (which gives the picture to be drawn) and a rectangle (which determines the position on the screen). We already have the texture, now we need to create the position information.

XNA uses a coordinate system that puts the origin (0,0) of any drawing operations in the top left hand corner of the screen. This is not same as a conventional graph, where we would expect the origin to be the bottom left hand corner of the graph. In other words, in an XNA game if you increase the value of `Y` for an object this causes the object to move down the screen.

The units used equate to pixels on the screen itself. Most Windows Phones have a screen resolution of 800 x 480 pixels. If we try to draw objects outside this area XNA will not complain, but it won't draw anything either.

XNA provides a `Rectangle` structure to define the position and size of an area on the screen.

```
Rectangle ballRectangle = new Rectangle(
    0, 0,
    ballTexture.Width, ballTexture.Height);
```

The above code creates a rectangle which is positioned at the top left hand corner of the screen. It is the same width and height as the ball texture.

Drawing Objects

Now that we have a sprite the next thing we can do is make the game draw this on the screen. To do this we have to fill in the `Draw` method.

```
protected override void Draw(GameTime gameTime)
{
    GraphicsDevice.Clear(Color.CornflowerBlue);

    base.Draw(gameTime);
}
```

This is the “empty” `Draw` method that is provided by Visual Studio when it creates a new game project. It just clears the screen to blue. We are going to add some code to the method that will draw our ball on the screen. Before we can do this we have to understand a little about how graphics hardware works.

Modern devices have specialised graphics hardware that does all the drawing of the screen. A Windows Phone has a Graphics Processor Unit (GPU) which is given the textures to be drawn, along with their positions, and then renders these for the game player to see. When a program wants to draw something it must send a batch of

drawing instructions to the GPU. For maximum efficiency it is best if the draw requests are be batched together so that they can be sent to the GPU in a single transaction. We do this by using the `SpriteBatch` class to do the batching for us.

```
protected override void Draw(GameTime gameTime)
{
    GraphicsDevice.Clear(Color.CornflowerBlue);

    spriteBatch.Begin();

    // Draw operations go here

    spriteBatch.End();

    base.Draw(gameTime);
}
```

When we come to write our game the only thing we have to remember is that our code must start and end a batch of drawing operations, otherwise the drawing will not work.

We now have our texture to draw and we know how the draw process will work, the next thing to do is position the draw operation on the screen in the correct place. Once we have done this we will have created a *sprite*. The `SpriteBatch` command that we use to draw on the screen is the `Draw` method:

```
spriteBatch.Draw(ballTexture, ballRectangle, Color.White);
```

The `Draw` method is given three parameters; the texture to draw, a `Rectangle` value that gives the draw position and the color of the “light” to shine on the texture when it is drawn. The complete `Draw` method looks like this

```
protected override void Draw(GameTime gameTime)
{
    GraphicsDevice.Clear(Color.CornflowerBlue);

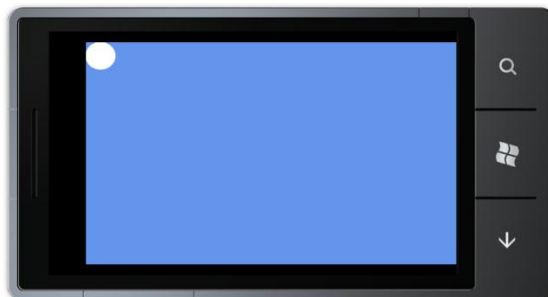
    spriteBatch.Begin();

    spriteBatch.Draw(ballTexture, ballRectangle, Color.White);

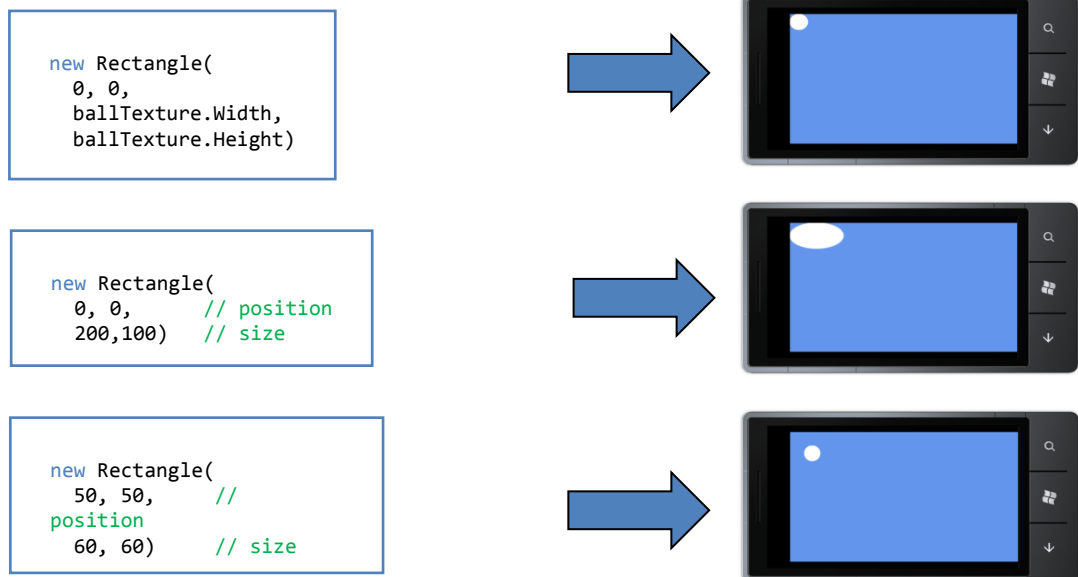
    spriteBatch.End();

    base.Draw(gameTime);
}
```

When we run this program we find that the ball is drawn in the top left hand corner of the screen on the phone:



Note that, by default, XNA games expect the player to hold the phone horizontally, in “landscape” mode, with the screen towards the left. We will see later that your game can change this arrangement if it needs to.



The first version of our program drew the ball with the same size as the original texture file. However, we can use whatever size and position we like for the drawing process and XNA will scale and move the drawing appropriately, as you can see above.

The solution in *Demo 01White Dot* contains a Windows Phone XNA game that just draws the white dot in the top left hand corner of the display.

Screen Sizes and Scaling

When you make a game it is important that it looks the same whenever it is played, irrespective of the size of the screen on the device used to play it. Windows Phone devices have a particular set of resolutions available to them, as do Windows PC and Xbox 360s. If we want to make a game that looks the same on each device the game must be able to determine the dimensions of the screen in use and then scale the display items appropriately.

An XNA game can obtain the size of the screen from the properties of the viewport used by the graphics adapter.

```

ballRectangle = new Rectangle(
    0, 0,
    GraphicsDevice.Viewport.Width / 20,
    GraphicsDevice.Viewport.Width / 20);

```

This code creates a ball rectangle that will be a 20th of the width of the display, irrespective of the size of the screen that is available.

Updating Gameplay

At the moment the game draws the ball in the same place every time. A game gives movement to the objects in it by changing their position each time they are drawn. The position of objects can be updated in the aptly named `Update` method. This is called thirty times a second when the game is running. We could add some code to update the position of the ball:

```
protected override void Update(GameTime gameTime)
{
    ballRectangle.X++;
    ballRectangle.Y++;

    base.Update(gameTime);
}
```

This version of `Update` just makes the `X` and `Y` positions of the ball rectangle one pixel bigger each time it is called. This causes the ball to move down (remember that the `Y` origin is the top of the screen) and across the screen at a rate of one pixel every thirtieth of a second. After fifteen seconds or so the ball has left the screen and will continue moving (although not visible) until we get bored and stop the game from running.

The solution in *Demo 02Moving Dot* contains a Windows Phone XNA game that draws a white dot that slowly moves off the display.

Using floating point positions

The code above moves the ball one pixel each time `Update` is called. This doesn't give a game very precise control over the movement of objects. We might want a way to move the ball very slowly. We do this by creating floating point variables that will hold the ball position:

```
float ballX;
float ballY;
```

We can update these very precisely and then convert them into pixel coordinates when we position the draw rectangle:

```
ballRectangle.X = (int)(ballX + 0.5f);
ballRectangle.Y = (int)(ballY + 0.5f);
```

The above statements take the floating point values and convert them into the nearest integers, rounding up if required. We can also use floating point values for the speed of the ball:

```
float ballXSpeed = 3;
float ballYSpeed = 3;
```

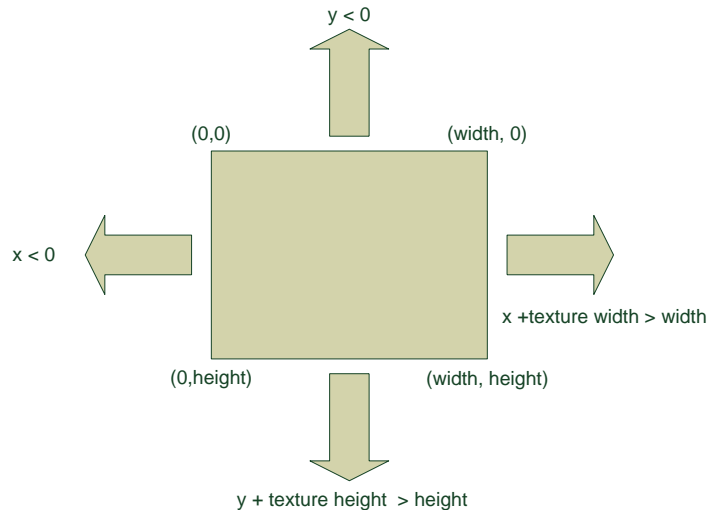
Each time the ball position is updated we now apply the speed values to the ball position:

```
ballX = ballX + ballXSpeed;
ballY = ballY + ballYSpeed;
```

Now that we can position the ball and control the speed very precisely we can think about making it bounce off the edge of the screen.

Making the ball bounce

Our first game just made the ball fly off the screen. If we are creating some kind of bat and ball game we need to make the ball "bounce" when it reaches the screen edges. To do this the game must detect when the ball reaches the edge of the display and update the direction of movement appropriately. If the ball is going off the screen in a horizontal direction the game must reverse the `X` component of the ball speed. If the ball is going off the top or the bottom of the screen the game must reverse the `Y` component of the ball speed.



The figure above shows the directions the ball can move off the screen and the conditions that become true when the ball moves in that direction.

```
if (ballX < 0 ||
    ballX + ballRectangle.Width > GraphicsDevice.Viewport.Width)
{
    ballXSpeed = -ballXSpeed;
}
```

The code above reverses the direction of the ball in the X axis when the ball moves off either the left or right hand edge of the display. We can use a similar arrangement to deal with movement in the Y direction.

The solution in *Demo 03 Bouncing Ball* contains a Windows Phone XNA game that draws a ball that bounces around the display.

Adding Paddles to Make a Bat and Ball Game

We now have a ball that will happily bounce around the screen. The next thing that we need is something to hit the ball with.



The paddle is a rectangular texture which is loaded and drawn in exactly the same way as the ball. The finished game uses two paddles, one for the left hand player and one for the right. In the first version of the game we are going to control the left paddle and the computer will control the right hand one. This means that we have three sprites on the screen. Each will need a texture variable to hold the image on the sprite and a rectangle value to hold the position of the sprite on the screen.

```

// Game World
Texture2D ballTexture;
Rectangle ballRectangle;
float ballX;
float ballY;
float ballXSpeed = 3;
float ballYSpeed = 3;

Texture2D lPaddleTexture;
Rectangle lPaddleRectangle;
float lPaddleSpeed = 4;
float lPaddleY;

Texture2D rPaddleTexture;
Rectangle rPaddleRectangle;
float rPaddleY;

// Distance of paddles from screen edge
int margin;

```

These variables represent our “Game World”. The `Update` method will update the variables and the `Draw` method will use their values to draw the paddles and ball on the screen in the correct position.

When the game starts running the `LoadContent` will fetch the images using the Content Manager and then set up the draw rectangles for the items on the screen.

```

protected override void LoadContent()
{
    // Create a new SpriteBatch, which is used to draw textures.
    spriteBatch = new SpriteBatch(GraphicsDevice);

    ballTexture = Content.Load<Texture2D>("ball");
    lPaddleTexture = Content.Load<Texture2D>("lpaddle");
    rPaddleTexture = Content.Load<Texture2D>("rpaddle");

    ballRectangle = new Rectangle(
        GraphicsDevice.Viewport.Width/2,
        GraphicsDevice.Viewport.Height/2,
        GraphicsDevice.Viewport.Width / 20,
        GraphicsDevice.Viewport.Width / 20);

    margin = GraphicsDevice.Viewport.Width / 20;


    lPaddleRectangle = new Rectangle(
        margin, 0,
        GraphicsDevice.Viewport.Width / 20,
        GraphicsDevice.Viewport.Height / 5);

    rPaddleRectangle = new Rectangle(
        GraphicsDevice.Viewport.Width -
            lPaddleRectangle.Width - margin, 0,
        GraphicsDevice.Viewport.Width / 20,
        GraphicsDevice.Viewport.Height / 5);

    lPaddleY = lPaddleRectangle.Y;
    rPaddleY = rPaddleRectangle.Y;
}

```

This `LoadContent` method repays careful study. It loads all the textures and then positions the paddle draw positions each side of the screen as shown below.



When the game starts the ball will move and the paddles will be controlled by the player.

Controlling Paddles Using the Touch Screen

We can use the Windows Phone touch screen to control the movement of a paddle. We can actually get very precise control of items on the screen, but we are going to start very simple. If the player touches the top half of the screen their paddle will move up. If they touch the bottom half of the screen their paddle will move down.

The Touch Panel returns a collection of touch locations for our program to use:

```
TouchCollection touches = TouchPanel.GetState();
```

Each TouchLocation value contains a number of properties that tell the game information about the location, including the location on the screen where the touch took place. We can use the Y component of this location to control the movement of a paddle.

```
if (touches.Count > 0)
{
    if (touches[0].Position.Y > GraphicsDevice.Viewport.Height / 2)
    {
        lPaddleY = lPaddleY + lPaddleSpeed;
    }
    else
    {
        lPaddleY = lPaddleY - lPaddleSpeed;
    }
}
```

The first statement in this block of code checks to see if there are any touch location values available. If there are it gets the location of the first touch location and checks to see if it is above or below the mid-point of the screen. It then updates the position of the paddle accordingly. If we run this program we find that we can control the position of the left hand paddle.

The solution in *Demo 04 Paddle Control* contains a Windows Phone XNA game that draws a ball that bounces around the display. You can also control the movement of the left hand paddle by touching the top or bottom of the screen area.

You can use the touch panel for much more advanced purposes than the example above. You can track the location and movement of a particular touch event very easily. You can also ask the touch panel to detect gestures that you are interested in.

Detecting Collisions

At the moment the bat and the ball are not “aware” of each other. The ball will pass straight through the bat rather than bounce off it. What we need is a way to determine when two rectangles intersect.

The `Rectangle` structure provides a method called `Intersects` that can do this for us:

```
if (ballRectangle.Intersects(1PaddleRectangle))
{
    ballXSpeed = -ballXSpeed;
}
```

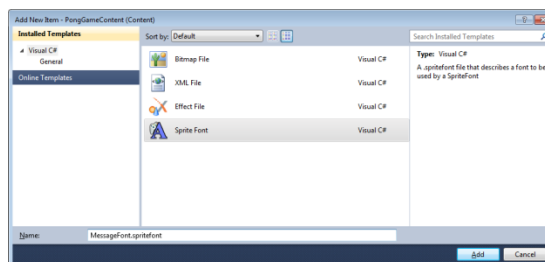
The above code tests to see if the ball has hit the left hand paddle. If it has the direction of movement of the ball is reversed to make the ball bounce off the paddle.

Keeping score

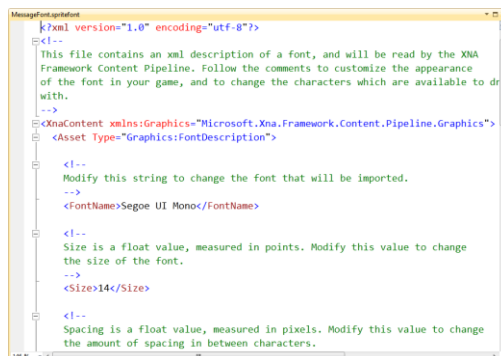
For the game to be worth playing it must also keep a score and display this for the players. Points are scored by hitting the back wall of the opponent player. We already have the code that does this. At the moment it just reverses the direction of movement of the ball, it is a simple matter to add a score variable for each player and increase this each time a “goal” is scored. The final thing we need to be able to do at this point is display the score for the players to see. To do this we need to find out how an XNA game draws text for a player.

