

Exact Schedulability Analysis of Global Fixed Priority Scheduling by Using Linear Hybrid Automata

Youcheng Sun¹, Giuseppe Lipari^{1 2}

¹Scuola Superiore Sant'Anna

²LSV - ENS Cachan and CNRS

October 8, 2014

- 1 Introduction
- 2 LHA models for multiprocessor G-FP scheduling
- 3 Weak Simulation Relation
 - Concrete state space
 - Symbolic state space
- 4 Experiments

- 1 Introduction
- 2 LHA models for multiprocessor G-FP scheduling
- 3 Weak Simulation Relation
 - Concrete state space
 - Symbolic state space
- 4 Experiments

Problem description

- m processors
- n sporadic tasks
- Global (task level) Fixed-Priority (G-FP) Preemptive Scheduling

Problem description

- m processors
- n sporadic tasks
- Global (task level) Fixed-Priority (G-FP) Preemptive Scheduling
- Is the taskset schedulable?

- A sporadic task τ_i s specified by a tuple (C, D, T)
 - C is the worst-case execution time
 - D is the relative deadline
 - T is the minimum time interval between two successive job releases of the task
 - $T \sim D$ with $\sim \in \{<, =, >\}$
- Fully Preemptive
- Tasks can migrate among different processors

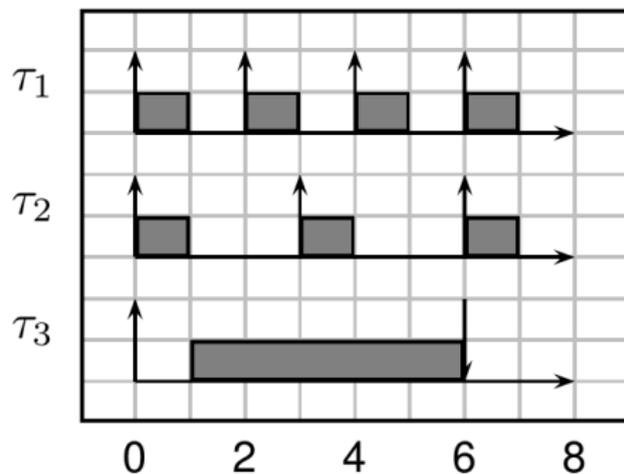
Critical instant for G-FP scheduling?

Critical instant for G-FP scheduling?

- Unknown

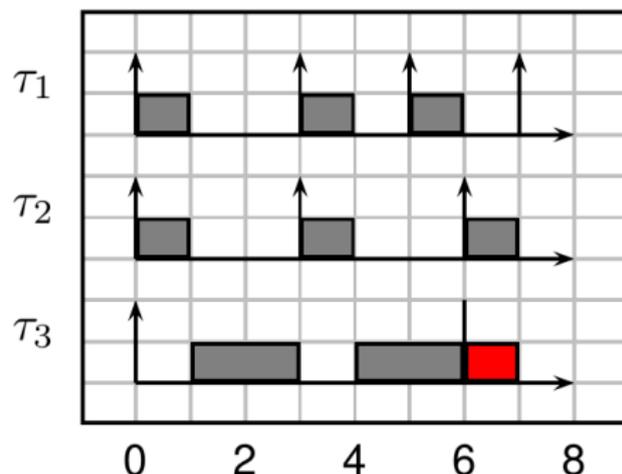
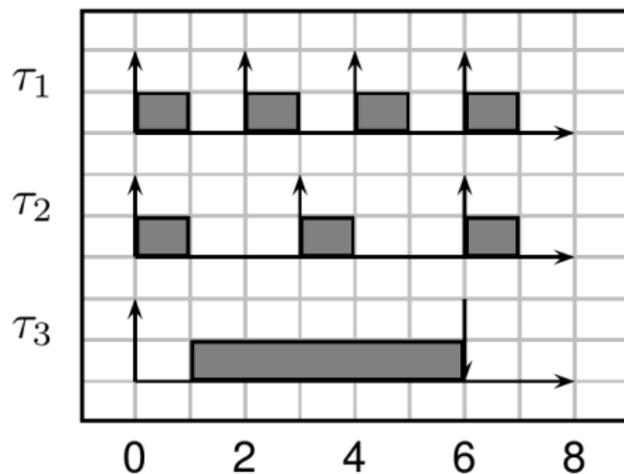
Critical instant for G-FP scheduling?

- Unknown
- An example (Baruah, 2007): $\tau_1 = (1, 2, 2)$, $\tau_2 = (1, 3, 3)$, and $\tau_3 = (5, 6, 6)$



Critical instant for G-FP scheduling?

- Unknown
- An example (Baruah, 2007): $\tau_1 = (1, 2, 2)$, $\tau_2 = (1, 3, 3)$, and $\tau_3 = (5, 6, 6)$



- Analytical (**over-approximate**) solutions: RTA-CE ([Sun et al., 2014](#)), RTA-LC ([Guan et al., 2009](#))

- Analytical (**over-approximate**) solutions: RTA-CE (Sun et al., 2014), RTA-LC (Guan et al., 2009)
- Model-based (**exact**) solutions : Geeraerts, Goossens, and Lindstrom, 2013 and Baker and Cirinei, 2007

- A Linear Hybrid Automaton (LHA) model for **exact** G-FP scheduling

- A Linear Hybrid Automaton (LHA) model for **exact** G-FP scheduling
- A [weak simulation relation](#) to simplify the state space exploration

- A Linear Hybrid Automaton (LHA) model for **exact** G-FP scheduling
- A **weak simulation relation** to simplify the state space exploration
- An evaluation on the pessimism of the state-of-the-art analytical G-FP schedulability analysis

Linear Hybrid Automata(LHA)

A Linear Hybrid Automaton is a tuple

$$H = \{V, D, L, \textit{init}, \textit{Lab}, T, \textit{Invar}\}$$

Linear Hybrid Automata(LHA)

A Linear Hybrid Automaton is a tuple

$$H = \{V, D, L, \text{init}, \text{Lab}, T, \text{Invar}\}$$

- 1 A finite set $V = \{x_1, \dots, x_n\}$ of continuous variables.
- 2 A labeling function D which linearly constrains variables' rate ($\dot{V} = \{\dot{x}_1, \dots, \dot{x}_N\}$) in each location.

