

Introduction to Real Time Operating Systems

Lecture 17

Josh Brake
Harvey Mudd College

Learning Objectives

By the end of this lecture you will be able to:

- Articulate the importance of multitasking and how computers are not true multitasking systems.
- Describe the basic concepts of a real-time operating system
- Compare and contrast the tradeoffs of various scheduling algorithms.

Outline

- Motivation for multitasking and real-time operating systems
- Introduction to key concepts
 - Tasks
 - Scheduling
 - Semaphores
 - Queues
 - Interrupts and Events
- Introduction to FreeRTOS

Multitasking Scenarios

- Printing information to a display in response to keyboard input
- Car
 - Airbag response
 - Braking system
- Robotics
 - Flight system in a drone
 - Control or signal processing algorithms

Soft vs. hard real-time requirements

- Soft real-time requirements are those that state a time deadline—but breaching the deadline would not render the system useless.
- Hard real-time requirements are those that state a time deadline—and breaching the deadline would result in absolute failure of the system.

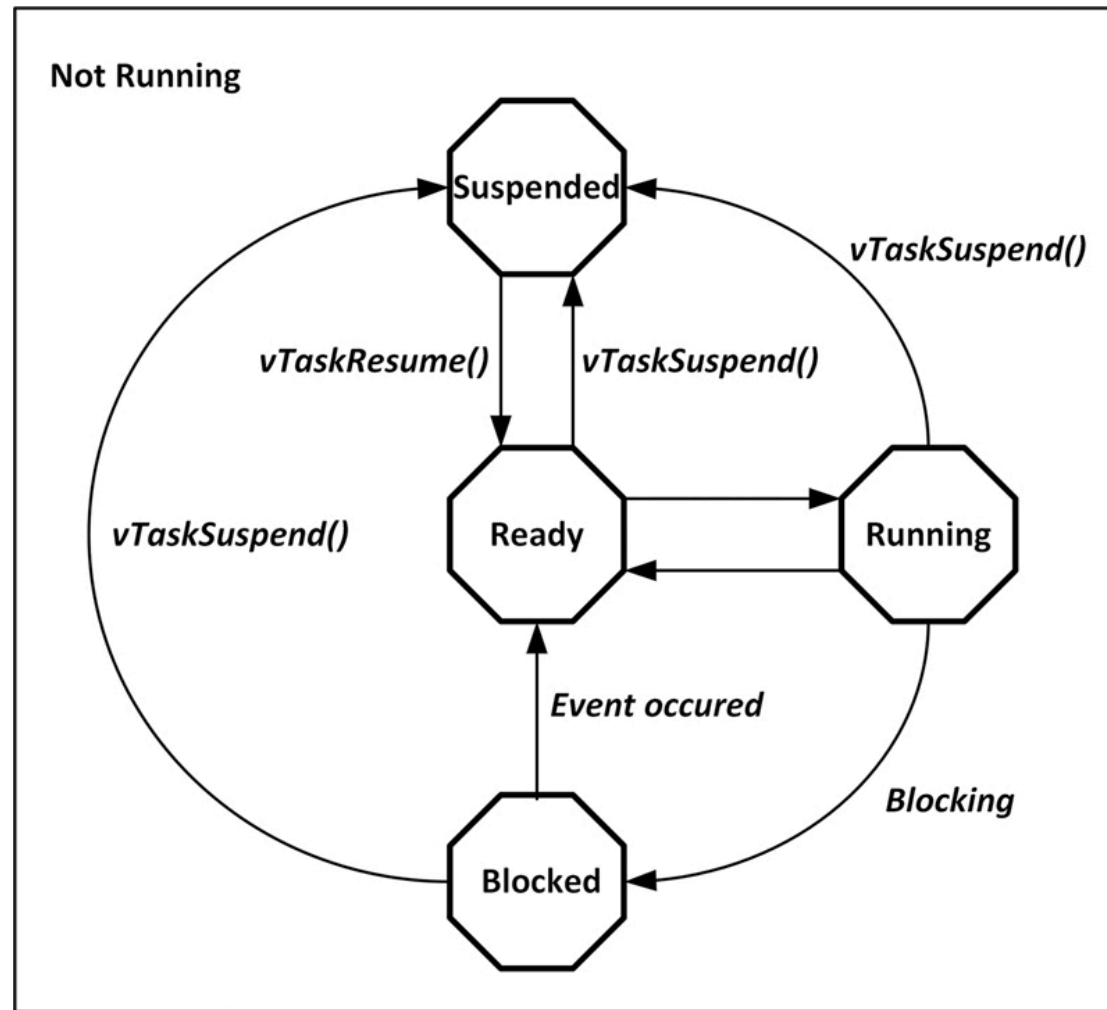
Why multitasking?

