Introduction to PIC Programming

Baseline Architecture and Assembly Language

by David Meiklejohn, Gooligum Electronics

Lesson 4: Reading Switches

The previous lessons have introduced simple digital output, by turning on or flashing an LED. That's more useful than you may think, since, with some circuit changes (such as adding transistors and relays), it can be readily adapted to turning on and off almost any electrical device.

But most systems need to interact with their environment in some way; to respond according to user commands or varying inputs. The simplest form of input is an on/off switch. This lesson shows how to read and respond to a simple pushbutton switch, or, equivalently, a slide or toggle switch, or even something more elaborate such as a mercury tilt switch – anything that makes or breaks a single connection.

This lesson covers:

- Reading digital inputs
- Conditional branching
- Using internal pull-ups
- Hardware and software approaches to switch debouncing

The Circuit

We need to add a pushbutton switch to the circuit used in lessons 1 to 3.

Luckily the Low Pin Count demo board used for these lessons includes a tact switch connected to pin GP3, as shown below. You should keep the LED from the previous lessons connected to GP1.

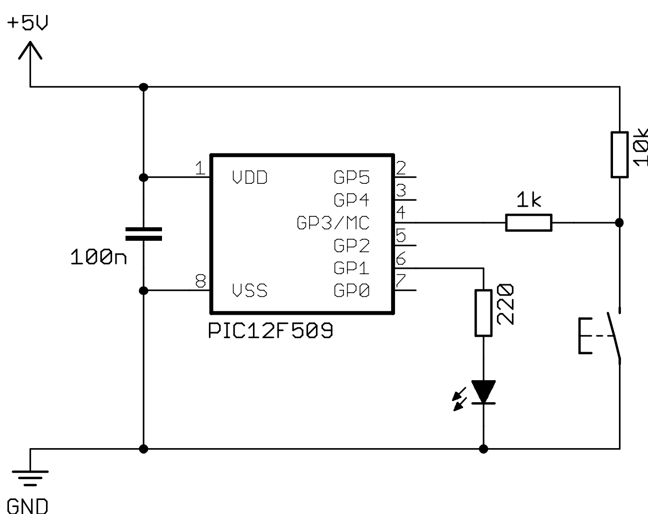

The pushbutton is connected to GP3 via a 1 kΩ resistor. This is good practice, but not strictly necessary. Such resistors are used to provide some isolation between the PIC and the external circuit, for example to limit the impact of over- or under-voltages on the input pin, electro-static discharge (*ESD*, which pushbuttons, among other devices, can be susceptible to), or to provide some protection against an input pin being inadvertently programmed as an output. If the switch was to be pressed while the pin was mistakenly configured as an output, set "high", the result is likely to be a dead PIC – unless there is a resistor to limit the current flowing to ground.

In this case, that scenario is impossible, because,

as mentioned in lesson 1, GP3 can only ever be an input. So why the resistor? Besides helping to protect the PIC from ESD, the resistor is necessary to allow the PIC to be safely and successfully programmed.

You might recall, from lesson 0, that the PICkit 2 is an In-Circuit Serial Programming (ICSP) programmer. The ICSP protocol allows the PICs that support it to be programmed while in-circuit. That is, they don't have to be removed from the circuit and placed into a separate, stand-alone programmer. That's very convenient, but it does place some restrictions on the circuit. The programmer must be able to set appropriate voltages on particular pins, without being affected by the rest of the circuit. That implies some isolation, and often a simple resistor, such as the 1 kΩ resistor here, is all that is needed.

To place a PIC12F508/9 into programming mode, a high voltage (around 12V) is applied to pin 4 – the same pin that is used for GP3. Imagine what would happen if, while the PIC was being programmed, with 12V applied to pin 4, that pin was grounded by someone pressing a pushbutton connected directly to it! The result in this case wouldn't be a dead PIC; it would be a dead PICkit 2 programmer…

But, if you are sure that you know what you are doing and understand the risks, you can leave out isolation or protection resistors, such as the 1 kΩ resistor on GP3.

The 10 kΩ resistor holds GP3 high while the switch is open. How can we be sure? According to the PIC12F509 data sheet, GP3 sinks up to 5 µA (parameter D061A). That equates to a voltage drop of up to 55 mV across the 10 kΩ and 1 kΩ resistors in series (5 µA × 11 kΩ), so, with the switch open, the voltage at GP3 will be a minimum of VDD – 55 mV. The minimum supply voltage is 2.0 V (parameter D001), so in the worst case, the voltage at GP3 = 2.0 – 55 mV = 1.945 V. The lowest input voltage guaranteed to be read as "high" is given as 0.25 VDD + 0.8 V (parameter D040A). For VDD = 2.0 V, this is 0.25 × 2.0 V + 0.8 V = 1.3 V. That's well below the worst-case "high" input to GP3 of 1.945 V, so with these resistors, the pin is guaranteed to read as "high", over the allowable supply voltage range.

In practice, you generally don't need to bother with such detailed analysis. As a rule of thumb, 10 kΩ is a good value for a *pull-up* resistor like this. But, it's good to know that the rule of thumb is supported by the characteristics specified in the data sheet.

When the switch is pressed, the pin is pulled to ground through the 1 kΩ resistor. According to the data sheet, GP3 sources up to 5 µA (parameter D061A). The voltage drop across the 1 kΩ resistor will be up to 5 mV (5 µA × 1 kΩ), so with the switch closed, the voltage at GP3 will be a maximum of 5 mV. The highest input voltage guaranteed to be read as a "low" is 0.15 VDD (parameter D030A). For VDD = 2.0 V (the worst case), this is 0.15 × 2.0 V = 300 mV. That's above the maximum "low" input to GP3 of 5mV, so the pin is guaranteed to read as "low" when the pushbutton is pressed.

Again, that's something you come to know as a rule of thumb. With just a little experience, you'll look at a circuit like this and see immediately that GP3 is normally held high, but is pulled low if the pushbutton is pressed.

Interference from $\overline{\text{MCLR}}$

There is a potential problem with using a pushbutton on GP3; as we have seen, the same pin can instead be configured (using the PIC's configuration word) as the processor reset line, $\overline{\text{MCLR}}$.

This is a problem because, by default, as we saw in lesson 1, when the PICkit 2 is used as a programmer from within MPLAB, it holds the $\overline{\text{MCLR}}$ line low after programming. You can then make the $\overline{\text{MCLR}}$ line go high by selecting "Release from Reset". Either way, the PICkit 2 is asserting control over the $\overline{\text{MCLR}}$ line, connected directly to pin 4, and, because of the 1 kΩ isolation resistor, the 10 kΩ pull-up resistor and the pushbutton cannot overcome the PICkit 2's control of that line.