Displaying Text

The XNA environment does not provide direct access to text drawing. To write text on the screen a game must load a character font and then use this to draw the characters. The font to be used is an item of game content which we need to add to the game. When the font is added the size and the font design is selected. When the game is built the Content Manager will build a set of character images and add these to the game. The game can then draw the text on the game screen along with the other game items.



The font, called a “SpriteFont”, is created as a new item and given a name. The actual font information is held in an XML file which you can edit with Visual Studio:



This file is where the font size and the name of the font are selected. Above you can see that the font “Segoe UI Mono” is being used with a font size of 14. Further down the XML file you will find settings for bold and italic versions of the font.

Once the `SpriteFont` has been added to the content in a program it can then be loaded when the program runs:

```
SpriteFont font;
```

```
font = Content.Load<SpriteFont>("MessageFont");
```

Note that this is exactly the same format as the `Load` we used to fetch the textures in our game only this time it is being used to fetch a `SpriteFont` rather than a `Texture2D`.

Once we have the font we can use it to draw text:

```
spriteBatch.DrawString( font, "Hello", new Vector2(50,100),
    Color.White);
```


The `DrawString` method is given four parameters. These are the font to use, the string to display, a vector to position the text on the screen and the colour of text to draw. A vector is a way that we can specify a position on the screen, in the example code above the text would be drawn 50 pixels across the screen and 100 pixels down (remember that the origin for Y is at the top of the screen).

The solution in *Demo 05 Complete Pong Game* contains a Windows Phone XNA game that implements a working pong game. The left hand paddle is controlled by the touch panel. The right hand paddle is controlled by an ultra-intelligent AI system that makes it completely unbeatable. Take a look at the code to discover how this top secret technology works. This game is not perfect. Because of the way the ball movement is managed it can sometimes get “stuck” on a paddle or off the edge of the screen. It is up to you to make a completed version of this.

8.3 Player interaction in games

The touch panel is not the only novel input device provided by the Windows Phone. Games can also make use of the accelerometer so that we can create games that are controlled by the tipping of the phone. The accelerometer can measure acceleration in three axes. While it can be used to measure acceleration (you could use your Windows Phone to measure the acceleration of your sports car if you like) it is most often used to measure the orientation of the phone. This is because the accelerometer is acted on by gravity, which gives a constant acceleration of 1 towards the centre of the earth.

You can visualize the accelerometer as a weight on the end of a spring, attached to the back of the phone. The picture below shows us how this might look.



If we hold the phone flat as shown, the weight will hang straight down underneath the phone. If we were to measure the distance in the X, Y and Z directions of the weight relative to the point where it is attached to the phone it would read 0, 0, -1, assuming

that the spring is length 1. The value of Z is -1 because at the moment the position of the weight is below the phone and the coordinate system being used has Z, the third dimension, increasing as we move up from the display.

If you tip the bottom of the phone up so that the far edge of the phone is now pointing towards your shoes the weight would swing away from you, increasing the value of Y that it has relative to the point where the string is attached. If you tip the bottom of the phone down, so that the phone tilts towards you and you can see the screen properly, the weight moves the other way, and the value of Y becomes less than 0. If the phone is vertical (looking a bit like a tombstone) the weight is directly below and in line with it. In this situation the value of Z will be 0, and the value of Y will be -1. Twisting the phone will make the weight move left or right, and will cause the value of X to change.

These are the values that we actually get from the accelerometer in the phone itself. So, at the moment the accelerometer seems to be measuring orientation (i.e. the way the phone is being held) not acceleration. If the phone starts to accelerate in other directions you can expect to see the values in the appropriate direction change as well. Our diagram above actually helps us visualise this too. If we push the phone away from us the weight will “lag” behind the phone for a second until it caught up resulting in a brief change in the Y value. We would see exactly the same effect in the readings from the accelerometer if we did move the phone in this way.

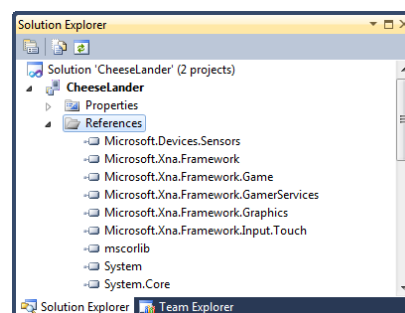
Getting Readings from the Accelerometer Class

When we used the touch panel from XNA we found that we just had to ask the panel for a list of active touch locations. It would be nice if we could do the same with the accelerometer but this is not how it works. The accelerometer driver has been written in a way that makes it useable in Silverlight programs as well. In Silverlight we were used to receiving messages from things when we wanted them to tell us something. The elements on a Silverlight page generate events when they are used and web request cause an event when they have data for us to use.

The accelerometer in Windows Phone works the same way. A program must make an instance of the `Accelerometer` class and connect a method to the `ReadingChanged` event that the class provides. Whenever the hardware has a new acceleration reading it will fire the event and deliver some new values for our program to use. If the program is written using Silverlight it can use those values directly and perhaps bind them to properties of elements on the screen. If the program is written using XNA it must make a local copy of the new values so that they used by code in the `Update` method next time it is called.

Using the Accelerometer Values

Before we can use the `Accelerometer` class we need to add the system library that contains it.



We can also add the appropriate namespace:

```
using Microsoft.Devices.Sensors;
```

Now the game can create an instance of the accelerometer, connect an event handler to it and then start the accelerometer running.


```
protected override void Initialize()
{
    Accelerometer acc = new Accelerometer();
    acc.ReadingChanged +=
        new EventHandler<AccelerometerReadingEventArgs>
            (acc_ReadingChanged);

    acc.Start();
    base.Initialize();
}
```

We can do this work in the `Initialize` method, as shown above. This is another method provided by XNA. It is called when a game starts running. It is a good place to put code to set up elements in our program. Next we have to create our event handler. This will run each time the accelerometer has a new value for the game.

```
Vector3 accelState = Vector3.Zero;
```

```
void acc_ReadingChanged
    (object sender, AccelerometerReadingEventArgs e)
{
    accelState.X = (float)e.X;
    accelState.Y = (float)e.Y;
    accelState.Z = (float)e.Z;
}
```

This method just copies the readings from the arguments to the method call into a `Vector3` value. The `Update` method in our game can then use these values to control the movement of the game paddle:

```
lPaddleY = lPaddleY - (accelState.X * lPaddleSpeed);
```

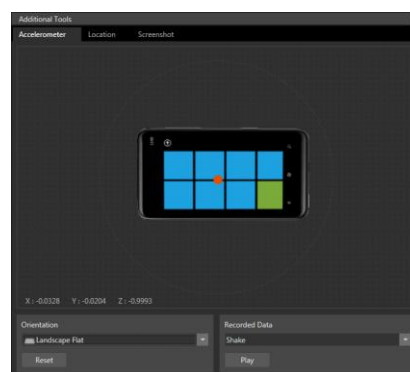
This code is very simple. It uses the `X` value of the accelerometer state to control the movement up and down of the left hand side paddle. You might think we should use the `Y` value from the accelerometer, but remember that our game is being played in *landscape* mode (with the phone held on its side) and the accelerometer values are always given as if the phone is held in *portrait* mode.

The accelerometer, in association with a very simple physics model, can be used to create “tipping” games, where the player has to guide a ball around a maze or past obstacles.

The solution in *Demo 06 Tipping Pong* contains a Windows Phone XNA game that implements a working pong game. The left hand paddle is controlled by tipping the phone to make the paddle move. You will notice that the further you tip the phone the faster the paddle moves, which is just how it should be.

Using the Accelerometer Emulation

If you start the program in the emulator you can use the Additional Tools menu to allow you to test the behaviour of the game by simulating the tipping of the emulated phone.



If you rest the mouse in the top left hand corner of the emulator a menu will appear and by clicking the >> on that menu you can open the Advanced Tools dialog. By selecting the Accelerometer tab and the Landscape Flat orientation you can get a 3D display of the phone which you can tip by dragging the red dot around. This allows you to get a feel of how the accelerometer input works. This is not a true guide to gameplay, but it does allow you to check that the game is using the correct axes.

Threads and Contention

The version of the accelerometer code above will work OK, but code like this can be vulnerable to a problem because of the way the program is written. In the case of this version of the game it would not drastically affect gameplay, but this issue is worth exploring because you may fall foul of similar problems when you start writing programs like these. Mistakes like these can give rise to the worst kind of programming bug imaginable, which is where the program works fine 99.999 per cent of the time but fails every now and then.

I hate bugs like these. I'm much happier when a program fails completely every time I run it because I can easily dive in and start looking for problems. If the program only fails once in a blue moon I have to wait around until I see the fault occur.

The problem has to do with the way that accelerometer readings are created and stored. We have two processes running working at the same time.

- The accelerometer is generating readings and storing them
- The Update method is using these readings to move the paddle

However, the Windows Phone only has one computer in it, which means that in reality only one of these processes can ever be active at any given time. The operating system in Windows Phone gives the illusion of multiple computer processors by switching rapidly between each active process. This is the cause of our problem. Consider the following sequence of events:

1. Update runs and reads the X value of the acceleration
2. The Accelerometer event fires and starts running. It generates and stores new values of X, Y and Z
3. Update reads the Y and Z values from the updated values
4. Update now has "scrambled" data made up of a mix of old and new readings.

In the case of our game we don't have much of a problem here, in that it only uses the X value of the accelerometer anyway. But if we had a more advanced game which used all three values the result would be that every now and then the Update method would be given information that wasn't correct. In an even larger and more complex program that used multiple processes this could cause huge problems.

Of course the problem will only happen every now and then, when the timing of the two events was very close together, but the longer the program runs the more chance there is of the problem arising.

The way to solve the problem is to recognise that there are some operations in a program that should not be interruptible. We need a way to stop the accelerometer process from being able to interrupt the Update process, and vice versa. The C# language provides a way of doing this. It is called a lock.

A given process can grab a particular lock object and start doing things. While that process has the lock object it is not possible for another process to grab that object. In the program we create a lock object which can be claimed by either the event handler or the code in Update that uses the accelerometer value:

```

object accelLock = new object();

void acc_ReadingChanged
    (object sender, AccelerometerReadingEventArgs e)
{
    lock (accelLock)
    {
        accelState.X = (float)e.X;
        accelState.Y = (float)e.Y;
        accelState.Z = (float)e.Z;
    }
}

```

The variable `accelLock` is of the simplest possible type, that of `object`. We are not actually storing any data in this object at all. Instead we are just using it as a token which can be held by one process or another. This is the new code in `Update` that uses the accelerometer reading and is also controlled by the same lock object

```

lock (accelLock)
{
    lPaddleY = lPaddleY - (accelState.X * lPaddleSpeed);
}

```

When a process tries to enter a block of code protected by a lock it will try to grab hold of the lock object. If the object is not available the process will be made to wait for the lock to be released. This means that it is now not possible for one method to interrupt another. What will happen instead is that the first process to arrive at the block will get the lock object and the second process will have to wait until the lock is released before it can continue.

Locks provide a way of making sure that processes do not end up fighting over data. It is important that any code protected by a lock will complete quickly, so that other processes do not have to spend a lot of time waiting for access to the lock.

It is also important that we avoid what is called the “Deadly Embrace” problem where process A has obtained lock X and is waiting for lock Y, and process B has obtained lock Y and is waiting for lock X. We can help prevent this by making a process either a “producer” which generates data or a “consumer” which uses data. If no processes are both consumers and producers it is unlikely that they will be stuck waiting for each other.

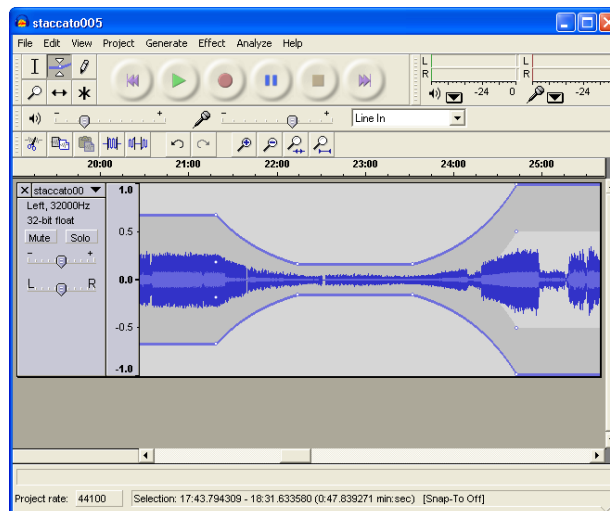
8.4 Adding sound to a game

We can make games much more interesting by adding sound effects to them. As far as XNA is concerned a sound in a game is just another form of game resource. There are two kinds of sounds in a game. There are sound effects which will accompany particular game events, for example the sound of a spaceship exploding when it is hit with a missile, and there is background music which plays underneath the gameplay.

Sound effects are loaded into the game as content. They start as WAV files and are stored in memory when the game runs. They are played instantly on request. Music files can also be supplied as game content but are played by the media player.

Creating Sounds

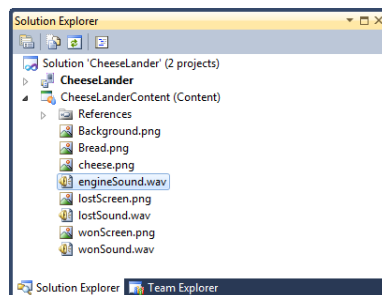
There are many programs that can be used to prepare sound files for inclusion in games. A good one is the program called Audacity. This provides sound capture and editing facilities. It can also convert sound files between different formats and re-sample sounds to reduce the size of sound files. Audacity is a free download from <http://audacity.sourceforge.net>



This shows a sound sample being edited in Audacity.

Sound Samples and Content

As far as an XNA game is concerned, a sound is just another item of content. It can be added alongside other content items and stored and managed by the content manager.



Within the game the `SoundEffect` class is used to hold the sound items:

```
// Sound effects
SoundEffect dingSound;
SoundEffect explodeSound;
```

These are loaded in the `LoadContent` method, as usual:

```
dingSound = Content.Load<SoundEffect>("ding");
explodeSound = Content.Load<SoundEffect>("explode");
```

A `SoundEffect` instance provides a `Play` method that triggers the playback of the sound:

```
if (ballRectangle.Intersects(rPaddleRectangle))
{
    ballXSpeed = -ballXSpeed;
    dingSound.Play();
}
```

This is the code that detects a collision of the bat with the ball. The ding sound effect is played when this happens. This is the simplest form of sound effect playback. The sound is made the instant the method is called. The program will continue running as the sound is played. Windows Phone can play multiple sound effects at the same time; the hardware supports up to 64 simultaneous sound channels.

Using the SoundEffectInstance Class

The `SoundEffect` class is very easy to use. It is a great way of playing sounds quickly. A game does not have to worry about the sound once it has started playing. However, sometimes a game needs to do more than just play the sound to completion. Some sound effects have to be made to play repeatedly, change in pitch or move to the left or the right. To get this extra control over a sound a game can create a `SoundEffectInstance` value from a particular sound effect. We can think of this as a handle onto a playing sound. For example we might want to create the sound of a continuously running engine. To do this we start with a `SoundEffect` which contains the sound sample.

```
SoundEffect engineSound;
```

This would be loaded from an item of content which contains the engine sound. Next we need to declare a `SoundEffectInstance` variable:

```
SoundEffectInstance engineInstance;
```

We can ask a `SoundEffect` to give us a playable instance of the sound:

```
engineInstance = engineSound.CreateInstance();
```

Now we can call methods on this to control the sound playback of the sound effect instance:

```
engineInstance.Play();
...
engineInstance.Pause();
...
engineInstance.Stop();
```

Note that the program can stop and start the playback of the sound effect instance, which is not something that is possible with a simple sound effect. A game can also control the volume, pitch and pan of a sound effect instance:

```
engineInstance.Volume = 1; // volume ranges from 0 to 1
engineInstance.Pitch = 0; // pitch ranges from -1 to +1
                        // -1 - octave lower
                        // +1 - octave higher
engineInstance.Pan = 0; // pan ranges from -1 to +1
                       // -1 - hard left
                       // +1 - hard right
```

The program can also make a sound loop, i.e. play repeatedly:

```
engineInstance.IsLooped = true;
```

Finally a program can test the state of the sound playback:

```
if (engineInstance.State == SoundState.Stopped)
{
    engineInstance.Play();
}
```

This would start the sound playing again if it was stopped. Using a `SoundEffectInstance` a game can make the pitch of the engine increase in frequency as the engine speeds up and even track the position of the car on the screen.