Linear Hybrid Automata(LHA)

A Linear Hybrid Automaton is a tuple

$$H = \{V, D, L, \textit{init}, \textit{Lab}, T, \textit{Invar}\}$$

- 1 A finite set $V = \{x_1, \dots, x_n\}$ of continuous variables.
- 2 A labeling function D which linearly constrains variables' rate ($\dot{V} = \{\dot{x}_1, \dots, \dot{x}_N\}$) in each location.
- 3 A finite set L of locations.
- 4 An initial function *init*.
- 5 A finite set *Lab* of synchronisation labels.
- 6 A finite set T of transitions (**every location has an outgoing stutter transition to itself**).
- 7 A labeling function *Invar* which assigns each location l an *invariant*.

- A **concrete state** $s = (l, \nu)$: l is a **location** and ν is a **valuation** over V

- A **concrete state** $s = (l, \nu)$: l is a **location** and ν is a **valuation** over V
 - A transition $s_1 \rightarrow s_2$; a sequence of transitions $s_1 \Rightarrow s_2$
 - The concrete state space of LHA : `space`

Concrete states and Symbolic States

- A **concrete state** $s = (l, \nu) : l$ is a **location** and ν is a **valuation** over V
 - A transition $s_1 \rightarrow s_2$; a sequence of transitions $s_1 \Rightarrow s_2$
 - The concrete state space of LHA : space
- A **symbolic state** $S = (l, \mathcal{C}) : l$ is a **location** and \mathcal{C} is a **linear constraint** and can be represented by a convex region

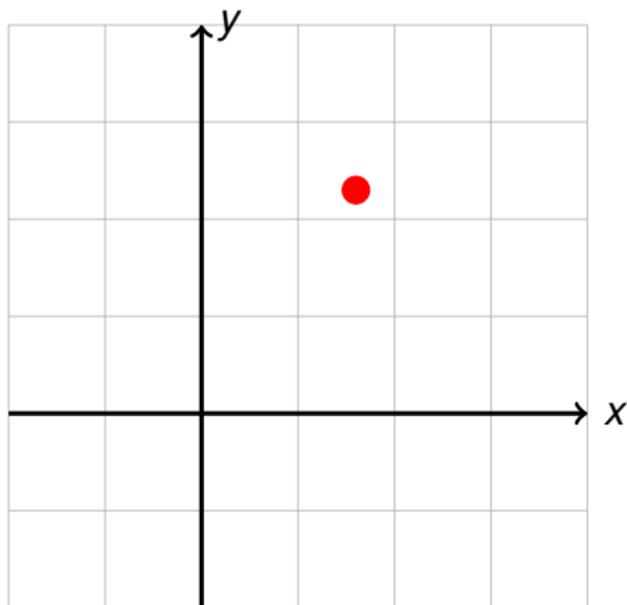
Concrete states and Symbolic States

- A **concrete state** $s = (l, \nu) : l$ is a **location** and ν is a **valuation** over V
 - A transition $s_1 \rightarrow s_2$; a sequence of transitions $s_1 \Rightarrow s_2$
 - The concrete state space of LHA : `space`
- A **symbolic state** $S = (l, \mathcal{C}) : l$ is a **location** and \mathcal{C} is a **linear constraint** and can be represented by a convex region
 - A transition $S_1 \rightarrow S_2$; a sequence of transitions $S_1 \Rightarrow S_2$
 - The symbolic state space of LHA : `space`

Examples

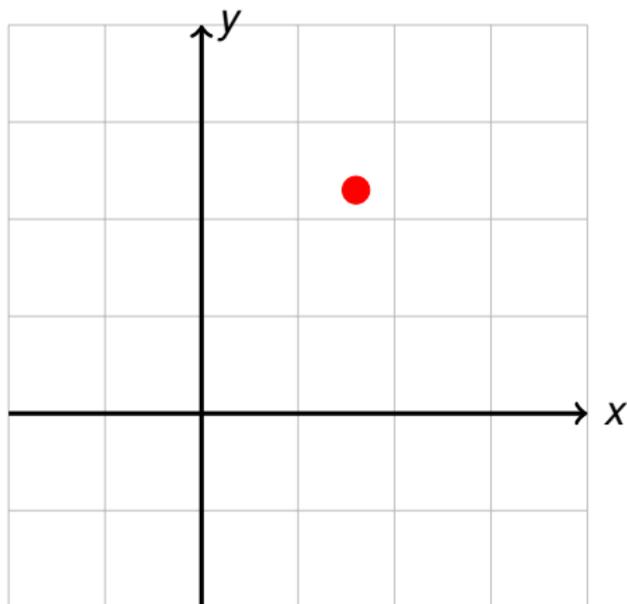
Examples

$s = (l, \nu)$ with
 $\nu = (1.6, 2.3)$

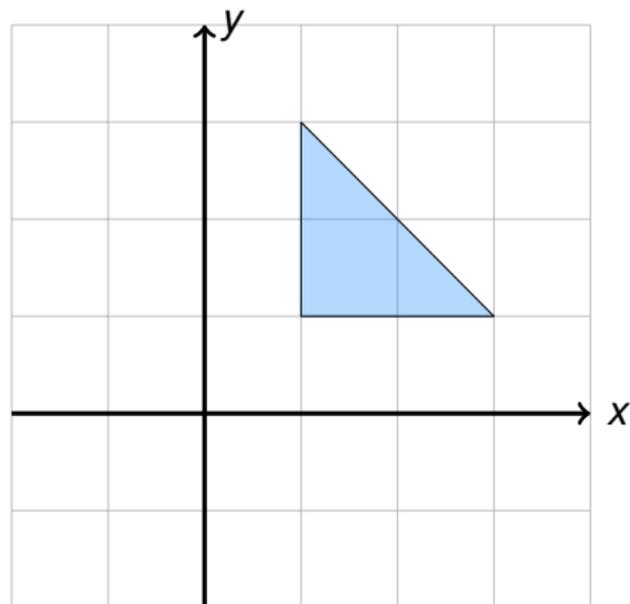


Examples

$s = (l, v)$ with
 $v = (1.6, 2.3)$



$S = (l, C)$ with
 $C = \{1 \leq x \leq 3 \wedge x + y \leq 4\}$

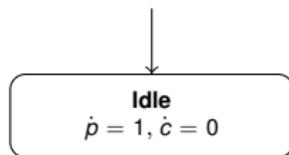


- 1 Introduction
- 2 LHA models for multiprocessor G-FP scheduling
- 3 Weak Simulation Relation
 - Concrete state space
 - Symbolic state space
- 4 Experiments

