

An Empirical Investigation of Eager and Lazy Preemption Approaches in Global Limited Preemptive Scheduling

Abhilash Thekkilakattil, Kaiqian Zhu, Yonggao Nie, Radu Dobrin and Sasikumar Punnekkat



MÄLARDALEN UNIVERSITY
SWEDEN



Vetenskapsrådet

We do Scheduling Everyday!



15 mins




20 mins



8 hrs




1 hr

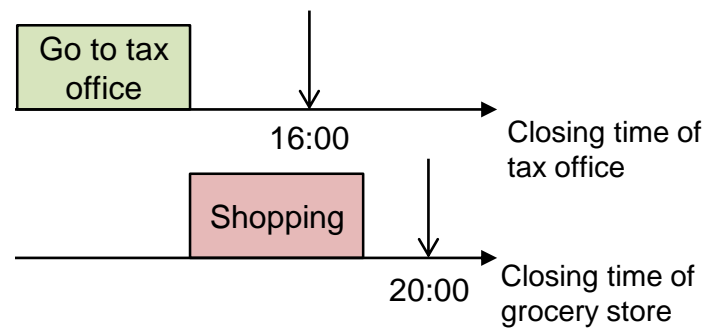


Priority based scheduling

Call from school: sick kid



Deadline based scheduling




Real-time Systems

Correctness of the system depends on both.





Correctness of the decision to inflate the airbags



2. Correctness of the time at which the decision is implemented

Computer programs called "tasks"

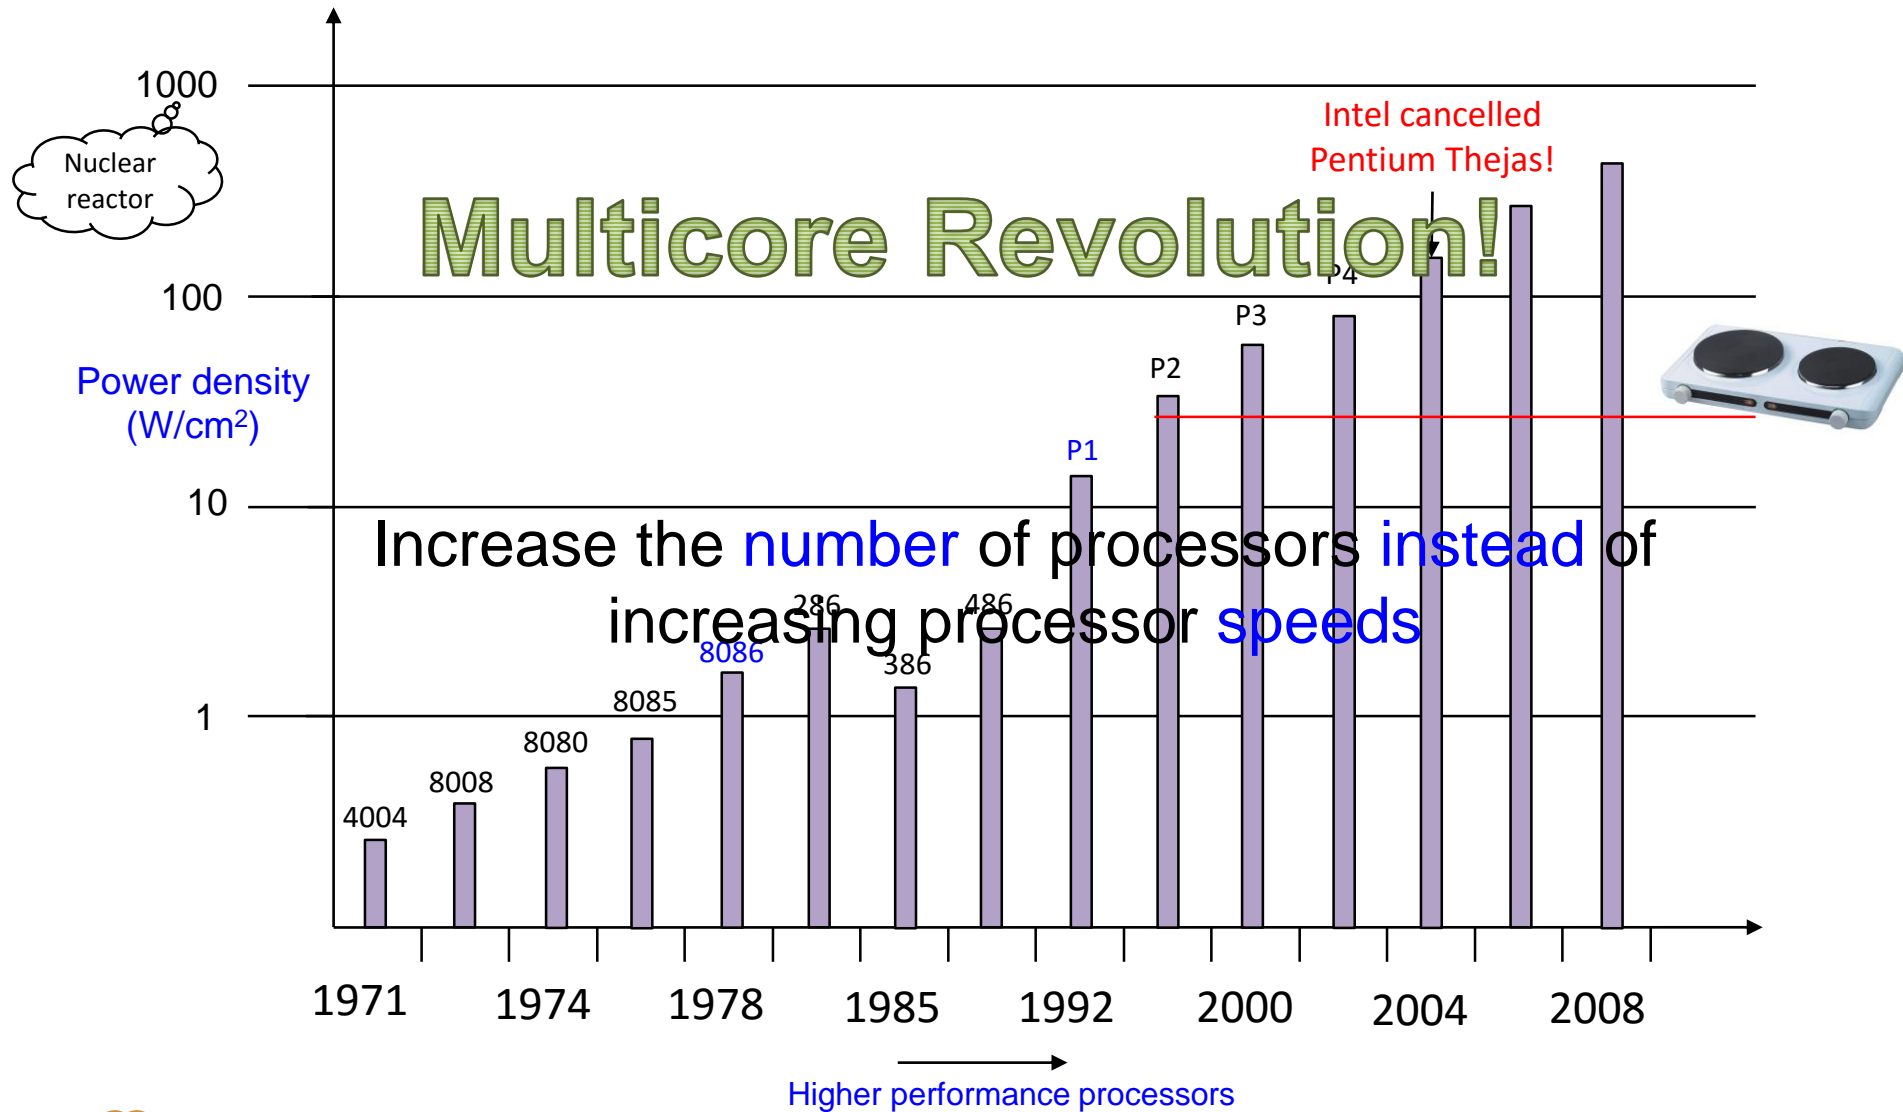


Real-time Scheduling

- Fewer computers than the number of tasks!
- Need to schedule the tasks such that deadlines are not missed
- For example tasks are typically executed highest priority first or earliest deadline first
- Schedulability test determines if a taskset meets deadlines under a given scheduling algorithm




Faster Computers vs. Power Demand



(adapted from Marko Bertogna's PhD thesis)




