



Introduction to FreeRTOS

**Kizito NKURIKIYEYU,
Ph.D.**

Readings

Read Chap 4 of Barry R. (2018) "Mastering the FreeRTOS™ Real Time Kernel A Hands-On Tutorial Guide"¹

Topics

- Task managements
- Task states
- Task priorities
- The Idle Task and the Idle Task Hook
- Scheduling Algorithms



¹This reference textbook is freely available at https://freertos.org/Documentation/161204_Mastering_the_FreeRTOS_Real_Time_Kernel-A_Hands-On_Tutorial_Guide.pdf

What is FreeRTOS?

- FreeRTOS is an open source real-time operating system (RTOS) for embedded systems.
- FreeRTOS supports many different architectures and compiler toolchains, and is designed to be small, simple, and easy to use.
- It provides many features: neat and readable source code, it is portable, scalable, provides preemptive and co-operative scheduling, has multitasking and interrupt management

What is FreeRTOS?

- An open source, real-time operating system for small, low-power edge devices easy to program, deploy, secure, connect, and manage.
- Includes libraries for connecting to Amazon Web Services (AWS) web services and other edge devices running AWS

What is FreeRTOS?

- FreeRTOS is the most popular RTOS for small embedded MCU²
- Originally developed by Richard Barry around 2003
- Since 2017 it is developed and maintained by Amazon Web Services.
- FreeRTOS is a real-time kernel targeting at hard real-time applications.
- Primarily written in C with few assembler functions
- FreeRTOS has a very small footprint^{3,4}
- It supports many compilers (CodeWarrior, GCC, IAR, etc.) as well as many processor architectures (ARM7, various PICs, 8051, x86, etc.).

²The most recent [Embedded Markets Study](#) shows that Embedded Linux and FreeRTOS continue to outpace other operating systems used in embedded development.

³The FreeRTOS scheduler requires 236 bytes of RAM and about 5 to 10 KBytes of ROM. Task creation requires an additional 64 bytes of stack. [FreeRTOS](#)




FIG 1. Most popular non-linux OS for IoT development^{5, 6}

⁵<https://iot.eclipse.org/community/resources/iot-surveys/assets/iot-developer-survey-2019.pdf>

⁶<https://www.andplus.com/blog/which-operating-system-should-you-use-for-your-iot-solution>

FreeRTOS features

- **Real-time**—actually, hard-real-time
- **Preemptive or cooperative operation**—The scheduler interrupts at regular frequency to increment the tick count
- **Dynamic scheduling**—Scheduler decision made at regular clock frequency.
- **Priority-based scheduling**—when more than one task exists at the highest priority, tasks are executed in round robin fashion.
- **Inter-Process Communication**—Tasks within FreeRTOS can communicate with each other through the use of queuing and synchronization mechanisms
- **Blocking and Deadlock avoidance**: Tasks are either non-blocking or will block with a fixed period of time.
- **Critical Section Protection**—handled by the disabling of interrupts
- **Memory Allocation**—FreeRTOS offers several heap

Some coding conventions⁸

- **pdTRUE**—boolean that represents a true value or 1 in C
- **pdFALSE**—boolean that represents a false value or 0 in C
- **pdPASS**—boolean that represents a successful operation or 1 in C
- **pdFAIL**—boolean that represents a failed operation or 0 in C
- **Datatype and function naming conventions**
 - **l**—denotes “long”
 - **u**—denotes “unsigned”
 - **s**—denotes “short”
 - **p**—denotes “pointers”
 - **v**—denotes “void”
 - **x**—variables of non stdint types are prefixed x
 - **e**—enumerated variables are prefixed e
- API functions are prefixed with their return type
- Macros are pre-fixed with the file in which they are defined. The pre-fix is lower case. e.g, **configUSE_PREEMPTION** is

