

Beginning Microcontrollers

with the

MSP430

Tutorial

Gustavo Litovsky
Version 0.4



This document is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

You May copy, transmit and distribute this document to anyone provided you attribute it to Gustavo Litovsky and the copyright notice remains. If you wish to use any part of the document for commercial purposes, please contact me at gustavo [at] glitovsky [dot] com

All Trademarks are the properties of their respective owners.

Original copies of this document are located at
www.glitovsky.com

What's new in this version?

Version 0.4

- A new chapter on Flash

Version 0.3

- A major update to the UART section describing more in detail how to configure the interface, including an algorithm on accurately calculating the registers. Also included are details on communicating with the Host PC such as framing.
- More information about the selection and use of crystals with the MSP430 and other micro-controllers
- Major editing improvements that better convey the information

Preface

I decided to write this tutorial after seeing many students struggling with the concepts of programming the MSP430 and being unable to realize their applications and projects. This was not because the MSP430 is hard to program. On the contrary, it adopts many advances in computing that has allowed us to get our application running quicker than ever. However, it is sometimes difficult for students to translate the knowledge they acquired when studying programming for more traditional platforms to embedded systems.

Programming for embedded systems (as is the case with the MSP430) is not more difficult than personal computers. In fact, it is much better in that it exposes us to the little details of how the system operates (the clocks, I/O) at a level that anyone can learn, as well as unparalleled flexibility. This, however, also forces us to make critical decisions that affect how the application runs.

The MSP430 microcontroller is an extremely versatile platform which supports many applications. With its ultra low power consumption and peripherals it enables the designing engineer to meet the goals of many projects. It has, of course, its limitations. It is geared mostly towards low energy and less intensive applications that operate with batteries, so processing capabilities and memory, among other things, are limited.

This tutorial will begin from the basics, introducing you to the theory necessary to manipulate binary digits and digital logic as used in the microcontroller. Using this you will be able to see how to manipulate the registers which control the operation of all microcontrollers. It will then cover the specifics of modules in the MSP430, from both the hardware and software perspective. I decided to follow this format primarily because you, the reader, might want to use this tutorial as a reference and just jump to a module you need help with. But I also wanted to keep the tutorial to be accessible for beginners and so the first part of the tutorial covers many of the essentials.

If you wish to begin programming immediately and understand code, you could skip to Chapter 4. Previous knowledge of the C programming language is assumed, although if you've done some Java programming, you will likely be familiar with the concepts. It is important to note that this tutorial should be supplemented by the Datasheet, User's Guide, Example Code and Application Notes for the specific MSP430 derivative used and any other component in the system. These are extremely useful in teaching you how to integrate all the ideas you have and apply them. All these documents are freely available at www.ti.com Don't Ignore them! They will answer many of your questions.

A companion file is included . This code has two components. The first is straight C code primarily based on the slaa325 application note from TI which was ported to the EZ430 and establishes very basic communications. The second part is based upon the TI Sensor Demo Code for the EZ430-RF2500.

Contents

Preface	v
1 Introduction	1
1.1 Why Microcontrollers?	1
1.2 What Hardware do I need?	1
1.2.1 EZ430-F2013 USB Stick Development Tool	2
1.2.2 EZ430-RF2500	3
1.2.3 Experimenter Boards	4
1.2.4 FET Debugger	4
1.2.5 Custom Hardware	4
2 Getting Started	7
2.1 Introduction	7
2.2 Running a Project using IAR	8
2.3 Running a Project using CCS	10
3 Microcontroller Basics	11
3.1 Data Types and Numeric Representation	11
3.2 Hexadecimal for MSP430	12
3.2.1 Conversion of Numbers	14
3.3 Digital Operations	14
3.4 Manipulating Module Registers	15
3.4.1 XOR Operator	20
3.5 ASCII	21
3.6 Conclusion	21
4 Beginning Programming for MSP430	23

5	MSP430 Clocks	27
5.1	Introduction	27
5.2	Internal Oscillators	28
5.3	External Crystals	28
5.4	Clock Sources	30
5.5	Clock Signals	30
5.6	Basic Clock Module In EZ430-RF2500	31
5.7	Considerations for using clocks	33
6	General Purpose Input Output - GPIO	35
6.0.1	Pin Multiplexing	35
6.1	Switches	38
6.1.1	Debouncing	39
6.2	LEDs	39
6.3	Bit banging	40
7	Flash Memory	41
7.1	Introduction	41
7.2	Flash and Memory Organization	41
7.3	Looking at the Flash	42
7.4	Programming the Flash	45
7.4.1	Configuring the Flash	46
7.4.2	Reading the Flash	47
7.4.3	Writing to Flash	48
7.4.4	Finalizing Flash Access	49
7.5	Flash Programming from the RAM	49
7.6	The Linker and Flash	50
7.7	Using the Flash for data storage	55
7.8	More Examples	55
8	Timers	57
8.1	Introduction	57
8.2	Setting up the clock	57
8.3	Timer Modes	58
8.4	Accuracy	58

9	Interrupts and Low Power Modes	59
9.1	Interrupts	59
9.2	Low Power Modes	65
9.2.1	Exiting Low Power Modes	67
10	Analog to Digital Converters - ADCs	69
10.1	Introduction	69
10.2	ADC Parameters	70
10.2.1	ADC Resolution	70
10.2.2	ADC Sampling Frequency	71
10.3	ADC Example - Temperature and Voltage	72
10.3.1	ADC Clock	73
10.3.2	ADC Modes	73
11	Digital Interfaces	75
11.1	Serial Peripheral Interface (SPI)	76
11.1.1	Configuring SPI	77
11.1.2	Using SPI for Communications	79
12	UART	81
12.1	Hardware Connection	81
12.1.1	UART connectivity on the MSP430F1611	81
12.1.2	UART connectivity on the MSP430F2274	82
12.2	Using UART	82
12.3	Configuring the UART	83
12.3.1	Selecting the UART Function for Pins	84
12.3.2	Enabling UART RX and TX	84
12.3.3	Select the character format	84
12.3.4	Selecting A Clock	85
12.3.5	Setting the Baud Rate Generator	86
12.3.6	Enabling the Module	87
12.3.7	Enabling Interrupts	87
12.4	Configuring the Host Computer	88
12.5	Sending and Receiving Information with the UART	88
12.5.1	Sending Bytes	88

12.5.2 ASCII	89
12.6 Sending Multiple bytes	89
12.6.1 Sending Strings	90
12.7 Receiving Data	91
12.8 Framing	93
12.9 Future Additions	95
13 Wireless Communications with CC2500	97
13.1 Introduction	97
13.2 SPI Interface	98

List of Figures

1.1	EZ430-F2013 Development Stick	2
1.2	EZ430-RF2500 Development Stick	3
2.1	Creating a new empty project in IAR	9
3.1	Converting Numbers Using the Windows Calculator	14
3.2	ADC10 Register Example	16
5.1	Loading Capacitors on Crystal	28
5.2	Selecting the DCO Frequency	32
6.1	MSP430F2274 Schematic	35
6.2	Connecting Switches to MSP430	38
7.1	MSP430 Memory Organization Examples	42
7.2	MSP430 Memory after writing 0xAA in IAR	44
7.3	MSP430 Memory after writing 0xFF in IAR	44
7.4	MSP430 Memory after writing 0x00 in IAR	45
7.5	MSP430 Memory after writing 0xAA in CCS	46
9.1	Low Power Mode Savings	66
10.1	3-bit ADC Example	71
11.1	SPI Configurations	77
12.1	Connecting UART	82
12.2	UART Bits	83

List of Tables

3.1	Data Types	12
3.2	Hexadecimal Number System	12
3.3	Extended Hexadecimal vs. Binary and Decimal	13
3.4	Digital Operations - OR and AND	14
3.5	Digital Operations - NOT and XOR	15
3.6	Hexadecimal representation of position	19
3.7	XOR Truth Table	20
3.8	Extended Hexadecimal vs. Binary and Decimal	21
11.1	SPI Signals	76
11.2	SPI Pins in the MSP430F2274	78

Chapter 1

Introduction

1.1 Why Microcontrollers?

To those who are unfamiliar with these devices, microcontrollers might seem extremely simple and rare as compared to personal computers. However, microcontrollers and microprocessors are embedded in many devices, with hundreds of them forming part of today's automobile, controlling everything from the engine to the sound system. Cellular phones include more sophisticated microprocessors, but these are not as different from the MSP430 that we will cover here in that the basics apply to both. The power of microcontrollers lies in their small size and adaptability. As opposed to fixed digital circuitry, microcontrollers can be programmed to perform many applications and can be later changed when improvement are required. This saves both time and money when a field upgrade is required (which you will discover to be a grand objective of any company). However, there are limitations with respect to processing power and memory (the two biggest problems you face the use of embedded processors). It is the job of the engineer to come up with the requirements of the application and select the proper device for the application. With the advances in processing capability, many more applications can be realized today with microcontrollers than ever before, especially due to their low power profile. Indeed, the flexibility of these devices will ensure their incorporation in designs many years into the future.

It is important to note that a wide array of microcontrollers exist, some rivaling or surpassing the capabilities of full fledged computers in the 70s, 80s, and maybe even 90s. UV Erasure of microcontroller and ROM are today mostly a thing of the past. With the advent of Flash memory, the microcontroller can be programmed hundred of thousands of times without any problems. Also, they incorporate a wide array of modules such Analog to Digital Converters, USB, PWM, and Wireless transceivers, enabling integration into any kind of application.

1.2 What Hardware do I need?

It is often the case students are confused about what exactly is needed for using a particular microcontroller. Indeed, many companies assume everyone will understand the basics and therefore skip this vital information. It used to be that companies provided the silicon (actual chip) and let the

application engineer sort out everything else. Today, most companies attempt to provide as much information as possible to the developing engineer since most engineers want to get the application working as soon as possible and avoid wasting time.

Normally, an embedded system is composed of the following:

- An Embedded Microcontroller
- A Programming/Debugging interface for the Microcontroller
- Microcontroller Support Circuitry
- Application Specific Circuitry

The Programming/Debugging interface is the most often ignored element of the system, but it is a very important factor. Without it, how is code supposed to be put in the Microcontroller? With just Programming capabilities, we can download a firmware image into the microcontroller. However, Debugging is often a necessity since no person can write perfectly good code and ensure it can run correctly. A common debugging interface is JTAG, which is often used in Microcontrollers. The MSP430 also uses this interface, but TI adds extra functionality whose information is available only under a Non Disclosure Agreement. Therefore, I would argue that selection of both the debugger(or programmer) and even the compiler will dictate much of the effectiveness of the time spent on the design. Do not settle for inferior tools! You will pay for them later in sweat and suffering.

Today's companies offer many platforms with which to develop and learn how to use their micro-controller. Such is the case with TI. Some of their most useful development platforms are:

1.2.1 EZ430-F2013 USB Stick Development Tool




Figure 1.1: EZ430-F2013 Development Stick

This is one of the first MSP430 development tools to be in the USB stick form factor, which is gaining popularity because it is so easy to use. Since it is very low cost(\$20)and has an integrated debugger, it allows very quick development. It is composed of two boards (both located inside the plastic casing): The programming board and the target board. The target board is the board

with the actual microcontroller (an MSP430F2013) and all the circuitry needed to use it. It can be detached from the programmer once the software has been downloaded to the MSP430.

The debugger board, with its USB connector, can allow programming on any computer (although there might be issues with a specific Operating System being unsupported). For more information, see the following links:

[EZ430-F2013 Home Page](#)

[MSP430F2013 Home Page](#)

[MSP430 Design Page - F2013 Contest](#)

This last website has the EZ430 contest that can provide you with real insight as to what can be done just with the EZ430.

1.2.2 EZ430-RF2500




Figure 1.2: EZ430-RF2500 Development Stick

This development board is very similar to the regular EZ430 but includes a RF2500 transceiver and a MSP430F2274 microcontroller (not the F2013). This is the kit that is the subject of this tutorials because it delivers great value - microcontroller, transceiver and debugger - at a very low price

TI also supplies the sensor demo code which can be quickly programmed on both devices to enable quick temperature monitoring. One target board has to be programmed with the End Device code and the other with the Access Point code. The End device is the connected to the battery board while the target board that has the Access Point software is left connected to the programmer board and to the PC.

The sensor demo is simple. The End Device takes a reading of its Temperature Sensor (integrated into the MSP430 and accessible as an ADC channel) and Voltage (the one applied to the MSP430F2274). It sends this data to the Access Point, which also takes its own temperature and voltage readings. The AP then integrates both of these and forwards them to the PC using the UART. This data can be read using a hyperterminal equivalent program (Putty, a free program, is recommended).

[EZ430-RF2500 Home Page](#)

1.2.3 Experimenter Boards

TI also provides more comprehensive boards, called Experimenter Boards. These are much larger and don't come with a built in Debugger. For that you need a FET Debugger(which is discussed below). The two boards that are available are:

[MSP430FG4618/F2013 Experimenter Board Home Page](#)

[MSP430F5438 Experimenter Board Home Page](#)

[Putty - Terminal Emulation Software Home Page](#)

1.2.4 FET Debugger

To access the JTAG pins and debug the MSP430, TI provides a debugger, called a FET Debugger. If you have purchased a board with a 14-pin JTAG connector, such as the experimenter boards or many other boards by third parties, you will likely need this debugger.

There are two versions: one that connects a parallel port and another with USB connectivity. The USB one is the most common one. This debugger allows step by step execution of code, breakpoints, and other advanced things.

For more information:

[MSP430 USB Debugging Interface Homepage](#)

There are a few other debuggers out there, mostly sold by a company called Olimex and some other third parties.

1.2.5 Custom Hardware

Most often, once you've developed your code and tested it on a development platform, you would like to create your own PCB with the microcontroller. The microcontroller itself is only an IC. It requires several things to operate (which every development kit as mentioned above provides):

- Supply voltage consistent with Electrical Specifications (1.8V - 3.6V for most MSP430)
- Decoupling capacitors to reduce noise on the supply voltage (no power supply is perfect)
- External Crystals (in some cases where they are needed for the clock generation)
- A Programmer/Debugger or Bootstrap Loader

The list above provides the basic elements most microcontrollers need, which are besides the specific parts and connections related to the application itself (such as sensors, transceivers, passive components etc).

Users of the Microcontroller must ensure they provide the correct power supply. Although the MSP430 family requires little current and thus can be operated by batteries or other low current

sources without any problems, it requires a specific voltage to be applied, and any deviation from the Maximum Recommended Specifications (available in the datasheet) can destroy the IC. Microcontrollers such as those made by Microchip and Atmel usually can tolerate 5V (and in some cases require it), but MSP430 generally accepts 1.8V to 3.6V, with possible damage to the IC when the voltage is over 3.9V or so. It is essential that if you use your custom designed circuit for the MSP430, you comply with the requirements set forth by the datasheet.

External Crystals should be used in applications that require more accuracy than is available from the microcontroller. They add more size and cost, but enable more accurate clocks. Chapter 5 provides some more information about the clocks and crystals.

Most students are familiar with computers and how to program them, using a compiler to generate the binary code and execute it. Microcontrollers, however, are different. They are programmed with code produced by a special compiler in a computer. Once the code is compiled and linked, it can be downloaded and debugged on the actual microcontroller (provided the interface to do so is available). The MSP430 devices use an interface called JTAG to perform these functions. JTAG allows a user to download and debug the code on the microcontroller, after being written with one of the several compilers that are available. This interface is accessed by using a FET Programmer that connects the computer to the microcontroller. For MSP430 devices that do not have many pins available, the programming interface available is called Spy-bi-Wire. This is proprietary to TI and makes JTAG use 2 lines instead of the usual 4 or 5.

Chapter 2

Getting Started

2.1 Introduction

I hear and I forget; I see and I remember; I do and I understand.

- Old Chinese Proverb

As the chinese proverb indicates, doing is the best way of understanding. This chapter was created to get you started compiling MSP430 programs, simply because it is the best way you will actually learn and understand what this tutorial teaches. In order to compile MSP430 applications, an MSP430 compiler is needed. There are several options:

- Code Composer Studio (CCS) - TI's own MSP430 development suite that is also used for many other TI devices. It is based on Eclipse and is a cost effective solution for development. Currently it is in version 5.1.

Code Composer Studio for MSP430

- IAR Embedded Workbench for MSP430 - IAR is a well established company and its compiler is very good. It produces output code that is sometimes smaller and faster than others. On the flip side, it is usually more expensive. A free version called the Kickstart edition is provided that allows up to 4kB or 8kB of code depending on the MSP430 device used. Another possibility is a 30-day Evaluation version that has no limitation.

IAR Embedded Workbench Kickstart for MSP430

- MSPGCC - A free MSP430 port of the well known GCC compiler. It requires some more work to get going but is a very attractive option. It can be used for professional development.

MSPGCC Wiki

If you are just beginning to program, I recommend that you get IAR Kickstart just to compile a few programs. You can then move to CCS or MSPGCC.

2.2 Running a Project using IAR

I assume that you've installed IAR. This should be straightforward. IAR uses the concept of project in a workspace. A workspace is simply a container of projects, and if you open IAR, a workspace is already created for you. Workspaces are actual files with the extension 'eww'. In order to compile, link and debug MSP430 code, we must create a project. To do this, go to Project-> Create New Project. Select MSP430 for the toolchain:

Once you press OK, you will be asked to save the project file. Please note that the file name will also be the project's name. Once the project is created, we must add the source code and configure it. Since we have created an empty project, no files are incorporated. Press the New Document button and IAR will open a new tab for a new file. save this file with the name main.c (the actual name can be almost anything but it helps to make it obvious).

Copy paste the hello world source code below in main.c and save it:

Listing 2.1: MSP430 Hello World - hello_world.c

```
1 #include "msp430x22x4.h"
2
3 volatile unsigned int i;      // volatile to prevent optimization
4
5 void main(void)
6 {
7     WDTCTL = WDTPW + WDTHOLD; // Stop watchdog timer
8     P1DIR |= 0x01;           // Set P1.0 to output direction
9
10    for (;;)
11    {
12        P1OUT ^= 0x01;       // Toggle P1.0 using exclusive-OR
13        i = 50000;          // Delay
14        do (i--);
15        while (i != 0);
16    }
17 }
```

We will cover the actual operation of the code later, but this code simply blinks the LED on board. Note that main.c is still not part of the project and cannot be compiled. To add it, right click on the project name in the workspace window, go to Add and then Add main.c. Note that this only works if main.c is the currently active tab. Else, select the Add files option and choose main.c.

We now need to configure the project options. To do this right click on the project name in the workspace browser and go to Options. A window will appear. This window includes all the options for the project. The most important one at the moment is to select the actual MSP430 device that will be debugged. In the Device selection box there is a button, clicking on it shows a list of MSP430 devices. Select here the MSP430 that you have on your board. ANother very important aspect is to ensure that the FET Debugger is used. For this go to the Debugger in the list and change " Simulator " that appears as default to " FET Debugger ".




Figure 2.1: Creating a new empty project in IAR

If while debugging your application, the download of your application is very slow, you can try to figure out why their application wasn't working only to discover that the application is not working.

Once you've finished, press OK and close the options dialog. Then right click on the project and press Rebuild all. IAR will compile and link the source code. You may then download and debug using the CTRL+D shortcut.

IAR will automatically download the application and if everything is OK, will stop at the first line of code. Pressing on the Go button will execute the code freely.

2.3 Running a Project using CCS

During the CCS installation, you should have chosen to install support for the MSP430 and the USB FET device.

Similarly to IAR, CCS also has the concept of a workspace where projects are located, though CCS uses Eclipse's method of storing a workspace which is based on a folder with the information. When you first open CCS, you'll be asked to give the path to the workspace. This can be almost any location. The first time CCS opens, it will request the mode in which you want to operate. The Code Size limited option is good for starting and evaluating until you must move up because of the code limitation. Once CCS opens, select the Project Menu and then New CCS Project. You can put any valid Project name. Ensure that the device family is MSP430 and select the variant and then the particular device. Finally select an Empty project.

CCS creates a new project and also a main.c file.

Copy paste the hello world source code below in main.c and save it:

Listing 2.2: MSP430 Hello World - hello_world.c

```
1 #include "msp430x22x4.h"
2
3 volatile unsigned int i;      // volatile to prevent optimization
4
5 void main(void)
6 {
7     WDTCTL = WDTPW + WDTHOLD; // Stop watchdog timer
8     P1DIR |= 0x01;           // Set P1.0 to output direction
9
10    for (;;)
11    {
12        P1OUT ^= 0x01;        // Toggle P1.0 using exclusive-OR
13        i = 50000;           // Delay
14        do (i--);
15        while (i != 0);
16    }
17 }
```

No configuration changes are needed. Simply go to Run->Debug or press F11. A new debug session will be open. You can then select run and the example will run.

