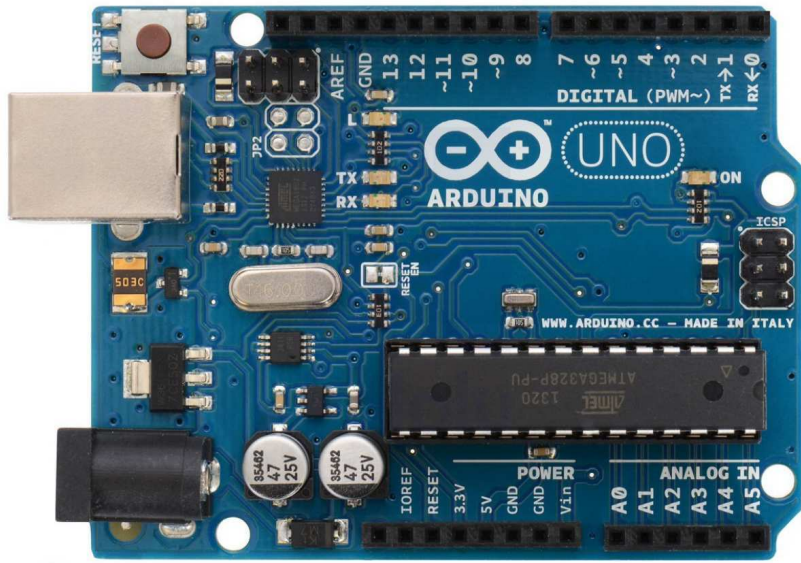


Example of a Target Embedded System



What is a Real-Time System?



Real-Time Systems have been defined as:

” Those systems in which the correctness of the system depends not only on the logical results of the computation, but also on the time at which the results are produced”

(J. Stankovic, "Misconceptions About Real-Time Computing", IEEE Computer, 21(10), October 1988)



What is a Real-Time System? (Cont.)

More precisely, a system is said to be a *Real-Time System* if:

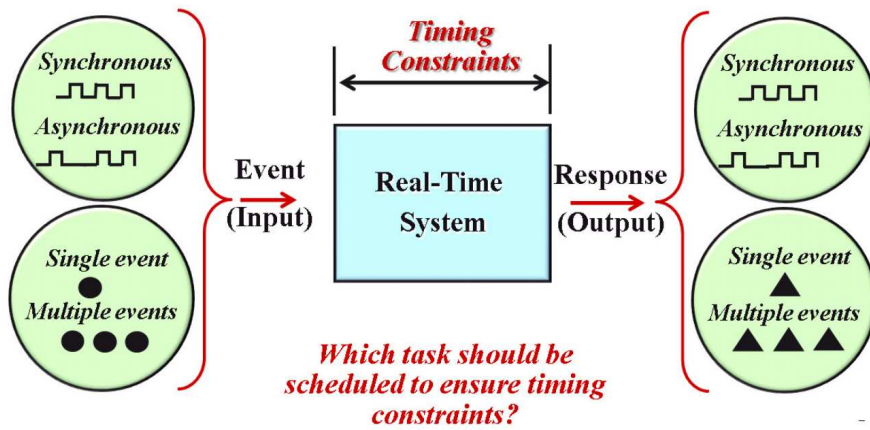
- It reacts to external stimuli
 - System interacting with its environment
- It reacts to these stimuli within a limit of time, regardless of the load of the system
 - **WCET = Worst Case Execution Time**
 - Timing correctness

Vs. usual systems

- Only logical correctness
- Manage the average case
 - "Best effort"



Response to External Events





Time Constraints

Time constraint =
Constraint imposed on the timing behavior of each job of a task

Release time

- Instant of time when a job becomes available for execution

Deadline

- Instant of time at which a job is required to be completed
- Absolute deadline* = release time + relative deadline

Response Time

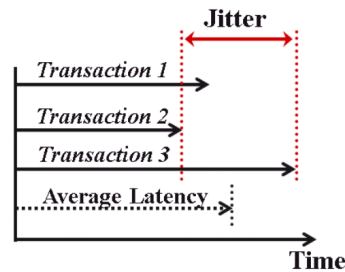
- Difference between the completion time and the release time of a job



Latency and Jitter

- System latency** = End-to-end delay between a change of state in the environment, and the corresponding output reaction produced by the system
 - Delay for scanning the environment, delay due to the OS, delay for executing calculations, delay for producing output results (communication delay)

- Jitter** = Incertitude on delays
 - Load of the system, etc.
 - Should be low with regards to the latency



