Simulation of Real-Time Scheduling with Various Execution Time Models

Maxime Chéramy, Pierre-Emmanuel Hladik,
Anne-Marie Déplanche and Sébastien Dubé

maxime.cheramy@laas.fr

June 19, 2014

9th IEEE International Symposium on Industrial Embedded Systems (SIES)

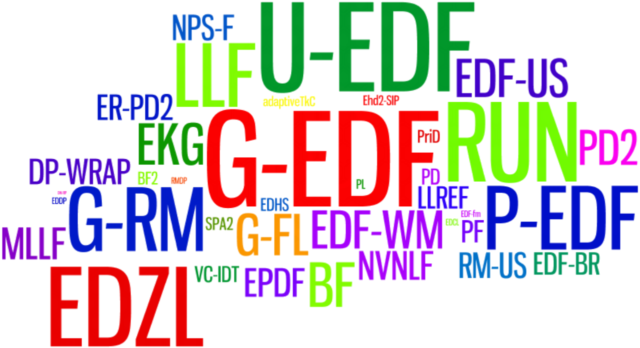
NPS-F
U-EDF
LLF
EDF-US
ER-PD2
adaptiveTkC
Ehd2-SIP
EKG
G-EDF
PriD
RUN
PD2
DP-WRAP
BF2
RMDP
PL
PD
EDDP
G-RM
EDHS
LLREF
P-EDF
SPA2
EDF-fm
MLLF
G-FL
EDF-WM
PF
NVNLF
RM-US
EDF-BR
EDZL
VC-IDT
EPDF
BF

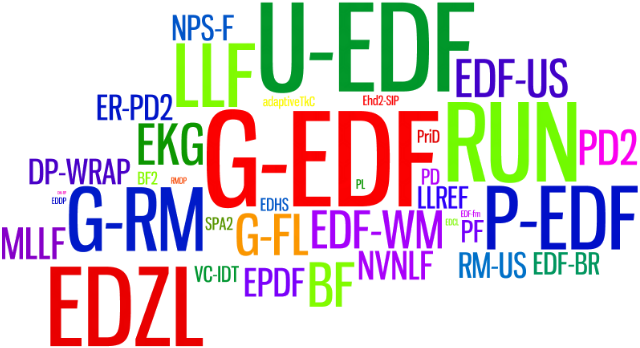

Requirement: evaluate their behavior and performance.

Many Real-Time Multiprocessor Scheduling Algorithms

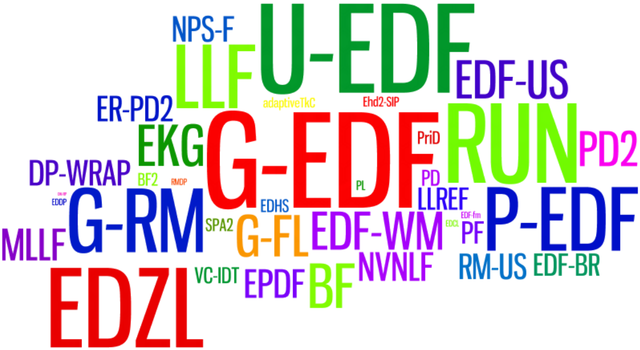

Requirement: evaluate their behavior and performance.

Ojective: propose a friendly and dedicated simulation-based tool.

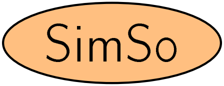
SimSo

SimSo: a Multiprocessor Real-Time Scheduling Simulator

Open Source
Easy to use

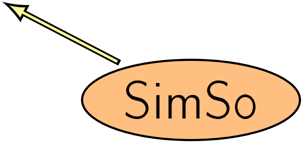

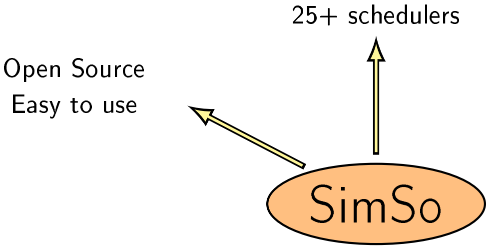
25+ schedulers
Open Source
Easy to use
SimSo

