




## Periodic scheduling

## REVIEW



- **Given:**
  - A set of periodic tasks  $\Gamma = \{\tau_1, \dots, \tau_n\}$  with
    - phases  $\Phi_i$  (arrival times of first instances of tasks),
    - periods  $T_i$  (time difference between two consecutive activations)
    - relative deadlines  $D_i$  (deadline relative to arrival times of instances)
    - computation times  $C_i$
  - ⇒  $j$ th instance  $\tau_{i,j}$  of task  $\tau_i$  with
    - arrival time  $a_{i,j} = \Phi_i + (j-1) T_i$ ,
    - deadline  $d_{i,j} = \Phi_i + (j-1) T_i + D_i$ ,
    - start time  $s_{i,j}$  and
    - finishing time  $f_{i,j}$

- **Find a feasible schedule**

## Rate monotonic scheduling (RM)

## REVIEW

- Rate monotonic scheduling (RM) (Liu, Layland '73):
  - Assign **fixed priorities** to tasks  $\tau_i$ :
    - $\text{priority}(\tau_i) = 1/T_i$
    - I.e., priority **reflects release rate**
  - **Always execute ready task with highest priority**
  - Preemptive: currently executing task is preempted by newly arrived task with shorter period.
- **Theorem (Liu, Layland, 1973):**  
RM is **optimal among all fixed-priority** scheduling algorithms.

BF - ES

- 3 -

## Schedulability check

## REVIEW

- $U_{lub} = n(2^{1/n} - 1)$
- **A set of tasks can be scheduled by RM if**  
 $U < U_{bnd}(RM) = \ln 2 \approx 0.69$
- But what can we tell about **schedulability** when processor utilization factor is **larger than**  $n(2^{1/n} - 1)$  ?
- Answer:  
We can compute a more precise result, if we make use of the knowledge of periods  $T_i$  and computation times  $C_i$ .

BF - ES

- 4 -

## Schedulability check

## REVIEW

- Remember:  
The response time  $R_{i,j}$  of an instance  $j$  of task  $i$  is the time (measured from the arrival time) at which the instance is finished:  $R_{i,j} = f_{i,j} - a_{i,j}$ .
- Compute an upper bound  $R_i$  on the response time:
  - Suppose that  $\tau_1, \dots, \tau_n$  are ordered with increasing periods (i.e. decreasing priorities).
  - Consider an arbitrary periodic task  $\tau_i$ .
  - At a critical instant  $t$ , when an instance of  $\tau_i$  arrives together with all higher priority tasks, we have:
    - $R_i = C_i + \sum_{k=1}^{i-1} (\# \text{ activations of } \tau_k \text{ during } [t, t + R_i]) \cdot C_k$   
 $= C_i + \sum_{k=1}^{i-1} \lceil R_i / T_k \rceil \cdot C_k$

BF - ES

- 5 -

## Schedulability check

## REVIEW

- Solution of fixed point equation?
- Compute the following sequence:
  - $R_i^{(0)} = C_i$ .
  - $R_i^{(j+1)} = C_i + \sum_{k=1}^{i-1} \lceil R_i^{(j)} / T_k \rceil \cdot C_k$ .
- It is easy to see that this sequence is monotonically increasing, i.e.,  $f(x) = C_i + \sum_{k=1}^{i-1} \lceil x / T_k \rceil \cdot C_k$  is monotonically increasing.
- $\Rightarrow$  If a least fixed point of  $f(x)$  exists, then the sequence converges to this fixed point.

BF - ES

- 6 -

## Schedulability check

REVIEW

⇒ Algorithm:

$\forall i: R_i^{(0)} = C_i$

**repeat**

$\forall i: R_i^{(j+1)} = C_i + \sum_{k=1}^{i-1} \lceil R_i^{(j)} / T_k \rceil \cdot C_k$

**until** ( $\exists i$  with  $R_i^{(j+1)} > D_i$ ) **or** ( $\forall i R_i^{(j+1)} = R_i^{(j)}$ );

**if** ( $\forall i R_i^{(j+1)} = R_i^{(j)}$ ) **then**

**report**("RM schedulable");

BF - ES

- 7 -

REVIEW

## Rate Monotonic Scheduling in Presence of Task Dependencies


BF - ES

- 8 -

## The priority inversion problem

## REVIEW

- **Priority inversion** can occur due to resource conflicts (exclusive use of shared resources) in fixed priority schedulers like RM:




- Here: Blocking time equal to length of critical section.

BF - ES

- 9 -

## The priority inversion problem

## REVIEW



- Blocking time equal to length of critical section + computation time of J<sub>2</sub>.
- **Unbounded time of priority inversion**, if J<sub>3</sub> is interrupted by tasks with priority between J<sub>1</sub> and J<sub>3</sub> during its critical region.

BF - ES

- 10 -

## Coping with priority inversion: The priority inheritance protocol

## REVIEW

### Idea of **priority inheritance protocol**:

- If a task  $J_h$  blocks, since another task  $J_l$  with lower priority owns the requested resource, then  $J_l$  inherits the priority of  $J_h$ .
- When  $J_l$  releases the resource, the priority inheritance from  $J_h$  is undone.
- Rule: Tasks always inherit the highest priority of tasks blocked by it.


BF - ES

- 11 -

## Direct vs. push-through blocking

## REVIEW

- **Direct blocking:** High-priority job tries to acquire resource already held by lower-priority job
- **Push-through blocking:** Medium-priority job is blocked by lower-priority job that has inherited a higher priority.




BF - ES

- 12 -

## Transitive priority inheritance

REVIEW



BF - ES

- 13 -

## Schedulability check

REVIEW

Let  $B_i$  be the maximum blocking time due to lower-priority jobs that a job  $J_i$  may experience.

