

# Lecture: Embedded Software Architectures

Jan Vitek

ECE/CS

Spring 2011

# Reading List

2

- **Mandatory Reading**
  - ▶ Chapter 5 of ECS textbook
- **Optional Reading**
  - *N/A*

# Software Architecture

3

- A *software architecture* gives the general structure of an embedded application independent of the actual computation performed
- Choice of architecture impacts issues such as:
  - ▶ development time / likelihood of software defects
  - ▶ responsiveness and latency
  - ▶ code size / complexity
- Rule of thumb:
  - ▶ *Select simplest architecture that meets application requirements*
  - ▶ Any extraneous complexity/generality costs additional development and verification effort

# Software Architectures

4

- Four well known choices:
  - ▶ Simple Round Robin
  - ▶ Round Robin with Interrupts
  - ▶ Round Robin with Interrupts and Function Queues
  - ▶ Real-time Operating System-based architectures
- The architectures are sorted in order of increasing complexity
- Round Robin (RR) architectures are also called *Cyclic Executives* in real-time literature
- The main different between RR and RTOS-based approaches is that in RR scheduling and admission control is done by the developer as opposed to leaving it to the OS

# Round Robin

5

- Simplest architecture, a single loop checks devices in predefined sequence and performs I/O right away

```
1. while(1) {
2.     if (device_1_ready()) { /*Perform D1 I/O and relate computation.*/ }
3.     if (device_2_ready()) { /*Perform D2 I/O and relate computation.*/ }
4.     ...
5.     if (device_N_ready()) { /*Perform DN I/O and relate computation.*/ }
6. }
```

- Works well for system with few devices, trivial timing constraints, proportionally small processing costs
- Response time of device  $i$  equal to WCET of the body of the loop

# Round Robin

6

- Periodic Round Robin

- ▶ In case the system must perform operations at different frequencies
- ▶ Add code to wait a variable amount of time

```
1.while(1) {
2.    waitForNextPeriod(10); // idle for up to 10 ms
3.    if (device_1_ready()) { /*Perform D1 I/O and relate computation.*/ }
4.    ...
```

- Exercise:

- ▶ *Think of how to implement a loop that runs every 10 ms and measures the drift*

# Round Robin

7

- **Limitations of the architecture:**
  - ▶ If some devices require small response times, while other have large WCET it will not be possible to guarantee that all timing constraints will be met
  - ▶ The architecture is fragile, adding a new task can easily cause missed deadlines
  
- ▶ **Question:**
  - Is the order in which devices appear significant?
  - Same question, but with code for devices having different processing times and timing constraints?

# Round Robin with Interrupts

8

- Hardware events requiring small response times handled by ISRs
- Typically ISRs do little more than set flags and copy data

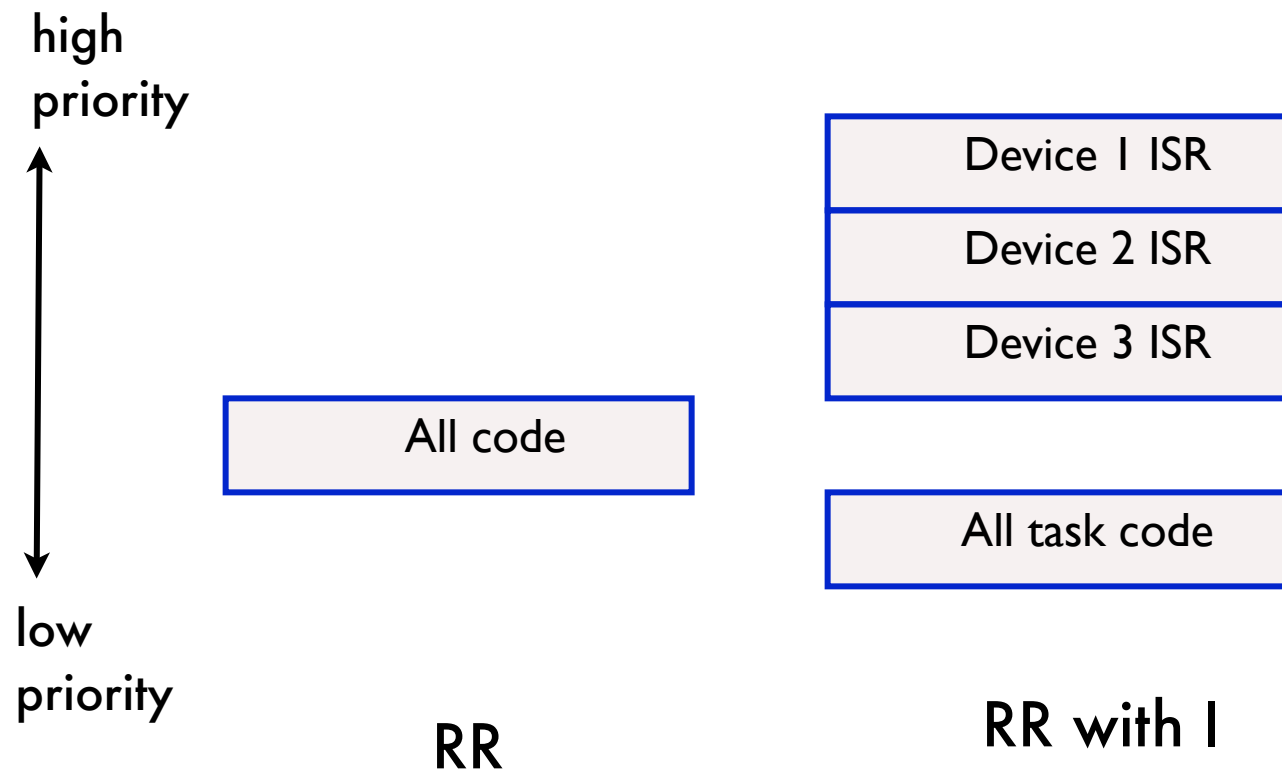
```
1. bool f_device_1 = FALSE;
2. bool f_device_2 = FALSE;
3.     ...
4. void interrupt handle_dev_1() {
5.     // handle device 1
6.     f_device_1 = TRUE;
7. }
8.     ...
9. void main() {
10.    while (1) {
11.        if(f_device_1) {
12.            f_device_1 = FALSE;
13.            // do processing related to device 1...
14.            if (f_device_2) {
15.                ...
16.            }
17. }
```