A `SoundEffectInstance` is implemented as a handle to a particular sound channel on the Windows Phone device itself. We have to be careful that a game doesn't create a very large number of these because this might cause the phone to run out of sound channels.

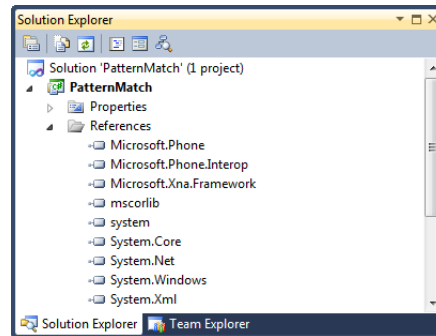
The solution in *Demo 07 Game with Sounds* contains a Windows Phone XNA game that implements a working pong game with sound effects.

8.5 Playing Sound in a Silverlight Program

A Silverlight program can use the XNA sound playback to provide sound effects. To do this it must include the XNA framework in the Silverlight program, load the sound from the project resources and then the sound can be played. Then, once the sound has started playing it must repeatedly update the XNA framework as the program runs so that sound playback is managed correctly.

Loading the XNA Namespace

The XNA namespace is not usually included in a Silverlight program but we can add it as a resource to a Silverlight project.



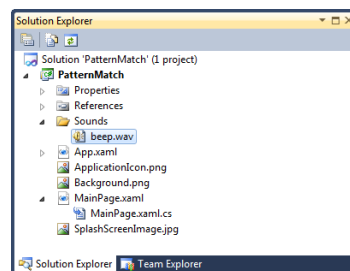
Once we have added the framework we can add using statements to make it easy to use the XNA classes in the Silverlight program:

```
using Microsoft.Xna.Framework;
using Microsoft.Xna.Framework.Audio;
```

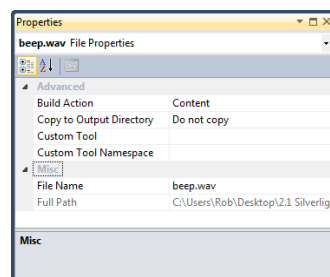
Note that we can't use many of the XNA elements, for example it would be impossible to use a `SpriteBatch` or a `Texture2D`, but we can use the audio classes.

Adding the Sound Resources

The sound resources are added to our Silverlight project as we would load any other.



Here you can see that I have created a folder called `Sounds` which is going to hold all the sound samples in the program. The resources themselves must be added as Content:



Now that the sound sample is part of the project, the next thing that the Silverlight program must do is load the sound data when the program starts running.

Loading a SoundEffect in a Silverlight program

Because there is no content management in a Silverlight program it must use a different way to load the `SoundEffect` value. A `SoundEffect` can be constructed using a method which loads the sound data from a stream. The following statement will do this for us:

```
SoundEffect beep;
...
beep = SoundEffect.FromStream(
    TitleContainer.OpenStream("Sounds/beep.wav"));
```

The `TitleContainer` class is part of the Silverlight solution, and can open any of the content items in the project as a stream. The result of this statement is that we now have a `SoundEffect` that we can play in a Silverlight program.

Playing the sound

A Silverlight program plays the sound in exactly the same way as an XNA program.

```
beep.Play();
```

A Silverlight program can also use `SoundEffectInstance` values to play more complicated sounds.

However, there is one important thing that a Silverlight system must do to make sure that sound playback works correctly. You will remember that in an XNA game the `Update` method is called 30 times a second. When an XNA game updates it also updates some of the system elements, including sound playback. A Silverlight program does not have this regular update behaviour, and so we need to add an update call to keep the framework running:

```
FrameworkDispatcher.Update();
```

This must be called around 30 times a second, otherwise sound playback may fail and the program itself may be stopped.

The best way to get this method called regularly is to create a timer to make the call for us. A timer is an object that will generate events at regular intervals. We can bind a method to the timer event and make the method call our dispatcher.

Our program can use a `DispatcherTimer`. This is part of the `System.Threading` namespace:

```
using System.Windows.Threading;
```

The code to create the timer and start it ticking is as follows:

```
DispatcherTimer timer = new DispatcherTimer();
timer.Tick += new EventHandler(timer_Tick);
timer.Interval = TimeSpan.FromTicks(333333);
timer.Start();
```

This code creates the timer, connects a method to the `Tick` event, sets the update rate to 30 Hz and then starts the timer running. The `timer_tick` method just has to update the XNA framework.

```
void timer_Tick(object sender, EventArgs e)
{
    FrameworkDispatcher.Update();
}
```

Timers are very useful things. A Silverlight program does have the regular update/draw method calls that an XNA game does, and so you can use a timer to allow you to create moving backgrounds or ticking clocks and the like.

The solution in *Demo 08 Silverlight with Sound* contains a Windows Phone Silverlight adding machine that plays a sound effect each time a new result is displayed.

8.6 Managing screen dimensions and orientation

We have already seen that a Windows Phone program can be used in more than one orientation. By default (i.e. unless we say otherwise) an XNA game is created that is played in landscape mode with the display towards the left. However, sometimes we want to play games in portrait mode. Our games can tell XNA the forms of orientation they can support and then receive events when the player tips the phone into a new orientation.

The graphics class in a game provides a number of properties that games can use to manage orientation and screen size:

```
graphics.SupportedOrientations = DisplayOrientation.Portrait |
    DisplayOrientation.LandscapeLeft |
    DisplayOrientation.LandscapeRight;
```

To allow a particular orientation we just have to add it to the values that are combined using the arithmetic OR operator, as shown above.

If we want our game to get control when the orientation of the screen changes we can bind a method to the event handler as shown below:

```
Window.OrientationChanged +=
    new EventHandler<EventArgs>(Window_OrientationChanged);
```

The `Window_OrientationChanged` method can then resize all the objects on the screen to reflect the new orientation of the game. We have already seen how a program can get the width and height of the display screen. The orientation handler method would use these values to set the new dimensions of the elements in the game.

```
void Window_OrientationChanged(object sender, EventArgs e)
{
    // resize the objects here
}
```

The required orientation is best set in the constructor for the game class.

Selecting a Screen Size

Games can also select a specific screen size:

```
graphics.PreferredBackBufferWidth = 480;
graphics.PreferredBackBufferHeight = 800;
```

If our game does this it forces the display to run at the requested resolution. It also forces the orientation too. The above width and height values would make the game work on portrait mode because the height of the game is greater than the width.

The clever thing about this is that when we set a particular display resolution the Windows Phone display hardware will make sure that we get that resolution, irrespective of the dimensions of the actual screen. In other words our game can operate as if the screen is that size and the hardware will make sure that it looks correct. This is a way you can make sure that your games will always look correct on the current version, and any new versions, of the Windows Phone hardware. No matter what new devices come along and whatever screen sizes they have the game will always look correct because the display will be scaled to fit the requested size.

We can use this to our advantage to improve the performance of our games:

```
graphics.PreferredBackBufferWidth = 240;
graphics.PreferredBackBufferHeight = 400;
```

The above statements ask for a screen which is smaller than the one fitted to most Windows Phone devices. In fact this screen is a quarter the size of the previous one. However, the hardware scaling in the Windows Phone will ensure that the screen looks correct on a larger display by performing scaling of the image. In a fast moving game

this is not usually noticeable and the fact that the display is only a quarter the size of the previous one will make a huge difference to the graphical performance.

Using the Full Screen

At the moment every XNA game that we have produced has been scaled slightly to fit a smaller screen. This is because by default an XNA game does not use the top bar of the phone display. This display area is set aside for status messages. If we want our game to use the entire screen we must explicitly request this:

```
graphics.IsFullScreen = true;
```

A game can set this property to true or false. If it is set to true this means that the game can use the whole of the Windows Phone screen. I always set this to true to give my games the maximum possible amount of display area.

Disabling the Screen Timeout

The owner of a Windows Phone can set a screen timeout for the device. The idea is that if the phone is left doing nothing for a while it will shut down the screen and lock itself to save battery life. The phone monitors the touch screen to detect user input and if there is no input for the prescribed time it will shut down. The phone does not check the accelerometer however, so if we make a game that is entirely controlled by tipping the phone our customers would become very upset when they played the game. They would just be getting to an interesting part and then find that their screen went blank.

A game can disable the screen timeout function by turning off the screen saver in the XNA Guide:

```
Guide.IsScreenSaverEnabled = false;
```

The XNA Guide is the part of XNA that has conversations with the player about their Xbox Live membership and achievements etc. If we tell it to stop the screen saver from appearing this means that the game can continue until the phone battery goes flat. You should use this option with care. If a user gets bored by your game and puts the phone down it will not shut down if the screensaver has been turned off. This may mean that the phone battery will go flat as the game has been left running.


If you do have a game that is controlled by the accelerometer I would also add some gameplay element where a player has to touch the screen every now and then to keep the phone awake. I would much rather have a game that worked in this way rather than one that ran the risk of flattening the phone battery.

The solution in *Demo 09 Full Screen Pong* plays pong on a lower resolution screen that covers the entire phone surface. It is noticeable that from a player's point of view it is hard to tell from the higher resolution version.

8.7 Combining XNA and Silverlight

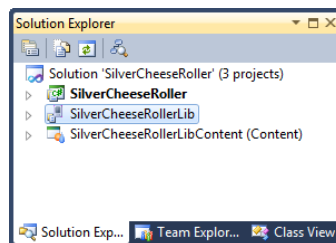
It is possible to create a single Visual Studio solution that contains a Silverlight application and also an XNA game component. This is very useful because Silverlight is very good for creating user interfaces but not very good at games, whereas XNA is very good for game creation but it is hard to create menus. We can use this to create high quality in game menus for games and we could also use this technique to add some 3D graphics to a Silverlight application.

The starting point for a combined application is the New Project dialog in Visual Studio. We can select the template that makes a combined project:

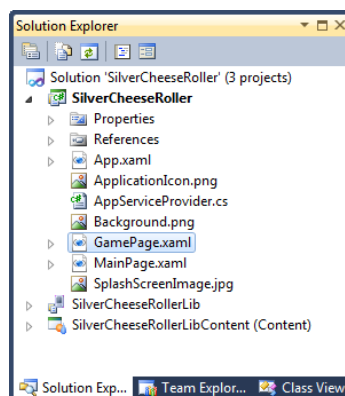


We can create combined applications from both the XNA and the Silverlight templates, in each case the application that is created is the same thing.

When we create the combined application we find that the solution now contains three projects.



The top project is the Silverlight application that provides the starting point. The second project contains a set of library routines that and the third is where the XNA content for the game part of the solution is stored.



When we create a combined application we also get a new Silverlight page that will contain the XNA game. The XNA gameplay is introduced as content on the GamePage. Later we will see how we can combine XNA and Silverlight controls on the same page, for now we will notice that if we look inside the GamePage.xaml page design there is a placeholder that indicates that there are no Silverlight components on this page.

Creating an XNA Game Environment in Silverlight

We have seen that when an XNA game runs a set of methods are called that are where the game initialises, loads its content and then repeatedly performs update (to move on the game world) and draw (to render the game world on the screen). When a game runs inside a Silverlight page it will need to have these behaviours provided for it.

```
protected override void OnNavigatedTo(NavigationEventArgs e)
protected override void OnNavigatedFrom(NavigationEventArgs e)
private void OnUpdate(object sender, GameTimerEventArgs e)
private void OnDraw(object sender, GameTimerEventArgs e)
```

These are the methods in the game page that we need to fill in. The `OnNavigatedTo` method takes the role of `Initialise` and `LoadContent` and the `OnUpdate` and `OnDraw` methods provide updating and drawing. The `OnNavigatedTo` method is always called when the user navigates to a page, and this is where a game should load content and get ready to run. This method also creates a timer that performs the calls of `OnUpdate` and `OnDraw` when the game is active.


Note that this is an XNA program running inside a Silverlight page, which means that it is possible the user may navigate away from the game and move to another page in the application. If this happens the `OnNavigatedFrom` method will be called. This could be used to make the game automatically pause when the user moves away from the screen.

In the template application the game will start when the user presses the

```
private void Button_Click(object sender, RoutedEventArgs e)
{
    NavigationService.Navigate(new Uri("/GamePage.xaml",
                                     UriKind.Relative));
}
```

When the button is pressed the `GamePage` is navigated to and the game begins.

The solution in *Demo 10 Silverlight Pong* contains a Silverlight application that contains a version of the Pong game.



If you press the button the game will start running. If you press the Back button on the phone during the game the Silverlight application will navigate back to the `MainPage` and the game will stop. Note that the game is restarted each time the page is navigated to. The present version of the game reloads all the textures and game assets each time the page is recreated. It might be more sensible to store the textures and game state in a separate class and cache the data so that the user can return to where they left off, rather than getting a new copy each time.

XNA and Silverlight together

It is possible to actually overlay Silverlight display elements on top of an XNA game. This makes it very easy to add buttons and text to the game by adding `Button` and `TextBlock` elements. To put some Silverlight elements on the game screen we first have to create the elements themselves and add them to `GamePage.xaml`:

```
<Grid x:Name="LayoutRoot">
  <StackPanel>
    <TextBlock x:Name="ScoreTextBlock" Text="0:0"
      TextAlignment="Center"
      Style="{StaticResource PhoneTextTitle1Style}" />
    <Button Content="Quit" Name="quitButton"
      Width="480" Click="quitButton_Click" />
  </StackPanel>
</Grid>
```

This creates a `StackPanel` containing a button and a textblock. The button will provide a quit behaviour and the textblock will display the game score.

Now that the page has some display elements, we need to get the page rendered on the display. The Silverlight element rendering is performed by a class with the fairly obvious name of `UIElementRenderer`.

```
UIElementRenderer elementRenderer;
```

We create an instance of the renderer when the page is loaded:

```
if (null == elementRenderer)
{
  elementRenderer = new UIElementRenderer(this,
    (int)ActualWidth, (int)ActualHeight);
}
```

This code checks to see if an element renderer exists and creates one if it is not present. Note that the constructor for this element is given a reference to the page it is rendering and the width and height of the draw area.

Once we have our renderer we can use it in the `OnDraw` method to render the Silverlight controls:

```
private void OnDraw(object sender, GameTimerEventArgs e)
{
  // Render the Silverlight controls using the UIElementRenderer.
  elementRenderer.Render();

  SharedGraphicsDeviceManager.Current.GraphicsDevice.Clear(
    Color.CornflowerBlue);

  spriteBatch.Begin();

  spriteBatch.Draw(ballTexture, ballRectangle, Color.White);

  ...


  spriteBatch.Draw(elementRenderer.Texture, Vector2.Zero,
    Color.White);

  spriteBatch.End();
}
```

The final drawing call in the above `OnDraw` method is where the Silverlight controls are rendered. Note that this means the controls will be drawn on top of the game elements. If we wanted to use the controls as the backdrop for the game we would draw the elements first.

When the XNA game runs the Silverlight elements respond to their inputs and we can connect event handlers to them. Note that this would be the case even if the elements were not drawn.

The solution in *Demo 11 Silverlight Elements* contains a Silverlight application that contains a version of the Pong game with Silverlight buttons and text on it.



If you press the Quit button on the game page it will return you to the previous page.

What We Have Learned

1. The XNA framework provides an environment for creating 2D and 3D games. The games can be created for Windows PC, Xbox 360 or Windows Phone.
2. XNA games operate in a completely different way from Silverlight applications. An XNA game contains methods to load content, update the game world and draw the game world which are called by the XNA framework when a game runs.
3. When Visual Studio creates a new XNA game it makes a class which contains empty LoadContent, Update and Draw methods that are filled in as a game is written. The LoadContent method should load all the required game content. The Update method is called by XNA 30 times a second to update the game content and the Draw method is called to render the game on the screen.
4. Any content (for example images or sounds) added to an XNA game is managed by the Content Management process that provides input filters for different file types and also output processors that prepare the content for deployment to the target device. Within an XNA game the way that content is loaded and used the same irrespective of the target hardware.
5. A sprite is made up of a texture which gives the image to be drawn and a position which determines where to draw the item and its size. The XNA Framework provides objects that can hold this information.
6. When drawing in 2D within the XNA system a given position on the screen is represented by pixel coordinates with the origin at the top left hand corner of the display.
7. The Windows Phone touch panel can provide a low level array of touch location information which can contain details of at least 4 touch events.
8. XNA games can display text by using spritefonts made from fonts on the host Windows PC
9. XNA and Silverlight programs can play sound effects. For greater control of sound effect playback a program can create a SoundEffectInstance value which is a handle to a playing sound.
10. XNA programs on Windows Phone can select a particular display resolution and orientation. If they select a resolution lower than that supported by the phone display the GPU will automatically scale the selected size to fit the screen of the device.