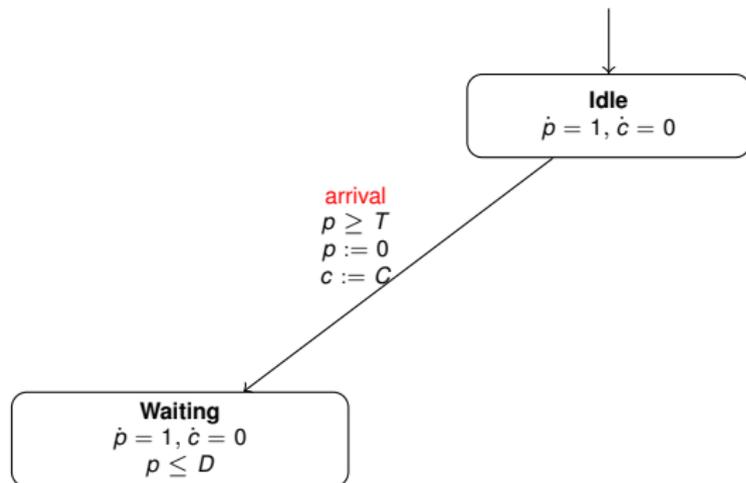
A task (C, D, T) is modeled by a LHA called `Task Automaton(TA)`

- Two continuous variables :
 - 1 p : the time passed since the latest task activation
 - 2 c : remaining computation time that needs to be executed

