

Verification of Real Time Systems - CS5270

3rd lecture

Hugh Anderson

National University of Singapore
School of Computing

January, 2007



A warning...



Outline

- 1 Administration
 - Assignment 1
 - Introduction to Uppaal
- 2 Scheduling
 - Scheduling concepts
 - Critical sections and Semaphores
- 3 Scheduling algorithms
 - RMS - Rate Monotonic Scheduling
 - Schedulability
 - EDF Earliest Deadline First



Assignment 1

Assignment number 1 is out

- Seven questions
- Some reading may be required?
- Hand in Feb 18



Uppaal

The website:



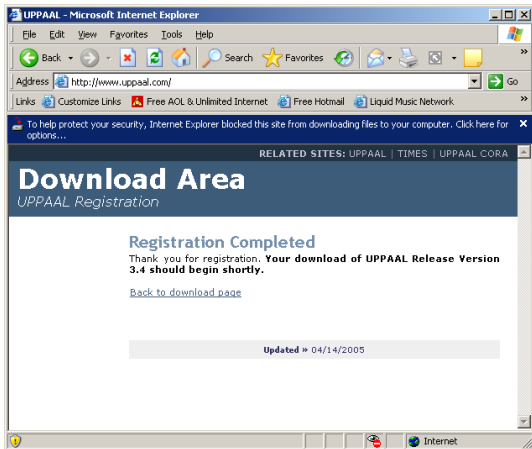
Uppaal

The license:



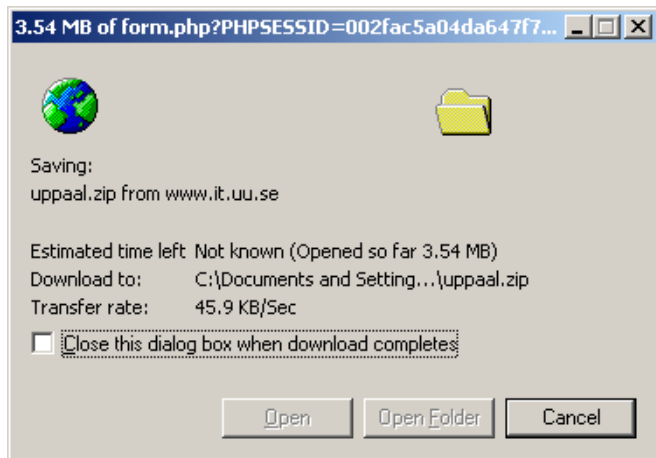
Uppaal

Complete registration:



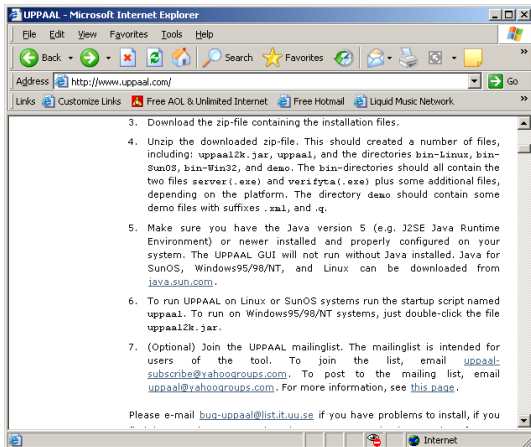
Uppaal

Download:



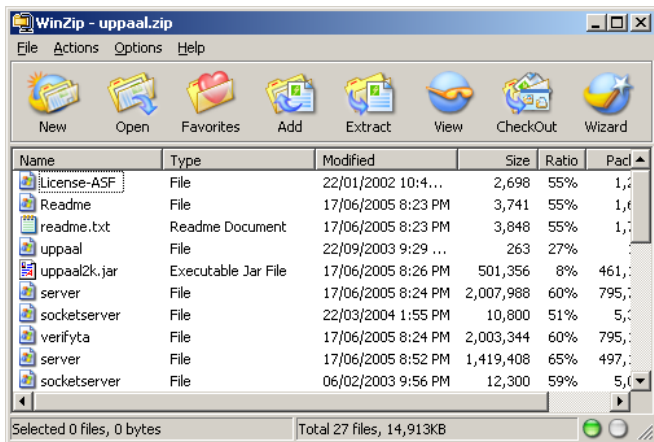
Uppaal

Instructions:



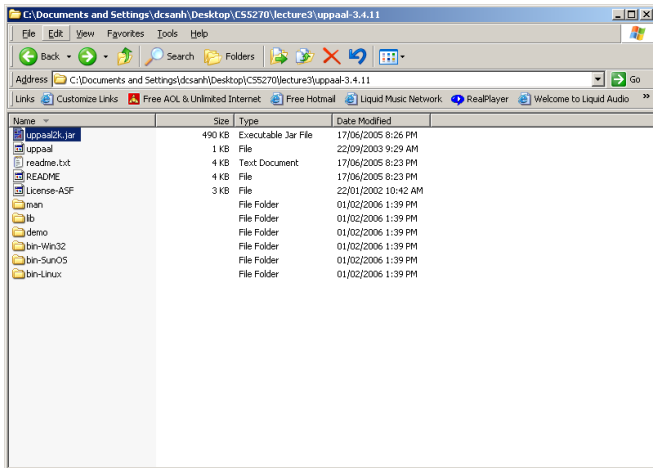
Uppaal

Extract/unzip:



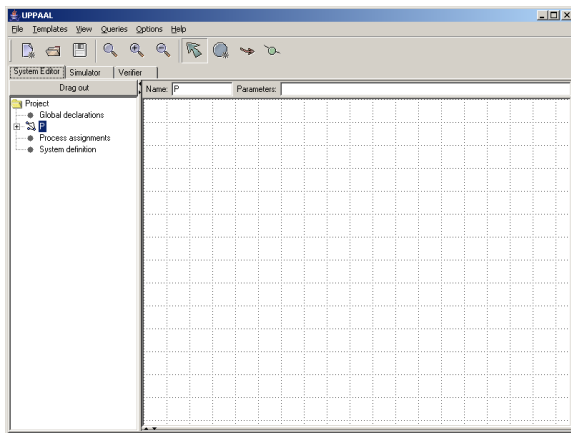
Uppaal

Click on jar file:



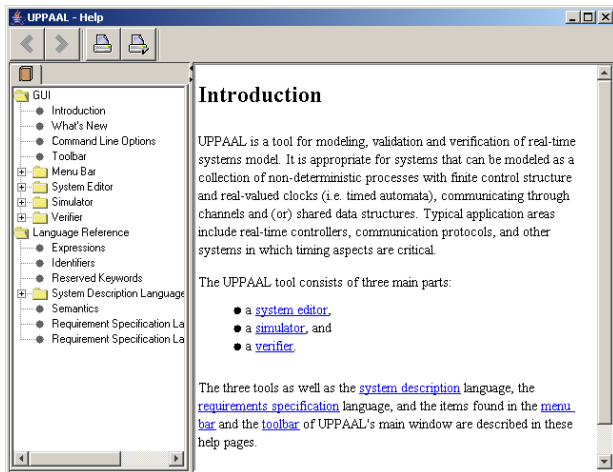
Uppaal

The application:



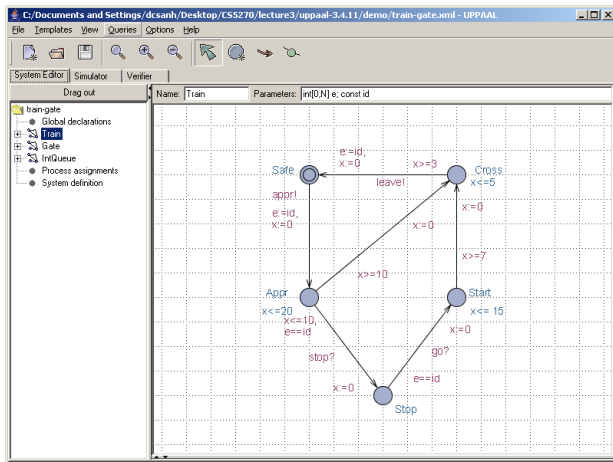
Uppaal

Help:



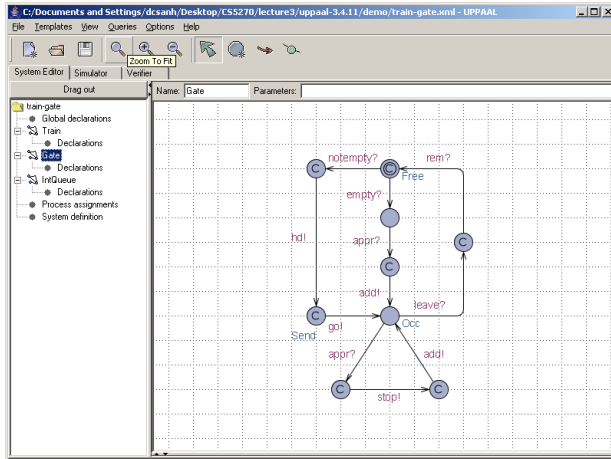
Uppaal

Load up a demo:



Uppaal

Look at TTS:



Uppaal

Simulation:

The screenshot displays the Uppaal simulation environment. The main window shows a state transition diagram for a system with four trains (Train1, Train2, Train3, Train4) and a Gate. The diagram includes states: Safe, Appr1, Appr, Start, and Stop. Transitions are labeled with guards and actions, such as $e1=4, x=0$ for the transition from Safe to Appr1, and $x=3$ for the transition from Appr1 to Cross. The diagram also shows a loop from Stop to Appr with guard $x=0$ and action $stop?$, and a transition from Appr to Stop with guard $x=0$ and action $stop?$. The diagram is titled "Train4" and "Gate".

On the left side, the "Drag out" panel shows "Enabled Transitions" with the expression $(Queue.4.empty \wedge Gate.2.empty?)$ and "Simulation Trace" with the expression $(Safe Safe Safe Safe Free Start)$. Below these are buttons for "Next", "Reset", "Prev", "Next", "Replay", "Open", "Save", and "Random", along with a speed slider from "Slow" to "Fast".

The "Variables" panel lists the following variables and their values:

```

e1 = 0
Queue.list()
Queue.list()
Queue.list()
Queue.list()
Queue.list()
Queue.len = 0
Queue.i = 0
Train1.x = 
Train2.x = 
Train3.x = 
Train4.x = 
Train2.x = 
Train3.x = 
Train4.x = 
Train3.x = 
Train4.x = 
Train4.x =
    
```

At the bottom, a control panel shows the status of the trains and the gate: Train1 (Safe), Train2 (Safe), Train3 (Safe), Train4 (Safe), Gate (Free), and Queue (Start).

Uppaal

Simulation:

The screenshot displays the Uppaal simulation environment. The main window shows a Petri net diagram with nodes: Safe, Appr, Cross, Start, and Stop. Transitions are labeled with events like 'leave?', 'go?', and 'stop?'. Place invariants are shown as $x=0$ or $x \leq 5$. The interface includes a menu bar (File, Templates, View, Queries, Options, Help), a toolbar, and three main panes: System Editor, Simulator, and Verifier.

System Editor: Shows 'Enabled Transitions' with the selected transition `[Gate.5.stop!, Train3.2.stop?]`. Below are 'Next' and 'Reset' buttons.

Simulation Trace: Lists the sequence of events:


```

    [Appr.3.add!, Gate.3.appr!, ...]
    [Gate.5.stop!, Train4.2.stop?]
    [Appr!, Safe, Safe, Stop, -, Start]
    [Gate.10.add!, Queue.5.add?]
    [Appr!, Safe, Safe, Stop, Occ, Start]
    [Train3.4.appr!, Gate.3.appr?]
    [Appr, Safe, Appr, Stop, -, Start]
    
```

 The current trace file is `Trace File:` with navigation buttons (Prev, Next, Replay, Open, Save, Random) and a speed slider from Slow to Fast.

Simulator: Shows the Petri net diagram with a 'Stop' place containing 4 tokens. A 'Gate' window is open below the diagram.

