

The Global Limited Preemptive Earliest Deadline First Feasibility of Sporadic Real-time Tasks

Abhilash Thekkilakattil, Sanjoy Baruah, Radu Dobrin and
Sasikumar Punnekkat



THE UNIVERSITY
of NORTH CAROLINA
at CHAPEL HILL



Vetenskapsrådet

Motivation

- Multi (-core) processors in real-time systems complicate the problems associated with fully preemptive schedulers
 - Complex hardware, e.g., different levels of caches
 - **Difficult** to perform **timing analysis**
 - Potentially large number of task migrations: implementation issues
 - **Difficult** to demonstrate **predictability**
 - **Difficult** to reason about **safety**
- Non-preemptive scheduling can be infeasible at arbitrarily small utilization
 - Long task problem: at least one task has execution time **greater** than the shortest deadline


Solution: Limit preemptions



Advantages of limiting preemptions

Combines **best of** preemptive and non-preemptive scheduling

- Control preemption related overheads
 - Context switch costs, cache related preemption delays, pipeline delays and bus contention costs
- Improve processor utilization
 - **Reduce** preemption related **costs** while **eliminating** infeasibility due to **blocking**



Anecdotal evidence: “*limiting preemptions improves safety and makes it easier to certify software for safety-critical applications*”





Limited preemptive scheduling landscape

Uniprocessor	Limited preemptive FPS (Yao et al., RTSJ'11)	Limited preemptive EDF (Baruah, ECRTS'05)
Multiprocessor	Global limited preemptive FPS (Marinho et al., RTSS'13)	?



G-LP-EDF scheduling model



Main contributions

1. Schedulability analysis for **Global Limited Preemptive Earliest Deadline First** (G-LP-EDF) scheduling of **sporadic** real-time tasks
2. Analysis of the effects of increasing the **processor speed** on G-LP-EDF feasibility



Methodology overview

A ~~necessary~~ **sufficient** schedulability condition:

Upper-bound on the work
generated under G-LP-EDF



Lower-bound on the work executed
under any work conserving algorithm



Methodology overview

A sufficient schedulability condition:


Upper-bound on the work
generated under G-LP-EDF

\leq


Lower-bound on the work executed
under any work conserving algorithm




Lower bound on the work done



$$\text{work done}(t_i, t_{i-1}) \geq m(t_{i-1} - t_i - x_i) + x_i$$

 J_i 's execution


 High priority job executions

 Low priority job executions (blocking)



Lower bound on the work done

Unschedulability scenario:



$$\text{work done}(t_i, t_{i-1}) \geq m(t_{i-1} - t_i - x_i) + x_i$$

$$\text{work done}(t_k, t_0) > (m - (m-1)s)(t_0 - t_k - L) + mL$$



Low priority job executions (blocking)



High priority job executions



Methodology overview

A sufficient schedulability condition:

Upper-bound on the work
generated under G-LP-EDF

\leq

Lower-bound on the work executed
under any work conserving algorithm



Upper bound on the work generated


..... under G-LP-EDF

In $[t_k, t_0)$, we consider the duration for which:

- a. Low priority tasks **block** high priority tasks
- b. **Higher priority** tasks execute



Maximum duration of blocking



$$\text{blocking}(t_k, t_0) : mL$$



Low priority job executions (blocking)




High priority job executions



Upper bound on the work done by task i

$$\text{Work done}(t_k, t_0) \leq \frac{1}{s} FF - DBF_i(t_0 - t_k, s) + mL$$



$FF - DBF_i(t, s) =$
(Forced-Forward Demand Bound Function)

$head_i$

+

$tail_i$

$$\hat{e} \hat{t} \hat{u} C_i$$

$$\hat{e} \hat{T}_i \hat{u}$$

$$R^3 D_i$$

$$C_i$$

$$D - \frac{C_i}{s} \in R < D_i$$

$$C_i - (D_i - R)s$$

otherwise

$$0$$



Methodology overview

A **sufficient schedulability** condition:

Upper-bound on the work
generated under G-LP-EDF



\leq

Lower-bound on the work executed
under any work conserving algorithm




$$\hat{a}_{FF-DBF_i}(t_0 - t_k, s) + mL \leq (m - (m-1)s)(t_0 - t_k - L) + mL$$

$$\hat{a}_{FF-DBF_i}(t_0 - t_k, s) \leq (m - (m-1)s)(t_0 - t_k - L)$$




Agenda

1. Schedulability analysis for **Global Limited Preemptive Earliest Deadline First** (G-LP-EDF) scheduling of **sporadic** real-time tasks
2. Analysis of the effects of increasing the **processor speed** on G-LP-EDF feasibility