SimSo: a Multiprocessor Real-Time Scheduling Simulator

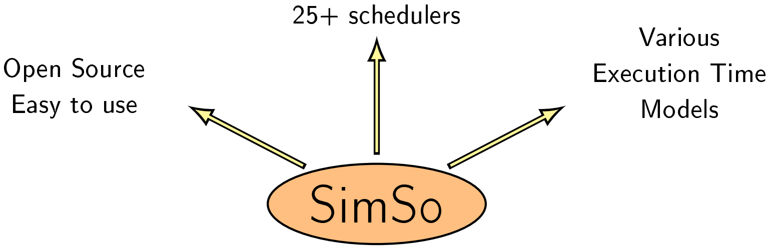

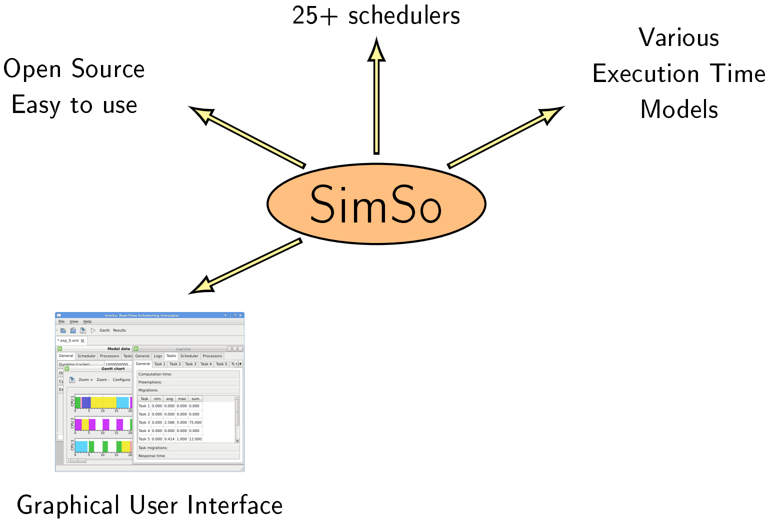
25+ schedulers
Open Source
Easy to use
Various
Execution Time
Models
SimSo
Graphical User Interface

25+ schedulers

Open Source
Easy to use

Various
Execution Time
Models

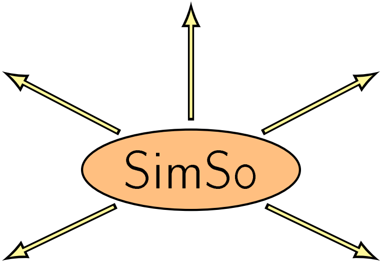

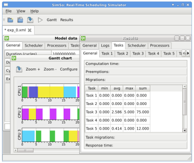

Graphical User Interface

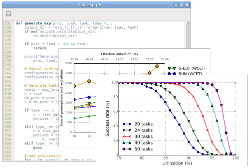

Script Mode

Execution Time Model 1: Using the WCET vs the ACET

Using the WCET is pessimistic

- Upper-bound rarely reached in practice
- Unlikely that all the jobs meet their WCET at the same time

Using the WCET is pessimistic

Upper-bound rarely reached in practice

Unlikely that all the jobs meet their WCET at the same time

It makes sense for schedulability analyses.

Execution Time Model 1: Using the WCET vs the ACET

Using the WCET is pessimistic

- Upper-bound rarely reached in practice
- Unlikely that all the jobs meet their WCET at the same time

It makes sense for schedulability analyses.

But we discuss its use for empirical studies focusing on other aspects, including the number of preemptions and migrations.

Execution Time Model 1: Using the WCET vs the ACET

We ran several simulations using the WCET and a random duration lower than the WCET.

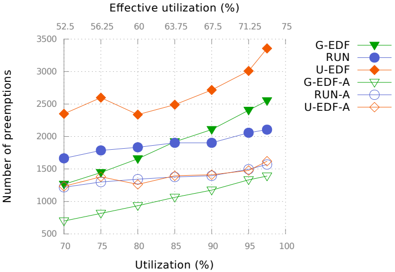

We ran several simulations using the WCET and a random duration lower than the WCET.

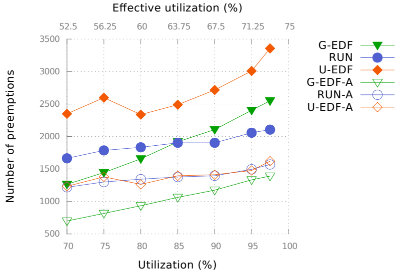

Using the WCET could introduce a bias in the evaluation

Gives an advantage to schedulers that use the WCET

Some schedulers could benefit from shorter execution times

Some schedulers are robust to punctual overloads

Execution Time Model 2: Simulating Cache Effects

Caches may have a significant impact on the computation time of the jobs, which depends on the scheduling decisions.

Execution Time Model 2: Simulating Cache Effects

Caches may have a significant impact on the computation time of the jobs, which depends on the scheduling decisions.

There are **cache-aware schedulers** that require to take into consideration these cache effects.

Execution Time Model 2: Simulating Cache Effects

Caches may have a significant impact on the computation time of the jobs, which depends on the scheduling decisions.

There are **cache-aware schedulers** that require to take into consideration these cache effects.

$\Rightarrow$ As a consequence, we want to simulate the cache effects.

Execution Time Model 2: Simulating Cache Effects

Extension of the Liu&Layland task model with:

- SDP: Stack Distance Profile[a]
- MIX / API: ratio of memory access per instruction
- CPI: Cycles per Instruction
- Number of executed instructions

[a] R.L. Mattson et al. "Evaluation techniques for storage hierarchies". In: *IBM Systems Journal* 9.2 (1970).