$\forall i: R_i^{(0)} = C_i$

**repeat**

$\forall i: R_i^{(j+1)} = C_i + B_i + \sum_{k=1}^{i-1} \lceil R_i^{(j)} / T_k \rceil \cdot C_k$

**until**  $(\exists i \text{ with } R_i^{(j+1)} > D_i)$  **or**  $(\forall i R_i^{(j+1)} = R_i^{(j)})$ ;

**if**  $(\forall i R_i^{(j+1)} = R_i^{(j)})$  **then**


**report**("RM schedulable");

BF - ES

- 14 -

## Problem: Chained Blocking

REVIEW





BF - ES

- 15 -

## Problem: Deadlock

REVIEW



BF - ES

- 16 -



## Priority Ceiling Protocol

## REVIEW


- Each semaphore  $S$  is assigned a **priority ceiling**:  
 $C(S)$  = priority of the highest-priority job that can lock  $S$
- The processor is assigned to a ready job  $J$  with highest priority.
- To enter a critical section,  $J$  needs priority  $> C(S^*)$ , where  $S^*$  is the currently locked semaphore with max  $C$ .  
 → otherwise  $J$  „blocks on semaphore“ and priority of  $J$  is inherited by job  $J'$  holding  $S^*$ .
- When  $J'$  exits critical section, its priority is updated to the highest priority of some job that is blocked by  $J'$  (or to the nominal priority if no such job exists).

BF - ES

- 17 -

## Example

## REVIEW




BF - ES

- 18 -

## Priority Ceiling Protocol

**Theorem (Sha/Rajkumar/Lehoczky):** Under the Priority Ceiling Protocol, a job can be blocked for at most the duration of one critical section.



- Suppose that  $j$  is blocked by lower-priority jobs  $j_1, j_2$
- Let  $j_1$  enter the critical section first with semaphore  $S_a$ ; let  $C_1$  be the max ceiling among the semaphores it blocks.
- Then  $j_2$  enters its crit. section  $\Rightarrow$  priority( $j_2$ ) >  $C_1$
- Since  $j$  uses  $S_a \Rightarrow$  priority( $j$ )  $\leq C_1$   
 $\Rightarrow$  priority( $j$ ) < priority( $j_2$ )  $\checkmark$

BF - ES

- 19 -

## Priority Ceiling Protocol

The Priority Ceiling Protocol prevents deadlocks.

Deadlock  $\Rightarrow$   $\exists$  jobs  $j_0, \dots, j_{n-1}$  s.t.

$j_0$  is blocked by  $j_1$ ,  $j_1$  is blocked by  $j_2, \dots$   
 $j_{n-1}$  is blocked by  $j_0$ .

Let  $j_i$  be the last job to enter the critical section  $\Rightarrow$  priority( $j_i$ ) > max ceiling of semaphores held by  $j_{i+1}, \dots, j_n$

However, since  $j_i$  is blocked by  $j_{i+1}, \dots, j_n$ ,  
 priority( $j_i$ )  $\leq$  ceiling of blocked semaphores.  $\checkmark$

BF - ES

- 20 -

## Incorporating a-periodic tasks

- In real systems, not all tasks are periodic
  - Environmental events to be processed
  - Exceptions raised
  - Background tasks running whenever CPU time budget permits
- Thus, real systems tend to be a combination of
  - periodic and
  - a-periodic tasksand of
  - hard real-time and
  - soft real-time tasks.

BF - ES

- 21 -

## A-periodic and periodic tasks together (1)

- A-periodic and periodic tasks together
  - can be handled by **dynamic-priority schedulers** like EDF
- Problem:
  - Off-line guarantees can not be given without **assumptions on a-periodic tasks**.
  - If deadlines for a-periodic tasks are hard, a-periodic tasks need to be characterized by a **minimum interarrival time between consecutive instances**
    - ⇒ bounds on the a-periodic load
  - A-periodic tasks with **maximum arrival rate** may be modeled as **periodic tasks** with this rate
    - ⇒ **periodic scheduling**
  - A-periodic tasks with maximum arrival rate are called **sporadic tasks**.

BF - ES

- 22 -

## A-periodic and periodic tasks together (2)

- Other solutions for the case that **periodic tasks have hard deadlines, a-periodic tasks have soft deadlines.**
  - Simplest solution: **Background scheduling**
    - A-periodic tasks are only executed when no periodic task is ready
    - Guarantees for periodic tasks do not change
    - Only applicable when load is not too high
  - Other solutions:
    - Define **new periodic tasks**, a so-called **server**
    - A-periodic tasks are executed during “execution time” of server process
    - Independent scheduling strategies possible for periodic tasks and a-periodic tasks “inside the server”

BF - ES

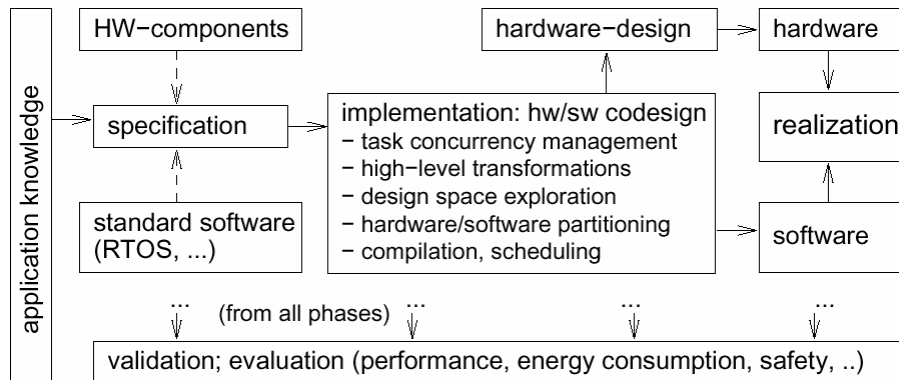
- 23 -

## REVIEW

BF - ES

- 24 -

## Overview



BF - ES

- 25 -

## Specification

Communication/ local computations	Shared memory	Message passing	
		Synchronous	Asynchronous
Communicating finite state machines	StateCharts, StateFlow		SDL, MSCs
Data flow model ⊂ Computational graphs			Kahn process networks, SDF  Petri nets
Von Neumann model	C, C++, Java	C, C++, Java with libraries CSP, ADA	
Discrete event (DE) model	VHDL, Simulink	Only experimental systems, e.g. distributed DE in Ptolemy	