TAB 1. FreeRTOS datatype mapping

| FreeRTOS type definition | Corresponding C/C++ data type |
|--------------------------|--|
| portCHAR | char |
| portSHORT | short |
| portLONG | long |
| portFLOAT | float |
| portDOUBLE | double |
| portSTACK_TYPE | uint32_t |
| portBASE_TYPE | Typically, this is a 32-bit type on a 32-bit architecture, a 16-bit type on a 16-bit architecture, and an 8-bit type on an 8-bit architecture. portBASE_TYPE is generally used for return types that can take only a very limited range of values, and for pdTRUE/pdFALSE type Booleans. |
| UBaseType_t | unsigned long |
| TickType_t | TickType_t is the data type used to hold the tick count value, and to specify times |

FreeRTOS Architecture overview

FreeRTOS's code breaks down into three main areas: tasks, communication, and hardware interfacing.

- `tasks.c`⁹ and `task.h`¹⁰ allow creating, scheduling, and maintaining tasks.
- about 40% of FreeRTOS's core code deals with communication. `queue.c`¹¹ and `queue.h`¹² handle FreeRTOS **communication**. Tasks and interrupts use queues to send data to each other and to signal the use of critical resources using **semaphores and mutexes**.
- The **Hardware Whisperer**—Much code in FreeRTOS kernel is hardware-independent. About 6% of FreeRTOS's core code acts a shim between the hardware-independent FreeRTOS **core and the hardware-dependent code** as shown in Fig. 2.

⁹<https://github.com/FreeRTOS/FreeRTOS-Kernel/blob/main/tasks.c>

¹⁰<https://github.com/FreeRTOS/FreeRTOS-Kernel/blob/main/include/task.h>

¹¹<https://github.com/FreeRTOS/FreeRTOS-Kernel/blob/main/queue.c>




FIG 2. Summary of the FREERTOS porting layer

Task Management

Tasks in FreeRTOS

- With FreeRTOS, application can be structured as a set of autonomous tasks
- Each task executes within its own context (e.g., stack) with no coincidental dependency on other tasks
- The scheduler starts, stops, swaps in, and swaps out tasks as needed
- Each task has a user-assigned priority
 - low priority —`taskIDLE_PRIORITY` or idle task priority. It's equal to zero
 - `configMAX_PRIORITIES-1` is the highest priority¹³
- It uses a "ready list" to keep track of all tasks that are currently ready to run.

```
1 static xList pxReadyTasksLists[  
    configMAX_PRIORITIES];
```

- `pxReadyTasksLists[0]` is a list of all ready priority 0 tasks
- `pxReadyTasksLists[1]` is a list of all ready priority 1 tasks and

The System Tick

- **System tick**—user-configurable heartbeat of a FreeRTOS system
- FreeRTOS configures the system to generate a periodic tick interrupt.
- User configurable but typically in the millisecond range.
- On each tick interrupt, the `vTaskSwitchContext()` function is called and selects the highest-priority ready task and puts it in the `pxCurrentTCB` variable

```
1 //Find the highest-priority queue that contains
  ready tasks.
2 while (listLIST_IS_EMPTY (& (pxReadyTasksLists[
  uxTopReadyPriority]))) {
3   configASSERT ( uxTopReadyPriority );
4   --uxTopReadyPriority;
5 }
6 //Make sure that tasks of the same priority get
```

Task Control Block

- The FreeRTOS kernel manages tasks via the **Task Control Block**
- The TCB exists for each task and contains all information (Table 2) necessary to completely describe the state of a task.

TABLE 2. Properties of FreeRTOS task control block

| | Explanation |
|------------------|--|
| Top of the stack | pointer to the last item placed on the stack for this task |
| Task state | list item that puts the TCB in the ready or blocked queues |
| Event list | list item used to place the TCB in the event lists |
| Priority | Task priority |
| Stack start | Pointer to the start of the process stack |
| TCB number | a filed used to help during debugging and tracing |
| Task name | the name of the task |
| Stack depth | the total depth of the stack in variables |

