Unit 7

Exceptions & Interrupts

Disclaimer 1

- This is just an introduction to the topic of interrupts. You are not meant to master these right now but just start to use them

- We will cover more about them as we investigate other modules that can make use of them

Exceptions

- In computer systems we may NOT know when
  - External hardware events will occur.
    - Can you think of an example?
  - Errors will occur

- Exception processing refers to
  - Handling events whose timing we cannot predict

- 3 questions to answer:
  - Q: Who detects these events and how? A: The hardware
  - Q: How do we respond? A: Calling a pre-defined SW function
  - Q: What is the set of possible events? A: Specific to each processor

An Analogy

- Scenario:
  - You're studying (i.e. listening to music and watching Netflix) but all of a sudden you get a text message. What do you do?
  - You stop what your doing and message back
  - When you're done you go back to studying (i.e. playing a video game or going to get coffee)
- This is what computers do when an exception/interrupt occurs

What are Exceptions?

- **Definition**: Any event that causes a break in normal execution
  - "Exceptions" is a broad term to catch many kinds of events that *interrupt* normal software execution
- Examples
  - **Hardware Interrupts / External Events [Focus for today]**
    - PC: Handling a keyboard press, mouse moving, USB data transfer, etc.
    - Arduino: Value change on a pin, ADC conversion done, Timers, etc.
  - Error Conditions [Focus for some other time]
    - Invalid address, illegal memory access, arithmetic error (e.g. divide by 0)
  - System Calls / Traps [Focus for some other time]
    - User applications calling OS code
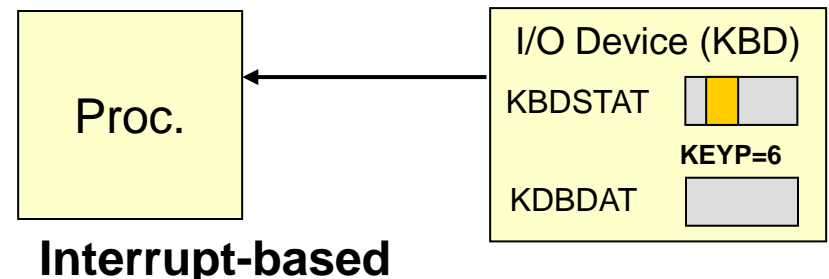
Interrupt Exceptions

- Two methods for processor and I/O devices to notify each other of events
  - Polling loop or "busy" loop (responsibility on proc.)
    - Processor has responsibility of checking each I/O device
    - Many I/O events happen infrequently (1-10 ms) with respect to the processors ability to execute instructions (1-100 ns) causing the loop to execute many times
  - Interrupts (responsibility on I/O device)
    - I/O device notifies processor only when it needs attention

Suppose: A Keyboard interface has a control/status register (KBDCSR) which sets a bit (say, the 6th) when a key is pressed (KEYP=6). How would we "poll" to check if a key is pressed.
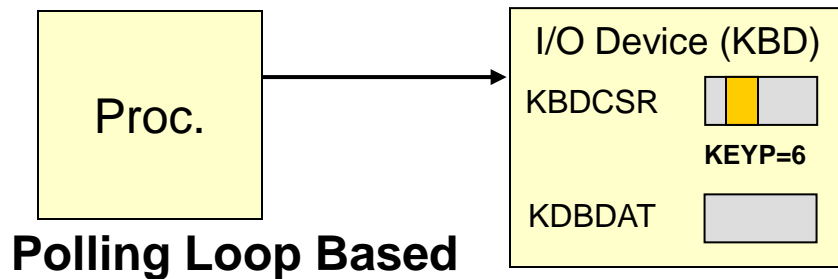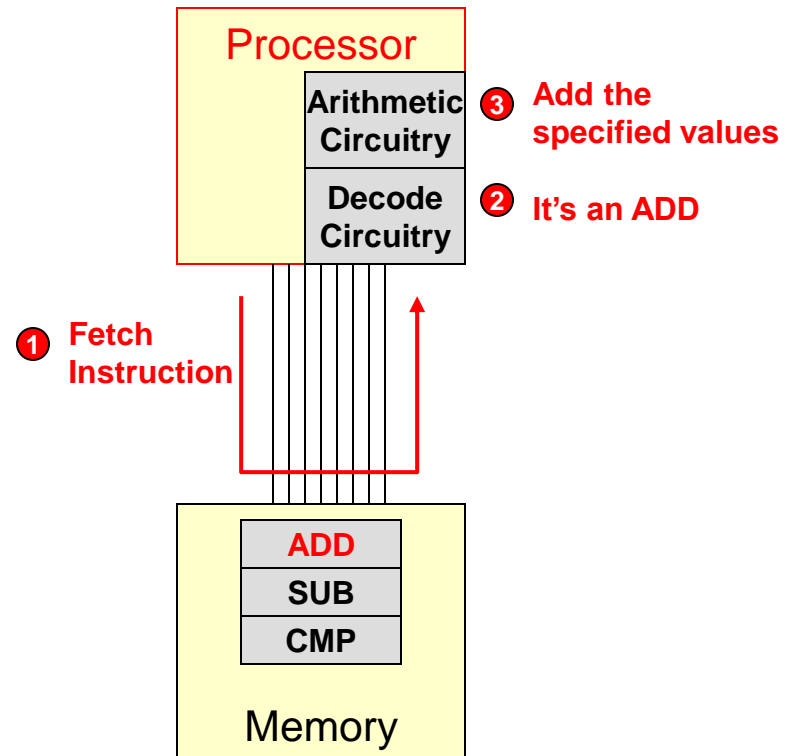
With Interrupts: We can ask the Keyboard to "interrupt" the processor when a key is pressed, so the processor doesn't have to sit there polling.

```
while ( (KBDCSR & (1 << KEYP)) == 0);
```

