

Embedded 1-12

Lecture 1

- Real Time System: Computer system that interacts w/ external events
- Embedded system: RT.S. that interfaces w/ external HW.

RTOS vs OS

Common: Both are interfaces between HW & SW.

Differ: RTOS gives user access to low level system devices.

RTOS processes

- Periodic: system monitoring, polling, sampling
- Sporadic: event driven (interrupt)

Lecture 2

Live lock: Tasks are allocated CPU time by the scheduler but don't execute further instructions.

Spin lock: Tasks are running & continuously polling on each other's unavailable resources.

Dead lock: Tasks are not running due to unavailable resources.

CPU usage	Resource usage	
	Yes	No
Yes	Spin	Live
No	Dead	Starvation

Lecture 3

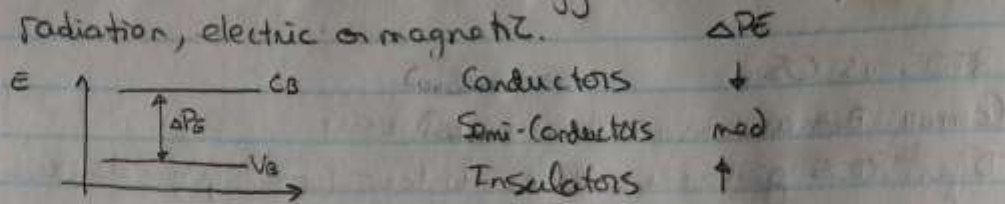
Sources of e- Ep.

- Electric field $\sim \vec{E} \cdot d\vec{l}$ \rightarrow Most commonly used
- Magnetic field $\sim \vec{\mu} \cdot \vec{B}$ \rightarrow Most efficient
- Heat $\sim nkT$ \rightarrow most problematic
- Radiation $\sim h\nu$ \rightarrow most flexible.

EMD

POI

tunneling: Quantum mechanical process used by transistors to allow electrons to jump from the valence band to the conduction band. The energy can come from heat, radiation, electric or magnetic.



$$E_T = E_p + E_k \rightarrow \frac{1}{2}mv^2 = \frac{1}{2} \frac{p^2}{m}$$

Lecture 4: Transistors

BJT: (A, E, C), bipolar since operation involves mobile e^- & p^+ . EB is forward biased & BC is reverse biased. Problem w/ heat dissipation & power loss.

MOSFET: Electronically varying the width of a conduction channel. Charge enter via source & leaves via drain. Voltage controlled by the gate. Advantages: gate insulated, so no current from gate to channel. Problems: SiO_2 so thin that susceptible to permanent damage from electric charges & does not perform well at weak RF signals.

Functions in

- depletion mode: As voltage applied, conduction ↓
- enhancement mode: Voltage applied = conduction ↑

	BJT	MOSFET	Power loss $\propto I^2 R$
Power (Heat)	more	less	
Switch speed	Slow	Fast	$V_{eff} \propto Power\ loss \propto \frac{dI}{dt}$
Gain speed	Fast	Slow	Fast switching speed $\propto \frac{dI}{dt}$ (BJT faster)
Current	mA	μA	
Susceptible to noise	Less	more	$V_{eff} = I_{eff} R \rightarrow \frac{L dI}{dt} \propto Heat$

EMB

POZ

CMOS: High speed, low noise, low power
No net current traveling across device \Rightarrow only power for switching = less heat

Memory cells

MRAM (Magnetoresistive) depends on SPINTRONICS (study of electron spin polarity) and uses spin transistors.
Fast speed, high density, low power, non-volatile, radiation hardness.

NRAM (Nanotubes): Can be insulator, semiconductor or conductor depending on Carbon gyration. Maintains charge w/o power.
Advantage: High density & speed.

FeRAM (Ferroelectric) 2 MOSFET + 2 Ferroelectric capacitors.
Polarization maintained w/o power. R/W faster than flash (1eV vs flash uses bombardment of electrons)

Architectural views

- Component view: large scale pieces & how they fit together
- Distribution view: Concurrency & communication protocols
- Resource view: Resource management for concurrent applications
- Reliability view: Redundancies & safety protocols, fault tolerance & recovery
- Deployment view: SW mapping on HW, memory mapping, task distributions, data flow.