Verifiers: Shows a list of variables for Train1 through Train4 and Gate, with values such as `Train1.x = 0`, `Train3.x = 0`, and `Gate.x = 0`.

Uppaal

Simulation:

The screenshot displays the Uppaal simulator interface for a train gate system. The main window shows a Petri net model with nodes: Safe, Cross, Appr, Start, and Stop. Transitions are labeled with guard conditions and actions, such as $el=4, x=0$ for the 'leave' transition and $x=3$ for the 'level' transition. The 'Variables' pane shows the state of the system, including an array `Queue.list` and counters `Queue.i` and `Train1.x`. The 'Simulation Trace' pane shows a sequence of events, with the current event being `Stop, Safe, Stop, Stop, Free, Shutdown`. The 'Gate' section at the bottom illustrates the state transitions of the gate, showing a sequence of states: Safe, a counter, Free, and Shutdown, with transitions labeled 'leave' and 'rem'.

Uppaal

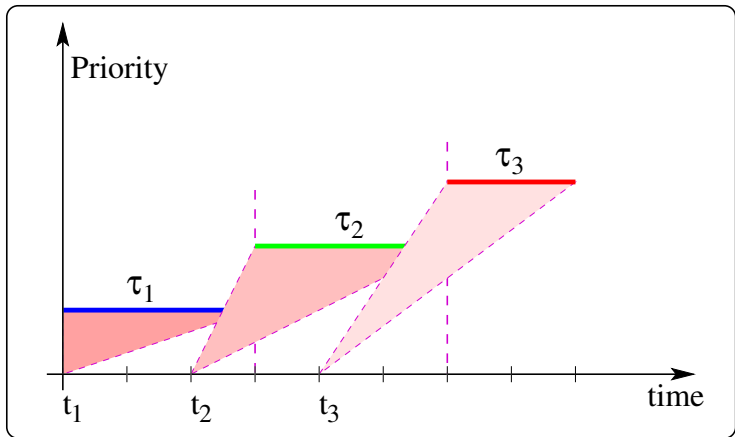
Verification:

The screenshot shows the Uppaal software interface with the following components:

- File Explorer:** Shows the file path `C:/Documents and Settings/dcsanh/Desktop/CS5270/lecture3/uppaal-3.4.11/demo/train-gate.xml - UPPAAL`.
- Menu Bar:** File, Templates, View, Queries, Options, Help.
- Toolbar:** Includes icons for file operations, search, and simulation.
- System Editor | Simulator | Verifier:** The active tab is the Verifier.
- Overview:**
 - Process list:
 - P0: (empty)
 - P1: `E<> Gate.0cc` (green status)
 - P2: `E<> Train1.Cross` (grey status)
 - P3: `E<> Train2.Cross` (green status)
 - Selected: `E<> Train1.Cross and Train2.Stop` (green status)
 - Buttons: Model Check, Insert, Remove, Comments.
- Query:**
 - Query text: `E<> Train1.Cross and Train2.Stop`
 - Comment: `Train 1 can be crossing bridge while Train 2 is waiting to cross.`
- Status:**
 - Established direct connection to local server.
 - `E<> Gate.0cc` (green): Property is satisfied.
 - `E<> Train2.Cross` (green): Property is satisfied.
 - `E<> Train1.Cross and Train2.Stop` (green): Property is satisfied.

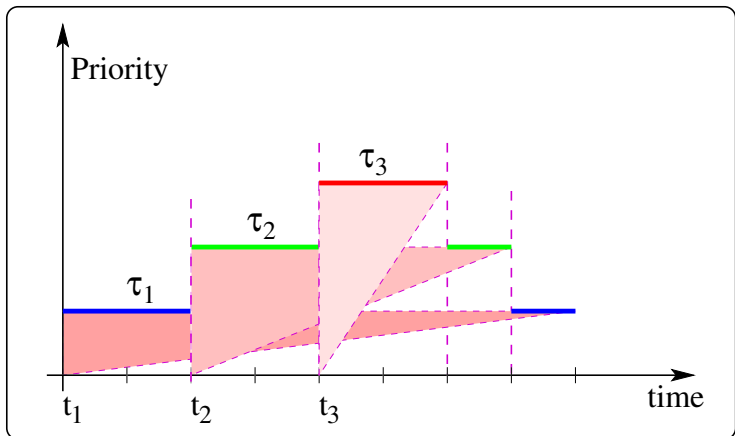
Non-preemptive scheduling

Tasks are delayed until other tasks complete:



Preemptive scheduling

Tasks preempt lower priority tasks:



Scheduling terms

Definitions:

Feasible: a schedule is termed feasible if all tasks can be completed within the constraints specified

Schedulable: a task set is schedulable if a particular scheduling algorithm produces a feasible schedule



Scheduling terms

Constraints found in various areas:

Timing (deadlines for tasks)

Precedence (which task comes first)

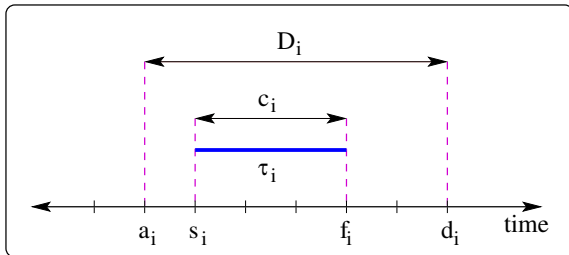
Resource (shared access)

Hard/Soft constraints



Scheduling terms

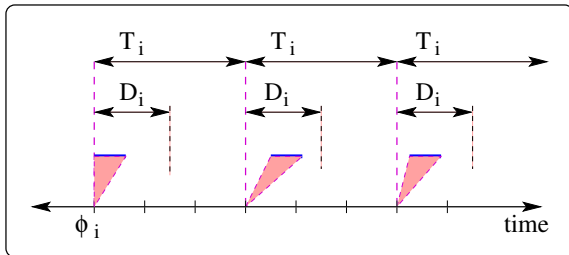
Deadlines:



- If a task t_i needs to finish before some time d_i , then this is called a **deadline**. A **relative deadline** D_i for the task is $D_i = d_i - a_i$.
- Tasks run for time c_j , and must complete before a deadline

Scheduling terms

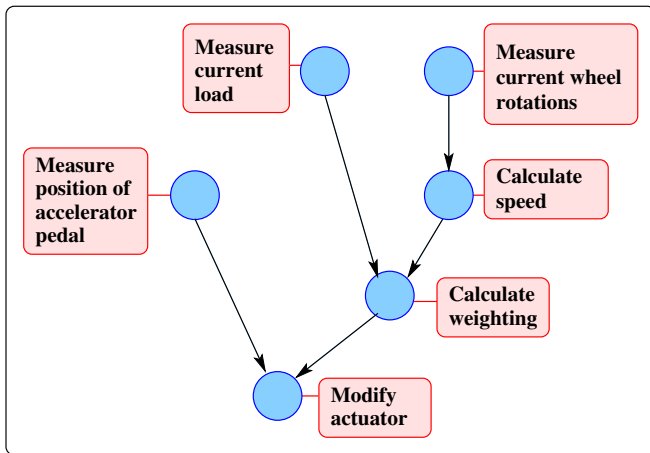
Periodic tasks:



- A **periodic** task is one that is regularly activated at a constant rate.
- Its **period** is T_i , and the time of first activation (its **phase**) is ϕ_i .

Scheduling terms

Precedence between tasks - visualize as a graph:



Scheduling terms

Resource access:

A *resource constraint*, may be some variable or device or some other structure in the system. Resources only become *critical resource constraints* when they are shared with other tasks.

- An *exclusive resource* is one which may require exclusion of all other tasks when the resource is accessed. This is called *mutual exclusion* (OS normally provide mechanisms to assist tasks to provide mutually exclusive access to a resource). The code which requires this mutually exclusive access is termed a *critical section* (CS).
- The most common mechanism for this purpose is called the *semaphore*, where a semaphore variable s_i is used to control access to an associated CS_i .