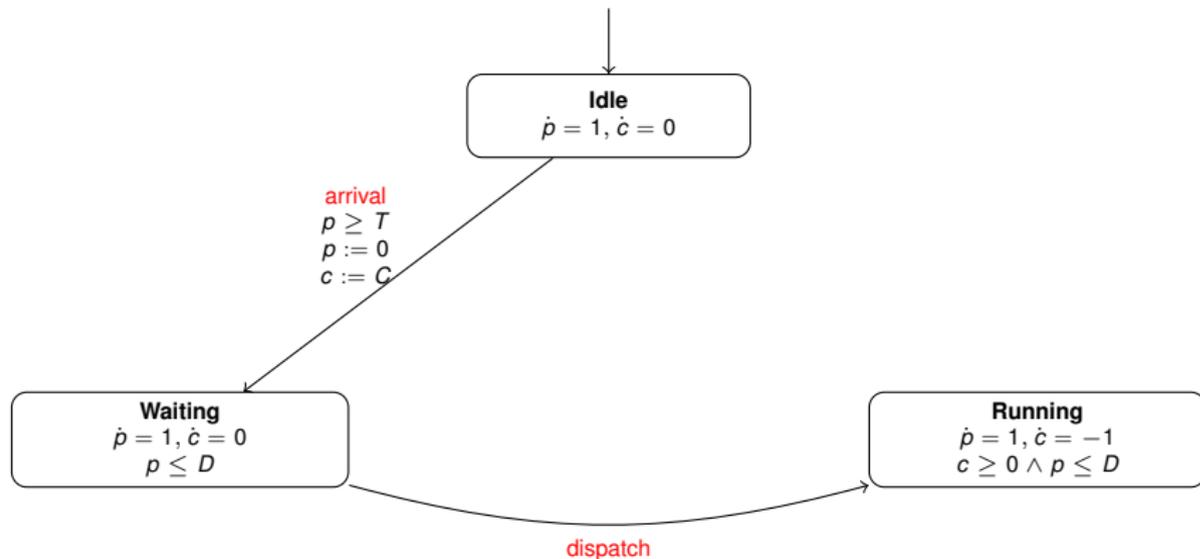
Task automaton



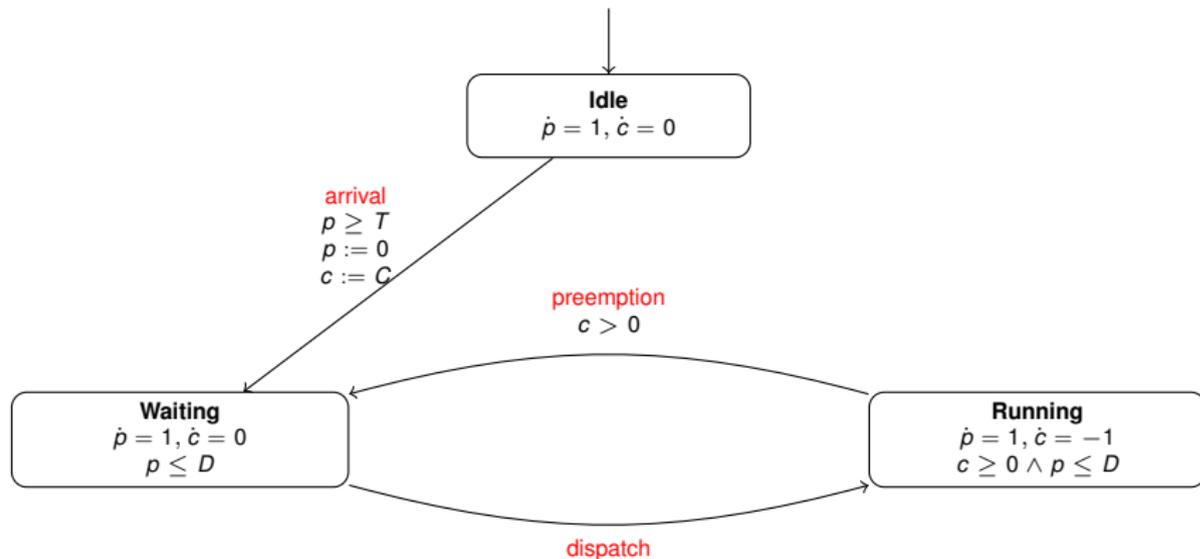
Task automaton



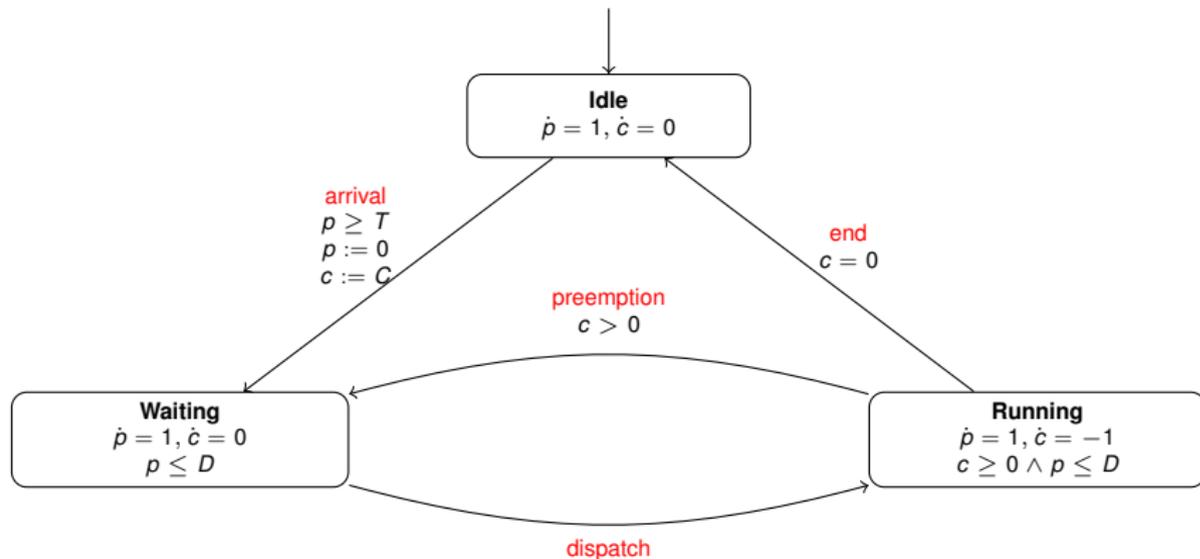
Task automaton



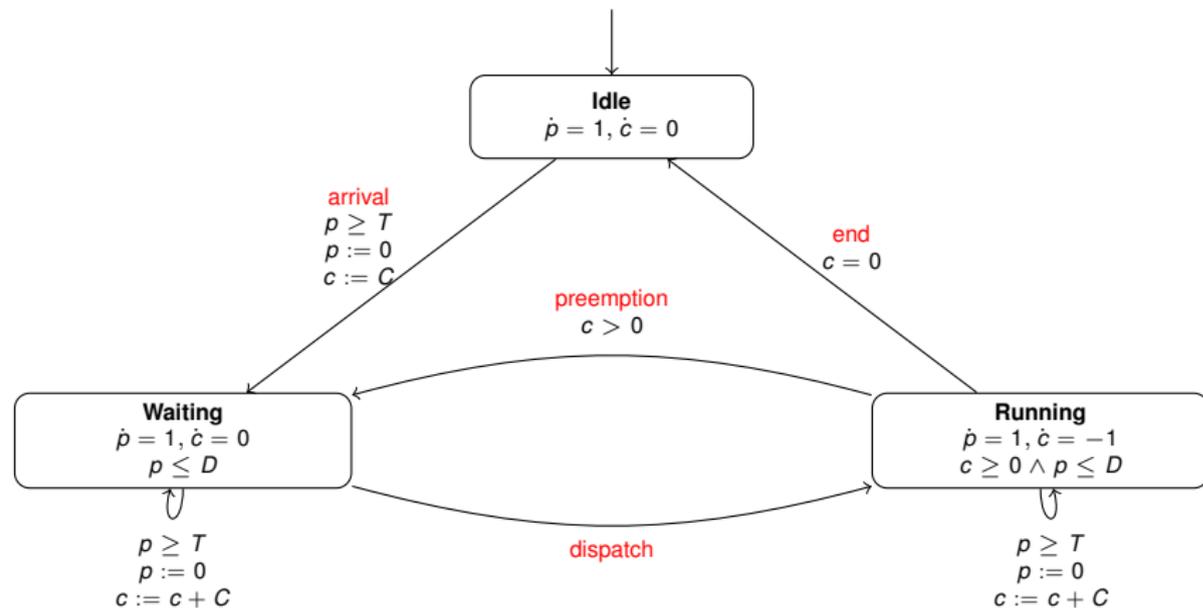
Task automaton



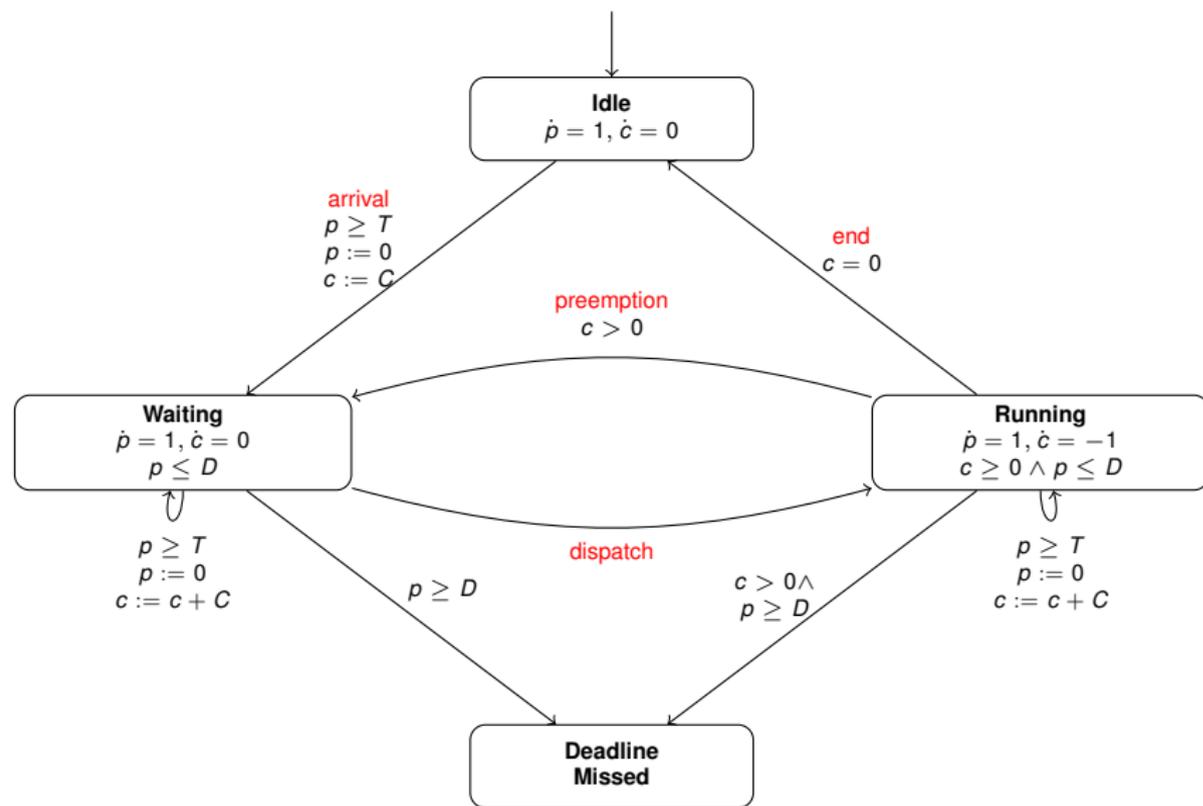
Task automaton



Task automaton



Task automaton



- The Scheduling Automaton(Sched)
 - It synchronises with TAs and decides which tasks to run and which tasks to wait.
 - It is a G-FP preemptive scheduler.
- The System Automaton(SA)
 - Composition of TAs and Sched : $Sched \times TA_1 \times \dots \times TA_n$

- The Scheduling Automaton(Sched)
 - It synchronises with TAs and decides which tasks to run and which tasks to wait.
 - It is a G-FP preemptive scheduler.
- The System Automaton(SA)
 - Composition of TAs and Sched : $Sched \times TA_1 \times \dots \times TA_n$
- The schedulability problem is now the reachability problem of DeadlineMissed in SA.

