



Rechnerstrukturen 2 - Übung 11

SCHEDULING

Today's Goal...

- **Closer look at these topics:**
 - Simulation and Analysis
 - Modelling
 - TDMA Scheduling
 - Round Robin Scheduling
 - Rate Monotonic Rate Scheduling
 - Static Priority Preemptive Scheduling



Simulation vs. Analysis


The problem with simulation

- Only covers one particular execution path
- Hard to simulate corner cases
- Does not scale with problem size


Analysis guarantees correctness

- Utilizes abstract models to describe system properties
- Ongoing research (SymTA/S tool developed at IDA)

State Space



Simulation



Single path

Analysis




Full coverage

Modelling - Event Streams


Task activation can be time triggered or event triggered

- In both ways modeled by event streams
- May trigger successor tasks by generating new events




Event streams

- Communication between tasks
- Properties depending on task behavior → multiple traces possible



Example traces at event stream:




Modelling - Event Streams - Characterization

Characterization of event streams

- Instead of investigating *all* individual event traces and their timing, formal schedulability analysis performs *one* analysis on the **event bounds**

All activating events within a time window of given size Δt

- $\eta^+(\Delta t)$: Maximum number of events in window Δt
- $\eta^-(\Delta t)$: Minimum number of events in window Δt




Modelling - Time Domain to Time Interval Domain Transformation (1/2)


Arrival Curves in time interval domain (Δt)

- $\eta^+(\Delta t)$: Maximum number of events in window Δt
- $\eta^-(\Delta t)$: Minimum number of events in window Δt

Have to be done for every possible event trace!



Δt	Min. #Events	Max. #Events
1P	0	3




Modelling - Time Domain to Time Interval Domain Transformation (2/2)


Arrival Curves in time interval domain (Δt)

- $\eta^+(\Delta t)$: Maximum number of events in window Δt
- $\eta^-(\Delta t)$: Minimum number of events in window Δt

Have to be done for every possible event trace!



Δt	Min. #Events	Max. #Events
1P	0	3
2P	1	4



Modelling - Event Models

Abstraction in time domain (t)


- Describe all traces on one event stream through a limited parameter set

Perform formal analysis on this parameter set: (P, J, d_{min})


P...Period

J...Jitter


d_{min} ...minimum distance between two consecutive events



Periodic events



Periodic events with jitter



Periodic events with jitter and burst

Allows conservative transformation from the actual event timing



- Only one analysis necessary instead of checking all event correlations

Modelling - Event Model: Periodic with Jitter ($\Delta t=1P$)

Event model: $(2,3,0) \leftarrow$ (Periode, Jitter, min. Distanz)

Analyze all possible event traces with $\Delta t=1P$

- Due to abstraction $(2,3,0)$, we know the best and worst case streams





Modelling - Event Model: Periodic with Jitter ($\Delta t=2P$)

Event model: (2,3,0)

Analyze all possible event streams with $\Delta t=2P$

- Due to abstraction (2,3,0), we know the best and worst case streams



Reminder: Scheduling

Task

- Set of instructions
- (Sub)program

Scheduling


- Temporal partitioning of processing or communication resources
- Sequenced task execution

Scheduler

- Assigns tasks to available (usually limited) resources

Preemption

- Temporary interruption of a task
- No cooperation required by the interrupted task
- Interrupted task is resumed eventually



Scheduling Analysis

What guarantees can be made given a certain environment?


- Worst-case
 - E.g. highest resource usage/bus congestion a certain task experiences

By Response time analysis

- Time from task activation to end of execution
- $t_{\text{resp}} \geq t_{\text{exec}}$

By Gantt-Charts


- Illustrate schedules
- Start and finish times
- Dependencies



TDMA Scheduling

Static time-driven scheduling

- Each task assigned to a specific time slot in reoccurring TDMA round
- Good analytical properties
- May lead to inefficient resource utilization



TDMA Analysis – Worst Case Response Time

$$R_i = C_i + \left(t_{TDMA} - t_{P_i} \right) \cdot \left\lceil \frac{C_i}{t_{P_i}} \right\rceil$$

Core execution time

Preemption by other tasks


Max. number of TDMA rounds needed for one Execution

Accumulated time of preemption until task finishes

Example:


$$C=10, t_p=3, t_{TDMA}=6$$

$$R = 10 + (6 - 3) \cdot \left\lceil \frac{10}{3} \right\rceil = 22$$




TDMA Output Jitter

- Different response times cause output jitter



TDMA Scheduling - Example



TDMA – Practical Issues

Clock synchronization in distributed embedded systems

- All tasks have to be synchronized
 - Phase and frequency adjustment




Quelle: <https://de.wikipedia.org/wiki/FlexRay>

Synchronization Methods

- Every task knows the entire TDMA schedule
- “Bus Sniffing”

After Synchronization

- No need for dedicated arbitration




Output event model

- May create output jitter due to differing RI


Round Robin Scheduling

- Dynamic time-driven scheduling
- Cyclic task execution
- Task release resources when execution finishes (no forced idle times)
 - Better resource utilization than TDMA



Round Robin - Practical Issues (1/2)

- Dynamic behavior makes analysis more difficult than TDMA
- Output event model
- May create output jitter and bursts
- Example: During execution P_3 generates an event every 5 clock cycles




Round Robin - Practical Issues (2/2)

Clock synchronization in distributed embedded systems

- Synchronization requirements (like TDMA)
- Implementation more challenging than TDMA
 - Does a task want to send data?
 - Next task to schedule?
- Control Overhead




Quelle: <https://de.wikipedia.org/wiki/FlexRay>



Quelle: https://www.nxp.com/files/static/abstract/overview_applications/FRWORKS.html

Rate Monotonic Scheduling (RMS)

- Priority-driven scheduling approach
- Deadlines at the end of each task's period
- Fixed priorities
 - The shorter the period, the higher the priority
- Optimal with regard to single processor scheduling
- Commonly used
- Little cost



RMS - Analysis

Processor utilization: $U(n)$

A system of n independent RMS scheduled processes always meets its deadlines if (sufficient):

(1)

$$\sum_{i=1}^n \underbrace{\frac{C_i}{T_i}}_{\text{Utilization of process } i} = U(n) \leq n(2^{\frac{1}{n}} - 1) \quad (\text{Liu/Layland '73})$$

Utilization
of process i

Where:

C_i : Core execution time of process i

T_i : Period of process i

→ **Sufficient** condition, but not necessary

→ If not met, no conclusion with respect to schedulability can be drawn

RMS – Example 1


$$T_1=100 \quad C_1=20$$

$$T_2=150 \quad C_2=40$$

$$T_3=350 \quad C_3=100$$

$$\sum_{i=1}^3 \frac{C_i}{T_i} = 75,24\% \leq 77,98\% = 3(2^{1/3} - 1)$$

Equation (1) met, system schedulable



RMS – Example 2 (1/2)


$$T_1=100 \quad C_1=30$$

$$T_2=150 \quad C_2=40$$

$$T_3=350 \quad C_3=100$$

$$\sum_{i=1}^3 \frac{C_i}{T_i} = 85,24\% > 77,98\% = 3(2^{1/3} - 1)$$

Equation (1) not met, schedule not guaranteed



- How to prove that all deadlines are still met in this case?

RMS - Workload

At each time t the accumulated requested workload is given by

$$W_n(t) = C_1 \left\lceil \frac{t}{T_1} \right\rceil + C_2 \left\lceil \frac{t}{T_2} \right\rceil + \dots + C_n \left\lceil \frac{t}{T_n} \right\rceil = \sum_{i=1}^n C_i \left\lceil \frac{t}{T_i} \right\rceil$$

- Evaluate at $t=nT_i$ for every task i
- If the workload $W_n(t)$ at time t is less than or equal to t , sufficient resources are available


RMS – Example 2 (2/2)

$$T_1=100 \quad C_1=30$$

$$T_2=150 \quad C_2=40$$

$$T_3=350 \quad C_3=100$$

$$\sum_{i=1}^3 \frac{C_i}{T_i} = 85,24\% > 77,98\% = 3(2^{1/3} - 1)$$



Workload for n=3:

$$W_3(t) = C_1 \left\lceil \frac{t}{T_1} \right\rceil + C_2 \left\lceil \frac{t}{T_2} \right\rceil + C_3 \left\lceil \frac{t}{T_3} \right\rceil$$

(with t=350)

$$= 30 \left\lceil \frac{350}{100} \right\rceil + 40 \left\lceil \frac{350}{150} \right\rceil + 100 \left\lceil \frac{350}{350} \right\rceil = 340 \leq 350$$

RMS – Example 3


$$T_1=100 \quad C_1=30$$

$$T_2=150 \quad C_2=40$$

$$T_3=250 \quad C_3=100$$

$$\sum_{i=1}^3 \frac{C_i}{T_i} = 96,67\% > 77,98\% = 3(2^{1/3} - 1)$$

Equation (1) not met, schedule not guaranteed



$$W_3(t) = C_1 \left\lceil \frac{t}{T_1} \right\rceil + C_2 \left\lceil \frac{t}{T_2} \right\rceil + C_3 \left\lceil \frac{t}{T_3} \right\rceil = 30 \left\lceil \frac{250}{100} \right\rceil + 40 \left\lceil \frac{250}{150} \right\rceil + 100 \left\lceil \frac{250}{250} \right\rceil = 270 \geq 250$$

RMS - Schedulability

Requirement: Worst case response time has to be smaller than deadline

How to trigger worst case behavior?

Critical instant = Situation in which a certain task experiences its worst case response time

- In case of RMS: A task is activated together with all higher priority tasks

Given a critical instant, if a task in a RMS scheduled system meets its first deadline then it will meet all deadlines

Recursive approach:

$$R_i^n = C_i + \underbrace{\sum_{j \in hp(i)} C_j \cdot \left\lceil \frac{R_i^{n-1}}{T_j} \right\rceil}_{\text{Core execution time of task } i} \leq D_i, \quad R_i^0 = 0$$

Core execution time of task i Preemption of task i by all higher priority tasks j in time window given by R_i^{n-1}

Recur until fixed point R_i reached or (due to monotocity)

$$R_i^n > D_i$$


Analysis with Arbitrary Deadlines

Arbitrary deadlines

- Now additional task activations during task execution/preemption possible

Goal: Find worst case response time and check $R_i \leq D_i$

Solution: Windowing technique



iterate over $q=1, \dots \{$

$$w_i(q) = q \cdot C_i + \sum_{j \in hp(i)} C_j \cdot \left\lceil \frac{w_i(q)}{T_j} \right\rceil \quad \} \text{ Workload for } q \text{ activations of task } i \text{ (Time to finish } q \text{ activations)}$$


$$R_i(q) = w_i(q) - (q-1)T_i \quad \} \text{ Response time of the } q\text{'th activation of task } i$$

$\} \text{ until } (w_i(q) \leq q \cdot T_i)$

$\} \text{ Iterate until } q\text{'th activation finishes before the } q+1\text{'th}$

$$R_{i,WorstCase} = \max_q \{ R_i(q) \}$$

Static Priority Preemptive Scheduling (1/4)




$$w_i(1) = C_i + \sum_{j \in hp(i)} C_j \cdot \left\lceil \frac{w_i(1)}{T_j} \right\rceil = 75 \quad \left\{ \begin{array}{l} w_3^0(1) = 0 \\ w_3^1(1) = 15 + 20 \left\lceil \frac{0}{75} \right\rceil + 40 \left\lceil \frac{0}{100} \right\rceil = 15 \\ w_3^2(1) = 15 + 20 \left\lceil \frac{15}{75} \right\rceil + 40 \left\lceil \frac{15}{100} \right\rceil = 75 \\ w_3^3(1) = 15 + 20 \left\lceil \frac{75}{75} \right\rceil + 40 \left\lceil \frac{75}{100} \right\rceil = 75 \end{array} \right. \quad \text{Fixed point reached}$$

$$R_i(1) = w_i(1) - (1-1)T_i = 75 - 0 = 75$$

$$w_i(1) = 75 > 55 = 1 \cdot T_i \quad \rightarrow \text{CONTINUE}$$

Static Priority Preemptive Scheduling (2/4)




$$w_i(2) = 2C_i + \sum_{j \in hp(i)} C_j \cdot \left\lceil \frac{w_i(2)}{T_j} \right\rceil = 150$$

$$R_i(2) = w_i(2) - (2-1)T_i = 150 - 55 = 95$$

$$w_i(2) = 150 > 110 = 2 \cdot T_i \quad \rightarrow \text{CONTINUE}$$

Static Priority Preemptive Scheduling (3/4)




$$w_i(3) = 3C_i + \sum_{j \in hp(i)} C_j \cdot \left\lceil \frac{w_i(3)}{T_j} \right\rceil = 185$$

$$R_i(3) = w_i(3) - (3-1)T_i = 185 - 2 \cdot 55 = 75$$

$$w_i(3) = 185 > 165 = 3 \cdot T_i \quad \rightarrow \text{CONTINUE}$$

Static Priority Preemptive Scheduling (4/4)



$$w_i(4) = 4C_i + \sum_{j \in hp(i)} C_j \cdot \left\lceil \frac{w_i(4)}{T_j} \right\rceil = 200$$

$$R_i(4) = w_i(4) - (4-1)T_i = 200 - 3 \cdot 55 = 35$$

$$w_i(4) = 200 < 220 = 4 \cdot T_i$$

STOP!


$$R_{i,WorstCase} = \max\{75, 95, 35\} = 95$$

Generalization

- Arbitrary priority assignment
- Arbitrary event models

$$R_i^n = C_i + \sum_{j=1}^{i-1} C_j \cdot \eta_j^+(R_i^{n-1}) \leq D_i$$

$$R_i^0 = 0$$



Reminder $\eta^+(\Delta t)$:
Max. number of
events that can occur
in time interval Δt .

Generalizing the Windowing Technique


Arbitrary event models and arbitrary deadlines

$$w_i(q) = q \cdot C_i + \sum_{j \in hp(i)} C_j \cdot \eta_j^+(w_i(q))$$

$$R_i(q) = w_i(q) - \delta_i^-(q)$$

$$w_i(q) \leq \delta_i^-(q+1)$$

$\delta^-(n)$ is the inverse of $\eta^+(\Delta t)$:
Smallest interval in which any n events may occur



END

