Periodic Resource Model for Compositional Real-Time Guarantees

Insik Shin and Insup Lee in RTSS'03

presented by Harald Röck

University of Salzburg, Department of Computer Sciences

8. November 2007

Outline

1. Introduction
2. Periodic Resource Model
3. Schedulability Analysis
4. Schedulability Bounds
5. Compositional Real-Time Guarantees

Introduction

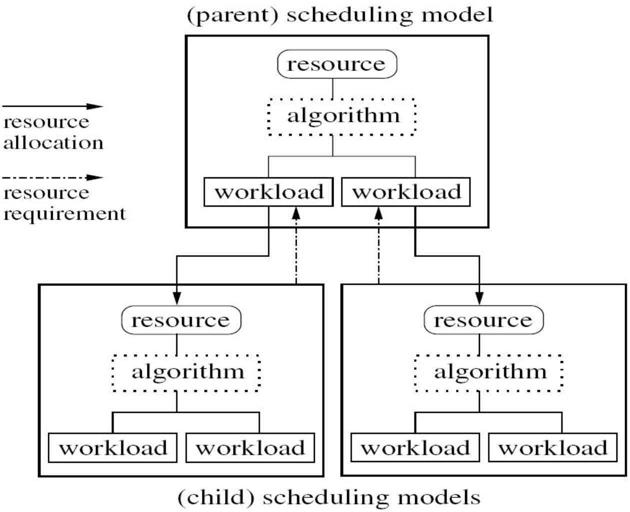

System Model

- Scheduling Model $M = (W, R, A)$
  - W the workload model set
  - R the resource model
  - A the scheduling algorithm
- Limitations:
  - Periodic tasks, i.e., $T(p, e)$
  - Tasks are independent and preemptive
  - Scheduling algorithms are limited to EDF and RM

System Model

- Scheduling Model $M = (W, R, A)$
  - W the workload model set
  - R the resource model
  - A the scheduling algorithm
- Limitations:
  - Periodic tasks, i.e., $T(p, e)$
  - Tasks are independent and preemptive
  - Scheduling algorithms are limited to EDF and RM

Definition

A scheduling model $M(W, R, A)$ is *schedulable* if a set of periodic workloads W is schedulable under a scheduling algorithm A with a partitioned resource R.

Existing Resource Models

Definition

Fractional resource $R_F(U_F)$: resource R is always available but only at a fractional capacity of U_F

Existing Resource Models

Definition

Fractional resource $R_F(U_F)$: resource R is always available but only at a fractional capacity of U_F

Definition

Bounded-delay resource partition model $R_B(U_B, D_B)$: resource R is available at its full capacity at some times but not available at the other times. U_B is the availability factor (rate), and D_B measures the largest deviation of a partition on any time interval.

Example

Two periodic tasks

- Two periodic tasks $T_1(7, 3)$ and $T_2(21, 1)$
- EDF scheduling algorithm
- Partitioned resource R is available 3 time units every 5 time units
- From Definitions 4 and 7 in [5], $R_B(U_B, D_B)$ has $U_B = 0.6$ and $D_B = 4$
- Scheduling model is $M(\{T_1, T_2\}, R_B(0.6, 4), EDF)$.

Example

Two periodic tasks

- Two periodic tasks $T_1(7,3)$ and $T_2(21,1)$
- EDF scheduling algorithm
- Partitioned resource R is available 3 time units every 5 time units
- From Definitions 4 and 7 in [5], $R_B(U_B, D_B)$ has $U_B = 0.6$ and $D_B = 4$
- Scheduling model is $M(\{T_1, T_2\}, R_B(0.6, 4), EDF)$.

Schedulability of *M* according to [5]

Over the fractional resource $U_B = 0.6$, T_1 and T_2 finish their execution at least 2 time units earlier then their deadlines. According to Theorem 1 in [5] the schedulability of M is inconclusive, because $2 < 4 = D_B$.