- Bare-metal programming
 - Setup and initialization – runs once
 - Infinite loop – runs continuously and handles main tasks
- Bare-metal + interrupts
 - Can now incorporate additional functionality which quickly responds to inputs and can guarantee that we don't miss important events.
 - Examples: Receiving and processing UART data, catching button inputs
- But as we build more and more complicated programs, it is hard to guarantee specific timing constraints are met.
 - Many different scenarios
 - Lots of edge cases which makes things very difficult to debug.

Why multitasking?

- Enter the concept of multitasking
 - Multitasking: means that several tasks (or programs) are processed in parallel on the same CPU
- Only have a single core on your microcontroller so this is not true parallelism, just swapping different tasks in and out
- Operations inside your bare-metal infinite loops are tasks.
 - For example, consider that you want to blink 2 LEDs at different frequencies. In bare-metal, you could have your infinite loop use timers to poll a timer and then toggle the LEDs based on the current time.
 - Works, but inefficient as the processor is always running.

Task States



Ibrahim, Dogan. *ARM-Based Microcontroller Multitasking Projects: Using the FreeRTOS Multitasking Kernel*.

The Mudd Multitasking Kernel

- You are taking E155, Clinic, and an HSA. In addition, you want to sleep 8 hours a night and have time to hang out with your friends playing board games over Zoom.
- Imagine you must manage your clinic mid-year report, MicroPs final project, and must read a book and write a paper for your HSA in addition to chatting with your friends. You only have one brain.
- What are different ways you can manage your tasks?

Scheduling Algorithms

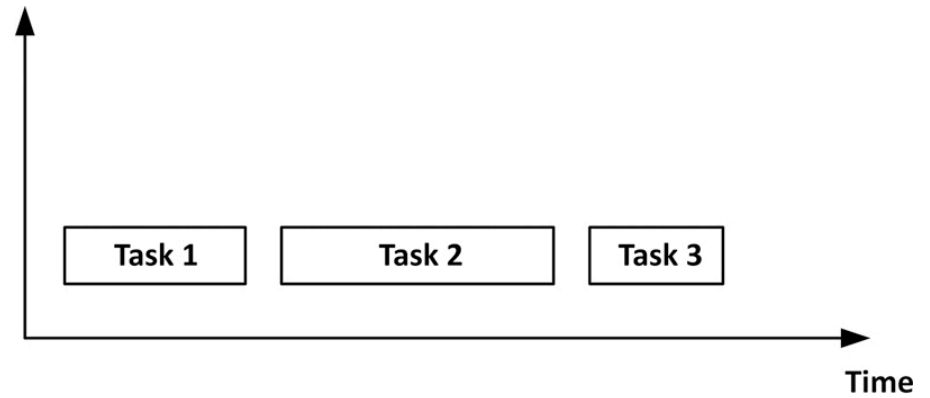
The scheduling algorithm decides which task is running on the core.

Three main algorithms:

- _____ scheduling
- _____ scheduling
- _____ scheduling

Co-operative scheduling

```
1 Task1() {  
2   // Task 1 code  
3 }  
4  
5 Task2() {  
6   // Task 2 code  
7 }  
8  
9 Task3() {  
10  // Task 3 code  
11 }  
12  
13 while(1) {  
14   Task1();  
15   Task2();  
16   Task3();  
17 }
```



Ibrahim, Dogan. ARM-Based Microcontroller Multitasking Projects: Using the FreeRTOS Multitasking Kernel.