LECTURE 5


	$T_{sec} = \frac{1}{f_{clk}} \times \text{ins eff.}$	
DSP	$10 \text{ GIPS} = 10 \text{ GHz}$	100%
CRC	$\frac{1}{2} \text{ GIPS} = 1 \text{ GHz}$	50%
RISC	$2 \text{ GIPS} = 1 \text{ GHz}$	200%

EMB

P03

Boundary Conditions

Reflection  Short 



$$V = V_1 + V_2 \quad I = I_1 - I_2 \quad \text{Also, } I = I_2 - I_1$$
$$IR = (I_1 R) + (I_2 2R) \quad I_2 - I_1 = I_1 + 2I_2$$
$$I = I_1 - 2I_2 \quad \boxed{-2I_1 = I_2}$$

Real time monitoring

- SW monitor to breakpoint

Pros: No extra HW Support needed

Cons: Requires target mem & CPU time

- On chip background debugger program

Pros: Reduces CPU overhead for RT applications

Cons: Eats up target resources

- In-circuit Emulator (ICE) or VxWorks

Pros: No overhead on target, non-disruptive

Cons: Expensive, hard to simulate RT, unreliable socket contacts.

Addressing

Direct: IR 2 bytes (op code + Addy). More space, but more flexibility

Inherent: IR 1 byte (op code), CPU already knows target Data Reg. More efficient, faster, less flexible.

Lecture 6


System Types

- Distributed: Slow, large
- Mission Critical: Reliable, 8
- Signal processing: Fast

Project lifecycle phases.

1. Conception stage: goals, requirements for interface, mem, interrupts, systems, CPU mem throughput & bottlenecks.
Output: MOU (Memorandum of understanding).
2. Requirements phase: Customer specific goals, technical requirements,
Output: SRS.
3. Design: build drivers & sub systems, ISRs, use simulation like FPGAs, VHDL
Output: SDS (Charts & diagrams)
4. Implementation: Code, build & test
Output: HW devices, SW Drivers & modules
5. Integration: unit test every modules, put them together and system test
Output: System test cases document (tests that proves that every requirements are met).
6. Release & Maintenance: Fix bugs, upkeep & upgrades.


Class diagram




ETB

POS.


Deployment diagram: models the physical aspect of Object Oriented SW System. Models a static view of Run-time configuration distribution of components & applications.



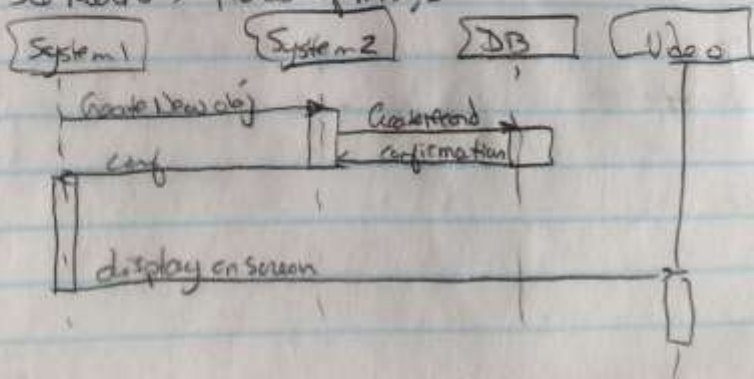
Component Diagram: Components wired together. Illustrates the service provider & consumer relationships.



Activity diagram: Detailed behavior inside a functional requirement



Msg Sequence diagram: behavior of a single path of a use case Scenario \Rightarrow flow of msgs.



EYB
p 06

LECTURE 8: VHDL

ENTITY abc IS

PORT (

a, b : IN STD_LOGIC;

c : OUT STD_LOGIC);

END abc

ARCHITECTURE guts OF abc IS

Component ... IS

port (...)

Signal ...