Problem Statement

1. Exact schedulability analysis: necessary and sufficient way to determine whether or not $M(W, \Gamma, A)$ is schedulable.

Problem Statement

1. Exact schedulability analysis: necessary and sufficient way to determine whether or not $M(W, \Gamma, A)$ is schedulable.
2. Periodic capacity bound: given W, A, and Π, find the smallest possible periodic capacity bound (Θ^* / Π) such that $M(W, \Gamma(\Pi, \Theta), A)$ is schedulable if $\Theta \geq \Theta^*$.
3. Utilization bound: given Γ and A find the largest possible utilization bound UB such that $M(W, \Gamma, A)$ is schedulable if $\sum_{T_i \in W} \frac{e_i}{p_i} \leq UB$

Problem Statement

1. Exact schedulability analysis: necessary and sufficient way to determine whether or not $M(W, \Gamma, A)$ is schedulable.
2. Periodic capacity bound: given W, A, and Π, find the smallest possible periodic capacity bound (Θ^* / Π) such that $M(W, \Gamma(\Pi, \Theta), A)$ is schedulable if $\Theta \geq \Theta^*$.
3. Utilization bound: given Γ and A find the largest possible utilization bound UB such that $M(W, \Gamma, A)$ is schedulable if $\sum_{T_i \in W} \frac{e_i}{p_i} \leq UB$
4. Algorithm set: given W and Γ find a set of algorithms $\mathcal{A}$ such that $M(W, \Gamma, A)$ is schedulable if $A \in \mathcal{A}$.

Problem Statement

1. Exact schedulability analysis: necessary and sufficient way to determine whether or not $M(W, \Gamma, A)$ is schedulable.
2. Periodic capacity bound: given W, A, and Π, find the smallest possible periodic capacity bound (Θ^*/Π) such that $M(W, \Gamma(\Pi, \Theta), A)$ is schedulable if $\Theta \geq \Theta^*$.
3. Utilization bound: given Γ and A find the largest possible utilization bound UB such that $M(W, \Gamma, A)$ is schedulable if $\sum_{T_i \in W} \frac{e_i}{p_i} \leq UB$
4. Algorithm set: given W and Γ find a set of algorithms $\mathcal{A}$ such that $M(W, \Gamma, A)$ is schedulable if $A \in \mathcal{A}$.
5. Compositional guarantee: given n scheduling models, derive a new parent scheduling model from the n scheduling models such that the new scheduling model is schedulable if and only if, the n child models are schedulable.

Periodic Resource Model

- Periodic resource model $\Gamma(\Pi, \Theta)$ characterizes a partitioned resource that guarantees allocations of Θ time units every Π time units.
- In the previous example we had $\Gamma(5, 3)$.

Definition

The *resource supply* of a resource is defined as the amount of resource allocation that the resource provides

Resource supply bound

Definition

For a periodic resource $\Gamma(\Pi, \Theta)$ a resource supply bound function $sbf_\Gamma(t)$ of a time interval length t calculates the minimum resource supply of Γ during t times units. (exact math in the paper)

- sbf_Γ is a non-decreasing step function
- Introduce a linear lower bound function as $lsbf_\Gamma$ such that $lsbf_\Gamma(t) \leq sbf_\Gamma(t)$. (exact math in the paper)

Resource supply

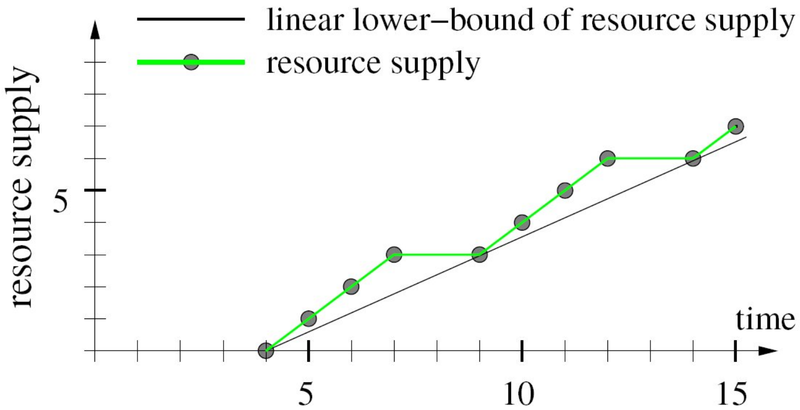

Service time

Definition

The *service time* of a resource is defined as the duration that it takes for the resource to provide a resource supply.

Service time

Definition

The *service time* of a resource is defined as the duration that it takes for the resource to provide a resource supply.

- For a periodic resource $\Gamma(\Pi, \Theta)$, the service time bound function $tbf_{\Gamma(t)}$ calculates the maximum service time of Γ for a t-time-unit resource supply.
- tbf_{Γ} is a non-decreasing step function
- $ltbf_{\Gamma}(t)$ is a linear service time bound function that upper-bounds $tbf_{\Gamma}(t)$

Service time

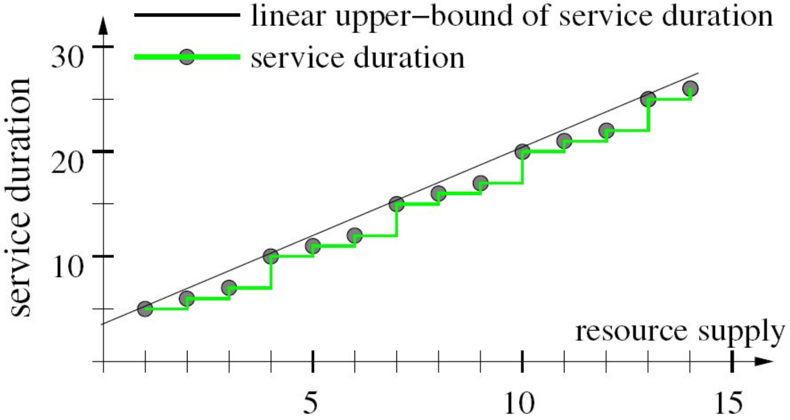

Resource Demand

Definition

The *resource demand* of a workload set is defined as the amount of resource allocation that the workload set requests.

Resource Demand

Definition

The *resource demand* of a workload set is defined as the amount of resource allocation that the workload set requests.

- For a periodic workload set W the resource demand bound function $dbf_W(t)$ of a time interval length t calculates the maximum resource demand of W during t time units.
- $dbf_W(t)$ is a discrete step function.
- A linear demand bound function $ldbf_W(t)$ upper-bounds $dbf_W(t)$.

Resource Demand

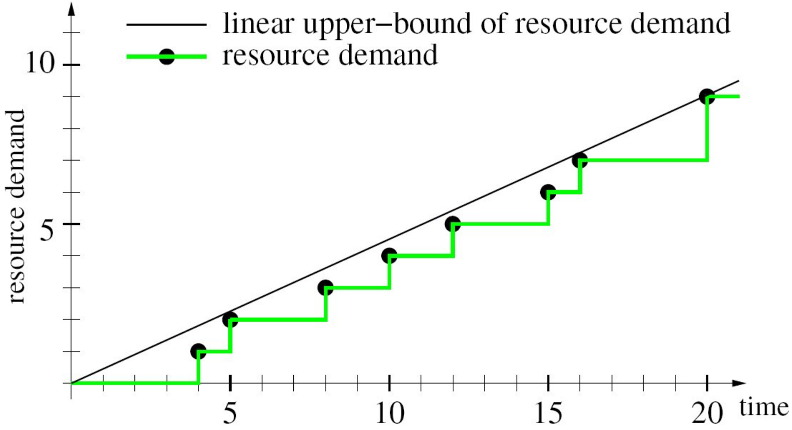

EDF Schedulability Analysis

Theorem

For a given scheduling model $M(W, \Gamma, EDF)$, M is schedulable if and only if the resource demand of W during a time interval is no greater than the resource supply of Γ during the same time interval for all time intervals during a hyperperiod, i.e.,

$$\forall 0 < t \leq 2 \cdot LCM_W : dbf_W(t) \leq sbf_\Gamma(t)$$

EDF Schedulability Analysis

Theorem

For a given scheduling model $M(W, \Gamma, EDF)$, M is schedulable if and only if the resource demand of W during a time interval is no greater than the resource supply of Γ during the same time interval for all time intervals during a hyperperiod, i.e.,

$$\forall 0 < t \leq 2 \cdot LCM_W : dbf_W(t) \leq sbf_\Gamma(t)$$

- Proof in the paper.

Example

Example

- Scheduling model $M(W, \Gamma(5,3), EDF)$
- $W = \{T_1(7,3), T_2(21,1)\}$

Example

Example

- Scheduling model $M(W, \Gamma(5,3), EDF)$
- $W = \{T_1(7,3), T_2(21,1)\}$

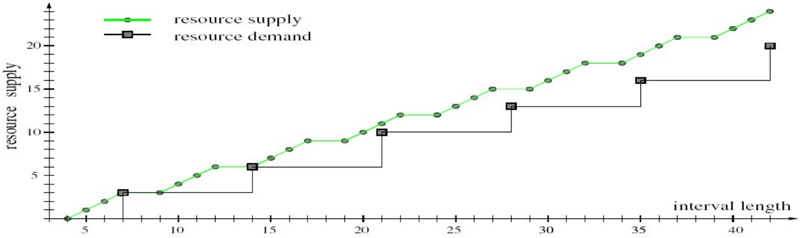

- M is schedulable, because $dbf_W(t) \leq sbf_\Gamma(t), 0 < t \leq 42$

Schedulability Bounds

- Scheduling model $M(W, \Gamma, A)$ provides only two of its elements, but not the other element
- Find a bound for the missing element such that M is schedulable
  - *Periodic capacity bound* (PCB): $M(W, \Gamma, A)$ provides W and A, calculate a bound for Γ.
  - *Utilization bound* (UB): $M(W, \Gamma, A)$ provides Γ and A, calculate a bound for W.

Periodic Capacity Bounds

Definition

The *periodic capacity bound* $PCB_W(\Pi)$ of a resource period Π is a number such that a scheduling model $M(W, \Gamma(\Pi, \Theta), A)$ is schedulable if

$$PCB_W(\Pi) \leq \frac{\Theta}{\Pi}$$

Periodic Capacity Bounds

Definition

The *periodic capacity bound* $PCB_W(\Pi)$ of a resource period Π is a number such that a scheduling model $M(W, \Gamma(\Pi, \Theta), A)$ is schedulable if

$$PCB_W(\Pi) \leq \frac{\Theta}{\Pi}$$

- Resource capacity:
  - Capacity of a periodic resource $\Gamma(\Pi, \Theta)$ is $\frac{\Theta}{\Pi}$
- How to get the optimal PCB_W?
  - Assume period Π is given, and $PCB_W(\Pi, EDF) = \frac{\Theta}{\Pi}$.
  - Find the smallest possible Θ such that $\forall 0 < t \leq 2 \cdot LCM_W : dbf_W(t) \leq sbf_\Gamma(t)$.
- Use the $lsbf_\Gamma(t)$ function to find a PCB_W numerically. (formula in the paper)

Utilization Bound

Definition

The *utilization bound* UB_Γ of a given resource Γ is a number such that a scheduling model $M(W, \Gamma, A)$ is schedulable if

$$\sum_{T_i \in W} \frac{e_i}{p_i} \leq UB_\Gamma$$

Utilization Bound

Definition

The *utilization bound* UB_Γ of a given resource Γ is a number such that a scheduling model $M(W, \Gamma, A)$ is schedulable if

$$\sum_{T_i \in W} \frac{e_i}{p_i} \leq UB_\Gamma$$

- Used to perform admission tests of a periodic workload set W over a given periodic resource Γ.
- Details of the formula and the proof are in the paper