BF - ES

- 26 -

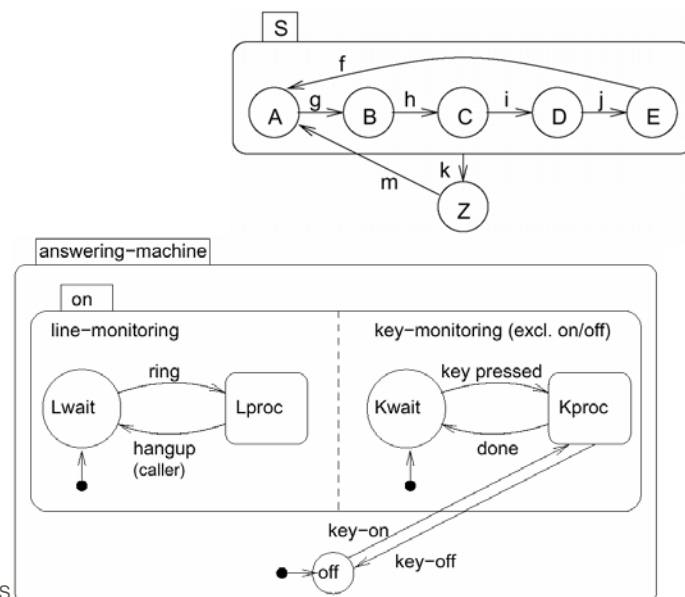
## StateCharts

- Why?
  - Concise models of complex systems:  
**StateCharts = FSMs + Hierarchy + Orthogonality + Broadcast communication**
  - Commercial tools (StateMate, StateFlow, ...)
- What?
  - Semantics
  - Virtual Prototyping (→ Matlab/Simulink/Stateflow)

BF - ES

- 27 -



## StateCharts: Hierarchy + Orthogonality



BF - ES

- 28 -

## StateCharts: Default-state mechanism



BF - ES

- 29 -


## StateCharts: Semantics (Statemate)

- Two stages
  - Preparation (for timeout events and scheduled actions)
  - Execution
- Preparation stage:
  - Fix scheduled actions that will be executed
  - Fix timeout events that will be generated
- Execution stage:
  - Determine the set of transitions to be taken based on internal and external events and on values of internal and external variables
  - Compute the next states and the reactions (evaluate right hand sides of assignments)
  - Transitions become effective, variables obtain new values.

BF - ES

- 30 -

## Hybrid Systems

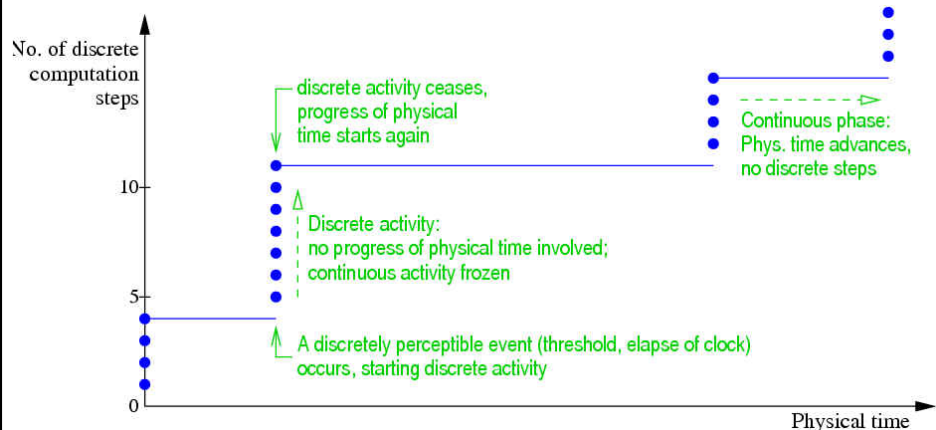


BF - ES

- 31 -

## Hybrid Systems: The super-step time model

- Two-dimensional time:



- Assumption: Computation time is negligible compared to dynamics of the environment.

BF - ES

- 32 -



## Petri Nets

- Why?
  - Modeling causal dependencies
  - Distributed systems
- What?
  - Reachability graph
  - Invariant generation
  - Deadlocks
- Advanced material:
  - Fairness
  - Nets with priorities
  - Predicate/transition nets


BF - ES

- 33 -

## Activated transitions

Transition  $t$  is „activated“ iff

$$(\forall p \in \bullet t : M(p) \geq W(p,t)) \wedge (\forall p \in t^\bullet : M(p) + W(t,p) \leq K(p))$$



Activated transitions can „take place“ or „fire“, but don't have to.  
The order in which activated transitions fire is not fixed (it is non-deterministic).

BF - ES

- 34 -

## Shorthand for changes of markings

Firing transition:

$$M'(p) = \begin{cases} M(p) - W(p,t), & \text{if } p \in \bullet t \setminus t^\bullet \\ M(p) + W(t,p), & \text{if } p \in t^\bullet \setminus \bullet t \\ M(p) - W(p,t) + W(t,p), & \text{if } p \in \bullet t \cap t^\bullet \\ M(p) & \text{otherwise} \end{cases}$$