Processor to Memory Performance Gap

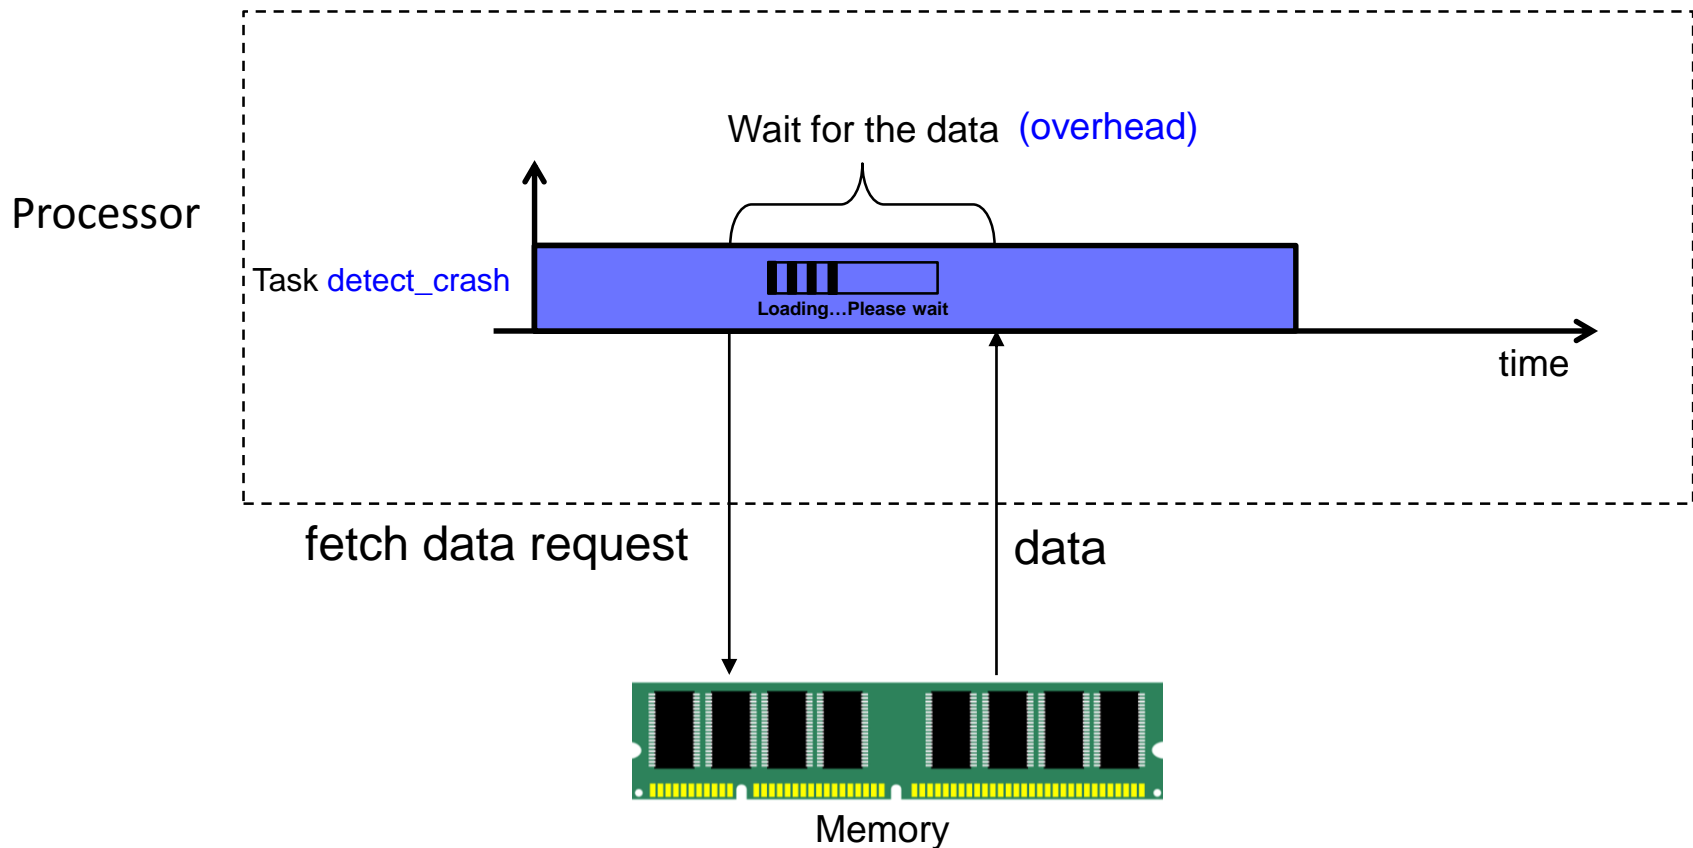


Hardware features to bridge this performance gap
e.g., caches

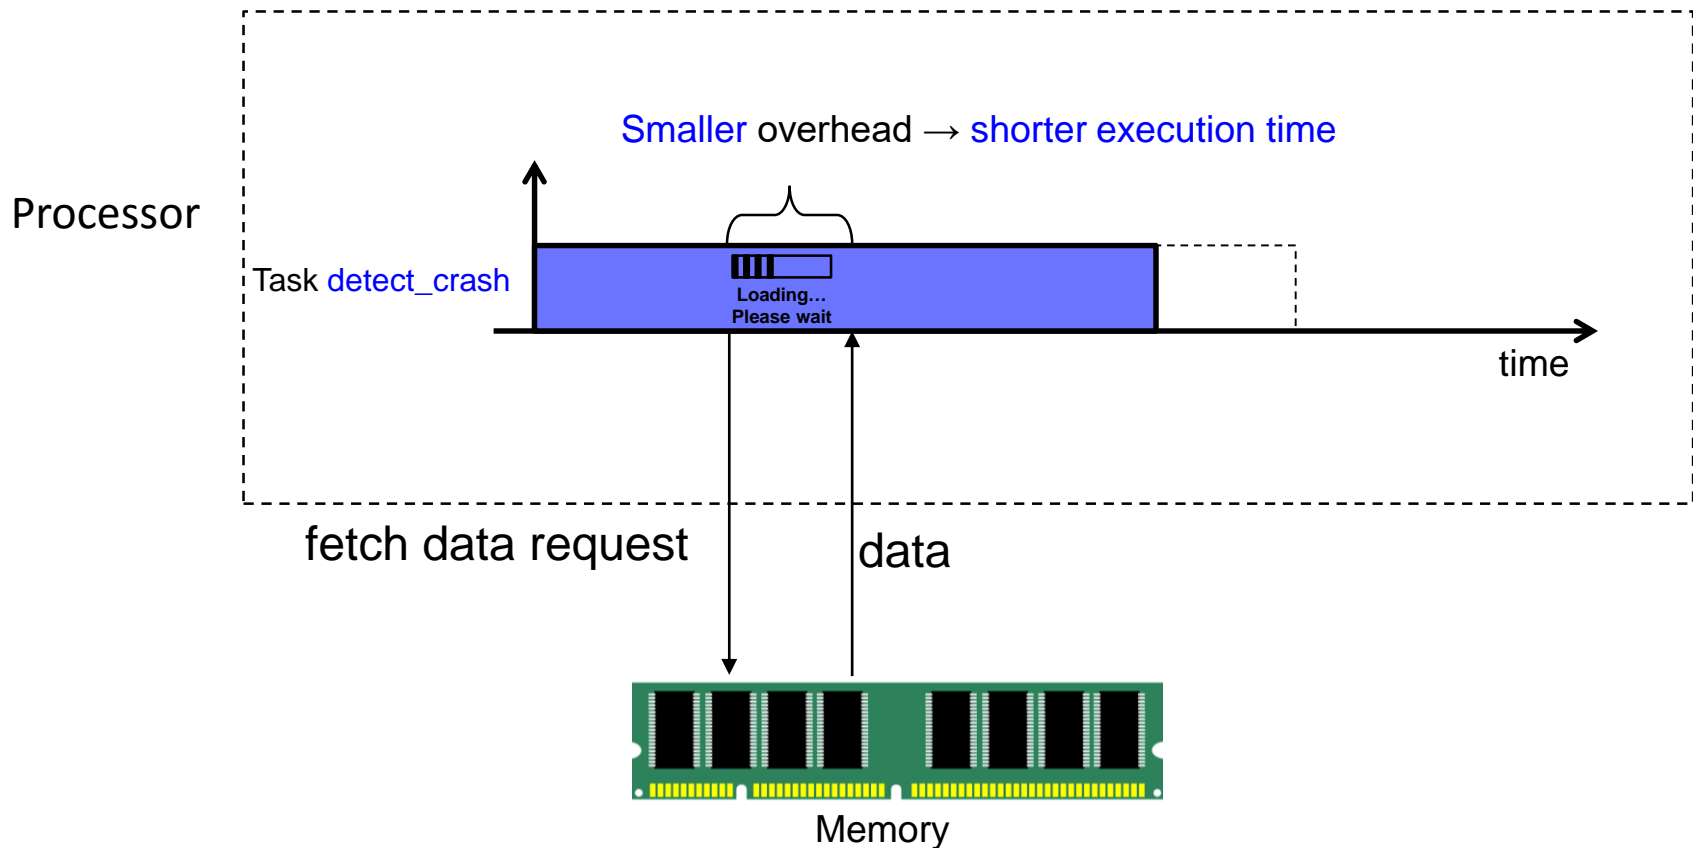
Predictability challenge!



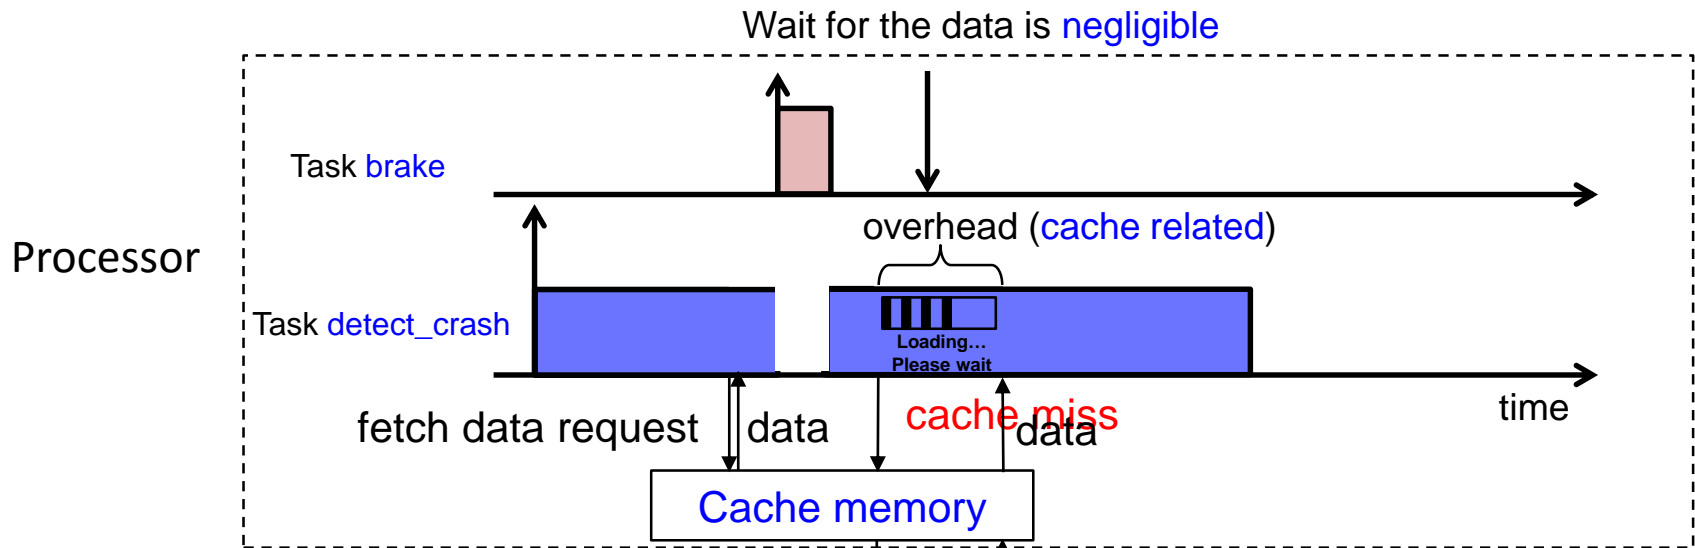
Processor to Memory Performance Gap




Processor to Memory Performance Gap



Bridging the Processor to Memory Performance Gap



Alternative 1: build predictable hardware e.g., PRETS



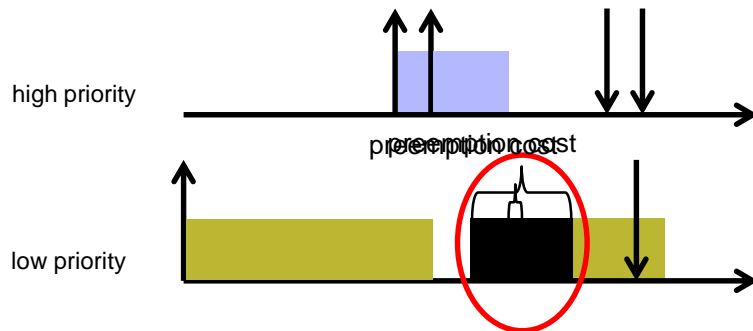
Memory

Alternative 2: efficient resource allocation & overhead accounting

Motivation


Preemptive Scheduling

- ☺ Zero blocking: high schedulable utilization
- ☹ High runtime overhead: preemption costs
- ☹ Difficult to perform timing analysis
- ☹ Difficult to demonstrate predictability and safety: multicores



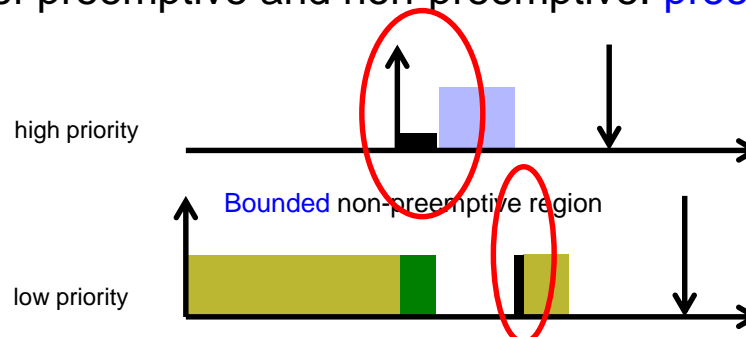
Non-Preemptive Scheduling

- ☺ Low runtime overhead: zero preemption costs
- ☺ Relatively easier to perform timing analysis
- ☺ Preferred by many safety standards
- ☹ Increased blocking: low schedulable utilization



Solution: Limited-Preemptive Scheduling

- ☺ Best of preemptive and non-preemptive: preempt only when necessary



Limited Preemptive Scheduling Models

- Preemption Threshold Scheduling (ThreadX, Wang and Saksena, 1999)

- Fixed Preemption Point Scheduling (Burns, 1994, Bril *et al.*, 2009, Yao *et al.*, 2011)

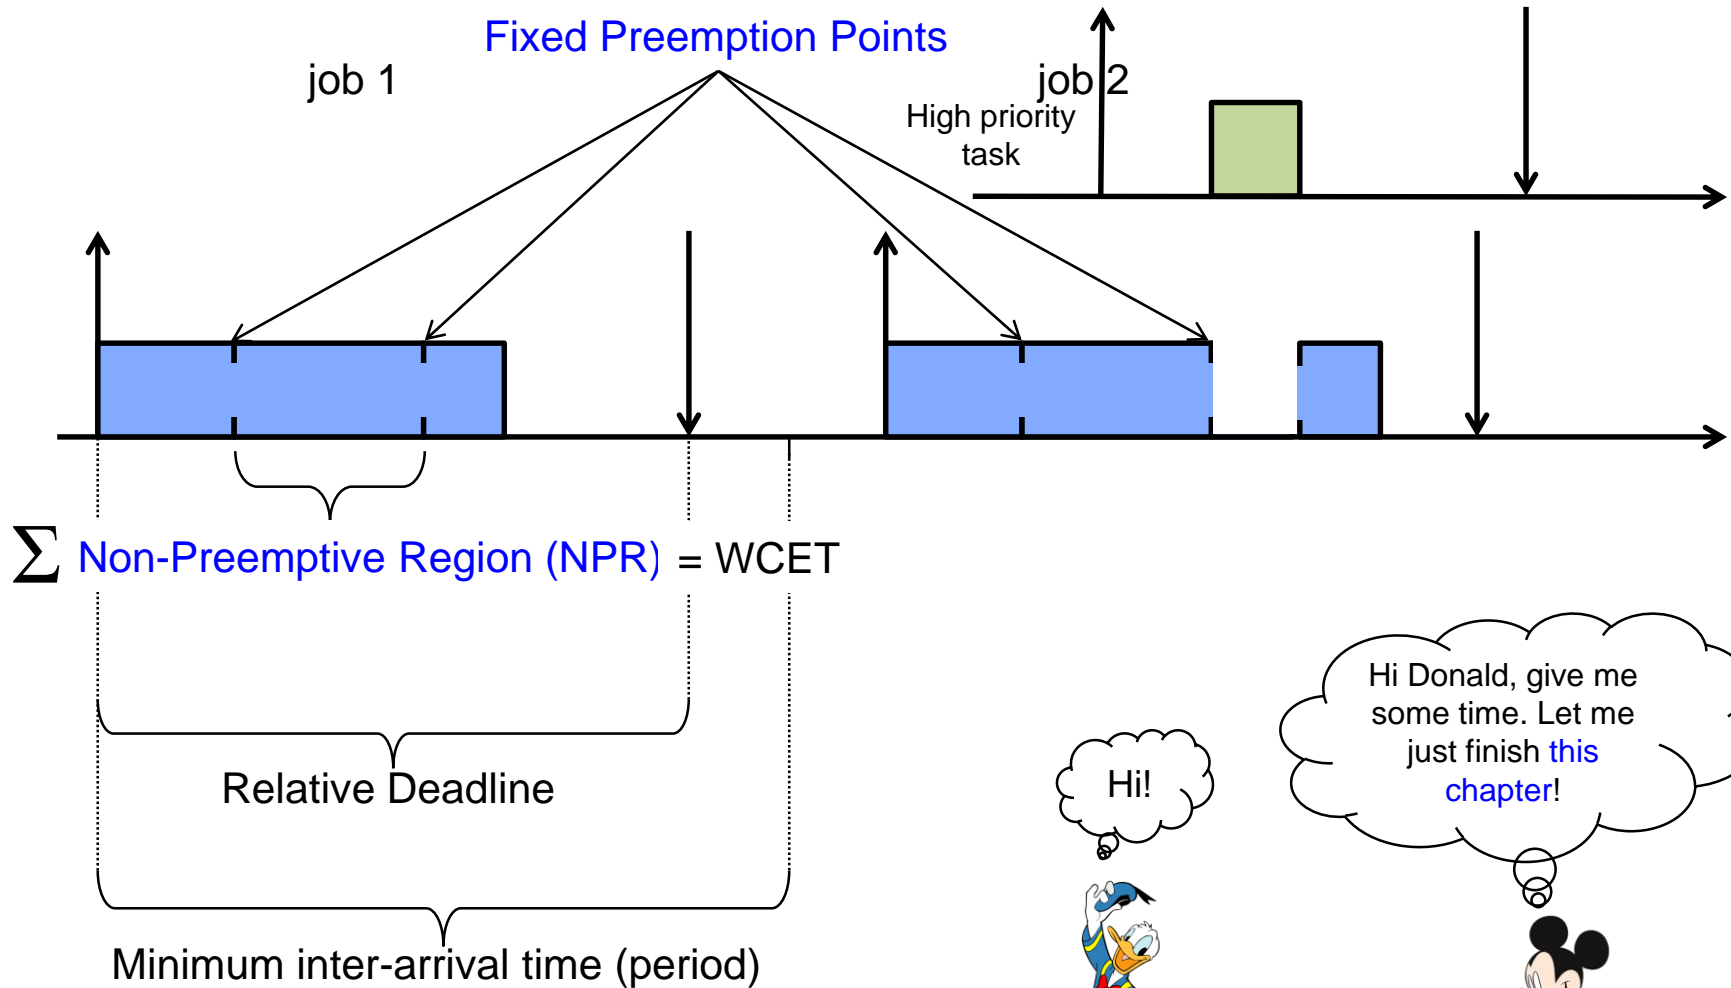
Our focus

- Floating Non-Preemptive Region Scheduling (Baruah, 2005)

- ...

