

# Saving Power with an RTOS

## INTRODUCTION

With increasing customer demand for green, energy efficient products, and the amazing growth in mobile devices where long battery life is highly desirable, engineers now have to consider how best to reduce the power consumed by their designs. This can be partly achieved by selecting components that consume less power, but the software architecture also contribute to power saving.

This article discusses approaches to saving power by using features commonly found in embedded Real Time Operating Systems.

## IDLE TASK SLEEP MODE

The first, easy, step to save power is to allow the RTOS to place the processor into a low power mode when it enters the Idle Task. To remain responsive, the processor will wake up on the next interrupt. If no external interrupt is triggered, this will be the next RTOS Tick interrupt.

For example, with OPENRTOS® or SAFERTOS® this simple

## RTOS SCHEDULING EFFICIENCY

It is a commonly held belief that adding an RTOS to an embedded design will create additional processing overhead, resulting in greater power usage. For simple designs, where a polling 'super loop' architecture may be more appropriate, this could be true. However, for more advanced, complex designs, deploying an RTOS using an event based scheduling algorithm typically reduces the amount of processing time required to run your application, allowing the selection of a smaller processor or adding functionality.

In a software project containing a priority based, pre-emptive RTOS, Tasks are executed in order of priority, i.e the highest priority Task ready to run is always allocated processor time. When there are no Tasks currently available, for example, they are all blocked waiting for events, the RTOS executes a background Task called the Idle Task. The Idle Task may implement some background processing, but typically the system will remain in the Idle Task waiting for an event to bring a Task back to life. The more effectively the software design exploits the RTOS, the greater the time spent in the Idle Task waiting.

power saving feature is implemented by placing the processor into a low power mode from within the Idle Task hook function, as shown in Figure 1.

The power saving that can be achieved by this simple method is limited by the necessity to periodically exit and then re-enter the low power mode to process RTOS Tick interrupts.