I/O Device (KBD)

Proc.

KBDCSR

KEYP=6

KDBDAT

Polling Loop Based

Proc.

I/O Device (KBD)
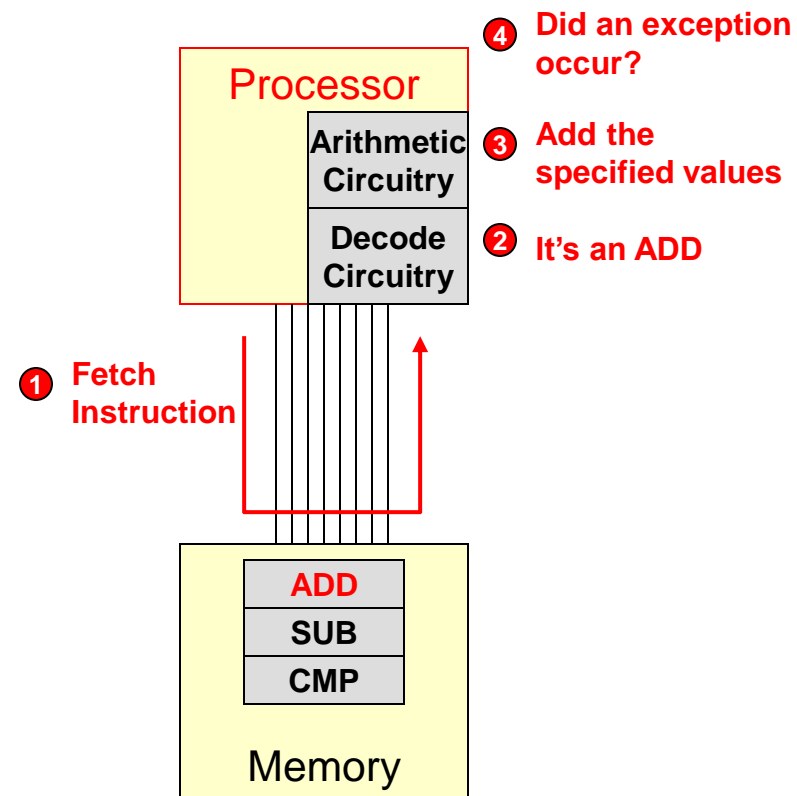
KBDSTAT

KEYP=6

KDBDAT

Interrupt-based

Recall: Instruction Cycle

- Processor hardware performs the same 3-step process over and over again as it executes a software program

  - Fetch an instruction from memory

  - Decode the instruction
    - Is it an ADD, SUB, etc.?

