Multiprocessor Global Schedulability Analysis in Continuous Time

Youcheng Sun

Grenoble, 2nd February 2016







► Youcheng Sun and Giuseppe Lipari.

"A Pre-Order Relation for Exact Schedulability Test of Sporadic Tasks on Multiprocessor Global Fixed-Priority Scheduling". @Real-Time Systems journal, 2015.

► Youcheng Sun.

"Real-Time Schedulability Analysis with Formal Techniques (Chapter 9)". PhD thesis, Scuola Superiore Sant'Anna, 2015.

1. Introduction

2. Multiprocessor G-FP Scheduling in LHA

3. Exact Schedulability Analysis

4. Approximate Schedulability Analysis

System model



- $m (m \le n)$ identical processors
- the Global scheduling policy
 - Global Fixed-Priority (G-FP)

Uniprocessor vs multiprocessor

- The Worst-Case Scenario is known upon uniprocessor
 but not true for multiprocessor
- An example from [Baruah@RTSS07]



Multiprocessor global schedulability analysis

- Exact schedulability analysis needs to consider a large amount possible arrival patterns
 - There is no worst-case scenario
 - Sporadic task activations add further complexity

State of the art

- ▶ [Baker and Cirinei@PDS2007]
- [Bonifaci and Marchetti-Spaccamela @Algorithmica2012]
- ▶ [Geeraerts et al@RTSJ2013]
- ▶ [Burmyakov and Bini@RTNS2015]

- 1. All in discrete time domain
- 2. Exact analysis in discrete time may be not always a good/safe idea

Continuous time vs discrete time

- Given a task set T that is schedulable by G-FP, is the schedulability preserved after scaling task parameters?
 - $\mathcal{T} = \{\tau_1 = (1, 4, 4), \tau_2 = (2, 6, 6)\} \rightarrow$ 10 · $\mathcal{T} = \{10 · \tau_1 = (10, 40, 40), 10 · \tau_2 = (20, 60, 60)\}$
 - Continuous time domain: YES!
 - Discrete time domain: ?

Example

- ► $\tau_1 = (1, 4, 4), \tau_2 = (1, 3, 3), \tau_3 = (1, 3, 3), \tau_4 = (1, 2, 2)$ and $m = 2 \rightarrow$ schedulable (under discrete time assumption)
- ► $\tau_1 = (10, 40, 40), \tau_2 = (10, 30, 30), \tau_3 = (10, 30, 30), \tau_4 = (10, 20, 20) \rightarrow \text{not schedulable}$



1. Introduction

2. Multiprocessor G-FP Scheduling in LHA

3. Exact Schedulability Analysis

4. Approximate Schedulability Analysis

Linear Hybrid Automata (LHA)



- continuous variables: x
 - its rate x
- locations: I_1 and I_2
- ► transitions: l₁ → l₂ with the synchronisation label (syncAct), guard (x ≤ 4) and update of the variable (x := 0)
- invariants: e.g. $x \ge 1$ in I_1

Concrete state and symbolic state

- A concrete state $s = (I, \nu)$
 - / is a location
 - ν is a valuation on variables
- A symbolic state S = (I, C)
 - / is a location
 - C is a linear constraint (represented by a convex region)
- *s_i* → *s_j*: a step change of states due to time elapse or a transition
- $s_i \Rightarrow s_j$: a sequence of step changes

Example: Var = $\{x, y\}$



Scheduling problem in LHA

- One automaton, named task automaton, per task for modeling the task's behaviour
- One (G-FP) scheduler automaton for making scheduling decisions

$$\begin{array}{c} \downarrow \\ \text{Idle} \\ \dot{\boldsymbol{p}} = 1, \, \dot{\boldsymbol{c}} = 0 \end{array}$$

Waiting
$$p = 1, c = 0$$

 $p \le D$

$$\begin{array}{c} \text{Running} \\ \dot{p} = 1, \, \dot{c} = -1 \\ c \ge 0 \land p \le \text{D} \end{array}$$



arrival

$$p \ge T$$

 $p := 0$
 $c := C$
Waiting
 $p = 1, c = 0$
 $p = 1, c = 0$
 $p = 1, c = -1$
 $c \ge 0 \land p \le D$
Running
 $p = 1, c = -1$
 $c \ge 0 \land p \le D$



















System automaton

- A system automaton SA = (T, Sched)
 - $T = {TA_1, ..., TA_n}$ is a set of *n* task automata;
 - Sched is the scheduler automaton;
 - $\ SA = Sched \times TA_1 \times \cdots \times TA_n.$
- ► The schedulability problem is now the reachability problem of DeadlineMissed in SA.

1. Introduction

2. Multiprocessor G-FP Scheduling in LHA

3. Exact Schedulability Analysis

4. Approximate Schedulability Analysis

Slack-time pre-order

For the SA automaton, the *slack-time* pre-order is defined as follows: $\forall s_1, s_2, s_1 \succeq_{st} s_2$ if and only if

 $\forall \tau_i: \mathbf{s}_1.\mathbf{p}_i \geq \mathbf{s}_2.\mathbf{p}_i \land \mathbf{s}_1.\mathbf{c}_i \geq \mathbf{s}_2.\mathbf{c}_i$

and we say s_1 dominates s_2 .