Compositional Real-Time Guarantees

- Given multiple scheduling models $M_1, \ldots, M_n$ derive a parent scheduling model $M_P(W_P, \Gamma_P, A_P)$ from $M_1, \ldots, M_n$ as follows:

Compositional Real-Time Guarantees

- Given multiple scheduling models $M_1, \ldots, M_n$ derive a parent scheduling model $M_P(W_P, \Gamma_P, A_P)$ from $M_1, \ldots, M_n$ as follows:
  - A_P and Π_P are given

- Given multiple scheduling models $M_1, \ldots, M_n$ derive a parent scheduling model $M_P(W_P, \Gamma_P, A_P)$ from $M_1, \ldots, M_n$ as follows:
  - A_P and Π_P are given
  - derive W_P from the child resource models $\Gamma_i(\Pi_i, \Theta_i)$ such that $W_P = \{T_1(\Pi_1, \Theta_1), \cdots, T_n(\Pi_n, \Theta_n)\}$.

Compositional Real-Time Guarantees

- Given multiple scheduling models $M_1, \ldots, M_n$ derive a parent scheduling model $M_P(W_P, \Gamma_P, A_P)$ from $M_1, \ldots, M_n$ as follows:
  - A_P and Π_P are given
  - derive W_P from the child resource models $\Gamma_i(\Pi_i, \Theta_i)$ such that $W_P = \{T_1(\Pi_1, \Theta_1), \cdots, T_n(\Pi_n, \Theta_n)\}$.
  - derive $PCB_{W_P}(\Pi_P, A_P)$

Compositional Real-Time Guarantees

- Given multiple scheduling models $M_1, \ldots, M_n$ derive a parent scheduling model $M_P(W_P, \Gamma_P, A_P)$ from $M_1, \ldots, M_n$ as follows:
  - A_P and Π_P are given
  - derive W_P from the child resource models $\Gamma_i(\Pi_i, \Theta_i)$ such that $W_P = \{T_1(\Pi_1, \Theta_1), \cdots, T_n(\Pi_n, \Theta_n)\}$.
  - derive $PCB_{W_P}(\Pi_P, A_P)$
  - compute Θ_P such that $\Theta_P = \Pi_P \cdot PCB_{W_P}(\Pi_P, A_P)$.

Theorem

M_P is schedulable if and only if $M_1, \cdots, M_n$ are schedulable.

Conclusion

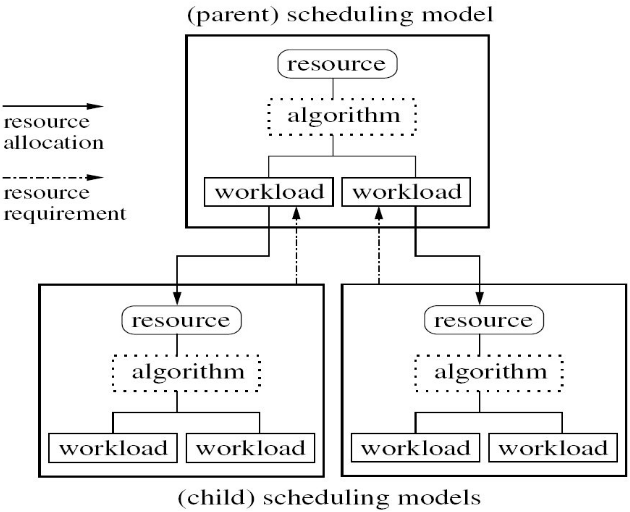
(parent) scheduling model
resource
algorithm
workload
workload
resource allocation
resource requirement
resource
algorithm
workload
workload
resource
algorithm
workload
workload
(child) scheduling models