Let

$$\underline{t}(p) = \begin{cases} -W(p,t) & \text{if } p \in \bullet t \setminus t^\bullet \\ +W(t,p) & \text{if } p \in t^\bullet \setminus \bullet t \\ -W(p,t) + W(t,p) & \text{if } p \in \bullet t \cap t^\bullet \\ 0 & \text{otherwise} \end{cases}$$

$$\Rightarrow \quad \forall p \in P: M'(p) = M(p) + \underline{t}(p)$$

$$\Rightarrow \quad M' = M + \underline{t} \quad \text{+ : vector add}$$

BF - ES

- 35 -

## Reachability graph

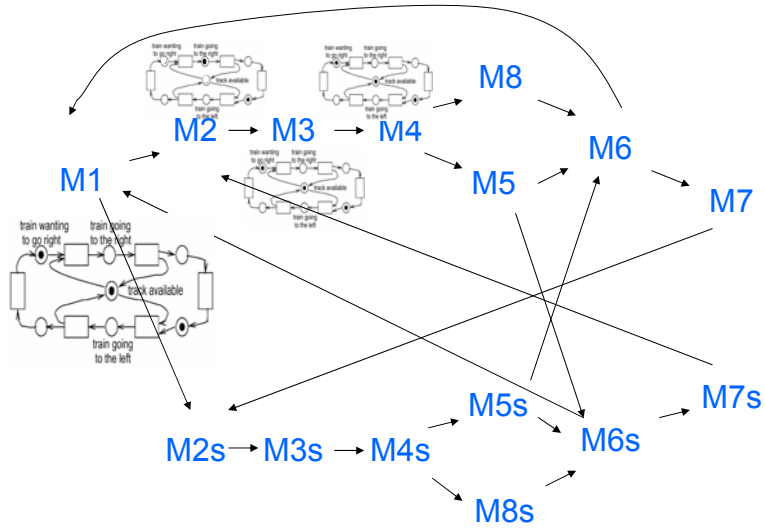
- $M [t > M' \text{ iff } M' = M + \underline{t}$
- $M [\varepsilon > M' \text{ iff } M = M'$
- $M [qt > M' \text{ iff } \exists M'' . M [q > M'' \text{ and } M'' [t > M'$
- $M [^* > M' \text{ iff } \exists q . M [q > M'$

- **Reachability set**  $R(M) = \{M' \mid M [^* > M'\}$
- **Reachability graph**  $RG(M)$ :  
nodes  $R(M)$ , edges  $\{(M, t, M') \mid M [t > M'\}$

BF - ES

- 36 -



## Reachable markings



BF - ES

- 37 -

## Unbounded Petri net



BF - ES

- 38 -

## Boundedness


- A P/T net is unbounded iff there exist two reachable markings  $M, M'$ , such that  $M \ll^* M'$  and  $M' > M$ .
- Every infinite sequence of markings  $(M_i)$  contains a weakly monotonically growing infinite subsequence  $(M'_j)$ , i.e., for  $j < k$ ,  $M'_j \leq M'_k$ .

BF - ES

- 39 -

## Algorithm for deciding boundedness

- Explore  $RG(M_0)$  depth-first:
  - If there exists a marking  $M'$  on the stack such that  $M' < M$ , stop with result UNBOUNDED;
- If entire graph explored, return BOUNDED.




BF - ES

- 40 -

## Invariant Generation

	$p_1$	$p_2$	$p_3$	$p_{-1}$	$p_{-2}$	$p_{-3}$	$p_0$
$t_1$	+1					-1	
$t_2$	-1	+1					-1
$t_3$		-1	+1				+1
$t_4$			-1	+1			
$t_5$				-1	+1		-1
$t_6$					-1	+1	+1
A $t_1$	+1					-1	
B $t_1+t_2$	+1				-1	-1	
C $t_2+t_3$		+1			-1		
D $t_3+t_4$			+1		-1		
D $t_4+t_5$				+1	-1	-1	

BF - ES



$$c_0 = 0, c_{-3} = 1$$

$$\Rightarrow c_{-2} = 1$$

$$c_{-1} = 1, c_1 = 1$$

$$c_2 = 1, c_3 = 1$$

$$c_0 = 1, c_3 = 0$$

$$\Rightarrow c_{-2} = 1$$

$$c_{-1} = 0, c_1 = 0$$

$$c_2 = 1, c_3 = 0$$

## Invariants & boundedness

### Theorem:


- If  $R$  is an invariant and  $p \in R$ , then  $p$  is bounded.
- If a net is covered by invariants then it is bounded.

BF - ES

- 42 -

## Deadlock

- A **dead marking (deadlock)** is a marking where no transition can fire.
- A Petri net is **deadlock-free** if no dead marking is reachable.




BF - ES

- 43 -

## Structural properties: deadlock-traps

- A place set  $S$  is a **(static) deadlock** if every transition that adds tokens to  $S$  also removes tokens from  $S$ .
- A place set  $S$  is a **trap** if every transition that removes tokens from  $S$  also adds tokens to  $S$ .
- A P/T net has the **deadlock-trap property**, if every (static) deadlock contains a trap that is sufficiently marked in  $M_0$ .
- Every homogeneous P/T net with non-blocking weights that has the deadlock-trap property is deadlock-free.



BF - ES

- 44 -

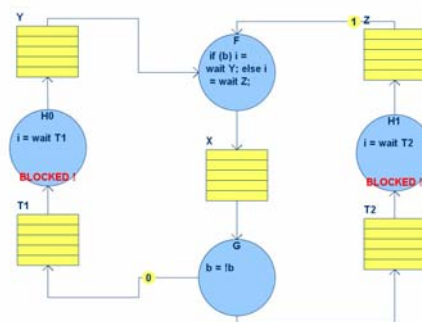
## Data flow models

- Why?
  - Many applications can be specified in the form of a set of communicating processes.
  - Communication exclusively through FIFOs
  - Describe local behavior + dependencies without worrying about global control
- What?
  - Kahn process networks
    - Park's runtime scheduling algorithm
  - Synchronous data flow (SDF)
    - Lee/Messerschmitt's static scheduling algorithm

BF - ES

- 45 -

## Kahn process networks



- Each node corresponds to one program/task;
- Communication is only via channels;
- Channels include FIFOs as large as needed;
- **Send operations are non-blocking, reads are blocking.**

BF - ES

- 46 -


## Kahn process networks are deterministic

- There is only one sender per channel.
- A process cannot check whether data is available before attempting a read.
- A process cannot wait for data for more than one port at a time.
- Therefore, the order of reads depends only on data, not on the arrival time.
- Therefore, Kahn process networks are deterministic (!), for a given input, the result will always be the same, regardless of the speed of the nodes.