Critical section

A critical section is:

- A piece of code belonging to task executed under mutual exclusion constraints.
- Mutual exclusion is enforced by **semaphores**.
 - **wait(s)**
 - Blocked if $s = 0$.
 - **signal(s)**
 - s is set to 1 when **signal(s)** executes.



Critical sections

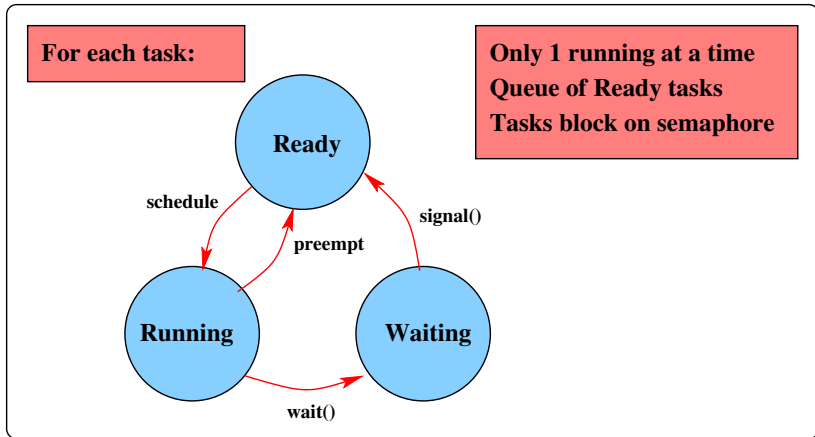
CS and blocking:

- A task waiting for an exclusive resource is **blocked** on that resource.
- Tasks blocked on the same resource are kept in a **wait queue** associated with the **semaphore** protecting the resource.
- A task in the running state executing **wait(s)** on a locked semaphore ($s = 0$) enters the **waiting** state.
- When a task currently using the resource executes **signal(s)**, the semaphore is released.
- When a task leaves its waiting state (because the semaphore has been released) it goes into the **ready** state:



Semaphores and blocking

In an OS, tasks block when waiting for a resource:



The scheduling problem

The general scheduling problem is NP-complete:

- There is a non-deterministic **Turing Machine** TM and a **polynomial** in one variable $p(n)$ such that for each problem instance of size n , TM determines if there exists a schedule and if so outputs one in at most $p(n)$ steps. Any non-deterministic polynomial time problem can be transformed in deterministic polynomial time to the general scheduling problem, and only **exponential** time deterministic algorithms are known.

Hence we must find imperfect but efficient solutions to scheduling problems. A great variety of algorithms exist, with various assumptions, and with different complexities.



Assumptions for RMS

In RMS:

- assume a set of tasks $\{\tau_1, \dots, \tau_m\}$ with periods T_1, \dots, T_m , $\phi_i = 0$ and $D_i = T_i$ for each task.
 - We allow preemption,
 - there is only a single processor, and
 - we have no precedence constraints.



RMS

The RMS algorithm:

- Assign a **static priority** to the tasks according to their periods.
- Priority of a task does not change during execution.
- Tasks with **shorter periods** have **higher priorities**.
- Preemption policy:
 - If T_i is executing and T_j arrives which has higher priority (shorter period), then preempt T_i and start executing T_j .



Assumptions RMS

From the Liu article:

(A1) The requests for all tasks for which hard deadlines exist are periodic, with constant interval between requests.

(A2) Deadlines consist of run-ability constraints only - i.e. each task must be completed before the next request for it occurs.

(A3) The tasks are independent in that requests for a certain task do not depend on the initiation or the completion of requests for other tasks.

(A4) Run-time for each task is constant for that task and does not vary with time.

(A5) Any nonperiodic tasks in the system are special; they are initialization or failure-recovery routines; they displace periodic tasks while they themselves are being run, and do not themselves have hard, critical deadlines.