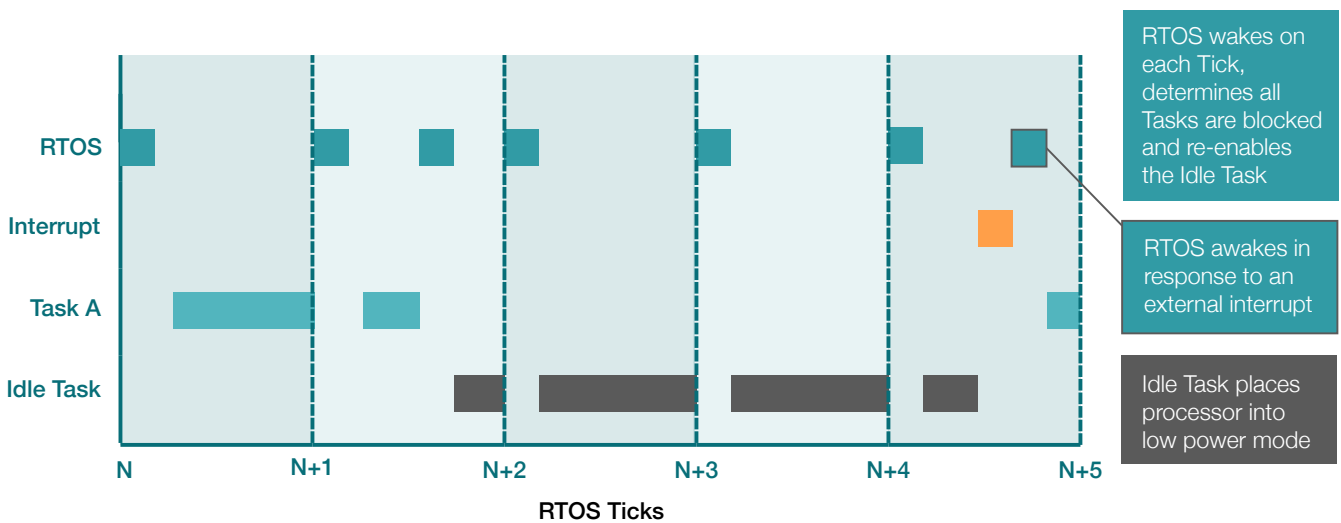


Figure 1. The interaction of the RTOS and Idle Task during Sleep Mode

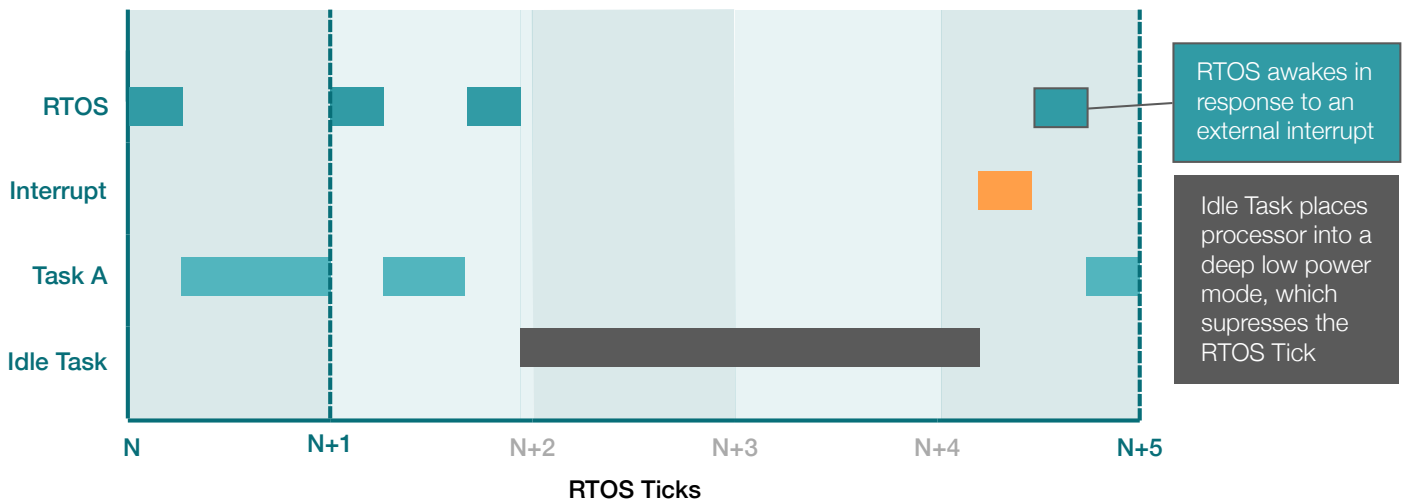


Figure 2. The interaction of the RTOS and Idle Task during Ultra-Low Power Mode

### SAFERTOS ULTRA-LOW POWER MODE

When supported by the processor, **SAFERTOS** can be purchased supporting a tickless Ultra-low Power Mode. This tickless sleep mode extends the power save strategy previously presented further. Here power is saved by completely stopping the RTOS tick interrupt during idle periods and placing the system in a power saving mode. Stopping the tick interrupt allows the processor to remain in a deep power saving mode longer until either an interrupt occurs, or it is time for the RTOS to re-activate a Task, as shown in Figure 2.

### WHEN USING SAFERTOS, TICKLESS SLEEP MODE CAN BE ENTERED WHEN:

1. The Idle Task is the only task able to run because all the application tasks are either in the Blocked state or in the Suspended state.
2. At least 'N' further complete tick periods will pass before the kernel is due to transition an application task out of the Blocked state, where 'N' is a user-defined value. This is done to avoid having to calculate and reprogram the clock on each tick.

In Tickless Sleep Mode the processor will be put to sleep until the next Task Block Time expires, or the next Software Timer is triggered, whichever is sooner. If the following conditions are true, **SAFERTOS** will place the processor in deep sleep mode at least until the next external interrupt.

1. Software timers are not being used, or are not due to expire, hence **SAFERTOS** is not due to execute a Timer Callback function at any time in the future.
2. All the application Tasks are either in the Suspended state, or in the Blocked state with an infinite timeout (a timeout value of portMAX\_DELAY), hence **SAFERTOS** is not due to transition a task out of the Blocked state at any fixed time in the future.

**SAFERTOS** knows the longest time it can possibly sleep without the risk of missing a Task Block Time expiring, or a Software Timer executing. However, it cannot predict when it will "actually" exit the low power sleep mode because it can't predict asynchronous interrupts being accepted. Therefore, when **SAFERTOS** exits a low power sleep, it cannot make any assumptions about how long it was asleep for and so calculates the actual sleep time before adjusting the system time accordingly.

WITTENSTEIN high integrity systems

Worldwide Sales and Support  
 Americas: +1 408 625 4712  
 ROTW: +44 1275 395 600  
 Email: sales@highintegritysystems.com  
 Web: www.highintegritysystems.com



WITTENSTEIN