Execution Time Model 2: Simulating Cache Effects

Extension of the Liu&Layland task model with:

- SDP: Stack Distance Profile[a]
- MIX / API: ratio of memory access per instruction
- CPI: Cycles per Instruction
- Number of executed instructions

[a] R.L. Mattson et al. "Evaluation techniques for storage hierarchies". In: *IBM Systems Journal* 9.2 (1970).

Integration of statistical cache models[a] to determine the duration of the jobs dynamically, during the simulation.

[a] D. Chandra et al. "Predicting inter-thread cache contention on a chip multi-processor architecture". In: *Proc. of HPCA*. 2005.

Poster

Thank you for your attention.

Any question?

→ Come to the poster session,
I'd be pleased to discuss this further with you.

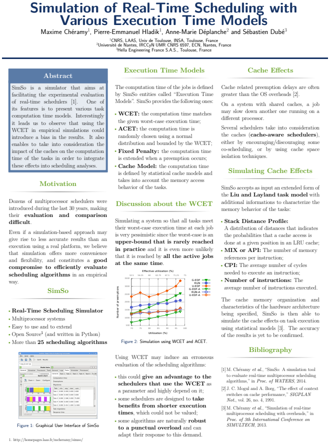

Simulation of Real-Time Scheduling with Various Execution Time Models

Maxime Chéramy[1], Pierre-Emmanuel Hladik[1], Anne-Marie Déplanche[2] and Sébastien Dubé[3]

[1]CNRS, LAAS, Univ de Toulouse, INSA, Toulouse, France
[2]Université de Nantes, IRCCyN UMR CNRS 6597, ECN, Nantes, France
[3]Hella Engineering France S.A.S., Toulouse, France

Abstract

SimSo is a simulator that aims at facilitating the experimental evaluation of real-time schedulers [1]. One of its features is to present various task computation time models. Interestingly it leads us to observe that using the WCET in empirical simulations could introduce a bias in the results. It also enables to take into consideration the impact of the caches on the computation time of the tasks in order to integrate these effects into scheduling analyses.

Motivation

Dozens of multiprocessor schedulers were introduced during the last 30 years, making their **evaluation and comparison difficult**.

Even if a simulation-based approach may give rise to less accurate results than an execution using a real platform, we believe that simulation offers more convenience and flexibility, and constitutes a **good compromise to efficiently evaluate scheduling algorithms** in an empirical way.

SimSo

- **Real-Time Scheduling Simulator**
- Multiprocessor systems
- Easy to use and to extend
- Open Source[1] (and written in Python)
- More than **25 scheduling algorithms**

Figure 1: Graphical User Interface of SimSo

1. http://homepages.laas.fr/mcheramy/simso/

Execution Time Models

The computation time of the jobs is defined by SimSo entities called "Execution Time Models". SimSo provides the following ones:

- **WCET:** the computation time matches the given worst-case execution time;
- **ACET:** the computation time is randomly chosen using a normal distribution and bounded by the WCET;
- **Fixed Penalty:** the computation time is extended when a preemption occurs;
- **Cache Model:** the computation time is defined by statistical cache models and takes into account the memory access behavior of the tasks.

Discussion about the WCET

Simulating a system so that all tasks meet their worst-case execution time at each job is very pessimistic since the worst-case is an **upper-bound that is rarely reached in practice** and it is even more unlikely that it is reached by **all the active jobs at the same time**.

Figure 2: Simulation using WCET and ACET.

Using WCET may induce an erroneous evaluation of the scheduling algorithms:

- this could **give an advantage to the schedulers that use the WCET** as a parameter and highly depend on it;
- some schedulers are designed to **take benefits from shorter execution times**, which could not be valued;
- some algorithms are naturally **robust to a punctual overload** and can adapt their response to this demand.

Cache Effects

Cache related preemption delays are often greater than the OS overheads [2].

On a system with shared caches, a job may slow down another one running on a different processor.

Several schedulers take into consideration the caches (**cache-aware schedulers**), either by encouraging/discouraging some co-scheduling, or by using cache space isolation techniques.

Simulating Cache Effects

SimSo accepts as input an extended form of the **Liu and Layland task model** with additional informations to characterize the memory behavior of the tasks:

- **Stack Distance Profile:** A distribution of distances that indicates the probabilities that a cache access is done at a given position in an LRU cache;
- **MIX or API:** The number of memory references per instruction;
- **CPI:** The average number of cycles needed to execute an instruction;
- **Number of instructions:** The average number of instructions executed.

The cache memory organization and characteristics of the hardware architecture being specified, SimSo is then able to simulate the cache effects on task execution using statistical models [3]. The accuracy of the results is yet to be confirmed.

Bibliography

[1] M. Chéramy *et al.*, "SimSo: A simulation tool to evaluate real-time multiprocessor scheduling algorithms," in *Proc. of WATERS*, 2014.

[2] J. C. Mogul and A. Borg, "The effect of context switches on cache performance," *SIGPLAN Not.*, vol. 26, no. 4, 1991.

[3] M. Chéramy *et al.*, "Simulation of real-time multiprocessor scheduling with overheads," in *Proc. of 3th International Conference on SIMULTECH*, 2013.