# Round Robin with Interrupts

9

- Latency of an ISR is function of response time of higher priority ISRs
- Lower bound on latency of RR loop is response time of the ISRs



# Round Robin with Interrupts

10

- Drawbacks

- ▶ All task code executes at same priority

- One can test some flags multiples times within loop body to reduce latency

- ▶ Shared data bugs

- ▶ Question:

- *What if one of the device requires large amount of processing time (larger than the time constraint of others?)*

# RR+I and Function Queues

11

- Rather than fixed order, program manages order of execution

```
1. #define DEV_1 1
2. #define DEV_2 2
3. ...
4. void interrupt handle_DEV_1() {
5.     // deal with device
6.     enqueue(DEV_1);
7. }
8. ...
9. void main() {
10.    while (1) {
11.        switch (dequeue()) {
12.            case DEV_1:    // process DEV_1
13.                break;
14.            ...
15.            default:      // empty queue nothing to do
16.        }
17.    }
18.}
```

# RR+I and Function Queues

12

- One could use function pointers, but they add complexity
  - ▶ *FP are useful if one does not want to hardwire the devices in the main loop*
- `enqueue()` reorders queue to improve latency of high priority devices
- For long running functions: break them up into multiple smaller units
  - ▶ *Question: does that improve latency?*
  
- Question
  - ▶ *Consider implementation of `dequeue()`, what kind of data structure would you use (why), is special care needed?*