Agenda

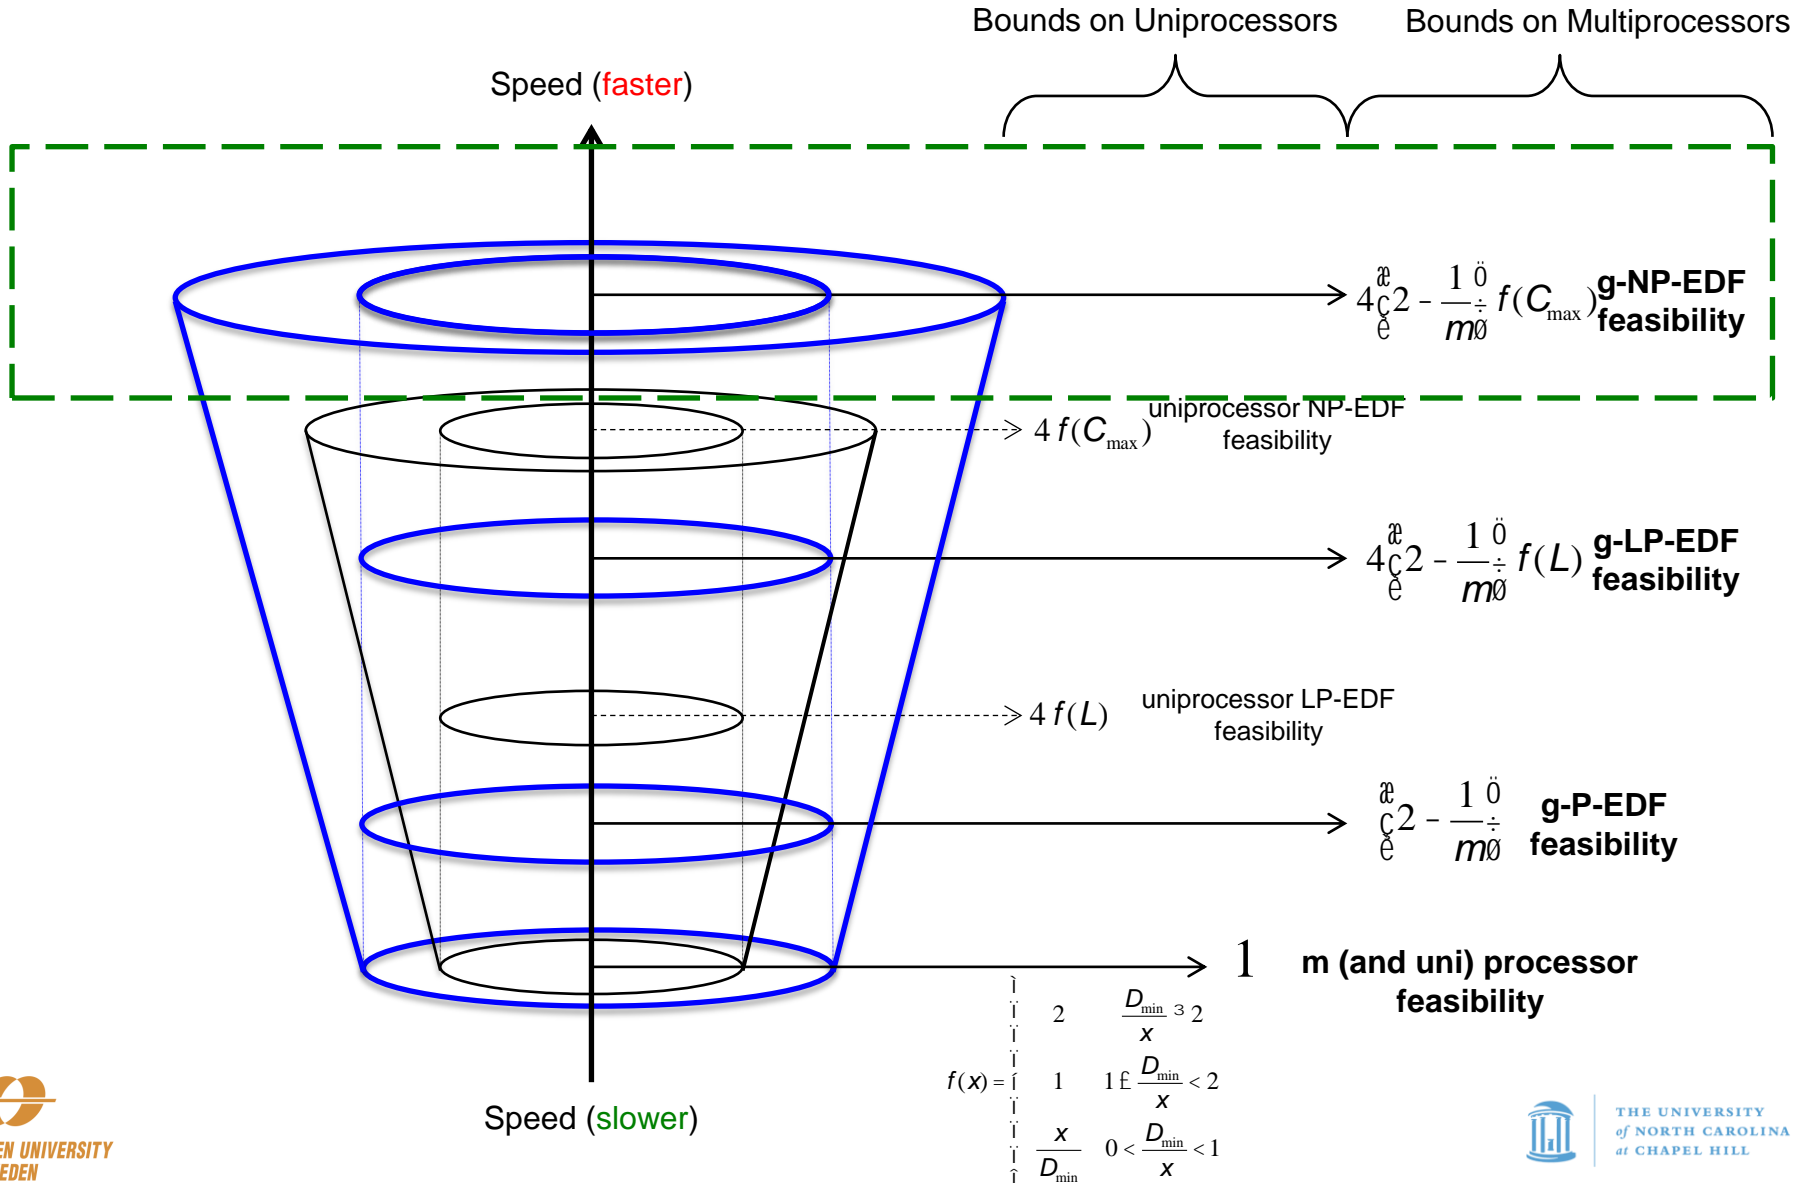
1. Schedulability analysis for **Global Limited Preemptive Earliest Deadline First** (G-LP-EDF) scheduling of **sporadic** real-time tasks
2. Analysis of the effects of increasing the **processor speed** on G-LP-EDF feasibility



Feasibility bucket




Processor speed vs. LP-EDF feasibility



Long task problem

Non-preemptive infeasibility arising from at least one task having **WCET greater than shortest deadline**





A solution: **code-refactoring** at task level

Our speed-up factor quantifies the **extent** to which **code-refactoring** must be done to **enable** non-preemptive feasibility

Agenda

1. Schedulability analysis for **Global Limited Preemptive Earliest Deadline First** (G-LP-EDF) scheduling of **sporadic** real-time tasks
2. Analyze the effects of increasing the **processor speed** on G-LP-EDF feasibility



Conclusions

- Global limited preemptive EDF feasibility analysis
 - To **control** preemption related overheads
 - Enables better reasoning about **predictability** of multi (-core) processor real-time systems
- Processor speed vs. preemptive behavior
 - Quantifies the extent to which **code-refactoring** must be performed to address the **long task problem**
 - **Sub-optimality** of G-NP-EDF



Future work

- Compare G-LP-EDF and G-P-EDF in presence of overheads
- Perform **trade-offs**: number of extra processors vs. speed-up
- Partition tasks comprising of **non-preemptive chunks**
- Accounting for suspensions



Thank you !



Questions ?