BEGIN

Internal Delay Model

~~Q<= R NOR NO AFTER INS;~~

~~Q<= S NOR Q AFTER INS;~~

only happens if hold time
is respected

Transport Model

~~Q<= TRANSPORT R NOR NO AFTER INS~~

~~Q<= TRANSPORT S NOR Q AFTER INS~~

happens no matter what
(detects rising edge)

LECTURE 9: MEMORY MGMT

- Swapping: Program size too large. When changing context, put process into secondary storage device. (slow)
- Overlaying: virtual memory. (slow & memory corruption)
- Partitioning: Mem is partitioned in blocks @ compile time. Fast, but results in fragmentation
- Paging: Fixed size blocks of program segment are stored in non-contiguous mem frames. Involves page table = overhead = slow
- Mem locking. (good for RT) enhances performance.
- Mem leak: allocated mem is not freed after use
- Contention: 2+ tasks try to access same mem @ same time

EMB

PO9

Allocation	Stack	Heap
Queue	Complete	Runtime
Organization	LIFO	Garbage dump
	Daisy Chained	Allocated on the fly so all over

Queueing -

- FIFO: loss free, sender & receiver isolated, no data loss
 ⇒ Miss on critical: No corruption, safer
- Ring buffer: Circular FIFO, fast but can lead to corruption
 ⇒ Signal processing: FAST
- Swing buffer: Series of ring buffers in parallel, R/W sequentially
 ⇒ FIFO of rings or 2+ ring buffers, lost prone & non blocking
 ⇒ Distributed: bigger mem space

Lecture 10: Memory Communication

- Pipes: I/O between parent & child processes. P is frozen until pipe is cleared
- Semaphores: integer w/ 2 functions, Vcs) Verhogen = increment, Pcs) = pdecren = decr.

Queue	Mailbox	Socket	Pipe
Faster	Slower	Async	Sync
R/W @ Same place	Copy msg	1 to many	Unrelated I/O
Less mem	more mem	inter-platforms	Shared memory (single platform)
Sync	Async		

Lecture 11: Jitter

Jitter: time between arrival & expected arrival.

$$J = \sqrt{\sum j_i^2}$$

END

P10.

Measuring Exec time

polling	Interruption	Logic Analyzer
Cheap	Cheap	Expensive
Slow	Faster	Fastest
CPU off	CPU on	No off
Src of Jitter	Src of Jitter	Instant.

Clock Requirements

- Correctness: $clock - RT \leq \epsilon$ (is it the real time?)
- Bounded drift: $\frac{dt}{dt} < k$ (does it drift over time)
- Monotonicity: Take 2 measures. if $t_1 > t_2$, then $b_1 > b_2$
- Chronoscopia: Measure something some day wrt some other day

Lecture 10: Smart Cards.

wired	RFID
cheap	expensive
Short Range	Long Range (flexibility)
less mem	More mem
Secure	Unsecure
Integrity	error.

Transmission

Async	Sync
Separate or No CLK	Shared CLK
Common	less common
Cheaper	exp.
Slower (hand shaking)	Faster
Higher error	less error

RFID: Smart cards where communication & power are transferred via RF
3 components: Interrogator (reader), tag, host computer

Tag: μ Chip w/ $\sim 2\text{Kbit}$.

- passive: only power from induction from RF of Reader
- Active: Own power to amplify return signal
- Semi-passive: Power src to turn on, but not to amp signal

Tag is composed of an antenna coil & Silicon chip (CPU, RAM, EEPROM).

Reader: Control R/W on tag. Generates RF field around antenna directed at the tag \rightarrow Carrier signal \rightarrow Tag gets power & outputs ID \rightarrow backscattering signal.

Inductive magnetic coupling @ 125kHz - 13.6MHz, up to 900MHz. Higher frequency = better range, metal penetration, faster, more power

Optical sensor: Temperature sensing of fingerprint
cross sensor \rightarrow Cells contain a photodiode that converts light to electrons by photoelectric effect

$E_e = E_{\text{photon}} - [E_{\text{conduction}} - E_v]$ Temperature is sensed by the photodiode noticing wavelength of light detected.