- 1 Introduction
- 2 LHA models for multiprocessor G-FP scheduling
- 3 Weak Simulation Relation**
 - Concrete state space
 - Symbolic state space
- 4 Experiments

- 1 Introduction
- 2 LHA models for multiprocessor G-FP scheduling
- 3 Weak Simulation Relation**
 - **Concrete state space**
 - Symbolic state space
- 4 Experiments

Definition

A weak simulation relation in concrete state space of SA is a preorder $\succeq \subseteq \text{space} \times \text{space}$ such that :

1 $\forall s_1, s_2, s_4$ s.t. $s_1 \succeq s_2, s_2 \rightarrow s_4$ there exists s_3 s.t.

$$s_1 \Rightarrow s_3 \text{ and } s_3 \succeq s_4.$$

2 $\forall s_1, s_2$ s.t. $s_1 \succeq s_2$:

s_2 in DeadlineMissed implies s_1 in DeadlineMissed

Whenever $s_1 \succeq s_2$, we say that s_1 (weak) simulates s_2 .

Definition

For the SA with a G-FP preemptive scheduler, its slack-time pre-order relation $\preceq_{st} \subseteq \text{space} \times \text{space}$ is defined such that $\forall s_1, s_2, s_1 \preceq_{st} s_2$ iff

$$\forall T_i : s_1.p_i \geq s_2.p_i \quad \wedge \quad s_1.c_i \geq s_2.c_i$$

The slack-time pre-order relation \preceq_{st}

Definition

For the SA with a G-FP preemptive scheduler, its slack-time pre-order relation $\preceq_{st} \subseteq \text{space} \times \text{space}$ is defined such that $\forall s_1, s_2, s_1 \preceq_{st} s_2$ iff

$$\forall \tau_i : s_1.p_i \geq s_2.p_i \quad \wedge \quad s_1.c_i \geq s_2.c_i$$

Theorem

\preceq_{st} is indeed a weak simulation relation in SA.

- 1 Introduction
- 2 LHA models for multiprocessor G-FP scheduling
- 3 Weak Simulation Relation**
 - Concrete state space
 - Symbolic state space**
- 4 Experiments

- A symbolic state $S = (I, \mathcal{C})$ abstracts a set of concrete states.
- For two symbolic states S_1 and S_2 , we say S_1 simulates S_2 if

$$\forall s_2 \in \mathcal{S}_2, \exists s_1 \in \mathcal{S}_1 \quad \text{s.t.} \quad s_1 \succeq s_2$$

A preliminary concept : convex region domination

- Assume a N-dimensional space
- Given two valuations $\nu = (\nu_1, \dots, \nu_N)$ and $\nu' = (\nu'_1, \dots, \nu'_N)$, we say ν **dominates** ν' , denoted as $\nu \geq \nu'$, if $\forall i \in [1, N], \nu_i \geq \nu'_i$.

A preliminary concept : convex region domination

- Assume a N-dimensional space
- Given two valuations $\nu = (\nu_1, \dots, \nu_N)$ and $\nu' = (\nu'_1, \dots, \nu'_N)$, we say ν **dominates** ν' , denoted as $\nu \geq \nu'$, if $\forall i \in [1, N], \nu_i \geq \nu'_i$.
- Given two convex regions \mathcal{C}_1 and \mathcal{C}_2 , we say \mathcal{C}_1 **dominates** \mathcal{C}_2 , denoted as $\mathcal{C}_1 \geq \mathcal{C}_2$, if for any valuation ν' in \mathcal{C}_2 , there exists a valuation $\nu \in \mathcal{C}_1$ such that $\nu \geq \nu'$.

Definition

For the SA with a G-FP preemptive scheduler, its slack-time pre-order relation $\succ_{st} \subseteq \text{Space} \times \text{Space}$ is defined such that $\forall S_1, S_2, S_1 \succ_{st} S_2$ iff $S_1.C$ dominates $S_2.C$.

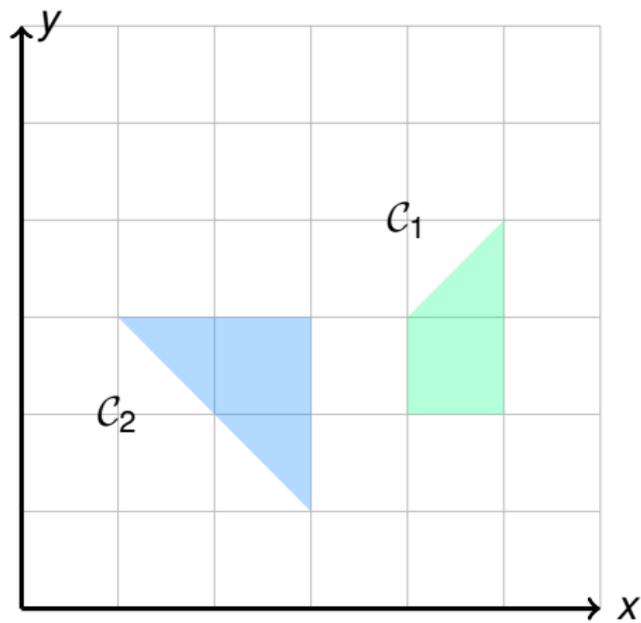
Definition

For the SA with a G-FP preemptive scheduler, its slack-time pre-order relation $\succ_{st} \subseteq \text{Space} \times \text{Space}$ is defined such that $\forall S_1, S_2, S_1 \succ_{st} S_2$ iff $S_1.C$ dominates $S_2.C$.

Theorem

$\succ_{st} \subseteq \text{Space} \times \text{Space}$ is a weak simulation relation.

How to decide $C_1 \geq C_2$?