TABLE 3. Detailed contents of a FreeRTOS task control bloc

| Field | Purpose | Optional |
|-------------------|---------------------------------|----------|
| pcTaskName [] | Human-readable task name | - |
| uxTCBNumber | Task identification number | * |
| uxPriority | Active priority | - |
| uxBasePriority | Baseline priority | * |
| pxStack | Lowest task stack address | - |
| pxEndOfStack | Highest task stack address | * |
| pxTopOfStack | Current top of the task stack | - |
| xGenericListItem | Link to the scheduler's lists | - |
| xEventListItem | Link to an event wait list | - |
| uxCriticalNesting | Interrupt disable nesting level | * |
| ulRunTimeCounter | CPU time consumed | * |
| pxTaskTag | User-defined, per-task pointer | * |
| xMPUSettings | Memory protection information | * |

```
1 typedef struct tskTaskControlBlock{
2   /*The location of the last item placed on the
  stack */
3   volatile StackType_t * pxTopOfStack;
4   /*The list that the state (Ready, Blocked,
  Suspended). */
5   ListItem_t xStateListItem;
6   /*Used to reference a task from an event list. */
7   ListItem_t xEventListItem;
8   /*The priority of the task. 0 is the lowest
  priority. */
9   UBaseType_t uxPriority;
10  /*Points to the start of the stack. */
11  StackType_t * pxStack;
12  /*Descriptive name given to the task to
  facilitates debugging*/
13  char pcTaskName[ configMAX_TASK_NAME_LEN ];
14 } tskTCB;
```

Tasks and task states

- A task is functions in RTOS
- A task can have the following states:




FIG 3. Common RTOS task states

- **Running**—the task is using microprocessor to execute instructions
- **Ready**—the task has instructions for microprocessor to execute, but is not yet executing, perhaps because another higher priority task is running
- **Blocked**—the task has nothing for




FIG 4. Task state diagram in the FREERTOS operating system.¹⁴

¹⁴Adapted from Bertolotti, I. C., & Hu, T. (2017). Embedded Software Development.

Tasks and task states

- An application can consist of many tasks
- **Only one** task of the application can be executed at any given time on the microcontroller (single core)
- Thus, a task can exist in one of two states: **Running** or **Not Running**
- Only the scheduler can decide which task should enter the Running state
- A task is said to have been **switched in** or **swapped in** when transitioned from the Not Running to the Running state
- A task is said to have been **switched out** or **swapped out** when transitioned from the running state to the not-running state

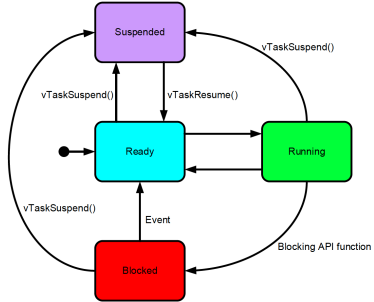


FIG 5. FreeRTOS tasks and task states

Tasks and task states

- New tasks are placed immediately in the Ready state
- The Ready list is arranged in order of priority with tasks of equal priority being serviced on a round-robin basis.
- The implementation of FreeRTOS actually uses multiple Ready lists – one at each priority level (Fig. 6)




FIG 6. Illustration of FreeRTOS Ready Task List Data Structure

Tasks Functions

- Tasks are implemented as C functions
- The prototype of a task function must return void and take a void pointer parameter
- A task will typically execute indefinitely in an infinite loop: must never terminate by attempting to return to its caller
- If required, a task can delete itself prior to reaching the function end
- A single task function definition can be used to create any number of tasks
 - Each created task is a separate execution instance
 - Each created task has its own stack
 - Each created task has its own copy of any automatic variables defined within the task itself

Tasks Functions

- Tasks are defined as a functions that return void and takes a void pointer as its only parameter.
- The parameter can be used to pass information of any type into the task
- Task functions should never return so are typically implemented as a continuous loop.
- Attempting to do so will result in a `configASSERT()` being called if it is defined.
- Call `vTaskDelete(NULL)` to ensure its exit is clean.

```
1 void vTaskFunction(void *pvParameters) {
2     while (true)
3     {
4         -- Task application code here. --
5     }
6     //In normal conditions, this code should never execute
```

Remarks

In FreeRTOS, all tasks share the same address space among them and with the operating system itself.