11. XNA programs on Windows Phone can disable the status bar at the top of the screen so that they can use the entire screen area. They can also disable the screensaver so that they are not timed out by the phone.
12. XNA and Silverlight can be combined in a single solution. The XNA game runs within a Silverlight page which can also contain Silverlight components.


9 Creating Windows Phone Applications

We now know enough to make applications that run on the Windows Phone device. In this section we will take a look at what takes a program and turns it into a “proper” application. This includes a variety of topics, from how to give your program a custom splash screen and icons to how an application can create and manage background tasks.


9.1 The Windows Phone Icons and Splash Screens

At the moment all the programs that we have written have had the same icon and “splash screen”. These have been the “empty” ones that Visual Studio created when it made a new project. There are two images that Windows Phone uses to identify your program when it is stored on the phone. To understand how these are used we have to learn a little about how Windows Phone users find and run programs on the device.

A Windows Phone has a “Start” screen which is displayed whenever the user presses the “Start” button on the device:




The user can then touch any of the large tiles on the start screen to run the program behind that tile. Alternatively they can slide the start screen to the left (or press the little arrow at the top right of the screen) to move to a list of applications that are installed on the device.



Then the user can scroll up and down the application list and select the program that they want to run. If the user wants to be able to start one of their applications from the Start screen they can “pin” it to the start screen.

Silverlight icons

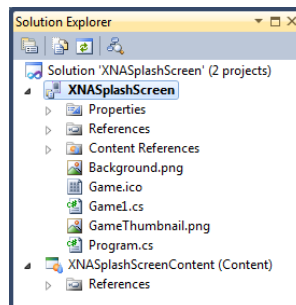
If you create a new Silverlight application Visual Studio will create a project that contains “default” versions of these two icons:



The file `ApplicationIcon.png` holds the image to be used in the application list. The file `Background.png` holds the image to be used when the application has been “pinned” to the Start menu.

XNA icons

If you create a new XNA application Visual Studio will create a slightly different arrangement of programs and the icon files are named differently too:



The file `GameThumbnail.png` holds the image to be used in the application list. The file `Background.png` holds the image to be used when the application has been “pinned” to the Start menu.

Customising Icons

If you want to customise the program icons (and you should) you can open and edit these files with any paint program that you like. Remember that they are stored in the project folder along with all the project files. However, you must make sure that you retain the file type and the image dimensions of the originals.

From my experience it turns out that designing icons like this (particularly small ones) is something of an art form. For this reason if you can find an artist to do the design for you this will probably be a good idea. If you have to do the work yourself I would advise you to make the design as simple as possible and make sure that it gives the user a good idea of what your program does. Remember that the icon for your program is also used to promote it in the Marketplace, so the better and slicker the icon looks the more chance there is of people downloading and buying your program.




If you use the Marketplace test kit to prepare your application for sale you can load all your artwork and icons into this. It will also check the size and format of each image.

Splash Screens

A “splash screen” is something that a program “splashes” onto the display as it starts up. These can be used to “brand” your application by showing your company logo when your program starts. Splash screens are also useful to improve the experience of program users when your program starts running. At the very least a splash screen will give them something nice to look at as your program loads.

Silverlight Splash Screens

The Silverlight run time system has a splash screen behaviour built in. A new Silverlight project actually contains an image file which is displayed during the interval between a program being started and `MainPage` being displayed on the screen. If you have run programs on the emulator you may have seen this for a brief instant at the very start of a program run.



If your programs become larger, with more elements on the page and more resources to be loaded, this page may be displayed for more time. You can make your own version by modifying the `SplashScreenImage.jpg` file in your Silverlight project.

XNA Splash Screens

XNA programs do not have splash screens provided as part of the game project. We have already seen that the way that the display is produced is quite different. When an XNA program starts running the `LoadContent` method is called to fetch assets from storage and make them available to the game. Once the content has been loaded the `Draw` and `Update` methods begin to be called. This can lead to problems. The XNA framework insists on a game displaying something on the screen within five seconds of the program starting. If this does not happen the framework assumes that the game has got stuck and aborts it. This would mean that your game never actually gets to run if the content loading takes too long.

Fortunately an XNA does not have to load all the content at the start. A game can use the content manager to load items at any time. A program can also fire off threads to perform background loading of items while the game runs. The upshot of all this is that for XNA you need to develop a strategy for loading content. Perhaps the game `LoadContent` method could just load a splashscreen texture and this could be displayed while the rest of the content items are fetched.

It turns out that this is a problem common to games on all platforms. A major limiting factor in the development of ever larger console games is the limited access speed and transfer rate of the optical drives used to load information from the storage media.

9.2 Fast Application Switching

The Windows Phone is a very powerful device. Computers of this power filled entire rooms just a few years ago. However, when compared with desktop systems it is somewhat constrained. It does not have a processor which is quite as powerful, and it is restricted by the limited storage capacity of the battery. While the phone can do a lot of processing, it is best if this is kept to a minimum to make battery life as long as possible.

Now we are going to find out how an application can be made to provide a good user experience within the single tasking environment provided on the Windows Phone. This provides a good insight into how programmers often have to make use of systems that are compromised by underlying limitations of the platform and operating system.

Task Navigation in Windows Phone

The Windows Phone operating system is inherently multi-tasking. This means that it can easily run more than one process at the same time. This is how it can play music using the Zune program while you use the phone to check your email. However, for a number of very sound technical reasons this multi-tasking ability is not extended to applications running on the system. The creators of the Windows Phone system took a deliberate design decision not to allow more than one application to be active at any given time.

In some ways this is a good thing. It stops our programs from looking bad just because another application running in the phone is stealing all the memory or processor time. But it does give us some programming challenges that we must address when we write Windows Phone applications.

When we write an application the aim should be to make it appear that the application never stops when a user leaves it and goes off to do something else. This “leaving and going off to do something else” is a behaviour that is strongly encouraged by the design of the Windows Phone user interface. We know how an application can store data using the mass storage on the phone, what we have to do is store and retrieve data to give the illusion that our program was never paused.

Later on we will see how we can create applications that do perform some background processing on behalf of the user. However, we will find that this ability is carefully managed to make sure that anything our program does in the background does not compromise the experience that the user gets.

The Start and Back buttons

While Windows Phone does not have multi-tasking it does have features which have been designed to make it easy for a user to move in and out of programs. These are based on the hardware **Back** and **Start** buttons that are fitted to every Windows phone device.

Put simply, **Start** moves the user forward to a new program that they may wish to run and **Back** takes the user back to a previous one. A user can press **Start** at any time to open the Start screen and select a different program to run. When they have finished using that different program they can use **Back** to return to the original one.

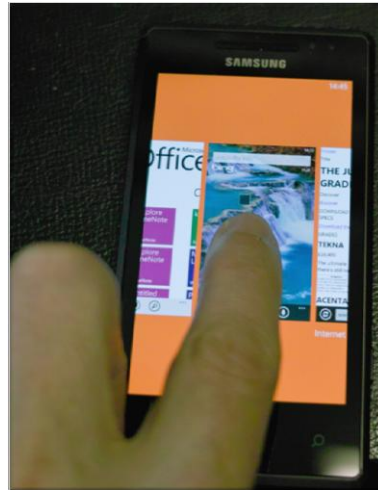
Used in combination these buttons greatly enhance the user experience. It is easy to press **Start** to send a quick SMS message and then press **Back** to return to whatever you were doing before. The **Back** and **Start** buttons allow the user to navigate through a stack of recently used applications in a similar way to browsing the internet, clicking on links and using the Back command to return to the previous page.

Pressing the **Start** button does not cause an application to be removed from memory to make way for a new one; instead it is “deactivated”. We can regard a deactivated application as “waiting in the wings”, ready to be called back onto the stage at some later time. Of course the call back to the stage may never come (just for real life

performers) and so when the application is deactivated it must store any important data just in case it the user never returns to it.

The Back Button Long Press

In addition to pressing and releasing the **Back** button a user of the phone can hold it down for half a second or so to activate a display of applications that are presently “waiting in the wings”. If the user selects one of these applications it is restarted exactly as if the user had returned to the application via the Back button.



This shows how waiting applications are selected. I am panning between the currently active applications (these include Office and the Bing search screen) to select the one that I want to go back to.

Creating a “Well Behaved” Application

The result of all this is that an application must be made adept at saving its state and then restoring itself to the same state when the user returns to it. The Windows Phone system helps us do this by sending an application messages that indicate what is going to happen to it. Next we are going to explore the situations in which the messages are sent and how we can make our applications and games respond correctly to give the appearance that they are always active.

Understanding Fast Application Switching

The name for this set of behaviours is “Fast Application Switching”. Because deactivated programs are usually held in memory the process of switching between them should be very quick. If we are going to understand how to make programs that behave properly we are going to have to learn some of the terminology used by the Windows Phone designers to describe what happens to make this work.

The Fast Application Switching Event Methods

The `App.xaml.cs` file in a Silverlight application contains methods that are called when the various fast application switching events occur:

```
// Code to execute when the application is launching (eg, from Start)
// This code will not execute when the application is reactivated
private void Application_Launching(object sender, LaunchingEventArgs e)
{
}

// Code to execute when the application is activated (brought to foreground)
// This code will not execute when the application is first launched
private void Application_Activated(object sender, ActivatedEventArgs e)
{
}

// Code to execute when the application is deactivated (sent to background)
// This code will not execute when the application is closing
private void Application_Deactivated(object sender, DeactivatedEventArgs e)
{
}

// Code to execute when the application is closing (eg, user hit Back)
// This code will not execute when the application is deactivated
private void Application_Closing(object sender, ClosingEventArgs e)
{
}
```

We first saw the file `App.xaml.cs` in section 4.5 *Pages and Navigation* when we used it as a place to create variables that can be shared by all the pages in a Silverlight application. It is the class that controls the creation and management of the Silverlight pages. Our job is to make our applications use these events to give the user the impression that our program is always ready for use. A similar set of messages are available to XNA programs so that games can be made to save their state.

The Windows Phone Application LifeCycle

Now that we know what event messages there are, and how the user moves from one application to another, we can consider how the messages work in the lifecycle of an application.

Program Launching

A user launches a program by selecting it from those available on the phone. When a program is launched the program files are loaded into memory and executed by the Windows Phone processor. During the launch process the `Application_Launching` method is called

```
private void Application_Launching(object sender,
                                   LaunchingEventArgs e)
{
}
```

This method should load any data from isolated storage. However, if there is a large amount of data to be read it might be best to load this in a separate thread (i.e. in the background). This is because the application will not display anything on the screen until this method has finished running. If the program spends a long time loading data this means your program will seem very slow to load.

There are two things that the user can do next. The user can exit the program by pressing the Back button at the top level menu of the program or the user can move to another program by pressing the Start button.

Pressing Back to Close an Application

Windows Phone applications do not have an “exit” option that can be used to close them. However, this does not mean that it is impossible to close an application. If the

user presses the Back button at the very “top level” of the application (i.e. the screen that is displayed when the application was started) the application will be closed at this point. When the application closes it is removed from memory and the previous application (if there is one) is resumed. This gives the behaviour where a user can “back” out of an application and return to the one they interrupted.

Just like a Windows PC application, a Windows Phone application uses the closing event so that it can save any open files.

```
private void Application_Closing(object sender, ClosingEventArgs e)
{
}
```

It is even possible to allow the user the chance to confirm that the program is being closed. The Microsoft Word application on the phone will display a “Do you want to save your text” message if you try to back out of it after you have entered some text.

Unfortunately for us users will not always close a program “cleanly” like this. They are much more likely to press the Start button and go off and do something else rather than close one application to run another. When we design a Windows Phone application we need to take care to make sure that we give proper consideration to how the program will behave when the user leaves it and returns to it. We must make sure that the application always does what the user expects, and does not suddenly discard changes or lose data.

Deactivating an Application by pressing Start

When the user presses Start the currently active program is deactivated. This means it is initially made “dormant”. Dormant programs are still in memory, but are not presently active. They are ready to be resumed using “Fast Application Switching”.

```
private void Application_Deactivated(object sender,
                                     DeactivatedEventArgs e)
{
}
```

At the point that an application is being made “dormant” it is sent a “deactivated” message to indicate that it is being “moved to the wings” to make way for another performer on the stage. This is the point at which the application must save the state of the current interaction and also save data in isolated storage in case it is not resumed. A program can also be made “dormant” if the phone lock screen is displayed because the phone has not been used and has timed out. A program may also be made dormant if the user presses the Search button, or starts using the camera to take a photograph.

Converting a Dormant Application into a Tombstoned one

A real theatre would have problems if a large number of performers were sent to the wings, the wings might get very crowded. If this happened some of the performers would be sent back to their dressing rooms. If these performers are ever needed again they will need to be called back, which means it will take longer for them to resume their performance.

Exactly the same thing happens in Windows Phone. As more and more applications are deactivated and made dormant the memory in the phone will eventually fill up. At this point some of the applications need to be removed from memory to make room for others. In Windows Phone terms they are converted from “dormant” (in memory and available for Fast Application Switching) to “tombstoned” (removed from memory). The user can still resume a tombstoned application if they press Back enough times, or select it from the “Long Press” menu but it will take slightly longer to reactivate.

If a tombstoned program is ever reactivated it will be loaded into memory and all the objects in the program will need to be constructed from scratch, just as they would be if the program was loaded for the first time.

At this point we might be forgiven for thinking that being tombstoned is just like being exited. However, this is not quite the case. The great thing about a tombstone is that we

can write things on it. The Windows Phone system provides some space in memory where an application can store state information if it is ever deactivated. This information can then be loaded back into the application if it is ever reactivated. The information is stored when the application is deactivated and maintained whether the application is made dormant or tombstoned.

The Windows Phone system also uses this state memory. If you deactivate an application at a particular page you will find that if you reactivate it you are returned to the same page. The Silverlight system uses the state storage in memory to record the active page when it is deactivated.

Resuming a Dormant or Tombstoned Application

When an application is being started it will receive either a “Launching” message if it is being launched as new or an “Activated” message if it is being reactivated.

The application can then test to see if it is being reactivated from being “dormant”, or if it is being reactivated from being “tombstoned”.

```
private void Application_Activated(object sender,
                                   ActivatedEventArgs e)
{
    if (e.IsApplicationInstancePreserved)
    {
        // Dormant - objects in memory intact
    }
    else
    {
        // Tombstoned - need to reload our objects
    }
}
```

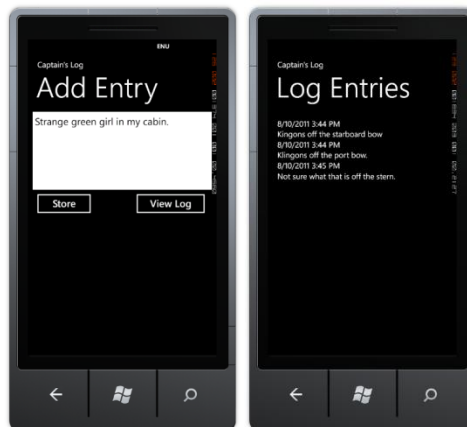
This method runs when the application is activated and tests which form of activation is being performed. When an application is restarted from the dormant state it will have all its memory objects intact and can just resume running. When an application is restarted from the tombstoned state it must reload all its data as it will have been removed from memory.

Note that an application does not get a message when it is moved from the dormant to the tombstoned state, it can only tell which state it was in if it is ever restarted.

Fast Application Switching in an application

We are now going to take a look at an application and consider how we would make it into one which works correctly with respect to Fast Application Switching. In Chapter 5 we created an application called “JotPad” that allowed the user to store text jottings on the phone. To investigate Fast Application Switching we are going to use an enhanced version of JotPad called “The Captain’s Log”.

An Introduction to “The Captain’s Log”



The Captain’s Log allows the user to store a sequence of entries in a log file. Each entry is timestamped when it is entered and then added to the log itself, which is a long string of text. When the program is started the Add Entry page is displayed. The user can then type in some text and press “Store” to add an entry or press “View Log” to see the log entries. If they press “View Log” the program navigates to the second page, which shows the accumulated log entries. Each time a log entry is added it is annotated with the current date and time, as shown above.