  - Execute the instruction
    - Perform the specified operation

- This process is known as the **Instruction Cycle**

Processor

| Arithmetic Circuitry | ❸ Add the specified values |
| Decode Circuitry | ❷ It's an ADD |

❶ Fetch Instruction

| ADD |
| SUB |
| CMP |

Memory

HW Detects Exceptions

- There's actually a 4th step
- After finishing each instruction the processor hardware checks for interrupts or errors automatically (i.e. this is built into the hardware)
  - Fetch an instruction from memory
  - Decode the instruction
    - Is it an ADD, SUB, etc.?
  - Execute the instruction
    - Perform the specified operation
  - Check for exceptions
    - If so, pause the current program and go execute other software to deal with the exception

Processor

Arithmetic Circuitry

Decode Circuitry

4 Did an exception occur?

3 Add the specified values

2 It's an ADD

1 Fetch Instruction

ADD

SUB

CMP

Memory

SW Handles Exceptions

- When exceptions occur, what should happen?
  - We could be anywhere in our software program…who knows where

- Common approach…
  - 1. Save place in current code and disable other interrupts
  - 2. Automatically have the processor call some function/subroutine to handle the issue (a.k.a. **Interrupt Service Routine** or **ISR**)
  - 3. Reenable interrupts & resume normal processing back in original code

If an interrupt happens here…

…the processor will automatically call a predetermined function (a.k.a. ISR)

…then resume the code it was executing previously

```c
#include<avr/io.h>
#include<avr/interrupt.h>

void codeToHandleInterrupt();

int main()
{
   // this is just generic code
   // for a normal application
   PORTC |= (1 << PC2);
   int cnt = 0;
   while(1){
     if( PINC & (1 << PC2) ) {
       cnt++;
       PORTD = segments[cnt];
     }
   }
   return 0;
}

ISR()
{
  // do something in response
  // to the event
}
```

Important Point:
HW detects exceptions.
Software handles exceptions.

When Exceptions Occur…

- How does the processor know which function to call "automatically" when an interrupt occurs

- We must tell the processor in advance which function to associate (i.e. call) with the various exceptions it will check for

- Just like a waiver forms asks for an emergency contact to call if something bad happens, we indicate what function to call when an interrupt occurs

This is fictitious "keyboard" code

If an interrupt happens here…

PCINT0_vect is not an argument. It identifies the ISR as being from a change on PORTB

…and the processor will automatically call a predetermined function

```c
#include<avr/io.h>
#include<avr/interrupt.h>

unsigned char value = 0;

int main()
{
  // this is just generic code
  lcd_init();
  // enable KBD interrupts
  KEYCTRL |= (1 << KIE);
  sei();
  while(1)
    { /* do useful work here */}
  return 0;
}

ISR(PCINT0_vect)
{
  // get the key value (ASCII)
  char value = KEYDAT;
  // echo the value to the LCD
  lcd_writedata(value);
}
```

Function Calls vs. Interrupts

Normal function calls

- **Synchronous**: Called whenever the program reaches that point in the code
- Programmer can pass arguments and receive return values

Interrupts

- **Asynchronous**: Called whenever an event occurs (can be anywhere in our program when the ISR needs to be called)
  - Requires us to know in advance which ISR to call for each possible exception/interrupt
  - Use ISR(*interrupt_type*) naming scheme in the Arduino to make this association
- No arguments or return values
  - How would we know what to pass if we don't know when it will occur
  - Generally interrupts update some global variables

AVR INTERRUPT SOURCES

Interrupt Sources

- An AVR processor like the ATmega328P has numerous sources of possible interrupts, many of which you will use in upcoming labs

- Communications modules
  - The AVR has several serial communications modules built in (think of these like forerunners of modern USB interfaces)
  - Interrupts can be configured to occur when data is received, sent, etc.

- Analog-to-Digital Converter (ADC) module
  - The ADC can generate an interrupt once it's done converting the analog voltage (i.e. you start it and then it will "interrupt" you when its done) to a digital number

- External Interrupts and Pin Change Interrupts *(See next slides)*
  - Can be used to connect 3rd party devices to the system and have them generate interrupts on 0->1 or 1->0 transitions