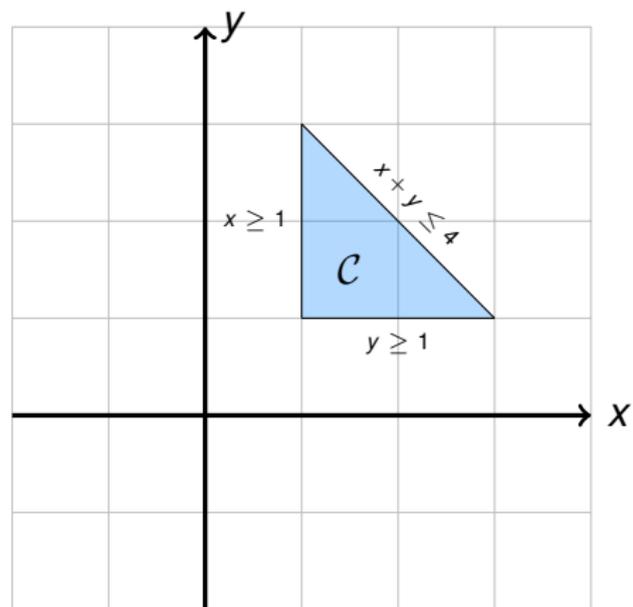
A widening operator ∇

- Given a convex region \mathcal{C}
- It's widening region $\nabla(\mathcal{C})$ is constructed as follows:
 - 1 Construct linear constraints \mathcal{C}' in $2 \times N$ dimensional space $(x_1, \dots, x_N, y_1, \dots, y_N)$ such that

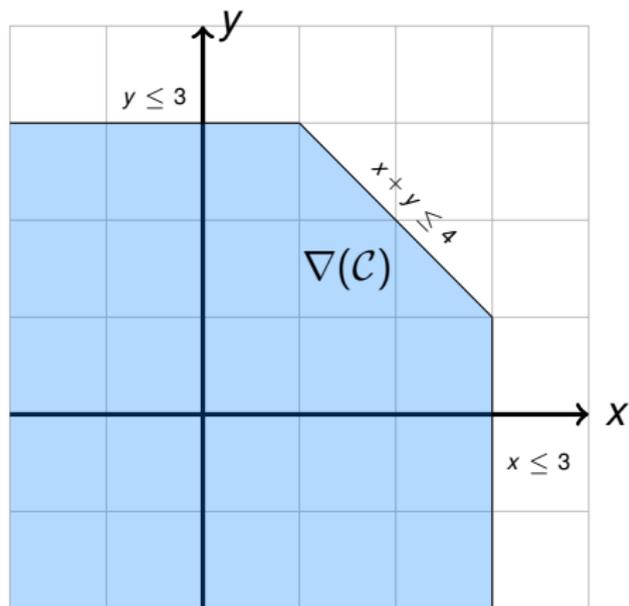
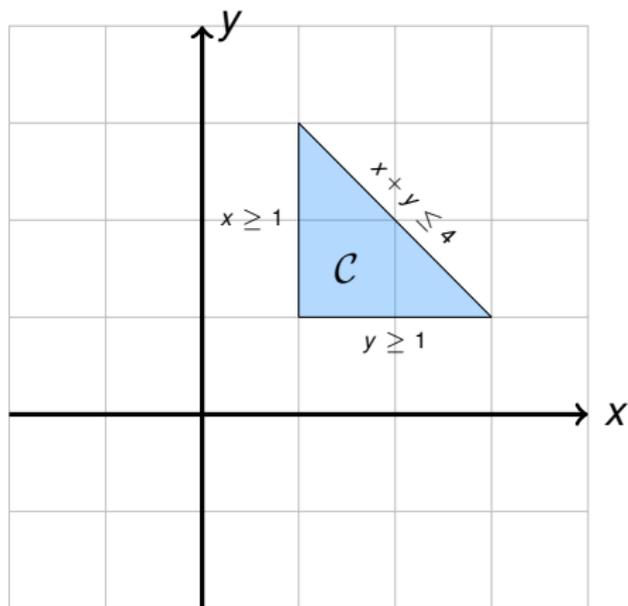
$$(y_1, \dots, y_N) \models \mathcal{C} \quad \wedge \quad \forall i, x_i \leq y_i$$

- 2 Remove the space dimensions higher than N in \mathcal{C}' .

An example of widening



An example of widening



Lemma

Given two convex regions \mathcal{C}_1 and \mathcal{C}_2 , $\mathcal{C}_1 \geq \mathcal{C}_2$ if and only if $\nabla(\mathcal{C}_1)$ includes $\nabla(\mathcal{C}_2)$.

Proof.



Algorithm 1 Schedulability Analysis in SA (SA-SA)