BF - ES

- 47 -

## Scheduling may be impossible



BF - ES

- 48 -



## Parks' Scheduling Algorithm (1995)


- Set a capacity on each channel ← *FIFOs always bounded!*
- Block a write if the channel is full
- Repeat
  - Run until deadlock occurs
  - If there are no blocking writes → terminate
  - Among the channels that block writes, select the channel with least capacity and increase capacity until producer can fire.

BF - ES

- 49 -

## Synchronous data flow (SDF)

- Asynchronous message passing = tasks do not have to wait until output is accepted.
- Synchronous data flow = all tokens are consumed at the same time.

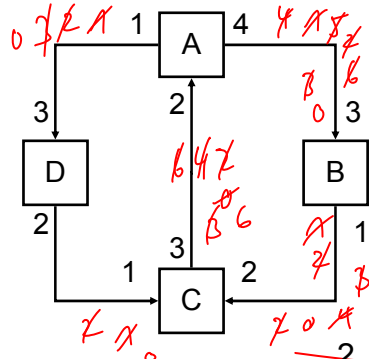


BF - ES

- 50 -

## SDF Scheduling Algorithm Lee/Messerschmitt 1987

1. Establish relative execution rates



$$\begin{pmatrix} AB & -4 & 3 \\ BC & -1 & 2 \\ DC & 1 & -2 \\ CA & 2 & -3 \\ AD & -1 & 3 \end{pmatrix} \vec{c} = 0$$

$$\begin{aligned} c_C &= 2c_D \\ c_B &= 4c_D \\ c_A &= 3c_D \\ c_D &= 1, c_A = 3, c_B = 4, c_C = 2 \end{aligned}$$

2. Determine periodic schedule

Schedule: ABABABBDCC

d(CA)=6


## Message Sequence Charts

- Why?
  - Modeling scenarios (instead of state-based behavior)
  - Parital-order semantics
  - ITU-T Standard Z.120
  - Integrated as *sequence diagrams* in UML
- What?
  - Synchronous vs. asynchronous concatenation
  - Regularity
- Advanced material:
  - Live sequence charts

## Message Sequence Charts

$Ch = (E, \leq, \lambda)$

- the elements of  $E_p$  are arranged along a life-line with the earlier elements appearing above the later elements.




BF - ES

- 53 -

## Language = Set of linearizations

- A **linearization** of a basic MSC is a sequence of actions  $\lambda(e_0), \lambda(e_1), \dots, \lambda(e_n)$  such that  $E = \{e_0, e_1, \dots, e_n\}$  and  $e_0 \leq e_1 \leq \dots \leq e_n$ .
- Each basic MSC  $Ch = (E, \leq, \lambda)$  defines a set of linearizations:  $lin(Ch) \subseteq \Sigma^*$




$lin(Ch) =$   
 $\{ U!I(req) \ I?U(req)$   
 $I!R(req) \ R?I(req)$   
 $R!I(grant) \ I?R(grant)$   
 $I!U(yes) \ U?I(yes) \}$

BF - ES

- 54 -

## Message Sequence Graphs



BF - ES

- 55 -

## Synchronous vs. asynchronous concatenation

- **edges** in an MSG represent chart **concatenation**:
- **Synchronous concatenation**  $Ch:Ch'$  means that *all* the events in  $Ch$  must finish before *any* event in  $Ch'$  can occur.
- **Asynchronous concatenation**  $Ch1 \circ Ch2$  is carried out at the level of life-lines.
- asynchronous concatenation of two charts is also a chart.
- synchronous concatenation of two charts **may not result in a chart.**
- asynchronous concatenation may lead to **non-regular languages**

BF - ES

- 56 -


## Communication-boundedness

Communication-boundedness is a sufficient condition for regularity.

- The **communication graph** of a basic MSC is a directed graph, where the nodes are the processes, edge  $p \rightarrow q$  if  $p!q(m)$  for some  $m$  in chart.
- MSC is **communication-bounded** iff communication graph consists of a single strongly-connected component (+ isolated nodes)
- MSG is **communication-bounded** iff communication graph of all loops is communication-bounded.

BF - ES


- 57 -



$p \rightarrow q$  not strongly connected  
 $\Rightarrow$  not communication-bounded



$p \leftrightarrow q$  ✓  
 $\Rightarrow$  communication-bounded



$p \leftrightarrow q$   $r \leftrightarrow s$  two components!  
 $\Rightarrow$  not communication-bounded

BF - ES

- 58 -

## VHDL

- Why?
  - Describing, simulating, synthesizing hardware
  - Standard in (European) industry
- What?
  - Entities, architectures
  - Multi-valued logic
  - Semantics
    - Transport delay model
- Advanced material:
  - IEEE 1164
  - Parameterized hardware

BF - ES

- 59 -

```
entity full_adder is
  port(a, b, carry_in: in Bit; -- input ports
        sum,carry_out: out Bit); --output ports
end full_adder;

architecture behavior of full_adder is
begin
  sum    <= (a xor b) xor carry_in after 10 Ns;
  carry_out <= (a and b) or (a and carry_in) or
              (b and carry_in)    after 10 Ns;
end behavior;

architecture structure of full_adder is
  component half_adder
    port (in1,in2:in Bit; carry:out Bit; sum:out Bit);
  end component;
  component or_gate
    port (in1, in2:in Bit; o:out Bit);
  end component;
  signal x, y, z: Bit; -- local signals
begin -- port map section
  i1: half_adder port map (a, b, x, y);
  i2: half_adder port map (y, carry_in, z, sum);
  i3: or_gate   port map (x, z, carry_out);
end structure;
```

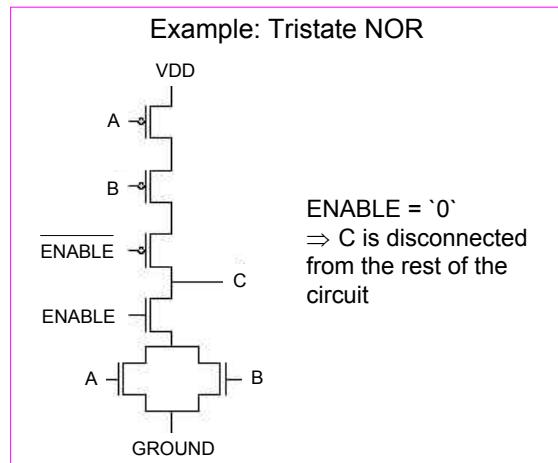
BF - ES

- Architectures describe implementations of entities.
- Architectures and their components can define a hierarchy of arbitrary depth.