Chapter 3

Microcontroller Basics

Microcontrollers are binary computers and so they operate on the basis of binary numbers. Binary numbers consist of only 1 and 0. Binary, however, is unnatural for humans to use. Assembly language is one step above binary. It is therefore the most basic language for controlling computers since it represents binary directly and it is easier to understand. Knowledge of assembly is not completely necessary to program the MSP430, but it is useful in optimizing routines to get the maximum performance (in terms of speed or memory). The C programming language is the primary language used and will be followed throughout this tutorial. In general, a compiler translates the C code into the binary code and we will not worry about how this is performed initially.

A microcontroller is not just a CPU. It usually incorporates a range of peripherals, besides the memory and storage needed to operate. Simply put, it is a computer with some specialized functions. Of course, it cannot compare to a modern PC in aspects such as speed and processing capability for all but the high performance CPUs, but it is useful in a wide variety of applications where larger processors are not feasible. We begin by discussing the most important part of any computer: numbers and computation.

3.1 Data Types and Numeric Representation

Using binary to represent numbers is sometimes confusing. I will attempt to clear this up. This section is critical to understanding how we operate on registers and control the microcontroller. The first thing to understand is that the C code can accept many data types and that the compiler assigns them a specific number of bits. The table below summarizes these data types and their capacity to represent values. If this is not strictly followed, code problems, overflows and errors will occur. Table 3.1 shows the data types and their sizes as used in the IAR compiler:

Therefore, when a variable is declared to be unsigned char, for example, only values from 0 to 255 can be assigned to it. If we tried to assign a higher number, overflow would occur. These data type representation are also compiler dependent. You should check your compiler's documentation for what data types are available. Unsigned and Signed simply determine whether numbers represented can be negative. In the case of unsigned, all numbers are assumed to be positive (0 to 255 for char, for example), while if the char variable was declared to be a Signed char its values go from -128 to 127.

Table 3.1: Data Types

Data Type	Bits	Decimal Range	Hex Range
Unsigned Char	8 bits	0 to 255	0x00 - 0xFF
Signed Char	8 bits	-128 to 127	0x00 - 0xFF
Unsigned Int	16 bits	0 - 65535	0x0000 - 0xFFFF
Signed Int	16 bits	-32768 to 32767	0x0000 - 0xFFFF
Unsigned Long	32 bits	-2^{31} to $2^{31}-1$	0x00000000 - 0xFFFFFFFF
Signed Long	32 bits	0 to $2^{32}-1$	0x00000000 - 0xFFFFFFFF

It is possible to specify whether by default a variable is Signed or Unsigned. This is done at the compiler's options menu.

When creating a variable, you must be aware of how large a number you will need to store and choose the data type appropriately.

3.2 Hexadecimal for MSP430

Hexadecimal notation is essential because it is so widely used in the microcontroller as a direct form of representing binary numbers. Hexadecimal is a number system which uses 16 symbols to represent numbers. These symbols are 0 to 9 followed by A to F. Table 3.2 shows these symbols and their equivalent in the decimal and binary number system.

Table 3.2: Hexadecimal Number System

Hexadecimal	Decimal	Binary
0	0	0
1	1	1
2	2	10
3	3	11
4	4	100
5	5	101
6	6	110
7	7	111
8	8	1000
9	9	1001
A	10	1010
B	11	1011
C	12	1100
D	13	1101
E	14	1110
F	15	1111

The decimal system only has symbols for 10 numbers(0 to 9), after which we must use two symbols to represent numbers (00 to 99). After 100 we must use 3 symbols(000 to 999), and so on. Hexadecimal numbers are usually denoted with *0x* before the actual number. This is standard notation

although sometimes *h* is used. Since Hexadecimal has 16 symbols(0x0 to 0xF), any number above 0xF must be represented by at least 2 symbols, and any number above 0xFF must be represented by at least three, etc.

It's important to note that leading zeros don't change the actual value, but they do indicate the total range available. In the case above, a register which contains 0x000F is 16 bits (also known as 2 bytes, or 4 nibbles). However, the upper 3 nibbles (upper 12 bits) are all 0 either by default or because they've been changed.

Table 3.3 shows an extended list of hexadecimal numbers. Binary and Decimal are also shown for comparison. As you can see, Hexadecimal can simplify writing and prevent errors since any binary number can easily be represented using less symbols. 11111111 is nothing more than 8 bits, which can be represented more easily by 0xFF. In fact, going from 0x00 and 0xFF we can represent $2^8 = 256$ numbers. This is because each position in hex represent 4 bits (a nibble) and two of them represents 8 bits (a byte). This is a compact form and very convenient in programming. Some software calculators, such as those in Windows or Linux, can easily convert between numbers in different number systems and are a big help while programming. Notice that sometimes they will not accept leading zeros and these zeros must be accounted for when programming.

Table 3.3: Extended Hexadecimal vs. Binary and Decimal

Hexadecimal	Decimal	Binary
0	0	0
1	1	1
2	2	10
3	3	11
4	4	100
5	5	101
6	6	110
7	7	111
8	8	1000
9	9	1001
A	10	1010
B	11	1011
C	12	1100
D	13	1101
E	14	1110
F	15	1111
0x10	16	10000
0x11	17	10001
0x12	18	10010
0x13	19	10011
.	.	
.	.	
.	.	
0x1F	31	111111
0x20	32	100000

3.2.1 Conversion of Numbers

To convert a number from one base to the other, first ensure the calculator is in **Scientific Mode**. To convert a number simply press the radio button corresponding to the base you want to convert from, enter the number and press the radio button corresponding to the base you want the result in. The number is automatically converted. The radio buttons are shown in Figure 3.1.




Figure 3.1: Converting Numbers Using the Windows Calculator

3.3 Digital Operations

To be able to manipulate the registers of the MSP430, some knowledge of digital logic is necessary. The most basic operations we will need to be familiar with are AND, OR, XOR, and NOT and their truth tables are shown in Tables 3.4 and 3.5.

A	B	A B	A	B	A & B
0	0	0	0	0	0
0	1	1	0	1	0
1	0	1	1	0	0
1	1	1	1	1	1

(a) OR

(b) AND

Table 3.4: Digital Operations - OR and AND

A	\bar{A}
0	1
1	0

(a)
NOT

A	B	$A \wedge B$
0	0	0
0	1	1
1	0	1
1	1	0

(b) XOR

Table 3.5: Digital Operations - NOT and XOR

3.4 Manipulating Module Registers

We previously discussed variables, which are collections of bits that can be used to represent numbers (although usually you don't deal with the individual bits and use decimal). Most programmers do not have a problem understanding their use. If we have to store the results of an operation, we can easily assign it to a variable and use it later, regardless of what number system we use. Registers are very similar to variables in that they are also a collection of bits that can be read and modified by you. As opposed to variables, however, registers are variables internal to the microcontroller and its modules. They do not reside in RAM and therefore aren't allocated by the compiler.

There are two types of registers : **CPU Registers** and **Module Registers**. We will not discuss the CPU's registers that are used to store data and manipulate it. Since we are not programming in assembly, we do not need to concern ourselves with them as the compiler takes care of all of this. Rather, we are talking about Module Registers that control many of the peripherals in the microcontroller. The reason that these registers are important will become more and more apparent as you will read the tutorial.

In the MSP430, some registers are 8-bits (one byte) and some are 16-bits (2 bytes). We usually represent them visually by empty cells, with each cell capable of having a value of 1 or 0.


7	6	5	4	3	2	1	0

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

The numbers above each cell simply represent the position, with position seven representing the **Most Significant Bit** (MSB) and position 0 representing the **Least Significant Bit** (LSB).

The register is simply a collection of bits. Each of the little boxes can have 1 or 0 in it used to control something or represent data. How do we control the above registers? We do so by using the Digital Operations we discussed in section 3.3. Before we discuss this, we must cover the defines used by the MSP430 code which we will be referencing. We have all these registers in the MSP430 and the information about them can be easily found in the User's Guide of each MSP430. At the end of each module's description will be a reference of the registers that control that module, their name, and the purpose of each bit. Figure 3.2 shows an example of a register from a User's Guide:

ADC10CTL0, ADC10 Control Register 0



SREFx	Bits 15-13	Select reference 000 $V_{R+} = V_{CC}$ and $V_{R-} = V_{SS}$ 001 $V_{R+} = V_{REF+}$ and $V_{R-} = V_{SS}$ 010 $V_{R+} = V_{REF+}$ and $V_{R-} = V_{SS}$ 011 $V_{R+} = \text{Buffered } V_{REF+}$ and $V_{R-} = V_{SS}$ 100 $V_{R+} = V_{CC}$ and $V_{R-} = V_{REF+} / V_{REF-}$ 101 $V_{R+} = V_{REF+}$ and $V_{R-} = V_{REF+} / V_{REF-}$ 110 $V_{R+} = V_{REF+}$ and $V_{R-} = V_{REF+} / V_{REF-}$ 111 $V_{R+} = \text{Buffered } V_{REF+}$ and $V_{R-} = V_{REF+} / V_{REF-}$
ADC10 SHTx	Bits 12-11	ADC10 sample-and-hold time 00 4 x ADC10CLKs 01 8 x ADC10CLKs 10 16 x ADC10CLKs 11 64 x ADC10CLKs
ADC10SR	Bit 10	ADC10 sampling rate. This bit selects the reference buffer drive capability for the maximum sampling rate. Setting ADC10SR reduces the current consumption of the reference buffer. 0 Reference buffer supports up to ~200 ksp/s 1 Reference buffer supports up to ~50 ksp/s
REFOUT	Bit 9	Reference output 0 Reference output off 1 Reference output on
REFBURST	Bit 8	Reference burst. 0 Reference buffer on continuously 1 Reference buffer on only during sample-and-conversion

Figure 3.2: ADC10 Register Example

The register in this case is 16-bit. Visually, bits are grouped together when they share a common function. Therefore, Bits 15 to 13 are designated as SREFx. Notice that below each cell there is the designation rw. This means that you can both read and write to these bits using the manipulation techniques we'll discuss in the following section. It is possible for a register to be read only. This

is usually reserved for information about a particular module, indicating its status or something similar. Looking at the register description, we see a wealth of information. First of all, this register configures several aspects of the ADC such as the input channel (the physical pin), the sample and hold time, and the references.

A very important point to make is that the register is called ADC10CTL0. You might be asking yourself how can a register be named. In fact, the name is a convenient reference point. It's much easier to refer to it by ADC10CTL0 than by The register starting at 04Ah. Remember, in programming we normally refer to variables by their names and we don't directly deal with their actual address since the compiler does that for us. It is good programming practice to replace odd numbers with names since humans are much more comfortable with them (have you ever visited Amazon.com by going to it's IP number? Not likely). We haven't declared these registers, so how do we know their names and how does the compiler know? The compiler provides a header file which includes the defines for all registers in all peripherals and shortcuts to utilizing them. This header file is included in every program we use and makes our life much easier. It is unlikely we'll ever refer to the register by its position. Rather, we simply refer to it as ADC10CTL0. Another important point to make is all those names in the bit cells (such as SREFx, ADC10SHTx, etc are also defined and can be easily used to make code more readable. The header file specifically contains these definitions for SREF0 to SREF16.

We will leave aside ADC10CTL0 for now and will use another very useful register called P1OUT. The Input/Output (I/O) module contains registers to control the output of the microcontroller pins (to drive them HIGH or LOW). Bit 0 of P1OUT controls P1.0, bit 1 controls P1.1 and so forth. Suppose we want to make P1OUT equal to 0xFF, which would set all 8 bits in it to 1 and will cause all pins in Port 1 to go HIGH, then the register would look as follows:

7	6	5	4	3	2	1	0
1	1	1	1	1	1	1	1

The equivalent C code would be:

```
P1OUT = 0xFF;           // Set all the bits of P1OUT to 1
```

On the other hand we can also take the new register full of 1's and make it all zeros again (this will make all pins in port 1 to become LOW):

```
P1OUT = 0x00;           // Set all the bits of P1OUT to 0
```

Setting the register as above is just fine when we're changing all the bits, but it won't work if previously one of the bits was set to 1 or 0 and we want to leave it that way, (the whole register will be in effect erased) Perhaps another part of the program controls one of the bits in the register and if we use the assignment (=) symbol we will affect it adversely.

Lets suppose P1OUT is all 0 except position 0 (bit 0), which contains a 1, as follows:

If we want to set bit 1 of the register to 1, this would be represented as follows in binary:

7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	1

$$\begin{array}{r}
 00000001 \\
 \text{OR} \\
 00000010 \\
 = \\
 00000011
 \end{array}$$

Remember that OR is performed bit by bit, each position with the same position, so that position 0 is added to the other position 0, position 1 with 1, and so forth. Our previous 1 in position 0 still remains unaffected.

We could, by virtue of what we learned previously, replace the binary with hex notation, which would make it easier to read and write:

$$\begin{array}{r}
 0x01 \\
 \text{OR} \\
 0x02 \\
 = \\
 0x03
 \end{array}$$

Much easier to write obviously. We are still left with how to write this down in C so that the compiler would understand this is what we want to do. Remember, we are referring to P1OUT, which the compiler knows represents a register of 8 bits. Now in the MSP430 C programming, the addition used is the OR operand (represented as the pipe symbol, a vertical line):

```
P1OUT |= 0x02;
```

This results in adding the 1 to position 1, while leaving the register as it was elsewhere. It's important to note that if the position to which 1 is added is already 1, there is no overflow (because we're performing an OR). Therefore:

```
0x02 | 0x02 = 0x02;
```

The OR is not a summation. It is a bit by bit assignment as per its truth table.

It is fairly obvious that if we wanted to set something to 0 we can't just use the OR operator since adding a zero will leave the bit unchanged. In this case the AND operand (In C it is represented by the ampersand symbol) is used to set something to 0 since

```
0xFF AND 0x02 = 0x02
```

or

```
0xFF & 0x02 = 0x02
```

Table 3.6: Hexadecimal representation of position

Bit	7	6	5	4	3	2	1	0
1	0x80	0x40	0x20	0x10	0x08	0x04	0x02	0x01
0	0x7F	0xBF	0xDF	0xEF	0xF7	0xFB	0xFD	0xFE

0xFF represents 11111111, so we're not affecting anything. Using the combination of the OR and the AND we can control any bit/bits in the register.

Diagram in table 3.6 shows the value of each position. for example, a '1' in position 1 is represented by 0x02.

Using the table, if we wanted to add a '1' in position 7, we just need to OR the register with 0x80. If we wanted to put a 0 in position 7, we would AND the register with the inverse of 0x80 which is 0x7F (a 0 in position 7 and 1 in the rest). These can also be added together to control multiple bits. For example, we can add a bit in position 7 and 6 by doing an OR with 0x80 and 0x40:

```
P1OUT |= 0x80 + 0x40;
```

To make these operations easier, the MSP430 header includes the exact definitions for the bits up to 16 bits:

```
#define BIT0          (0x0001)
#define BIT1          (0x0002)
#define BIT2          (0x0004)
#define BIT3          (0x0008)
#define BIT4          (0x0010)
#define BIT5          (0x0020)
#define BIT6          (0x0040)
#define BIT7          (0x0080)
#define BIT8          (0x0100)
#define BIT9          (0x0200)
#define BITA          (0x0400)
#define BITB          (0x0800)
#define BITC          (0x1000)
#define BITD          (0x2000)
#define BITE          (0x4000)
#define BITF          (0x8000)
```

We can use these definitions so we don't have to memorize the hexadecimal notations. We can use them as follows:

```
P1OUT |= BIT0
```

The preprocessor (a program that runs before the compiler) goes through and replaces every instance of BIT0 with 0x0001. You don't see this happen (unless you look at some intermediate files). After the preprocessor runs, the above listing becomes:

```
P1OUT |= 0x0001;
```

The above definition for each uses 4 nibbles (4 positions or 2 bytes) because we are representing 16 bits. So we can perform the following:

```
P1OUT |= BIT4;    \\ Make P1.4 High(1)
```

This will turn on pin P1.4. To combine we can also do:

```
P1OUT |= BIT4 + BIT3;    \\ Make P1.4 and P1.3 High(1)
```

We can put a zero in the position without explicitly inverting. The **tilde** is the equivalent of the logical NOT operation in C:

```
P1OUT &= ~BIT4;    \\ Make P1.4 Low (0)
```

All this together makes register manipulation much easier. By the way, Binary can be input to the code by first putting 0b before the bits representing the number:

```
P1OUT |= 0b11111111;
```

3.4.1 XOR Operator

Finally, there exists also the XOR operator whose truth table is show below:

A	B	A ^ B
0	0	0
0	1	1
1	0	1
1	1	0

Table 3.7: XOR Truth Table

This table shows that if we always XOR with 1 we flip a bit from 0 to 1 or from 1 to 0 . In both cases, either when A is 0 or A is 1, XORing with 1 results in the bit flipping. This is convenient with elements such LEDs since it provides a toggle without needing to check what was the previous value.

```
P1OUT ^= BIT4;    \\ Toggle P1.4
```

XORing with 0 leaves the bit as it is and is therefore not very useful operation.

The operations discussed above form the basic elements for bit manipulation necessary to control all the registers of the microcontroller and its peripherals, even allowing communication with other devices connected using SPI,I²C or any other bus.

3.5 ASCII

Some knowledge of ASCII is useful since the UART information (communications with PC) can be interpreted as ASCII characters. ASCII is a character encoding scheme. This means that symbols such as letters and numbers are represented in a binary, hex, or decimal (which are related between themselves). Please refer to freely available ASCII tables which can show you, for example, that the number 0 is represented by 0x30. Therefore, if a 0x30 is sent to the PC using the UART module and the PC is in ASCII mode, it will show the number '0' on screen.

Table 3.8 shows some ASCII characters and their equivalents in Hex, Binary and Decimal:

Table 3.8: Extended Hexadecimal vs. Binary and Decimal

Glyph	Hexadecimal	Decimal	Binary
A	41	65	100 0001
B	42	66	100 0010
C	43	67	100 0011

For a more extended list, you can look up ASCII in Wikipedia or another other source. ASCII differentiates between each glyph (letter, characters, etc), and even between their lowercase and uppercase equivalents. Also, ASCII defines other elements such as space or newline.

The use of the UART to send ASCII characters and strings will be covered in the chapter on UART. A simple example is when we want to send the string "Hello World" we would send each of the characters, with the space between the words also sent. At the end of the line most likely we would send the newline character.

3.6 Conclusion

Now that we know how to modify registers, we must understand why and when to modify them. Much of this depends on the applications, but some general guidelines can be established that will enable you to do as needed for your application.

Chapter 4

Beginning Programming for MSP430

We will now begin discussing the basics of how to start programming the MSP430. The code will be based for the EZ430-RF2500 platform and be primarily geared towards the IAR compiler. It will likely run on CCE or other compilers, but some porting might be needed. As we go through the next chapters, we'll learn to integrate more and more of the modules into our design for a variety of applications. We can't possibly cover all uses of these modules but we will cover enough so that you will be able to leverage these ideas and further learn about the modules on your own.

A very basic piece of code that can be compiled and ran in an MSP430 microcontroller is:

Listing 4.1: MSP430 Hello World - hello_world.c

```
1 #include "msp430x22x4.h"
2
3 volatile unsigned int i;      // volatile to prevent optimization
4
5 void main(void)
6 {
7     WDTCTL = WDTPW + WDTHOLD; // Stop watchdog timer
8     P1DIR |= 0x01;           // Set P1.0 to output direction
9
10    for (;;)
11    {
12        P1OUT ^= 0x01;        // Toggle P1.0 using exclusive-OR
13        i = 50000;           // Delay
14        do (i--);
15        while (i != 0);
16    }
17 }
```

This code has some elements that were previously discussed and it is very simple. The first line includes the header file with definitions that simplify programming. We don't have to address P1DIR by its hexadecimal number (the address which points to it). Rather, we simply use P1DIR.

In this case we included msp430x22x4.h, but other header files will be needed for other MSP430 derivatives. The simplest way to see which one to use is to lookup the header file used in the example files of that MSP430. It is also possible to search the compiler directories for these files and guess given a particular derivative. Notice the letter x in the header name can be one or multiple numbers or a letter (commonly F for Flash).

Following this, there is a variable declaration. `i` is declared to be unsigned int. The volatile modifier tells the compiler not to eliminate it in its optimizations routines. If a variable is declared and never modified in `main()` the compiler will likely eliminate it because it doesn't know about an interrupt routine accessing it.

Now comes the main routine of the code. `main()` is common to among all programs. When code execution begins (when debugging for example), the first line to be executed is the first line in `main()`. Something that might not be obvious is what is actually set when the program begins. That is, are any clocks running before we execute the code for setting the speed?

The answer is that it must be running. The CPU cannot operate without a clock. However, what is the speed and who sets it? When the MSP430 starts, it has defaults for many of the peripherals, including clocks, I/O, and others. Check the Datasheet and user's guide for information about the defaults. The Master Clock (MCLK), which is the source clock of the CPU, starts up at a frequency specified in the datasheet. . This is done behind the scenes, but is very important. Although we won't address it yet, it is important to change the clock frequency and make sure it is adapted to our application and ensure that all clocks and crystals are running correctly.

One of these is the Watchdog timer, which must be disabled or else the microcontroller will restart endlessly. The first line of `main()` stops the Watchdog timer. This is a line that will be standard in all MSP430 programs, unless you will adopt something else (such as using the Watchdog timer for another purpose).

A possible problem that is not apparent at first should be discussed. What happens when we require a lot of initialization time before `main()` starts? This can happen when we, for example, declare a very large array to be used in our program. The issue is that the Watchdog timer will expire before initialization has ended and this will restart the system, in an endless loop. This can happen if we have a lot of memory, greater than 2kB, and we are declaring an array with many elements.

To avoid this problem, we can declare a function for low level initialization, as follows:

```
1 void __low_level_init(void)
2 {
3     WDTCTL = WDTPW+WDTHOLD;
4 }
```

This code will run before any variable initialization and will stop the Watchdog regardless of any initialization. However, this function is not available in all compilers. IAR and CCE do include it.

The rest of code in Listing 4.1 is simply meant to flash an LED connected to P1.0. The first step in this is to change the direction of pin 0 in port 1 to an output. This only needs to be done once. We then include an infinite loop. There is no condition for breaking so it continues forever. We could have easily used

```
1 while(1)
2 {
3     P1OUT ^= 0x01;    // Toggle P1.0 using exclusive-OR
4     i = 50000;       // Delay
5     do (i--);
6     while (i != 0);
7 }
```

The only way then to stop the loop is to use the break command. The rest of the code is simple. We XOR the value that is in position 0 of P1OUT with one, toggling it(from 1 to 0 and from 0 to

1). We then create a delay with a do... while loop. The timing of the delay is not always easy to establish (unless steps are taken to measure the number of cycles it takes). We simply have a loop that goes 50000 times and executes nothing (except the check for the condition). Once this loop is done, we begin again with the XOR of P1OUT.

Usually, when we write an MSP430 application we do as follows:

1. Stop the Watchdog Timer
2. Setup the Oscillators and clocks and check for their correct operation using the appropriate handlers
3. Setup The I/O pins
4. Setup the rest of the modules to be used
5. start the application specific code

Information about the Watchdog timer is included in the User's Guide. The next chapter will add upon the elements we will need to create a fully functional application.

Chapter 5

MSP430 Clocks

5.1 Introduction

Although most users of Personal Computers would not be aware of this, clocks are at the heart of any synchronous digital system. CPUs require clocks to run the CPU since asynchronous operation is not possible in computer processing (it presents many difficulties such as when the data to be processed comes from different places at different speeds). In PCs, the selection of clock speeds is determined by various factors. Unless you are overclocking, you will never deal with them directly. Microcontrollers, on the other hand, are usually very flexible with clocks and require that the designer specify what clocks will be used and at what speeds. Usually this means that both hardware and software aspects of clocks be considered during the design stage.

As an example, the MSP430 accepts a low frequency crystal (typically 32.768Khz), which must be added externally. Upon initializing the microcontroller, the clock system of the MSP430 must be configured to take advantage of this clock. In the EZ430-RF2500, no crystal is available and the only oscillator used is the internal oscillator. This saves space but clock precision is reduced.

Why are clocks important? One important effect is on the frequency of the CPU. The speed at which it runs depends on a clock called MCLK. Because of this, the speed of instruction execution will depend on the clock. Do you have an application which needs to move data fast? The slow default speed might be too slow for your needs. On the other hand, the power consumption of the microcontroller is very dependent on the CPU speed. Although peripherals do consume current, the CPU, when running, is in most cases the major offender. Current consumption usually varies linearly with clock speed and therefore one way to keep consumption to a minimum is to set the clock speed as low as possible (or turn it off completely when not needed). MSP430 Microcontrollers generally use two categories of clocks, fast and slow. The fast clocks source the CPU and most modules and varies usually from several hundred kHz to Several MHz (25MHz currently in the new MSP430F5xx family). The slow clocks belong to the low kHz range and are usually either 12kHz or 32.768kHz.

How do we generate the clocks? There are many ways to do so, but in the MSP430 (and many other microcontrollers) there are generally three types of clock sources:

- Internal Oscillators

- External Crystals
- External Oscillators

5.2 Internal Oscillators

Internal Oscillators are usually an RC network with circuitry to try and improve the accuracy of the clock. Trimming and Calibration by the manufacturer can significantly improve the precision of the clock. Remember, however, that temperature does affect the clock frequency and if this dependency isn't reduced by some sort of calibration, it can wreak havoc on an application. Therefore, it is usually important to select The benefit of this type of oscillators is that their frequency can be easily changed and they don't occupy any more space on the PCB. On the MSP430, A fast Digitally Controlled Oscillator (DCO) oscillator is available, with another slow 12kHz oscillator optional and usually available in newer MSP430 variants.

5.3 External Crystals

External Crystals add a large measure of accuracy to oscillators and should be used as much as possible unless cost and area considerations are more important. They are not tunable but with PLL or other circuitry can generate other frequencies based on the crystal.

The selection and use of crystals is not always so straightforward because they are specified using several parameters. Aside from the frequency, accuracy, drift and aging factors, the most overlooked part of adding a crystal to the circuit is the loading capacitance. A detailed explanation for the need and purpose of the loading capacitance is available from most crystal manufacturers.

In general, every crystal used in a parallel configuration, which is the case in the MSP430 and many other microcontrollers, requires capacitors to provide the capacitance that is necessary for the oscillator to oscillate at the right frequency. These capacitors are usually connected to the two crystal pins as follows:




Figure 5.1: Loading Capacitors on Crystal

The effective loading required is specified by the crystal manufacturer, not by the microcontroller manufacturer (although the microcontroller manufacturer could make recommendations or force you to use a certain loading because of built in capacitors). This loading is what the crystal needs to oscillate at the frequency it is specified at. Any change in the effective loading changes the crystal's effective frequency, affecting the system using the crystal. The loading capacitors provide this capacitance, but they're not the only ones. Every pair of parallel plates with voltage across

them forms a capacitor. Two parallel pins such as microcontroller pins have voltage and therefore introduce parasitic capacitance from the circuit. Internal microcontroller circuitry does the same because of the way the oscillator is constructed. The PCB on which the microcontroller is mounted also adds parasitic capacitance. From this you should realize that to get the effective loading capacitance you must take into account all these extra parasitics. Some of these are specified in the datasheet and others are usually assumed to be in a certain range.

Lets begin a simple exercise by assuming there are no parasitics. If we assume that we selected a 4MHz crystal with a parasitic capacitance of 7pF. We need the combination of C1 and C2, which are usually selected to have the same value, to give a loading of 7pF.

$$C_{Leff} = \frac{C1 * C2}{C1 + C2} = \frac{C1 * C2}{C1 + C2}$$

Therefore if C1 = C2, we have

$$C_{Leff} = \frac{C1^2}{2C1} = \frac{C1}{2}$$

$$C1 = 2 * C_{Leff} = 2 * 7 = 14pF$$

Therefore, both C1 and C2 should be 14pF to give 7pF of loading capacitance.

As was previously mentioned, this does not take into account the parasitics and the microcontroller effect. The oscillator inside the microcontroller introduces two capacitances generally called C_{XIN} and C_{XOUT} . Each appears in parallel with one of the external loading capacitors, resulting in the effective load capacitance of:

$$C_{Leff} = \frac{(C1 + C_{XIN}) * (C2 + C_{XOUT})}{(C1 + C_{XIN}) + (C2 + C_{XOUT})}$$

We are not finished, however. The PCB itself and the pins introduce a parasitic capacitance that should be taken into account as well. This parasitic capacitance appears in parallel to the previous parasitic capacitance and therefore is simply added (as is done for parallel capacitors):

$$C_{Leff} = \frac{(C1 + C_{XIN}) * (C2 + C_{XOUT})}{(C1 + C_{XIN}) + (C2 + C_{XOUT})} + C_{par}$$

PCB Parasitic capacitances are not always easy to measure. Because of this, it is common to assume a PCB parasitic capacitance of 2-5 pF.

It is worth the note that in the work above we assumed that we must provide external capacitors to load the crystal. This is not always true. The MSP430 accepts a 32kHz crystal for low power application. If you look at several designs with the 32kHz crystal you will notice there are no loading crystals. In these cases the MSP430 provides the loading capacitance internally to save space. The internal capacitance provided can be fixed, as is the case with MSP430 derivatives such as the MSP430F1611, or they can be software selectable from among several values. For example, the MSP430F1611 datasheet page 38 specifies C_{XIN} and C_{XOUT} , the Integrated input capacitances, to be both 12pF, meaning that the effective loading capacitance is 6pF and the crystal

required should have a loading capacitance as close to 6pF as possible if parasitics are not taken into account.

Note that a 32 kHz crystal actually has a 32.768 kHz frequency (which is useful because 32768 is a power of 2, the base for binary numbers). The 32 kHz crystal to be used is usually of the watch crystal type. Mouser and Digikey have several acceptable crystals available in several packages. Cost for this crystal can be \$0.30, and will usually be less than \$1.

References:

MSP430 32-kHz Crystal Oscillators Application Note - slaa322b

Crystal Oscillators Using MSP430

Pierce-gate oscillator crystal load calculation

What is frequency at load capacitance?

5.4 Clock Sources

Clock sources are the term encompassing both crystals and internal oscillators. Their availability depends on your particular MSP430 derivative. For example, The MSP430F2274 does not include the XT2CLK (probably due to the low number of pins in the package). The following is a summary of the nomenclature of the clock sources:

- LFXT1CLK: Low-frequency/high-frequency oscillator that can be used with low-frequency watch crystals or external clock sources of 32,768-Hz. or with standard crystals, resonators, or external clock sources in the 400-kHz to 16-MHz range.
- VLOCLK: Internal low frequency oscillator with 12-kHz nominal frequency. Its low frequency means very low power.
- DCOCLK: Internal digitally controlled oscillator (DCO).
- XT2CLK: Optional high-frequency oscillator that can be used with standard crystals, resonators, or external clock sources in the 400-kHz to 16-MHz range.

All these sources give you great flexibility. You could avoid using crystals altogether by using the internal DCO and VLO, or, if you need more precision, use the external crystals at the expense of PCB space and some money. It is standard practice to use LFXT1 with a 32.768 kHz crystal, leaving XT2 to be used with a high frequency crystal.

5.5 Clock Signals

The clocks available are actually the last element in a long chain. That is, first we begin with the clock generator, followed by some circuitry which further helps improve clock performance and change its parameters (such as dividing the frequency by some constant). Finally, clock distribution in the MSP430 allows for different clock sources. As an example, a Timer Module in the MSP430

can be source from either the fast clock or the slow one. The CPU, however, can only be sourced from the fast clock called MCLK. The User's Guide section for each peripheral will specify what are the allowable clock sources.

Note that the flexibility of the Clocks is limited. In the MSP430F2274 we have 3 clock signals available to CPU and Modules:

- **ACLK:** Auxiliary clock. ACLK is software selectable as LFXT1CLK or VLOCLK. ACLK is divided by 1, 2, 4, or 8. ACLK is software selectable for individual peripheral modules.
- **MCLK:** Master clock. MCLK is software selectable as LFXT1CLK, VLOCLK, XT2CLK (if available on-chip), or DCOCLK. MCLK is divided by 1, 2, 4, or 8. MCLK is used by the CPU and system.
- **SMCLK:** Sub-main clock. SMCLK is software selectable as LFXT1CLK, VLOCLK, XT2CLK (if available on-chip), or DCOCLK. SMCLK is divided by 1, 2, 4, or 8. SMCLK is software selectable for individual peripheral modules.

The CPU is always sourced by MCLK. Other peripherals are sourced by either SMCLK, MCLK, and ACLK. It is important to consider all the modules in the system and their frequency requirements to select their source. ADC sampling at high speeds cannot be done with ACLK.

5.6 Basic Clock Module In EZ430-RF2500

Since the EZ430-RF2500 uses an MSP430F2274, we must refer to MSP430F2274 documents, especially example code and User's Guide. An important feature of the MSP430F2274 is that Texas Instruments has included very useful calibration constants that reside in the Information Flash Memory (It is similar to where the program code resides but it exists to allow the user to store things like calibration constants and not get erased when we erase code and flash it).

The first consideration is to check which Clock Sources we have for the EZ430-RF2500. Inspecting the board (and checking the schematic) shows that there is in fact no crystal for the MSP430F2274. Therefore, all our sources for the Basic Clock Module are internal oscillators:

- **DCOCLK** - Internal High Speed Oscillator up to 16MHz
- **VLOCLK** - Very Low Frequency (12kHz) oscillator

The DCOCLK is the main clock for the system and supplies the CPU. The VLOCLK is extremely useful when our application goes to one of the sleep modes (which we discuss in 9.1). This clock can maintain timers for waking up the system at a certain point in the future. Selecting the frequency for DCOCLK is a matter of changing the settings of some registers, as shown in Figure 5.2.

As you can see, It is difficult to select the correct frequency precisely. To avoid issues, TI provides a set of calibrated constants that have a 1% deviation from the desired frequency.

A general function to initialize DCOCLK to 8MHz and enable the VLOCLK is provided next.




Figure 5.2: Selecting the DCO Frequency

Listing 5.1: Configuring EZ430 Clocks

```

1 void configureClocks()
2 {
3     // Set system DCO to 8MHz
4     BCSCTL1 = CALBC1_8MHZ;
5     DCOCTL = CALDCO_8MHZ;
6
7     // Set LFXT1 to the VLO @ 12kHz
8     BCSCTL3 |= LFXT1S_2;
9 }

```

The following are the list of all calibration constant provided in the Flash:

```

CALDCO_16MHZ
CALDCO_12MHZ
CALDCO_8MHZ
CALDCO_1MHZ

```

Although these are only 4 constants (they don't include other possibly useful frequencies), they will suffice for many applications.

Listing 5.2: MSP430 Hello World with Clock Control

```

1 #include "msp430x22x4.h"
2
3 void configureClocks();
4 volatile unsigned int i; // volatile to prevent optimization
5
6 void main(void)
7 {
8     WDTCTL = WDTPW + WDTHOLD; // Stop watchdog timer
9     configureClocks();
10    P1DIR |= 0x01; // Set P1.0 to output direction
11
12    for (;;)
13    {
14        P1OUT ^= 0x01; // Toggle P1.0 using exclusive-OR
15        i = 50000; // Delay
16        do (i--);
17        while (i != 0);
18    }

```

```
19 }
20
21 void configureClocks()
22 {
23     // Set system DCO to 8MHz
24     BCSCTL1 = CALBC1_8MHZ;
25     DCOCTL = CALDCO_8MHZ;
26
27     // Set LFXT1 to the VLO @ 12kHz
28     BCSCTL3 |= LFXT1S_2;
29 }
```

Listing 5.2 shows the Hello World application with the added clock routine. If you run this code you will immediately notice that the LED will flash much faster. The reason for this is that the default DCO clock frequency is much lower than 8MHz. It is the CPU which tells pin 1.0 to toggle, and so the speed of the CPU determines the toggle speed (although there is a small delay because of the I/O Module and other things calls). The function itself adds some overhead.

8MHz will be the standard speed we will use in our EZ430-RF2500 application. This is fast enough to run most things, and is good for communicating with the CC2500. If you look at the CC2500 datasheet, you'll see that they specify the maximum clock frequency possible. The 8MHz clock will source the SPI module that communicates with the radio and although it can be divided further to reduce the clock speed to comply with the required specifications if we have a faster clock, 8MHz will work without any modifications.

5.7 Considerations for using clocks

Clocks are among the most critical aspects of microcontrollers and the MSP430 because they sometimes require significant engineering effort to use. This is especially true for battery powered applications where we don't have the luxury of an electrical outlet. After all, these are the kind of applications for which the MSP430 was designed.

As previously mentioned, there is a direct tradeoff between clock speed and energy consumption. During the development process a bottom-up and top-down approaches are usually used. We start the initial development at a limited clock speed and scale up as necessary to achieve each one of our tasks. On the other hand, the top-down approach begins with a fast clock and scales the clock down after functionality is proven in an attempt to conserve energy. In many cases, it is immediately apparent where we should begin. A hand held watch would require a bottom-up approach, while audio sampling would require top-down. In many cases, the use of modules and external devices complicates the clock selection process because of conflicting requirements. The flexible clock system of the MSP430 simplifies things, but tradeoffs will still be required.

Although this topic will be covered in later chapters, energy can be saved by turning clocks off. This will prevent us from using any modules for which the clock has been turned off, but should be done in cases where it's advantageous.

Chapter 6

General Purpose Input Output - GPIO

Digital pins are the main communication element that a microcontroller has with the outside world. They can go HIGH (1) or LOW(0) depending primarily on the system voltage (VCC). The datasheet for your particular MSP430 will provide details about the voltage swing used to represent these two positions, and they are usually several hundred millivolts from the system voltage.

6.0.1 Pin Multiplexing




Figure 6.1: MSP430F2274 Schematic

Looking at the schematic representation of the MSP430 (MSP430F2274 in Figure 6.1) you will notice that each pin has a long name, with several designations separated by a slash. Because microcontrollers have a limited number of pins, the manufacturer has to multiplex the pins among the internal modules. That is, each pin can only do one thing at a time and you have to select what that function will be (upon startup there are some defaults which you will likely override with your code). Multiplexing can become a problem if you're in short supply of I/O pins and you're using

others for specialized functions. In your design stage you should take all of this into account and plan which pins will be used (and with what functionality).

Most pins on the MSP430 devices include GPIO functionality. In their naming, this is indicated by PX.Y, where the X represents the port number and the Y the pin number in that port. MSP430 devices have several different ports, depending on the actual device, and these ports each have 8 pins (making byte access easy). For example, Port 1 is represented by pins P1.0 P1.1, P1.2, P1.3, ..., P1.7.

Each port is assigned several 8 bit registers that control them and provide information about their current status. Each register is designated PxYYY, where the lowercase x represents the port number and YYY represent the name of the functionality of the register.

- PxSEL - This register selects what functionality is shown at the pin, GPIO or an internal Module
- PxDIR - If the pin is set for GPIO in PxSEL, this register selects the pin direction - Input or Output
- PxOUT - If the pin is set for GPIO in PxSEL and If the pin has output direction as set by PxDIR, this selects whether it is HIGH or LOW
- PxIN - If the pin is set for GPIO in PxSEL and has input direction as set by PxDIR, this represents the level at the input (HIGH or LOW)
- PxIFG - Interrupt Flag for the corresponding I/O pin and it is set when that pin experiences the appropriate signal edge transition(either high to low or low to high).
- PxIES -Interrupt transition selection register. If a bit is 0, the pin will generate an interrupt and a flag is the transition is low-to-high. A 1 in the bit will generate an interrupt and a flag if the transition is high-to-low.
- PxIES - Interrupt enable register. If a bit is set to 1, a pin will generate an interrupt if the appropriate transition has occurred.

In addition, for those MSP430 devices that support it, there are registers related to pullup/pulldown resistors:

- PxREN - Internal Pullup/pulldown resistor is enabled for the specified pin if a bit is set to 1. A 0 disables the resistor.
- PxOUT - If a pins pullup/down resistor is enabled, the corresponding bit in the PxOUT register selects pull-up(bit = 1) or pull-down(bit = 0).