- Timer interrupts (See next slides)
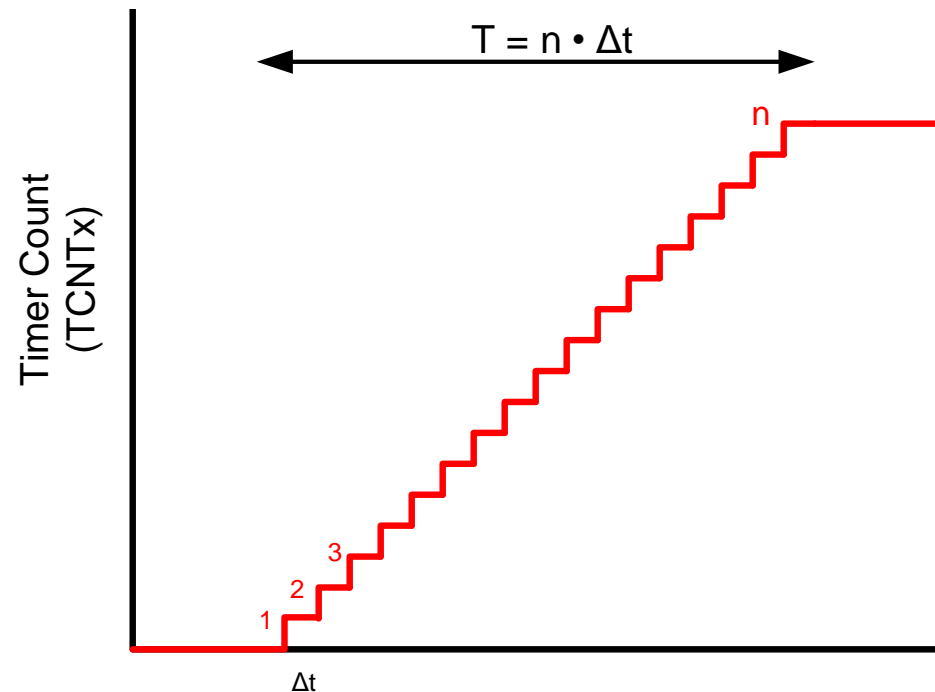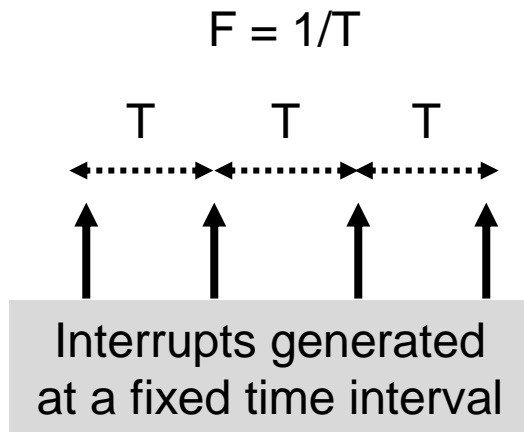 - Generate an interrupt at a regular interval
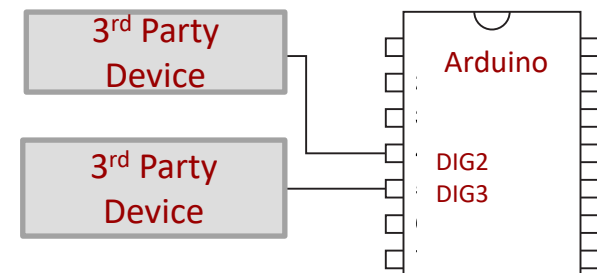
Pin Change Interrupts

- Pin Change Interrupt can detect if any pin that is part of a particular PORT (i.e. B, C, D) has changed its value
  - Interrupt if a pin changes state ($0\rightarrow1$ or $1\rightarrow0$)
  - 3 individual pin change interrupts
    - Pin Change Interrupt 0 = any bits on PORTB change
    - Pin Change Interrupt 1 = any bits on PORTC change
    - Pin Change Interrupt 2 = any bits on PORTD change
  - Interrupt only says *a* pin of the port has changed but not *which* one
    - The function that gets called can figure out what happened by reading the PINx register and AND'ing it appropriately just like we did in previous labs
  - Useful if you have to monitor a number of external sources for changes

Counter/Timer Interrupts (1)

- Most processors have some hardware counters that count at some known rate or frequency (i.e. 1 KHz) and a register that can be loaded with some upper limit, **n**.

- HW counter starts counting at 0 and generates an interrupt when it reaches the upper limit (**n**)

$F = 1/T$

Interrupts generated at a fixed time interval

$T = n \cdot \Delta t$

Timer Count (TCNTx)

n

3

2

1

Δt

Counter/Timer Interrupts (2)

- Suppose we se the timer to count at a frequency of 500 KHz.