- This approach keeps the operating system as simple as possible, and makes it easy and efficient to share variables, because all tasks automatically and implicitly share memory.
- However, it's error prone since it becomes harder to protect tasks from each other with respect to illegal memory accesses¹⁵
- FreeRTOS can offer memory protection on some processor architectures, through its Memory Protection Unit, if the processor allows it.
- Practically, the operating system is implemented as a library of object modules. The application and the operating system modules are therefore bundled together in the resulting executable image.

¹⁵In many cases, the lack of memory protection is not an issue since many

Create a Task

```
1 BaseType_t xTaskCreate(  
2 /* Function that implements the task. */  
3 TaskFunction_t pvTaskCode,  
4 /* Text name for the task. */  
5 const char * const pcName,  
6 /* Stack size in words, not bytes. */  
7 configSTACK_DEPTH_TYPE usStackDepth,  
8 /* Parameter passed into the task. */  
9 void *pvParameters,  
10 /* Priority at which the task is created. */  
11 UBaseType_t uxPriority,  
12 /* Used to pass out the created task's handle. */  
13 TaskHandle_t *pxCreatedTask  
14 );
```

LISTING 3: Excerpt of the FreeRTOS task control block

TAB 4. Parameters of the FreeRTOS's xTaskCreate function

| | |
|---------------|--|
| pvTaskCode | Pointer to the task entry function (just the name of the function that implements the task). |
| pcName | A descriptive name for the task. This is mainly used to facilitate debugging, but can also be used to obtain a task handle. The maximum length of a task's name is defined by configMAX_TASK_NAME_LEN in FreeRTOSConfig.h. |
| usStackDepth | The number of words (not bytes!) to allocate for use as the task's stack. For example, if the stack is 16-bits wide and usStackDepth is 100, then 200 bytes will be allocated for use as the task's stack. |
| pvParameters | A value that is passed as the parameter to the created task. If pvParameters is set to the address of a variable then the variable must still exist when the created task executes - so it is not valid to pass the address of a stack variable. |
| uxPriority | The priority at which the created task will execute. Priorities are asserted to be less than configMAX_PRIORITIES. If configASSERT is undefined, priorities are silently capped at (configMAX_PRIORITIES - 1). |
| pxCreatedTask | Used to pass a handle to the created task out of the xTaskCreate() function. pxCreatedTask is optional and can be set to NULL. |

Return values

- **pdTRUE**—when the task was created successfully
- **errCOULD_NOT_ALLOCATE_REQUIRED_MEMORY**—the task could not be created because there was insufficient heap memory

Not Running State

- **Blocked state**—A task that is waiting for an event
- Tasks can enter the blocked state to wait for two different types of event
 - Time related events where a delay period expiring or an absolute time being reached. For example, when the **vTaskDelay()**¹⁶ function is called
 - Synchronization events where the events originate from another task or interrupt. For example, wait for data to arrive on a queue
- It is possible for a task to block on a synchronization event with a timeout. For example, wait for a maximum of 10 ms for data to arrive on a queue
- The Suspended state: tasks in the Suspended state are not available to the scheduler.
 - A task can be suspended through a call to the **vTaskSuspend()**¹⁷ API function
 - A suspended task can be resumed via a call to the

Not Running State

- **Ready State**—tasks that are able to execute (they are not in the Blocked or Suspended state) but are not currently executing because a different task of equal or higher priority is already in the Running state. Tasks are able to run, and therefore “ready” to run, but not currently in the Running state
- When a task is first created and made ready to run, the kernel puts it into the ready state. In this state, the task actively competes with all other ready tasks for the processor's execution time.
- Ready tasks can only move to the running state. Because many tasks might be in the ready state, the kernel's scheduler uses the priority of each task to determine which task to move to the running state.
- Tasks in the ready state cannot move directly to the blocked state. They first need to run so they can make a blocking call
- Because many tasks might be in the ready state, the kernel's

Time & delay

- Functions like the Arduino `delay()`²⁰ generate delay using a null loop and poll a counter until it reaches a fixed value. This waste processor cycles

```
1 void delay(unsigned long ms) {
2   uint16_t start = (uint16_t)micros();
3   while (ms > 0) {
4     yield();
5     if (((uint16_t)micros() - start) >= 1000) {
6       ms--;
7       start += 1000;
8     }
9   }
10 }
```

- Such functions are called **blocking delays** because they block code execution until they returns—which means nothing else can run while they are running.

Time & delay

- FreeRTOS's `vTaskDelay()`²¹ API function is non-blocking
 - `vTaskDelay()` places the calling task into the Blocked state for a fixed number of tick interrupts
 - The task in the Blocked state will not use any processing time at all

```
1 void vTaskDelay( const TickType_t
2   xTicksToDelay );
```

- `xTicksToDelay` is the number of tick interrupts that the calling task should remain in the Blocked state before being transitioned back into the Ready state.
- The constant `portTICK_RATE_MS` stores the time in milliseconds of a tick, which can be used to convert milliseconds into ticks.

```
1 // delay for 500ms.
2 vTaskDelay(500/portTICK_RATE_MS);
```

- An alternative method is to use the `vTaskDelayUntil()`²² API




FIG 7. The execution sequence when the tasks use `vTaskDelay()`

Time & delay

- The length of a tick depends on the operating system configuration and, to some extent, on hardware capabilities.
- The configuration macro `configTICK_RATE_HZ` represents the configured tick frequency in Hertz and may be useful to convert back and forth between the usual time measurement units and clock ticks.
- For instance, the quantity `1000000/configTICK_RATE_HZ` is the tick length, approximated as an integer number of microsecond.
- The function `TickType_t xTaskGetTickCount(void)` returns the current time as the number of ticks elapsed since the operating system scheduler was started.

TAB 5. Time-Related Primitives of FREERTOS

| Function | Purpose | Optional |
|--------------------------------|-------------------------------|----------|
| <code>xTaskGetTickCount</code> | Return current time, in ticks | - |

vTaskDelayUntil() API²³

```
1 void vTaskDelayUntil( TickType_t *pxPreviousWakeTime  
  , const TickType_t xTimeIncrement );
```

- **pxPreviousWakeTime**—Pointer to a variable that holds the time at which the task was last unblocked. The variable must be initialised with the current time prior to its first use (see the example below). Following this the variable is automatically updated within vTaskDelayUntil().
- **xTimeIncrement**—The cycle time period. The task will be unblocked at time (***pxPreviousWakeTime + xTimeIncrement**). Calling vTaskDelayUntil with the same xTimeIncrement parameter value will cause the task to execute with a fixed interval period.

²³INCLUDE_vTaskDelayUntil must be defined as 1 for this function to be available. See the RTOS Configuration documentation for more information.

vTaskDelayUntil() API

```
1 void vTaskBlinkLed(void* pvParameters) {  
2     // set pin PB0 of PORTB for output  
3     DDRB |= (1<<PB0);  
4     TickType_t xLastWakeTime = xTaskGetTickCount();  
5     while (true) {  
6         // Toggle the 1st bit on PORT B (i.e. PB0)  
7         PORTB ^= (1<<PB0);  
8         vTaskDelayUntil(&xLastWakeTime, (1000/  
9             portTICK_PERIOD_MS));  
10    }  
11 }
```

LISTING 4: Example of using the vTaskDelayUntil() and blink an LED every 500ms

Remarks

- The tick counter data type is either a 16- or 32-bit unsigned integer
- Therefore, it is important to consider that the FREERTOS time counter will overflow²⁴ and restart counting from zero²⁵.
- The vTaskDelay() function already handle time counter wraparound automatically, and hence, no special care is needed to use them.
- The vTaskDelay() function implements a relative time delay to the time of the call, which is implicitly taken as a time reference point. The amount of delay from that reference point is fixed and is equal to xTicksToDelay.
- The function vTaskDelayUntil() implements an absolute time delay and the delay between wake-up events is kept constant and equal to xTimeIncrement, regardless of what the task did in between (see Fig. 8 for details).