> *When the PICkit 2 is used as a programmer within MPLAB, it will, by default, assert control of the* $\overline{\text{MCLR}}$ *line, overriding the pushbutton switch on the Low Pin Count Demo Board.*

If you are using MPLAB 8.10 or later, this problem can be overcome by changing the PICkit 2 programming settings, to tri-state the PICkit 2's $\overline{\text{MCLR}}$ output (effectively disconnecting it) when it is not being used to hold the PIC in reset.

To do this, select the PICkit 2 as a programmer (using the "Programmer → Select Programmer" submenu) and then use the "Programmer → Settings" menu item to display the PICkit 2 Settings dialog window, shown on the right.

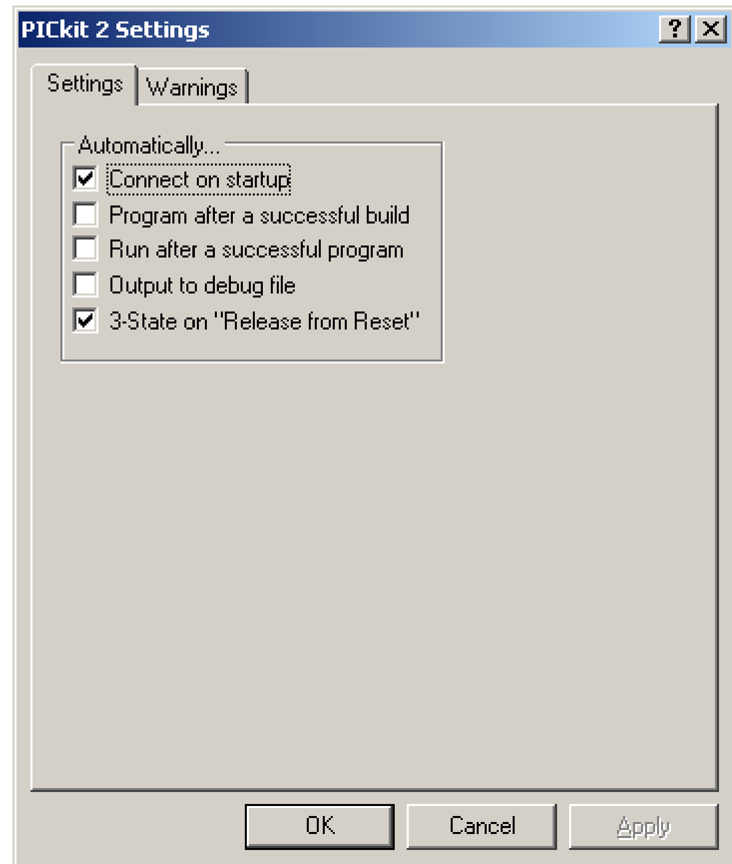

Select the '3-State on "Release from Reset"' option in the "Settings" tab and then click on the "OK" button.

After using the PICkit 2 to program your device, the PICkit 2 will assert control of the $\overline{\text{MCLR}}$ line, holding it low. If your application is not configured for external reset, the GP3 input will be held low, overriding the pushbutton on the LPC Demo Board.

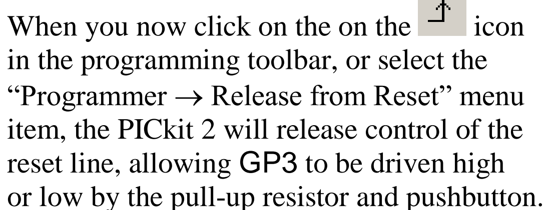

When you now click on the on the icon in the programming toolbar, or select the "Programmer → Release from Reset" menu item, the PICkit 2 will release control of the reset line, allowing GP3 to be driven high or low by the pull-up resistor and pushbutton.

Reading the Switch

We'll start with a short program that simply turns the LED on when the pushbutton is pressed.

Of course, that's a waste of a microcontroller. To get the same effect, you could leave the PIC out and build the circuit shown on the right! But, this simple example avoids having to deal with the problem of switch contact bounce, which we'll look at later.

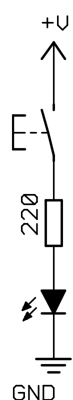

In general, to read a pin, we need to:

- Configure the pin as an input
- Read or test the bit corresponding to the pin

Recall, from lesson 1, that the pins on the 12F508 and 12F509 are only *digital* inputs or outputs. They can be turned on or off, but nothing in between. Similarly, they can read only a voltage as being "high" or "low". As mentioned above, the data sheet defines input voltage ranges where the pin is guaranteed to read as "high" or "low". For voltages between these ranges, the pin might read as either; the input behaviour for intermediate voltages is *undefined*.

As you might expect, a "high" input voltage reads as a '1', and a "low" reads as a '0'.

Normally, to configure a pin as an input, you would set the corresponding TRIS bit to '1'. However, this circuit uses GP3 as an input. As discussed above, GP3 is a special pin, in that it can be configured as an external reset. If it is not configured as a reset, it is always an input. So, when using GP3 as an input, there's no need to set the TRIS bit. Although for clarity, you may as well do so.

An instruction such as ‘movf GPIO,w’ will read the bit corresponding to GP3. The problem is that it reads all the pins in GPIO, not just GP3. If you want to act only on a single bit, you need to separate it from the rest, which can be done with logical masking and shift instructions, but there’s a much easier way – use the bit test instructions. There are two:

‘btfsc f,b’ tests bit ‘b’ in register ‘f’. If it is ‘0’, the following instruction is skipped – “**b**it **t**est **f**ile register, **s**kip if **c**lear”.

‘btfss f,b’ tests bit ‘b’ in register ‘f’. If it is ‘1’, the following instruction is skipped – “**b**it **t**est **f**ile register, **s**kip if **s**et”.

Their use is illustrated in the following code:

```
start
        movlw   b'111101'       ; configure GP1 (only) as an output
        tris    GPIO            ; (GP3 is an input)

        clrf    GPIO            ; start with GPIO clear (GP1 low)
loop
        btfss   GPIO,3          ; if button pressed (GP3 low)
        bsf     GPIO,1          ;   turn on LED
        btfsc   GPIO,3          ; if button up (GP3 high)
        bcf     GPIO,1          ;   turn off LED

        goto    loop            ; repeat forever
```

Note that the logic seems to be inverse; the LED is turned on if GP3 is clear, yet the ‘btfss’ instruction tests for the GP3 bit being set. Since the bit test instructions skip the next instruction if the bit test condition is met, the instruction following a bit test is executed only if the condition is *not* met. Often, following a bit test instruction, you’ll place a ‘goto’ or ‘call’ to jump to a block of code that is to be executed if the bit test condition is not met. In this case, there is no need, as the LED can be turned on or off with single instructions, which we have (surprisingly) not seen before:

‘bsf f,b’ sets bit ‘b’ in register ‘f’ to ‘1’ – “**b**it **s**et **f**ile register”.