- We can then load the upper limit register (usually referred to as OCRx) with the value, **n**.

- The timer has an internal register/variable, TCNTx, that will start at 0 and increment at the specified frequency (e.g. 500KHz). When the count reaches the upper limit (TCNTx == OCRx), an interrupt will be generated

  - If OCRx = 1000 and the frequency is 500Khz, an interrupt will occur after 1000 / 500E3 = 2 milliseconds

- The timers can be set to immediately restart at 0 (i.e. TCNTx = 0) again to generate interrupts at a regular interval

- ATmega328P has **<u>three</u>** such timers that can be used

Who You Gonna Call?

- The HW maintains a table/array (a.k.a. **interrupt vector table**) in memory
  - Each location in the table is associated with a specific interrupt
  - Each entry specifies which function to call when that interrupt occurs
- When a certain interrupt occurs, the HW automatically looks up the ISR/function to call in the table and then calls it
- More on this in your OS class (CS 350) or Architecture course (EE 457)

Page 58 on http://www.atmel.com/images/Atmel-8271-8-bit-AVR-Microcontroller-ATmega48A-48PA-88A-88PA-168A-168PA-328-328P_datasheet_Complete.pdf

Interrupt that HW Associates with this entry…	Table Entry	…Which User Defined Function to Call
Reset	0	
External Interrupt Request 0	1	ISR(INT0_vect)
External Interrupt Request 1	2	ISR(INT1_vect)
Pin Change Interrupt Request 0 (Port B)	3	ISR(PCINT0_vect)
Pin Change Interrupt Request 1 (Port C)	4	ISR(PCINT1_vect)
Pin Change Interrupt Request 2 (Port D)	5	ISR(PCINT2_vect)
Watchdog Time-out Interrupt	6	
Timer/Counter2 Compare Match A	7	
…	…	
Timer/Counter0 Compare Match A	14	ISR(TIMER0_COMPA_vect)
Timer/Counter0 Compare Match B	15	ISR(TIMER0_COMPB_vect)
Timer/Counter0 Overflow	16	ISR(TIMER0_OVF_vect)
SPI Serial Transfer Complete	17	
USART Rx Complete	18	ISR(USART_RX_vect)
USART Data Register Empty	19	ISR(USART_UDRE_vect)
USART Tx Complete	20	ISR(USART_TX_vect)
ADC Conversion Complete	21	ISR(ADC_vect)
EEPROM Ready	22	
Analog Comparator	23	
Two-wire Serial Interface	24	
Store Program Memory Read	25	

Interrupt Service Routines

- An ISR is written like any other function (almost)
  - Must be declared as an ISR for a specific interrupt by using a special name [e.g. **ISR(TIMER1_COMPA_vect)**]. This tells the compiler to fill in the **interrupt vector table** to call this code when an ADC interrupt occurs.
  - No arguments can be passed
  - No values can be returned
  - Must include the **avr/interrupt.h** header

- ISRs have access to other functions and global variables like any other function

```c
#include <avr/io.h>
#include <avr/interrupt.h>

int main()
{
  ...
}


ISR(TIMER1_COMPA_vect)
{
  // get the key value (ASCII)
  char value = KEYDAT;
  // echo the value to the LCD
  lcd_writedata(value);
}
```

To Use Arduino Interrupts

- Define the ISR in your software program
  - If an interrupt occurs for which there is no ISR, the Arduino will reset/reboot!
- During initialization you must enable the interrupt source
- Either…
  - Wait for the interrupt to occur (e.g. wait for a pin to change or the timer to reach the modulus count)
  - or Invoke behavior that will eventually lead to an interrupt (start the timer so that eventually it will generate an interrupt when the count reaches the upper limit)

Enabling Interrupts

- Each interrupt source is DISABLED by default and must be ENABLED
- For an interrupt to be handled, **<u>two</u>** "enablers" need to agree
  - **Enabler 1**: A separate "local" interrupt enable bit per source (i.e. ADC, timer, pin change, etc.)
  - **Enabler 2**: A single "global" interrupt enable bit (1-bit for entire processor called the I-bit)
- Analogy: Local judge per state but 1 supreme court for entire nation
  - Both local judge and supreme court must agree (be set to '1') for the interrupt to occur. If either are '0' then the interrupt will not occur.