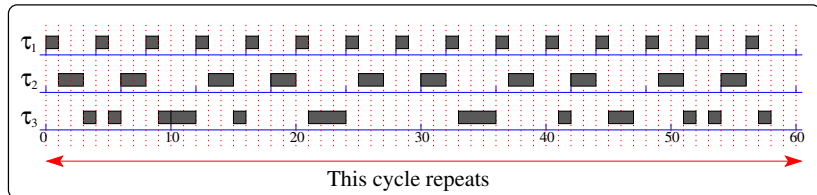


RMS

Given this task set:

	τ_1	τ_2	τ_3
C_i	1	2	3
T_i	4	6	10

An RMS schedule is:



RMS

Properties of RMS:

- RMS is **optimal** (Given the previous constraints)
 - If a set of periodic tasks (satisfying the assumptions set out previously) is not schedulable under RMS then no static priority algorithm can schedule this set of tasks.
- RMS requires very little run time processing.



Schedulability terms

Definition of PUF, the Processor Utilization Factor:

The **processor utilization factor** U is the fraction of processor time spent in the task set:

$$U = \sum_{i=1}^m \frac{C_i}{T_i}$$

If this factor is *greater* than 1 then of course, the task set can not be scheduled. However if $U \leq 1$ then it is possible that it may be RMS-schedulable. If a particular set of tasks has a **feasible** RMS schedule, and any increase of the runtime of any task would render the particular set infeasible, then the processor is said to be **fully utilized**.

Schedulability terms

Example of PUF calculation:

	τ_1	τ_2	τ_3
C_i	1	2	3
T_i	4	6	10

the processor utilization factor is $U = \sum_{i=1}^m \frac{C_i}{T_i} = 0.833$.



Schedulability

The least upper bound of processor utilization:

The **least upper bound** U_{lub} is the minimum of the U over all sets of tasks that fully utilize the processor.

- If $U \leq U_{lub}$, then the set of tasks is guaranteed to be schedulable.
- Table gives a sufficient value for U_{lub} for different numbers of tasks for RMS ($U_{lub} = m(2^{\frac{1}{m}} - 1)$), but note that it may be possible to schedule a task set even if the criterion fails.

m	1	2	3	4	5	6	∞
U_{lub}	1.000	0.828	0.780	0.757	0.743	0.735	0.690

Schedulability

Using our example:

	τ_1	τ_2	τ_3
C_i	1	2	3
T_i	4	6	10

the processor utilization factor is $U = \sum_{i=1}^m \frac{C_i}{T_i} = 0.833$.

- The least upper bound for rate monotonic scheduling for 3 tasks is given in the table as $U_{\text{lub}} = 0.780$, and since $U_{\text{lub}} < U$, we cannot guarantee that this task set is schedulable..



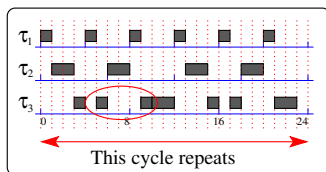
Schedulability

Another example:

	τ_1	τ_2	τ_3
C_i	1	2	3
T_i	4	6	8

the processor utilization factor is $U = \sum_{i=1}^m \frac{C_i}{T_i} = 0.95833$.

- With this set of tasks, we have that task τ_3 fails to complete within its period (8), task set is not schedulable using RMS.



Earliest Deadline First

The policy:

- Tasks with **earlier deadlines** will have **higher priorities**.
- Applies to both periodic and aperiodic tasks.
- EDF is optimal for dynamic priority algorithms.
- A set of periodic tasks is schedulable with EDF iff the utilization factor is not greater than 1.



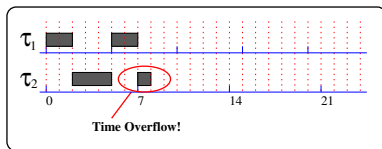
Earliest Deadline First

RMS fails:

	τ_1	τ_2
C_i	2	4
T_i	5	7

From $U = 0.4 + 0.57 = 0.97 \leq 1$ we know that it is guaranteed to be schedulable under EDF, and might be schedulable under RMS.

- It is not RMS schedulable:



Earliest Deadline First

EDF can guarantee deadlines in the system at higher loading:

	τ_1	τ_2
C_i	2	4
T_i	5	7

From $U = 0.4 + 0.57 = 0.97 \leq 1$ we know that it is guaranteed to be schedulable under EDF.

- EDF scheduling succeeds:

