

Designing Responsive and Real-Time Systems

Chapter 10

Renesas Electronics America Inc.
Embedded Systems using the RX63N

Rev. 1.0

Learning Objectives

- Most Embedded Systems have multiple real time tasks running at the same time.
- “Which task should it perform first?” determines how responsive the Processor is.
- Response of a Processor determines many important matrices while choosing the right processor for our application.
- In this chapter we will see how to design Responsive and Real time Embedded Systems along with some task scheduling approaches.

Scheduling

- Topic—How do we make the processor do things at the right times?
- There are various methods, the best fit depends on:
 - System requirements: response time
 - Software complexity: number of threads of execution
 - Resources: RAM, interrupts, energy available

Scheduling

- How do we schedule the tasks on the CPU?
 - An infinite loop in main.
 - Real-time operating system.
 - Is there anything else available?

Scheduling

■ A Simple Scheduling Example

- Consider an embedded system which controls a doorbell in a house. When someone presses a switch, the doorbell should ring. The system's **responsiveness** describes how long it takes from pressing the switch to sounding the bell.

– Simple code to realize the above example:

```
1. void main (void){
2.     init_system();
3.     while(1){
4.         if(switch == PRESSED){
5.             Ring_The_Bell();
6.         }
7.     }
8. }
```

Scheduling

- The doorbell in the previous example is very responsive. What if we decide to throw in burglar detector and a smoke detector?
 - Our while loop should look like this:

```
3. while(1){
4.     if(switch == PRESSED){
5.         Ring_The_Doorbell();
6.     }
7.     if(Burglar_Detected() == TRUE){
8.         Sound_The_Burglar_Alarm();
9.     }
10.    if(Smoke_Detected() == TRUE){
11.        Sound_The_Fire_Alarm();
12.    }
13. }
```

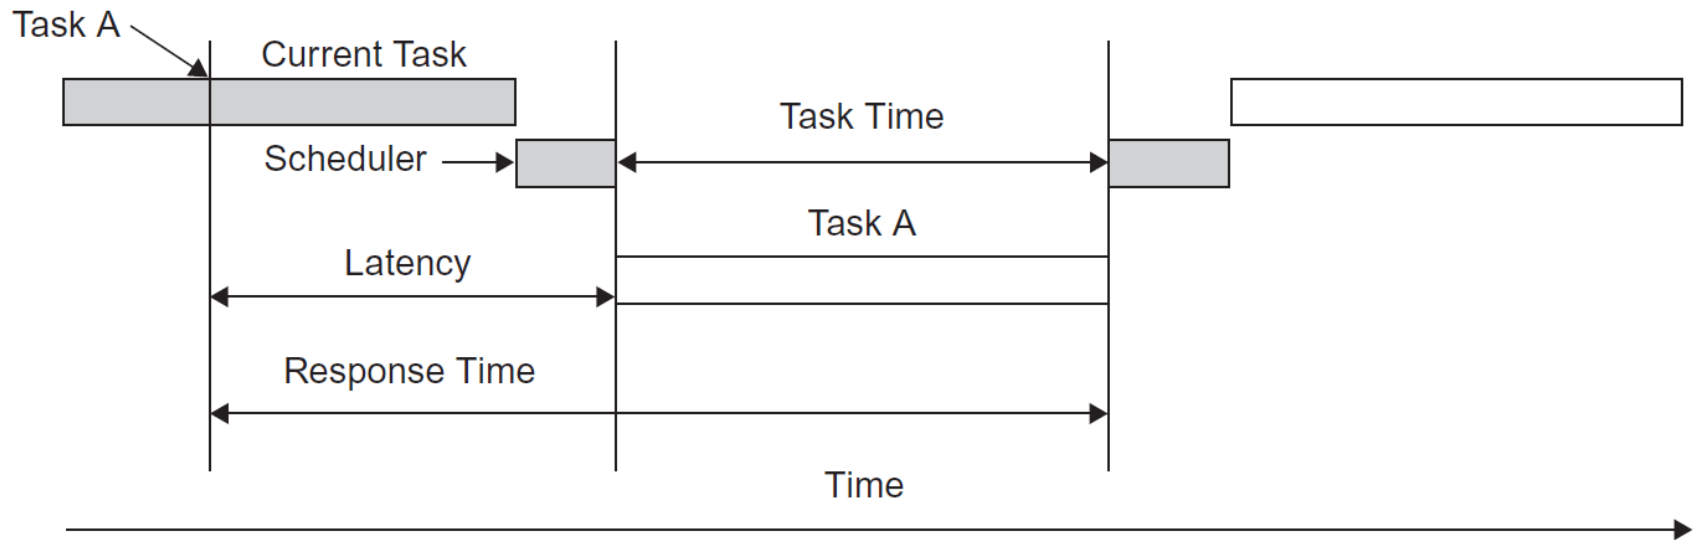
Scheduling

- How should we share the processor's time between these tasks?
- It depends on various factors:
 - Delay between smoke detection and alarm going off?
 - What if smoke is detected while someone is at the door and someone is trying to break thorough the window? What gets precedence?
- Now that we have to share the processor, we have to worry about how long the bell rings and the alarms sound. If doorbell rings for 30 seconds, it will run for 30 seconds. During this time, we won't know if our house is burning or being robbed.

Scheduling

- This example reveals the two fundamental issues in scheduling for responsive system:
 - If we have multiple tasks ready to run, which one do we run first? This decision defines the **ordering** of task execution.
 - Do we allow one task to interrupt or **preempt** another task (or even itself)?
- Both of these decisions will determine the system's **responsiveness** (measured by response time).

Task Ordering



- $T_{\text{release}}(i)$ = Time at which task i becomes ready to run.
- $T_{\text{response}}(i)$ = Delay between request for service and completion of service for task i .
- $T_{\text{task}}(i)$ = Time needed to perform computations for task i .
- $T_{\text{ISR}}(i)$ = Time needed to perform interrupt service routine i .

Going Dynamic

- We can change the order based on current conditions (e.g., if the house is on fire) using a **dynamic** schedule.
- We will try to take a step forward and hard code the priorities in our earlier home alarm system example.

– Code:

```
3. while(1){
4.     if(Smoke_Detected() == TRUE){
5.         Sound_The_Fire_Alarm();
6.     }
7.     else if (Burglar_Detected() == TRUE) Sound_The_Burglar_Alarm();
8.     }
9.     else if (switch == PRESSED) { Ring_The_Doorbell();
10.    }
11. }
```

Going Dynamic

- The else clauses will change the schedule to a dynamic one.
 - As long as smoke is detected, the burglar alarm and doorbell will be ignored. Similarly, burglar detection will disable the doorbell.
- This **strict prioritization** may or may not be appropriate for a given system. We may want to ensure some **fairness**, perhaps by limiting how often a task can run.

Task Preemption

- Can one of the tasks interrupt a lower priority task? Or do all tasks run to completion?
- What if a burglar breaks a window a second after the doorbell switch is pressed? We wait 30 seconds till we know there is an intruder?
- No. We can just allow the smoke and burglar detection code to preempt Ring_The_Doorbell. We will need to use a more sophisticated task scheduler which can:
 - (1) preempt and resume tasks.
 - (2) detect events which trigger switching and starting tasks.

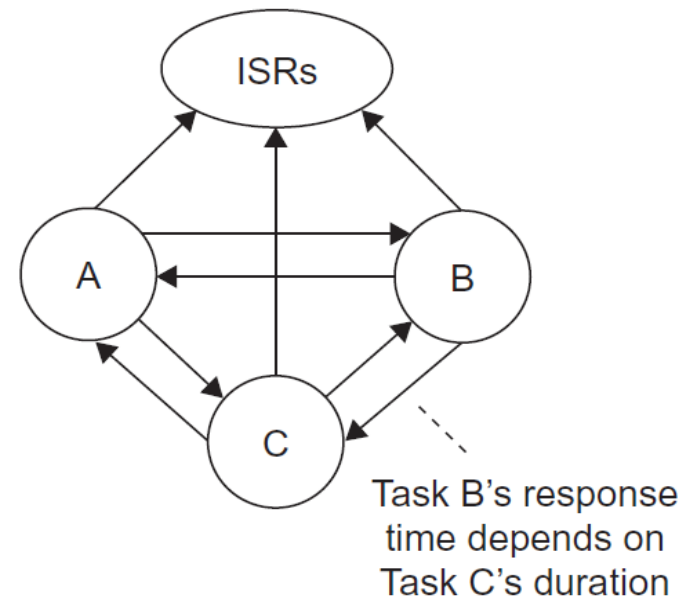
Is it Fair?

- Prioritizing some tasks over others can lead to starvation of lower priority tasks? They might never get a chance to run.
- This scenario is called as starvation.
- Some ways to feed these lower priority tasks some processor time and keep them happy.
 - We can allow multiple tasks to share the same priority level. If both tasks are ready to run, we alternate between executing each of them (whether by allowing each task to run to completion, or by preempting each periodically).
 - We can limit how often each task can run by defining the task frequency. This is the common approach used for designers of real-time systems. Note that we can still allow only one task per priority level.

Response Time and Preemption

- With the non-preemptive static scheduler, each task's response time depends on the duration of all other tasks and ISRs, so there are nine dependences.

Non-Preemptive
Static Scheduling



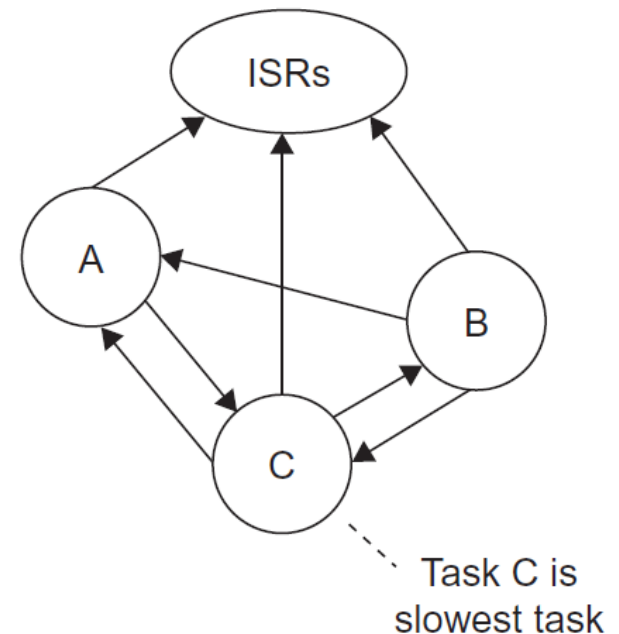
9 dependences

- Higher priority tasks and ISRs
- Lower priority tasks

Response Time and Preemption

- With the non-preemptive dynamic scheduler, we assign priorities to tasks (A B C).
- Any task no longer depends on lower priority tasks, so we have more timing independence and isolation.
- If C starts running, it means a delay for other tasks, A and B each which results in a total of eight dependences.

Non-Preemptive
Dynamic Scheduling




8 dependences

- Higher priority tasks and ISRs
- **Slowest** task

Response Time and Preemption

- With the preemptive dynamic scheduler, we also prioritize the tasks (A B C). Because a task can preempt any lower priority task.
- As a result we have only six dependences.
- This means that in order to determine the response time for a task, we only need to consider **higher priority tasks**.

Preemptive Dynamic Scheduling



6 dependences

- Only higher priority tasks and ISRs

Stack Memory Requirements

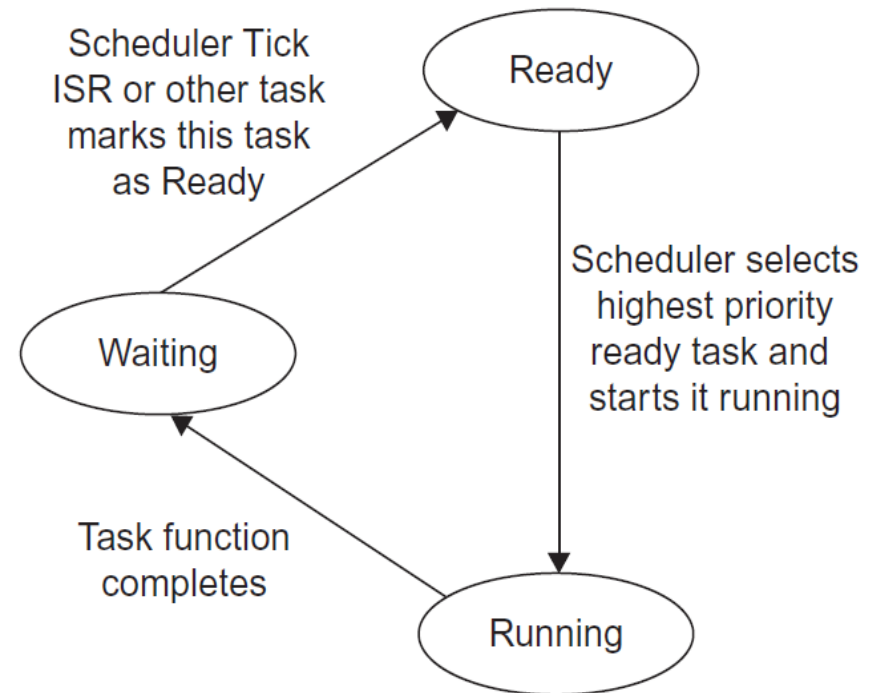
- The non-preemptive approach requires only **one call stack**, while a preemptive approach typically requires **one call stack per task**.
- The function call stack holds a function's state information, such as return address and limited lifetime variables.
- Without task preemption, task execution does not overlap in time, so all tasks can share the same stack.

Task Management

- A task will be in one of several possible states. The scheduler and the task code itself both affect which state is active.
- With a **dynamic scheduler**, a task can be in any one of the Ready, Waiting, or Running states.
- Task management consists of:
 - Changing states
 - Context switching
 - Sharing Data between tasks

Task States

- **Waiting** for the scheduler to decide that this task is ready to run.
- **Ready to start running** but not running yet. There may be a higher-priority task which is running.
- **Running** on the processor.



Non-preemptive Dynamic Scheduler

Transitions between States

■ The transition from **ready to running**:

- In a non-preemptive system, when the scheduler is ready to run a task, it selects the highest priority ready task and moves it to the running state.
- In a preemptive system, when the kernel is ready to run a task, it selects the highest priority ready task and moves it to the running state by restoring its context to the processor.

■ The transition from **running to waiting**:

- In a non-preemptive system, the only way a task can move from running to waiting is if it **completes**.
- In a preemptive system, the task can yield the processor and request a delay.

Transitions between States

■ The transition from **waiting to ready**:

- In a non-preemptive system the timer tick ISR moves the task by setting the run flag. Alternatively, another task can set the run flag to request for this task to run.
- In a preemptive system, the kernel is notified that some event has occurred so it moves that particular task from the waiting state to the ready state.

■ The transition from **running to ready**:

- In a non-preemptive system this transition does not exist, as a task cannot be preempted.
- In a preemptive system, when the kernel determines a higher priority task is ready to run, it will save the context of the currently running task, and move that task to the ready state.

Context Switching for Preemptive Systems

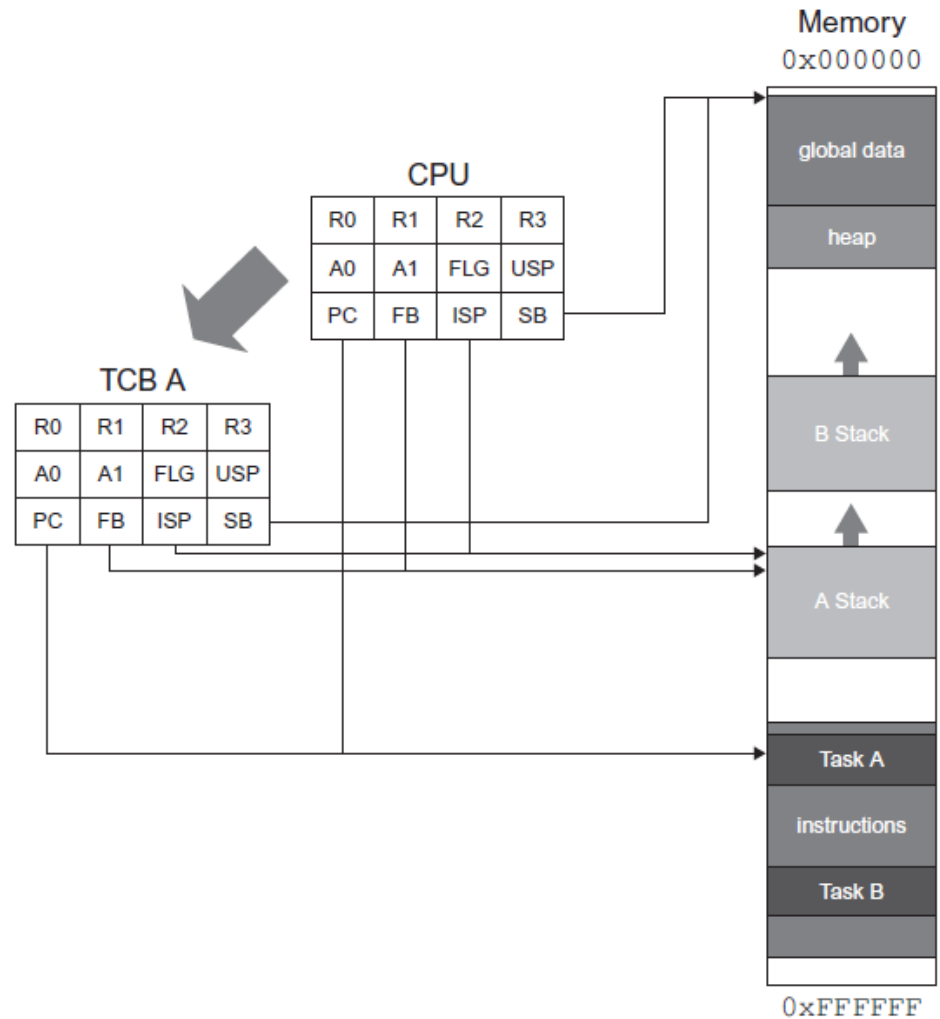
- In preemptive systems, some of these state transitions require the scheduler to save a task's execution context and restore another task's context to ensure programs execute correctly.
- This is called Context Switch.

Context Switching

- In order to perform a context switch from one task to another correctly, we must first copy all of this task-specific processor register information to a storage location (TCB of first task).
- Next, we must copy all of the data from TCB of second task into the CPU's registers.
- Now the CPU will be able to resume execution of the Second task where it left off.

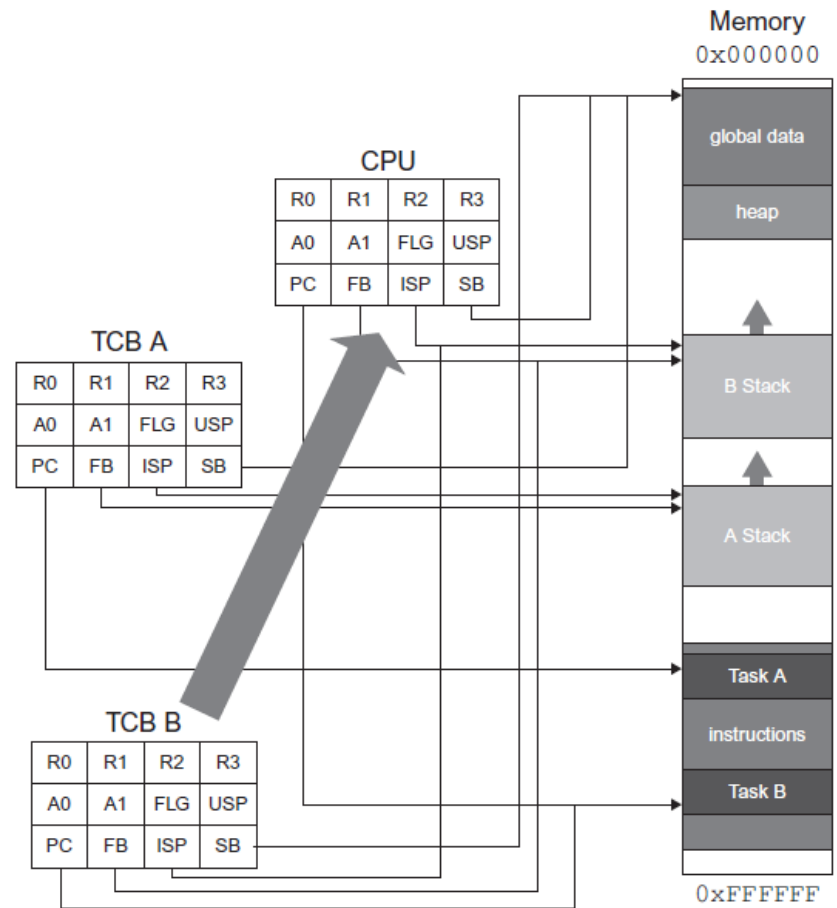
Context Switching

- Saving Task A's context:



Context Switching

- Restoring Task B's context:



Sharing Data

- Preemption among tasks introduces a vulnerability to data race conditions.
- The system can fail in new ways when:
 - Multiple tasks or ISRs share data, or
 - Multiple instances of a function can execute concurrently
- In order to prevent these failures we need to be careful when considering shared data.

Data Shared Objects

- If a data object is accessed by code which can be interrupted (is not *atomic*), then there is a risk of data corruption.
- *Atomic* code is the smallest part of a program that executes without interruption.

Data Shared Objects

– Code:

```
1. unsigned time_minutes, time_seconds;
2. void task1 (void){
3.     time_seconds++;
4.     if(time_seconds >= 60){
5.         time_minutes++;
6.         time_seconds = 0;
7.     }
8. }

9. void task2 (void){
10.     unsigned elapsed_sec;
11.     elapsed_seconds = time_minutes * 60 + time_seconds;
12. }
```

Data Shared Objects

- In the code example our shared object is a pair of variables which measure the current time in minutes and seconds.
- If task1 is preempted between lines 4 and 5 or lines 5 and 6, then when task2 runs it will only have a partially updated version of the current time, and elapsed_seconds will be incorrect.
- If task2 is preempted during line 11, then it is possible that time_minutes is read **before** task1 updates it and time_seconds is read **after** task 1 updates it. Again, this leads to a corrupted elapsed_seconds value.

Function Reentrancy

- Another type of shared data problem comes with the use of non-reentrant functions.

– Code:

```
1. void task1 ( ){
2     . . . . .
3     swap(&x, &y);
4     . . . . .
5. }
6. void task2 ( ){
7     . . . . .
8     swap(&p, &q);
9     . . . . .
10. }
11. int Temp;
12. void swap (*i, *j ){
13.     Temp = *j;
14.     *j = *i;
15.     *i = Temp;
16. }
```

Function Reentrancy

- Such functions are called **non-reentrant**. The code which can have multiple simultaneous, interleaved, or nested invocations which will not interfere with each other is called reentrant code.
- In a multi-processing environment, the non-reentrant functions should be eliminated.
- A function can be checked for its reentrancy based on these three rules:
 - A reentrant function **may not use variables in a non-atomic way** unless they are stored on the stack of the calling task or are the private variables of that task.
 - A reentrant function **may not call other functions** which are not reentrant.
 - A reentrant function **may not use the hardware in a non-atomic way**.

Shared-Data Solutions and Protection

- ***Disable Interrupts***

- Disable the interrupts during the critical section of the task.

