FreeRTOS & Tasks

Introduction to FreeRTOS

Objective

To introduce what, why, when, and how to use Real Time Operating Systems (RTOS) as well as get you started using it with the SJSU-Dev environment.

I would like to note that this page is mostly an aggregation of information from Wikipedia and the FreeRTOS website.

What is an OS?

Operating system (OS) is system software that manages computer hardware and software resources and provides common services for computer programs. - Wikipedia

Operating systems like Linux or Windows

They have services to make communicating with Networking devices and files systems possible without having to understand how the hardware works. Operating systems may also have a means to multi-tasking by allow multiple processes to share the CPU at a time. They may also have means for allowing processes to communicate together.

What is an RTOS?

An RTOS is an operating system that meant for real time applications. They typically have fewer services such as the following:

- Parallel Task Scheduler
- Task communication (Queues or Mailboxes)
- Task synchronization (Semaphores)

Why use an RTOS?

"

You do not need to use an RTOS to write good embedded software. At some point though, as your application grows in size or complexity, the services of an RTOS might become beneficial for one or more of the reasons listed below. These are not absolutes, but opinion. As with everything else, selecting the right tools for the job in hand is an important first step in any project. In brief:

Abstract out timing information

The real time scheduler is effectively a piece of code that allows you to specify the timing characteristics of your application - permitting greatly simplified, smaller (and therefore easier to understand) application code.

• Maintainability/Extensibility

Not having the timing information within your code allows for greater maintainability and extensibility as there will be fewer interdependencies between your software modules. Changing one module should not effect the temporal behavior of another module (depending on the prioritization of your tasks). The software will also be less susceptible to changes in the hardware. For example, code can be written such that it is temporally unaffected by a change in the processor frequency (within reasonable limits).

• Modularity

Organizing your application as a set of autonomous tasks permits more effective modularity. Tasks should be loosely coupled and functionally cohesive units that within

themselves execute in a sequential manner. For example, there will be no need to break functions up into mini state machines to prevent them taking too long to execute

to completion.

Cleaner interfaces

Well defined inter task communication interfaces facilitates design and team development.

• Easier testing (in some cases)

Task interfaces can be exercised without the need to add instrumentation that may have changed the behavior of the module under test.

Code reuse

Greater modularity and less module interdependencies facilitates code reuse across projects. The tasks themselves facilitate code reuse within a project. For an example of the latter, consider an application that receives connections from a TCP/IP stack - the same task code can be spawned to handle each connection - one task per connection.

• Improved efficiency?

Using FreeRTOS permits a task to block on events - be they temporal or external to the system. This means that no time is wasted polling or checking timers when there are actually no events that require processing. This can result in huge savings in processor utilization. Code only executes when it needs to. Counter to that however is

the need to run the RTOS tick and the time taken to switch between tasks. Whether the saving outweighs the overhead or vice versa is dependent of the application. Most

applications will run some form of tick anyway, so making use of this with a tick hook function removes any additional overhead.

• Idle time

It is easy to measure the processor loading when using FreeRTOS.org. Whenever the

idle task is running you know that the processor has nothing else to do. The idle task also provides a very simple and automatic method of placing the processor into a low power mode.

Flexible interrupt handling

Deferring the processing triggered by an interrupt to the task level permits the interrupt

handler itself to be very short - and for interrupts to remain enabled while the task level

processing completes. Also, processing at the task level permits flexible prioritization -

more so than might be achievable by using the priority assigned to each peripheral by the hardware itself (depending on the architecture being used).

Mixed processing requirements

Simple design patterns can be used to achieve a mix of periodic, continuous and event driven processing within your application. In addition, hard and soft real time requirements can be met though the use of interrupt and task prioritisation.

• Easier control over peripherals

Gatekeeper tasks facilitate serialization of access to peripherals - and provide a good



• Etcetera

- FreeRTOS Website (https://www.freertos.org/FAQWhat.html)

Design Scenario

Building a controllable assembly conveyor belt



Think about the following system. Reasonable complex, right?

Without a scheduler

- ? Small code size.
- **?** No reliance on third party source code.
- ? No RTOS RAM, ROM or processing overhead.
- ? Difficult to cater for complex timing requirements.
- ? Does not scale well without a large increase in complexity.
- ? Timing hard to evaluate or maintain due to the inter-dependencies between the different functions.

With a scheduler

? Simple, segmented, flexible, maintainable design with few inter-dependencies.

? Processor utilization is automatically switched from task to task on a most urgent need basis with no explicit action required within the application source code.

? The event driven structure ensures that no CPU time is wasted polling for events that have not occurred.

Processing is only performed when there is work needing to be done.

* Power consumption can be reduced if the idle task places the processor into power save (sleep) mode, but may also be wasted as the tick interrupt will sometimes wake the processor unnecessarily.