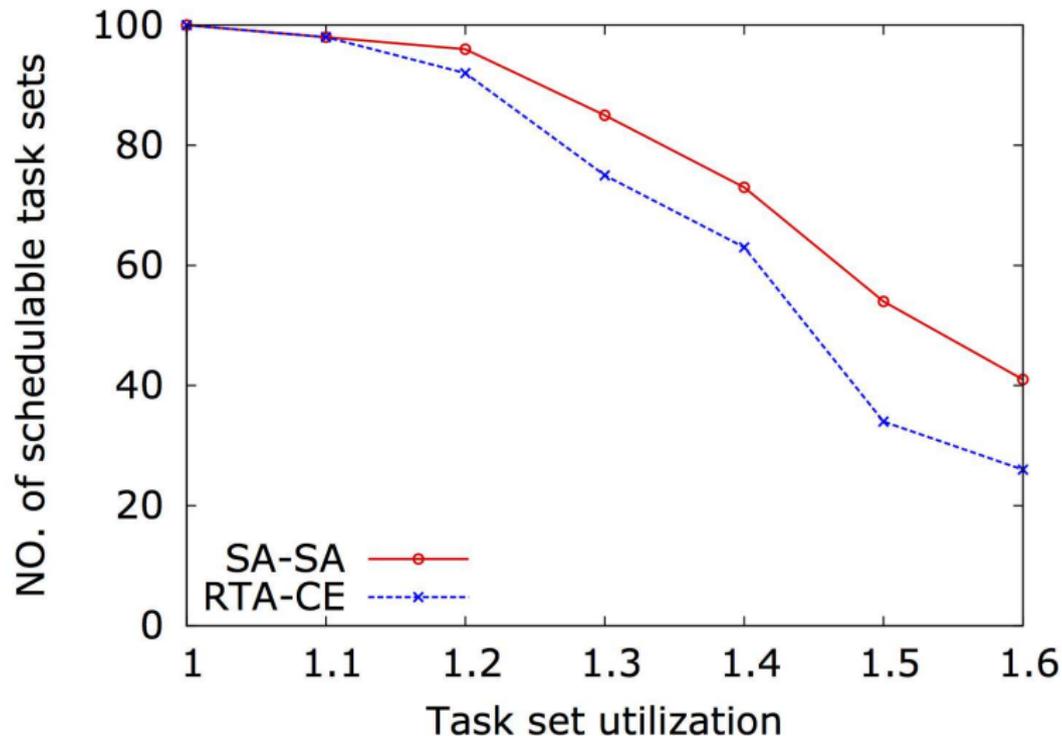
```
1:  $R \leftarrow \{S_0\}$ 
2: while true do
3:    $P \leftarrow \text{Post}(R)$ 
4:   if  $P \cap F \neq \emptyset$  then
5:     return NOT schedulable
6:   end if
7:    $R' \leftarrow R \cup P$ 
8:    $R' \leftarrow \text{Max}^{\preceq}(R')$ 
9:   if  $R' = R$  then
10:    return schedulable
11: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

- 1 Introduction
- 2 LHA models for multiprocessor G-FP scheduling
- 3 Weak Simulation Relation
 - Concrete state space
 - Symbolic state space
- 4 Experiments

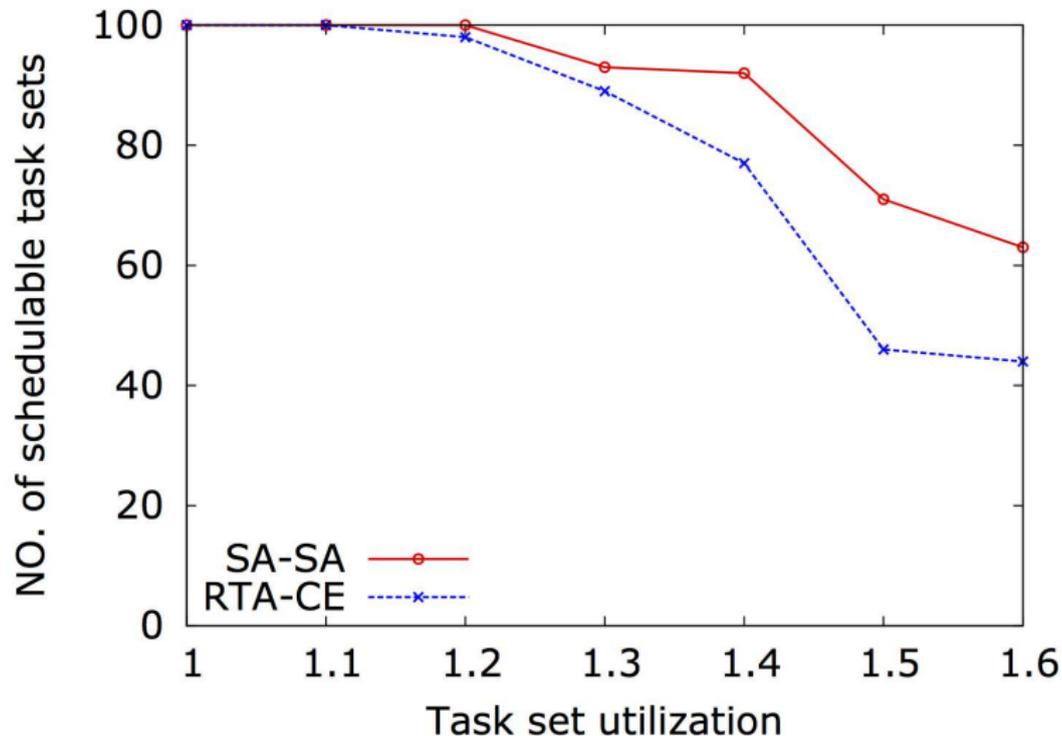
- RTA-CE : Response Time Analysis with Carry-in Enumeration
([Sun et al., 2014](#))

- RTA-CE : Response Time Analysis with Carry-in Enumeration (Sun et al., 2014)
- $m \in \{2, 3\}$ and $n = 5$