- ***Use a Lock***

- ***RTOS-Provided Semaphore***

- ***RTOS-Provided Messages***

- ***Disable Task Switching***

Disabling Interrupts

- One of the easiest methods is to disable the interrupts during the critical section of the task.
- Once the critical section of the code is executed, the interrupt masking can be restored to its previous state.

```
1. #define TRUE 1
2. #define FALSE 0
3. static int error;
4. static int error_count;
5. void error_counter ( ) {
6.     if(error == TRUE) {
7.         SAVE_INT_STATE;
8.         DISABLE_INTS;
9.         error_count++;
10.        error = FALSE;
11.        RESTORE_INT_STATE;
12.    }
13. }
```

Use a Lock

- Associate every shared variable with a lock variable, which is also declared globally.
- If a function uses the shared variable, then it sets the lock variable; and once it has finished process, it resets the lock variable.
- Every function must test the lock variable before accessing it.

```
1. unsigned int var;  
2. char lock_var;  
3. void task_var ( ){  
4.     unsigned int sum;  
5.     if(lock_var == 0){  
6.         lock_var = 1;  
7.         var = var + sum;  
8.         lock_var = 0;  
9.     }  
10.    else { //message to scheduler to check var  
12.    }  
14. }
```

RTOS-Provided Semaphore

- Most operating systems provide locks to shared variables through the use of *semaphores*.
- A semaphore is a mechanism that uses most of the multitasking kernels to protect shared data and to synchronize two tasks.

Example of Semaphore

```
1. typedef int semaphore;
2. float temp;
3. semaphore var_temp
4. void Task1 (void){
5.     wait (var_temp);
6.     temp = (9/5)(temp + 32); /* Celsius into Fahrenheit */
7.     Signal (var_temp);
8. }
9. void Task2 (void){
10.    wait(var_temp);
11.    temp = ADDR0; /* Read ADC value from thermistor */
12.    temp = ADCtotemp_conversion();
13.    Signal (var_temp);
14. }
```

- Just before the tasks enter the critical section, they request the semaphore and only then perform the operation on the shared variable.

Nonpreemptive Dynamic Scheduler

- The scheduler has three fundamental parts.
 - **Task Table:** This table holds information on each task, including:
 - The address of the task's root function.
 - The period with which the task should run (e.g., 10 ticks).
 - The time delay until the next time the task should run (measured in ticks).
 - A flag indicating whether the task is ready to run.
 - **Tick ISR:** Once per time tick (say each 1 millisecond) a hardware timer triggers an interrupt.
 - **Task Dispatcher:** It is simply an infinite loop which examines each task's run flag. If it finds a task with the **run** flag set to 1, the scheduler will clear the **run** flag back to 0, execute the task, and then go back to examining the **run** flags.

Nonpreemptive Dynamic Scheduler

	Priority		Length		Frequency																						
Task 1		2		1		20																					
Task 2		1		2		10																					
Task 3		3		1		5																					
Task 4		0		1		3																					
Elapsed time	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	
Task executed				T4		T3	T4			T4	T2	T4	T3		T4	T3		T4			T2	T4	T1	T4	T3		
Task T1	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	20	19	18	17	16	15	
Task T2	10	9	8	7	6	5	4	3	2	1	10	9	8	7	6	5	4	3	2	1	10	9	8	7	6	5	
Task T3	5	4	3	2	1	5	4	3	2	1	5	4	3	2	1	5	4	3	2	1	5	4	3	2	1	5	
Task T4	3	2	1	3	2	1	3	2	1	3	2	1	3	2	1	3	2	1	3	2	1	3	2	1	3	2	
Run T1																						W	W	W	R		
Run T2											R											R					
Run T3						R					W	W	W	R		W	R					W	W	W	W	W	R
Run T4				R			R			R			R			R			R				W	R		R	

Example Application Using RTC Scheduler

- Let's use the RTC scheduler to create a toy with red and green LEDs flashing at various frequencies. The grnLED task toggles a green LED (on board LED 6) every one second, and the redLED task toggles a red LED (on board LED 12) every 0.25 seconds. The grn_redLED task toggles one red LED and one green LED (on board LED 7 and LED 8 respectively) every 0.5 seconds.