We want to make the Captains Log program work in the way that the user expects. It must give the appearance of being available at all times, so that the user is not aware of way that data is being stored and loaded as required. When the user returns to the application the data that they had previously entered should be present, leaving the impression that it has always been there.

The “Captain’s Log” and fast application switching

When we create an application we have to decide how it should behave when the user moves away from it. Should it store all the data in isolated storage each time, or should it only persist things if the user is closing the program? This decision has an impact on the user experience. The user might not expect information to be stored if they “jump out” of an application using the Start button. In fact the user might think they can press Start and re-launch the Captain’s Log application as a quick way of abandoning an unfinished log entry they have decided not to post.

Note that these are **not** technical questions. If you are making a Windows Phone program for a customer you must consult with the customer to find out what they want the program to do. The worst thing you can do is make an assumption about how it is supposed to work.

When a user presses Start to move from an application it might be best if the data is stored as a “work in progress” and not actually committed to main storage. That is how we are going to make the Captain’s Log program work.

Data in the “Captain’s Log” program

The Captain’s Log program only contains two data elements which are both strings. They are the complete log text which is displayed on the “Log Entries” page and any half-finished log entry which is displayed on the “Add Entry” page.

```
// This is shared amongst all the pages
// It is the contents of the log itself, as a large string
public string LogText;

// This is also shared, only used by the log entry page
public string LogEntry = "";
```

The best place to put these variables is in the `App` class in the `App.xaml.cs` file. This file is the “container” for the entire application and is a good place to put data that is shared by all the pages in a program. This is the class that contains the methods that are used to manage Fast Application Switching so these methods can save and load the program data, as we shall see later.

Page Navigation in the Captain's Log Program

This application has two pages which the user will navigate between. It is important that a page always shows the correct data. We know that Silverlight pages are stateless (i.e. they cannot be used to record any data) and so when a page is displayed it is important to make sure that the content on the page is correct. We can use the `OnNavigatedTo` method to do this:

```
protected override void OnNavigatedTo
    (System.Windows.Navigation.NavigationEventArgs e)
{
    // When we navigate to the page -
    //put the semi-completed log text in the display

    // Get a reference to the parent application
    App thisApp = App.Current as App;

    // Put the text onto the display
    logTextBox.Text = thisApp.LogEntry;
}
```

Silverlight will call this method for us when the user of the program navigates to the page. The code in the method gets a reference to the `App` class for this application and then loads the `LogEntry` text from this class and displays it in a `TextBox` on the screen. This makes sure that every time the page is displayed it will hold the latest application data.

We can do something similar when the user moves off the page. In this case we want to make sure that any changes that the user has made to the contents of the `TextBox` on the screen are reflected in the data our program is working with.

```
protected override void OnNavigatedFrom
    (System.Windows.Navigation.NavigationEventArgs e)
{
    base.OnNavigatedFrom(e);

    // When we leave the page -
    // store the semi-completed log text in the application

    // Get a reference to the parent application
    App thisApp = App.Current as App;

    // Store the text in the application
    thisApp.LogEntry = logTextBox.Text;
}
```

This method is called by Silverlight when the user navigates away from the page. It copies the text from the textbox into the application data. In effect it is the reverse of the previous code.

Note that this technique is not just important in the context of Fast Application switching; it is also how we ensure that when the user moves between the two pages in our application the pages always display the correct information.

Saving Data when the user closes the Captain's Log

We know that when the user presses the Back button at the “top” menu in an application (in this case the “Add Entry” page) the Silverlight system will close the application. When an application is closed the Silverlight system will call two methods:

1. The `OnNavigatedFrom` method is called on the page the user is presently viewing.
2. The `Application_Closing` method is called to allow the application to save data.

The `OnNavigatedFrom` method copies any data from the current page into our application, the `Application_Closing` method must then save this data in isolated storage.

```
private void Application_Closing(object sender, ClosingEventArgs e)
{
    SaveToIsolatedStorage();
}
```

The application contains a method called `SaveToIsolatedStorage` that puts the values of `LogText` and `LogEntry` into isolated storage.

Loading Data when the user opens the Captain's Log

Now that we have code that will store data when the application closes, we have to add some code that will load the program data when it starts running. We can do this by adding to the `Application_Activated` method:

```
private void Application_Launching
    (object sender, LaunchingEventArgs e)
{
    LoadFromIsolatedStorage();
}
```

When the application launches it now fetches its data from isolated storage. In the case of the Captain's Log this version of the method is fine.

Once this method has been completed the Silverlight system will load the Silverlight page that the user was at and then call the `OnNavigatedTo` method on that page. This will then take the application data and display it for the user to see.

Handling the Start button

The code that we have written so far is how we would make an application for a Windows PC. When the application starts it loads data from a file, and when it is closed it puts the data back. However, for a Windows Phone application we have a bit more work to do. This is because a phone user might press the Start button when our program is running and cause our application to be made dormant.

When the Start button is pressed the Silverlight system calls the `OnNavigatedFrom` method and then calls the `Application_Deactivated` method. We can put code in here to save data so that the user can return to the application.

```
private void Application_Deactivated
    (object sender, DeactivatedEventArgs e)
{
    SaveToIsolatedStorage();
    SaveToStateObject();
}
```

This method seems a bit strange, in that it stores the program data in two places, isolated storage and state storage. However, it turns out that this is exactly the right thing to do. We need to store the data in isolated storage because the user may never return to the program, and they would be very upset if it lost their data. We store the data in the State object because we want to have a memory copy of the data which we can use to restore application state if the user reactivates the program.

Handling the user reactivating a program

If a user finds their way back to our program it must look as if they have never gone away from it. The job of the `Activated` method is to reload data if required. We need to remember that there are two ways that an application can be activated

- It can be activated from the dormant state. In this case all the data in the program is intact and there is no need to restore anything since it is all still in memory.
- It can be activated from the tombstoned state. In this case none of the data is present and the program must reload it. It can do this by using data that it stored in the State Object when it was deactivated.

The following method does this.

```
private void Application_Activated
    (object sender, ActivatedEventArgs e)
{
    if (!e.IsApplicationInstancePreserved)
    {
        // We are coming back from a tombstone -
        // need to restore from the state object
        LoadFromStateObject();
    }
}
```

It tests the flag that is set to true if the application is being resumed from the dormant state. If this flag is false (i.e. it is being resumed from the tombstoned state) then it loads all the program state from the State object.

Using State memory

A program can always use isolated storage to hold state information, even when it is being made dormant. However this is not always the best place to store small amounts of data that are needed to hold the state of the user interface. To make it easy and quick to hold this kind of data the Windows Phone system provides a way a program can persist small amounts of data in memory. The phone calls this “state” storage and each program has its own state storage area. This memory is kept intact even if a program is tombstoned.

The Captain’s Log program can use this area to hold data when it is tombstoned rather than back to isolated storage. Then if the program is re-activated it can load the information from this area rather than using isolated storage.

```
private void Application_Deactivated
    (object sender, DeactivatedEventArgs e)
{
    SaveToIsolatedStorage();
    SaveToStateObject();
}
```

The method above is the one that is called when a program is “tomstoned”. It calls a `SaveToStateObject` method in `App.xaml.cs` to save the state of the program to memory. This method `SaveToStateObject` saves the entire state of the log by using a helper method that can save strings in the state object:

```
private void SaveToStateObject()
{
    SaveTextToStateObject("Log", LogText);
    SaveTextToStateObject("Entry", LogEntry);
}
```

The `SaveTextToStateObject` method takes two parameters, the name of the item to be stored, and the string:

```
private void SaveTextToStateObject(string filename, string text)
{
    IDictionary<string, object> stateStore =
        PhoneApplicationService.Current.State;

    stateStore.Remove(filename);

    stateStore.Add(filename, text);
}
```

The state storage dictionary object for each application is stored in the `PhoneApplicationService.Current.State` property. To gain access to this your program must include the `Microsoft.Phone.Shell` namespace:

```
using Microsoft.Phone.Shell;
```

The state storage works like a pure dictionary of objects indexed by a string and so our program to save the data can create a local variable of this type and then just store the new item in it, making sure to remove any existing one first. Loading from the state storage works in exactly the same way too:

```
private bool LoadTextFromStateObject(string filename,
                                     out string result)
{
    IDictionary<string, object> stateStore =
        PhoneApplicationService.Current.State;

    result = "";

    if (!stateStore.ContainsKey(filename)) return false;

    result = (string)stateStore[filename];

    return true;
}
```

This code is virtually identical to the code we wrote to load information from the isolated setting storage in the phone. The only change is that because the state storage is implemented as a pure dictionary we can use the `ContainsKey` method to see if the data was stored by the application when it was tombstoned. The fact that these methods look and function in a very similar way is actually good design.

Form Navigation and Fast Application Switching

The Silverlight system maintains a record of the most recently used page for an application, along with the “back stack” of pages visited to reach that page. When a user returns to an application the correct page is displayed and the user will be able to perform page navigation back to previous pages in the correct way.

However, we need to remember that although page navigation is retained, the actual content of each page is lost. Fortunately it is very easy for us to make sure that a page has the correct data content by putting appropriate code into the `OnNavigatedTo` method on each page.

Fast Application Switching and Development

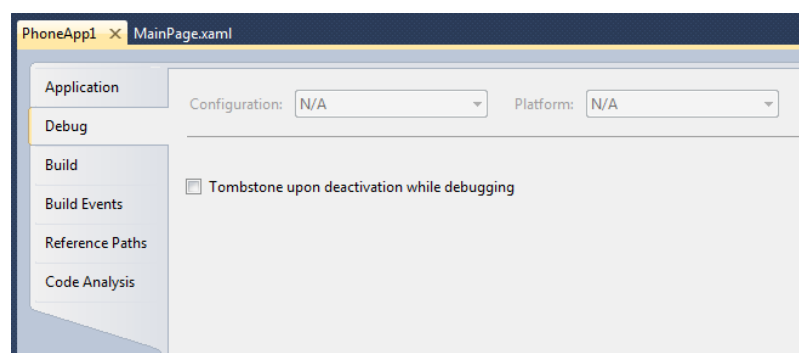
We can switch from an application even when it is running via Visual Studio. If the application is restarted any debugging sessions are automatically resumed in the development environment. This works even if the application is being debugged within a Windows Phone device.

You can see this at work if you press the Start button in the emulator while one of your programs is running. The Start page is opened but Visual Studio does not end

debugging. If you then press the Back button you will find that the program will resume execution after being reactivated.

Forcing a Program to be removed from memory

When we are testing your program we must make sure that it works correctly even when it is removed from memory and then reactivated. We know that the Windows Phone will only remove programs when memory is short. Rather than force us to try to load up memory and force our programs to be removed, Visual Studio provides a way that we can request programs to be removed from memory as soon as they are made dormant.



In the Debug tab of the properties for a project there is a textbox that, if checked, requests that a program be removed from memory when it is deactivated while debugging. This allows us to test for proper behaviour of a program even when it is completely removed from memory.

Fast Application Switching and Design

It is unfortunate that we have to do so many complicated things to give the user the impression that our application never actually stops. However, it is putting in the effort to make sure that your applications work correctly in this respect as it adds a great deal to the user experience. It is worth spending some time working with this as it is a good way to brush up on your understanding of events and dictionary storage, amongst other things. It is also an interesting way to discover that the road to producing a good user experience is not always a smooth one, particularly if you are writing programs in a constrained environment.

An XNA game can also connect to events that will be generated when it is stopped or deactivated and you can make games that store game state and high scores when the user moves away from them in exactly the same way.

The solution in *Demo 01 Complete Captains Log* contains a Windows Phone Silverlight log program that works correctly when deactivated.

9.3 Launchers and Choosers

The next part of our application development is concerned with how we use the Launchers and Choosers in the phone. These are the built in behaviours that are provided to allow our programs to perform standard actions. A Launcher is means by which our program can fire off a phone feature which does not return any data. A Chooser is where a phone feature is used to select something. Your application can then do something with the item returned.

Launchers and choosers make use of deactivation in that when a program activates launcher or chooser it is automatically deactivated. When the launcher or chooser finishes the program is reactivated.

Using a Launcher

These are all the launchers, along with a brief explanation of what each does:

PhoneCallTask	starts the Phone application with a particular phone number and display name selected. Note that the program cannot place the call, the user must initiate this
EmailComposeTask	our program can set properties on an email and then launch a task to allow the user to send the message.
SmsComposeTask	starts the Messaging application and display a new SMS message. Note that the message is not sent automatically.
SearchTask	starts the Search application using a query you provide
WebBrowserTask	starts the Web browser and displays the URL you provide.
MediaPlayerLauncher	starts the Media Player application and play a given media file
MarketplaceReviewTask	your program can start the review process of your application
MarketplaceDetailTask	your program can show details of a particular application in the market place
MarketplaceHubTask	starts the Windows Phone Marketplace client
MarketplaceSearchTask	your program can search the Windows Phone Marketplace for a particular application and then show search results
BingMapsDirectionsTask	your program can create a set of Bing Maps directions between to places, for driving or walking
BingMapsTask	your program can display a map centred on the current location of the phone or a specified location
ConnectionSettingsTask	your program can use this task to allow users to change the connection settings of the phone
ShareStatusTask	your program can fill in a status message that the user can then send to social networks of their choice
ShareLinkTask	your program can set up a url that can then be shared by the user over the social network of their choice

Each of these tasks can be started by your application and when the task is complete your application will be re-activated. Of course there is no guarantee that our application will regain control at any point in the future. The user may instead go off and do lots of other things instead of returning to the program.

Adding an email feature to JotPad



We might decide to add a Mail feature to the Jotpad program. By pressing the Mail button a user can send the contents of the jotpad as an email. The image above shows how this might work.

```
private void mailButton_Click(object sender, RoutedEventArgs e)
{
    sendMail("From JotPad", jotTextBox.Text);
}
```

When the mail button is clicked the event handler takes the text out of the textbox and calls the `sendMail` method to send this. It also adds a From message as well. The `sendMail` method is very simple:

```
private void sendMail(string subject, string body)
{
    EmailComposeTask email = new EmailComposeTask();

    email.Body = body;
    email.Subject = subject;
    email.Show();
}
```

This creates an instance of the `EmailComposeTask` class and then sets the `Body` and `Subject` properties to the parameters that were provided. It then calls the `Show` method on the task instance. This is the point at which the JotPad application is tombstoned to make way for the email application. When the user has finished sending the email the JotPad program will resume.

To make use of the tasks a program must include the `Tasks` namespace:

```
using Microsoft.Phone.Tasks;
```

The solution in *Demo 02 Email JotPad* contains a Windows Phone Silverlight jotter pad that can send the jotting as an email. Note that this will not work properly in the Windows Phone emulator as this does not have an email client set up. However, you can see the deactivated behaviour in action as the program tries to send an email and displays a warning instead.


Using a Chooser

These are all the choosers that are available:

CameraCaptureTask	starts the Camera application for the user to take a photo
EmailAddressChooserTask	starts the Contacts application and allows the user to select a contact's email address
PhoneNumberChooserTask	starts the Contacts application and allows the user to select a contact's phone number.
PhotoChooserTask	starts the Photo Picker application for the user to choose a photo.
SaveEmailAddressTask	saves the provided email address to the Contacts list Returns whether or not the save was completed.
SavePhoneNumberTask	saves the provided phone number to the Contacts list. Returns whether or not the save was completed
AddressChooserTask	your program can request that the user select an address from the contact the phone
GameInviteTask	your program can invite another phone user to a multiplayer gaming session
SaveContactTask	your program can populate a contacts entry and allow the user to save this to a contact list on the phone. Returns whether or not the save was completed.
SaveRingtoneTask	your program can give the user the option to save an audio file in an appropriate format as a ringtone for the phone

A chooser works in exactly the same way as a launcher with just one difference. A chooser can return a result to the application that invoked it. This is done in a manner we are very used to by now; our application binds a handler method to the completed event that the chooser provides.

Picture display



We can explore how this works by creating a simple picture display application this will allow the user to choose a picture from the media library on the phone and then display this on the screen. We must create the chooser in the constructor for the MainPage of our application:

```

PhotoChooserTask photoChooser;

// Constructor
public MainPage()
{
    InitializeComponent();

    photoChooser = new PhotoChooserTask();

    photoChooser.Completed +=
        new EventHandler<PhotoResult>(photoChooser_Completed);
}

```

This constructor creates a new `PhotoChooser` value and binds a handler to the completed event that it generates. The completed event handler uses the result returned by the chooser and displays this in an image.

```

void photoChooser_Completed(object sender, PhotoResult e)
{
    if (e.TaskResult == TaskResult.OK)
    {
        selectedImage.Source =
            new BitmapImage(new Uri(e.OriginalFileName));
    }
}

```

This method sets the source of the image to a new `BitmapImage` that is created using the address of the photo result. Note that the `TaskResult` property of the result allows our program to determine whether or not the user chose something to return.