- 60 -

## Multi-valued logic

- Many subcircuits can be effectively disconnected from the rest of the circuit (they provide „high impedance“ values to the rest of the circuit).
- Example: subcircuits with tri-state outputs.

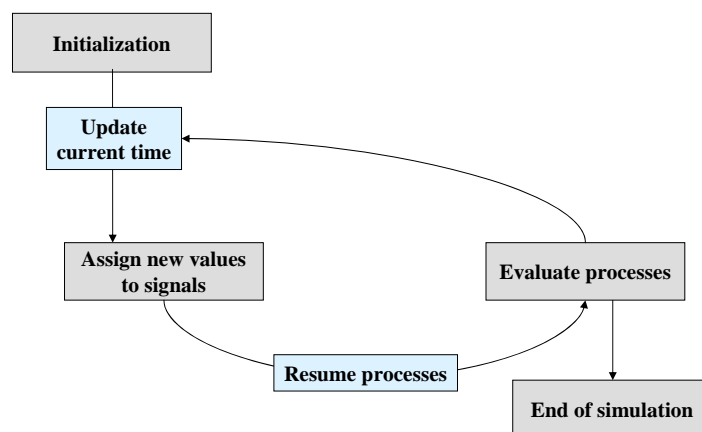


☞ We introduce signal value 'Z', meaning „high impedance“

BF - ES

- 61 -


## Semantics



BF - ES



- 62 -

## Inertial vs. transport delay model



- INERTIAL is the default  
Output <= NOT input AFTER 10 ns;

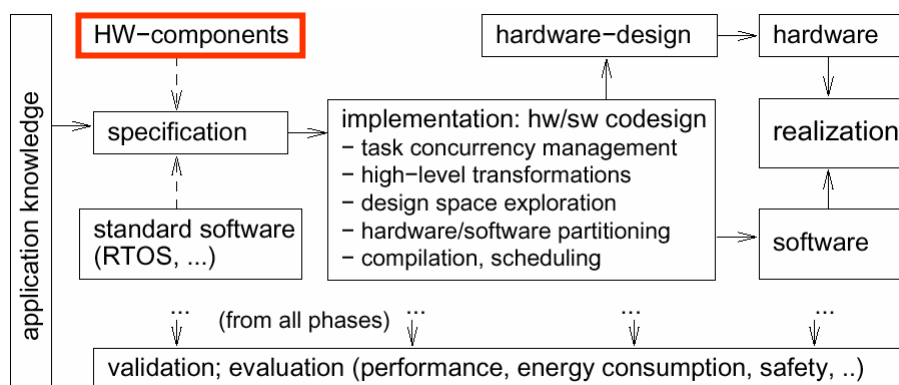
-- TRANSPORT must be specified  
Output <= TRANSPORT NOT input AFTER 10 ns;



BF - ES

- 63 -

## Overview of embedded systems design



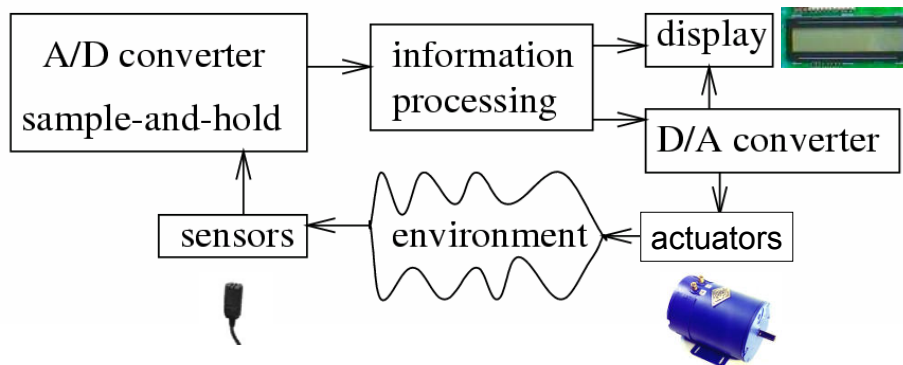
BF - ES

- 64 -



## Embedded System Hardware

- Embedded system hardware is frequently used in a loop („*hardware in a loop*“):



BF - ES

- 65 -


## Sensors, A/D + D/A converters

- Why?
  - Embedded systems interact with physical environment.
- What?
  - Sample & hold circuits
  - A/D converters
  - D/A converters
- Advanced material:
  - Image sensors, ...

BF - ES

- 66 -

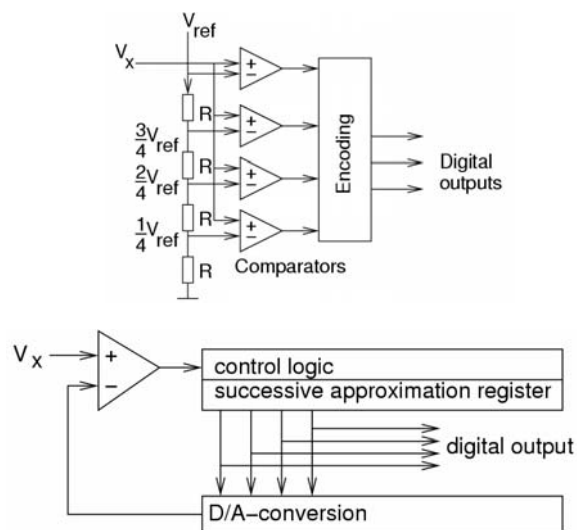
## Sample and Hold



BF - ES

- 67 -

## Flash conversion vs. successive approximation



BF - ES

- 68 -

## Information Processing

- Why?
  - Embedded systems must be efficient
  - Embedded processors need not be instruction set compatible with standard PCs
- What?
  - Power/energy efficiency
  - Code-size efficiency
  - Runtime efficiency
- Advanced material:
  - Reconfigurable logic, Multimedia processors, scratch pad memory, ...