* The kernel functionality will use processing resources. The extent of this will depend on the chosen kernel tick frequency.

? This solution requires a lot of tasks, each of which require their own stack, and many of which require a queue on which events can be received. This solution therefore uses a lot of RAM.

? Frequent context switching between tasks of the same priority will waste processor cycles.

FreeRTOS Tasks

What is an FreeRTOS Task?

A FreeRTOS task is a function that is added to the FreeRTOS scheduler using the xCreateTask() API call.

A task will have the following:

```
1. A Priority level
```

- 2. Memory allocation
- 3. Singular input parameter (optional)
- 4. A name
- 5. A handler (optional): A data structure that can be used to reference the task later.

A FreeRTOS task declaration and definition looks like the following:

```
void vTaskCode( void * pvParameters )
{
    /* Grab Parameter */
    uint32_t c = (uint32_t)(pvParameters);
    /* Define Constants Here */
    const uint32_t COUNT_INCREMENT = 20;
    /* Define Local Variables */
    uint32_t counter = 0;
```

```
/* Initialization Code */
initTIMER();
/* Code Loop */
while(1)
{
    /* Insert Loop Code */
}
/* Only necessary if above loop has a condition */
xTaskDelete(NULL);}
```

Rules for an RTOS Task

- The highest priority ready tasks ALWAYS runs
 If two or more have equal priority, then they are time sliced
- Low priority tasks only get CPU allocation when:
 - All higher priority tasks are sleeping, blocked, or suspended
- Tasks can sleep in various ways, a few are the following:
 - Explicit "task sleep" using API call vTaskDelay();
 - Sleeping on a semaphore
 - Sleeping on an empty queue (reading)
 - Sleeping on a full queue (writing)

Adding a Task to the Scheduler and Starting the Scheduler

The following code example shows how to use **xTaskCreate()** and how to start the scheduler using **vTaskStartScheduler()**

```
tskIDLE_PRIORITY,/* Priority at which the task is created. */
    &xHandle ); /* Used to pass out the created task's handle. */
/* Start Scheduler */
vTaskStartScheduler();
return 0;}
```

Task Priorities



In the above situation, the high priority task task never sleeps, so it is always running. In this situation where the low priority task never gets CPU time, we consider that task to be **starved**.



In the above situation, the two tasks have the same priority, thus they share the CPU. The time each task is allowed to run depends on the **OS tick frequency**. The **OS Tick Frequency** is the frequency that the FreeRTOS scheduler is called in order to decide which task should run next. The **OS Tick is a hardware interrupt** that calls the RTOS scheduler. Such a call to the scheduler is called a **preemptive context switch**.

Context Switching

When the RTOS scheduler switches from one task to another task, this is called a Context Switch.

What needs to be stored for a Context switch to happen

In order for a task, or really any executable, to run, the following need to exist and be accessible and storable:

• Program Counter (PC)

- This holds the position for which line in your executable the CPU is currently executing.
- Adding to it moves you one more instruction.
- Changing it jumps you to another section of code.

• Stack Pointer (SP)

- This register holds the current position of the call stack, with regard to the currently executing program. The stack holds information such as local variables for functions, return addresses and [sometimes] function return values.
- General Purpose Registers
 - These registers are to do computation.
 - In ARM:
 - ° R0 R15
 - In MIPS:
 - \$v0, \$v1
 - ∘ \$a0 \$a3
 - \$t0 \$t7
 - \$s0 \$s7
 - **\$t8 \$t9**
 - Intel 8086
 - AX
 - ∘ BX
 - $\circ CX$
 - DX
 - ∘ SI
 - DI
 - ∘ BP

How does Preemptive Context Switch work?

- 1. A hardware timer interrupt or repetitive interrupt is required for this preemptive context switch.
 - 1. This is independent from an RTOS
 - 2. Typically 1ms or 10ms
- 2. The OS needs hardware capability to have a chance to STOP synchronous software flow and enter the OS "tick" interrupt.
 - 1. This is called the "Operating System Kernel Interrupt"
 - 2. We will refer to this as the OS Tick ISR (interrupt service routine)

- 3. Timer interrupt calls RTOS Scheduler
 - 1. RTOS will store the previous PC, SP, and registers for that task.
 - 2. Scheduler picks the highest priority task that is ready to run.
 - 3. Scheduler places that task's PC, SP, and registers into the CPU.
 - 4. Scheduler interrupt returns, and the newly chosen task runs as if it never stopped.

NOTE: Context switching takes time. The reason why most OS ticks are 1ms to 10ms, because any shorter means that there is less time for your code to run. If a context switch takes 100uS to do, then with a OS tick of 1ms, your code can only run for 900uS. If an OS Tick is only 150uS, then your code may only have enough time to run a few instructions. You spend more CPU time context switching then you do performing actual work.

Revision #1 Created 2 years ago by Khalil Estell Updated 2 years ago by Khalil Estell