‘bcf f,b’ clears bit ‘b’ in register ‘f’ to ‘0’ – “**b**it **c**lear **f**ile register”.

Previously, we have set, cleared and toggled bits by operating on the whole GPIO port at once. That’s what these bit set and clear instructions are doing behind the scenes; they read the entire port, set or clear the designated bit, and then rewrite the result. They are examples of ‘read-modify-write’ instructions, as discussed in lesson 2, and their use can lead to unintended effects – you may find that bits other than the designated one are also being changed. This unwanted effect often occurs when sequential bit set/clear instructions are performed on the same port. Trouble can be avoided by separating sequential ‘bsf’ and ‘bcf’ instructions with a ‘nop’.

Although unlikely to be necessary in this case, since the bit set/clear instructions are not sequential, a shadow register (see lesson 2) could be used as follows:

```
start
        movlw   b'111101'       ; configure GP1 (only) as an output
        tris    GPIO            ; (GP3 is an input)

        clrf    GPIO            ; start with GPIO clear (LED off)
        clrf    sGPIO           ; update shadow copy

loop
        btfss   GPIO,3          ; if button pressed (GP3 low)
        bsf     sGPIO,1         ;   turn on LED
        btfsc   GPIO,3          ; if button up (GP3 high)
        bcf     sGPIO,1         ;  turn off LED
```

```
        movf    sGPIO,w         ; copy shadow to GPIO
        movwf   GPIO

        goto    loop            ; repeat forever
```

It's possible to optimise this a little. There is no need to test for button up as well as button down; it will be either one or the other, so we can instead write a value to the shadow register, assuming the button is up (or down), and then test just once, updating the shadow if the button is found to be down (or up).

It's also not really necessary to initialise GPIO at the start; whatever it is initialised to, it will be updated the first time the loop completes, a few µs later – much too fast to see. If setting the initial values of output pins correctly is important, to avoid power-on glitches that may affect circuits connected to them, the correct values should be written to the port registers before configuring the pins as outputs, i.e. initialise GPIO before TRIS.

But when dealing with human perception, it's not important, so the following code is acceptable:

```
;***** VARIABLE DEFINITIONS
        UDATA_SHR
sGPIO   res 1                   ; shadow copy of GPIO


;************************************************************************
RESET   CODE    0x000           ; effective reset vector
        movwf   OSCCAL          ; update OSCCAL with factory cal value


;***** MAIN PROGRAM

;***** Initialisation
start
        movlw   b'111101'       ; configure GP1 (only) as an output
        tris    GPIO            ; (GP3 is an input)


;***** Main loop
loop
        clrf    sGPIO           ; assume button up -> LED off
        btfss   GPIO,3          ; if button pressed (GP3 low)
        bsf     sGPIO,1         ;   turn on LED

        movf    sGPIO,w         ; copy shadow to GPIO
        movwf   GPIO

        goto    loop            ; repeat forever
```

If you didn't use a shadow register, but tried to take the same approach – assuming a state (e.g. "button up"), setting GPIO, then reading the button and changing GPIO accordingly – it would mean that the LED would be flickering on and off, albeit too fast to see. Using a shadow register is a neat solution that avoids this problem, as well as any read-modify-write concerns, since the physical register (GPIO) is only ever updated with the correctly determined value.

Debouncing

In most applications, you want your code to respond to transitions; some action should be triggered when a button is pressed or a switch is toggled. This presents a problem when interacting with real, physical switches, because their contacts *bounce*. When most switches change, the contacts in the switch will make and break a number of times before settling into the new position. This contact bounce is generally too fast for the human eye to see, but microcontrollers are fast enough to react to each of these rapid, unwanted transitions.

Dealing with this problem is called switch *debouncing*.

The following picture is a recording of an actual switch bounce, using a common pushbutton switch:

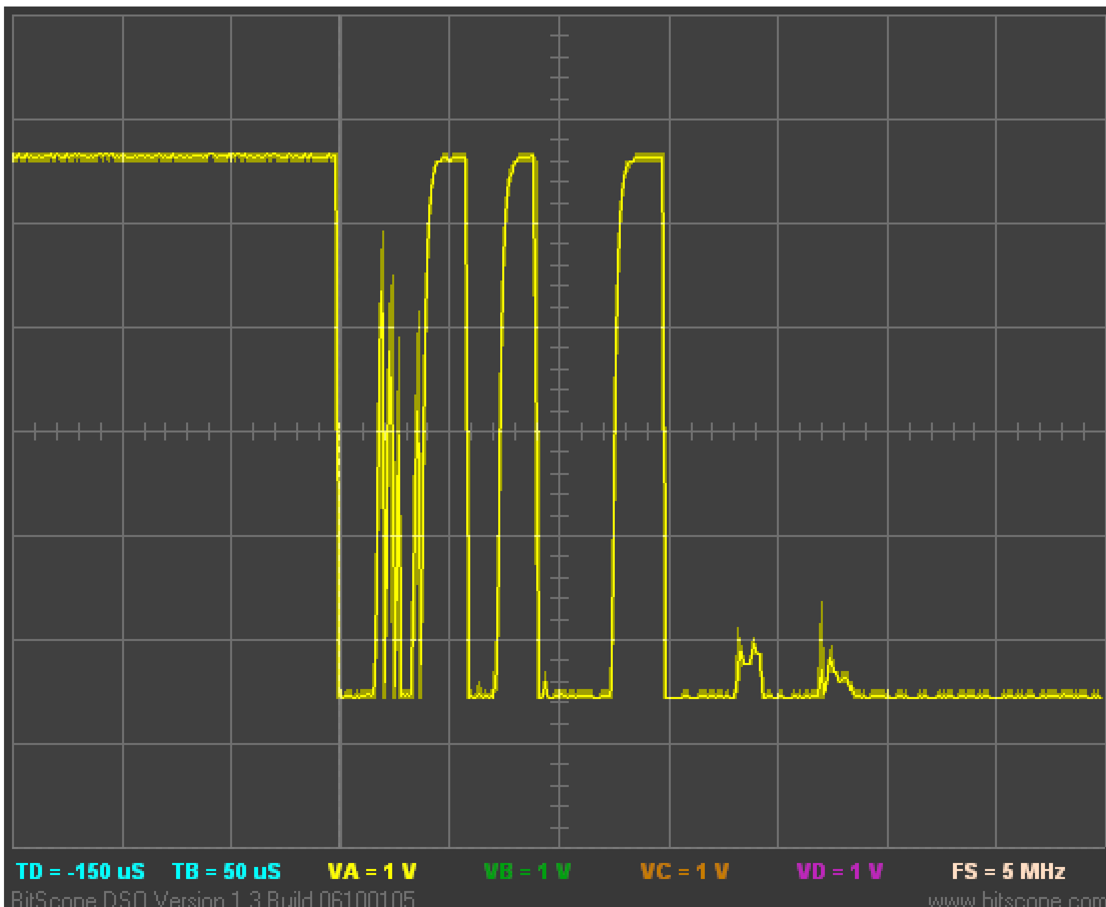

The switch transitions several times before settling into the new state (low), after around 250 µs.

A similar problem can be caused by *electromagnetic interference (EMI)*. Unwanted spikes may appear on an input line, due to electromagnetic noise, especially (but not only) when switches or sensors are some distance from the microcontroller. But any solution which deals effectively with contact bounce will generally also remove or ignore input spikes caused by EMI.

Hardware debouncing

Debouncing is effectively a filtering problem; you want to filter out fast transitions, leaving only the slower changes that result from intentional human input.

That suggests a low-pass filter; the simplest of which consists of a resistor and a capacitor (an "*RC filter*").

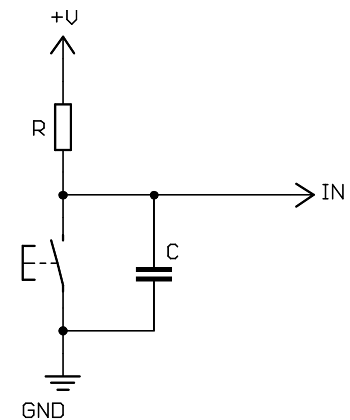

To debounce a normally-open pushbutton switch, pulled high by a resistor, the simplest hardware solution is to place a capacitor directly across the switch, as shown at right.

In theory, that's all that's needed. The idea is that, when the switch is pressed, the capacitor is immediately discharged, and the input will go instantly to 0 V. When the contacts bounce open, the capacitor will begin to charge, through the resistor. The voltage at the input pin is the voltage across the capacitor:

$$V_{in} = V_{DD}\left(1 - e^{-t/RC}\right)$$

This is an exponential function with a *time constant* equal to the product RC.

The general I/O pins on the PIC12F508/9 act as TTL inputs: given a 5 V power supply, any input voltage between 0 V and 0.8 V reads as a '0' (according to parameter D030 in the data sheet).

As long as the input voltage remains below 0.8 V, the PIC will continue to read '0', which is what we want, to avoid transitions to '1' due to switch bounce.

Solving the above equation for V_{DD} = 5.0 V and V_{in} = 0.8 V gives t = 0.174RC.

This is the maximum time that the capacitor can charge, before the input voltage goes higher than that allowed for a logical '0'. That is, it's the longest 'high' bounce that will be filtered out.

In the pushbutton press recorded above, the longest 'high' bounce is approx. 25 µs. Assuming a pull-up resistance of 10 kΩ, as in the original circuit, we can solve for C = 25 µs ÷ (0.174 × 10 kΩ) = 14.4 nF. So, in theory, any capacitor 15 nF or more could be used to effectively filter out these bounces.

In practice, you don't really need all these calculations. As a rule of thumb, if you choose a time constant several times longer than the maximum settling time (250 µs in the switch press above), the debouncing will be effective. So, for example, 1 ms would be a reasonable time constant to aim for here – it's a few times longer than the settling time, but still well below human perception (no one will notice a 1 ms delay after pressing a button).

To create a time constant of 1 ms, you can use a 10 kΩ pull-up resistor with a 100 nF capacitor:

10 kΩ × 100 nF = 1 ms

Testing the above circuit, with R = 10 kΩ, C = 100 nF and using the same pushbutton switch as before, gave the following response:

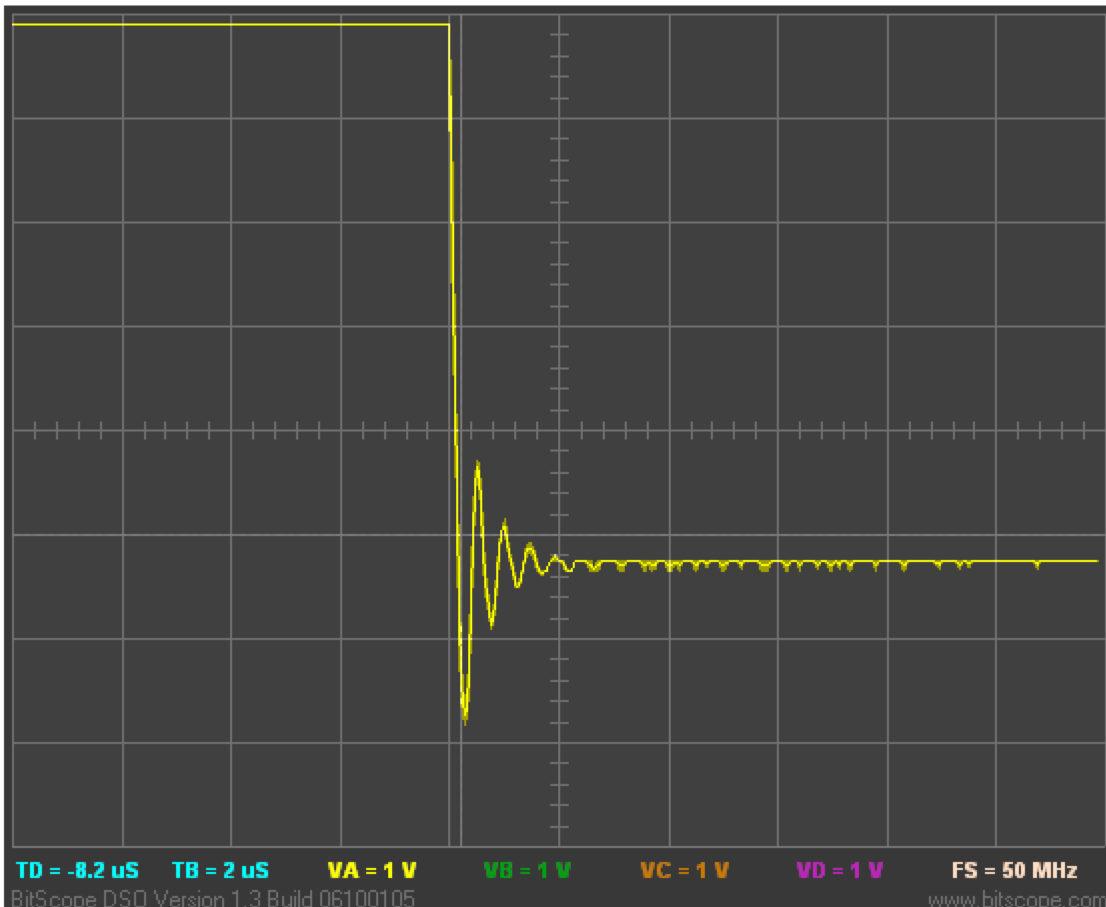

Sure enough, the bounces are all gone, but there is now an overshoot – a ringing at around 2 MHz, decaying in around 2 µs. What's going on?