If a DeadlineMissed location is reachable from s_2 , so is from s_1 .

Extend \succeq_{st} to symbolic states

- ► A symbolic state S = (I, C) abstracts a (possibly infinite) set of concrete states
- For two symbolic states S₁ and S₂, we say S₁ dominates S₂ if

$$\forall \ \mathbf{s}_2 \in \mathbf{S}_2, \ \exists \ \mathbf{s}_1 \in \mathbf{S}_1 \ \mathsf{s.t.} \ \mathbf{s}_1 \succeq_{\mathbf{st}} \mathbf{s}_2$$

$$\bullet \ \mathbf{S}_1.\mathcal{C} \supseteq \mathbf{S}_2.\mathcal{C} \Rightarrow \mathbf{S}_1 \succeq_{\mathbf{st}} \mathbf{S}_2$$

More general case

► $(I_1, C_1) \succeq_{st} (I_2, C_2)$ if



A widening operator ∇



 $(I_1, C_1) \succeq_{st} (I_2, C_2)$ if and only if $\nabla(C_1)$ includes $\nabla(C_2)$.

Schedulability analysis in system automaton

Algorithm 1 Schedulability Analysis in SA (SA-SA)

- 1: $R \leftarrow \{S_0\}$
- 2: while true do
- 3: $P \leftarrow \mathsf{Post}(R)$
- 4: if $P \cap F \neq \emptyset$ then
- 5: return NOT schedulable
- 6: end if
- 7: $R' \leftarrow R \cup P$
- 8: $\mathbf{R}' \leftarrow \mathsf{Max}^{\succeq}(\mathbf{R}')$
- 9: if R' = R then
- 10: return schedulable
- 11: else
- 12: $R \leftarrow R'$
- 13: end if
- 14: end while

Decidability interval (when $D \leq T$)

For a set $\mathcal{T} = \{\tau_1, \dots, \tau_n\}$ of *n* sporadic tasks running on *m* processors, it is schedulable iff there is no deadline miss happens within [0, L] with $L = \sum_{1 \le i \le m} C_i + \sum_{m < i \le n} D_i$.

Proof sketch

(Assumption) $\mathcal{T}_{k} = \{\tau_{1}, \ldots, \tau_{k}\} \subseteq \mathcal{T}$ and there is a dominant time interval $[0, L_{k}]$: $\forall t' > L_{k} \exists t \in [0, L_{k}]$ s.t.

 $\forall \tau_i \in \mathcal{T}_k \ \ \boldsymbol{c}_i(t) \geq \boldsymbol{c}_i(t') \land \boldsymbol{p}_i(t) \geq \boldsymbol{p}_i(t')$

 $\implies [0, L_k + D_{k+1}]$ is enough for detecting if τ_{k+1} misses a deadline ($c_{k+1} > 0 \land p_{k+1} = D_{k+1}$)

The base:
$$L_m = \sum_{1 \le i \le m} C_i \Longrightarrow L_n = L_m + \sum_{m < i \le n} D_i$$



The termination of algorithm SA-SA is guaranteed.

Some simulation results

With or WithOut Simulation relation

State space size



Runtime complexity

• m = 2 and n = 6



Exact analysis vs approximate analysis

• m = 2 and n = 5



Summary

- An exact G-FP schedulability test in continuous time
 - Deal with general task parameters
- ► A pre-order relation to for faster reachability analysis
 - Open source tool: FOrmal Real-Time Scheduler (FORTS)
- The decidability interval for multiprocessor global scheduling of sporadic tasks
- Complexity remains high...

1. Introduction

2. Multiprocessor G-FP Scheduling in LHA

3. Exact Schedulability Analysis

4. Approximate Schedulability Analysis

















- RTA tests
 - e.g. [Guan et al@RTSS09], [Sun et al@RTCSA14]

$$X \leftarrow \left\lfloor \frac{\Omega_k(X)}{m} \right\rfloor + C_k$$

RTA (cont.)

 τ_k can complete its execution within a time interval of length x if

$$\left\lfloor \frac{\Omega_k(\boldsymbol{x})}{\boldsymbol{m}} \right\rfloor + \boldsymbol{C}_k \leq \boldsymbol{x}$$

RTA (cont.)

 τ_k can complete its execution within a time interval of length *x* if

$$\left\lfloor \frac{\Omega_k(\boldsymbol{x})}{\boldsymbol{m}} \right\rfloor + \boldsymbol{C}_k \leq \boldsymbol{x}$$

 Given a task set that is schedulable by the previous RTA, the schedulability may NOT be preserved under the same test after scaling task parameters.

• hint:
$$\lfloor \frac{15}{2} \rfloor = 7$$
, $\lfloor \frac{15 \cdot 10}{2} \rfloor = 75 > 7 \cdot 10$

Suggestions

 τ_k can complete its execution within a time interval of length ${\pmb x}$ if

$$\frac{\Omega_k(\boldsymbol{x})}{\boldsymbol{m}} + \boldsymbol{C}_k \leq \boldsymbol{x}$$

IRTA (Integer Response Time Analysis)

$$oldsymbol{X} \leftarrow \left\lceil rac{\Omega_k(oldsymbol{X})}{oldsymbol{m}}
ight
ceil + oldsymbol{C}_k$$

Thank you!