Scheduler Code for Adding a New Task

```
1. int Add_Task(void (*task)(void), int time, int priority){
2.     //Check for valid priority
3.     if(priority >= MAX_TASKS || priority < 0)
4.         return 0;
5.     //Check to see if we are overwriting an already scheduled
task
6.     if(GBL_task_table [priority].task != NULL)
7.         return 0;
8.     // Schedule the task
9.     GBL_task_table[priority].task = task;
10.    GBL_task_table[priority].run = 0;
11.    GBL_task_table[priority].timer = time;
12.    GBL_task_table[priority].enabled = 1;
13.    GBL_task_table[priority].initialTimerValue = time;
14.    return 1;
15. }
```


Scheduler Code for Task Timer Initialization

```
1. void init_Task_Timers(void){
2.     int i;
3.     // Initialize all tasks
4.     for(i = 0; i < MAX_TASKS; i++){
5.         GBL_task_table[i].initialTimerValue = 0;
6.         GBL_task_table[i].run = 0;
7.         GBL_task_table[i].timer = 0;
8.         GBL_task_table[i].enabled = 0;
9.         GBL_task_table[i].task = NULL;
10.    }
11. }
```

Scheduler Code Initialization

```
1. void Init_RTC_Scheduler(void){
2.     IEN(TMR0,CMIA0) = 0;
3.     IPR(TMR0,CMIA0) = 3;
4.     MSTP(TMR0) = 0;
5.     TMR0.TCNT = 0x00;
6.     TMR0.TCORA = 78;
7.     TMR0.TCSR.BYTE = 0xE2;
8.     IEN(TMR0,CMIA0) = 1;
9.     TMR0.TCCR.BYTE = 0x0C;
10.    TMR0.TCR.BYTE = 0x48;
11. }
```

- A timer is set up to generate an interrupt at regular intervals. Within the interrupt service routine the timer value for each task is decremented.
- When the timer value reaches zero, the task becomes ready to run.

Schedular code for Running Tasks

```
1. void Run_RTC_Scheduler(void){
2.     int i;
3.     // Loop forever
4.     while(1){
5.         // Check each task
6.         for(i = 0; i < MAX_TASKS; i++){
7.             // check if valid task
8.             if(GBL_task_table[i].task != NULL){
9.                 // check if enabled
10.                if(GBL_task_table[i].enabled == 1){
11.                    // check if ready to run
12.                    if(GBL_task_table[i].run == 1){
13.                        // Reset the run flag
14.                        GBL_task_table[i].run = 0;
15.                        // Run the task
16.                        GBL_task_table[i].task();
17.                        // break out of loop to start at entry 0
18.                        break;
19.                    }
20.                }
21.            }
22.        }
23.    }
24. }
```

Code – Tasks and Main

```
1. void grnLED(void){
2.     if(LED3 == LED_ON)
3.         LED3 = LED_OFF;
4.     else
5.         LED3 = LED_ON;
6. }

7. void redLED(void){
8.     if(LED4 == LED_ON)
9.         LED4 = LED_OFF;
10.    else
11.        LED4 = LED_ON;
12.}

13.void grn_redLED(void){
14.    if(LED5 == LED_ON)
15.        LED5 = LED9 = LED_OFF;
16.    else
17.        LED5 = LED9 = LED_ON;
18.}
```

Code – Tasks and Main

```
1. void main(void){  
2.     ENABLE_LEDS;  
3.     init_Task_Timers();  
4.     Add_Task(grnLED,10000,0);  
5.     Add_Task(redLED,2500,1);  
6.     Add_Task(grn_redLED,5000,2)  
7.     Init_RTC_Scheduler();  
8.     Run_RTC_Scheduler();  
9. }
```

In This Chapter We Learned

- Scheduling
- Task Ordering
- Dynamic Scheduling
- Task Preemption and Response time
- Task Management
- Task States
- Transitioning between states
- Context Switching
- Shared Data problems and its solutions
- Example: Non preemptive dynamic scheduler



Renesas Electronics America Inc.

© 2014 Renesas Electronics America Inc. All rights reserved.