²⁴https://en.wikipedia.org/wiki/Integer_overflow




FIG 8. Relative versus absolute delays

Task Priority

- Higher priorities task run before lower priority task
- Tasks with the same priority share CPU time (in time slices)
- The scheduler runs at the end of each time slice to select the next task to run
- The length of the time slice is set by the SysTick interrupt frequency by configuring the `configTICK_RATE_HZ` while setting the compile configuration in `FreeRTOSConfig.h`
- The initial priority of a task is assigned when created by using `xTaskCreate()`
- The task priority can be queried by using the `xTaskPriorityGet()`
- The task priority can be changed by using `vTaskPrioritySet()`
- Low numeric values denote low priority tasks (0 is the lowest priority)




FIG 9. Comparison between relative and absolute time delays, as implemented by `vTaskDelay` and `vTaskDelayUntil`.—An absolute delay is better when a task has to carry out an operation periodically because it guarantees that the period will stay constant even if the response time of the task—the grey rectangles in Figure—varies from one instance to another

Idle Task and Idle Task Hook Functions

- The idle task is executed when no application tasks are in running state²⁶
- The idle task is automatically created by the scheduler when `vTaskStartScheduler()` is called
- The idle task has the **lowest possible priority** (priority 0) to ensure it never prevents a higher priority application task
- Application specific functionality can be directly added into the idle task via an idle hook (or call-back) function
 - An idle hook function is automatically called per iteration of the idle task loop
 - Can be utilized to execute low priority, background or continuous processing

```

1 void vApplicationIdleHook () {}
2 void vApplicationTickHook () {}
3 void vApplicationAssertHook () {}
    
```

TAB 6. Summary of the task-related primitives of FreeRTOS²⁷

| Function | Purpose | Optional |
|-------------------------------------|---------------------------------------|----------|
| <code>vTaskStartScheduler</code> | Start the scheduler | - |
| <code>vTaskEndScheduler</code> | Stop the scheduler | - |
| <code>xTaskCreate</code> | Create a new task | * |
| <code>vTaskDelete</code> | Delete a task given its handle | * |
| <code>uxTaskPriorityGet</code> | Get the priority of a task | * |
| <code>vTaskPrioritySet</code> | Set the priority of a task | * |
| <code>vTaskSuspend</code> | Suspend a specific task | * |
| <code>vTaskResume</code> | Resume a specific task | * |
| <code>xTaskResumeFromISR</code> | Resume a specific task from an ISR | * |
| <code>xTaskIsTaskSuspended</code> | Check whether a task is suspended | * |
| <code>vTaskSuspendAll</code> | Suspend all tasks but the running one | - |
| <code>xTaskResumeAll</code> | Resume all tasks | - |
| <code>uxTaskGetNumberOfTasks</code> | Return current number of tasks | - |

²⁷The elements marked with * are optional; they may or may not be present, depending on the FreeRTOS configuration

Remarks

- When `main()` executes, the FreeRTOS's scheduler is not yet active. It can be explicitly started and stopped by means of the following function calls

```
1 void vTaskStartScheduler (void);  
2 void vTaskEndScheduler (void);
```

- The function `vTaskStartScheduler()` reports errors back to the caller in an unusual way:
 - When it is able to start the scheduler successfully, `vTaskStartScheduler` does not return to the caller at all
 - Otherwise, `vTaskStartScheduler()` may return to the caller for two distinct reasons:
 - 1 The scheduler was successfully started, but one of the tasks then stopped it by invoking `vTaskEndScheduler()`.
 - 2 The scheduler was not started at all because an error occurred.
- After creation, it is possible to change the priority of a task or retrieve it by means of the functions

```
UBaseType_t priority);  
2 UBaseType_t uxTaskPriorityGet (TaskHandle_t xTask  
);
```

The end