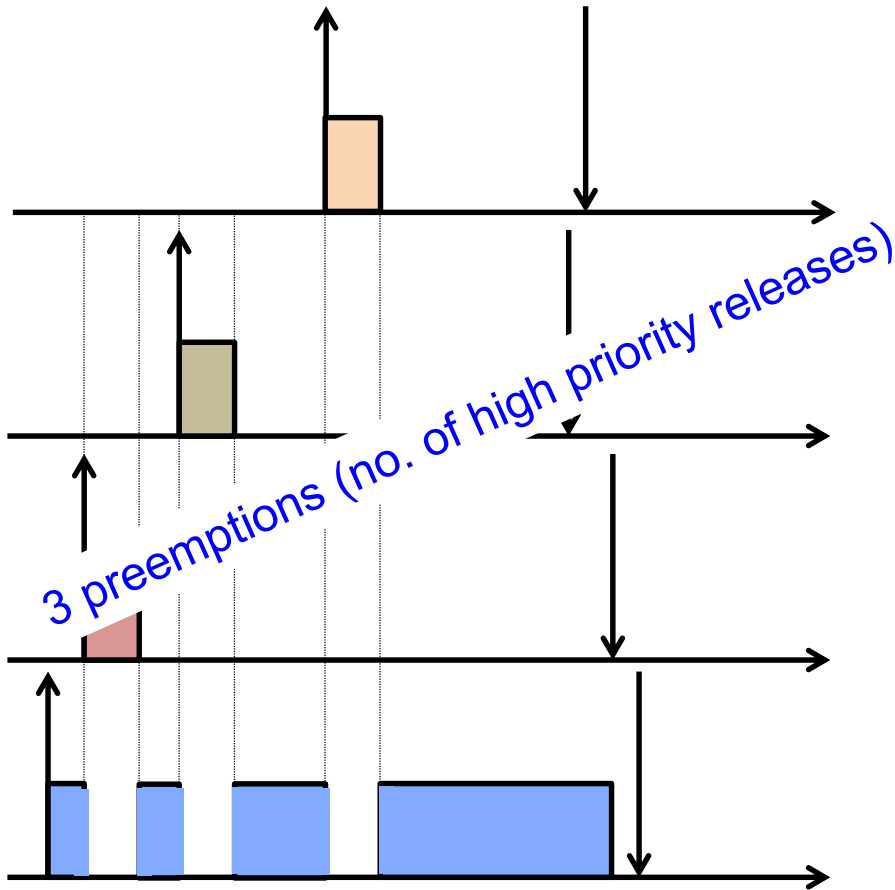


Fixed Preemption Points Scheduling

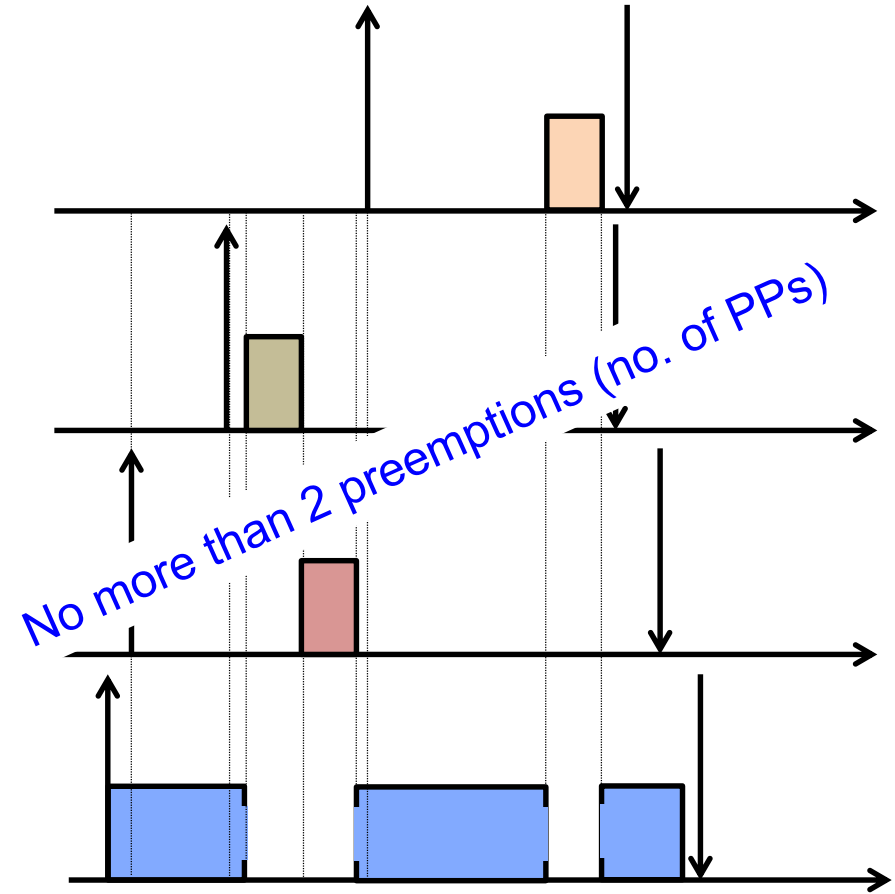


Reduction in Preemptions

Fully Preemptive Scheduling



Fixed Preemption Points Scheduling



Limited Preemptive Scheduling Landscape

Fixed priority based scheduling

Deadline based scheduling

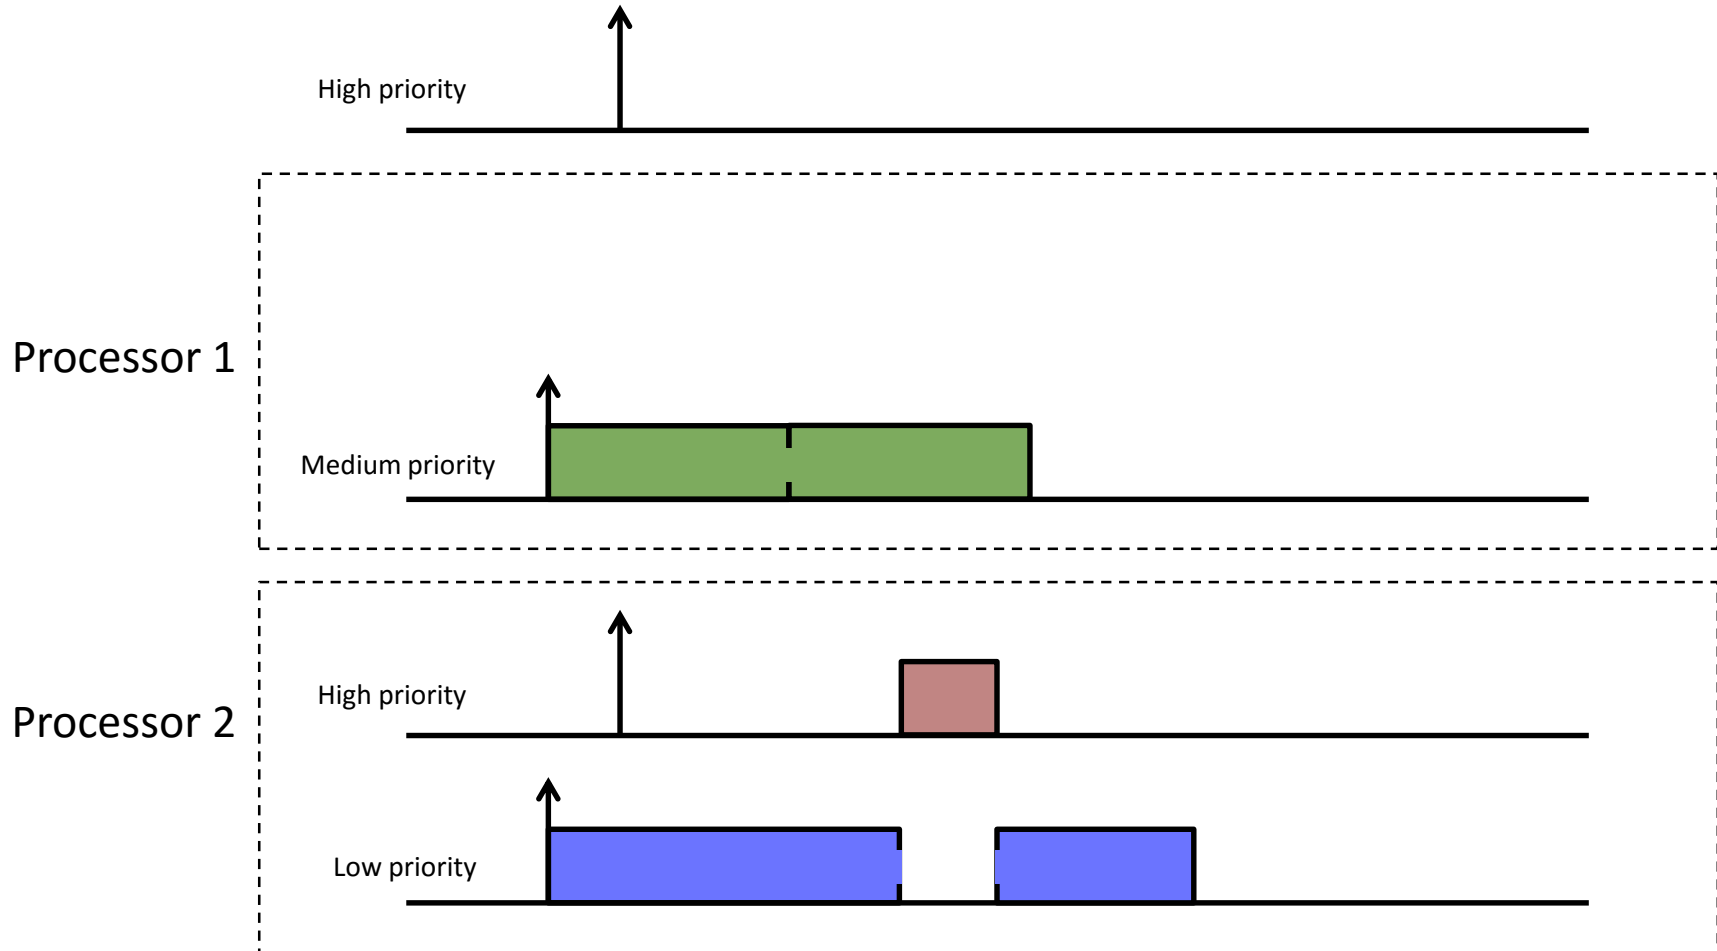
Uniprocessor	Limited Preemptive FPS	Limited Preemptive EDF
Multiprocessor	Global Limited Preemptive FPS	Global Limited Preemptive EDF

Handling preemptions in global LP scheduling:

1. **Eager** Preemption Approach
2. **Lazy** Preemption Approach

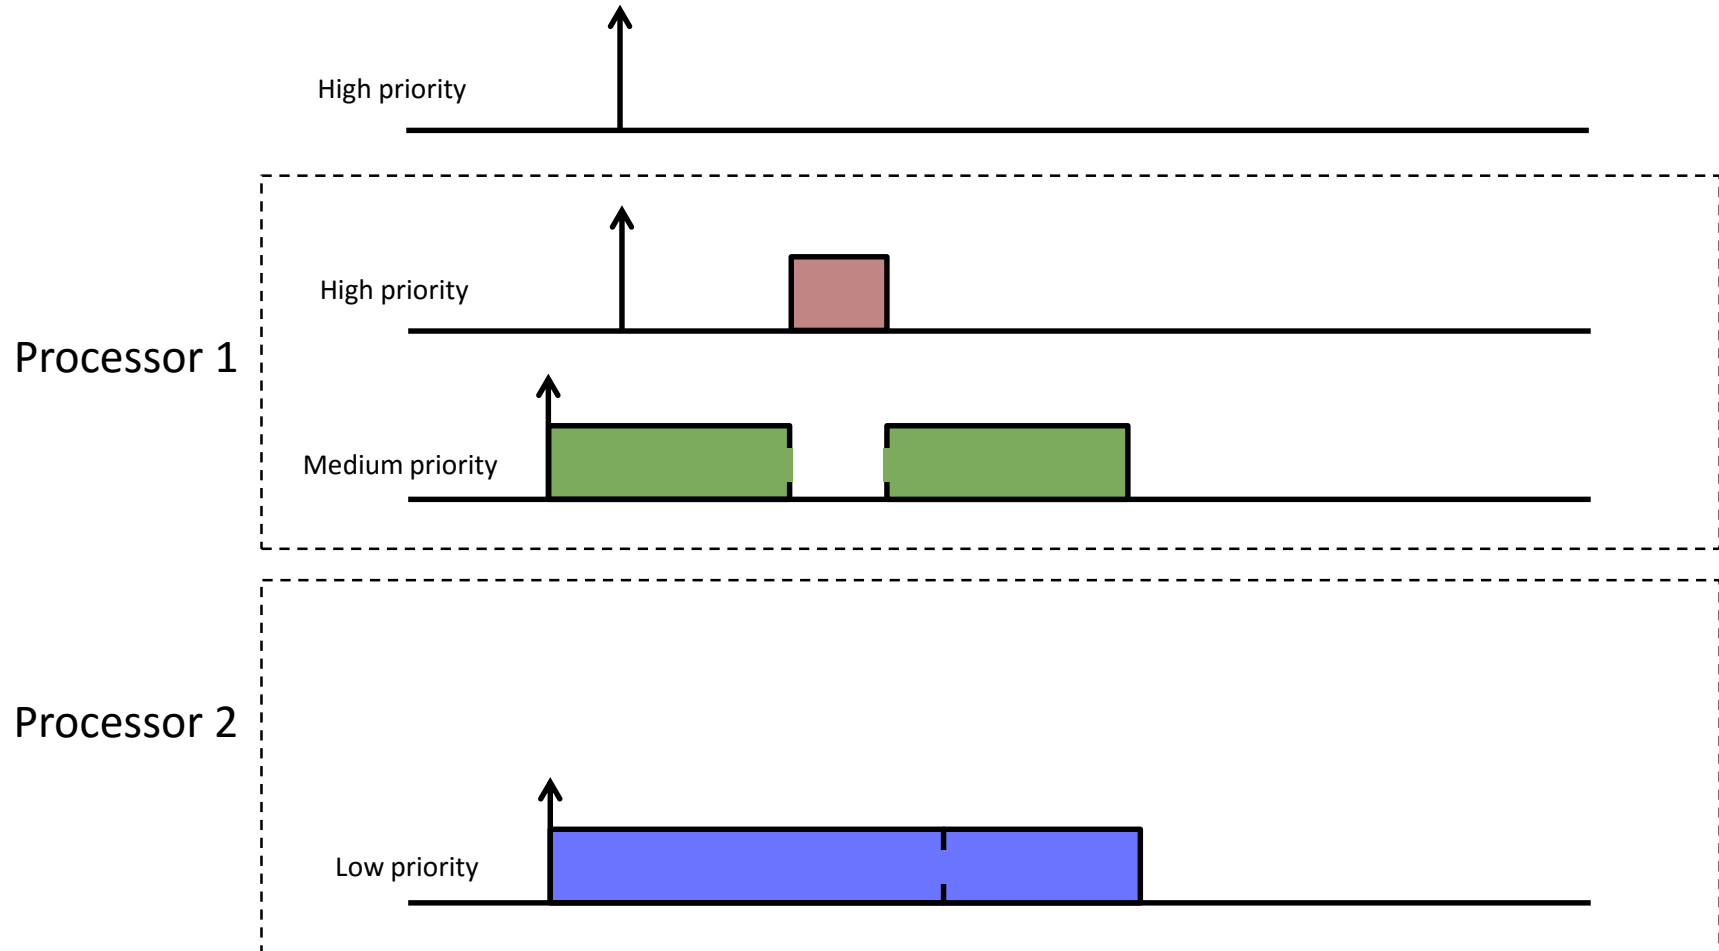
Global Limited Preemptive Scheduling

Lazy Preemption Approach: wait to preempt the lowest priority task



Global Limited Preemptive Scheduling

Eager Preemption Approach: preempt the **lower** priority task that **first finishes its NPR** (not necessarily the lowest)



Advances in the Limited Preemptive Scheduling Landscape

Uniprocessor	Limited preemptive FPS Burns'94, Bril <i>et al.</i> , RTSJ'09, Yao <i>et al.</i> , RTSJ'11	Limited preemptive EDF Baruah, ECRTS'05
Multiprocessor	Global Limited Preemptive FPS Marinho <i>et al.</i> , RTSS'13, Davis <i>et al.</i> , TECS'15, Thekkilakattil <i>et al.</i> , RTNS'15, Marinho, PhD Thesis'15, Serrano <i>et al.</i> , DATE'16	Global Limited Preemptive EDF Block <i>et al.</i> , RTCSA'07 Thekkilakattil <i>et al.</i> , ECRTS'14, Chattopadhyay <i>et al.</i> , RTNS'14, Marinho, PhD Thesis'15

Focus on schedulability!