**"Local" Interrupt Enable
(1 per interrupt source)**

**"Global" Interrupt Enable
(1 for entire system)**
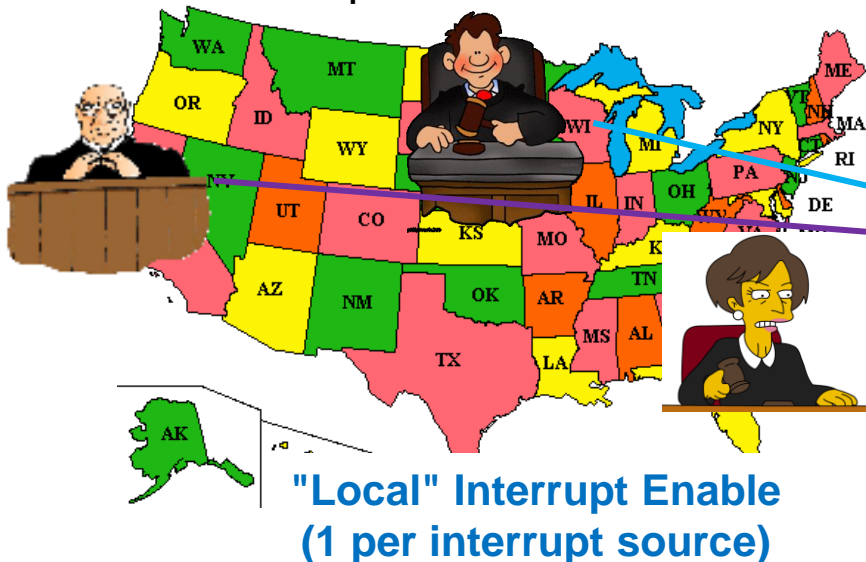
Enabling Interrupts

- All interrupt sources must be enabled before they can be used

- Each source of an interrupt has its own interrupt enable bit
  - Located in one of the control registers for the module
  - 0 = Don't use interrupts, 1 = Can interrupt
  - Example: `TIMSK1 |= (1 << OCIE1A);`

TIMSK1 Register

-	-	ICIE1	-	-	OCIE1 B	OCIE1 A	TOIE1

- Processor has a global interrupt enable bit in the status register
  - I-bit = 0 ⇒ all interrupts are ignored
  - I-bit = 1 ⇒ interrupts are allowed
  - Set or clear in C with the **`sei()`** and **`cli()`** function calls.

- **Summary**: For a module to generate an interrupt
  - The global I-bit must be a one
  - The local interrupt enable bit must be a one
  - Something must happen to cause the interrupt
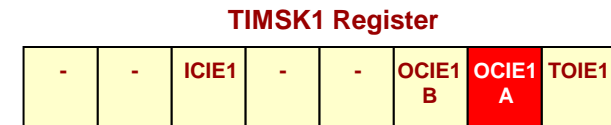
Interrupt Example

- Example: Updating a variable or printing to the LCD at a certain interval
- Tracking time without interrupts
  - Use delays
  - Is this an accurate way to track time?
- Better to use interrupts
  - Setup timer to generate an interrupt at a fixed period
  - Enable interrupts and start the hardware timer
  - The program is now free to do other things
  - After the specified time, an ISR will be called
  - The program can start several tasks, (e.g. multiple timers, an ADC conversion, etc.) and handle each when they finish via an ISR

```
int qsecs = 0;
while(1)
{
  qsecs++;
  lcd_moveto(0,0);
  lcd_stringout("0.25s elapsed");
  _delay_ms(250);
}
// is "0.25s elapsed" printed
// every 250 ms?
```