BF - ES

- 69 -

## Dynamic voltage scaling (DVS)

Power consumption of CMOS circuits (ignoring leakage):

$$P = \alpha C_L V_{dd}^2 f \text{ with}$$

$\alpha$  : switching activity

$C_L$  : load capacitance

$V_{dd}$  : supply voltage

$f$  : clock frequency

Delay for CMOS circuits:

$$\tau = k C_L \frac{V_{dd}}{(V_{dd} - V_t)^2} \text{ with}$$

$V_t$  : threshold voltage

( $V_t < \text{than } V_{dd}$ )


☞ Decreasing  $V_{dd}$  reduces  $P$  quadratically, while the run-time of algorithms is only linearly increased  
 $E = P \times t$  decreases linearly (ignoring the effects of the memory system and  $V_t$ )

BF - ES

- 70 -

## Code-size efficiency

- **CISC machines:** RISC machines designed for run-time-, not for code-size-efficiency
- **Compression techniques:** key idea



BF - ES

- 71 -

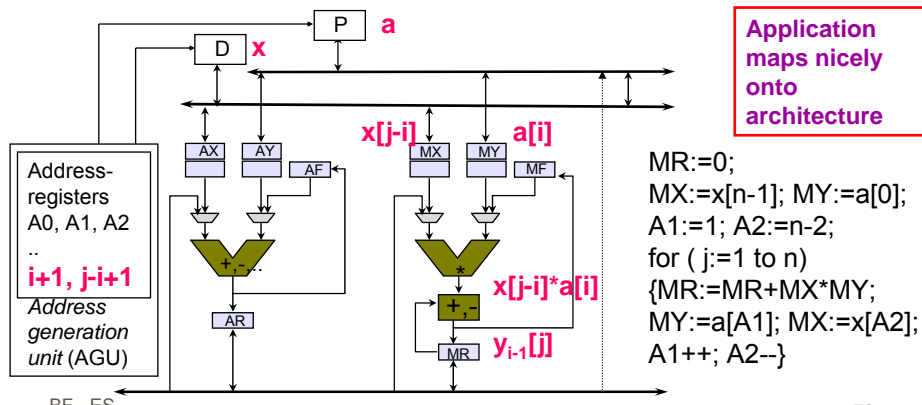
## Run-time efficiency

### - Domain-oriented architectures -

**Application:**  $y[j] = \sum_{i=0}^{n-1} x[j-i] * a[i]$

$\forall i: 0 \leq i \leq n-1: y_i[j] = y_{i-1}[j] + x[j-i] * a[i]$

**Architecture:** Example: Data path ADSP210x



BF - ES

- 72 -

## Real-time communication

- Why?
  - Modular system development, support and evolution
  - Single network vs. wiring harness
- What?
  - Bus-master approach
  - TDMA, CSMA
  - Collision handling
- Advanced material:
  - FlexRay

BF - ES

- 73 -

## TDMA - Time Division Multipl. Access

Operational principle:

- Progress of time is divided into TDMA rounds, within which the individual nodes have private time slots with different phase delay to start of the round; the slots are non-overlapping

Problems:

- Private slots waste bandwidth
- Need for global clock synchronisation
- Number of nodes and their worst-case message lengths need to be fixed a priori
- This leads to either designs using huge safety margins or to complex interference between node performance and TDMA setup (lacking separation of concerns between computation and communication)

BF - ES

- 74 -

## Carrier Sense Multiple Access (CSMA)

Operational principle:

- If communication medium is idle then send a message (node decides on its own — no global authority)

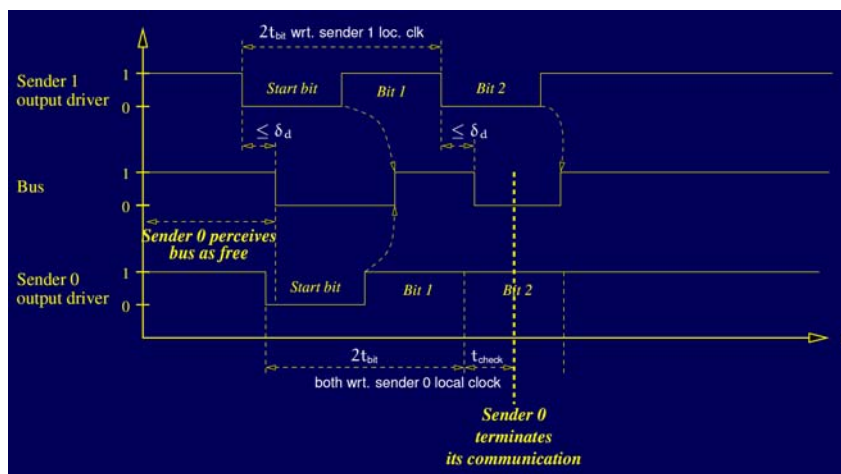
Problems:

- Multiple nodes may start almost synchronously, leading to collision on the medium
- Message may be crippled
- Message may be overwritten and thus not delivered
- If message delivery is vital (std. in ES) then collision has to be resolved
  - Collision detection or collision avoidance, arbiting the bus such that at least one of the colliding messages is delivered uncrippled