The final link we need to provide is the event handler for the Load button that starts the process off.

```

private void loadButton_Click(object sender, RoutedEventArgs e)
{
    photoChooser.Show();
}

```

This method just calls the `Show` method on the chooser to start the choosing process running.

The solution in *Demo 03 PictureDisplay* contains a Windows Phone program that you can use to select an image from the media storage. This program will not work on the Windows Phone emulator which does not have any images in its media store. It will also not work on the Windows Phone device if you use the Zune software to connect the phone to Visual Studio. This is because when the phone is connected to Zune the media content is not available. However, if you disconnect the phone and then run the application it will work correctly.

9.4 Background Processing

From the previous sections we know that a Windows Phone will only allow one program at a time to be active on the phone. However, there are occasions when it would be very useful for some part of an application to be active when it is not visible to the user. The Windows Phone operating system provides this in the form of “background processing” features, where a developer can create code which will run when their application is not active.

Note that this is **not** the same as allowing a given program to run in the background. The background components are entirely separate, have no user interface and have strict limitations placed on what they can and can’t do. They only run in specific situations so that the phone operating system can ensure that they will have an impact on the user experience. The phone operating system will shut down background tasks if

memory, battery or network connectivity is compromised. The user can also disable background tasks for a given application.

We should therefore regard a background task as a piece of “icing” on an application. It should not be a fundamental part of how the application works, since it may never get to run and the user can disable it. Examples might be sales applications that display a message when a requested item appears in stock, or news applications that display the latest headlines on a tile on the Start screen.

Background and Periodic Scheduled Tasks

There are a number of different types of background process; we are going to start by considering the two general purpose background tasks, the *periodic* and the *resource intensive* tasks. We create these tasks in the same way, however the situations in which the code in them gets to run is slightly different.

- A periodic task is allowed to run for a short amount of time (up to 15 seconds) at regular intervals - usually every half an hour or so
- A resource intensive task is allowed to run for longer times (up to 10 minutes) when the phone is locked (i.e. the user is not doing anything with it), connected to a power source and has high quality network connectivity available

The periodic option is useful for updating things like weather displays, status information or location tracking. The resource intensive option is useful for downloading/uploading updates to databases, performing data processing or compression.

Agents and Tasks

When talking about these kinds of agents we need to sort out terminology. A *Task* is the container that is managed by the operating system and is run when circumstances allow. An *Agent* is the actual code payload that does the work. When we create an agent we make a single item which will run in either periodic or resource intensive mode. When the agent gets control it can determine which mode it is running in.

The Captain’s Log Location Tracker

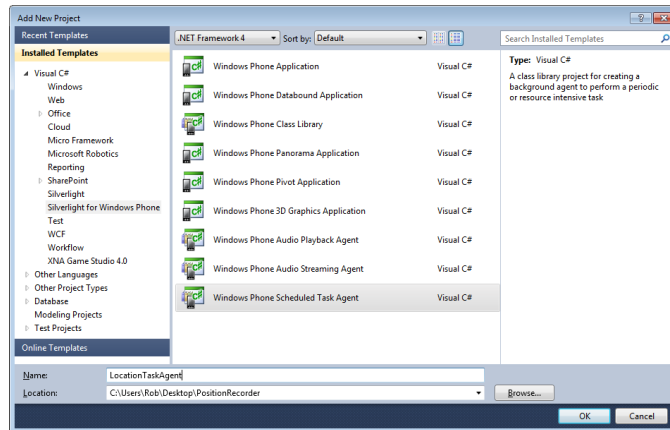
To investigate how to use an agent in an application we could return to the “Captain’s Log” program.




The program can be improved by adding a location tracking feature. When this is enabled it will save the location of the phone using a periodic task to regularly read position information from the GPS device in the phone, add a timestamp and save it in the log file. The idea is that the background task will save the phone position information when the Captain’s Log program is not active.

Adding a Background Task to Captains Log

A background task is added as an extra project to a Windows Phone solution.




Note that it is added as a Scheduled Task item, which could be either periodic, resource intensive, or both. Once it has been added the background task will show up as another project in the solution for the application.



A given application can only have one Scheduled Task associated with it. When the application is built the project that contains the code for the scheduled task is added to the other parts of the application.

The main Captain's Log program will create a scheduled task using the code in the agent project. For this to work we need to make the main application able to use the agent code. In Visual Studio this means adding the output of the LocationTaskAgent project to the references used by the foreground project. We have already seen how to do this when we considered solutions made up of multiple projects, this is another example of the technique in action.



This is how the project is added to the references, and to the right you can see how the agent code has been added to the references for this project.

Background Task Agent Code

When we create a background task Visual Studio makes an empty template that we can fill with agent code.

```
namespace LocationTaskAgent
{
    public class ScheduledAgent : ScheduledTaskAgent
    {
        protected override void OnInvoke(ScheduledTask task)
        {
            //TODO: Add code to perform your task in background
            NotifyComplete();
        }
    }
}
```

We write the code that performs the task in the `OnInvoke` method. When the method finishes it must call the `NotifyComplete` method to stop this run of the task. Note that we don't ever control when, or even if, this code gets to run. When we register the Agent with the phone we will request either periodic, resource intensive, or both execution scenarios. Then, if the user has allowed background tasks and the time is right, our program gets to run.

Our agent code can determine what kind of task is running by checking the type of the `ScheduledTask` that was passed to `OnInvoke`:

```
if (task is PeriodicTask)
{
    // Is running as a periodic task
}
```

Loading Application Data in a Background Task Agent

In the case of the Captain's Log, the task must load the log string from isolated storage, add location information to it and then store the log string back again. The application and its background task both share the same isolated storage area. Note that there is no question of the two fighting over access to this resource, as a background task will never run if the application is already active.

```
protected override void OnInvoke(ScheduledTask task)
{
    string message = "";

    string logString = "";

    if (loadTextFromIsolatedStorage("Log", out logString))
    {
        message = "Loaded ";
    }
    else
    {
        message = "Initialised ";
    }

    //When we get here we have a log string - add location to it

    NotifyComplete();
}
```

This part of the background task agent will call a method to load the log text from isolated storage. This is the same code that is used to store and load the data when the foreground program runs. If the log is initially empty (i.e. the agent is running before the application has stored anything in isolated storage) then the string in the message variable is set appropriately.

Accessing Location Information in a Background Task Agent

Now that we have the log string, the next thing to do is find the position of the phone and add this to the log text.

```
protected override void OnInvoke(ScheduledTask task)
{
    //When we get here we have a log string - add location to it

    GeoCoordinateWatcher watcher = new GeoCoordinateWatcher();
    GeoPosition<GeoCoordinate> position = null;

    watcher.StatusChanged +=
        new EventHandler<GeoPositionStatusChangedEventArgs>
        (delegate(object sender, GeoPositionStatusChangedEventArgs e)
        {
            if (e.Status == GeoPositionStatus.Ready) {
                position = watcher.Position;
                GPSDoneFlag.Set();
            }
        });

    GPSDoneFlag.Reset();

    watcher.Start();

    // Wait for 10 seconds to give time to get position
    GPSDoneFlag.WaitOne(10000);

    // When we get here we the position or we timed out
}
```

This part of the background agent creates a new `GeoCoordinateWatcher` and starts it running. The `GeoCoordinateWatcher` exposes a `StatusChanged` message that we can connect to. This event is fired whenever the GPS reader changes state. The code above binds a handler to this method that looks for a change to the `Ready` state. When the GPS reaches this state it reads the position value and then sets a flag to allow a waiting thread to continue. We used the same technique when our programs were waiting for TCP network requests to complete in Chapter 7. It allows us to write programs that can wait for the completion of a system action. If the GPS position is available the value is copied into the `position` variable.

Note that for a background task agent the behaviour of the GPS is slightly different. Since it would be very inefficient to start the GPS system for every background process the system instead makes a copy of the location information at regular intervals and uses this for any agents that request position information. This means that rather than having to wait for an event to indicate that position information is available (which is usually how a GPS watcher would be used) this program can read the position information immediately to create the log string.

Note that the code above works fine on the Windows Phone device but you might notice problems using the emulator.

Building a Log Message

Now that we have the position information the next thing we have to do is create a log message and store it back in isolated storage. At the moment the data is stored as a simple pair of values, it would be easy to save it as XML instead.

```
protected override void OnInvoke(ScheduledTask task)
{
    // When we get here we got the position or we timed out
    string positionString ;

    if (position != null)
        // If we get here we got the position or we timed out
        positionString = position.Location.ToString();
    else
        positionString = "Not Known";

    DateTime timeStamp = DateTime.Now;
    string timeStampString = timeStamp.ToShortDateString() +
        " " + timeStamp.ToShortTimeString() +
        System.Environment.NewLine;

    logString = logString + timeStampString + positionString;

    saveTextToIsolatedStorage("Log", logString);

    // When we get here we can display a toast message
}
```

The above code builds a text message which contains the position value or “Not Known” if this was not available. It then adds this to the end of the log string and saves this back to isolated storage.

Showing a “Toast” Message from a Background Agent

Once the location has been stored the location logger could tell the user that this has been performed. A background task is not able to communicate directly with the user, but it can display popup messages or change the appearance of tiles on the Start screen.



The screenshot above shows a simple popup that was produced by the above code. The phone will play an audible warning and the user can start our application by tapping the popup. The code to display a toast popup is very simple. Once the instance has been created we just have to set the title and the message to be displayed and then call the Show method.

```
protected override void OnInvoke(ScheduledTask task)
{
    // When we get here we can display a toast message

    ShellToast toast = new ShellToast();
    toast.Title = "Captain's Log";
    toast.Content = message + positionString;
    toast.Show();

    // Now we need to schedule the task for debugging
}
```

We can modify the toast to give it a uri which identifies a particular page in our application which should be navigated to if the user selects it. This would make it possible for an alert from a background task to take the user directly to the part of the program concerned with that alert.

Debugging with a Background Agent

You might think that debugging a background agent would be difficult, but actually this is not the case. Visual Studio will allow us to use breakpoints and single step through background agent code in exactly the same way as we would for any other part of our program.

When we are debugging we don't want to have to wait until the conditions are right for a background task to run. For this reason the program provides a method that we can call to launch a background task for testing. We can use conditional compilation to trigger a call of a method that will run a background agent when we are debugging.

```
#define DEBUG_AGENT
```

By defining this symbol we can cause the compiler to process code that will schedule the agent to run.

```
protected override void OnInvoke(ScheduledTask task)
{
    // Now we need to schedule the task for debugging

    #if DEBUG_AGENT
        ScheduledActionService.LaunchForTest(task.Name,
                                              TimeSpan.FromSeconds(60));
    #endif

    NotifyComplete();
}
```

The code above would make the background agent run again in 60 seconds time. Note that this does not request repeated execution of the agent, although because each time the agent runs it requests another run in 60 seconds, it has that effect.

The final call of `NotifyComplete` is how the background task agent informs the Windows Phone runtime system that it has completed processing.

Managing the Background Agent Task

Now we have created the code for our background agent we can now consider how we are going to get it running. The Windows Phone operating system contains a class called `ScheduledActionService` that manages all scheduled tasks. We can call methods on this class to search for active tasks and start and stop them. An active task has a particular name that we look for. Of course we can only control the tasks that our application runs; we are not allowed to see tasks started by other applications.

```

PeriodicTask periodicTask = null;
string periodicTaskName = "CaptainTracker";

private void StartTracking()
{
    periodicTask = ScheduledActionService.Find(periodicTaskName)
        as PeriodicTask;

    // If the task already exists and the IsEnabled property
    // is false, background agents have been disabled by the user
    if (periodicTask != null && !periodicTask.IsEnabled) {
        MessageBox.Show(
            "Background agents disabled by the user.");
        return;
    }

    // If the task already exists and background agents are enabled
    // for the application, we must remove the task and then add it
    // again to update the schedule
    if (periodicTask != null && periodicTask.IsEnabled) {
        RemoveAgent(periodicTaskName);
    }


    periodicTask = new PeriodicTask(periodicTaskName);

    // The description is required for periodic agents. This is the
    // string that the user will see in the background services
    // Settings page on the device.
    periodicTask.Description = "Log Tracker";
    ScheduledActionService.Add(periodicTask);

    // If debugging is enabled, use LaunchForTest to launch the
    // agent in one minute.
    #if(DEBUG_AGENT)
        ScheduledActionService.LaunchForTest(periodicTaskName,
            TimeSpan.FromSeconds(60));
    #endif
}

```

The above code will start an agent as a periodic task and then cause it to run in 60 seconds if we are debugging. Note that the task must have a particular name and that the user can disable background agents for a particular application if they wish. We could start a resource intensive task by creating a [ResourceIntensiveTask](#) instance instead of a [PeriodicTask](#). If we want to create a task that runs in both situations we can create a [ScheduledTask](#).



The solution in *Demo 04 Location Logger* contains a Windows Phone program that performs location logging using a background task. When the program starts running it checks for an existing tracking agent and allows the user to stop or start the tracking. It uses the “LaunchForTest” debugging mode to update the location every minute.

File Transfer Tasks

It is also possible for us to create a background task to transfer files to and from our application's isolated storage area. The transfers will take place when the application is not running and when an application is restarted it can monitor the state of any active downloads and display their status. Files can be fetched from HTTP or HTTPs hosts but at the moment FTP is not supported by this system.

Background File Transfer Policies

There are a number of policies that govern the amount of data that can be transferred by an application using the background file transfer agent:

- Maximum Upload file size: 5Mb
- Maximum Download file size over cellular (mobile phone) data: 20Mb
- Maximum Download file size over WiFi: 100Mb

A given transfer can modify these values using `TransferPreferences` if different settings are required for it.

Creating a Background Transfer

The background transfer services are all defined in the background transfer namespace:

`using Microsoft.Phone.BackgroundTransfer;`

We start a transfer by creating a `BackgroundTransferRequest` object and then adding it to the ones presently managed by the `BackgroundTransferService`. The transfer is given a name so that next time the application runs it can locate the transfer and find out the progress of the transfer.

It is also possible to bind to events that the transfer generates, so that a running application can display progress information.

Note that it is up to our application to manage the active transfers and remove ones that have completed.



The solution in *Demo 05 File Transfer* contains a Windows Phone program that will fetch an image file from the internet and display it. The program also displays status information as the file is downloaded.

Scheduled Notifications

While it is not possible for an application to create calendar appointments it can create scheduled notifications. These are retained when the phone is switched off and can cause a toast popup or a change to a Live Tile. Note that a scheduled notification can't start your application running, but if the user taps on the notification your application will be started. The notifications can be made to fire once, or repeat at given intervals.

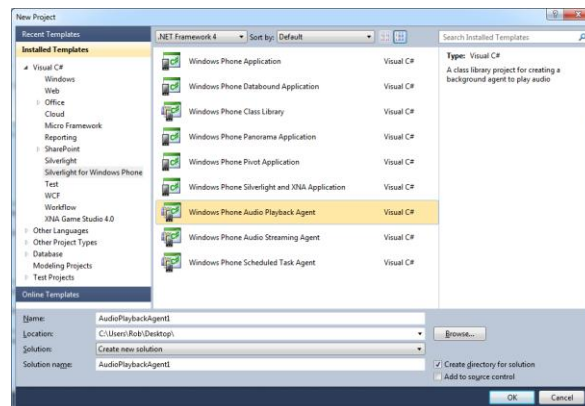
Reminders are created in a similar way to File Transfers, and can be located and managed by our applications by names that we give them.



The solution in *Demo 06 EggTimer* contains a Windows Phone program that will time from 1 to 5 minutes.

Audio Playback Agent

An application can also create an audio playback agent. This integrates with the normal playback controls and allows an application to play music when it is not running. A background agent project must be created and added to a Windows Phone solution in the same way as a background task.



The playback agent can play music from the isolated storage for the application but it does not have access to the media library on the phone. However it can be used to playback from a network media source.