Interrupt Example

- Interrupt method:
  - #include <avr/interrupt.h>
  - Initialize the on-board hardware module
  - Enable interrupts
  - Start the timer hardware module
  - Loop using the results
    - In most application, should check if ISR actually executed by using a flag mechanism (more on this in later slides)
  - As the time elapses, the ISR associated with the timer will execute and the timer will start tracking the next time period
- Be sure to follow the special syntax for how to declare the ISR which starts with ISR and has the name of the interrupt vector (e.g. **ISR(TIMER1_COMPA_vect)**)
- **<u>Note</u>**: If you enable an interrupt but have no ISR() written the Arduino will reboot immediately when the interrupt occurs

```c
#include<avr/io.h>
#include<avr/interrupt.h>
volatile int qsecs = 0;
volatile char qsec_flag = 0;
int main() {
  // Set to CTC (repeat) mode
  TCCR1B |= (1 << WGM12);

  // Load the MAX count. Assuming prescalar
  // of 256, counting to 15625 = 0.25s
  // w/ 16 MHz clock
  OCR1A = 15625;

  // Local Timer Interrupt & global enable
  TIMSK1 |= (1 << OCIE1A);
  sei();

  // Set prescalar = 256 and start counter
  TCCR1B |= (1 << CS12);

  while (1) {
    // Do real work while waiting for interrupt
    if(qsec_flag) {
      // use qsec data variable
      qsec_flag = 0;
    }
} }
ISR(TIMER1_COMPA_vect){
  qsecs++;        // will increments every 0.25s
  qsec_flag = 1;
}
```

Interrupt Flag Bits (Skip)

- Modules usually contain an interrupt flag (IF) bit in the same register as the interrupt enable (IE) bits
  - Flag is set when the module wants to generate an interrupt.
  - Flag is cleared when the ISR is called
  - Allows the program to see if interrupts would have occurred if they were enabled (i.e. If we aren't using interrupts for the timer we can still look at the OCF1A bit to see if it would have *tried* to generate an interrupt)

TIFR1 Register

-	-	-	-	-	OCF1B	OCF1A	TOV0
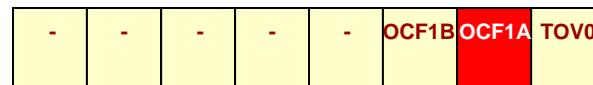
Disclaimer 2

- All processors handle interrupts differently.
  - The AVR is typical in some ways, not in others.
  - If working with a different processor, don't assume it works the same as the AVR.
  - READ THE MANUAL!

Flags and volatile variables

COMMUNICATING WITH ISRS

Communicating with ISRs & Other Code

- Global variables can be shared between main code and ISR's
- ISR's can modify the contents of a global variable and other code and check for changes in that global variable
- Common idiom: a "flag" variable to indicate a desired event has happened
  - ISR checks on every interrupt to see if the desired event occurred and only then sets a flag to 1
  - Main or other code checks the flag variable then resets it to 0

```c
#include<avr/io.h>
#include <avr/interrupt.h>
int flag; // shared variable
main()
{
  flag = 0;

  // Loop waiting for interrupt to occur
  while (flag == 0);

  // Do something
}


ISR(SOME_INTERRUPT_vect)
{
  flag = 1;
}
```

Using a Flag Variable

- A common idiom is to use a "flag" variable to indicate a desired event has happened

- Approach

  - Initialize a global variable to 0

  - ISR checks on every change to see if the desired event occurred and only then sets a flag to 1

  - Main or other code checks the flag variable then resets it to 0 awaiting the next time the event occurs

```c
int pb2flag; // shared variable

main() {
  pb2flag = 0;  // Initialize to 0
  while(1){
    // Loop waiting for flag to be set
    if (pb2flag == 1){
      pb2flag = 0; // reset flag to 0 so we
                   // can detect the next push
      stringout("button push!");
    }
    // Check for other things or do work
  }
}
ISR(PCINT0_vect)
{
  // Some bit changed, see if it is PB2
  if( (PINB & (1 << PB2)) == 0){
    pb2flag = 1;
  }
}
```

ISR Timing

- Why not just do the work in the ISR?

- Because an interrupt can not interrupt another interrupt!
  - That's a mouthful
  - When you are in an ISR no other interrupts can occur possibly delaying important events, or even losing information (e.g. a "high-speed" communications link with limited space)

- **Solution**: Never do long latency work in an ISR

```
main() {
  while(1){
    // Check for other things or do work
  }
}
ISR(PCINT0_vect)
{
  // Some bit changed, see if it is PB2
  if( (PINB & (1 << PB2) ) == 0){
    stringout("button push!");
  }
}
```

Main Point:
Get in and out of an ISR quickly.