BF - ES

- 75 -

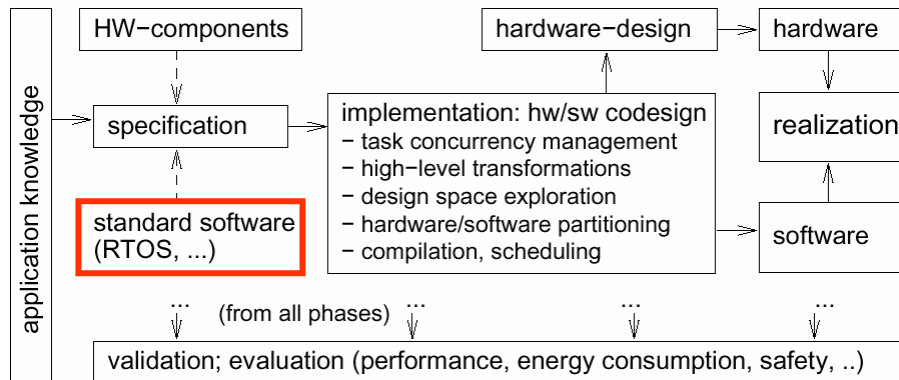
## Collision handling



BF - ES

- 76 -

## Overview of embedded systems design



BF - ES

- 77 -

## Scheduling

- Why?
  - Scheduling key issue in implementing RT-systems
  - Different algorithms with different assumptions and cost
- What?
  - Aperiodic scheduling
  - Periodic scheduling
  - Scheduling with resource constraints
- Advanced material:
  - Priority ceiling protocol

BF - ES

- 78 -

## Aperiodic scheduling: EDF – Earliest Deadline First

- EDF: At every instant execute the task with the earliest absolute deadline among all the ready tasks.
- **Theorem (Horn '74):**  
Given a set of  $n$  independent task with arbitrary arrival times, any algorithm that at every instant executes the task with the earliest absolute deadline among all the ready tasks is optimal with respect to minimizing the maximum lateness.

BF - ES

- 79 -

## Aperiodic scheduling: Non-preemptive version

- **Theorem (Jeffay et al. '91):** EDF is an optimal *non-idle* scheduling algorithm also in a *non-preemptive* task model.
- When idle schedules are allowed: problem is NP-hard.
- Possible approaches:
  - Heuristics
  - Bratley's algorithm: branch-and-bound

BF - ES

- 80 -



## Aperiodic scheduling: : Scheduling with precedence constraints

- Non-preemptive scheduling with non-synchronous arrival times, deadlines and precedence constraints is NP-hard.
- Restrictions:
  - Consider synchronous arrival times (all tasks arrive at 0)
  - Allow preemption.
- **Theorem (Lawler 73):**  
LDF (Latest Deadline First) is optimal wrt. maximum lateness.

BF - ES

- 81 -

## Periodic scheduling: EDF

- **Theorem:** A set of periodic tasks  $\tau_1, \dots, \tau_n$  with  $D_i = T_i$  is schedulable with EDF iff  $U \leq 1$ .
- EDF is applicable to both periodic and a-periodic tasks.
- If there are only periodic tasks, priority-based schemes like “rate monotonic scheduling (RM)” (see later) are often preferred, since
  - They are simpler due to fixed priorities  
 $\Rightarrow$  use in “standard OS” possible
  - sorting wrt. to deadlines **at run time** is not needed

BF - ES

- 82 -


## Rate monotonic scheduling (RM)

- Rate monotonic scheduling (RM) (Liu, Layland '73):
  - Assign **fixed priorities** to tasks  $\tau_i$ :
    - $\text{priority}(\tau_i) = 1/T_i$
    - I.e., priority **reflects release rate**
  - **Always execute ready task with highest priority**
  - Preemptive: currently executing task is preempted by newly arrived task with shorter period.
- **Theorem (Liu, Layland, 1973):**  
RM is **optimal among all fixed-priority** scheduling algorithms.

BF - ES

- 83 -

## The priority inversion problem



- Blocking time equal to length of critical section + computation time of  $J_2$ .
- **Unbounded time of priority inversion**, if  $J_3$  is interrupted by tasks with priority between  $J_1$  and  $J_3$  during its critical region.

BF - ES

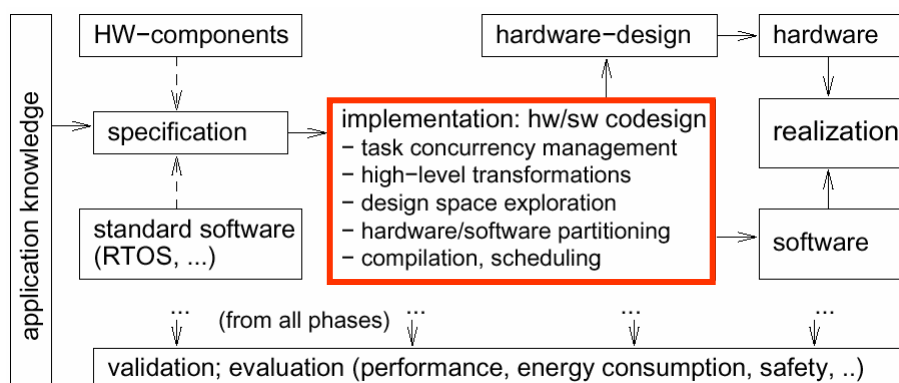
- 84 -

## Coping with priority inversion: The priority inheritance protocol

### Idea of **priority inheritance protocol**:

- If a task  $J_h$  blocks, since another task  $J_l$  with lower priority owns the requested resource, then  $J_l$  inherits the priority of  $J_h$ .
- When  $J_l$  releases the resource, the priority inheritance from  $J_h$  is undone.
- Rule: Tasks always inherit the highest priority of tasks blocked by it.

## Coming up...



## Thursday

- Midterm exam: December 18, 2008, HS I, Math building, 16:15 - 17:45
- Open book: bring any handwritten or printed notes, or any books you like.
- Please bring your ID.