XY addressing

CCD: Camera store charge according to the intensity of the image being detected. # of electrons collected proportional to light intensity. Light collected over entire image simultaneously then transferred

EM3

PI2

Optical

CCD

photoelectric
measures
wavelength &
charge
cheaper
XY addressing
faster
More noise
Less power

photoelectric
measures
light intensity &
charge produced
expensive
FIFO
Slower
less noise
More power
More jitter because of FIFO.

L14: Reliability

RT fundamental requirement is to meet its deadlines

Sync

RT blocking until buffer cleared
 less OH.
 Shared clock
 Avoids race condition
 less flexible

Async (preferred for RT)

RT done in parallel
 Need a I/O control block
 Independent clocks
 More computation
 Avoid missed deadlines *
 More corruption


Layers:

Service provisional point (SPP) \Rightarrow Drives the SAP \updownarrow layers.
 Service Access points (SAP) \Rightarrow Services the SPP.


Reliability




$$R = R_A * R_B$$



$$R = R_A + R_B - R_A * R_B$$



$$R = R_D * [R_A + R_B + R_C - R_A R_B - R_A R_C - R_B R_C + R_A R_B R_C]$$



$$R = [A + B + C - AB - AC - BC + ABC] * D$$

L15 Attributes of RTS.

- watchdog {
- Reliability: must meet its deadlines.
 - Predictability: Provide range bound and worst case exec. times.
 - Scheduling: Requires knowledge of timing constraints \Rightarrow Makes the system reliable. Must know if deadlines are hard or soft.
- clock.

Priority or deadline \neq importance. Dynamic scheduling = OH and more unstable. Unbounded operations are not permitted in RTS. Timeout mechanisms can cause dead locks.

EMB

P14

- **Concurrency**: management of, using Semaphores & Mutex
 - Synchronisation
 - Communication
 - Resource Sharing

- **Visibility**: Allows user access to communication devices and task priority assignments.

⇒ **Priority inversion**: most likely source of missed deadline

- **Priority inheritance protocol**: task priority raised to the level of resource it owns (highest priority of tasks waiting on the resource) ⇒ Not prevent deadlocks.
- **Priority ceiling protocol**: tasks suspended except the one owning the highest priority resource that can run
- **Priority disinheritance**: shift back to normal

LC-Kernels

1) OS top layer

- file mgmt / security
- Task control
- UI / shell

APIs, DLL

FP support

2) Executive

- I/O
- Memory management

task pointers

Stacks, Cache mgmt

Drivers

3) Kernel Top layer

- Communication drivers

Network mgmt

Task communication tools

Task Scheduling

Clocks, timing, Interrupts

4) micro kernel

- Scheduler
- ISR - Interrupts

raw task management

5) Nano Kernel

- Task dispatching
- Task bookkeeping

EMP

PL5

L17 Interrupts

External: Timers, devices

(Generally HW)

Internal: divide 0, Scheduler

(Generally SW)

Software interrupts

- Typically sync machine language op. (no, undefined op)
- Defined in CPU HW so OS independent

Hardware interrupts

- Typically Async external events. (OS & Scheduler have no control)
- Vectored by CPU HW.
- Noise, jitter, ground loops may cause stack overflow.

Interrupt vector: Array of ptrs to ISR. (16 bit / interrupt).

Non-vectored: ctrl is transferred to one single routine that decides how to handle the interrupt.

ISR

- CPU vector to ISR in HW (independent of OS).
- re-entrant: can call itself, can be called at the same time by many processes (use semaphores)

Non maskable interrupts (NMI)

- highest priority
- Can't be stopped (avoid if possible)
- Almost all OS SW interrupts are NMI

Simultaneous Interrupts Handling

- Masking prevents other interrupts from being serviced while one is.
- Pri-Emptying/nesting: higher priority int. can int. lower priority int.
- Queuing: Serviced sequentially.

ETB
p16.

L18 - Interrupt Analysis.

Interrupt sequence

- 1) Update Status register & make internal copy
- 2) Determine priority & mask lower priority interrupt lines
- 3) Ack interrupt priority level
- 4) Save PC & registers on stack
- 5) Set PC to head of ISR & service
- 6) Return from exception & restore task context

Timing interference

$t(s)$ = time for task to run on itself.