What We Have Learned

1. A Windows Phone application is represented on the device by two icon files. One is used to represent the application in the Application list on the phone and the other, larger, icon is used if the application is pinned to the Start menu on the phone. Good icons are important as they are both used on the phone and in the Windows Phone Marketplace to brand an application.
2. Silverlight applications also have a “splash screen” image that is displayed while the application starts. The default splash screen is provided as part of a new Silverlight project but application writers are advised to customise this. XNA projects do not have a built in splash screen mechanism but game developers are advised to create one if their content takes more than a few seconds to load.
3. The Windows Phone operating system provides a single tasking model of execution for applications. An application may be deactivated at any time and replaced by another. The Start and Back keys on the phone device allow a user to deactivate one running application and start another by pressing Start,

and then return to the deactivated application by pressing Back. The user can also hold down the Back button to display a list of deactivated applications from which to choose.

4. A deactivated application can either be in memory (dormant) or not (tombstoned). When a dormant application is reactivated it has all memory intact. A tombstoned application must reload all data when it is reactivated.
5. There are four events which are produced in an application as it runs on the phone. They are “Launched” (fired when the application starts), “Closed” (fired when the application is closed), “Deactivated” (fired when the application is deactivated) and “Activated” (fired when the application is activated from a dormant or tombstoned state).
6. Windows Phone provides in memory storage mechanism which can be used to store live data when an application is deactivated. This can be read back if the application is ever reactivated.
7. Programs can use Launchers and Choosers to interact with phone subsystems. A Launcher starts an external application (for example and email send) and a chooser launches an external application that returns a result (for example selecting a picture from the media store). A program may be deactivated while the external application runs.
8. An application can create background agents that can run when the application is not active. These can be *periodic* (run for a short time every 30 minutes) or *resource intensive* (run when the phone is not in use, externally powered and has a good network connection). The agent is a separate project that is added to the application solution.
9. An application can create background file transfers that will run when the application is not active.
10. An application can create scheduled notifications which will run at particular intervals. These will appear even when the application is not running.

10 Windows Phone Marketplace

We now know enough to make complete applications and games that will work correctly within the Phone environment and use the built-in features of the phone system. We are now going to take a look at how we can make sure that our programs are ready for the marketplace, then we are going to look at the submission procedure and finally we are going to consider some things we can do to make our programs stand out, and hopefully increase sales.

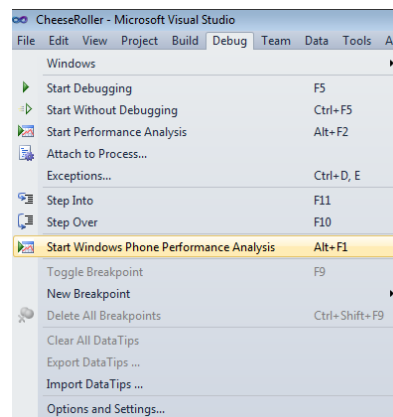
10.1 Preparing an Application for Sale

In this section we are going to focus on the things that we can do to prepare an application for sale.

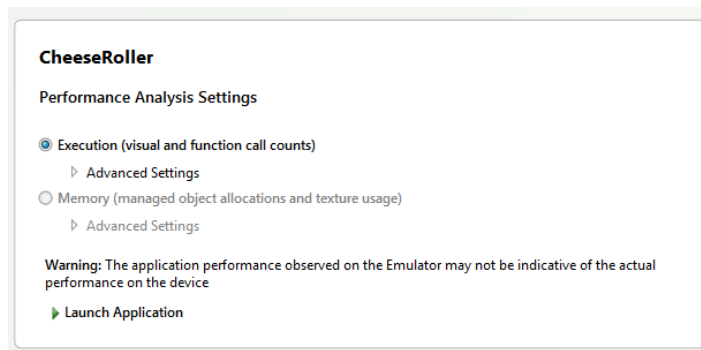
Performance Analysis

Modern computers are so powerful that they are often able to compensate for badly written software. It is frequently cost effective to buy a faster computer rather than optimise a program. However, when developing for Windows Phone we do have to worry about performance a lot more than we would for a desktop application. This is because the processor power is limited and the device itself has less memory than a PC.

The Windows Phone SDK provides a set of performance analysis tools that we can use to find out what our program is doing and how will it performs. They can give an idea of the loading our programs place on the processor, the screen update rate (how rapidly our program is updating the screen) and memory usage. We can even find out very low level detail about which particular methods are being called in our program. This is very useful as it means we can focus our optimisation on those parts of the program that are consuming most of the processor time.




The analysis tools are started from the Debug menu in Visual Studio. When they are opened they display a configuration screen:



There are a number of different options we can select, depending on whether we are interested in performance or the memory usage. When we have selected our options the program will be run within the performance analysis too. The program can run inside the emulator or a hardware device. In either case the system produces a log file of the session which is stored as part of the Visual Studio Project. This makes it easy to look back at previous performance tests and determine the effect of changes that we make.

When the program is stopped at the end of a performance analysis session Visual Studio will write the log file and then allow us to explore the results.

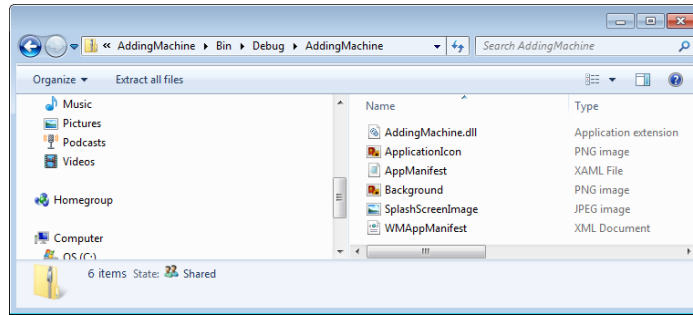


This is the log file for a Silverlight application that performs image processing. We can see that for the first few seconds the program is initialising the camera device. Then it starts to process the video stream and the CPU usage increases. The program also shows the frame rate that the application is achieving. The program itself does not use much memory, although it allocates some video buffers once the camera has initialised.

This is a very high level view of the information, it is possible to drill down and take a much more detailed view. Things to watch for are sudden jumps in CPU usage, a steady increase in memory use or a large number of Garbage Collection (GC) events. These might indicate that an application is not being very sensible with the way that it uses the phone resources.

Creating a XAP File for Application Distribution

We have already seen that an application is described by a Visual Studio solution. In Chapter 3 “Running Windows Phone Applications” we saw that an entire application is stored in a single XAP file. This is the file that is transferred into a Windows Phone device when the program is deployed. It is also the file that is submitted to the Windows Phone Marketplace when we want to sell our application.



This file contains all the program assemblies and resources along with manifest files that describe the application. The `WMAppManifest.xml` file lists the capabilities of the phone that the application uses and also sets the genre of the application. This determines where on the phone the program is stored. Programs with the Game genre are stored in the game hub, applications are stored in the applications hub.

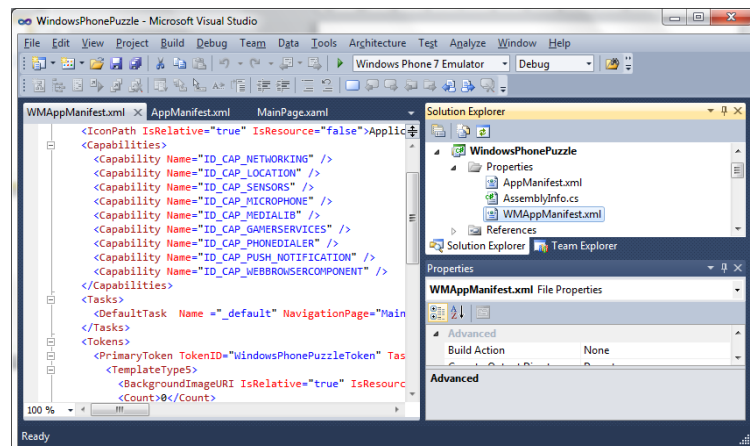
```
<?xml version="1.0" encoding="utf-8"?>
<Deployment xmlns="http://schemas.microsoft.com/windowsphone/2009/deployment"
AppPlatformVersion="7.0">
  <App xmlns="" ProductID="{b92ce770-1b22-4d52-998e-113c2784067a}" Title="PictureDisplay"
RuntimeType="Silverlight" Version="1.0.0.0" Genre="apps.normal"
Author="PictureDisplay author" Description="Sample description" Publisher="PictureDisplay">
  <IconPath IsRelative="true" IsResource="false">ApplicationIcon.png</IconPath>
  <Capabilities>
    <Capability Name="ID_CAP_GAMERSERVICES"/>
    <Capability Name="ID_CAP_IDENTITY_DEVICE"/>
    <Capability Name="ID_CAP_IDENTITY_USER"/>
    <Capability Name="ID_CAP_LOCATION"/>
    <Capability Name="ID_CAP_MEDIALIB"/>
    <Capability Name="ID_CAP_MICROPHONE"/>
    <Capability Name="ID_CAP_NETWORKING"/>
    <Capability Name="ID_CAP_PHONEDIALER"/>
    <Capability Name="ID_CAP_PUSH_NOTIFICATION"/>
    <Capability Name="ID_CAP_SENSORS"/>
    <Capability Name="ID_CAP_WEBBROWSERCOMPONENT"/>
  </Capabilities>
  <Tasks>
    <DefaultTask Name="_default" NavigationPage="MainPage.xaml"/>
  </Tasks>
  <Tokens>
    <PrimaryToken TokenID="PictureDisplayToken" TaskName="_default">
      <TemplateType5>
        <BackgroundImageURI IsRelative="true"
          IsResource="false">Background.png</BackgroundImageURI>
        <Count>0</Count>
        <Title>PictureDisplay</Title>
      </TemplateType5>
    </PrimaryToken>
  </Tokens>
</App>
</Deployment>
```

Above you can see the default `WMAppManifest.xml` file that was created for the `PictureDisplay` application. This version of the file indicates that the application will use all of the phone capabilities. Actually the only feature that it uses is the `ID_CAP_MEDIALIB` one. When an application is submitted to the Marketplace the content of this file is checked against the calls that the application makes. If the application uses resources that are not specified in this file it will fail validation.

When a customer installs a program they are told the features of the phone that it will use and are asked to confirm this. The idea is to stop programs with unexpected

behaviours being used by unsuspecting customers, for example a picture display program that also tracks the position of the user.

Before we submit this file we need to set the author, description and publisher fields.



We can edit the manifest file from within Visual Studio. When we build the application it is incorporated into the XAP file that is produced. We can find this in the bin folder that Visual Studio creates as part of the application project.

Note that when we produce a version of the program for upload to the marketplace we need to make sure that it has been compiled in Release mode, not Debug.

Creating Application Tiles and Artwork

Before we can submit an application we also need to prepare the tiles that will denote the application on the device and we also have to prepare some artwork for display on the Windows Phone Marketplace. The icons must have particular sizes and be stored in Portable Network Graphics (PNG) files. The actual sizes that you need are:

- A small mobile app tile icon (required) which is used in the phone Windows Phone Marketplace, 99 x 99 pixels in size.
- A large mobile app tile icon (optional) which is used in the phone Windows Phone Marketplace, 173 x 173 pixels in size.
- A large PC app tile icon (required) which is used in the phone Windows Phone Marketplace, 200 x 200 pixels in size.
- Background art (optional) which is used in the Background panorama in the Marketplace entry for your application, 1000 x 800 pixels in size.

You must also create at least one screenshot which is 800x480 pixels in size. You can actually upload several screenshots, this is a very good idea as it gives potential purchasers a better idea of what your program is like.

It is a very good idea to get the help of a graphic designer to prepare the icons and background art for your program.

Testing Your Application

Unlocking Devices

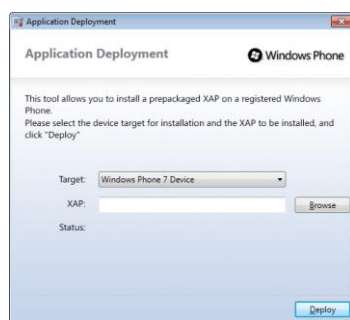
A Windows Phone can normally only run programs that have been obtained via the Marketplace. However a marketplace member can unlock a phone so that it can be used to run programs that are loaded into it from Visual Studio. Paid Marketplace members can unlock up to three devices, whereas student members who joined via DreamSpark can only unlock one. The unlocking is performed using an application that is installed as part of the Windows Phone software development kit. This application can lock and unlock phones



Applications that are deployed to the phone will stay in the phone memory and can be run from the application or game hubs in the same way as ones obtained from the Marketplace. However, you can only have up to 10 of your own programs on a given device at any time.

Distributing XAP files

If you want to send a program to another developer you can send them the XAP file and they can use the Application Deployment program to send the XAP file to their unlocked device or the emulator on their Windows PC.



This allows you to distribute applications for testing before you release the finished product.

Program Obfuscation

If you send someone a XAP file it is very easy for them to open this up and take a look at all things in it, including all the program code that you spend so long writing. Unfortunately it is a simple matter to open an assembly and take a look at how it works. We did this in chapter 3 using the ildasm program. This means that you should be careful when sending out programs that you don't give away all the hard work that you put into making the program. Whenever you send out a program you should take steps to make it more difficult for someone to unpick the contents.

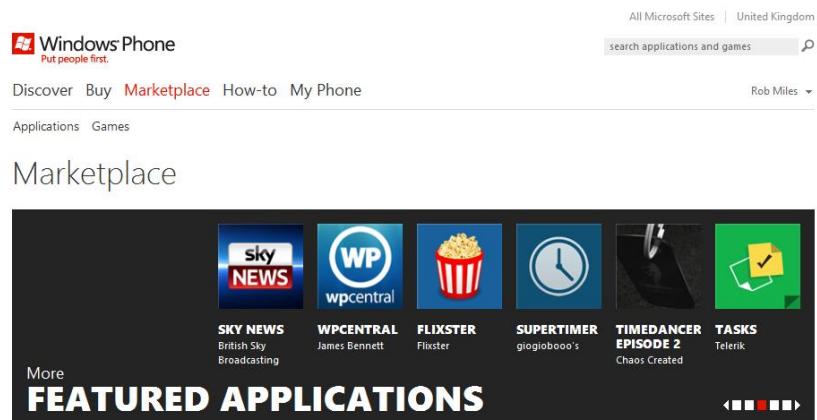
In software development this process is called *obfuscation*. It involves the use of a special tool that you use in the build process to make it very hard for someone looking at your program to deduce how it works and what it does. The tool changes variable names and adds spurious control flow to make a program very hard to understand.

There are a number of free obfuscation tools that you can use to make your program harder to understand. You could use these on your program code before you submit the XAP file. The Windows Phone developer programme also includes access to obfuscation tools from PreEmptive Solutions. These are being provided at very low cost (these systems are usually extremely expensive) and you should consider using them if you are concerned about this issue. These tools can also be used to instrument your code and allow you to find out what use is being made of the different parts of your program.

10.2 Distributing Windows Phone Applications and Games

The only way that a Windows Phone owner can get a program onto their device is by downloading it from the Windows Marketplace. Only programs that have been through the Marketplace validation process can be loaded onto a phone.

Obtaining Windows Phone Applications



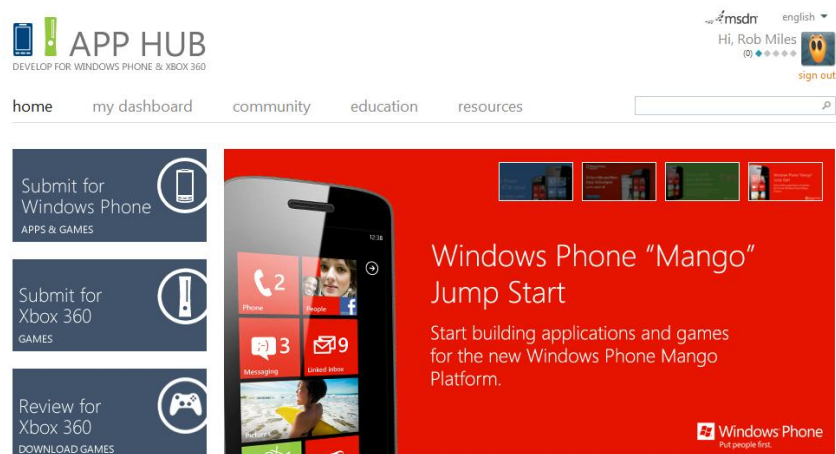
A Windows Phone owner can use their Windows Live ID to sign into their phone and gain access to services provided by Windows Live, Zune and Xbox Live along with the Marketplace.

Windows Phone owners can download applications via the Marketplace application on the device (via WIFI or the cellular network), the Zune software on the Windows PC or from the Windows Phone Marketplace tab on the Windows Phone website www.windowsphone.com. Very large applications cannot be downloaded via the cellular network and must be transferred using WIFI or the Zune program.