# Real-time Operating Systems

13

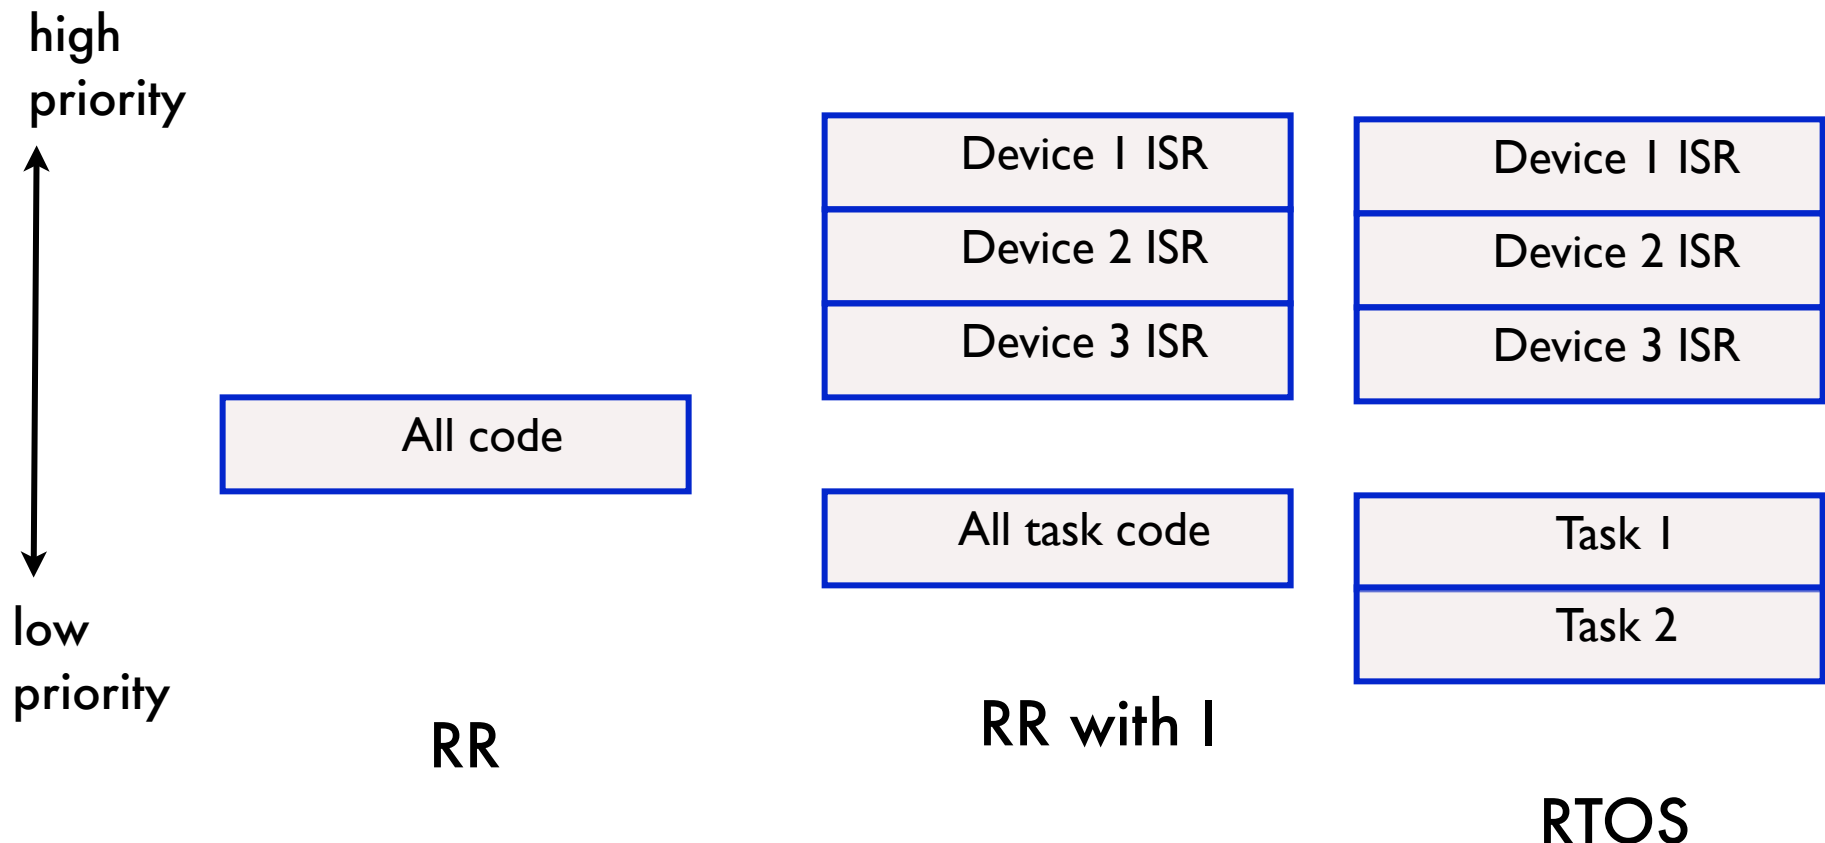
- Rely on the operating for scheduling tasks
- Leverage preemptive scheduling to ensure that deadlines are met

```
1.  static pthread_t thread_1;
2.  ...
3.  void interrupt handle_DEV_1() {
4.      // handle device
5.      CHECK( pthread_wakeup(thread_1) );
6.  }
7.  ...
8.  void task_1() {
9.      while (1) {
10.         pthread_suspend_np();
11.         // process device 1 I/O
12.     }
13. }
14. ...
```

# Real-time Operating Systems

14

- The scheduler in a RTOS takes care of scheduling all tasks according to their priority
- Long running, low priority, tasks can be preempted by higher priority ones



# Real-time Operating Systems

15

- **Services that an RTOS *could* provide:**
  - ▶ **Scheduling tasks**
    - create/terminate threads
    - timing threads operations
    - preemption
  - ▶ **Synchronization**
    - semaphores and locks
  - ▶ **Input/Output**
    - interrupt handling
  - ▶ **Memory management**
    - separate stacks
    - segmentation
    - allocation/deallocation
  - ▶ **File system**
    - persistent store
  - ▶ **Security**
    - user vs. kernel space
    - identity management

# Conclusion

16

- Software architectures describe the structure of a system independently of its function
  - ▶ *Round Robin* is a simple architecture for devices with few (or uniform) timing constraints
  - ▶ *Round Robin with Interrupts* extends RR with low-latency interrupt handling
  - ▶ *Round Robin with Interrupts and Function Queues* allows dynamic scheduling of tasks under programmatic control
  - ▶ *Real-time Operating Systems* relieve programmers from having to deal with scheduling and time management