Critical Note - In almost all MSP430 devices I've seen, only ports 1 and 2 have interrupt capability. Make sure to take this into account when designing the circuit. Following the convention, Port 1 has the registers P1SEL, P1DIR, P1OUT, and P1IN, among others. P2OUT controls the output level of each pin (1 or 0) in the port because they are mapped as follows:

The 8 bits each can be 1 (HIGH) or 0 (LOW), representing VCC or 0 (or close to it). The port is controlled by P4OUT. If, for example, we wished to turn on pin 0 of the port HIGH, i.e. P4.0, we would need to set :

BIT 7	P2.7
BIT 6	P2.6
BIT 5	P2.5
BIT 4	P2.4
BIT 3	P2.3
BIT 2	P2.2
BIT 1	P2.1
BIT 0	P2.0

```
P2OUT |= 0x01; // P4.0 High
```

This would set bit 0 to 1 and therefore VCC would appear on that pin. The grouping of the port makes it easy to control 8 pins. For example, if an IC were connected to the port, it could be used to seamlessly transfer bytes since the input register of Port 2 (P4IN) can be read as a whole byte and no bit manipulation is needed there. It would appear obvious that if we can set the pins to HIGH or LOW, there can be no input and therefore I/O (input output) is only Input at that moment. Each pin on the microcontroller is actually multiplexed. That is, each pin can be both input or output, but it can only be one of them at any one point in time. The selection of functionality is done using the PDIR register, where X represents the port. In PDIR a 0 indicates an input while a 1 indicates an output. Therefore, if we wanted to set pin 0 of Port 4 to be an input we would need to do:

```
P4DIR &= ~BIT0;
```

This leaves all the bits alone while making bit 0 to be 0. There is another level of complexity. Every pin can also have a secondary function besides I/O. This is microcontroller dependent and specified in the datasheet of the specific microcontroller. The complexity of the pin is shown hierarchically as follows:

First, we select between I/O and secondary function using PSEL. Then we select between Input and Output using PDIR. If we selected input, P4IN will contain the binary representation of the voltages at each pin in the port. If we selected output, P4OUT will enable setting the pins of each port to the corresponding HIGH or LOW, depending on the application.

Remember, these registers have start up defaults and they are specified in the data sheet.

For GPIO pins, there are several other important Registers. Some of these deal with interrupts and this topic will be discussed later in Chapter 9.

PxREN - This register controls the internal pullup and pulldown resistors of each pin. These can be important when using switches. Instead of using an external pullup resistor, we can simply configure everything internally, saving space. The corresponding bit in the PxOUT register selects if the pin is pulled up or pulled down.

In PxREN, a 0 in a bit means that the pullup/pulldown resistor is disabled and a 1 in any bit signifies that the pullup/pulldown resistor is enabled. Once you select any pin to have the pullup/pulldown resistor active (by putting a 1 in the corresponding bit), you must use PxOUT to select whether the pin is pulled down (a 0 in the corresponding bit) or pulled up (a 1 in the corresponding bit).

The pullup/pulldown is useful in many situations, and one of the most common one is switches connected to the MSP430.

6.1 Switches

Adding switches to a microcontroller is one of the most popular ways to communicate with it. Many situations require you to let the microcontroller perform some operation. It's relatively simple to do as well.

Consider first that your ultimate objective is to generate a clean high to low or low to high transition. The switch at an initial moment is either high or low and the pressing of the momentary button (you could use a different kind of switch as well) results in a change in the voltage level.




Figure 6.2: Connecting Switches to MSP430

Figures above show how a switch is connected to P2.0 (a GPIO with interrupt capability) of the MSP430. Let's think a bit about what happens when the system is powered up, as in case 6.2(a). Here, the switch is normally open and the P2.0 is left floating and this is very bad. When the switch is pressed, P2.0 is connected to ground. There is no clear voltage transition because the initial condition is floating. Similar designs have been used by many newcomers and should be avoided. A design such as that is possible if there is an internal pull-up resistor (more about this later).

Let's analyze the second case 6.2(b), where the designer included an external pullup resistor. Here, the switch is also open initially, but P2.0 is connected to VCC through the resistor. The resistor limits the current significantly (considering a resistor with 100k or so). When the user presses the switch, one end of the resistor and P2.0 get connected to ground. P2.0 experiences a transition from high to low. The resistor has current flowing through it and although it is wasted as heat, it is very small and lasts for only a few moments for the duration of the switch press.

Many microcontrollers include internal pull-up or even pull-down, which free the designer from having to include them. In this case, simply connect the button to either VCC and GND and it should work.

6.1.1 Debouncing

Up to now, we've assumed that the switch is perfect. In reality that is not so. The mechanical construction of the switch results in internal vibrations (bouncing) in which the metals bounce until they settle and make contact. During this time, a switch can result in many connections/disconnections that a microcontroller will interpret as multiple presses when in fact a user only pressed once. The end result is that multiple interrupts will be generated and can cause issues in the end application.

What to do? In general there are two approaches to doing debouncing: Analog and Software.

The analog approach uses a capacitor and a schmitt trigger. When a capacitor is added to the switch input, the capacitor is pulled high by the pull-up resistor and therefore is charged. When the key is pressed, the capacitor will resist the change in voltage (the switch is grounding the input). The capacitor will discharge, preventing the bouncing of the switch contact. This depends heavily on the size of the capacitor and resistor (on their RC time constant). The larger the RC time constant, the more bouncing will be reduced. However, the switches responsiveness will be affected. When the user releases the button, the capacitor slowly charges and therefore any bounces are also filtered. However, because of the slow transition, the I/O input will transition in an unknown state which could cause issues. A schmitt trigger is then required to avoid any problems. This is typically not an issue because the MSP430 includes a schmitt trigger at the I/O inputs.

The other solution that can be done is in software. Using interrupts we can detect the switches high to low and low to high transitions and ignore those that are inside of the bounce window by waiting for a set amount of time (using a timer or a for loop to waste cycles) before accepting new transitions as actual user switch presses. The amount of time to wait depends on the switch to be used, but 40ms is a safe bet in many cases. If you can't wait that long, measure the bouncing time using an oscilloscope and create a worse case estimate. We must also ensure that the first bounce of the switch was not a momentary malfunction (as could be the case if the switch is malfunctioning or the board was dropped creating an accidental contact).

What do I recommend? Do both. Adding a capacitor is easy and you can always not populate it if needed. It provides some improvement that can be enhanced with a software approach.

6.2 LEDs

Another important GPIO use is to control LEDs as indicators. The first way to perform this is to simply connect the led to the GPIO and drive it using the pin. A resistor of an appropriate value should be added in series with the LED to limit the current because the microcontroller will source (provide) the current and its capabilities in this respect are limited. Typical LEDs require on the order of tens of milliamperes. For low power microcontrollers this figure reaches the upper limit of the current sourcing capability of the microcontroller. Because this current is shared among the GPIO, adding LEDs can limit the drive capability of I/O. This might require buffering and complicate the circuitry. A solution commonly used is to connect the LED in reverse, as shown in figure ???. Here, the microcontroller is used to sink the current, which it is usually able to

do more readily. Unfortunately TI does not provide much information in this respect and some experimentation is in order.

The value of the series resistance varies depending on the current required to drive the LED and the VCC of the system. Typical values are 220, 330 and 470 Ohm.

To calculate the resistance required, the calculation is a simple application of Ohms and Kirchoff's laws. Assume that the voltage across the diode is the forward voltage specified in the datasheet. Assume also the forward current specified to obtain a specific luminosity of the LED. Then you can calculate the voltage drop across the resistor and its resistance.

6.3 Bit banging

Bit banging refers to the implementation of a digital communications protocol from GPIO primitives. In other words, using a predefined sequence of high and low on the GPIO, we can emulate a protocol such as SPI. When dedicated hardware is unavailable (such as when you've used up all possible modules that are available), bit banging may be one of few choices left without resorting to using a different device. Because the implementation is purely software, it is very inexpensive. However, it has several drawbacks, primarily its limited .

As you previously learned, each module on the MSP430 requires a clock which is used for timing. When you perform bit banging, the CPU is manually modifying the GPIO registers to perform the bit banging. Therefore, for each bit to be sent, several cycles are required. This is in contrast to the mostly single cycle nature of dedicated hardware. In effect, the rate at which bit banging can be performed is several times slower than dedicated hardware and results in increased latency.

Another major issue is the processing time required from the CPU. Instead of working on more useful things, the CPU spends its time communicating. This takes a huge toll, especially for large amounts of data.

It is advisable to stay away from bit banging as much as possible and only to use it as a last resort or if latency is not critical in the application. One application in which I've adopted bit banging is reading the serial ID from a one wire IC. Because this is only necessary at bootup, there's no load on the system while it's running and it simplifies the hardware design significantly.

Chapter 7

Flash Memory

7.1 Introduction

Flash memory is the most popular type of non-volatile storage, and it has been adopted by most microcontroller manufacturers. The ability to read/write hunder of thousands of times, even while the processor is running, allows a degree of flexibility not previously available with the older ROM and PROM technologies.

The flash memory in the MSP430 is used primarily to hold code. Because of the flexibility in its design, it is possible to write and read the flash at run time with the CPU. This allows you, for example, to reprogram the microcontroller with a new firmware image. It also enables you to store information that will survive a restart, making it an easy and inexpensive alternative to external EEPROM devices.

7.2 Flash and Memory Organization

The Flash memory in the MSP430 is usually divided into two sections, the main flash and the Information flash:

- Main Flash - Represents the bulk of the MSP430's flash and used for program code and constant data. Itself divided into several flash banks, with each bank further divided into segments. When a microcontroller is denoted to have "32kB" of Flash, it is referring to the Main flash.
- Information Flash - Several extra segments of flash separate from the Main flash. These are primarily intended for information such as calibration constants, but are much smaller than the main banks.

The location and amount of each of these sections depends on the specific MSP430 you're using, and the details are contained in the Memory Map included in the datasheet of the partical MSP430 you're using. Here is an example of the memory organization of several MSP430 devices, taken from their datasheets. Only Flash sections are included. This information was taken directly from

the datasheet of the devices. Please note that some details such as the breakdown of the flash segments was obtained from the linker file since the datasheet lacked these specifics. The use of the linker will be discussed later.

		MSP430F5438A	MSP430F2274	MSP430FG4618
Main Flash		256kB 045BFFh-005C00h	32kB 0FFFFh-08000h	116KB 01FFFFh - 003100h
Main Code Memory	Bank 3	64 KB 03FFFFh-030000h	N/A	N/A
	Bank 2	64 KB 02FFFFh-020000h		
	Bank 1	64 KB 01FFFFh-010000h		
	Bank 0	64 KB 0045BFFh-0040000h 00FFFFh-005C00h		
Info Memory	Info A	128B 0019FFh-001980h	64 Byte 10C0-10FF	64 Byte 1080-10FF
	Info B	128B 00197Fh-001900h	64 Byte 1080-10BF	64 Byte 1000-107F
	Info C	128B 0018FFh-001890h	64 Byte 1040-107F	None
	Info D	128B 00187Fh-001800h	64 Byte 1000-103F	None

Figure 7.1: MSP430 Memory Organization Examples

We can see that for the FG4618 device, we have a total main flash size of:

$$01FFFFh - 003100h = 1CEFFh = 118527 \text{ bytes}$$

This calculation is easily done use the Windows Calculator in hex mode (the h after each number denotes it is in hex). After the subtraction, convert 1CEFF to decimal.

7.3 Looking at the Flash

Before we get into the details of using the Flash, it is best to understand how to look at it. Knowing how to look at the flash of the MSP430 using the debugger is perhaps the most useful flash debugging technique. Using a debugger such as IAR Workbench or Code Composer Studio, it is possible to view the actual contents of the flash memory at runtime when halted (they cannot be viewed while the MSP430 is running)

Listing 7.1: flash_view.c - Software to view Flash

```

1  //*****
2  //  MSP430 Flash Viewing
3  //
4  //  MSP430F2274
5  //  Description: Running this software using CCS or IAR you can see
6  //                some of the behavior of Flash.
7  //
8  //  Gustavo Litovsky
9  //
10 //  Built with CCS Version 5.1 and IAR Embedded Workbench Version: 5.20
11 //*****
12
13 #include <msp430.h>
14
15 void main(void)
16 {
17     char *Flash_ptr = (char *)0x1040;    // Flash pointer
18
19     WDTCTL = WDTPW + WDTHOLD;           // Stop watchdog timer
20     BCSCTL1 = CALBC1_1MHZ;              // Set DCO to 1MHz
21     DCOCTL = CALDCO_1MHZ;
22     FCTL2 = FWKEY + FSSEL0 + FN1;       // MCLK/3 for Flash Timing Generator
23
24     __no_operation();                   // Place breakpoint here, observe 0x1040
25     FCTL3 = FWKEY;                       // Clear Lock bit
26     FCTL1 = FWKEY + WRT;                 // Set WRT bit for write operation
27
28     *Flash_ptr = 0xAA;                   // Write value to flash
29     __no_operation();                   // Place breakpoint here, observe 0x1040
30     *Flash_ptr = 0xFF;                   // Write value to flash. Fails because
31                                         // we can't set bits to '1'
32
33     __no_operation();                   // Place breakpoint here, observe 0x1040
34     *Flash_ptr = 0x00;                   // Write value to flash. All bits set to 0
35     __no_operation();                   // Place breakpoint here, observe 0x1040 with 0x00
36
37     // Ensure we don't exit from program
38     while(1)
39     {
40         __no_operation();
41     }
42 }

```

Add the code to a project, download and debug on a MSP430 board. Place a breakpoint on the `__no_operation()` lines if your MSP430 has enough breakpoints, or else as you go through the code, move the breakpoint forward as needed. To view the Flash Memory contents in IAR, you must be currently in an active debug session. Then, go to View then Memory. A new window will show. In the go to field type "0x1040" and press ENTER. The window should show memory locations at 0x1040 and around it. Going from left to right, each byte represents 0x1040 to 0x104F.

On the first breakpoint, you will observe that the flash memory at 0x1040 has the value 0xFF because during download the MSP430 was erased (although the Info flash might not have been erased, it is likely in the erased state).




Figure 7.2: MSP430 Memory after writing 0xAA in IAR




Figure 7.3: MSP430 Memory after writing 0xFF in IAR

Run the code again and break on the the second breakpoint. You should see that the flash contents has changed to 0xAA because we wrote it.




Figure 7.4: MSP430 Memory after writing 0x00 in IAR

After running the code again on the 3rd breakpoint you will see that 0xAA has remained, despite the fact that we attempted to set bits to '1' and that is not possible without an erase operation which the code doesn't perform.

Finally on the 4th breakpoint we will see that the flash contents are 0x00.

In CCS, the procedure is similar. Once in the debug session, go to View, then "Memory Browser". Then type "0x1040" and Press Go:

7.4 Programming the Flash

Although reading the flash is very easy, reprogramming it is not as straightforward. Because of the way flash is built, by default it stores a '1' bit in each cell. If you were to erase the flash, by default it will return to what we would consider a '1'. When we reprogram the flash, all that's being done is to change '1' in certain bits to '0'. You cannot change a '0' to a '1' except by using an erase procedure. This was the reason why in the previous example, attempting to set the byte to 0xFF did not work, since it required setting all bits back to '1'. This is the main difference between flash and other technologies such as EEPROM and a main source of confusion.

In general, writing to flash will require an erase. This is unless it was previously erased or we're lucky enough so that the bits we want to modify are all '1' and we want to make them a '0'.

Note that Flash erase is performed Segment by Segment, not byte by byte (another difference between it and EEPROM), which means that if you were to erase a segment, all of its bytes will be erased back to 0xFF. The main flash section banks are each divided into segments which are 512




Figure 7.5: MSP430 Memory after writing 0xAA in CCS

byte large. A segment is the smallest amount of flash that can be erased. The Information Flash is also divided into segments (but is not arranged into banks) and here each segment is also the smallest amount that can be erased. However, the Info flash segments are smaller than those of the Main info, with 64 or 128 bytes.

If you don't care about the information in that segment, you can erase it and write data to it. However, in most cases, we're just updating the flash. The easiest approach in this case is to read the flash memory into an array in RAM, modify that array in RAM and write it all back. An example of this is provided later.

The flash erase procedure involves applying high voltage to the flash cells, which resets them. This voltage is generated by a circuit known as a charge pump that takes the relatively low voltage supplied to the MSP430 and generates a much higher voltage. It is this circuit which usually results in flash taking more die space and a host of other issues (though the issue is almost disappearing as IC technology nodes become smaller). This high voltage has to be applied and maintained for a certain amount of time. In some derivatives of the MSP430, the user must configure a clock source to provide the appropriate timing to the erase circuitry. Newer MSP430 do this automatically and remove the possibility that the flash may be destroyed if the erase timing is too long.

7.4.1 Configuring the Flash

Configuring the Flash Controller is necessary for erase and write operation, but not for reading operations. This is especially true for the Flash Timing Generator which must be properly configured

if the MSP430's Flash Controller requires it and a write or erase procedure will be performed. The Timing Generator usually requires 257 kHz to 476 kHz, which can be provided by dividing any of the system clocks ACLK, MCLK, SMCLK. Please check your device's datasheet for the frequency range required for your flash. The divider allows a division of 1 to 64 and this can be configured by 6 bits FN0 to FN5 in the Flash Control Register. The formula for the divider bits are:

$$1 + FN0 + 2 * FN1 + 4 * FN2 + 8 * FN3 + 16 * FN4 + 32 * FN5$$

In example 7.1, the clock was divided by 3 to obtain roughly 333kHz by setting FN1.

FSSEL_0,FSSEL_1,FSSEL_2, and FSSEL_3 select the clock source from ACLK, MCLK, SMCLK and SMCLK.

Finally, access to the flash configuration registers is locked by means of a key. Any modification of the Flash Controller registers requires the 0A5h Flash Key (defined usually as FWKEY) to be provided to the register in the upper byte as well. This prevents random accesses to the register from wreaking havoc on the flash, such as an accidental erase. If the key is not used during the write to the register, a security key violation occurs. KEYV flag and triggers a PUC system reset

An interesting method sometimes used to forcefully restart an MSP430 is to write dummy byte to the Flash register without the key. This is a common technique in devices without the Power Management Module of the F5xx and later MSP430 which includes this ability.

7.4.2 Reading the Flash

How do we use the flash memory? Because of the MSP430's Von Neumann architecture, using and addressing the Flash memory is extremely simple. This is because it is very much related to pointers and arrays and regular memory.

If you remember from C, a pointer is a variable that holds the address of another variable. We can use a pointer to point to an address in flash, and then read the value at that location by dereferencing it:

for example, if we have:

```
char number[5] = {1, 2, 3, 4, 5};
char * ptr;
ptr = & number[0];
```

To read the data of the information memory, we simply take a pointer and point it to that segment and dereference it.

Therefore, we can read the first byte at the start of the Info memory:

```
char *Flash_ptr;
Flash_ptr = (char *)0x1000;
char num = *Flash_ptr;
```

where num would now contain the value of the byte at position 0x1000; of course we can read the whole segment by using a for loop (or using memcopy)

Listing 7.2: Reading a Flash segment

```

1 char *Flash_ptr (char *)0x1000;
2
3 char infMem[256];
4 int i;
5
6 for(i = 0; i < 256; i++)
7 {
8     infMem[i] = *Flash_ptr;
9     Flash_ptr++;
10 }
11
12 FCTL3 = FWKEY;           // Clear Lock bit
13 FCTL1 = FWKEY + ERASE;  // Set Erase bit
14 *Flash_ptr = 0;        // Dummy write to erase Flash seg
15
16
17 FCTL1 = FWKEY + WRT;    // Set WRT bit for write operation
18 for (i = 0; i < 64; i++)
19 {
20     *Flash_ptr++ = value; // Write value to flash
21 }

```

Note that when we increment the flash pointer, we are incrementing the address to which it refers to. Also note that *i* will go from 0 to 255 and not from 1 to 256.

inf_mem will now have an exact copy of the flash memory.

7.4.3 Writing to Flash

Writing to the flash is a little more involved than reading because, as previously mentioned, the flash controller must be properly configured with the clocks and their sources. Also, writing generally requires an erase procedure. A single segment of flash is erased using the following code

Listing 7.3: Erasing a Flash segment

```

1 char *Flash_ptr;           // Flash pointer
2 Flash_ptr = (char *)0x1040; // Initialize Flash pointer
3 FCTL3 = FWKEY;           // Clear Lock bit
4 FCTL1 = FWKEY + ERASE;  // Set Erase bit
5 *Flash_ptr = 0;        // Dummy write to erase Flash seg

```

Once this code is erased, the segment extending from 0x1040 until the end of the segment (depending on the MSP430 and the location, this could be 64, 128 or 512 bytes). Erasing multiple segments simply requires you to change the address the pointer points to to the next segment and erase it as well. This is very important when you will be using multiple segments contiguously since writing can be performed sequentially. All of those segments must be erased before written.

Normal (non Block) writes to the flash are either byte or word. To begin the write process, the flash controller must be configured for write

```
FCTL1 = FWKEY + WRT;           // Set WRT bit for write operation
```

```
for (i = 0; i < 64; i++)
{
    *Flash_ptr++ = value;    // Write value to flash
    value++;
}
```

The for loop in the code above will write 0 to 64 starting at the location `Flash_ptr` points to before the loop.