$t'(s)$ = time for task to run in the presence of I

T_i = interrupt execution time

f = frequency of interrupt

$$t'(s) = t(s) + t'(s) \times f \times T_i \quad \Rightarrow \quad t'(s) = \frac{t(s)}{(1 - fT_i)}$$

For many interrupts: $t'(s) = \frac{t(s)}{1 - \sum f_i T_i}$

Ex1 Int takes 1% of CPU time \Rightarrow find fT_i & $t'(s)$.

$$\left. \begin{aligned} t'(s) &= t(s) + t'(s) f T_i \\ 1.01 t(s) &= t(s) + 1.01 t(s) f T_i \end{aligned} \right\} \begin{aligned} t'(s) &= \frac{t(s)}{1 - f T_i} = \frac{t(s)}{0.99} \\ 0.01 &= f T_i \end{aligned}$$

Ex2 What is $t'(s)$ in terms of $t(s)$ if I_1 is 2% and I_2 is 4%

$$t'(s) = t(s) + \sum t'(s) f_i T_i$$

Nothing or Clearing

$$t' = \frac{t}{1 - \sum f_i T_i}$$

Harding

$$t' = \frac{t}{1 - \sum f_i T_i + \sum f_i T_i \frac{t'}{t}} = \frac{t}{1 - 0.02 - 0.04 + \frac{(0.02)(0.02) + (0.04)(0.04)}{0.96}}$$

EMB

P17

- Exception Handling: ISR of error condition
- Fault latency: time between fault occurrence & error occurrence
- Error latency: time btw error & its effect on system
- IL (Interrupt Latency): time btw hw interrupt & start of ISR
- CSL (Context Switch Latency): time between last instruction of T1 and first inst. of T2.
- DL (Dispatcher Latency): time between end of ISR & restart of original task (kernel rescheduling, set enable & context switch)
- WISRT (worst case ISR execution handling time): max time to execute ISR
- WIR (worst case interrupt response time): $IL + WISRT$
- WITR (worst case interrupt task response time): $WIR + DL$
 $= IL + WISRT + DL$
- TRR (Task Response Time): task time from start to finish



L19 - Priority.

- Pre Emption: higher priority T. stealing CPU
- Priority inversion: low priority blocking high from CPU because of a resource it owns.
- Unbounded priority inversion: priority inversion, but low is blocked by higher priority.



ERB
 PIB

p = priority
T = Task

Priority Inheritance Protocol: temporarily raise low P.T. to the level of a higher P.T. waiting on a resource the low P.T. currently owns.

Priority ceiling Protocol:

- 1) Assign priority to SR = priority of Task that locks Z.
- 2) Suspend tasks requesting a resource except T that owns the resource with the highest p.



Period transformation:

- Fixed rate monotonic scheduling \Rightarrow task assigned a fixed p. depending on its period higher frequency (freq) = higher priority (PT)

Round Robin: every one gets a turn


Pre-emptive

Generic Fixed P. automatically assigned & fixed.

Rate Monotonic (RM)

- Priority assigned by period.
- period = deadline (ignore other deadline)
- Most CPU efficient
- Total utilization $\mu_T = \sum e/p$.

EMB
P19.



Deadline Monotonic (DM)

- Shorter deadline has higher priority.
- Dead line utilization $\mu = \sum e/d \leq 1$ will work
 $\mu > 1$ may work.

Least Computation Time (LCT)

- Shortest e has higher priority

Fixed Utilization

- $\mu = \sum p$ Greatest u has highest priority

Earliest deadline first (EDF)

- At any time, T w/ closest deadline has highest P.
- $(d - t)_{\min}$
- Optimal algorithm for uniprocessor (CPU μ is maxed)
- All tasks schedulable if $\sum p \leq 1$

Shortest Completion Time (SCT)

- Least exec time remaining has highest p.
- $e_{\text{left}} = (e_{\text{tot}} - e_{\text{done}})_{\min}$

Least Slack Time (LST)

- Task has the least amount of time to meet deadline
- $[(d - t) - (e_{\text{tot}} - e_{\text{left}})]_{\min}$