Example #1: GSM Communication

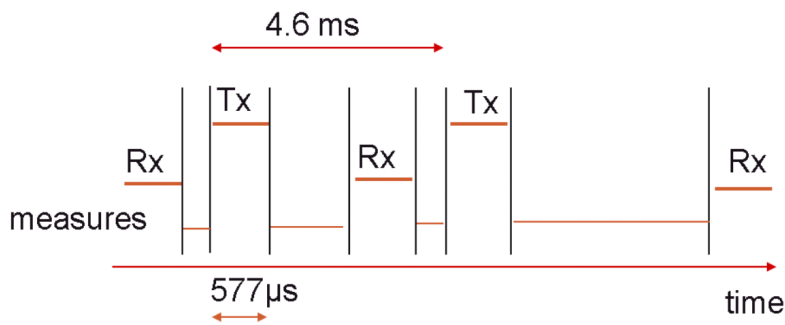


Specification

- Software is in charge:
 - Of some operations performed at physical level (receiving, sending, measuring the level of the electromagnetic field, etc.)
 - Of logical operations such as scanning for incoming calls, maintaining connections when changing of beams, etc.
- Sending and receiving voice data is a critical task
 - 577μs of voice data are transmitted every 4.6ms
 - 577μs of voice data are received every 4.6ms
- Management of the mobility issue
 - If the distance to the relay is increasing, data should be emitted earlier to remain synchronized with the relay (Doppler effect)



Example #1: A GSM Phone (Cont.)



Constraints

- Timing constraints on sending and receiving
- Global constraint: audio quality



Key Characteristics (Cont.)



Performance

- CPU performance: MIPS
- Throughput performance: bps

Compactness

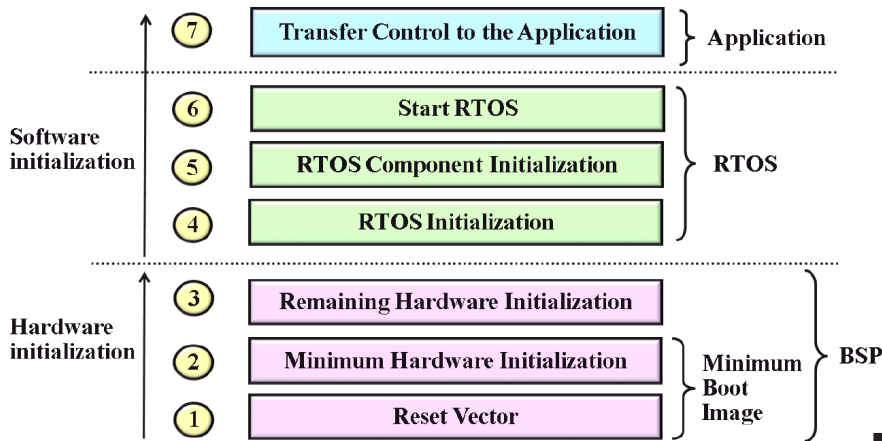
- RTOS memory footprint should be small

Scalability, i.e handling application-specific requirements

- Modular components
- File systems, protocol stacks, etc.



Target Initialization Sequence



Scheduling



Limitations of schedulers in general-purpose OS

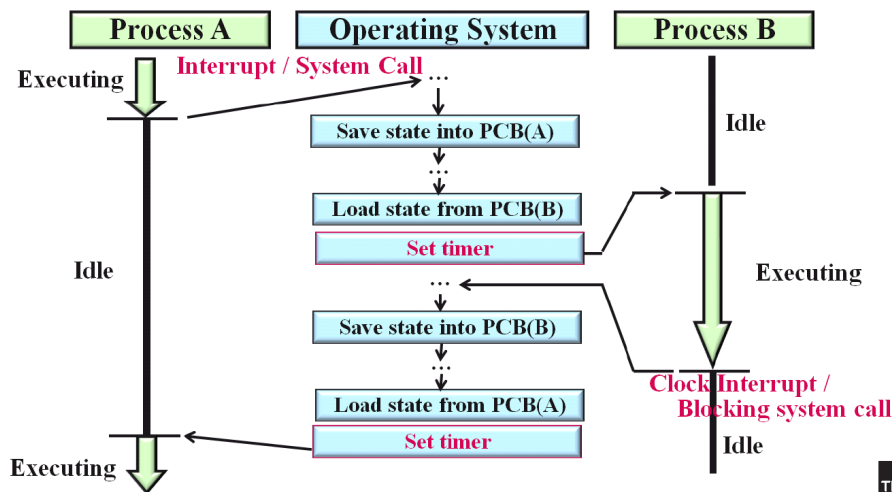
- Schedulers are defined with interactivity in mind
 - Dynamic priorities
 - Calculation tasks are penalized
- Timing constraints are not taken into account
- Schedulability analysis cannot be performed