Don't call functions that will take a long time to complete (e.g. LCD output)

Another Issue: Compiler Optimizations

- Example: When optimizing this code, compiler sees that "flag" is never modified in main (and doesn't see any "calls" to the ISR)

- Thus, the compiler will optimize the code to avoid reading "flag" from memory each time (since that is slow)

- Problem: Due to the compiler optimization our code won't work even if the ISR sets the flag to 1

- Solution: Tell the compiler that "flag" can change due to some ISR by declaring it as *volatile*

Original Code

```
int flag;
main() {
    flag = 0;
    // Loop waiting for flag non-zero
    while (flag == 0);

    // Do something
}

ISR(SOME_INTERRUPT_vect)
{
    flag = 1;
}
```

If you just look at main, would you expect this while loop to ever terminate?

Result of compiler optimization

```
int flag;
main() {
    flag = 0;
    // compiler optimized result
    while (1);

    // Do something
}
ISR(SOME_INTERRUPT_vect)
{
    flag = 1;
}
```

Another Issue: Compiler Optimizations

Original Code

- **Solution**: Tell the compiler that "flag" can change due to some ISR by declaring it as *volatile*
  - Declaring a global variable as volatile tells the compiler not to optimize the code but always get the latest value of the variable
- **Important Rule:** Use "volatile" for any global variable that is **updated in an ISR** and **used elsewhere in the code**
- **Corollary**: No need to use "volatile" for variables not used with ISRs (e.g. "buf" in the example at the right).
- Arrays are implicitly "volatile" (processor always gets latest values)

```c
char buf[17]; // not used in an ISR.
volatile int flag;
main() {
    flag = 0;
    // Loop waiting for flag non-zero
    while (flag == 0);

    // Do something
    snprintf(buf,3,"Hi");
}

ISR(SOME_INTERRUPT_vect)
{
    flag = 1;
}
```

Volatile declaration tells compiler to always look at the latest value of "flag"

Performing critical sections without be interrupted

NEED FOR ATOMIC OPERATIONS

Need for Atomic Operations

- Sometimes performing an operation requires several steps (ex. Copying bits into a register)

- If an interrupt occurs in the middle of the sequence it may see a strange value/state of the variable and do something we didn't expect

- Atomic operations are compound statements that should execute all together (not be interrupted)

```
#define MASK 0b00001111
main() {

    PORTD = 0x0f; // PORTD starts at 1's
    char x = 0x05;

    // copy lower 4 bits of x to PORTD
    PORTD &= ~MASK;
    PORTD |= (x & MASK);
    // both lines should be done "together"

    // we would expect PORTD to end w/ 5
    //  but what if interrupt occurred
    //  between these two lines

}

ISR(SOME_INTERRUPT_vect)
{
    // if lower 4 bits all = 0
    if( (PORTD & MASK) == 0){
        // do something
    }
}
```

Atomic Operations

- "Atomic" ⇒ Can't be interrupted while executing
- The problem gets worse at the assembly level since many C operations (one line of code) require multiple assembly language instructions, and interrupts can occur between them.
  - Even "x++" is actually 3 steps: Read old value of x, Add 1, Update x to new value
- Need a way to ensure all operations occur together and are not interrupted (e.g. ensure an interrupt doesn't occur in the middle)

Updated Code for Atomicity

- Solution to ensure "atomic" operation
  - **Disable interrupts using cli()**
  - **Perform the operation**
  - **Re-enable interrupts using sei()**
- The **code** between **cli()** and **sei()** that cannot be interrupted is called a **"critical section"** (since it must be done together)

```
#define MASK 0b00001111
main() {

    PORTD = 0x0f; // PORTD starts at 1's
    char x = 0x05;

    // blue code can't be interrupted
    cli();
    PORTD &= ~MASK;
    PORTD |= (x & MASK);
    sei();


    // now we can be interrupted

}

ISR(SOME_INTERRUPT_vect)
{
    // if lower 4 bits all = 0
    if( (PORTD&MASK) == 0){
        // do something
    }
}
```

Code between cli() and sei() is called a "critical section"

Can't happen during the critical section

Built-In Atomic Block

- In a larger program there are some issues that might arise
  - OK to disable interrupts, but shouldn't turn them back on if they were already disabled by some other code
- Solution: Use atomic.h and the ATOMIC_BLOCK()
- Turns interrupts off, then restores to previous state.

```
#include <util/atomic.h>

main() {

    ATOMIC_BLOCK()
    { /* like cli() */
        // interrupts now off
        // Do critical section
    } /* like sei() */

    // interrupt setting restored
}

ISR(SOME_INTERRUPT_vect)
{
}
```