The problem is that the description above is idealised. In the real world, capacitors and switches and the connections between them all have resistance, so the capacitor cannot discharge instantly (which is why the edge on the high → low transition shown above is not, and can never be, exactly vertical). More significantly, every component and interconnection has some inductance. The combination of inductance and capacitance leads to oscillation (the 2 MHz ringing). Inductance has the effect of maintaining current flow. When the switch contacts are closed, a high current rapidly discharges the capacitor. The inductance causes this current flow to continue beyond the point where the capacitor is fully discharged, slightly charging in the opposite direction, making V_{in} go (briefly) negative. Then it reverses, the capacitor discharging in the opposite direction, overshooting again, and so on – the oscillation losing energy to resistance and quickly dying away.

So – is this a problem? Yes!

According to the PIC12F508/9 data sheet, the absolute minimum voltage allowed on any input pin is -0.3 V. But the initial overshoot in the pushbutton press response, shown above, is approx. -1.5 V. That means that this simple debounce circuit is presenting voltages outside the 12F508/9's absolute maximum ratings. You might get away with it. Or you might end up with a fried PIC!

To avoid this problem, we need to limit the discharge current from the capacitor, since it is the high discharge current that is working through stray inductance to drive the input voltage to a dangerously low level.

In the circuit shown at right, the discharge current is limited by the addition of resistor R2.

We still want the capacitor to discharge much more quickly than it charges, since the circuit is intended to work essentially the same way as the first – a fast discharge to 0 V, followed by much slower charging during 'high' bounces. So we should have R2 much smaller than R1.

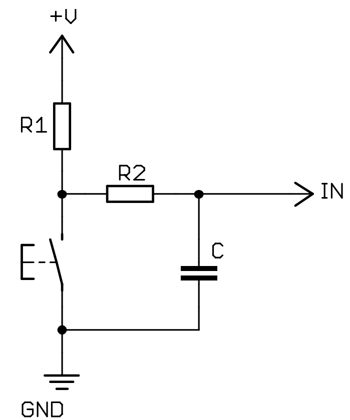

The following oscilloscope trace shows the same pushbutton response as before, with R1 = 10 kΩ, R2 = 100 Ω and C = 100 nF:

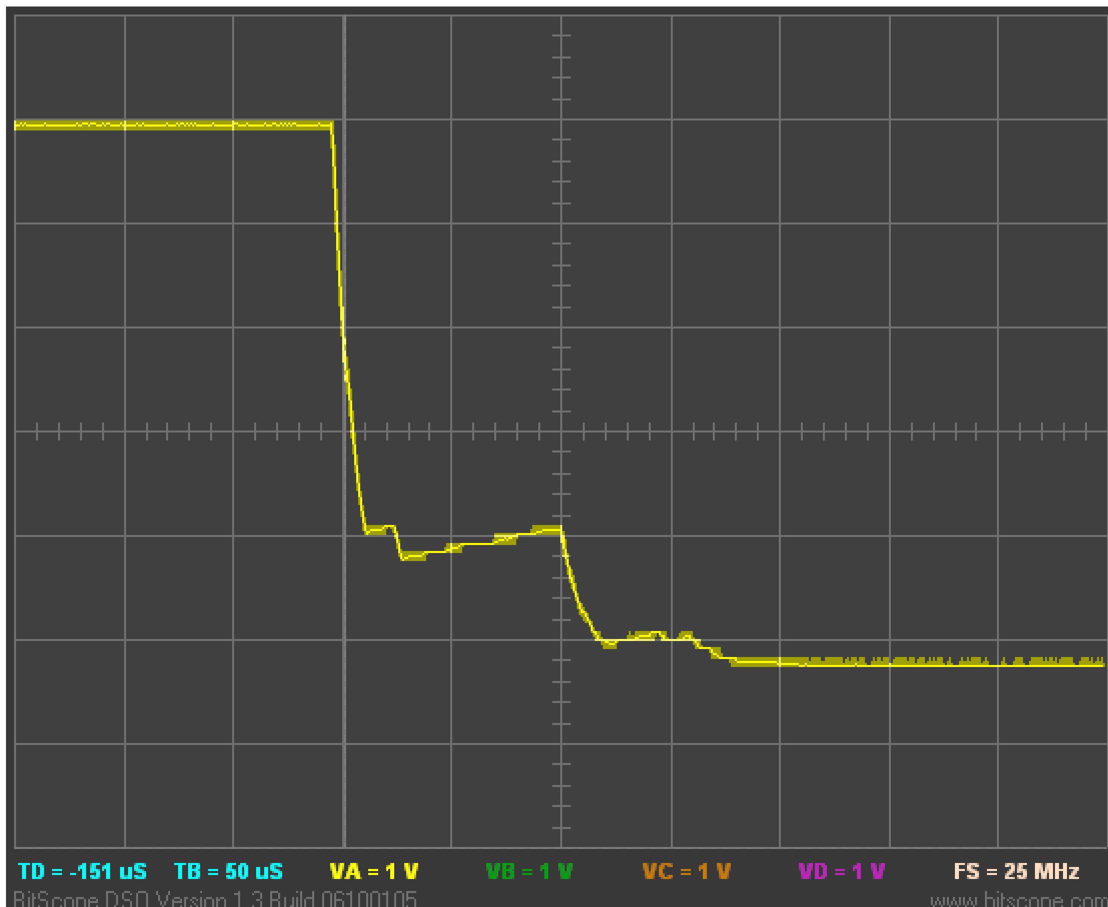

The ringing has been eliminated.

Instead of large steps from low to high, the bounces show as "ramps", of up to 75 µs, where the voltage rises by up to 0.4 V.

This effect could be reduced, and the decline from high to low made smoother, by adjusting the values of R1, R2 and C. But small movements up and down, of a fraction of a volt, will never be eliminated. And the fact that the high → low transition takes time to settle is a problem for the PIC's inputs.

With a 5 V power supply, according to the PIC12F508/9 data sheet, a voltage between 0.8 V and 2.0 V on a TTL-compatible input (any of the general I/O pins) is undefined. Voltages between 0.8 V and 2.0 V could read as either a '0' or a '1'. If we can't guarantee what value will be read, we can't say that the switch has been debounced; it's still an unknown.

An effective solution to this problem is to use a Schmitt trigger buffer, such as a 74LS14 inverter, as shown in the circuit on the right.