Each task must yield control or else it can starve all other tasks.

Co-operative scheduling

Rules

1. Tasks must not block the overall execution, for example, by using delays or waiting for some resources and not releasing the CPU.
2. The execution time of each tasks should be acceptable to other tasks.
3. Tasks should exit as soon as they complete their processing.
4. Tasks do not have to run to completion and they can exit for example before waiting for a resource to be available.
5. Tasks should resume their operations from the point after they release the CPU.

Round-robin scheduling

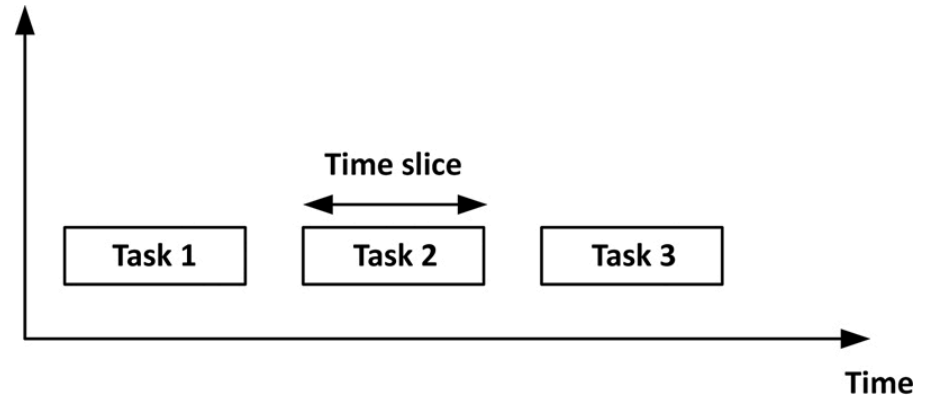
The scheduler creates a periodic time slice and equally divides CPU use between tasks.

Advantages:

- It is easy to implement.
- Every task gets an equal share of the CPU.
- Easy to compute the average response time.

Disadvantages:

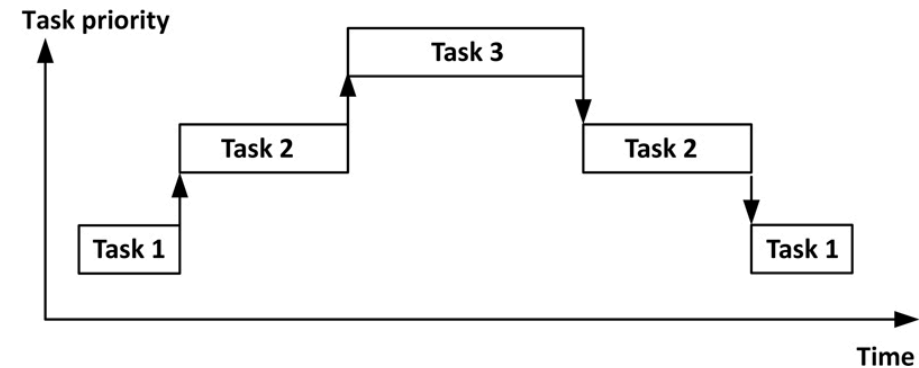
- It is not generally good to give the same CPU time to each task.
- Some important tasks may not run to completion.
- Not suitable for real-time systems where tasks usually have different processing requirements.



Ibrahim, Dogan. *ARM-Based Microcontroller Multitasking Projects: Using the FreeRTOS Multitasking Kernel*.

Preemptive Scheduling

- Most common scheduling algorithm in real-time systems
- Tasks are assigned priorities
- Higher priority tasks can preempt lower priority tasks to take the CPU
- Need to be careful to assign priorities appropriately or you can starve lower priority tasks



Ibrahim, Dogan. *ARM-Based Microcontroller Multitasking Projects: Using the FreeRTOS Multitasking Kernel*.