Preemptive Behavior of FPS and EDF

... under the preemptive and **limited preemptive** paradigms
for **eager** and **lazy** approaches

Uniprocessors	Buttazzo, RTSJ'05 (RM vs EDF: Judgement Day)
Multiprocessors	No significant work!

How does the number of preemptions vary with the **scheduling algorithm**
and **approach to preemption** on **multiprocessors**?

Experimental Setup

- **Weighted metric** to count preemptions:

$$W(p) = \frac{\sum_{\forall \Gamma} U(\Gamma) P(\Gamma, p)}{\sum_{\forall \Gamma} U(\Gamma)}$$

- Enables investigation of number of preemptions *w.r.t* a **second parameter** in addition to utilization
- Tasksets generated using Uunifast-Discard
- Periods in the range 5-500
- Utilizations in the range 1 to $m/2$ (in one case upto m)
- Schedule simulated for 10000 time units



Weighted Metric to Count Preemptions

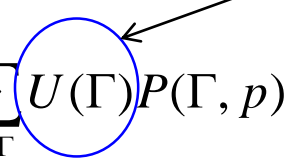
$$W(p) = \frac{\sum_{\forall \Gamma} U(\Gamma) P(\Gamma, p)}{\sum_{\forall \Gamma} U(\Gamma)}$$

Number of preemptions generated by taskset Γ w.r.t parameter p

Weighted Metric to Count Preemptions

$$W(p) = \frac{\sum_{\forall \Gamma} U(\Gamma) P(\Gamma, p)}{\sum_{\forall \Gamma} U(\Gamma)}$$

Utilization of taskset Γ



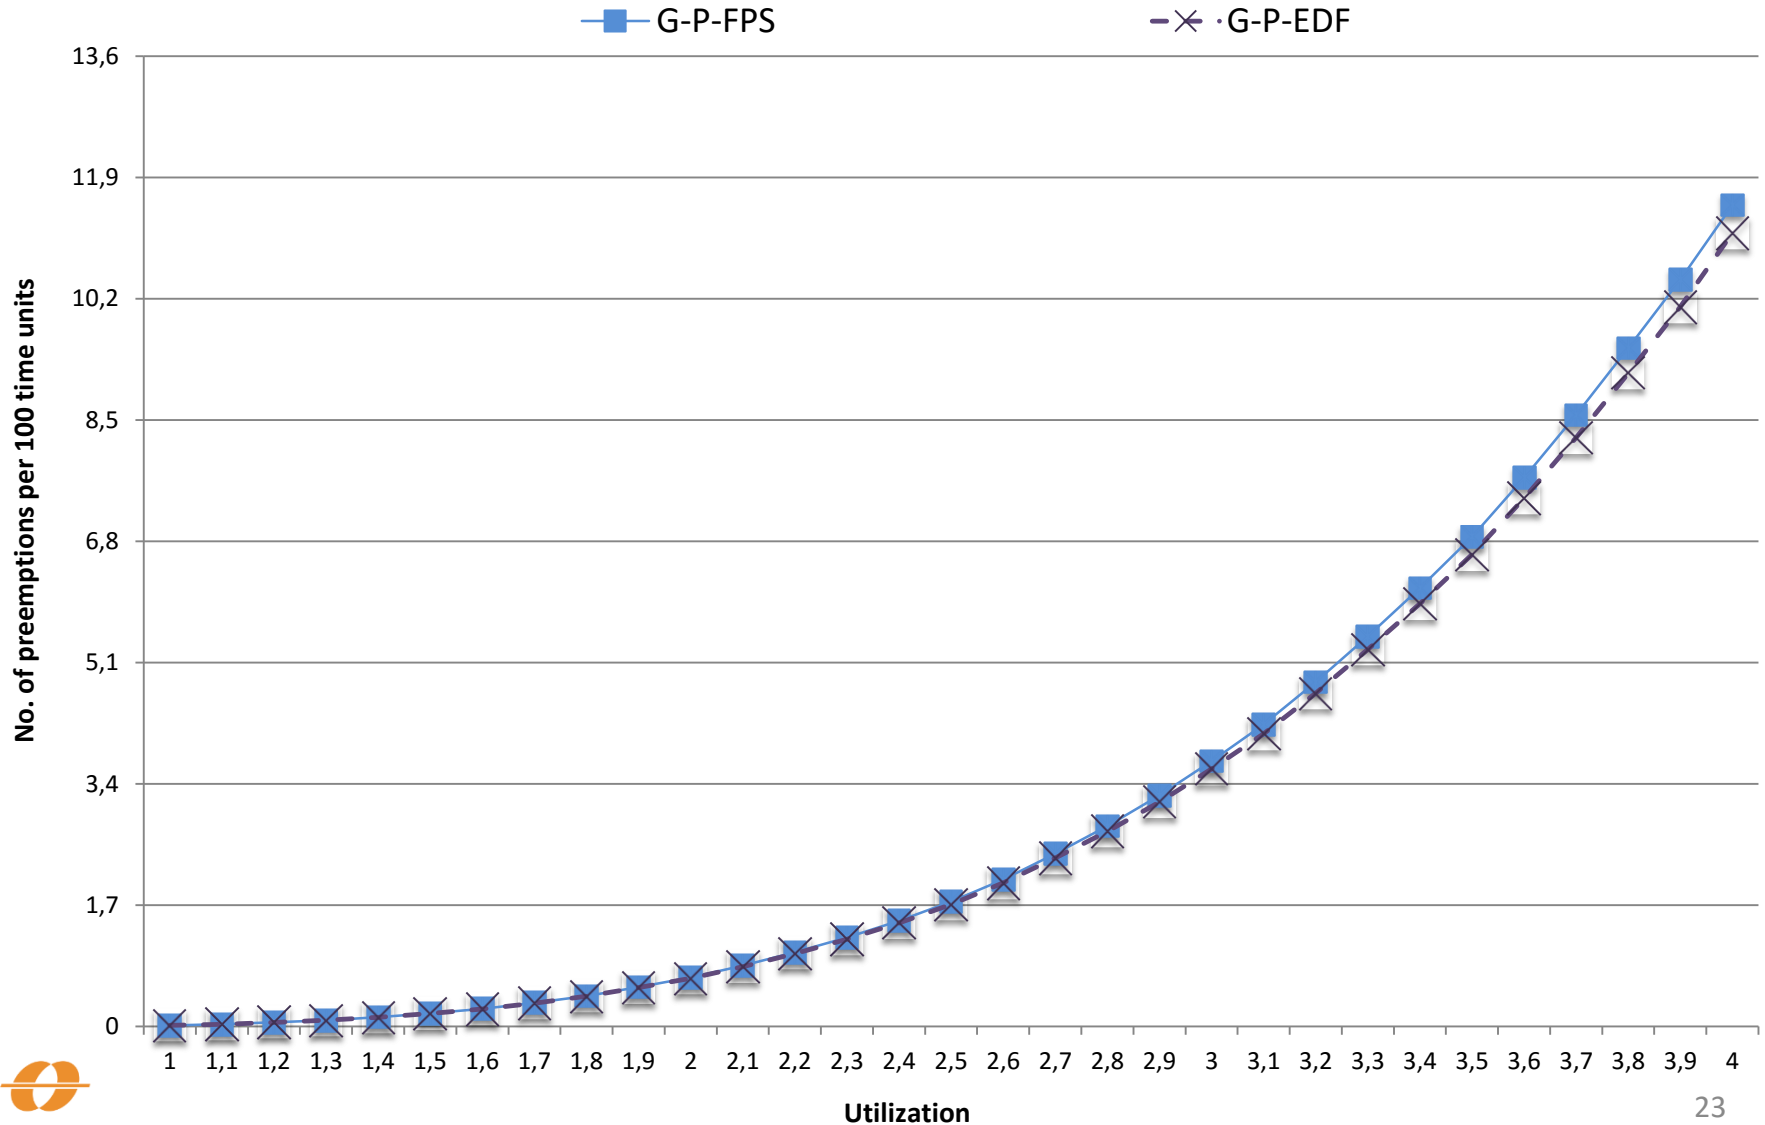
Experiments

We investigated how *weighted preemptions* vary with:

1. Varying utilizations
2. Varying number of tasks
3. Varying number of processors
4. Varying NPR lengths

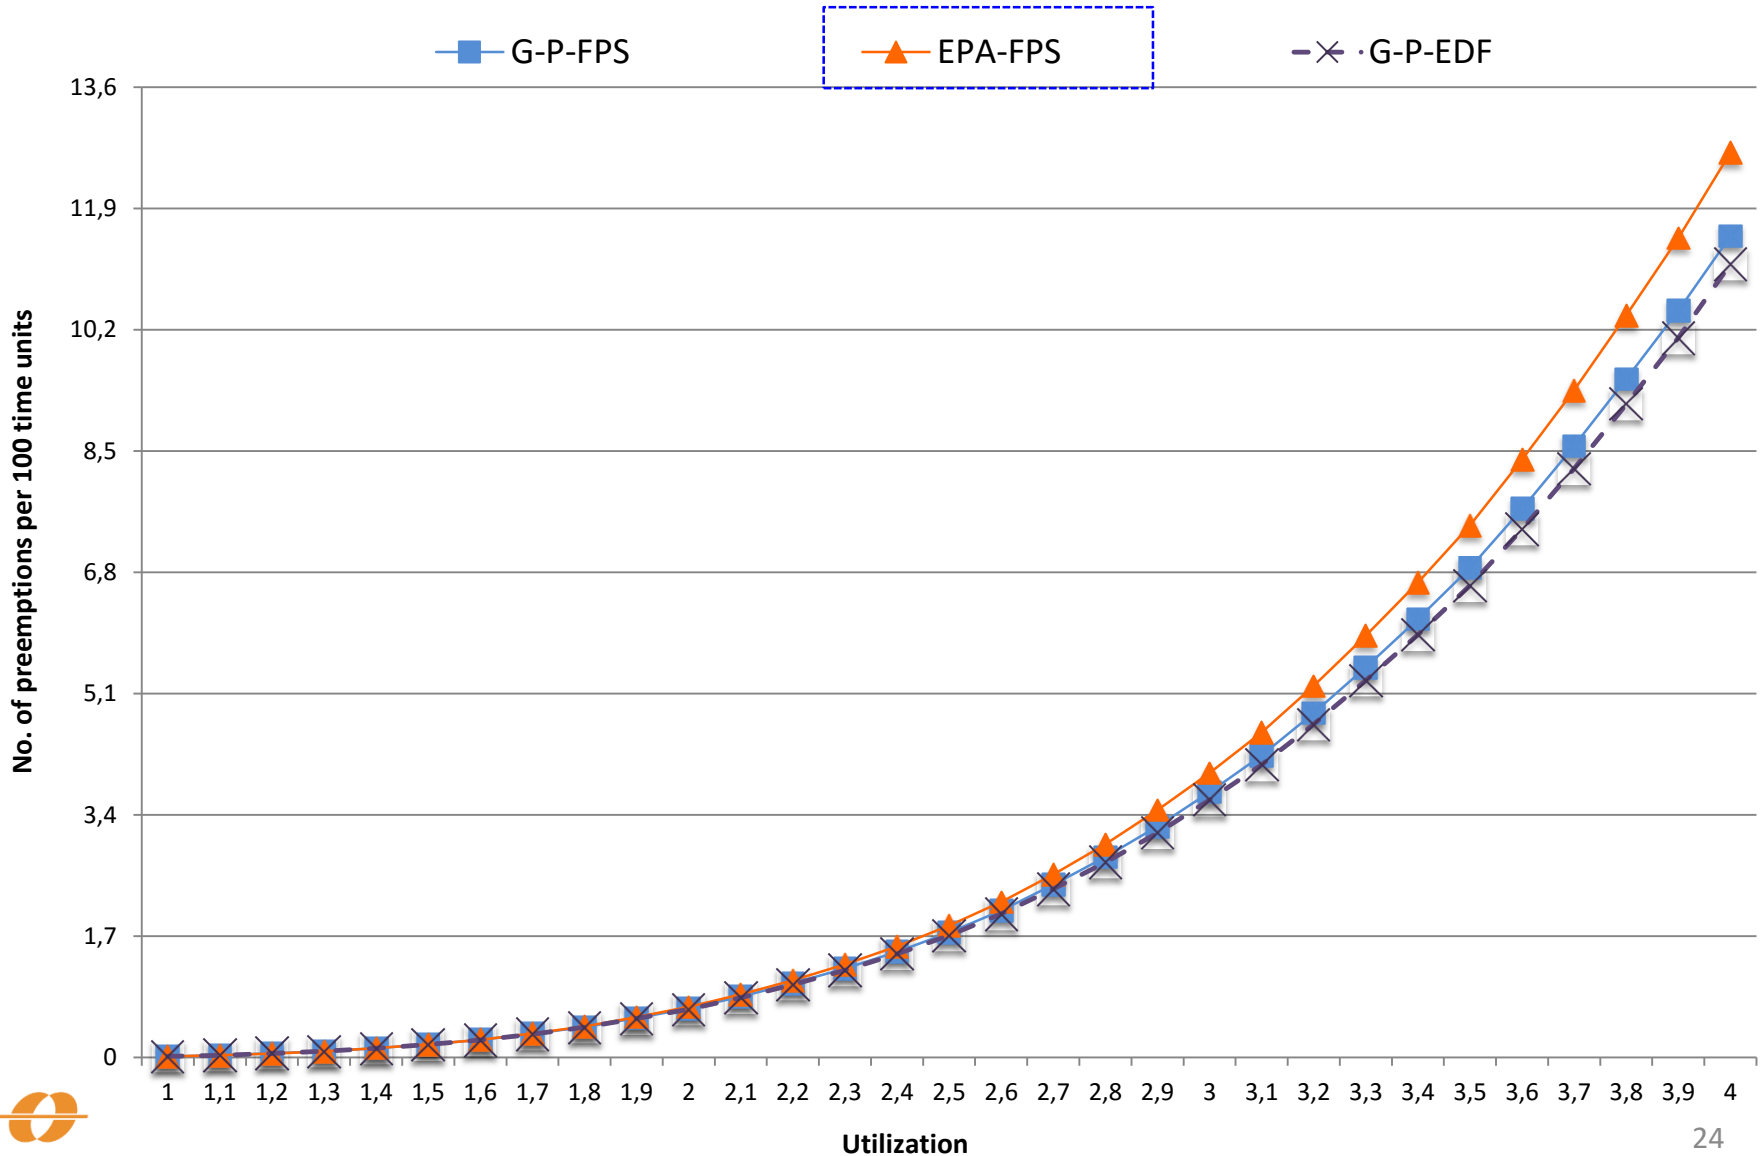
Varying Utilization

Period: 5 to 500, $n=25$ and $m=4$




Varying Utilization

Period: 5 to 500, n=25 and m=4




Varying Utilization

Period: 5 to 500, n=25 and m=4



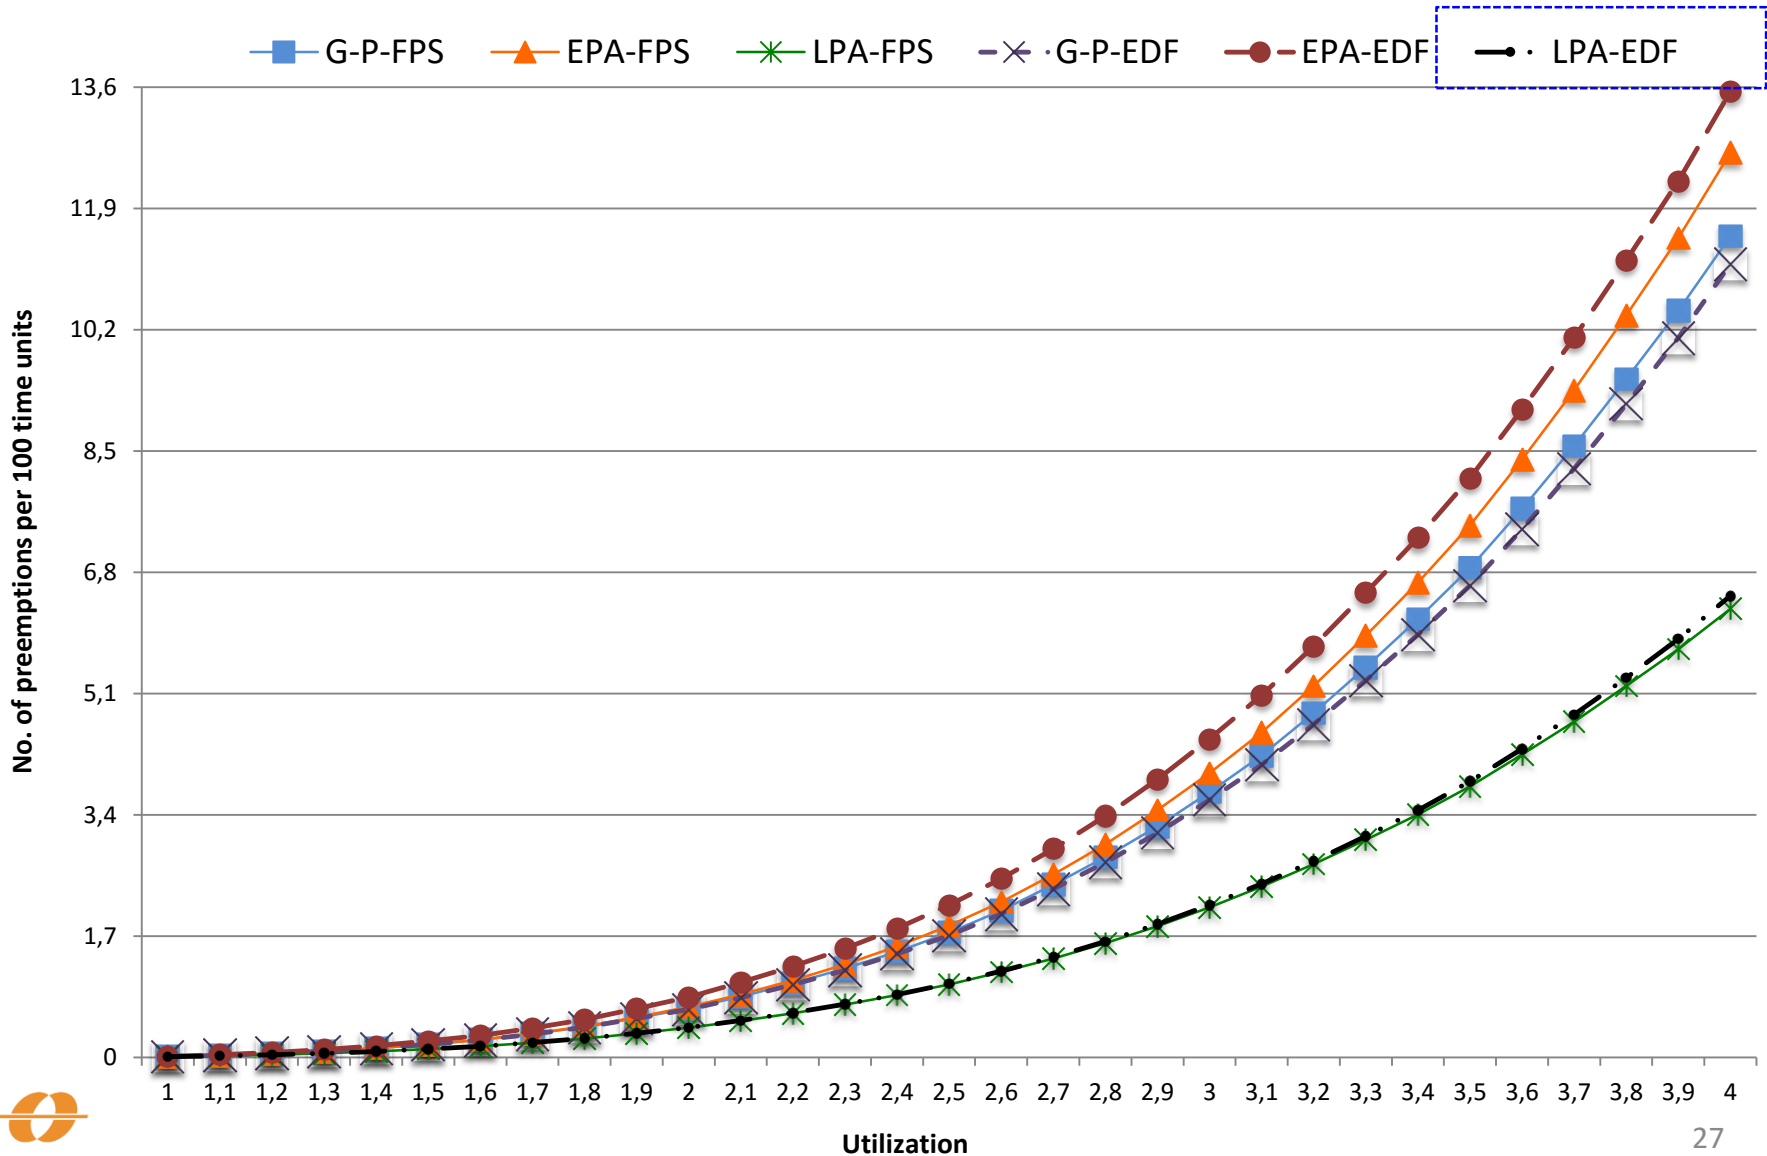
Varying Utilization

Period: 5 to 500, $n=25$ and $m=4$



Varying Utilization

Period: 5 to 500, $n=25$ and $m=4$



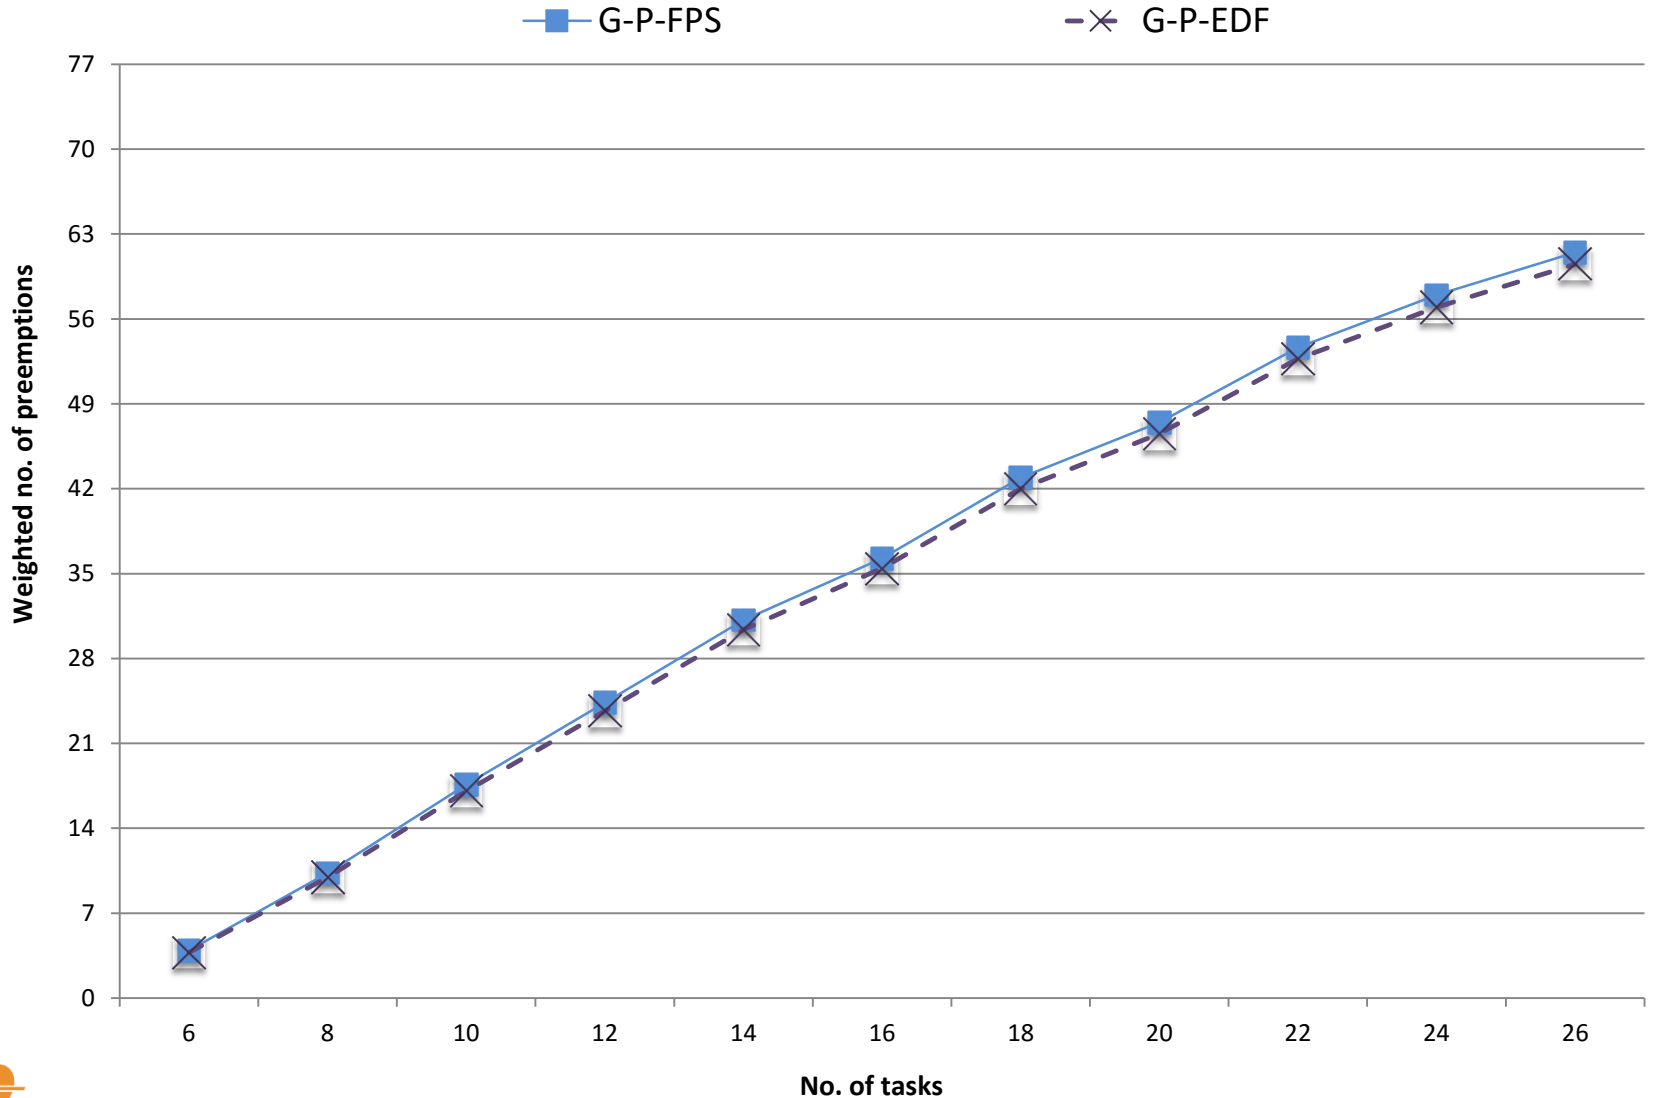
Experiments

We investigated how *weighted preemptions* vary with:

1. Varying utilizations
2. *Varying number of tasks*
3. Varying number of processors
4. Varying NPR lengths


Varying Number of Tasks

Period: 5 to 500, $m=4$ and NPR length=10%




Varying Number of Tasks

Period: 5 to 500, $m=4$ and NPR length=10%




Varying Number of Tasks

Period: 5 to 500, $m=4$ and NPR length=10%




Varying Number of Tasks

Period: 5 to 500, $m=4$ and NPR length=10%



Varying Number of Tasks

Period: 5 to 500, $m=4$ and NPR length=10%



Experiments

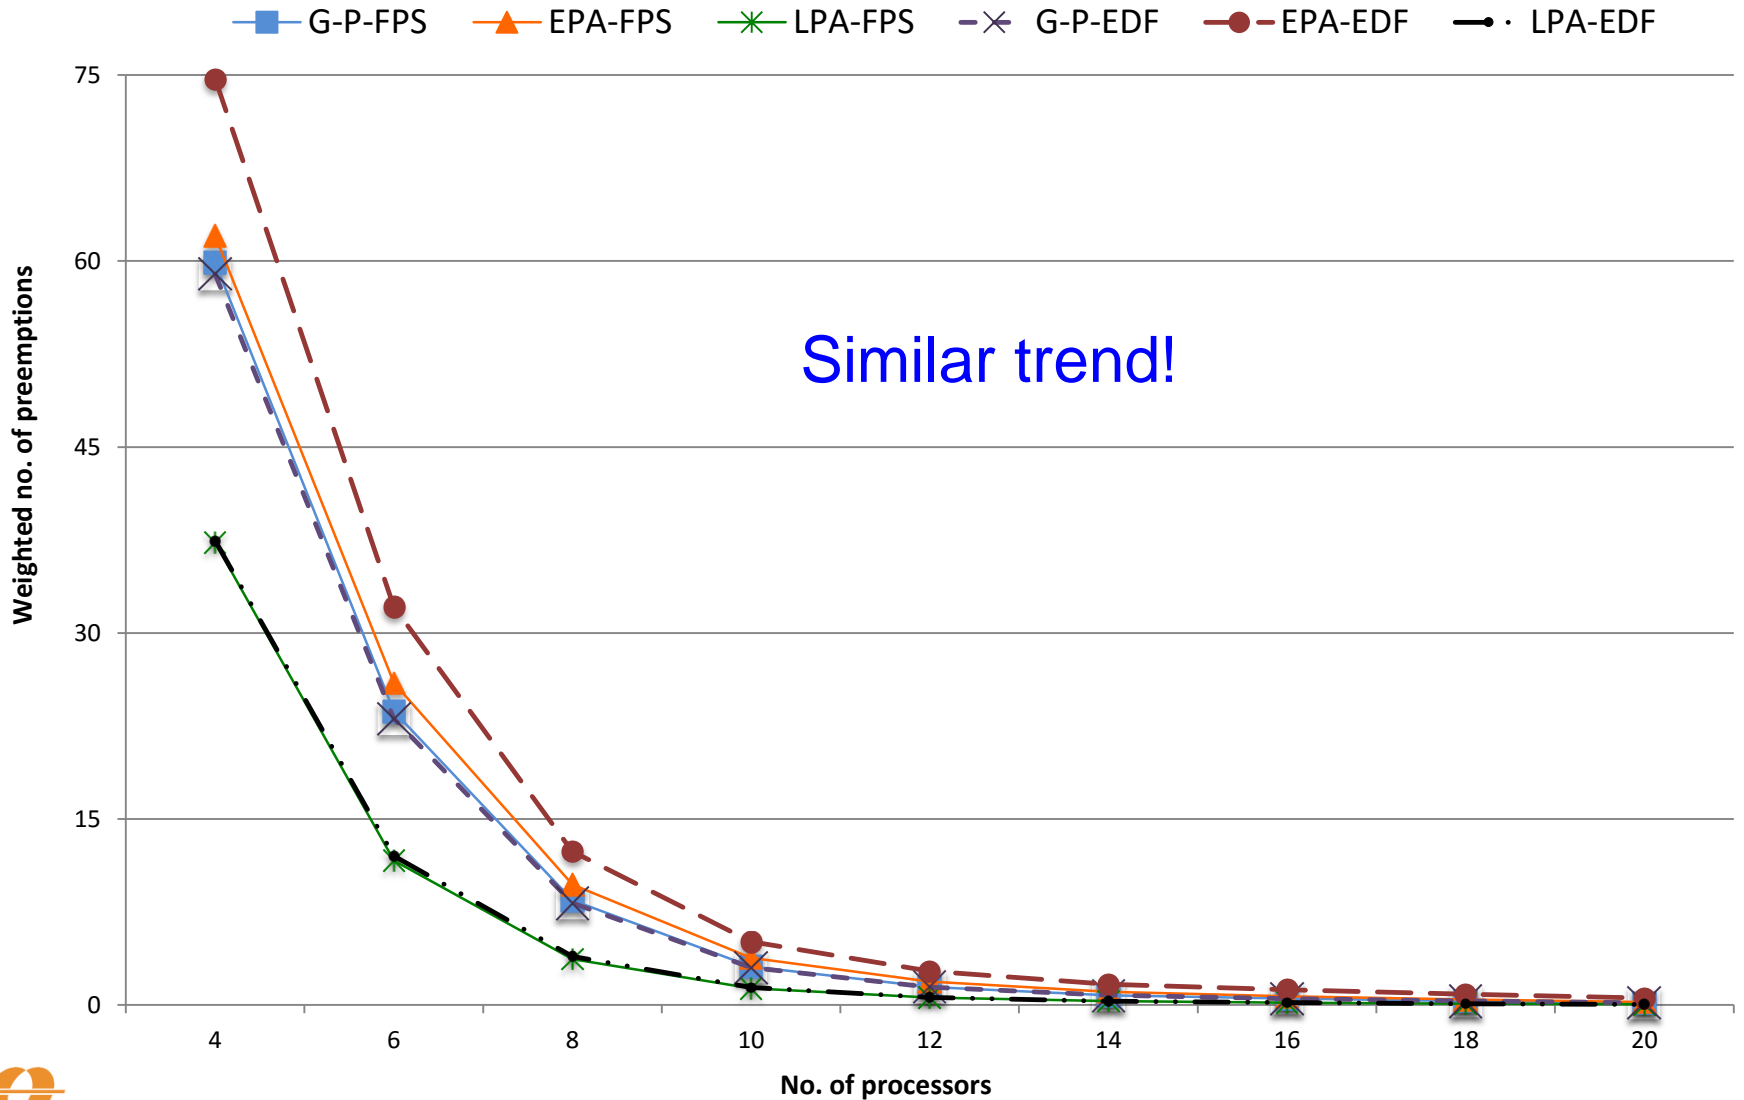
We investigated how *weighted preemptions* vary with:

1. Varying utilizations
2. Varying number of tasks
3. *Varying number of processors*
4. Varying NPR lengths



Varying Number of Processors

Period: 5 to 500, $n=25$ and NPR length=10%



Similar trend!



Experiments

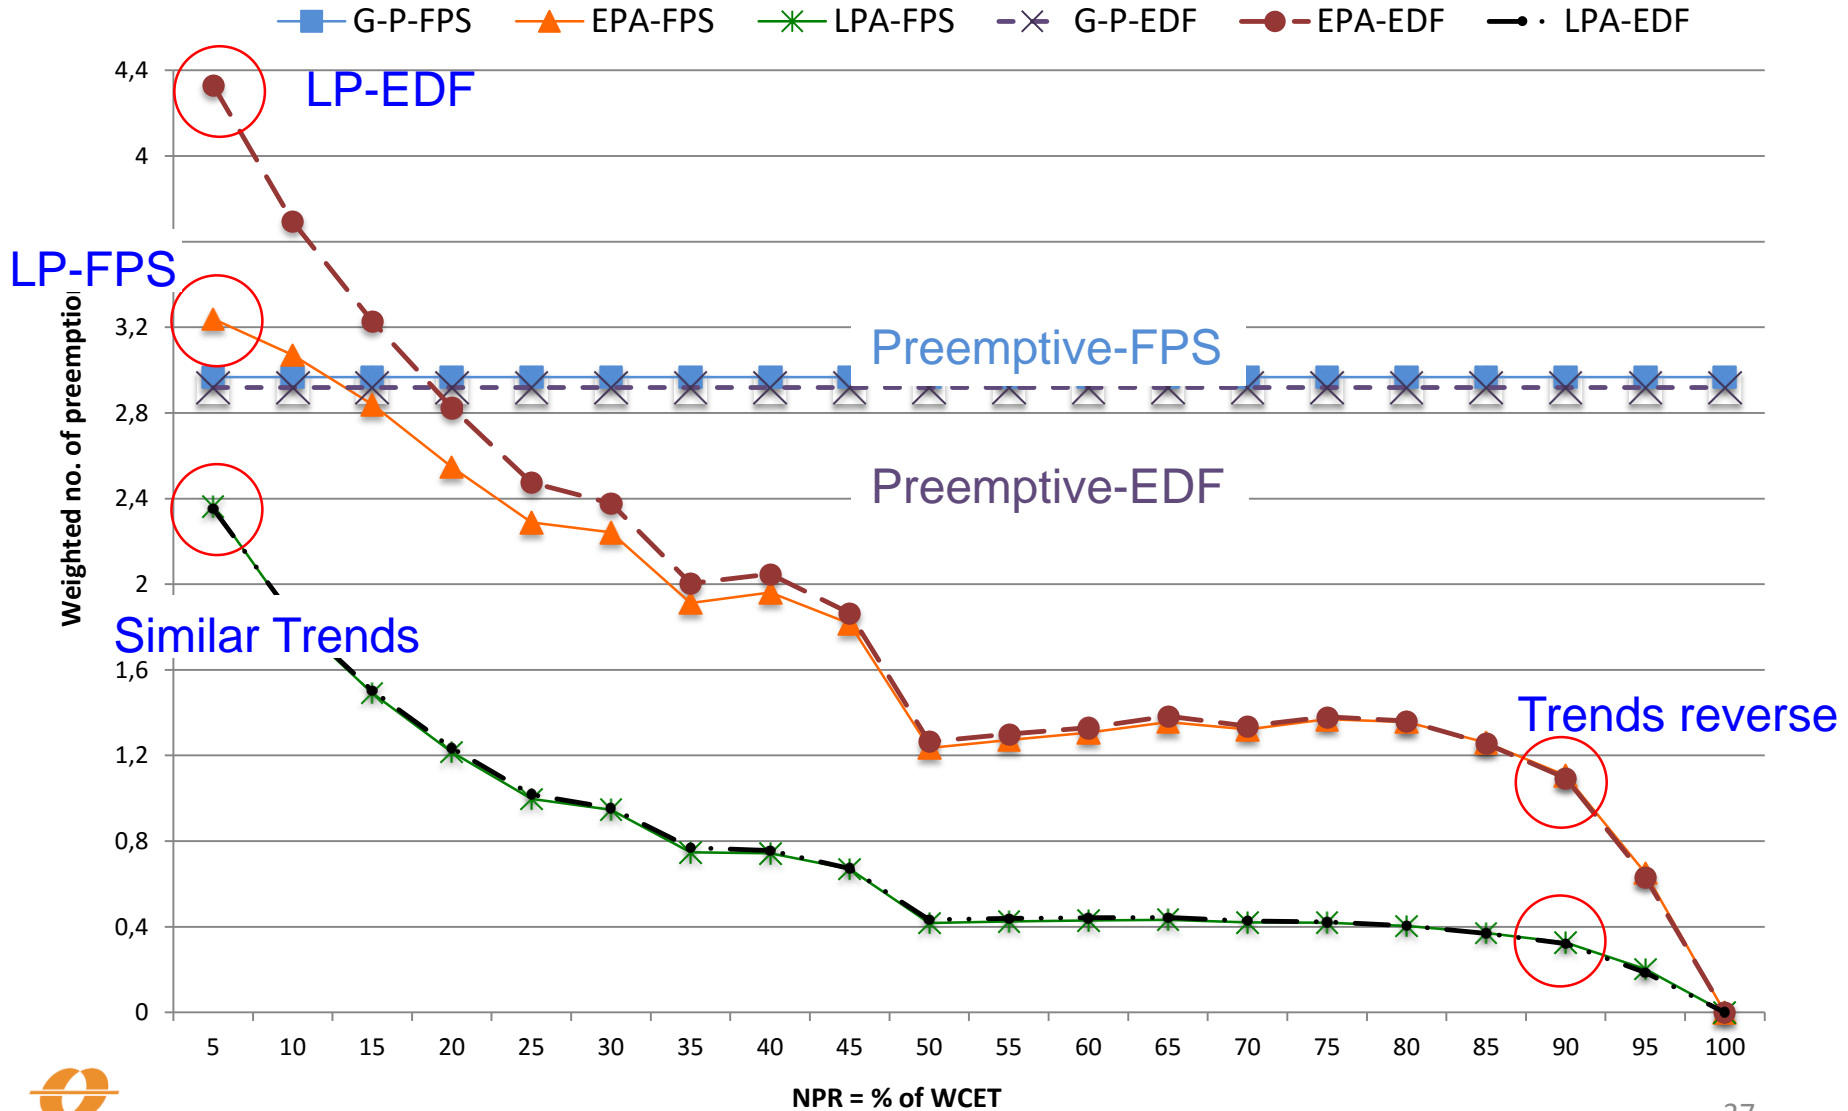
We investigated how *weighted preemptions* vary with:

1. Varying utilizations
2. Varying number of tasks
3. Varying number of processors
4. Varying NPR lengths



Varying Lengths of NPRs

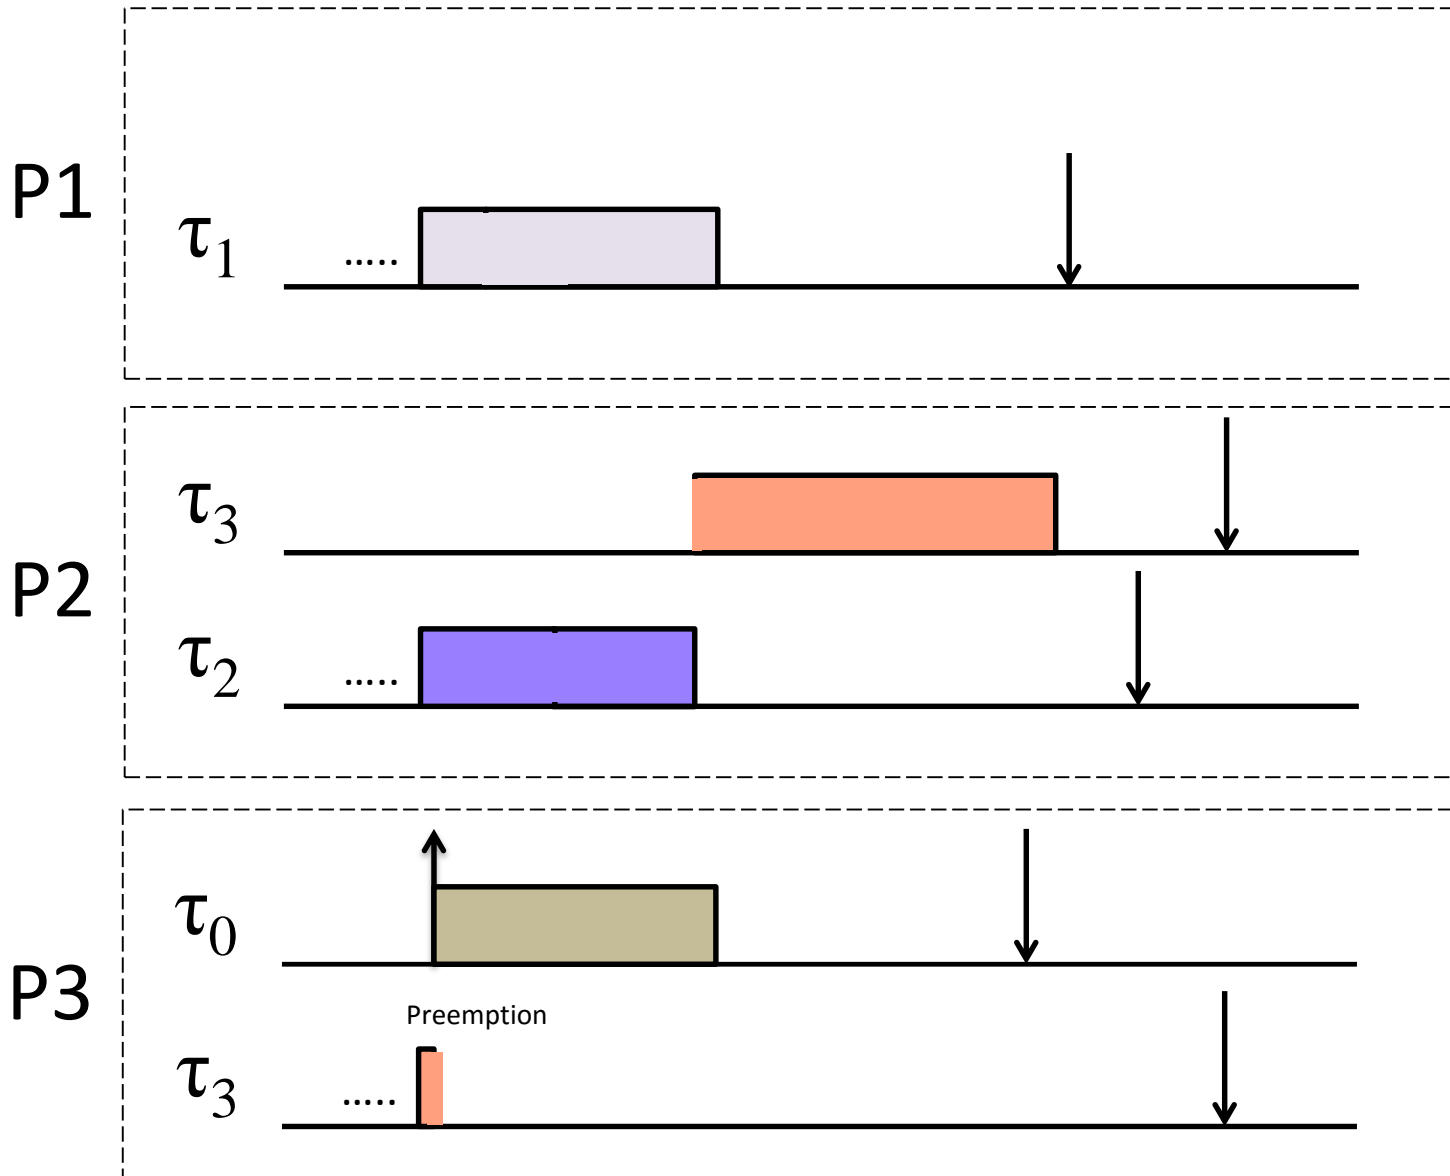
Period: 5 to 500, n=25 and m=4



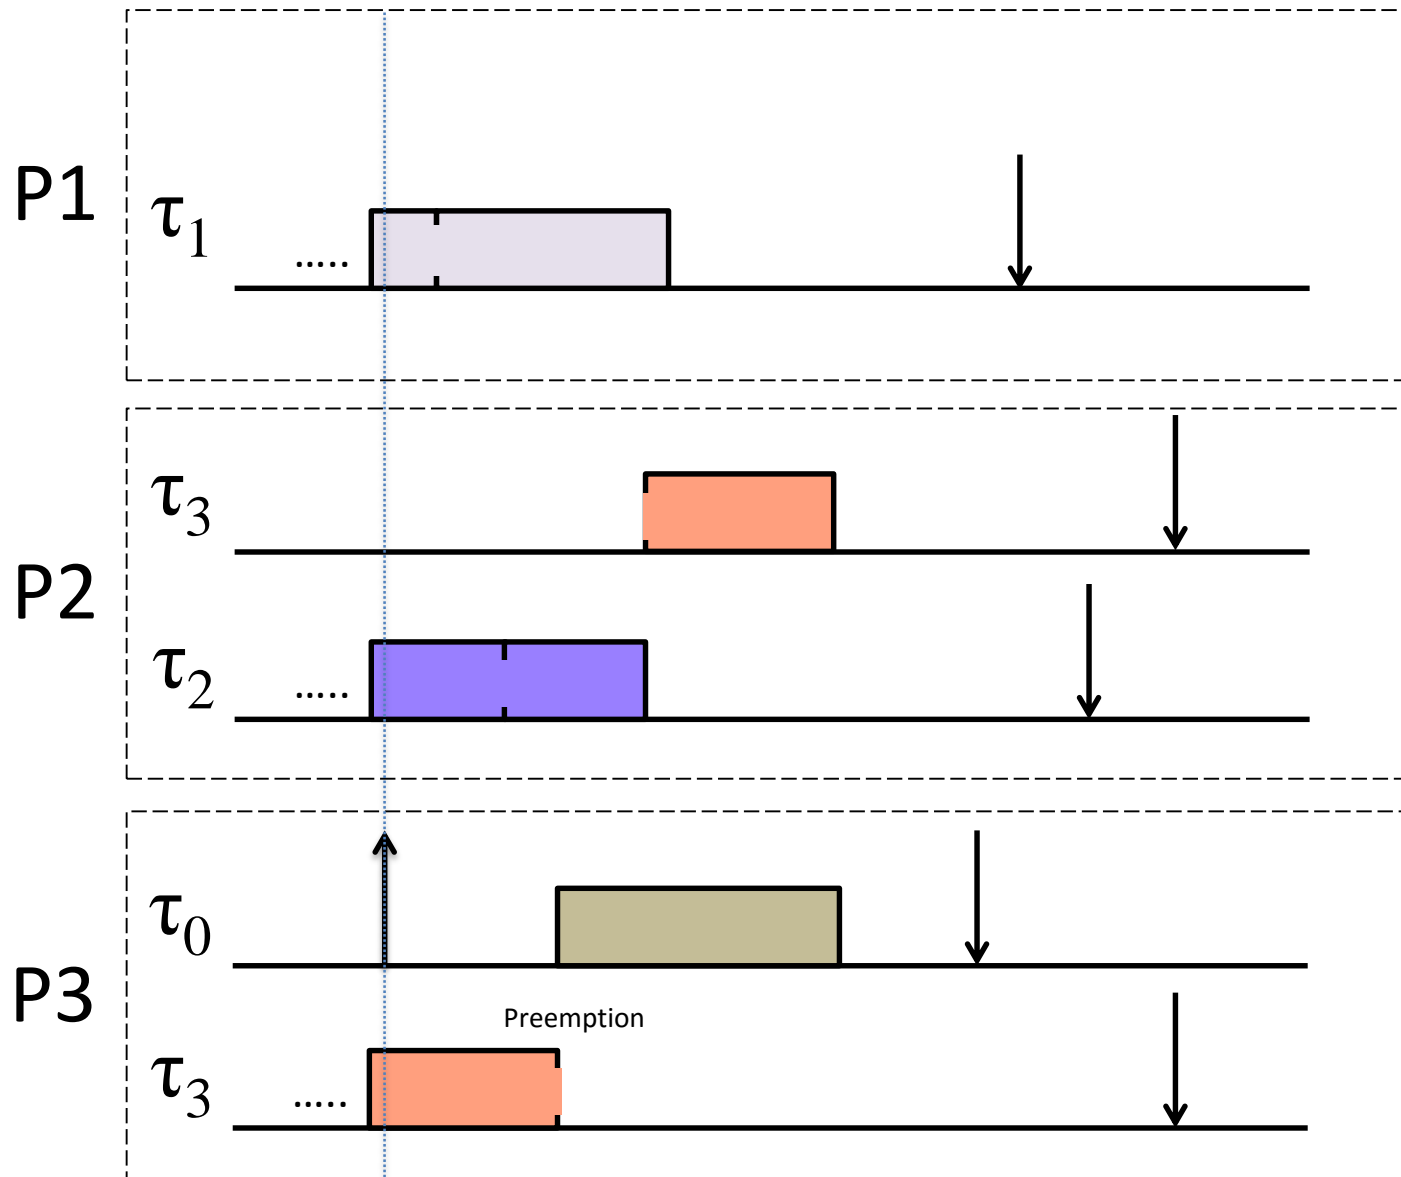
Observations

1. **Eager** preemption approach generates **most** preemptions while **lazy** preemption approach generates **least**