A given Windows Live ID can be used on up to five phone devices. If you remove an application from a device you can reload it later without buying it again as the Marketplace keeps a list of everything you have bought. The marketplace also provides a rating system which allows owners of applications to give ratings and feedback on them.

Creating Windows Phone Applications and Games

A registered Windows Phone developer can upload their applications and games to the Marketplace for distribution. They do this via the “App Hub” at create.msdn.com



Marketplace Approval Testing

Before an application can be distributed by the Marketplace it must first go through an approval process. This makes sure that the application behaves correctly in respect of things like the Start and Back buttons and also that the program doesn't do anything silly like grab all the memory or take twenty minutes to start running. This process also makes sure that programs that are offensive are not made generally available.

The testing is partly automatic and partly human. An engineer will run your program and ensure that it works correctly. The application will be tested with light and dark backgrounds on the phone to make sure that all the menus are visible.

If your application is not approved you will receive a written report that gives details of the parts of the program that need attention. When you resubmit the re-test will focus on the areas that were identified, not on the whole program again. Sometimes an application will be approved with some "warnings". This means that issues have been identified that need to be addressed, but the situation is not so serious as to require that the application cannot be approved. In that situation you would be expected to release an upgraded version later.

Once an application has been approved it is listed in one of the application categories and is available for download. A developer can produce an upgraded version of an application which, once it has been through the approvals process, will be made available as an upgrade to all existing owners.

Some applications are free and others are paid. When a phone owner buys a paid application the cost can be added to their phone bill or they can register a credit card to be used to pay for purchases. Many applications have a "trial" mode which is free. This can later be upgraded to a fully featured version. When a program is running it can easily test whether it is being used in "full" or "trial" mode.

Registered developers are allowed 100 submissions of free applications per year. If each of these is approved this means that they can make 100 free applications available. If a developer wants to submit any more free applications each submission will cost them \$20. A developer can submit an unlimited number of "paid" applications.

Applications can be loaded directly onto the phone "over the air" using the mobile phone operator or WiFi. They can also be loaded using the Zune software that also provides a media management for the phone as well. Marketplace membership

Before a developer can submit applications to the Marketplace they must first become a member. Membership is keyed to the Windows Live ID of the member. It costs \$99 per year to be member of the marketplace. If you stop being a member of the marketplace all your applications are removed. Students who have access to the DreamSpark programme can get free membership of the Windows Phone marketplace.

Members of the marketplace have their identity validated when they join and the Marketplace gives them a unique key which is used to sign their applications. The Marketplace also retains their bank and tax details so that they can be paid for application sales.

Members of the Marketplace can set the selling price of their application and receive 70% of the cost paid by the customer. This is paid directly into their bank account once they have reached a threshold of \$200 worth of sales.

Trial Mode

An application or game can be downloaded as a free demo. A program can easily determine whether or not it is operating in demo mode:

```
LicenseInformation info = new LicenseInformation();

if ( info.IsTrial() )
{
    // running in trial mode
}
}
```

The LicenseInformation class is in the Windows.Phone.Marketplace namespace. If the program is running in trial mode it can restrict the user actions, for example it could disable certain functions or stop the program running after a time limit. If the user wishes to upgrade the application it can use the Launcher mechanism to direct the user to the marketplace.

However, you should remember that a paid for program with a trial mode will still only be visible in the “paid” application part of Windows Phone Marketplace. This means that customers who only ever look for free applications will never see your program. Perhaps a more successful strategy might be to produce two versions of your program, a limited “free” one and a fully featured “paid” version. That way potential customers have more chance of finding and using your program.

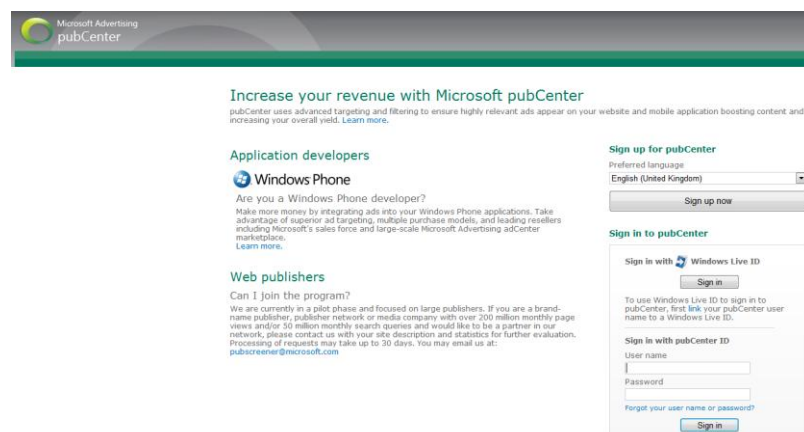
Adding Advertising

If you don’t want to sell your application, but you do want to make money from it, you can add in-application advertisements. These can be added to XNA or Silverlight programs and will display advertising content to users. The user can respond to the advertisement by linking through to the advertisers web site or placing a phone call to them. In either case you will receive 70% of the advertising revenue that is generated.



The advertisements are displayed in particular areas of the phone screen.

You can sign up for an account at the Microsoft PubCenter:



It is now possible to include advertising in applications in a large number of regions around the world. The Advertising SDK is part of the Windows Phone SDK and provides a set of test advert servers that you can use in XNA games and Silverlight applications.

The Submission and approval process

The submission process for programs is not hard to use. It is all managed from the developer site for Windows Phone and XNA Games:

<http://create.msdn.com>

Visitors to the site can download documentation and the development tools. Members of the Marketplace can use the dashboard pages to view the progress of submissions and submit new programs. There are a number of detailed walkthroughs that will take you through the membership and submission process, you should go through these to make sure you understand what is going on and what is required of you and any programs that you submit.

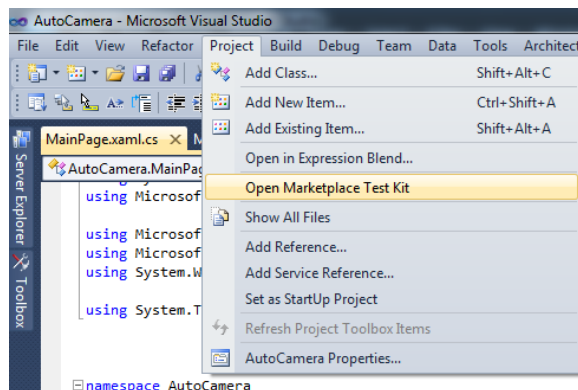
Windows Phone Application Certification Guidelines

The Windows Phone developer website provides access to a very detailed document that describes how to create applications. It is very important that you read the latest version of this document and follow the guidelines set out in it. It has been through numerous versions as the process has evolved, so you should make sure that you are using the most up to date version of the text. You can find the guidelines here:

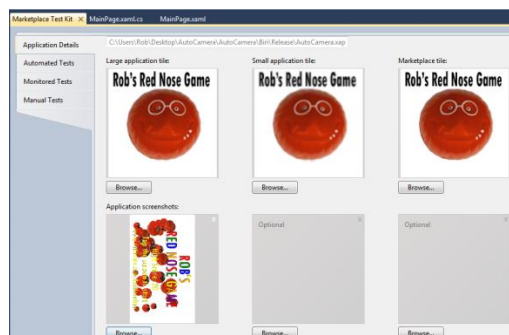
msdn.microsoft.com/en-us/library/hh184844.aspx

The Marketplace Test Kit

One way to make sure that an application passes the approvals process first time is to make use of the Marketplace Test Kit. This is built into the development tools and provides access to the same set of tools and procedures used by the testers. Visual Studio will check for changes to the procedures and update the Test Kit automatically so that it always reflects approvals policy.



It can be found in the Project menu.



The tests even check for the correct file type and size of the application tiles. There are also some automated tests and the Test Kit also provides an itemised walkthrough of the tests performed by the testers so that you can perform these tests before you submit your application.

Run Tests		
Passed: 2 Failed: 2		
Result	Test Name	Test Description
✔ Passed	XAP Package Requirements	Validation of XAP file size and content files
✔ Passed	Capability Validation	Validation of application capabilities
✘ Failed	Iconography	Validation of Application Icons
✘ Failed	Screenshots	Validation of Screenshots

This shows some tests having passed and others failed.

Private Beta Testing

If you want to get feedback from users before formally releasing your application you can create a “private” beta test release. Rather than making the application visible to all the customers on the marketplace, this form of release sends a weblink to up to 100 testers who can follow the link to a version of your program that can be downloaded onto a phone and used for up to 90 days, after which time the application is removed from their phone.

submit an app!

Let's get started. Distribute your app by giving it a name and uploading the app package. You can also learn what to expect during this [submission and certification process](#).

*** Required fields**

* App name for App Hub: App name only visible in App Hub

* Distribute to: Public Marketplace Private Beta Test. [Learn more about beta testing.](#)

* Browse to upload a file: Max size: 225 MB
Expected format: *.xap

* App version number: .

This option is set at the beginning of the submissions process.

10.3 Making your Application Stand Out

The Windows Phone Marketplace is now acquiring a fairly large number of applications and so if we want our programs to be the ones that are downloaded we have to do a bit of extra work to get noticed. Here are some tips that you might find useful.

Design to Sell

An attractive looking application will have a much greater chance of being noticed. The tiles, backgrounds and screenshots that are displayed on Marketplace should all be designed to help sell the application. You should also make sure that you provide lots of screenshots and also the large background graphic. Do not do the Marketplace description just before you publish the application. Make sure that you spend some time getting the words and the display just right. A really attractive application might even be picked as a specially selected one and highlighted on the Marketplace.

Target Different Localisations

While there might be more than one application like the one that you are selling, there might not be any in the German language. So, by producing a localised version of the program you could have that marketplace to yourself. There is very good localisation support available to Windows Phone developers, you can find out more here:

msdn.microsoft.com/en-us/hh336287

Use App Connect

The App Connect feature in Windows Phone makes it possible for customers to find your application even when they are searching for other things. For example a Bing search for “Knitting wool” could be linked through to your knitting pattern application. The process of linking an application in this way is free and all you have to do is add details of the categories under which your application can appear. You can find out more here:

msdn.microsoft.com/en-us/library/hh202969.aspx

Give your Program Away

If you provide a free version there is more chance that users will download it and use it. The free version could contain adverts or “nag” messages. Because it is easy to create a launcher that will take the user into the Windows Phone Marketplace you can provide a “buy upgrade” option that can target a paid version of the program or game.

Release Upgrades/Episodes

Rather than releasing a 15 level game or a fully featured program you could instead release a “starter” version and then provide upgrades at regular intervals. There are a number of reasons why this is a good idea.

- You will get to market earlier than your competition.
- Your application will appear regularly in the “new items” list on Marketplace. Since a lot of sales of an application are made in the few days after a new release, making more new releases will improve your sales.
- You can get the customers engaged in providing ideas for new content and features. This means that you can build a good relationship with them.

Change Categories

When you submit an application you will be asked to choose the category in which it is to be displayed in the Marketplace. It is often the case that one particular category doesn’t really work. A game could be thought of as a sports simulation, an arcade game or a puzzle. If you are not sure which is the best category, try them moving the game around and see how this affects sales. This will also get you exposure in other categories, which will also be beneficial.

Encourage Good Feedback

A good relationship with customers is something that you should work hard to cultivate. In the unlikely event of your application failing it should display a populated mail message and provide the user with an easy way of sending it. Then, if you receive such a message you should respond constructively. Building a relationship in this way makes customers into debuggers and advocates of your programs, which is really useful.

10.4 What To Do Next

At this point in the document you should have all the skills you need to make something that can be published on Windows Phone. Here is a list of things you can do to take the next step.

Register as a Developer

If you have not already done this you should head over to create.msdn.com and sign up. If you are a student you should stop off at dreamspark.com and get a free registration.

Get the Toolkit

The tools can be found at create.msdn.com and will provide you with everything you need in a single install. This includes Visual Studio, the emulator, Expression Blend, the Advertising SDK and programs to unlock your device.

Publish Something

Once you have got your development environment working you should put something out in Marketplace. It doesn't really matter whether it is something you are massively proud of, and you can always remove it later, but you never know, there might be someone out there desperate for the application that you make.

Make Suggestions

If you find something you don't like, or have an idea that could make things better, you can use the UserVoice forums to suggest things and vote on them. The Windows Phone team really do read and respond to these suggestions, and quite a few have found their way into the device. You can find the forums here: wpdev.uservoice.com

Resources

There are lots of good resources for phone developers.

Learning Windows Phone Development

Windows Phone Jump Start

create.msdn.com/en-US/education/catalog/article/wp7_jump_start

Windows Phone Code Samples

msdn.microsoft.com/en-us/library/ff431744.aspx

Application Features for Windows Phone

msdn.microsoft.com/en-us/library/ff402551.aspx

XNA

XNA Game Studio 4.0 on MSDN

msdn.microsoft.com/en-us/library/bb200104.aspx

XNA Game Development Resource Page

create.msdn.com/en-us/education/gamedevelopment

Farseer Physics Engine

farseerphysics.codeplex.com

Silverlight

Windows Phone Silverlight, Development Quickstarts

create.msdn.com/en-us/education/quickstarts

Royalty Free Icons

thenounproject.com

The Silverlight Toolkit

silverlight.codeplex.com

Design toolbox

www.microsoft.com/design/toolbox

Windows Phone Azure Toolkit

watoolkitwp7.codeplex.com

Program Ideas

We now know how to make programs for Windows Phone. We can create multi-page Silverlight interfaces which match those of the phone itself and we have also had a look at how to create games using XNA. We also know how to create connected applications that make use of data hosted on the internet. We have discovered some of the difficulties encountered by programmers when they write code for small, resource constrained, devices and we now also know how the Windows Phone operating system allows us to create useable applications in spite of these issues. Finally we know how to use the underlying Windows Phone system so that our applications can make use of features in the phone itself.

At this point we have a superb set of tools, next we have to find a problem to solve with them.

Application Ideas

Applications ideas do not always appear fully formed. Quite often you will start with a small idea for a solution and then add features or discover situations where the solution could be made even more useful. This is not a bad way to come up with something original, but you must beware of adding too many feature ideas before you have made something work, otherwise the whole thing might collapse under its own weight before you have built anything.

You can get good ideas for applications by talking to lots of people and trying to find out what would be useful to them. Not everyone is aware of just how much you can do with modern devices and so if you say that you can make something mobile that can track position, take pictures and make use of internet services they might have an idea of a situation where some of that power would be useful to them.

Game Ideas

In some ways game ideas are easier to come by than application ideas. Systems like XNA let you “doodle” with game code to find out what it does. I’ve already mentioned the importance of playtesting. This is where you give your game to other people to find out if the game is playable. Playtesting can also throw up ideas for game development. Often a playtester will suggest changes to the program that will make it more interesting and fun. Don’t be afraid to do silly things in a game (add 10 times as many aliens, invert gravity, make everything bigger/smaller etc etc) and see what happens to the play experience.

The great thing about a game is that rather than solving a specific problem, as with an application, a game just has to be fun to play.

Fun with Windows Phone

Windows Phone provides an amount of processing power that was unavailable in the world a few years ago, let alone in everyone's pocket. It also provides an amazing level of connectivity, high performance graphics and devices such as cameras and location sensors that make it possible to build truly novel applications.

Personally I reckon that just being able to program for this device, let alone sell the results, is inspiring enough. I hope you have learned enough from these notes to get yourself started on the fun you can have with this platform.

Rob Miles
October 2011

What We Have Learned

1. Windows Phone owners get their applications from the Windows Phone Marketplace.
2. Registered developers can submit applications for approval and distribution via the Marketplace.
3. Developers register and track the progress of their applications via the Windows Phone developer website.
4. Registering as a developer costs \$99 a year, students can register for free via the Dreamspark programme.
5. Developers can produce free or paid applications, but are limited to submitting 100 free applications per year. Further submissions of free applications cost \$20 each.
6. A registered developer can unlock phones so that they can run programs compiled in Visual Studio.
7. Applications are distributed as XAP files which contain all the resources and content along with a manifest file that describes the content and gives the phone setting information about the application.
8. XAP files can be loaded directly onto unlocked devices by developers.
9. Programmers should obfuscate their programs so that it is not easy for anyone obtaining a XAP file to find out how the code works. There is an obfuscation processor available via the Windows Phone developer web site.
10. A developer must read the Windows Phone Certification guidelines before submitting programs to the approval process.