Schedulers for Real-Time Systems

RMA, EDF, FCFS, Round-Robin, Fixed priority



Preemption Latency in a Priority-Based System



IPC - Inter Process Communications



Limitations for RT systems

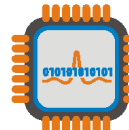
- No deterministic transmission delay for:
 - Signals, pipes, message passing, semaphores, (shared memory)

A Few Solutions

- Implementation with deterministic processing time
 - Example: RT-FIFO
 - RT-FIFO = FIFO with predefined maximum size and with pre-allocated messages



Input / Output Management



Limitations for RT systems

- Difficult or impossible to bound the time to perform an I/O operation

Solutions

- Drivers working closely with the kernel
- Deterministic delay in all inter-process communications
 - Communication structures with basic and deterministic algorithms
 - RT-FIFO, RT-signals, etc.





Reentrant Runtime Libraries

Issue

- Libraries offer services such as: Input / Output operations, dynamic memory allocations, etc.
- Services may have an internal state: buffer (I/O), pointer to the next chunk of free memory, etc.
- If preemption occurs when a service is being called, internal state of the service may be corrupted

Solutions

- Protection at application level: mutex, semaphore, etc.
 - Safe, but costly
- Use a reentrant version of the library (if available)
 - Safe, less costly (but more costly than non-reentrant libraries)



Exceptions and Interrupts

Exceptions

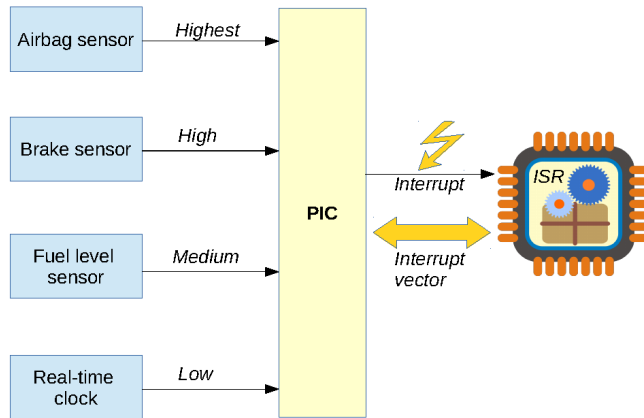
- *Events that disrupt the normal execution flow of the processor and force the processor to execute special instructions in a privileged state*
- Synchronous exceptions
 - Divide by zero
- Asynchronous exceptions
 - Pushing the reset button of the embedded system

Interrupts

- = *Asynchronous exceptions triggered by events external hardware devices generate*



Programmable Interrupt Controller



Programmable Interrupt Controller (Cont.)

- **Hardware-implemented**
- Prioritizing multiple interrupt sources so that at any time the highest priority interrupt is presented to the core CPU for processing
- PICs have a set of Interrupt Request Lines
 - Interrupt = physical signal on a given line
- Handlers specified at PIC level
 - **ESR** = **E**xception **S**ervice **R**outine
 - **ISR** = **I**nterrupt **S**ervice **R**outine



Programmable Interrupt Controller: Example

Source	Priority	Vector Address	IRQ	Max freq.	Description
Airbag sensor	Highest	14h	8	N/A	Detects the need to deploy airbags
Brake sensor	High	18h	7	N/A	Informs about the current braking power
Fuel Level Sensor	Med	1Bh	6	20Hz	Detects the level of fuel
RTC	Low	1Dh	5	100Hz	Clock runs at 10ms tick



Masking Interrupts



Why masking?

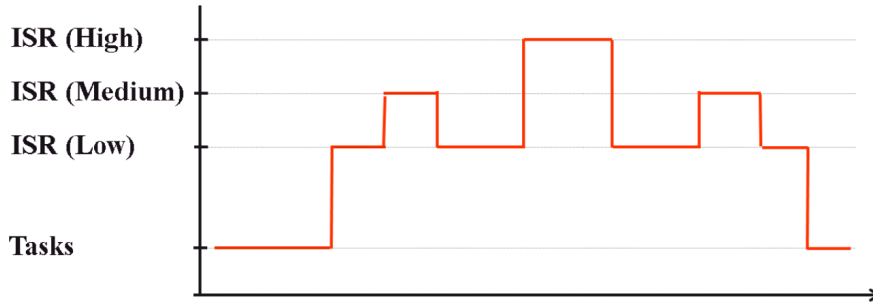
- Reduce the total number of interrupts raised by devices
 - Warning: Interrupts may be lost when masked
- Atomic operations

How is masking done?