Q: How are tasks and their priorities different than interrupts?

Scheduling Algorithm Goals

- Be _____ such that each process gets a fair share of the CPU.
- Be _____ by keeping the CPU busy.
- The scheduler should require minimal _____.
- Maximize _____ by minimizing the time users must wait.
- Be _____ so that same tasks take the same time when run multiple times.
- Minimize _____.
- Maximize _____.
- Enforce _____.
- Avoid _____.

Other scheduling algorithms

- First-come first-served
- Shortest time remaining first
- Longest time remaining first
- Multilevel queue scheduling
- Dynamic priority scheduling

FreeRTOS



Introduction to FreeRTOS

- We will use FreeRTOS as our example
 - Other popular RTOSes include Zephyr, NuttX, VxWorks. Varying licensing agreements.
 - Like a programming language: once you learn one RTOS, concepts transfer to others.
- FreeRTOS licensed under MIT license – very permissive.
 - Can be used in commercial applications and users retain all ownership of their IP.

Code Structure of FreeRTOS

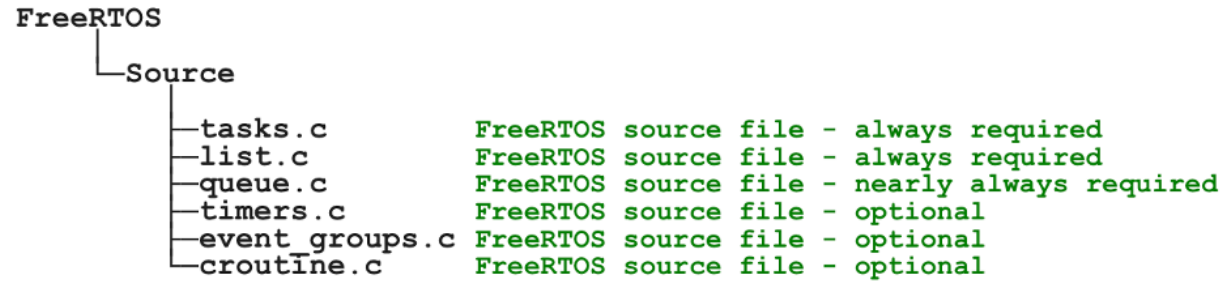


Figure 2. Core FreeRTOS source files within the FreeRTOS directory tree

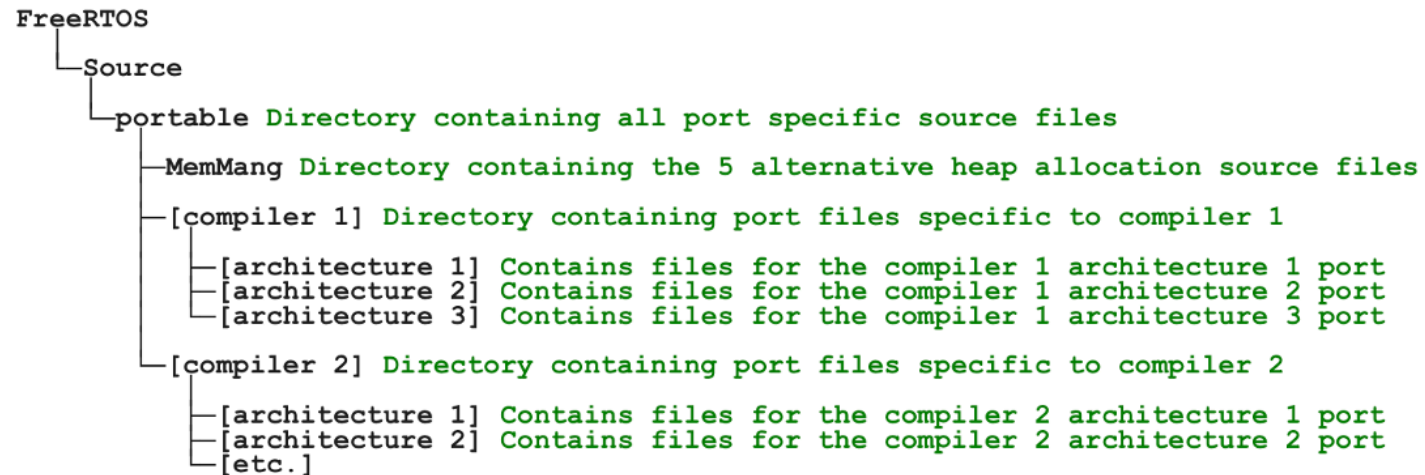


Figure 3. Port specific source files within the FreeRTOS directory tree

Data Types and Naming Conventions

- Port specific data types _____ - holds tick count value
- _____ – Base type value which is most efficient data type on a given system. Typically the word size.

Naming Conventions: Variable Names

Variable Names - prefixes tell their type

Prefix	Type
<code>c</code>	<code>char</code>
<code>s</code>	<code>int16_t</code> (short)
<code>i</code>	<code>int32_t</code> (long)
<code>x</code>	<code>BaseType_t</code> and other non-standard types (structs, task handles, queue handles, etc.)

Naming Conventions: Function Names

Function	Description
<code>vTaskPrioritySet()</code>	returns a void and is defined within <code>task.c</code> .
<code>xQueueReceive()</code>	returns a variable of type <code> BaseType_t </code> and is defined within <code>queue.c</code> .
<code>pvTimerGetTimerID()</code>	returns a pointer to void and is defined within <code>timers.c</code> .

Template Project

```
int main( void )
{
    /* Perform any hardware setup necessary. */
    prvSetupHardware();

    /* --- APPLICATION TASKS CAN BE CREATED HERE --- */

    /* Start the created tasks running. */
    vTaskStartScheduler();

    /* Execution will only reach here if there was insufficient heap to
    start the scheduler. */
    for( ;; );
    return 0;
}
```

Listing 1. The template for a new main() function