7.4.4 Finalizing Flash Access

Once write or erase operations on flash have been finalized, the flash must be locked again to prevent any erroneous modification. If a write has been performed, the Write Bit should be cleared:

```
FCTL1 = FWKEY;           // Clear WRT bit
```

Afterwards, the actual lock:

```
FCTL3 = FWKEY + LOCK;   // Set LOCK bit
```

7.5 Flash Programming from the RAM

Although I have not previously mentioned it, all previous code to control the flash was executed from the Flash itself, in particular, from the code flash where the code to perform the reprogramming is located. This is usually acceptable for writing to the Info flash. However, when the Flash is being reprogrammed, the CPU is held while the programming is ongoing. After this process is finished, the CPU will begin executing instructions from the instruction following the write. Clearly, a CPU in the Off state might not be desirable in many applications.

It is possible to execute code from RAM. This allows the CPU to continue executing instructions while the flash is being programmed (erased, written). However, while the flash is actually being written or erased and the `BUSY` bit is set, the CPU cannot read or write from/to the flash. An Access Violation occurs in this case and the results are unpredictable. Flash programming from RAM is necessary for several reasons:

- You require the CPU to continue running through the periods when the flash is being programmed (note the caveat below with MSP430)
- You wish to perform Mass Erase or Block-Write

Note that in many MSP430 devices (regardless of whether executing from Flash or RAM), interrupt handling during flash programming is not possible. This is because the interrupt vectors themselves are stored in Flash areas. You must therefore disable all interrupts (via the General Interrupt Enable or GIE bit), and re-enable them after all flash operations are complete and the flash is locked.

7.6 The Linker and Flash

We previously discussed how a compiler and linker work together to generate an output file. The linker is responsible for placing the code in the appropriate locations in the binary output so that they will be placed properly in the device, in this case the MSP430. But how does the linker know where to put this information?

When you create a project and select the MSP430 derivative, the linker selects a linker file from a database. For a compiler to support a particular device, it must also have the information about the linker file. Such a linker file that is distributed with IAR Embedded Workbench for MSP430, specifically for the MSP430F5438A, is shown below:

You may locate these linker XCL files in the IAR installation folder. For example, in Windows 7 these are located in:

```
C:\Program Files (x86)\IAR Systems\Embedded Workbench 6.0 Kickstart\430\con
```

You can create your own XCL file by copying the appropriate XCL file to another folder(do not modify these originals), and then going to the Project options and selecting a customer linker file. In IAR, right click on the project's name and go to options. Then go to the Linker. In the Config tab, select Override Default and select the new xcl file.

Listing 7.4: MSP430F5438A IAR Linker

```

1 // *****
2 //
3 // XLINK configuration file for MSP430F5438A
4 //
5 // Copyright 1996-2010 IAR Systems AB
6 //
7 // $Revision: $
8 //
9 // *****
10
11 // -----
12 // Description
13 //
14 //
15 //
16 // Usage:
17 //
18 //   xlink [file file ...] -f lnk430f5438a.xcl
19 //
20 // -----
21 // Device summary
22 //
23 //
24 //
25 // Core:                MSP430X
26 //
27 // Interrupt vectors:   64
28 //
29 // Peripheral units:    0-001FF
30 //
31 // Information memory (FLASH): 01800-019FF
32 //
33 // Read/write memory (RAM):    01C00-05BFF
34 //
35 // Read-only memory (FLASH):  05C00-0FFFF
36 //                            10000-45BFF
37 //
38 //
39 //
40 // -----
41 // Segments
42 //
43 //
44 // -----
45 // Data read/write segments (RAM)
46 //
47 //
48 //
49 // The following segments are available for both
50 // the DATA16 and DATA20 segment groups.
51 //
52 // segment           Usage
53 // -----           -----
54 // DATA<nn>_Z      Data initialized to zero

```

```

55 // DATA<nn>_I      Data initialized by copying from DATA<nn>_ID
56 // DATA<nn>_N      Data defined using __no_init
57 // DATA<nn>_HEAP   The heap used by 'malloc' and 'free'
58 //
59 // segment          Usage
60 // -----
61 // CSTACK           Runtime stack
62 //
63
64
65 // -----
66 // Program and data read-only segments (FLASH)
67 //
68
69 //
70 // The following segments are available for both
71 // the DATA16 and DATA20 segment groups.
72 //
73 // segment          Usage
74 // -----
75 // DATA<nn>_C      Constant data, including string literals
76 // DATA<nn>_ID     initializers for DATA<nn>_I
77 //
78 // segment          Usage
79 // -----
80 // INFO             Information memory
81 // INFOA            Information memory, bank A
82 // INFOB            Information memory, bank B
83 // INFOC            Information memory, bank C
84 // INFOD            Information memory, bank D
85 // CSTART           Program startup code
86 // CODE             Program code
87 // ISR_CODE         Program code for interrupt service routines
88 // DIFUNCT          Dynamic initialization vector used by C++
89 // CHECKSUM         Checksum byte(s) generated by the -J option
90 // INTVEC           Interrupt vectors
91 // RESET            The reset vector
92 //
93 // Notes:
94 //
95 // * The segments CSTART, ISR_CODE, and DIFUNCT, as well as the segments in
96 //   the DATA16 segment group must be placed in in the range 0000-FFFD.
97 //
98 // * The INFOx and INFO segments overlap, this allows data either to be
99 //   placed in a specific bank or anywhere in the info memory.
100 //
101 // * The INTVEC and RESET segments overlap. This allows an application to
102 //   either use the reset vector provided by the runtime library, or
103 //   provide a reset function by defining an interrupt function associated
104 //   with the reset vector.
105 //
106
107
108 // -----
109 // Configuration

```

```
110 //
111
112 // -----
113 // Stack and heap sizes
114 //
115
116 // Uncomment for command line use
117 //-D_STACK_SIZE=80
118 //-D_DATA16_HEAP_SIZE=80
119 //-D_DATA20_HEAP_SIZE=80
120
121
122 // -----
123 // Define cpu
124 //
125
126 -cmsp430
127
128
129 // -----
130 // Support for placing functions in read/write memory
131 //
132
133 -QCODE_I=CODE_ID
134
135
136 // -----
137 // Placement directives
138 //
139
140 // -----
141 // Read/write memory
142 //
143
144 -Z (DATA) DATA16_I, DATA16_Z, DATA16_N, DATA16_HEAP+_DATA16_HEAP_SIZE=1C00-5BFF
145 -Z (DATA) CODE_I
146 -Z (DATA) DATA20_I, DATA20_Z, DATA20_N, DATA20_HEAP+_DATA20_HEAP_SIZE
147 -Z (DATA) CSTACK+_STACK_SIZE#
148
149
150 // -----
151 // Read-only memory
152 //
153
154 // -----
155 // Information memory
156 //
157
158 -Z (CONST) INFO=1800-19FF
159 -Z (CONST) INFOA=1980-19FF
160 -Z (CONST) INFOB=1900-197F
161 -Z (CONST) INFOC=1880-18FF
162 -Z (CONST) INFOD=1800-187F
163
164
```

```

165
166 // -----
167 // Low memory 0-0FFFF
168 //
169
170 // -----
171 // Code
172 //
173
174 -Z (CODE) CSTART, ISR_CODE=5C00-FF7F
175
176 // -----
177 // Constant data
178 //
179
180 -Z (CONST) DATA16_C, DATA16_ID, DIFUNCT, CHECKSUM=5C00-FF7F
181
182
183 // -----
184 // All memory 0-FFFFFF
185 //
186
187 // -----
188 // Code
189 //
190
191 -P (CODE) CODE=5C00-FF7F, 10000-45BFF
192 -Z (CODE) CODE_ID
193
194 // -----
195 // Constant data
196 //
197
198 -Z (CONST) DATA20_C, DATA20_ID=5C00-FF7F, 10000-45BFF
199
200
201 // -----
202 // Interrupt vectors
203 //
204
205 -Z (CODE) INTVEC=FF80-FFFF
206 -Z (CODE) RESET=FFFE-FFFF

```

The linker file contains many comments which assist in explaining it. The linker is a mapping between sections and memory addresses. For example, you can see on line 191 the definition for the section in which code is placed. Note that this definition is exactly the memory range assigned for banks A to D in the datasheet, which was shown in Figure 7.1. The information memory sections are also defined in lines 158 to 162.

It is important to understand the linker file and the section that are used because it might need to be modified. For example, if you wish to use a bank of main flash for a particular purpose, it would be wise to exclude it from the code section and reserve it in its own section. Much more details on the linker file are provided in the linker documentation. In IAR this is accessible by going to Help → IAR Linker and Library Reference Guide.

7.7 Using the Flash for data storage

Many applications require placing data in flash. The easiest way to do so is to declare a global constant array with some data:

```
__root const char calibrationData[5] = {1, 2, 3, 4, 5};
```

In IAR, the `__root` keyword is used to tell the compiler to include the array. If the keyword isn't used, the compiler might not actually add the array. When the compiler will compile the code, it will place `calibrationData` in flash, though since you've declared it, you can access it by name without another pointer.

Although the array was declared `const` and through the C language you cannot modify it, you can use a pointer and the MSP430's Flash controller to modify the data in flash and effectively change the data in the array.

Please note that the compiler can place this anywhere in flash (TODO: usually `infoflash`??). This might not be a good idea. In general, you want to know specifically where the compiler will place it so that if you decide to erase that segment you're sure that it will not affect anything else.

7.8 More Examples

TI provides a code examples zip file for each family of microcontrollers. Inside of these files, usually files containing the word `flashwrite` are specific examples of manipulating the Flash Memory.

Chapter 8

Timers

8.1 Introduction

Timers are common on all microcontrollers . They usually come in a variety of flavors and with different capabilities, however their working is relatively simple.

The timer is provided a clock source and counts a certain amounts of steps. When it reaches the limit of those steps it is usually set to fire”. The number of steps and the frequency of the clock source determine the time that it takes the timer to fire. For example, if our timer is sourced using a 32768Hz crystal and we were to count 32768 times it would fire in 1 second.

Using a certain clock frequency we can create a trigger that is a subset of that frequency. In the EZ430, no external clock is provided. There are two sources: the VLO and the fast DCO. The VLO is usually running at 12kHz. If we were to count 12000 times, it would take us one second. Note that the timer is a peripheral and does not require the CPU.

At the heart of MSP430 Timer, there is the TAR register. This is the counting register

8.2 Setting up the clock

The timer requires a clock and it may be sourced in a variety of ways. ACLK and SMCLK are the two most obvious options. Some MSP430 derivatives also enable sourcing though a pin, usually called TACLK. Clearly, the type of clock dictates how fast the timer can trigger.

The timer in the MSP430 is 16 bit, which means that it can only count from 0 up to 65536. This puts an upper limit as far as the frequency we can use. Using a 32768Hz clock source, we can trigger every 2 seconds. If we wanted a longer period, we would need to slow down the clock or find another way to slow it down (like using a function that counts more).

If the Timer is sourced from a slow clock, such as the 32768 Hz ACLK, it can only trigger as fast as 30us.

For a 8MHz clock for example, the timer could trigger every 0.125μs or so (in reality this is not possible because of other issues. However, as a whole the timer is much more accurate.

Because of this limitation, you should choose the clock source carefully.

8.3 Timer Modes

Although I previously said that the timer counts up to a value, in fact the timer in the MSP430 has more features. The timer can count UP to the value, it can count up to the max (65536 or 0xFFFF), it can count up and

8.4 Accuracy

A major issue with the selection of the clock source arises that affects precision. Using the VLO we are limited because of its speed because of the binary nature of the register that holds the counter. The fastest speed that the VLO would allow us to do is: $12000/1 = 12\text{kHz}$. In fact it will be a bit less because of execution time. However it's reasonable to see that 1 is the minimum number of counts.

Chapter 9

Interrupts and Low Power Modes

In this section I will not go into the details of the instruction set of the CPU, but rather will discuss some important features every MSP430 user should know about, the Interrupts and the Low Power Modes.

9.1 Interrupts

Interrupts represent an extremely important concept often not covered in programming since they are not readily used in a PC by most programmers. However, it is critical for anyone who programs microcontrollers to understand them and use them because of the advantages they bring.

Imagine a microcontroller waiting for a packet to be received by a transceiver connected to it. In this case, the microcontroller must constantly poll (check) the transceiver to see whether any packet was received. This means that the microcontroller will draw a lot of current since the CPU must act and use the communications module to do the polling (and the transceiver must do the same). Worse, the CPU is occupied taking care of the polling while it could be doing something more important (such as preparing to send a packet itself or processing other information). Clearly, if the CPU and the system are constantly polling the transceiver, both energy and CPU load are wasted. Interrupts represent an elegant solution to this problem. In the application above, a better solution would be to set things up so that the microcontroller would be informed (interrupted) when a packet has been received. The interrupt set in this case would cause the CPU to run a specific routing (code) to handle the received packet, such as sending it to the PC. In the meantime and while no interrupt has occurred, the CPU is left to do its business or be in low power mode and consume little current. This solution is much more efficient and, as we will later cover, allows for significant power savings. Interrupts can occur for many reasons, and are very flexible. Most microcontrollers have them and the MSP430 is no exception.

There are in general three types of interrupts:

- System Reset
- Non-maskable Interrupts (NMI)
- Maskable Interrupts

The interrupts are presented in decreasing order of system importance. The first two types are related to the microcontroller operating as it is supposed to. The last case is where all the modules allow interrupt capability and the user can take advantage of for the application. System Reset interrupts simply occur because of any of the conditions that resets the microcontroller (a reset switch, startup, etc.). These reset the microcontroller completely and are considered the most critical of interrupts. Usually you don't configure anything that they do because they simply restart the microcontroller. You have no or little control of these interrupts.

The second type of interrupts is the non maskable ones. Mask ability refers to the fact that these interrupts cannot be disabled collectively and must be disabled and enabled individually. These are interrupts in the category where the error can possibly be handled without a reset. Just like a normal PC, a microcontroller is a machine that has to be well oiled and taken care of. The supply voltage has to be satisfied, the clocks have to be right, etc. These interrupts occur for the following reasons:

- An edge on the RST/NMI pin when configured in NMI mode
- An oscillator fault has occurred because of failure
- An access violation to the flash memory occurred

These interrupts do not usually occur but they should be handled. By handled I mean code needs to be written to do something to deal with the problem. If an oscillator has failed, a smart thing to do would be to switch to another oscillator. The User's Guide for the MSP430F2274 provides more information about these type of interrupts. The last type of interrupts is the most common one and the one you have a large control over. These are the interrupts that are produced by the different modules of the MSP430 such as the I/O, the ADC, the Timers, etc. To use the interrupt, we need the following procedure:

1. Setup the interrupt you wish and its conditions
2. Enable the interrupt/s
3. Write the interrupt handler

When the interrupt happens, the CPU stops executing anything it is currently executing. It then stores information about what it was executing before the interrupt so it can return to it when the interrupt handler is done (unless the interrupt handler changes things). The interrupt handler is then called and the code in it is executed. Whenever the interrupt ends, the system goes back to its original condition executing the original code (unless we changed something). Another possibility is that the system was in a Low Power Mode, which means the CPU was off and not executing any instructions. The procedure is similar to the one detailed above except that once the interrupt handler has finished executing the MSP430 will return to the Low Power Mode. Note that the CPU is always sourced from the DCO (or an external high speed crystal) and this source must be on for the interrupt handler to be processed. The CPU will activate to run the interrupt handler. The process of going from CPU execution (or wakeup) to interrupt handler takes some time. This is especially true whenever the system is originally in sleep mode and must wakeup the DCO to

source the CPU. Normally, it is not critical but some applications might have issues with it taken longer than desired. We will now discuss a useful example: Using the switch of the EZ430-RF2500 to turn on and off the LED. This is best done by example. Two different listings are provided, the traditional way and an interrupt driver. You'll realize the immense benefits of using the interrupt solution quickly

Listing 9.1: EZ430-RF2500 Toggle LED using Switch - Polling

```

1 #include "msp430x22x4.h"
2
3 void configureClocks();
4 volatile unsigned int i;           // volatile to prevent optimization
5
6 void main(void)
7 {
8     WDTCTL = WDTPW + WDTHOLD;      // Stop watchdog timer
9     configureClocks();
10
11     // Configure LED on P1.0
12     P1DIR = BIT0;                  // P1.0 output
13     P1OUT &= ~BIT0;                // P1.0 output LOW, LED Off
14
15     // Configure Switch on P1.2
16     P1REN |= BIT2;                 // P1.2 Enable Pullup/Pulldown
17     P1OUT = BIT2;                  // P1.2 pullup
18
19     while(1)
20     {
21         if(P1IN & ~BIT2)           // P1.2 is Low?
22         {
23             P1OUT ^= BIT0;         // Toggle LED at P1.0
24         }
25     }
26 }
27
28 void configureClocks()
29 {
30     // Set system DCO to 8MHz
31     BCSCTL1 = CALBC1_8MHZ;
32     DCOCTL = CALDCO_8MHZ;
33
34     // Set LFXT1 to the VLO @ 12kHz
35     BCSCTL3 |= LFXT1S_2;
36 }

```

The code in listing 9.1 uses a technique called polling and is a terrible way to check if the switch has been pressed. The reason for this is apparent if you consider what is going on. The CPU is in a constant loop checking for whether the switch went low - remember that the switch is usually high until we press on it. It checks whether the pin is zero at any time. At the speed the CPU is going (8MHz), it will poll millions of times a second. There is no need to such a high rate of polling, but we're forced to do that because that's what the CPU is using. Second, let's assume that the program correctly detects that the pin has zero. It will toggle the LED just as we want. The problem is that a user's press of the switch will last many milliseconds. During this time the polling will continue rapidly and the LED will toggle again and again because the switch is still pressed, not something

we want. We can add some code for delay that will prevent this, but it causes issues. How long should the delay be? This is an added problem we wish to avoid. Another big issue is that during all this time, the CPU is continuously working and it doesn't do anything else, a complete waste of a microcontroller that is wasting a lot of energy. The next listing shows an improved program. In this case, we use the interrupt flag, the indicator that an interrupt has occurred.

Listing 9.2: EZ430-RF2500 Toggle LED using Switch - Interrupt Flag

```

1 #include "msp430x22x4.h"
2
3 void configureClocks();
4 volatile unsigned int i;    // volatile to prevent optimization
5
6 void main(void)
7 {
8     WDTCTL = WDTPW + WDTHOLD;    // Stop watchdog timer
9     configureClocks();
10
11     // Configure LED on P1.0
12     P1DIR = BIT0;                // P1.0 output
13     P1OUT &= ~BIT0;            // P1.0 output LOW, LED Off
14
15     // Configure Switch on P1.2
16     P1REN |= BIT2;              // P1.2 Enable Pullup/Pulldown
17     P1OUT = BIT2;               // P1.2 pullup
18     P1IES |= BIT2;              // P1.2 Hi/lo edge
19     P1IFG &= ~BIT2;            // P1.2 IFG cleared just in case
20
21     while(1)
22     {
23         if(P1IFG & BIT2)        // P1.2 IFG cleared
24         {
25             P1IFG &= ~BIT2;    // P1.2 IFG cleared
26             P1OUT ^= BIT0;     // Toggle LED at P1.0
27         }
28     }
29 }
30
31 void configureClocks()
32 {
33     // Set system DCO to 8MHz
34     BCSCTL1 = CALBC1_8MHZ;
35     DCOCTL = CALDCO_8MHZ;
36
37     // Set LFXT1 to the VLO @ 12kHz
38     BCSCTL3 |= LFXT1S_2;
39 }

```

In this case we don't have an interrupt handler that is called whenever the interrupt occurs. Rather, the CPU continuously polls for the flag to see if the interrupt occurred. There are several advantages to this technique: we correctly toggle the LED only when the switch is pressed and don't have the continuous toggle issue since not until the next transition do we toggle the LED. Another benefit is that, as opposed to interrupts, the CPU continuously executes code and doesn't go to an interrupt handler. This allows us to continue processing some information and check the flag later,

in which case we can act on it, rather than be forced to do so immediately. However, some of the previous disadvantages exist, mainly: The CPU still works too hard on something trivial like polling and we are still wasting energy. An important part of the code is clearing the interrupt flag for the pin. This ensures any previous flags set are cleared and we won't immediately jump into the interrupt handler when we activate the interrupts. Many programmers have made an error of not clearing the flags and then having the code run continuously. Always make sure you clear it when it is necessary.

The configuration of the LED pin is quite simple. It is the configuration of the switch that is important to us. First, we enable the pullup resistor by using P1REN and P1OUT. Afterwards, we enable the interrupt on P1.2 using P1IE (interrupt enable). Another important issue is that we can control the condition of the interrupt. P1IES controls whether the interrupt occurs on the rising edge (0 to 1) or on the falling edge (1 to 0). Only transitions cause interrupts not just static levels.

The next listing shows yet another listing that finally uses the interrupt method.