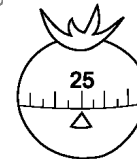
- Disable the device so that it cannot assert additional interrupts
- Mask the interrupts of equal or lower priority levels
- Disable the line between the PIC and the core processor
 - Interrupts of any priority level do not reach the processor



Nested Interrupts



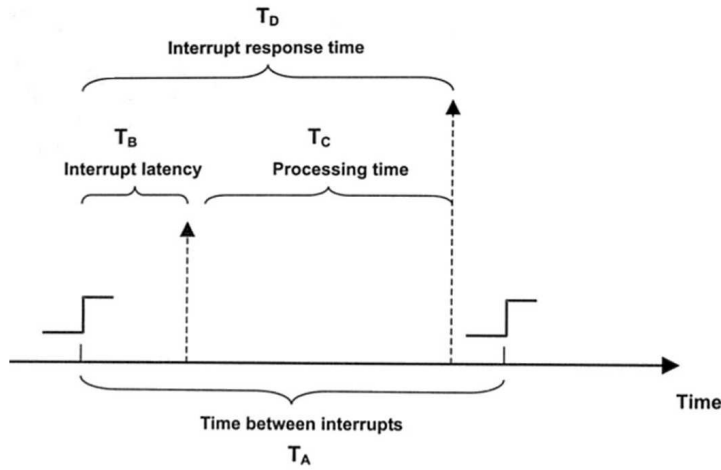
Exceptions Timing



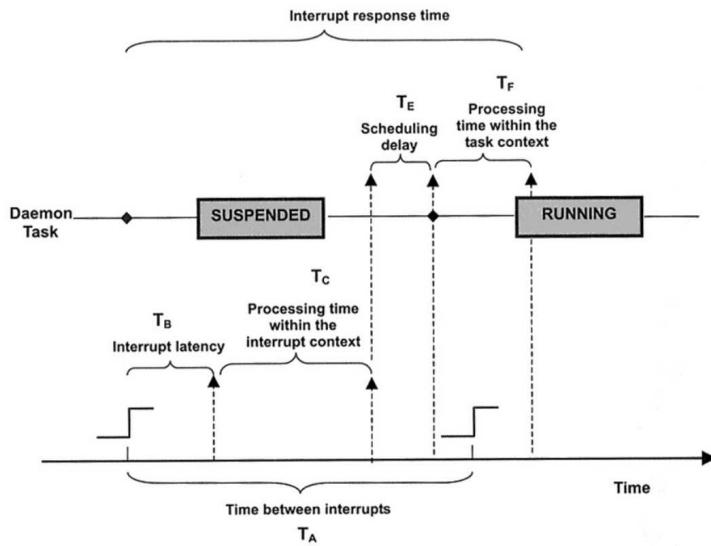
- ESR and ISR should be short
- Designers
 1. Should avoid situations where interrupts could be missed
 2. Should be aware of the interrupt frequency of each device
 - Maximum allowed duration can be deduced from this information
 3. Should take into account the fact that real-time tasks cannot execute when interrupts are being processed



Exceptions Timing (Cont.)



Cutting ISRs in Two Parts



Conclusion on ISRs

Benefits about cutting ISRs in two parts

- Lower priority interrupts can be handled
- Reduces the chance of missing interrupts
- Increase concurrency because devices can start working on next operations earlier
- Urgent tasks have a chance to be executed sooner

General guide for implementing ISRs

- An ISR should disable interrupts of the same level
- An ISR should mask all interrupts if it needs to execute a sequence of code as one atomic operation
- An ISR should avoid calling non-reentrant functions and blocking system calls

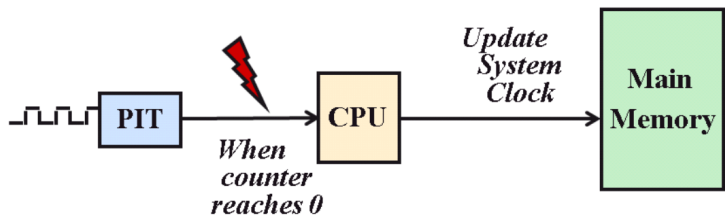


Timer and Timer Services

Use of Timers

Applications often have to perform jobs after a given delay has elapsed

Programmable Interval Timers (PIT)



Timer Interrupt Service Routine

- Updating the system clock
 - Absolute time (calendar)
 - Elapsed time (since power up)
- Calling a registered kernel function to notify the occurrence of a pre-programmed period
 - `announce_time_tick()`
 - Calls the Scheduler and the *Soft-timer* facility
- Acknowledging the interrupt, reinitializing the necessary timer control register and returning from interrupt



Steps in Servicing Timer Interrupt