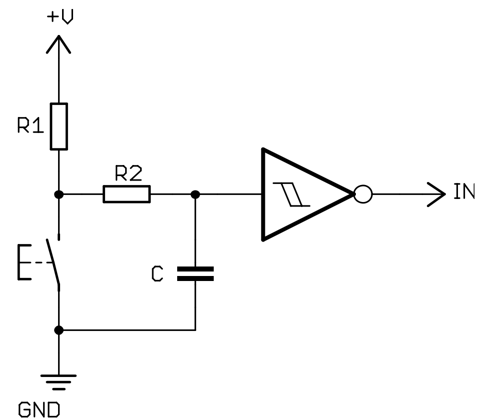

A Schmitt trigger input displays *hysteresis*; on the high → low transition, the buffer will not change state until the input falls to a low voltage threshold (say 0.8 V). It will not change state again, until the input rises to a high voltage threshold (say 1.8 V).

That means that, given a slowly changing input signal, which is generally falling, with some small rises on the way down (as in the trace above), only a single transition will appear at the buffer's output. Similarly, a Schmitt trigger will clean up slowly rising, noisy input signal, producing a single sharp transition, at the correct TTL levels, suitable for interfacing directly to the PIC.

Of course, if you use a Schmitt trigger inverter, as shown here, you must reverse your program's logic: when the switch is pressed, the PIC will see a '1' instead of a '0'.

Note that when some of the PIC12F508/9's pins are configured for special function inputs, instead of general purpose inputs, they use Schmitt trigger inputs. For example, pin 4 of the 12F508/9 can be configured as an external reset line ($\overline{\text{MCLR}}$) instead of GP3. When connecting a switch for external $\overline{\text{MCLR}}$, you only need an RC filter for debouncing; the Schmitt trigger is built into the reset circuitry on the PIC.

Software debouncing

One of the reasons to use microcontrollers is that they allow you to solve what would otherwise be a hardware problem, in software. A good example is switch debouncing.

If the software can ignore input transitions due to contact bounce or EMI, while responding to genuine switch changes, no external debounce circuitry is needed. As with the hardware approach, the problem is essentially one of filtering; we need to ignore any transitions too short to be 'real'.

But to illustrate the problem, and provide a base to build on, we'll start with some code with no debouncing at all.

Suppose the task is to toggle the LED on GP1, once, each time the button on GP3 is pressed.

In pseudo-code, this could be expressed as:

```
do forever
      wait for button press
      toggle LED
      wait for button release
end
```

Note that it is necessary to wait for the button to be released before restarting the loop, so that the LED should only toggle once per button press. If we didn't wait for the button to be released before continuing, the LED would continue to toggle as long as the button was held down; not the desired behaviour.

Here is some code which implements this:

```
;***** VARIABLE DEFINITIONS
        UDATA_SHR
sGPIO   res 1                   ; shadow copy of GPIO


;************************************************************************
RESET   CODE    0x000           ; effective reset vector
        movwf   OSCCAL          ; update OSCCAL with factory cal value
```

```
;***** MAIN PROGRAM

;***** Initialisation
start
        clrf    GPIO            ; start with LED off
        clrf    sGPIO           ;   update shadow
        movlw   b'111101'       ; configure GP1 (only) as an output
        tris    GPIO            ; (GP3 is an input)

;***** Main loop
loop

waitdn  btfsc   GPIO,3          ; wait until button pressed (GP3 low)
        goto    waitdn

        movf    sGPIO,w
        xorlw   b'000010'       ; toggle LED on GP1
        movwf   sGPIO           ;   using shadow register
        movwf   GPIO

waitup  btfss   GPIO,3          ; then released (GP3 high)
        goto    waitup

        goto    loop            ; repeat forever
```

If you build this program and test it, you will find that it is difficult to reliably change the LED when you press the button; sometimes it will change, other times not. This is due to contact bounce.

Debounce delay

The simplest approach to software debouncing is to not attempt to detect the bounces at all.

Instead, you can estimate the maximum time the switch could possibly take to settle, and then simply wait at least that long, after detecting the first transition. If the wait time, or delay, is longer than the maximum possible settling time, then you can be sure that, after the delay, the switch will have finished bouncing.

It's simply a matter of adding a suitable debounce delay, after each transition is detected, as in the following pseudo-code:

```
do forever
      wait for button press
      toggle LED
      delay debounce_time
      wait for button release
      delay debounce_time
end
```

Note that the LED is toggled immediately after the button press is detected. There's no need to wait for debouncing. By acting on the button press as soon as it is detected, the user will experience as fast a response as possible.

But it is important to ensure that the "button pressed" state is stable (debounced), before waiting for button release. Otherwise, the first bounce after the button press would be seen as a release.

The necessary minimum delay time depends on the characteristics of the switch. For example, the switch tested above was seen to settle in around 250 µs. Repeated testing showed no settling time greater than 1 ms, but it's difficult to be sure of that, and perhaps a different switch, say that used in production hardware, rather than the prototype, may behave differently. So it's best to err on the safe side, and choose the longest delay we can get away with. People don't notice delays of 20 ms or less (flicker is only barely perceptible at 50Hz, corresponding to a 20 ms delay), so a good choice is probably 20 ms.

As you can see, choosing a suitable debounce delay is not an exact science!

The above code can be modified to call the 10 ms delay module we developed in lesson 3, as follows:

```
loop

waitdn  btfsc   GPIO,3          ; wait until button pressed (GP3 low)
        goto    waitdn

        movf    sGPIO,w
        xorlw   b'000010'       ; toggle LED on GP1
        movwf   sGPIO           ;   using shadow register
        movwf   GPIO

        movlw   .2
        pagesel delay10
        call    delay10         ; delay 20ms to debounce (GP3 low)
        pagesel $

waitup  btfss   GPIO,3          ; wait until button released (GP3 high)
        goto    waitup          ;   before continuing

        movlw   .2
        pagesel delay10
        call    delay10         ; delay 20ms to debounce (GP3 high)
        pagesel $

        goto    loop            ; repeat forever
```

If you build and test this code, you should find that the LED now reliably changes state every time you press the button.

Counting algorithm

There are a couple of problems with using a fixed length delay for debouncing.

Firstly, the need to be "better safe than sorry" means making the delay as long as possible, and probably slowing the response to switch changes more than is really necessary, potentially affecting the feel of the device you're designing.

More importantly, the delay approach cannot differentiate between a glitch and the start of a switch change. As discussed, spurious transitions can be caused be EMI, or electrical noise – or a momentary change in pressure while a button is held down.

A commonly used approach, which avoids these problems, is to regularly read (or *sample*) the input, and only accept that the switch is in a new state, when the input has remained in that state for some number of times in a row. If the new state isn't maintained for enough consecutive times, it's considered to be a glitch or a bounce, and is ignored.

For example, you could sample the input every 1 ms, and only accept a new state if it is seen 10 times in a row; i.e. high or low for a continuous 10 ms.

To do this, set a counter to zero when the first transition is seen. Then, for each sample period (say every 1 ms), check to see if the input is still in the desired state and, if it is, increment the counter before checking again. If the input has changed state, that means the switch is still bouncing (or there was a glitch), so the counter is set back to zero and the process restarts. The process finishes when the final count is reached, indicating that the switch has settled into the new state.

The algorithm can be expressed in pseudo-code as:

```
count = 0
while count < max_samples
      delay sample_time
      if input = required_state
            count = count + 1
      else
            count = 0
end
```

Here is the modified "toggle an LED" main loop, illustrating the use of this counting debounce algorithm:

```
loop
        banksel db_cnt
db_dn   clrf    db_cnt          ; wait until button pressed (GP3 low)
        clrf    dc1             ; debounce by counting:
dn_dly  incfsz  dc1,f           ;   delay 256x3 = 768us.
        goto    dn_dly
        btfsc   GPIO,3          ;   if button up (GP3 set),
        goto    db_dn           ;       restart count
        incf    db_cnt,f        ;   else increment count
        movlw   .13             ;   max count = 10ms/768us = 13
        xorwf   db_cnt,w        ;   repeat until max count reached
        btfss   STATUS,Z
        goto    dn_dly

        movf    sGPIO,w
        xorlw   b'000010'       ; toggle LED on GP1
        movwf   sGPIO           ;   using shadow register
        movwf   GPIO

db_up   clrf    db_cnt          ; wait until button released (GP3 high)
        clrf    dc1             ; debounce by counting:
up_dly  incfsz  dc1,f           ;   delay 256x3 = 768us.
        goto    up_dly
        btfss   GPIO,3          ;   if button down (GP3 clear),
        goto    db_up           ;       restart count
        incf    db_cnt,f        ;   else increment count
        movlw   .13             ;   max count = 10ms/768us = 13
        xorwf   db_cnt,w        ;   repeat until max count reached
        btfss   STATUS,Z
        goto    up_dly

        goto    loop            ; repeat forever
```

There are two debounce routines here; one for the button press, the other for button release. The program first waits for a pushbutton press, debounces the press, then toggles the LED before waiting for the pushbutton to be released, and then debouncing the release.

The only difference between the two debounce routines is the input test: '`btfsc GPIO,3`' when testing for button up, versus '`btfss GPIO,3`' to test for button down.

The above code demonstrates one method for counting up to a given value (13 in this case):

The count is zeroed at the start of the routine.

It is incremented within the loop, using the '`incf`' instruction – "**inc**rement **f**ile register". As with many other instructions, the incremented result can be written back to the register, by specifying '`,f`' as the destination, or to W, by specifying '`,w`' – but normally you would use it as shown, with '`,f`', so that the

count in the register is incremented. The baseline PICs also provide a 'decf' instruction – "**dec**rement **f**ile register", which is similar to 'incf', except that it performs a decrement instead of increment.

We've seen the 'xorwf' instruction before, but not used in quite this way. The result of exclusive-oring any binary number with itself is zero. If any dissimilar binary numbers are exclusive-ored, the result will be non-zero. Thus, XOR can be used to test for equality, which is how it is being used here. First, the maximum count value is loaded into W, and then this max count value in W is xor'd with the loop count. If the loop counter has reached the maximum value, the result of the XOR will be zero. Note that we do not care what the result of the XOR actually is, only whether it is zero or not. And we certainly do not want to overwrite the loop counter with the result, so we specify ',w' as the destination of the 'xorwf' instruction – writing the result to W, effectively discarding it.

To check whether the result of the XOR was zero (which will be true if the count has reached the maximum value), we use the 'btfss' instruction to test the zero flag bit, Z, in the STATUS register.

Alternatively, the debounce loop could have been coded by initialising the loop counter to the maximum value at the start of the loop, and using 'decfsz' at the end of the loop, as follows:

```
db_dn   ; wait until button pressed (GP3 low), debounce by counting:
        movlw   .13             ; max count = 10ms/768us = 13
        movwf   db_cnt
        clrf    dc1
dn_dly  incfsz  dc1,f           ; delay 256x3 = 768us.
        goto    dn_dly
        btfsc   GPIO,3          ; if button up (GP3 set),
        goto    db_dn           ;   restart count
        decfsz  db_cnt,f        ; else repeat until max count reached
        goto    dn_dly
```

That's two instructions shorter, and at least as clear, so it's a better way to code this routine.

But in certain other situations it really is better count up to a given value, so it's worth knowing how to do that, including the use of XOR to test for equality, as shown above.

Internal Pull-ups

The use of pull-up resistors is so common that most modern PICs make them available internally, on at least some of the pins.

By using internal pull-ups, we can do away with the external pull-up resistor, as shown in the circuit on the right.

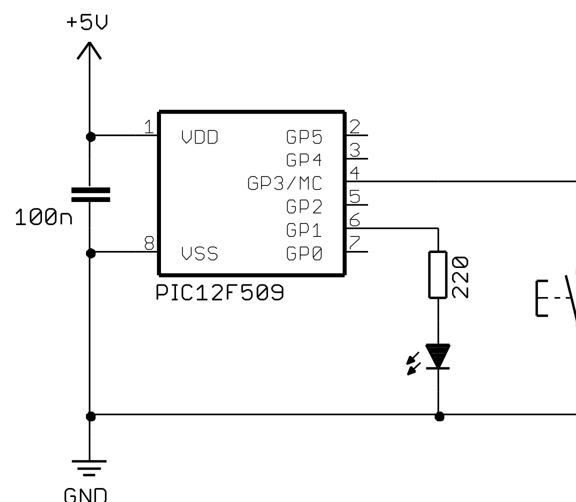

Strictly speaking, the internal pull-ups are not simple resistors. Microchip refer to them as "weak pull-ups"; they provide a small current which is enough to hold a disconnected, or *floating*, input high, but does not strongly resist any external signal trying to drive the input low. This current is typically 250 µA on most input pins (parameter D070), or up to 30 µA on GP3, when configured with internal pull-ups enabled.

The internal weak pull-ups are controlled by the $\overline{\text{GPPU}}$ bit in the OPTION register:

	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
OPTION	$\overline{\text{GPWU}}$	$\overline{\text{GPPU}}$	T0CS	T0SE	PSA	PS2	PS1	PS0

The OPTION register is used to control various aspects of the PIC's operation, including the timer (which will be introduced in the next lesson) and of course weak pull-ups.

Like TRIS, OPTION is not directly addressable, is write-only, and can only be written using a special instruction: '`option`' – "load **option** register".