Listing 9.3: EZ430-RF2500 Toggle LED using Switch - Interrupt

```

1 #include "msp430x22x4.h"
2
3 void configureClocks();
4 volatile unsigned int i;           // volatile to prevent optimization
5
6 void main(void)
7 {
8     WDTCTL = WDTPW + WDTHOLD;      // Stop watchdog timer
9     configureClocks();
10
11     // Configure LED on P1.0
12     P1DIR = BIT0;                  // P1.0 output
13     P1OUT &= ~BIT0;                // P1.0 output LOW, LED Off
14
15     // Configure Switch on P1.2
16     P1REN |= BIT2;                 // P1.2 Enable Pullup/Pulldown
17     P1OUT = BIT2;                  // P1.2 pullup
18     P1IE |= BIT2;                  // P1.2 interrupt enabled
19     P1IES |= BIT2;                 // P1.2 Hi/lo falling edge
20     P1IFG &= ~BIT2;               // P1.2 IFG cleared just in case
21
22     _EINT();
23
24     while(1)
25     {
26         // Execute some other code
27     }
28 }
29
30 void configureClocks()
31 {
32     // Set system DCO to 8MHz
33     BCSCTL1 = CALBC1_8MHZ;
34     DCOCTL = CALDCO_8MHZ;
35
36     // Set LFXT1 to the VLO @ 12kHz
37     BCSCTL3 |= LFXT1S_2;

```

```
38 }
39
40 // Port 1 interrupt service routine
41 #pragma vector=PORT1_VECTOR
42 __interrupt void Port_1(void)
43 {
44     P1OUT ^= BIT0;           // Toggle LED at P1.0
45     P1IFG &= ~BIT2;        // P1.2 IFG cleared
46 }
```

Listing 9.3 shows the interrupt solution. The interrupts are activated using `eint()`; At this point you might be asking yourself: Wait, I activated the interrupt when I called the line with `P1IE`. Why do I have to enable it again? The answer lies in the name of these interrupts. These interrupts are maskable, which means that although they are controller individually, a second layer exists which can enable or disable them together. Therefore, we must disable the masking to allow our interrupts to run. The major benefit of the interrupt solution is that code continues to be executed after the interrupt is enabled. Rather, the while loop is executed. Whenever the user presses on the button, the interrupt handler is executed and then the CPU returns to its state of execution before the interrupt occurred.

An important note regarding switches is in order. Switches are not perfect. It is possible that the mechanical system inside it will toggle several times when the user presses on it. This is of course undesirable and caused by the internal mechanism vibrating and connecting/disconnecting several time. This effect is appropriately called bouncing. To avoid this, most programmers add code called a Debouncer. Simply put, once the interrupt for the switch fires, you wait a certain amount of time before you allow any other interrupt to be called. You may also wait to perform that actual action of turning on the LED, but simply avoiding multiple interrupt calls within a short period of time that no person can press on a with multiple times is the objective. I have not implemented a de-bouncer since the performance of the switch is not that bad, but a small delay of 50mS would probably suffice to prevent and bouncing problems.

Something we have not discussed yet is the case of multiple interrupts. In your application you'll likely use many interrupts and you can't predict when they occur. This means that you can't ensure by design that one interrupt will not be called while an interrupt handler is being executed. This issue is called nested interrupts and is a complex topic. Program execution can spiral out of control and unexpected things happen. Keeping interrupt routines short limits the time they are active and so avoids any nesting. However, this might not be sufficient. A simple solution, which might or might not be acceptable depending on your application is to disable the interrupts during an interrupt handler and re-enable them after the interrupt handler is done. If the loss of an interrupt during that time is acceptable, this represents a simple solution. To do this we can use the following:

```
_DINT(); // GIE flag clear disabling general interrupts
```

This solution might not acceptable. Improving on this technique is to create flags which, upon the return of an interrupt, are checked and acted upon. These flags can be variables that can indicate if and how many times a certain interrupt occurred. To do this simply make interrupt handlers increment the variable. This is very quick and is similar to what some operating systems use. As I mentioned, this topic is very complex and will not be further addressed.

We work next on augmenting the above software by reducing the power consumption. To do this we take advantage of the Low Power Modes provided by

9.2 Low Power Modes

Low Power Modes (LPMs) represents the ability to scale the microcontroller's power usage by shutting off parts of the MSP430. The CPU and several other modules such as clocks are not always needed. Many applications which wait until a sensor detects a certain event can benefit from turning off unused (but still running) parts of the microcontroller to save energy, turning them back on quickly when needed.

Each subsequent LPM in the MSP430 turns off more and more modules or parts of the microcontroller, the CPU, clocks and . T covers the LPM modes available in the MSP430F1611, which are similar to the ones available in most MSP430 derivatives. For the most accurate information, refer to the User's Guide and Datasheet of your particular derivative. Not mentioned here are LPM5 and some other specialized LPMs. These are available only in select MSP430 devices (in particular the newer F5xx series). However, the basic idea behind LPMs is the same , i.e. gradual shut down of the microcontroller segments to achieve power reduction.

- Active - Nothing is turned off (except maybe individual peripheral modules). No power savings
- LPM0 - CPU and MCLK are disabled while SMCLK and ACLK remain active
- LPM1 - CPU and MCLK are disabled, and DCO and DC generator are disabled if the DCO is not used for SMCLK. ACLK is active
- LPM2- CPU, MCLK, SMCLK, DCO are disabled DC generator remains enabled. ACLK is active
- LPM3 - CPU, MCLK, SMCLK, DCO are disabled, DC generator is disabled, ACLK is active
- LPM4 - CPU and all clocks disabled

It is important to note that once parts of the microcontroller are shut off, they will not operate until specifically turned on again. However, we can exit a low power mode (and turn these parts back on), and interrupts are extremely useful in this respect. Another Low Power Mode, LPM5, is available in some derivatives of the MSP430

As an example of the savings that can be achieved by incorporating Low Power Modes into the software design, I present Figure 9.1, which is shown in the MSP430F2274 User's Guide:

We usually enter a low power mode with the interrupts enabled. If we don't, we will not be able to wake up the system.

```
_BIS_SR(LPMx_bits + GIE);
```




Figure 9.1: Low Power Mode Savings

The x represents LPMs from 0 to 4 in the MSP430F2274. Another important thing to know is that the low power modes don't power down any modules such as the ADC or timers and these must be turned off individually. This is because we might want to combine the low power mode with them. A simple example is going in a Low Power Mode while the ADC is sampling and then waking up from the ADC interrupt when the sample has been obtained.

Example of Interrupts and Low Power Modes: The following listing is a modification of Listing 9.3. It incorporates entering a LPM and gaining significant energy savings.

Listing 9.4: EZ430-RF2500 Toggle LED using Switch - Low Power Mode

```

1 #include "msp430x22x4.h"
2
3 void configureClocks();
4 volatile unsigned int i; // volatile to prevent optimization
5
6 void main(void)
7 {
8     WDTCTL = WDTPW + WDTHOLD; // Stop watchdog timer
9     configureClocks();
10
11     // Configure LED on P1.0
12     P1DIR = BIT0; // P1.0 output
13     P1OUT &= ~BIT0; // P1.0 output LOW, LED Off
14
15     // Configure Switch on P1.2
16     P1REN |= BIT2; // P1.2 Enable Pullup/Pulldown
17     P1OUT = BIT2; // P1.2 pullup
18     P1IE |= BIT2; // P1.2 interrupt enabled
19     P1IES |= BIT2; // P1.2 Hi/lo falling edge
20     P1IFG &= ~BIT2; // P1.2 IFG cleared just in case
21
22     // Enter Low Power Mode 4 (LPM4) w/interrupt
23     __bis_SR_register(LPM4_bits + GIE);
24 }
25
26 void configureClocks()
27 {
28     // Set system DCO to 8MHz
29     BCSCTL1 = CALBC1_8MHZ;
30     DCOCTL = CALDCO_8MHZ;
31
32     // Set LFXT1 to the VLO @ 12kHz
33     BCSCTL3 |= LFXT1S_2;
34 }

```

```

35 |
36 | // Port 1 interrupt service routine
37 | #pragma vector=PORT1_VECTOR
38 | __interrupt void Port_1(void)
39 | {
40 |     P1OUT ^= BIT0;           // Toggle LED at P1.0
41 |     P1IFG &= ~BIT2;        // P1.2 IFG cleared
42 | }

```

In this case, we use the interrupt and we go into a low power mode. You can see that after entering the LPM there is not code (such as a while loop).

Why did we enter LPM4? Consider what we need during the microcontroller's time in the Low Power Mode. We don't need anything running really (in this application). There are no timers or any other peripherals that need a clock and so we can safely turn all of them off. LPM3 leaves the VLOCLK on (VLOCLK is the source for ACLK), which we don't need. LPM4 disables all clocks and the only interrupts that can activate the system from sleep are GPIO interrupts. If you go to LPM4 expecting that a Timer will activate the interrupt, it will never happen! The Timer's clock source was disabled. All you'll have is an MSP430 that's stuck in LPM forever. This is not our case so we can use LPM4 without any problems.

9.2.1 Exiting Low Power Modes

As previously mentioned, after executing the interrupt handler the system returns to the previous state. If the CPU was on and executing instructions, it will continue from where it left. If the system was called to a Low Power Mode, it will return. Suppose we don't want to return to a Low Power Mode and that perhaps we want to wakeup to do a lot of processing that wouldn't be good to put in an interrupt. To do this we must exist the low power mode in the interrupt handler. This is done with the following code:

```
__bic_SR_register_on_exit(LPMx_bits);    // Clear LPMx bits from 0 (SR)
```

Where x again represents the LPM that we want to exit from.

An example of this application is shown in the following listing.

Listing 9.5: EZ430-RF2500 Toggle LED using Switch - Low Power Mode 2

```

1 | #include "msp430x22x4.h"
2 |
3 | void configureClocks();
4 | volatile unsigned int i;    // volatile to prevent optimization
5 |
6 | void main(void)
7 | {
8 |     WDTCTL = WDTPW + WDTHOLD; // Stop watchdog timer
9 |     configureClocks();
10 |
11 |     // Configure LED on P1.0
12 |     P1DIR = BIT0;           // P1.0 output
13 |     P1OUT &= ~BIT0;        // P1.0 output LOW, LED Off

```

```

14
15 // Configure Switch on P1.2
16 P1REN |= BIT2; // P1.2 Enable Pullup/Pulldown
17 P1OUT = BIT2; // P1.2 pullup
18 P1IE |= BIT2; // P1.2 interrupt enabled
19 P1IES |= BIT2; // P1.2 Hi/lo falling edge
20 P1IFG &= ~BIT2; // P1.2 IFG cleared just in case
21
22 // Enter Low Power Mode 4 (LPM4) w/interrupt
23 __bis_SR_register(LPM4_bits + GIE);
24
25 // Interrupt ran and we left LPM4, we end up here
26 while(1)
27 {
28 // Execute some important code
29 }
30 }
31
32 void configureClocks()
33 {
34 // Set system DCO to 8MHz
35 BCSCTL1 = CALBC1_8MHZ;
36 DCOCTL = CALDCO_8MHZ;
37
38 // Set LFXT1 to the VLO @ 12kHz
39 BCSCTL3 |= LFXT1S_2;
40 }
41
42 // Port 1 interrupt service routine
43 #pragma vector=PORT1_VECTOR
44 __interrupt void Port_1(void)
45 {
46 P1OUT ^= BIT0; // Toggle LED at P1.0
47 P1IFG &= ~BIT2; // P1.2 IFG cleared
48 __bic_SR_register_on_exit(LPM4_bits); // Exit LPM4
49 }

```

The above code functions as follows: After setting up the LED and switch we go to LPM4, where energy is greatly conserved. The MSP430 waits for the interrupt. Once the switch is pressed, the LED is toggled and we exit LPM4. The next code that is executed is the while loop immediately after the LPM4 line. Note that in this case I just put a while loop to . You can put any code that you want.

Chapter 10

Analog to Digital Converters - ADCs

10.1 Introduction

The ADC is arguably one of the most important parts of a Microcontroller. Why is this? The real world is analog and to be able to manipulate it digitally some sort of conversion is required. Although many techniques are possible, using an Analog to Digital Converter (ADC or A2D) is the most common when good signal fidelity is necessary. A more simple interface could be some threshold circuit that would allow you to know whether a signal passed a certain threshold. The resolution is obviously very poor and would not be used to sample real signals unless it is required. The ADC is also probably the most difficult module to master. It is so because despite the simplicity that is assumed during courses that cover it (in which we just assume the ADC samples at a particular frequency and it does so perfectly to the available resolution.).

We will cover the ADC present in the MSP430F2274, a 10-bit SAR ADC. Most of the information will easily carry over to other ADCs, even if these ADCs have different features.

What is an ADC? It is a device which converts the analog waveform at the input to the equivalent representation in the digital domain. It does so by a variety of techniques which depend on the architecture of the ADC itself. SAR stands for Successive Approximation and this name indicates the general idea behind the conversion. There are many architectures out there.

The ADC samples and quantizes the input waveform. Sampling means that it measures the voltage of the waveform at a particular moment in time. In reality, the ADC cannot use an impulse but rather a larger time frame to sample so that the value is stable. When we talk about the sampling frequency, we refer to the fact that the ADC will be continuously sample. This is the most common mode discussed in DSP courses. However, an ADC can sample one or in many other ways, depending on the application. Quantization means it assigns a particular value, usually binary, to represent the analog voltage present. Both of these important concepts will be discussed next.

10.2 ADC Parameters

10.2.1 ADC Resolution

Probably the most important and most cited parameter of the ADC is its resolution, usually expressed in bits. The MSP430F2274 in our case has 10-bits of resolution. The resolution is an extremely important parameter because it dictates how the lowest difference we can differentiate. The higher the resolution, the more we can distinguish between minute differences in the voltage.

The resolution is specified in usually bits. This tells us how many different voltage levels we can represent.

$$\text{Quantization Levels} = 2^{\text{bits}}$$

For a 10-bit ADC we have a maximum of

$$2^{10} = 1024$$

possible levels of voltages we can represent. That's why higher bit ADCs such as 16 and even 24 are used: you can readily distinguish between smaller changes in the signal and your digital representation better approximates the original. For 16-bits we have $2^{16} = 65536$ and for 24-bit we have more than 16 million different levels, a huge increase due to the exponential nature. Note however that noise is a prevalent issue in ADCs and becomes the major constraint at a certain point

The question now comes to how the number of levels is related to the input voltage? This is simple. The number of levels possible is divided equally for the voltage range. If we have a 1V on the ADC's input pin and assuming the ADC is operating with a voltage range of 3V, each successive level from 0V to 3V is:

$$\frac{3V}{1024 - 1} = 0.00293V = 3mV$$

If the signal changes by less than 3mV we can't distinguish the difference, but every 3mV change in signal level translates into an addition of 1 to the binary number. If we have 0V at the input (without any noise), then we have 10 zeros (since it's a 10-bit ADC and no signal at the input). However, if we have 3mV at the input then we have 9 zeros followed by a one.

$$0V \rightarrow 0000000000 = 0x00$$

$$3mV \rightarrow 0000000001 = 0x01$$

You can think about the ADC as a staircase. For each range of voltages we assign an equivalent binary number usually referred to as the "code". The binary representation is usually replaced by hex since it is easier to read, as discussed in previous chapter. Note however that this is the ideal case and no offset or errors are taken into account.

An example of this is shown in the next figure

The following shows a simple 3-bit ADC. Such ADC does not exist but serves to show how each new level is represented in binary:

Unfortunately, we cannot cover all the important details of ADCs. The interested reader is referred to the following excellent sources of information:

3V	_____	111
2.625V	_____	110
2.25V	_____	101
1.875V	_____	100
1.5V	_____	011
1.125V	_____	010
0.75V	_____	001
0.375V	_____	000
0V	_____	

Figure 10.1: 3-bit ADC Example

- Application Note: SLAA013: "Understanding Data Converters". Texas Instruments : <http://www.ti.com/litv>
- Book: "Data Converter Handbook". Newnes 2004. ISBN - 0750678410

The general equation of the MSP430F2274 (which can be assumed for most other ADCs given a few changes) is:

$$N_{ADC} = 1023 \times \frac{V_{in} - V_{R-}}{V_{R+} - V_{R-}}$$

Notice that 1023 is $2^{10} - 1$.

A simple example follows: Lets assume that the input voltage is 1V. V_{R+} and V_{R-} represent the upper and lower limits of the voltage range. In our case we will have $V_{R+} = V_{CC}$ and $V_{R-} = V_{SS} = 0V$.

Using the equation above we have:

$$N_{ADC} = 1023 \times \frac{1V - 0V}{3V - 0V} = 1023 \times 0.333 = 341 = 0x155$$

The binary representation is not shown because it is large and hexadecimal represents it nicely. Notice that 1V is exactly a third of the voltage range and therefore the output code will be a third of all the maximum number, 1023 in this case.

10.2.2 ADC Sampling Frequency

The Sampling Frequency is another critical parameter of the systems. Usually the ADC manufacturer specifies this in SPS (samples per second).

The waveform at the input of the ADC might not stay constant. It might remain constant for long periods of time or change continuously (as with a voice at the input). When we sample the data, if

it doesn't change, we can sample once. However, it is more likely that we will setup the ADC to sample regularly. Different signals at the input of the ADC will require different sampling speeds. It is possible to oversample by sampling at a very high speed and therefore be able to capture different signals. Realize that if not sampling fast enough, very quick events can not be detected. Speech signals for example are generally considered to have important information for the region of 300Hz to 4000Hz. Using Nyquist's theorem we realize that we must therefore sample at a frequency twice the highest frequency (This is for low pass signals, i.e. signals where the lowest frequency is relatively close to 0Hz).

10.3 ADC Example - Temperature and Voltage

Instead of going over the details of the ADC, we will first see how a simple ADC application can be done using the MSP430F2274. This is done to give a good general picture of how the ADC can be quickly configured. Later more advanced features will be discussed.

The code to sample the integrated temperature sensor of the MSP430F2274 and the VCC is listed next:

Listing 10.1: EZ430-RF2500 ADC Sampling of Temperature and VCC

```

1 #include "msp430x22x4.h"
2
3
4 volatile unsigned int i;          // volatile to prevent optimization
5
6 void main(void)
7 {
8     WDTCTL = WDTPW + WDTHOLD;    // Stop watchdog timer
9
10    // Configure MCU Clocks
11    BCSCTL1 = CALBC1_1MHZ;
12    DCOCTL = CALDCO_1MHZ;
13    // Set LFXT1 to the VLO @ 12kHz
14    BCSCTL3 |= LFXT1S_2;
15
16    // Configure LED on P1.0
17    P1DIR = BIT0;                 // P1.0 output
18    P1OUT &= ~BIT0;              // P1.0 output LOW, LED Off
19
20    volatile long temp;
21    int delay;
22    int temp, volt;
23
24    // Configure ADC to Sample Temperature Sensor
25    ADC10CTL1 = INCH_10 + ADC10DIV_4;    // Temp Sensor ADC10CLK/5
26    ADC10CTL0 = SREF_1 + ADC10SHT_3 + REFON + ADC10ON + ADC10IE + ADC10SR;
27    for( delay = 240; delay > 0; delay-- );    // delay to allow reference to settle
28    ADC10CTL0 |= ENC + ADC10SC;           // Sampling and conversion start
29    __bis_SR_register(CPUOFF + GIE);     // LPM0 with interrupts enabled
30
31    // Read the conversion data
32    temp = ADC10MEM;

```

```

33     ADC10CTL0 &= ~ENC;           // Stop ADC Conversion
34
35     // Reconfigure ADC to Sample VCC
36     ADC10CTL1 = INCH_11;           // AVcc/2
37     ADC10CTL0 = SREF_1 + ADC10SHT_2 + REFON + ADC10ON + ADC10IE + REF2_5V;
38     for( degC = 240; degC > 0; degC-- ); // delay to allow reference to settle
39     ADC10CTL0 |= ENC + ADC10SC;     // Sampling and conversion start
40     __bis_SR_register(CPUOFF + GIE); // LPM0 with interrupts enabled
41     volt = ADC10MEM;
42     ADC10CTL0 &= ~ENC;
43     ADC10CTL0 &= ~(REFON + ADC10ON); // turn off A/D to save power
44 }

```

Listing 10.1 is heavily commented for easier understanding. As with any program, we begin by setting up the basic elements, clocks and I/Os. Line 25 is where the actual ADC use begin. We first must setup various parameters. In this case the channel and the clock of the ADC. The channel refers to the pin which will be sampled. The temperature sensor is integrated internally and therefore in this case does not belong to any pin. Channel 10 is the channel for the temperature sensor. On the same line we also configure the clock divider. This takes the clock source of the ADC and divides it by 5. More on why this is done later.

Line 26 configures other parameters. In this case we select the reference to be 1.

10.3.1 ADC Clock

In the case that we are using a simple resistive sensor that forms As with any digital device, the ADC itself also requires a clock, which is used for timing purposes.

The actual voltage used by the ADC as reference can vary and can be the supply voltage of the MSP430, an internal reference, or an external reference provided by the user. This results in a great flexibility

The ADC inputs of the MSP430 are usually denoted by an Ax, with the x representing the channel number. Most commonly, 8 or 16 channels are available (A0 - A15).

10.3.2 ADC Modes

The ADC with most MSP430 devices is very flexible. Although normally we think of sampling a single source (channel), the MSP430 allows us several modes which enable:

- Sample one channel once
- Sample one channel multiple times sequentially
- Sample Multiple channels (a range of channels) once
- Sample multiple channels multiple times sequentially

This can greatly simplify applications since we do not need to sample each channel individually, but can group them together. In the case of a Tri axial analog output accelerometer, for example, all axes (X,Y , and Z) can be sampled one after the other sequentially, saving time in the process.

Chapter 11

Digital Interfaces

Imagine two people try to communicate. Do you tell the other person "Hello" by spelling out "H" "E" "L" "L" "O", each letter separately? Of course not. Doing so would be slow and inefficient (and for humans difficult to understand). Similarly, when two digital devices attempt to communicate (to transmit and receive information and commands), they do so by exchanging their binary alphabet (0 and 1). However, to make communication more efficient, they group these binary numbers into groups. Moreover, there is a certain protocol used. The protocol is an agreed mechanism by which two parties communicate. It could be possible that we were to say the word hello by saying "O" "L" "L" "E" "H". We don't do that of course, but it is possible (it is done in computers where we determine whether to send the Most significant bit or Least significant bit first in the group of bits). Another reason this is done is because we can represent more things with more bits (one bit can represent two things and this is a limitation). This will become more obvious when we discuss the use of the SPI interface to access the registers of other devices (and when there are more than two registers, as is most always the case, we must be able to represent them all). The protocol allows two devices to communicate by agreeing on a set of common rules. This is the basic idea about digital interfaces.

However, in most cases, you don't get to choose any rules (you could do what's called bit banging which is simulating a protocol by simple I/O manipulation). There are a distinct set of protocols available and you will choose from what is available. For example, the CC2500 transceiver in the EZ430-RF2500 uses the SPI interfaces, which is a serial interface (bits are sent serially), is synchronous (uses a clock to indicate when bits are sent/received, and allows multiple devices on the bus. For the transceiver to successfully understand what we are telling it (isn't that the purpose of all communications), we must configure the MSP430 to the parameters which the transceiver will understand (and which the MSP430 will understand back). We will cover this in the following section.

Three of the most common digital interfaces (and the two that are available in most MSP430 devices (and in most embedded microcontrollers in general) are Serial Peripheral Interface (SPI), the Universal asynchronous receiver/transmitter (UART), and I²C.

Another popular interface is USB, which is available in almost all computers. Increasingly, USB is being added to microcontrollers because of the ubiquity of connecting to a computer using it. However, it is outside of our scope and won't be covered. There are many sources out there about it.

Another note of importance. The MSP430 (and other microcontrollers) group digital interfaces such as SPI, UART and I²C together in one module (or even in different modules). Therefore, don't be surprised if one of these communication modules can do only one type of communications at a time (although it is possible to switch). As mentioned before, microcontrollers have limited pins and therefore require multiplexing between different functionalities.

In the MSP430, the digital interface modules are called Universal Serial Communication Interface (USCI) and Serial Communication Interface (USART). Different derivatives have different combinations of these two modules. The MSP430F2274 in the EZ430-RF2500 includes two of these modules: USCLA0 and USCIB0. USCLA0 is capable of running UART, LIN, IrDA, and SPI USCIB0 is capable of running SPI and I2C.

However, TI designed the EZ430-RF2500 and already assigned the functionality of the modules. USCLA0 is allocated for UART communications with the PC, while USCIB0 is connected to the CC2500 and needs to be in SPI mode.

11.1 Serial Peripheral Interface (SPI)

SPI is a synchronous data link included in the USART module. Using the 4 wires (sometimes 3) it allows a microcontroller to communicate with any number of sensors or devices that support it. As previously mentioned, SPI works on the basis of bytes (8 bits) and not just bits.

The SPI protocol deals with the concept of master and slaves. A master is the device (like a microcontroller) which initiates the communications and provides the clock for synchronization. The slave is a device (like the CC2500) which responds to the communications, receives commands, and sends data, all as requested by the microcontroller. Note that it is possible for a microcontroller to be a slave (as long as it supports that functionality). Such can be done for two microcontrollers that are communicating with each other. One will necessarily be the master and the other the slave.

The following shows a diagram of the signals used by SPI:

Table 11.1: SPI Signals

Signal Name	Alternative Names	Description
SCLK	SCK, CLK	Serial Clock
MOSI/SIMO	SDI, DI, SI	Master Output Slave Input
MISO/SOMI	SDO, DO, SO	Slave Output Master Input
SS	nCS, CS, nSS, STE, CE	Slave Select, Chip Enable

The serial clock signal is simply a clock provided to synchronize the communications between the master and slave. SOMI and SIMO are very descriptive names. They indicate exactly how they should be connected (The Master IC's output to the slave's input). Finally, a fourth signal, which can be tied to low (Ground) at the slave, is necessary to select the slave with which we want to communicate. SPI can't communicate with all the slaves at once, so the slave select signal is used. Only the currently active slave receives the information because its slave select line is selected to low. The rest of the slaves must have their Slave select line high or else they will receive the commands. Figures 11.1 shows SPI with only 1 slave and with multiple slave, so that the use of the Slave Select (or Chip Enable) is apparent.




Figure 11.1: SPI Configurations

11.1.1 Configuring SPI

Before we can use SPI to communicate with another device, we must configure the MSP430's USART module. Depending on the MSP430, it may have 1 or 2 of such modules (or in small versions it might not include one at all). Check with the User's Guide and Datasheet to ensure that the USART module is SPI capable. In the datasheet of the MSP430F2274 we find the following pins to be the ones for SPI:

Table 11.2: SPI Pins in the MSP430F2274

Pin	Peripheral Module Functionality(Multiplexed)
P3.0/ UCB0STE/ UCA0CLK/ A5	General-purpose digital I/O pin USCLB0 slave transmit enable USCLA0 clock input/output ADC10, analog input A5
P3.1/ UCB0SIMO/ UCB0SDA	General-purpose digital I/O pin USCLB0 slave in/master out in SPI mode SDA I ² C data in I ² C mode
P3.2/ UCB0SOMI/ UCB0SCL	General-purpose digital I/O pin USCLB0 slave out/master in in SPI mode SCL I ² C clock in I ² C mode
P3.3/ UCB0CLK/ UCA0STE	General-purpose digital I/O pin USCLB0 clock input/output USCLA0 slave transmit enable
P3.4/ UCA0TXD/ UCA0SIMO	General-purpose digital I/O pin USCLA0 transmit data output in UART mode slave in/master out in SPI mode
P3.5/ UCA0RXD/ UCA0SOMI	General-purpose digital I/O pin USCLA0 receive data input in UART mode slave out/master in in SPI mode

All the pins are also multiplexed and can be either a general I/O or serve the function of a peripheral module. In this microcontroller, there is an USCLA0 and USCLB0 modules, both of which are capable of SPI. To configure the SPI interface we need to make use of the register which controls it. We will assume that we are using the ez430-RF2500, which uses SPI only on USCLB0 since USCLA0 is used with the UART to communicate with the PC. The ez430-RF2500 uses a 4-wire configuration (the Chip Select pin is part of the module). This is incompatible with multiple slaves on the bus. In that case, we will need to configure the SPI to use 3-wire with P3.0 as a general purpose pin. This will be discussed later. This configuration is the most common one, as the 4-wire configuration allows for multiple masters, which we won't deal with. Pin 3.0 is used as a general I/O to control the CS_n (same as slave select). Pins P3.1 and P3.2 are used as SIMO and SOMI, respectively. Pin P3.3 is the clock which allows synchronization between the master (MSP430) and the slave. We will first show the code and then explain it in detail:

```

1 P3SEL |= 0x0C;           // P3.3,2 USCI_B0 option select
2 P3DIR |= 0x01;         // P3.0 output direction
3 UCB0CTL0 |= UCMSB + UCMST + UCSYNC; // 3-pin, 8-bit SPI mstr, MSB 1st,
4 UCB0CTL1 |= UCSSEL_2;   // SMCLK as clock source
5 UCB0BR0 = 0x02;
6 UCB0BR1 = 0;
7 UCB0CTL1 &= ~UCSWRST;   // **Initialize USCI state machine**

```

Some of the comments explain what is done, but we will explain in more detail. The first thing done is that the pins' functionality has to be determined. P3.2 and P3.3 are selected to be the Primary peripheral module function (in this case USCI_B0). This is necessary because the default functionality for these pins is general I/O. P3.1 is presumably already set for SPI. P3.0 needs to be set as an output because it is the pin that controls the Chip Select (Slave select). The code afterwards sets settings for the SPI. These are conveniently put together on one line. UCMSB selects the Most Significant Bit(MSB) first. UCMST selects this SPI module as the master. UCSYNC simply indicates that we are selecting a synchronous mode. By default, the UCB0CTL0 register is set for 3-wire communications.

As with most MSP430 peripherals, the USCI needs a clock to source it. In this case, the code selects SMCLK, which has to be previously set by the code. UCB0BR0 and UCB0BR1 are set the same for most any SPI code, so we will ignore those statements. Finally, we enable the SPI using UCSWRST.

11.1.2 Using SPI for Communications

The first thing we must do to send information over SPI is to enable the chip. In this case P3.0 is the I/O pin connected to the CSn (Chip Select) of the CC2500. For this we must put the CSn low (logic 0):

```

1 P3OUT &= ~0x01;           // Pull CS low to enable SPI

```

Once this is done, we can simply send the information by filling the buffer:

```

1 UCB0TXBUF = 0x00;         // Dummy write to start SPI

```

Finally when we are finished using SPI we must pull Chip select high to deselect the CC2500:

```

1 P3OUT |= 0x01;           // Pull CS high to disable SPI

```

This turning on and off of the Chip Select doesn't have to be done after every writing to the buffer. It is done when we are finished with the chip.

Let's assume we have the CC2500 RF IC using the SPI and that we've enabled CSn line by pulling it low. The following code writes to a register:

```

1 void TI_CC_SPIWriteReg(char addr, char value)
2 {
3     TI_CC_CSxOUT &= ~TI_CC_CSxPIN; // /CS enable by pulling CSn low
4     while (TI_CC_SPI_USCIB0_PxIN&TI_CC_SPI_USCIB0_SOMI); // Wait for CCxxxx ready
5     IFG2 &= ~UCB0RXIFG;           // Clear flag
6     UCB0TXBUF = addr;             // Send address by putting in SPI buffer

```

```

7 |   while (!(IFG2&UCB0RXIFG));           // Wait for TX to finish (status byte received)
8 |   IFG2 &= ~UCB0RXIFG;                 // Clear flag
9 |   UCB0TXBUF = value;                   // Send data
10|   while (!(IFG2&UCB0RXIFG));          // Wait for TX to finish (status byte received)
11|   TI_CC_CSxOUT |= TI_CC_CSxPIN;       // /CS disable
12| }

```

The function accepts two bytes as parameters. The first is the address of the register and the second is the value we want to put in the register. From the CC2500 datasheet you can see that there is a register called IOCFG2. It's address is 0x00. This means that if we wanted to change the value of this register, we would call the function as follows:

```

1 | TI_CC_SPIWriteReg(0x00, 0x0F);

```

An analysis of the function is now discussed. As in all previous code, the first thing to do is to enable the CS line by pulling it low. Afterwards, the while loop waits until the slave is ready by checking the input pin. The CC2500 pulls SOMI low when it is ready to receive. We clear the flag of the SPI because perhaps we received a byte before and we want to know when we receive one. Again as before, we send the address by just filling the buffer with the byte of the address. Now we wait for the flag that a byte has been received. This is because when writing to registers, the status byte is sent on the SO pin each time a header byte or data byte is transmitted on the SI pin. so in fact the CC2500 sends this byte back and we just check for its presence to know that our byte has been sent. Once we know this happened, we again clear the flag so that we will know the next time it is sent. We fill the buffer, this time with the actual data to be written to the register and wait to receive the status byte again (indicating the status byte was received. After this we pull CSn (Slave or chip enable) high to turn off the SPI interface.

The code covers much of what needs to be done. It is important to read the datasheet of the IC to be communicated, as each has its own requirements. More information is available at the CC2500 datasheet from the TI website. Also code is available as part of the ez430-RF2500 demo code.

Chapter 12

UART

In this chapter we will cover the UART module in the MSP430 to enable communications with a computer. The Universal Asynchronous Receiver/Transmitter (UART) module is part of both the Universal Serial Communication Interface (USCI) Module and the universal synchronous/asynchronous receive/transmit(USART) Module. Depending on which derivative of the MSP430 you are using, you will use either USCI or USART and enable one of the modes of operation such as UART, SPI, I2C, etc. In this section we will focus on the UART functionality in either of these. Note that you may only use the module to do one type of communication at a time. However, it might be possible to have several connections to the USART/USCI Modules and switch between protocols, but this can be problematic at times and will not be addressed. Also note that although the names of registers may differ when using the UART in either the USCI or USART modules, the major approaches are the same and only minor modifications should be required such as register names.

12.1 Hardware Connection

UART is simple to connect and it uses 2 lines: TX (Transmit) and RX (Receive). No clock is required because it's asynchronous.

When interfaced to another module that uses UART, you will need to cross connect the lines. The TX of one device will connect to the RX of another device, while the RX will be connected to the TX of the second device, as shown in figure 12.1.

RS232 is a common connection between PC and devices, and it can be easily converted to RS232 using a level converter. Another popular solution is to connect the UART of the MSP430 to a USB to UART converter such as the FT232RL and FT2232 from FTDICHIP. These ICs creates a virtual COM port on the host machine that allows you to seamlessly transfer bytes between the host machine and the MSP430.

12.1.1 UART connectivity on the MSP430F1611

The MSP430F1611 includes two USART modules, USART0 and USART1. Both are capable of using UART independently of each other.




Figure 12.1: Connecting UART

UART0 uses P3.4/UTXD0 as TX and P3.5/URXD0 as RX

UART1 uses P3.6/UTXD1 as TX and P3.7/URXD1 as RX

12.1.2 UART connectivity on the MSP430F2274

The MSP430F2274 includes two USCI modules, USCLA0 and USCIB0. Only USCLA0 has UART functionality and uses P3.4/UCA_0TXD /UCA_0SIMO as TX and P3.5/UCA_0RXD /UCA_0SOMI as RX.

12.2 Using UART

From the name of the peripheral you can easily guess that UART sends bit serially, one bit after the other, as opposed to parallel. However we don't usually send individual bits because of the fact that they convey little information by themselves. Rather, the UART uses a buffer that is one byte (8 bits) long to which information is written.

On a deeper level, the module sends more than just the byte a user places in the buffer. UART Typically sends a start bit, seven or eight data bits, an even/odd/no parity bit, an address bit (address-bit mode), and one or two stop bits. The extra bits depend on the configuration of the UART module, but the most common configuration is 8 data bits, one stop bit and no parity. Extra bits that are sent are ultimately discarded and not stored at the destination buffer. The result of the parity check is however available. Figure 12.2 shows the bit arrangement in the MSP430's UART Module:




Figure 12.2: UART Bits

12.3 Configuring the UART

Before we can use the UART successfully to send data, both the host computer and the UART Module must be configured. Because the communication is asynchronous, if the two are not configured correctly then errors will occur because the interpretation of the timing of the bits will be off.

To better explain how to configure the UART module, we will use the following code as an example.

Listing 12.1: Configuring UART

```

1 P3SEL |= 0x30;           // P3.4,5 = USART0 TXD/RXD
2 ME1  |= UTXE0 + URXE0;  // Enable USART0 TXD/RXD
3 UCTL0 |= CHAR;          // 8-bit character
4 UTCTL0 |= SSEL0;        // UCLK = ACLK = 32.768kHz
5 UBR00 = 0x03;           // 32.768kHz/9600 = 3.41
6 UBR10 = 0x00;
7 UMCTL0 = 0x4a;          // Modulation
8 UCTL0 &= ~SWRST;        // Initialize USART state machine
9 IE1  |= URXIE0 + UTXIE0; // Enable USART0 RX/TX interrupt

```

The code above demonstrates that the following steps are needed:

1. Select the I/O pins for UART
2. Enable UART TX, RX, or Both
3. Select the character format - 7 or 8 bits
4. Source the UART Module with a clock of a certain frequency
5. Set the baud generator correctly so as to get a correct baud rate from the clock sourced
6. Enable the UART module state machine
7. Enable Interrupts if required

We will take the code and improve upon it to understand it and adapt it to most any MSP430 UART modules.

12.3.1 Selecting the UART Function for Pins

As we mentioned previously, the MSP430 (and many other microcontrollers) multiplexes the functionality of various modules at the physical pins. In our case we can select between general I/O and the internal modules. This means that we can either use a pin as a General Purpose Input/Output (HIGH or LOW) or one of the internal modules.

The functionality selected is determined by PxSEL registers, where x can represent 1,2,3 ... ports, depending on which port the UART module is located. Each of the 8 bits in the register controls the functionality of the pin. As an example, the MSP430F1611 has two modules which support UART functionality. From the MSP430F1611 datasheet, pages 7 and 8, you can see that the UART pins that can be used:

```
P3.4/UTXD0
P3.5/URXD0
P3.6/UTXD1
P3.7/URXD1
```

P3.4 indicates pin 4 on port 3, while P3.5 indicates the same port but pin 5 and as you can see, they can be a General Purpose digital I/O pin or a UART pin. Because the pins are in port 3, we need to modify P3SEL. The question is how should we modify it? In the User's Guide for the MSP430F1611 we find that:

Bit = 0: I/O Function is selected for the pin

Bit = 1: Peripheral module function is selected for the pin

To select UART functionality for P3.4 and P3.5 we use the following c language statement that sets the two bits for pins 4 and 5, selecting them for the UART module:

```
P3SEL |= BIT4 + BIT5;
```

12.3.2 Enabling UART RX and TX

The UART module is very flexible. It can control the RX and TX functionalities individually. In our case we will enable both. To do this the following code is used:

```
ME1 |= UTXE0 + URXE0;    // Enable USART0 TXD/RXD
```

Of course, for different UART modules we will need to use a different register settings. This can be easily found by looking at the user's guide for the particular MSP430.

12.3.3 Select the character format

Different settings are available for the communications between the two devices. Among these are the number of data bits, parity and stop bits. However, the most common setting, designated 8N1,

is 8 data bits, no parity and 1 stop bit. Because of the default settings of the registers, parity is disabled by default and the stop bit is 1. The only thing left to do is to make the data bits 8 and not 7 as follows:

```
UCTL0 |= CHAR;    // 8-bit character
```

Some other options include PENA for enabling parity, PEV for selecting Odd or Even parity, or SPB for selecting two stop bits. An 8-bit character is normally used because it allows us to use ASCII, as well as allowing for a larger range of numbers to be represented

12.3.4 Selecting A Clock

Just like most modules in the MSP430, a clock must be chosen for the UART, and this is arguably the most critical task. Remember that the UART module sends bits every x amount of time. The timing for the bits is determined from the baud rate, which is derived from a division of the clock time. We will assume the clock has already stabilized and is working. Also, it should operate at a specific frequency (which we need to know).

In general, the clock frequency dictates the maximum baud rate allowed. ACLK, which is the auxiliary clock and is usually 32kHz or less (although in some cases it can be sourced from an external high frequency crystal). This means that in reality, and to keep errors low, we should select a baud rate of 9600 baud or less. Higher baudrates require a higher clock and this is normally sourced from SMCLK and MCLK, which themselves are sourced by the internal DCO or an external high frequency crystal. A crystal is always preferable because it has a much better frequency stability and would almost guarantee better performance as far as error rates. It's only drawback is the added price and space requirements.

To select the clock source for BRCLK we need to set the SSELx bits (bits 5 and 4) of the UTCTL1 (for USART1) or UTCTL0 (for USART0). As indicated in the datasheet, bits 5 and 4 will cause the following selection:

Binary	Hex	Define	Clock Source
00	0x00	SSEL0	UCLKI
01	0x10	SSEL1	ACLK
10	0x20	SSEL2	SMCLK
11	0x30	SSEL3	SMCLK

The define Column indicates an easy to use definitions defined in the header file msp430x16x.h that simplifies reading the code.