- $T_i \in [100, 1000]$ and a series of taskset utilisation levels U seperated by 0.1
- For each (m, n, U) configuration 100 tasksets are generated by Randomfixedsum (Emberson et al., 2010)

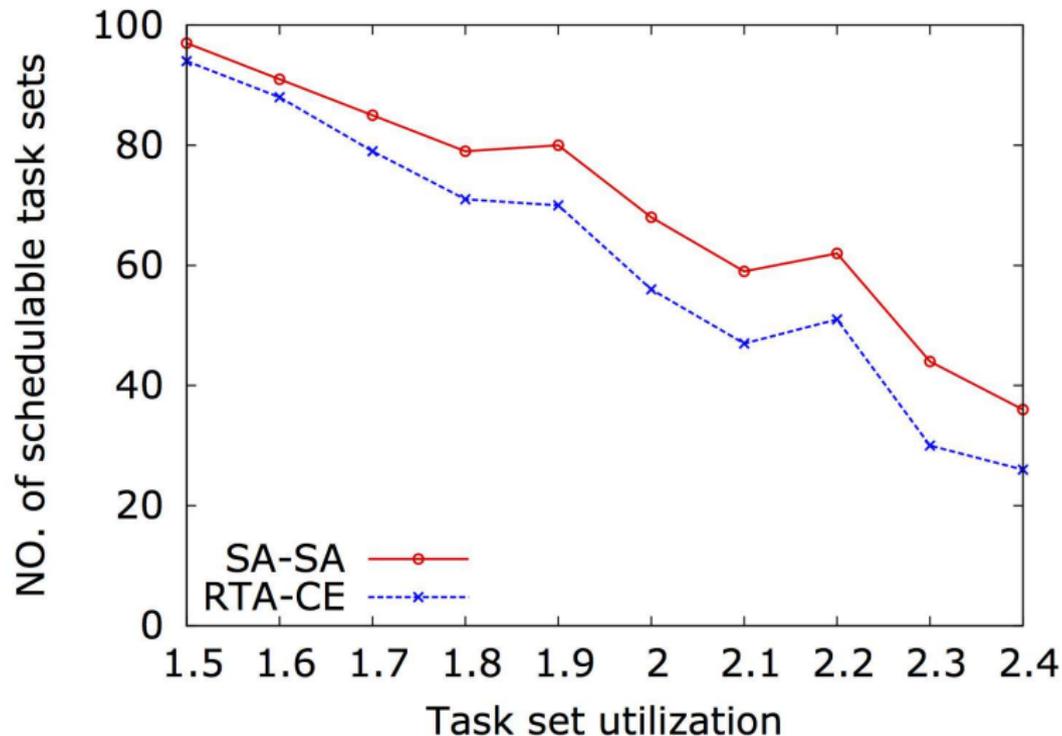
Results I: $m = 2, n = 5, \frac{D_i}{T_i} \in [0.8, 1]$



Results II: $m = 2, n = 5, \frac{D_i}{T_i} \in [0.8, 1.2]$

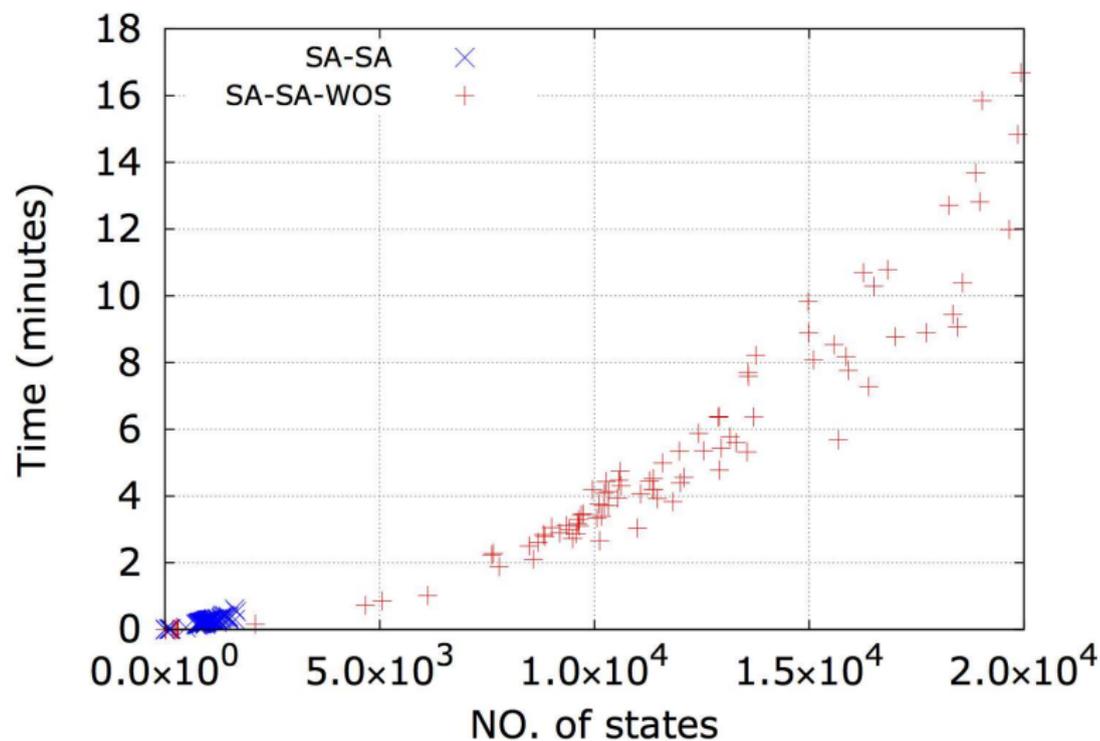


Results IV: $m = 3, n = 5, \frac{D_i}{T_i} \in [0.8, 1]$



- SA-SA-WOS: SA-SA WithOut Simulation
- $m = 2, n = 5, \frac{D_i}{T_i} \in [0.8, 1],$ and $U \in [1, 1.6]$
- 100 task sets

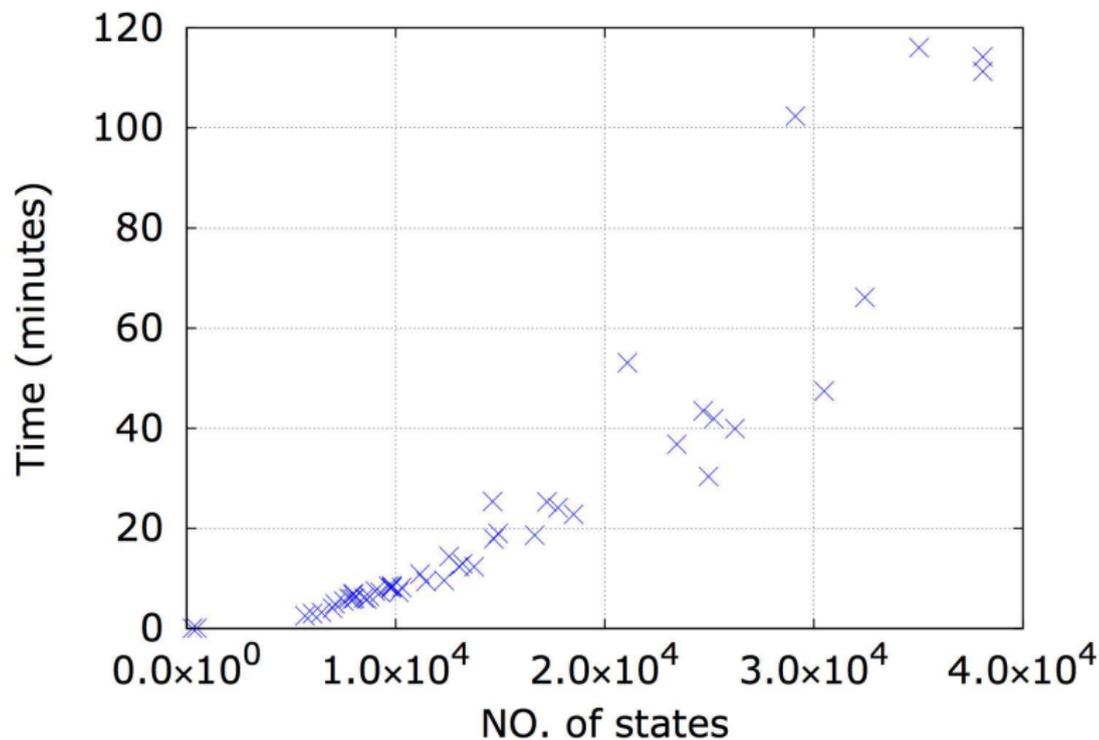
Results



SA-SA: another experiment on complexity

- $m = 2, n = 6, \frac{D_i}{T_i} \in [0.8, 2],$ and $U \in [1, 2]$
- 50 task sets

Results



Questions?