By default (after a power-on or reset), $\overline{\text{GPPU}}$ = 1 and the internal pull-ups are disabled.
To enable internal pull-ups, clear $\overline{\text{GPPU}}$.

Assuming no other options are being set (leaving all the other bits at the default value of '1'), internal pull-ups are enabled by:

```
        movlw   b'10111111'     ; enable internal pull-ups
                ; -0------          pullups enabled (/GPPU = 0)
        option
```

Note the way that this has been commented: the line with '`-0------`' makes it clear that only bit 6 ($\overline{\text{GPPU}}$) is relevant to enabling or disabling the internal pull-ups, and that they are enabled by clearing this bit.

In the PIC12F508/9, internal pull-ups are only available on GP0, GP1 and GP3.

Internal pull-ups on the baseline PICs are not selectable by pin; they are either all on, or all off. However, if a pin is configured as an output, the internal pull-up is disabled for that pin (preventing excess current being drawn).

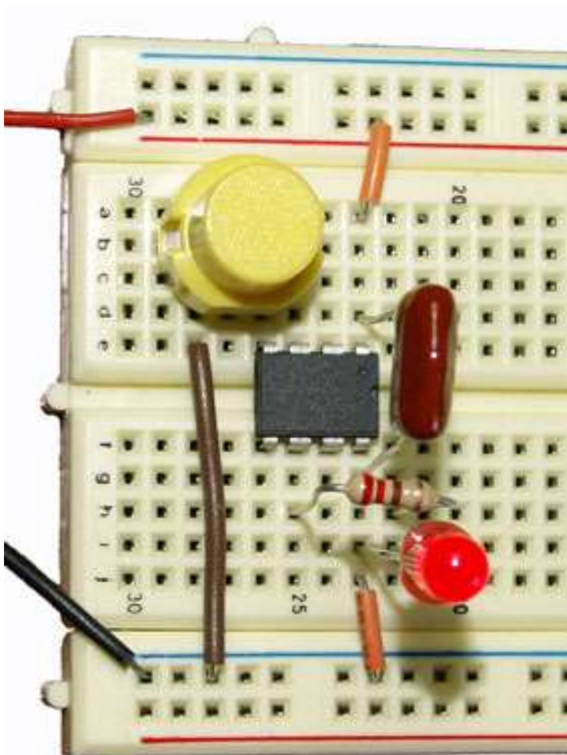

If you want to test that the internal pull-ups are working, you will need to remove the PIC from the low pin count demo board, and place it in your own circuit, with no external pull-up resistor.

For example, you could build it using prototyping breadboard, as illustrated on the left.

Note that this minimal circuit, diagrammed above and illustrated here, does not include a current-limiting resistor between GP3 and the pushbutton. As discussed earlier, that's generally ok, but to be safe and protect against such things as electrostatic discharge, it's good practice to include a current limiting resistor, of around 1 kΩ, between the PIC pin and the pushbutton.

But as this example illustrates, functional PIC-based circuits really can need very few external components!

Complete program

Here's the complete "Toggle an LED" program, illustrating how to read and debounce a simple switch connected through an internal pull-up:

```
;************************************************************************
;                                                                       *
;   Filename:       BA_L4-Toggle_LED-int_pu.asm                         *
;   Date:           16/9/07                                             *
;   File Version:   1.0                                                 *
;                                                                       *
;   Author:         David Meiklejohn                                    *
;   Company:        Gooligum Electronics                                *
;                                                                       *
;************************************************************************
;                                                                       *
;   Architecture:   Baseline PIC                                        *
;   Processor:      12F508/509                                          *
;                                                                       *
;************************************************************************
;                                                                       *
;   Files required: none                                                *
;                                                                       *
;************************************************************************
;                                                                       *
;   Description:    Lesson 4, example 5                                 *
;                                                                       *
;   Demonstrates use of internal pull-ups plus debouncing               *
;                                                                       *
;   Toggles LED when pushbutton is pressed (low) then released (high)   *
;   Uses counting algorithm for switch debounce                         *
;                                                                       *
;************************************************************************
;                                                                       *
;   Pin assignments:                                                    *
;       GP1 - LED                                                       *
;       GP3 - pushbutton switch                                         *
;                                                                       *
;************************************************************************

    list        p=12F509
    #include    <p12F509.inc>

                ; int reset, no code protect, no watchdog, 4Mhz int clock
    __CONFIG    _MCLRE_OFF & _CP_OFF & _WDT_OFF & _IntRC_OSC


;***** VARIABLE DEFINITIONS
        UDATA_SHR
sGPIO   res 1                   ; shadow copy of GPIO

        UDATA
db_cnt  res 1                   ; debounce counter
dc1     res 1                   ; delay counter


;************************************************************************
RESET   CODE    0x000           ; effective reset vector
        movwf   OSCCAL          ; update OSCCAL with factory cal value
```

```
;***** MAIN PROGRAM

;***** Initialisation
start
        movlw   b'10111111'     ; enable internal pull-ups
                ; -0------          pullups enabled (/GPPU = 0)
        option
        clrf    GPIO            ; start with LED off
        clrf    sGPIO           ;   update shadow
        movlw   b'111101'       ; configure GP1 (only) as an output
        tris    GPIO            ; (GP3 is an input)

;***** Main loop
loop
        banksel db_cnt          ; select data bank for this section

db_dn   ; wait until button pressed (GP3 low), debounce by counting:
        movlw   .13             ; max count = 10ms/768us = 13
        movwf   db_cnt
        clrf    dc1
dn_dly  incfsz  dc1,f           ; delay 256x3 = 768us.
        goto    dn_dly
        btfsc   GPIO,3          ; if button up (GP3 set),
        goto    db_dn           ;   restart count
        decfsz  db_cnt,f        ; else repeat until max count reached
        goto    dn_dly

        ; toggle LED on GP1
        movf    sGPIO,w
        xorlw   b'000010'       ; toggle shadow register
        movwf   sGPIO
        movwf   GPIO            ; write to port

db_up   ; wait until button released (GP3 high), debounce by counting:
        movlw   .13             ; max count = 10ms/768us = 13
        movwf   db_cnt
        clrf    dc1
up_dly  incfsz  dc1,f           ; delay 256x3 = 768us.
        goto    up_dly
        btfss   GPIO,3          ; if button down (GP3 clear),
        goto    db_up           ;   restart count
        decfsz  db_cnt,f        ; else repeat until max count reached
        goto    up_dly

        ; repeat forever
        goto    loop

        END
```

That’s it for reading switches for now. There’s plenty more to explore, of course, such as reading keypads and debouncing multiple switch inputs – topics to explore later.

But in the next lesson we’ll look at the PIC12F508/9’s timer module.