To select the 32kHz ACLK source we simply do:

```
UTCTL0 |= SSEL0;    // UCLK = ACLK = 32.768kHz
```

If we wanted to select SMCLK we could do:

```
UTCTL0 |= SSEL1;    // UCLK = ACLK = 32.768kHz
```

```
UTCTL0 |= SSEL1;    // UCLK = SMCLK
```

UCLKI allows reversing the polarity, which is useful in certain situations but in fact rarely used.

12.3.5 Setting the Baud Rate Generator

The baud rate we get is derived from the clock sourced to the UART module. Standard baud rates are used and include:

Baud Rate
1200
2400
9600
19200
57600
115200
230400

Using the desired baud rate and the clock rate we can calculate the divisors required to generate the baud rate. Extensive information about this is available in the User's Guide and other sources. A simple way to calculate the registers is to use one of the many available calculators:

<http://msp gcc.sourceforge.net/baudrate.html>

These calculators use the same methodology that will be presented.

Given a baud rate **b = 9600** and a clock speed **c = 32768** we have:

$$\text{divisor} = c/b = 32768/9600 = 3.413$$

The integer part of this divisor is easily implementable and you can see so in the code:

```
UBR00 = 0x03;    // 32.768kHz/9600 -> 3 integer part
```

16-bit Integer divisors are possible because there are two registers, UBR00 and UBR01, both 8-bit. UBR00 contains the upper byte while UBR01 contains the lower byte (lower 8 bits). In this case 3 is not greater than 255 so that we place 0x00 in the upper byte and 3 (or its hex equivalent 0x03) on the lower byte.

However, we need to get as close as possible to 3.413 or else our baudrate will be off and we would like it to be as close as possible. In cases where there is no fractional part we ideally don't need to deal with it, but it occurs in most every case because the clock used is never perfect. Using the modulation register we can approximate the fractional part of 3.413 to some degree of accuracy.

To calculate the number of bits we need to set in the modulation register we multiply the fractional part by 8: $8 * 0.413 = 3.304$ Rounding gives us 3, so that tell us 3 of the 8 bits in the modulation register need to be set. To determine which bits to turn on we to successively multiply the fractional part:

The important thing to note is whether, for every addition of one to the multiplying factor, we jump to the next integer. For example, multiplying by 2 resulted in 0.826 which has an integer of 0. Multiplying by 3 means 1.239 which has an integer part of 1. Because they are not the same we say that we set it that bit because it carried. On the other hand from 3 to 4 we still have 1 as the integer part, so we don't set. Notice that we have a total of 8 multiplications so we can set the 8 bits of the modulation register. Starting from the multiplication by 2, we consider the first bit to

Multiplication	Result	Bit Result
0.413 * 2	0.826	0
0.413 * 3	1.239	1
0.413 * 4	1.652	0
0.413 * 5	2.065	1
0.413 * 6	2.478	0
0.413 * 7	2.891	0
0.413 * 8	3.304	1
0.413 * 9	3.717	0

be 0 because multiplication by 1 previously has the same integer as multiplying by 2. The second multiplication by 3 changes the integer part so we do set that one. The multiplication by 4 doesn't so we don't set that bit.

Note that when we first multiplied $8 * 0.413 = 3.304$ we said that we needed to have 3 bits set (1) when rounding down in the register and from the table it is apparent that we have obtained only 3 bits set. All we need to do is to consider where do we start with the bits, from the bottom or top? The answer is that the bit for which we multiplied by 2 is the Least Significant, i.e. we start from the bottom of the table and go up, resulting with the following:

$$01001010 = 0x4A$$

NOTE: Some sources state that we should start from the top, but this does not seem to comply with TI's own algorithm.

12.3.6 Enabling the Module

Once all the settings are complete, we can enable the module as follows:

```
UCTL0 &= ~SWRST;    // Initialize USART state machine
```

12.3.7 Enabling Interrupts

Although it is possible to perform UART operations like as receiving bytes using polling techniques, it is most commonly done using interrupts and these must be enabled. The UART reset state does not have to be enabled to turn them on or off. This can be done as follows:

```
IE1 |= URXIE0 + UTXIE0;    // Enable USART0 RX/TX interrupt
```

Although both interrupts are enabled in this case, it is more common to use the RX interrupt.

Name	URL
Putty	http://www.chiark.greenend.org.uk/~sgtatham/putty/
RealTerm	http://realterm.sourceforge.net/

12.4 Configuring the Host Computer

Once the MSP430's UART is configured, we need to configure the software on the Host computer, if it is used. This varies depending on the software. Some free software options are:

Both of these programs are free and open source, although RealTerm is more terminal oriented and has many more useful options than Putty.

To configure the terminal, find the COM Port that is connected to the MSP430. In Linux this is usually done a command such as `ls /dev/ttyUSB*` Which will show a list of all USB VCP ports. Note that in some cases the chipset used maps differently, but the above technique works well in most cases, especially for FTDI Chip USB to UART ICs. In Windows finding the COM port is done either by using the Device Manager in the Control Panel. RealTerm itself shows all possible COM Ports.

Set the terminal to this port and to use the baud rate you've previously selected. Remember to set the data bits to 8, the parity to none and the stop bits to 1 unless you've configured the UART differently.

12.5 Sending and Receiving Information with the UART

Now that the UART is correctly configured and the hardware is set (including the PC software), we can send data using the UART. This is in fact very simple for single bytes. Each UART module contains two buffers, a receive buffer and a transmit buffer. For the sending buffer, the user places data (a byte in this case) and the module sends it by serializing it and sending the appropriate bits. For the receive buffer it is the opposite, the bits are deserialized and converted into a byte.

12.5.1 Sending Bytes

To send a byte of data we simple put it in TXBUF0 or TXBUF1:

```
TXBUF0 = 0x41;
```

```
TXBUF1 = 0x41;
```

The UART module takes the data put into the buffer and sends it automatically without any further intervention. Sending multiple bytes isn't much harder. We simply have to ensure that the sending of the previous byte was finalized before putting new information in the buffer, and this is done by checking the appropriate flag:

```
while (!(IFG1 & UTXIFG0));    // USART0 TX buffer ready?
```

The MSP430 will keep executing the while loop (do nothing because of the colon) until the sending flag has been cleared.

This means we can put this together to send 2 bytes, one after the other:

```
1 while (!(IFG1 & UTXIFG0));    // USART0 TX buffer ready?
2 TXBUF0 = 0x41;
3 while (!(IFG1 & UTXIFG0));    // USART0 TX buffer ready?
4 TXBUF0 = 0x42;
```

Even when sending an individual byte, always use the while loop to ensure the buffer is ready before trying to transmit it.

12.5.2 ASCII

Since the UART buffer is one byte (when it is configured to use 8 bits), we can send any number from 0 to 255, or 0x00 to 0xFF. A major use of sending data is to send readable information. ASCII is a standard encoding supported by terminal emulators and the C compilers for the MSP430, which allows us to send readable characters. In this case ASCII maps a byte with a specific value to a specific character. For example, the lowercase 'a' corresponds to 0x61 while 'b' corresponds to 0x62.

It is not necessary to learn the ASCII table. Any compiler will automatically replace 'a' with the corresponding number without the user needing to explicitly specify that value. In the case of IAR and other compilers, it is also legal to perform the following comparison:

```
1 char buf = 'a';
2 if(buf == 'a')
3 // Successful -This will execute
4 else
5 // This will not run
6
7 // Some other code
```

More will be discussed about using ASCII to send readable strings to the Host computer.

12.6 Sending Multiple bytes

Sending single bytes can be useful in some cases, but real applications require sending many bytes. We will delve into how to send and receive multiple bytes and strings. It is your job as the designer to create a protocol if necessary to send useful data. Some of this will be covered later.

Say for example we wanted to automate and send the alphabet from A to Z (no small caps).

This can be done as follows:

```

1 for(int i = 'A'; i < 'Z'; i++)
2 {
3     TXBUF0 = i;
4     while (!(IFG1 & UTXIFG0));    // USART0 TX buffer ready?
5 }

```

A simple for loop can be used to send bytes.

This similar approach can be used to iterate through an array. Suppose we have an array called information with the following information in it:

Value	4	8	15	16	23	42
Byte	0	1	2	3	4	5

The above array is a collection of decimal numbers with 4 at position 0 and 42 at position 5. The array has 6 elements and to send it we can use the following code:

```

1 for(int i = 0; i < 6; i++)
2 {
3     TXBUF0 = information[i];
4     while (!(IFG1 & UTXIFG0));    // USART0 TX buffer ready?
5 }

```

The for loop accesses the array iteratively and sends the information, waits for the byte to be sent, sends the next one and so on.

12.6.1 Sending Strings

Sending strings to the PC is useful in several situations. One of them is when you wish to use a terminal program such as Putty to receive data or debug the microcontroller by printing readable information. This eliminates the need to build a custom application to read and write information to the port. This requires time and is not always possible(although since Putty and RealTerm are open source, adapting the code simpler than starting from scratch). To simplify our work of sending the strings we will create a simple function:

```

1 void sendstring(char * str, int len)
2 {
3     int i = 0;
4     for(i = 0; i < len; i++)
5     {
6         while (!(IFG1 & UTXIFG0));    // USART0 TX buffer ready?
7         TXBUF0 = str[i];
8     }
9 }

```

We can use the function above to send the string Hello World as follows:

```
sendstring("Hello World\n", 12 );
```

Note that

n is not two characters but rather one character, the newline character. Therefore only 12 bytes should be sent and not 13.

The function can be simplified to avoid counting the length of the string by using the compiler's string functions, specifically strlen, which gives the size of the string.

Listing 12.2: Send String Function

```

1 void sendstring(char * str)
2 {
3     int i = 0;
4     for(i = 0; i < strlen(str); i++)
5     {
6         while (!(IFG1 & UTXIFG0));    // USART0 TX buffer ready?
7         TXBUF0 = str[i];
8     }
9 }

```

Be cautious in using this function. Calling the function as follows is fine: sendstring("Hello World\n"); However, sending a fixed array with characters will fail. This is because the compiler automatically appends a

0 to the end of the string, which is how strlen finds the length. if the parameter to the function is an array that does not contain

0, strlen will not find the end and the results are unpredictable (yet they will always be disastrous).

12.7 Receiving Data

Another function that is very common is receiving data. On the MSP430 this is done almost automatically if it is correctly configured and the UART RX functionality is enabled. The only issue that must be resolved is that multiple bytes received successively overwrite the buffer. If the host computer sends two bytes and the user does not fetch the first one from the RX buffer before the second byte arrives, it will be lost since the buffer will be overwritten with the second byte.

The complexity of handling this problem depends on your application. In all cases, however, every time a byte arrives it should be copied to another buffer (an array for example). Once enough bytes have been received, we process them as needed. The number of bytes received can be fixed (processing every 5 bytes received for example) or flexible (such as using delimiters to indicate the start and end).

We will now discuss how to receive individual bytes and show how we can place them in a buffer for processing. The USART module can generate an interrupt to immediately execute code when a byte is received. For this to occur we must first set the correct interrupt to be active, using the following code:

```
IE1 |= URXIE0;    // Enable USART0 RX interrupt
```

OR

```
IE2 |= URXIE1;    // Enable USART1 RX interrupt
```

When a character is received, the interrupt vector will be called and the code in the interrupt service routine (ISR) is executed. Such an interrupt vector can be, for example, made to echo a character (send the character that was received back):

```

1 #pragma vector=USART0RX_VECTOR
2 __interrupt void usart0_rx (void)
3 {
4     while (!(IFG1 & UTXIFG0));    // USART0 TX buffer ready?
5     TXBUF0 = RXBUF0;             // RXBUF0 to TXBUF0
6 }

```

where USART0RX_VECTOR can also be USART1RX_VECTOR or the USCI RX vectors. The interrupt does exactly what its name means, it interrupts the microcontroller by stopping the current statement it is processing, and redirects it to the interrupt code, returning to the original statement that was being executed once the interrupt handling is finished. This might or might not be done quickly enough, depending on what is done at the interrupt. It is usually best to keep interrupts as short as possible, raise a flag to indicate that the data needs to be processed and end the interrupt service routine.

If an interrupt vector is not the right solution, a flag can also be used. A flag is a bit set by a module to indicate that something has occurred. IFG1 and IFG2 each contain a bit, URXIFG0 and URXIFG1 respectively, that indicate whether a character has been received. this can be checked by:

```

if(IFG1 & URXIFG0)
{
    \\Do Something with the received character
}

```

OR

```

if(IFG2 & URXIFG1)
{
    \\Do Something with the received character
}

```

It is important to clear the flag after it has been checked because it is not automatically cleared.

```
IFG1 &= ~URXIFG0;    // Clear flag URXIFG0
```

OR

```
IFG2 &= ~URXIFG1;    // Clear flag URXIFG1
```

To demonstrate receiving of multiple bytes, we will assume everything is initialized and interrupts are enabled.

Listing 12.3: Buffering UART Data

```
1 char buffer[256];
2 int bufpos = 0;
3 char process_flag = 0;
4 void main()
5 {
6 // System and UART Initialization Here
7
8 __EINT();
9
10 while(1)
11 {
12     if( process_flag)
13     {
14         \\ Process bytes
15
16         process_flag = 0;
17     }
18 }
19 }
20
21 #pragma vector=USART0RX_VECTOR
22 __interrupt void usart0_rx (void)
23 {
24     buffer[bufpos] = RXBUF0;
25     if(bufpos == 255)
26     {
27         process_flag = 1;
28         bufpos = 0;
29     }
30 }
```

The code above is relatively simple. 256 bytes are collected, a flag raised and the information set for processing.

12.8 Framing

Using fixed sizes of bytes is not always a good idea. Whenever the size of the data to be communicated can vary, for example commands and data can be of different. In this case a more flexible method is desirable. This can be realized using delimiters, which is just using special characters before and after the actual data to indicate frame the data.

The methodology for using delimiters is in fact very similar to the Serial Line Internet Protocol (SLIP). Although we won't be using IP on top of the delimiters, it will be useful as a comparison. More information on SLIP is available here:

<http://en.wikipedia.org/wiki/SLIP> <http://tools.ietf.org/html/rfc1055>

We must first choose two bytes that will represent the start and end. This can be any value between 0x00 and 0xFF. It is better to use values that will not be used in any other way. 0x00 and 0xFF are all too common when faults and should be avoided. We will therefore choose the following and define them:

```
#define START_BYTE 0x7A
#define END_BYTE    0x7E
```

Listing 12.4: UART Framing

```
1 #define START_BYTE 0x7A
2 #define END_BYTE    0x7E
3 #define MAX_BUF_LEN 256
4
5 char uart_buffer[MAX_BUF_LEN];
6 int bufpos = 0;
7 char process_flag = 0;
8 int uart_packet_len = 0;
9 char frame_received_flag = 0;
10
11 void main()
12 {
13     // System and UART Initialization Here
14
15     __EINT();
16
17     while(1)
18     {
19         if(frame_received_flag)
20         {
21             \\ Process bytes
22
23             process_flag = 0;
24         }
25     }
26 }
27
28 #pragma vector=USART0RX_VECTOR
29 __interrupt void usart0_rx (void)
30 {
31
32     /* Protect against buffer overflow */
33     /* In this case overwrite */
34     if(buffer_counter == MAX_BUF_LEN-1)
35     {
36         buffer_counter = 0;
37     }
38
39     /* Received Starting Byte */
40     if(!packet_start_flg && (SERIAL_RX_BUF == START_BYTE))
41     {
42         packet_start_flg = 1;
43         buffer_counter = 0;
44     }
45     /* Store Frame Bytes without SFB and EFB */
46     else if(packet_start_flg && (SERIAL_RX_BUF != END_BYTE))
47     {
48         // Store in buffer
49         uart_buffer[buffer_counter] = SERIAL_RX_BUF;
50         buffer_counter++;
```

```
51 | }
52 | /* Received Stop Byte so process Frame */
53 | else if(packet_start_flg && (SERIAL_RX_BUF == END_BYTE))
54 | {
55 |     frame_received_flag = 1;
56 |     uart_packet_len = buffer_counter-2; // Discard START and END bytes
57 |     packet_start_flg = 0;
58 | }
59 | }
```

The code above is commented so it is easier to understand. The code starts by checking if the `START_BYTE` (0x7A in this case) was received. If it was, then every subsequent byte is stored in the buffer, as long as the byte to be stored is not the `END_BYTE`. If the `END_BYTE` is received, we set the flag indicating that the frame was received and reset the process of receiving the frames. In addition to this, the first for loop protects against frames larger than the buffer. This is necessary since writing past the array is dangerous and should never be done. To avoid losing frames we would need to increase the size of the buffer, which means more memory would be required.

12.9 Future Additions

The code presented above is of course only the beginning. There are, as always, issues with the code presented above:

- 1) Although we've received a frame, we don't know if all of the bytes received are correct. Using a Cyclic Redundancy Check (CRC) we can have much better certainty that the bytes have not been modified unintentionally.
- 2) If a frame is lost there is no retransmission.
- 3) It is possible that the `START` and `END` bytes are not unique and the frame delimiters detected are incorrect. Therefore, a more unique pattern such as using two bytes that are unlikely to occur should be considered.

Some of these options are likely necessary if the UART Error rate is higher than it should. If so, consider adding a crystal for a stable clock or reducing the baudrate.

Chapter 13

Wireless Communications with CC2500

13.1 Introduction

If you're using the EZ430-RF2500 you're likely attracted by its wireless transceiver. In fact, many of today's applications use some form of wireless communications, from WiFi to bluetooth. The part of the tutorial will discuss using the EZ430-RF2500's wireless capabilities at a basic level to transmit data from one node to another.

It is not possible to cover all the aspects of the transceiver, as there are many options. Rather, the main focus will be to discuss the most common features to transmit and receive data. The code provided is based on the SLAA325 application note from Texas Instruments. This application note concerns communications between an MSP430 and several of TI's transceivers which share many common aspects.

Although we will focus on using the CC2500, many issues in this section will be applicable to many digital sensors and devices that use the SPI interface or even another digital bus.

The EZ430-RF2500 includes the CC2500 transceiver. This is a 2.4GHz transceiver capable of up to 250kbps. It is an all digital transceiver that is accessed using the SPI bus.

Transmission of packets at its most basic level consists of:

1. Initialize the SPI interface
2. Configure the CC2500 with the correct frequency, output power
3. Fill the TX FIFO with the data to be sent
4. Tell the CC2500 to initiate transmission

For reception, an interrupt based method will be used as follows:

1. Initialize the SPI interface if necessary
2. Configure the CC2500 with the correct frequency, output power, and modulation
3. Configure the interrupt pins to alert MSP430 that a packet has arrived
4. When a packet is received, retrieve it from the RX FIFO

13.2 SPI Interface

Initialization of the SPI interface is similar to what was previously discussed. We must ensure that the SPI clock used does not go above the maximum frequency supported by the transceiver.

The transceiver has registers internally similar to the MSP430. However, since these are external to the MSP430, we must use the SPI interface to read and write to them. When a register is read, we send the name of the register to be read. The transceiver responds with whatever value that is stored in that register. When writing, we must provide two things, the register's address and the value which we want to write.

Physically the SPI interface of the EZ430-RF2500 consists of:

- CSN P3.0 UCB0STE
- SIMO P3.1 UCB0SIMO
- SOMI P3.2 UCB0SOMI
- SCLK P3.3 UCB0CLK

These 4 pins of port 3 of the MSP430F2274 are for the SPI interface connectivity using the USCLB0 module. Because of the limited number of I/O on the board, it is possible to use pins P3.1 to P3.3 for another SPI device, although this device will require a CSN line of its own and both devices can't be talking to the Microcontroller at the same time.

The `spi.h` and `spi.c` include all functions that are needed to communicate with the CC2500. The `TL_CC_SPISetup` configures the SPI interface.

Note that this implementation uses a 3-wire interface, leaving us to manually control the CSN line.

The CC2500 improves on this even further. Two access modes are available for reading and writing: Single and Burst. Single access means we can only read/write to one register, with each following read requiring a completely new operation. Burst Access is used to access multiple registers sequentially without the need to give the each register's position. What the transceiver does is simply assume that the next value you provide will go to the next register. Table 37 in page 60 of the CC2500 datasheet shows the different registers available, what modes of access is possible, and whether they are read/write or just read.

Two registers are special, these are the TX and RX FIFO. A FIFO can be thought of as a large array. The data to be transmitted or data that was received is placed here by the microcontroller. These FIFOs are fixed at 64 bytes. This can result in some problems when the data to be sent exceeds 64 or so bytes. A solution to this will be discussed later.