Barry, Richard. *Mastering the FreeRTOS Real Time Kernel: A Hands-On Tutorial Guide*. 2016.

Creating Tasks

```
BaseType_t xTaskCreate( TaskFunction_t pvTaskCode,  
                        const char * const pcName,  
                        uint16_t usStackDepth,  
                        void *pvParameters,  
                        UBaseType_t uxPriority,  
                        TaskHandle_t *pxCreatedTask );
```

Listing 13. The xTaskCreate() API function prototype

Return

- `pdPass` or `pdFail` - indicates if task was successfully created.

Parameters

- `pvTaskCode` - pointer to C function that implements that task
- `pcName` - Descriptive name for the task
- `usStackDepth` - size of stack to be allocated by the kernel when creating the stack (in words)
- `pvParameters` - pointer to void to pass in parameters. Need to cast void pointer to correct type inside the function to use it.
- `uxPriority` - Defines the priority of the task
- `pxCreatedTask` - handle to created task

Printing to Terminal

```
void vTask1( void *pvParameters )
{
const char *pcTaskName = "Task 1 is running\r\n";
volatile uint32_t ul; /* volatile to ensure ul is not optimized away. */

    /* As per most tasks, this task is implemented in an infinite loop. */
    for( ;; )
    {
        /* Print out the name of this task. */
        vPrintString( pcTaskName );

        /* Delay for a period. */
        for( ul = 0; ul < mainDELAY_LOOP_COUNT; ul++ )
        {
            /* This loop is just a very crude delay implementation. There is
            nothing to do in here. Later examples will replace this crude
            loop with a proper delay/sleep function. */
        }
    }
}
```

```
void vTask2( void *pvParameters )
{
const char *pcTaskName = "Task 2 is running\r\n";
volatile uint32_t ul; /* volatile to ensure ul is not optimized away. */

    /* As per most tasks, this task is implemented in an infinite loop. */
    for( ;; )
    {
        /* Print out the name of this task. */
        vPrintString( pcTaskName );

        /* Delay for a period. */
        for( ul = 0; ul < mainDELAY_LOOP_COUNT; ul++ )
        {
            /* This loop is just a very crude delay implementation. There is
            nothing to do in here. Later examples will replace this crude
            loop with a proper delay/sleep function. */
        }
    }
}
```

```
1 int main(void) {
2     xTaskCreate(vTask1, "Task 1", 1000, NULL, 1, NULL);
3     xTaskCreate(vTask2, "Task 2", 1000, NULL, 1, NULL);
4     vTaskStartScheduler();
5
6     for( ;; );
7 }
```

Barry, Richard. *Mastering the FreeRTOS Real Time Kernel: A Hands-On Tutorial Guide*. 2016.

Printing to Terminal Example Output

```
1 int main(void) {  
2     xTaskCreate(vTask1, "Task 1", 1000, NULL, 1, NULL);  
3     xTaskCreate(vTask2, "Task 2", 1000, NULL, 1, NULL);  
4     vTaskStartScheduler();  
5  
6     for( ;; );  
7 }
```

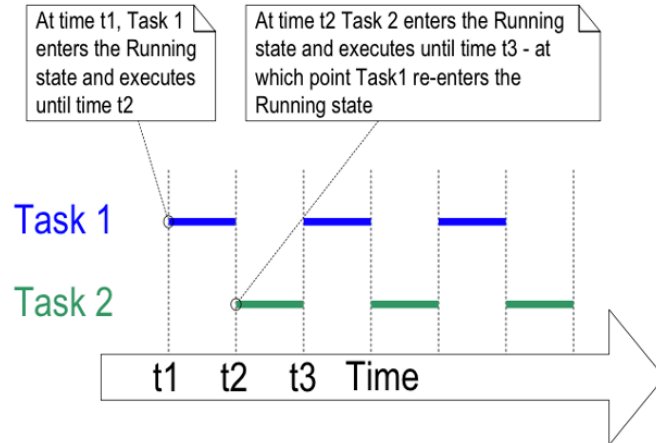


Figure 11. The actual execution pattern of the two Example 1 tasks

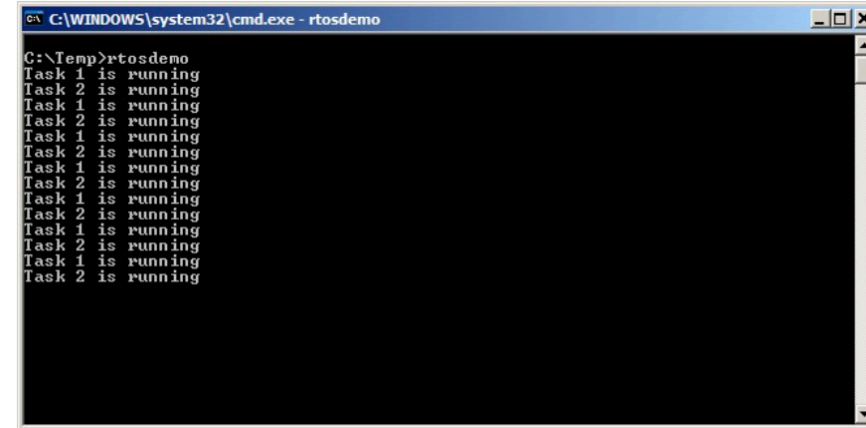


Figure 10. The output produced when Example 1 is executed¹

Example: Priorities

```
/* Define the strings that will be passed in as the task parameters. These are
defined const and not on the stack to ensure they remain valid when the tasks are
executing. */
static const char *pcTextForTask1 = "Task 1 is running\r\n";
static const char *pcTextForTask2 = "Task 2 is running\r\n";

int main( void )
{
    /* Create the first task at priority 1. The priority is the second to last
parameter. */
    xTaskCreate( vTaskFunction, "Task 1", 1000, (void*)pcTextForTask1, 1, NULL );

    /* Create the second task at priority 2, which is higher than a priority of 1.
The priority is the second to last parameter. */
    xTaskCreate( vTaskFunction, "Task 2", 1000, (void*)pcTextForTask2, 2, NULL );

    /* Start the scheduler so the tasks start executing. */
    vTaskStartScheduler();

    /* Will not reach here. */
    return 0;
}
```

Listing 21. Creating two tasks at different priorities

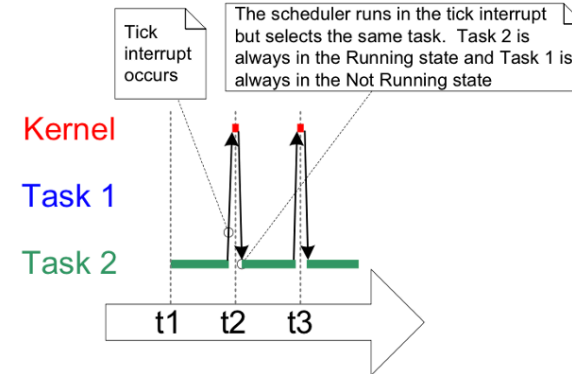


Figure 14. The execution pattern when one task has a higher priority than the other

Revisiting Not Running State

Three Options

- Suspended
- Ready
- Blocked

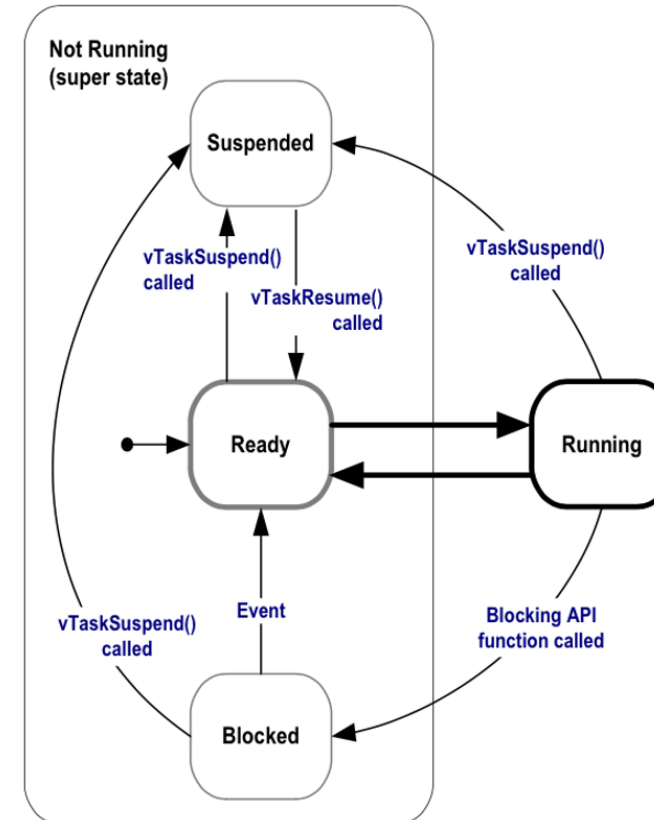


Figure 15. Full task state machine

Barry, Richard. *Mastering the FreeRTOS Real Time Kernel: A Hands-On Tutorial Guide*. 2016.

Example: Printing with better delay using blocked state

```
void vTaskFunction( void *pvParameters )
{
    char *pcTaskName;
    const TickType_t xDelay250ms = pdMS_TO_TICKS( 250 );

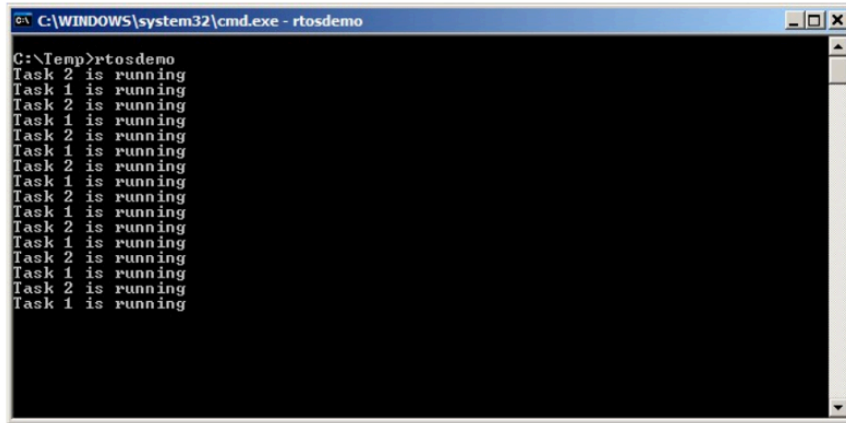
    /* The string to print out is passed in via the parameter.  Cast this to a
    character pointer. */
    pcTaskName = ( char * ) pvParameters;

    /* As per most tasks, this task is implemented in an infinite loop. */
    for( ;; )
    {
        /* Print out the name of this task. */
        vPrintString( pcTaskName );

        /* Delay for a period.  This time a call to vTaskDelay() is used which places
        the task into the Blocked state until the delay period has expired.  The
        parameter takes a time specified in 'ticks', and the pdMS_TO_TICKS() macro
        is used (where the xDelay250ms constant is declared) to convert 250
        milliseconds into an equivalent time in ticks. */
        vTaskDelay( xDelay250ms );
    }
}
```

Listing 23. The source code for the example task after the null loop delay has been replaced by a call to vTaskDelay()

Example: Sending tasks to blocked state



```
C:\WINDOWS\system32\cmd.exe - rtsdemo
C:\Temp>rtdemo
Task 2 is running
Task 1 is running
Task 2 is running
Task 1 is running
Task 2 is running
Task 1 is running
Task 2 is running
Task 1 is running
Task 2 is running
Task 1 is running
Task 2 is running
Task 1 is running
Task 2 is running
Task 1 is running
Task 2 is running
Task 1 is running
Task 2 is running
Task 1 is running
Task 2 is running
Task 1 is running
```

Figure 16. The output produced when Example 4 is executed

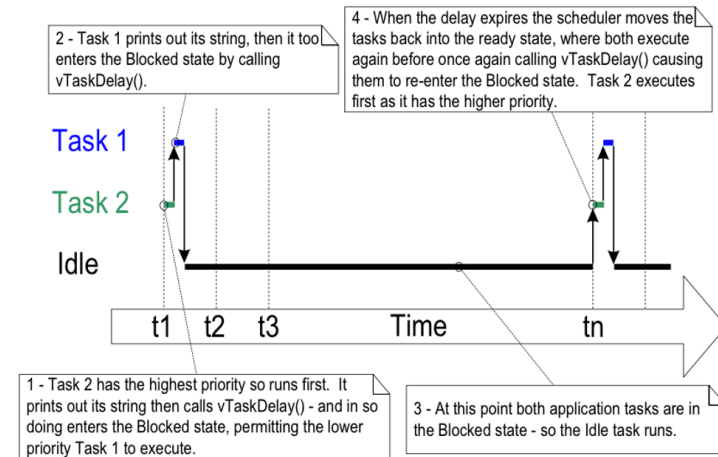


Figure 17. The execution sequence when the tasks use `vTaskDelay()` in place of the NULL loop

Key Terms in Real-time Systems

- **Tasks** – C functions which indicate things to do. Implemented as infinite loops.
- **Scheduling** – The process of determining what task is currently running
- **Semaphores** – an abstract data type which controls access to a resource used by multiple tasks
- **Queues** – A way to communicate between tasks (Chapter 4 of Mastering FreeRTOS)
- **Interrupts and Events** – how to safely integrate interrupts with the RTOS kernel (Chapter 6 of Mastering FreeRTOS)

Summary

- Multitasking is an important concept in advanced embedded systems
 - Have timing constraints that must be met (both soft and hard deadlines)
 - Hard to debug and manage systems with increasing complexity while guaranteeing all deadlines are met.
- Real-time operating systems introduce a scheduler which enables the programmer to efficiently use - CPU cycles while ensuring deadlines are met.
- FreeRTOS is an open and accessible platform to learn RTOS concepts like tasks, queues, semaphores, and resource management.