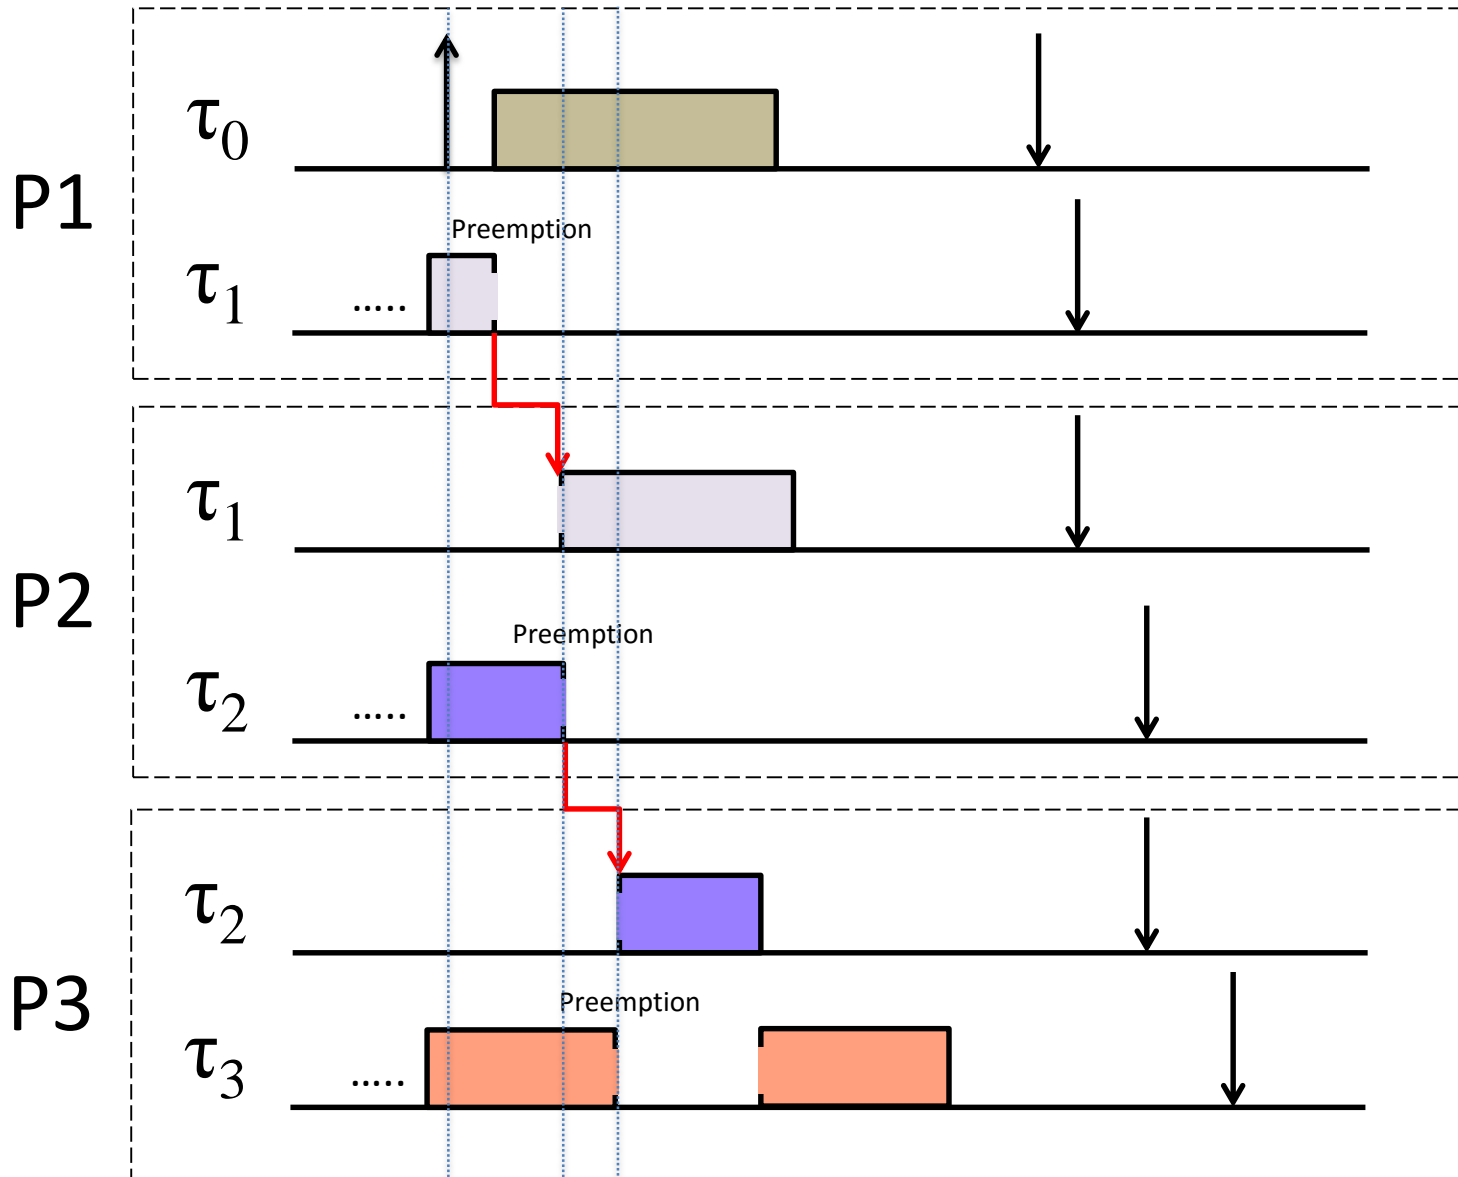
Preemptive Scheduling



Lazy Preemption Approach



Eager Preemption Approach

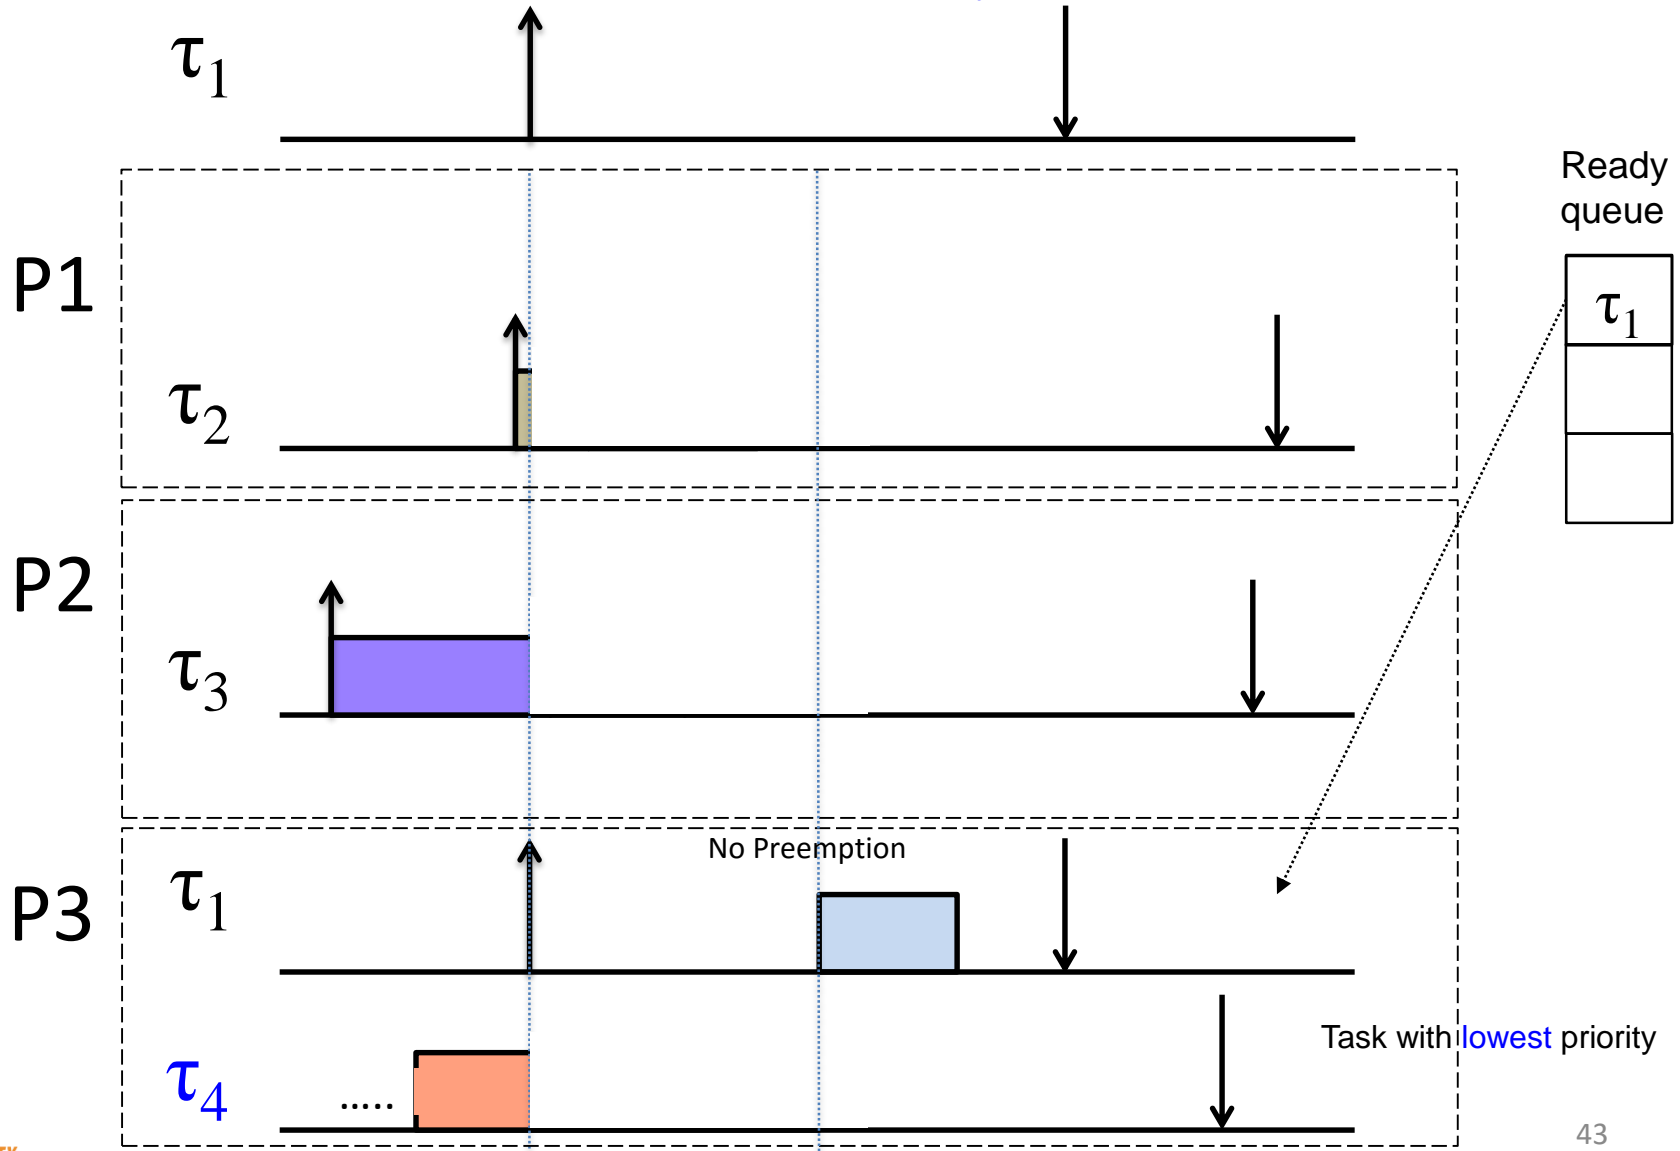


Observations

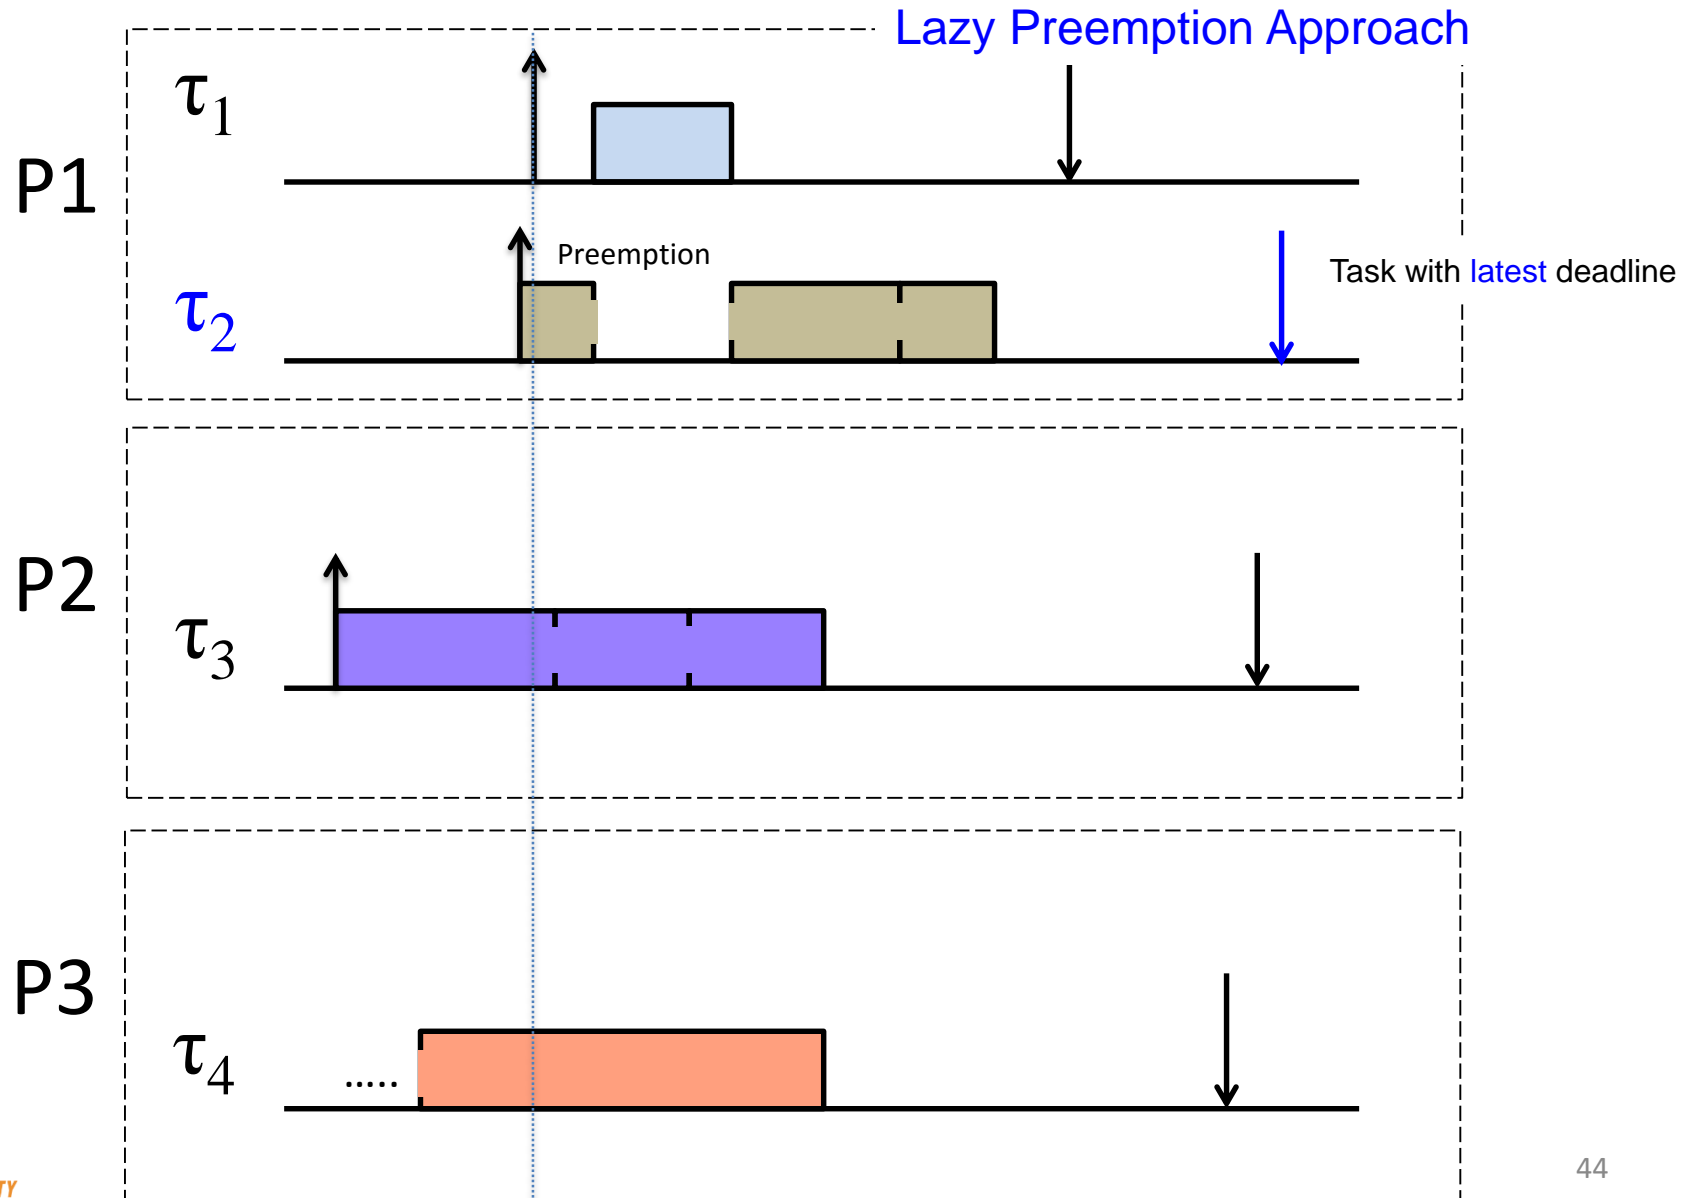
1. **Eager** preemption approach generates **most** preemptions while **lazy** preemption approach generates **least**
2. Preemptive EDF generates **fewer** preemptions than preemptive FPS
3. Limited preemptive EDF generates **more** preemptions than limited preemptive FPS

Preemptive Behavior under FPS

Lazy Preemption Approach



Preemptive Behavior under EDF



Conclusions

1. Limited preemptive scheduling on multiprocessors **may not reduce the number of preemptions**
 - Number of preemptions **larger than fully preemptive scheduling** under the eager approach
2. The **preemptive behavior of EDF** on uniprocessors **generalizes to multiprocessors**
 - G-P-EDF generates **fewer** preemptions than G-P-FPS
3. This, however, **does not** generalize to **global limited preemptive scheduling**
 - G-LP-EDF generates **more** preemptions than G-LP-FPS
4. G-LP-FPS with LPA generates the **least** number of preemptions

Future Work

- **Hybrid** approach to **manage** preemptions on multicores to get **best** of **eager** and **lazy** approaches: integration with preemption thresholds
- Probabilistic schedulability analysis for mixed-criticality tasks with fixed preemption points
- Improved preemption overhead accounting for tasks with fixed preemption points

Thank you !



Questions/Comments ?