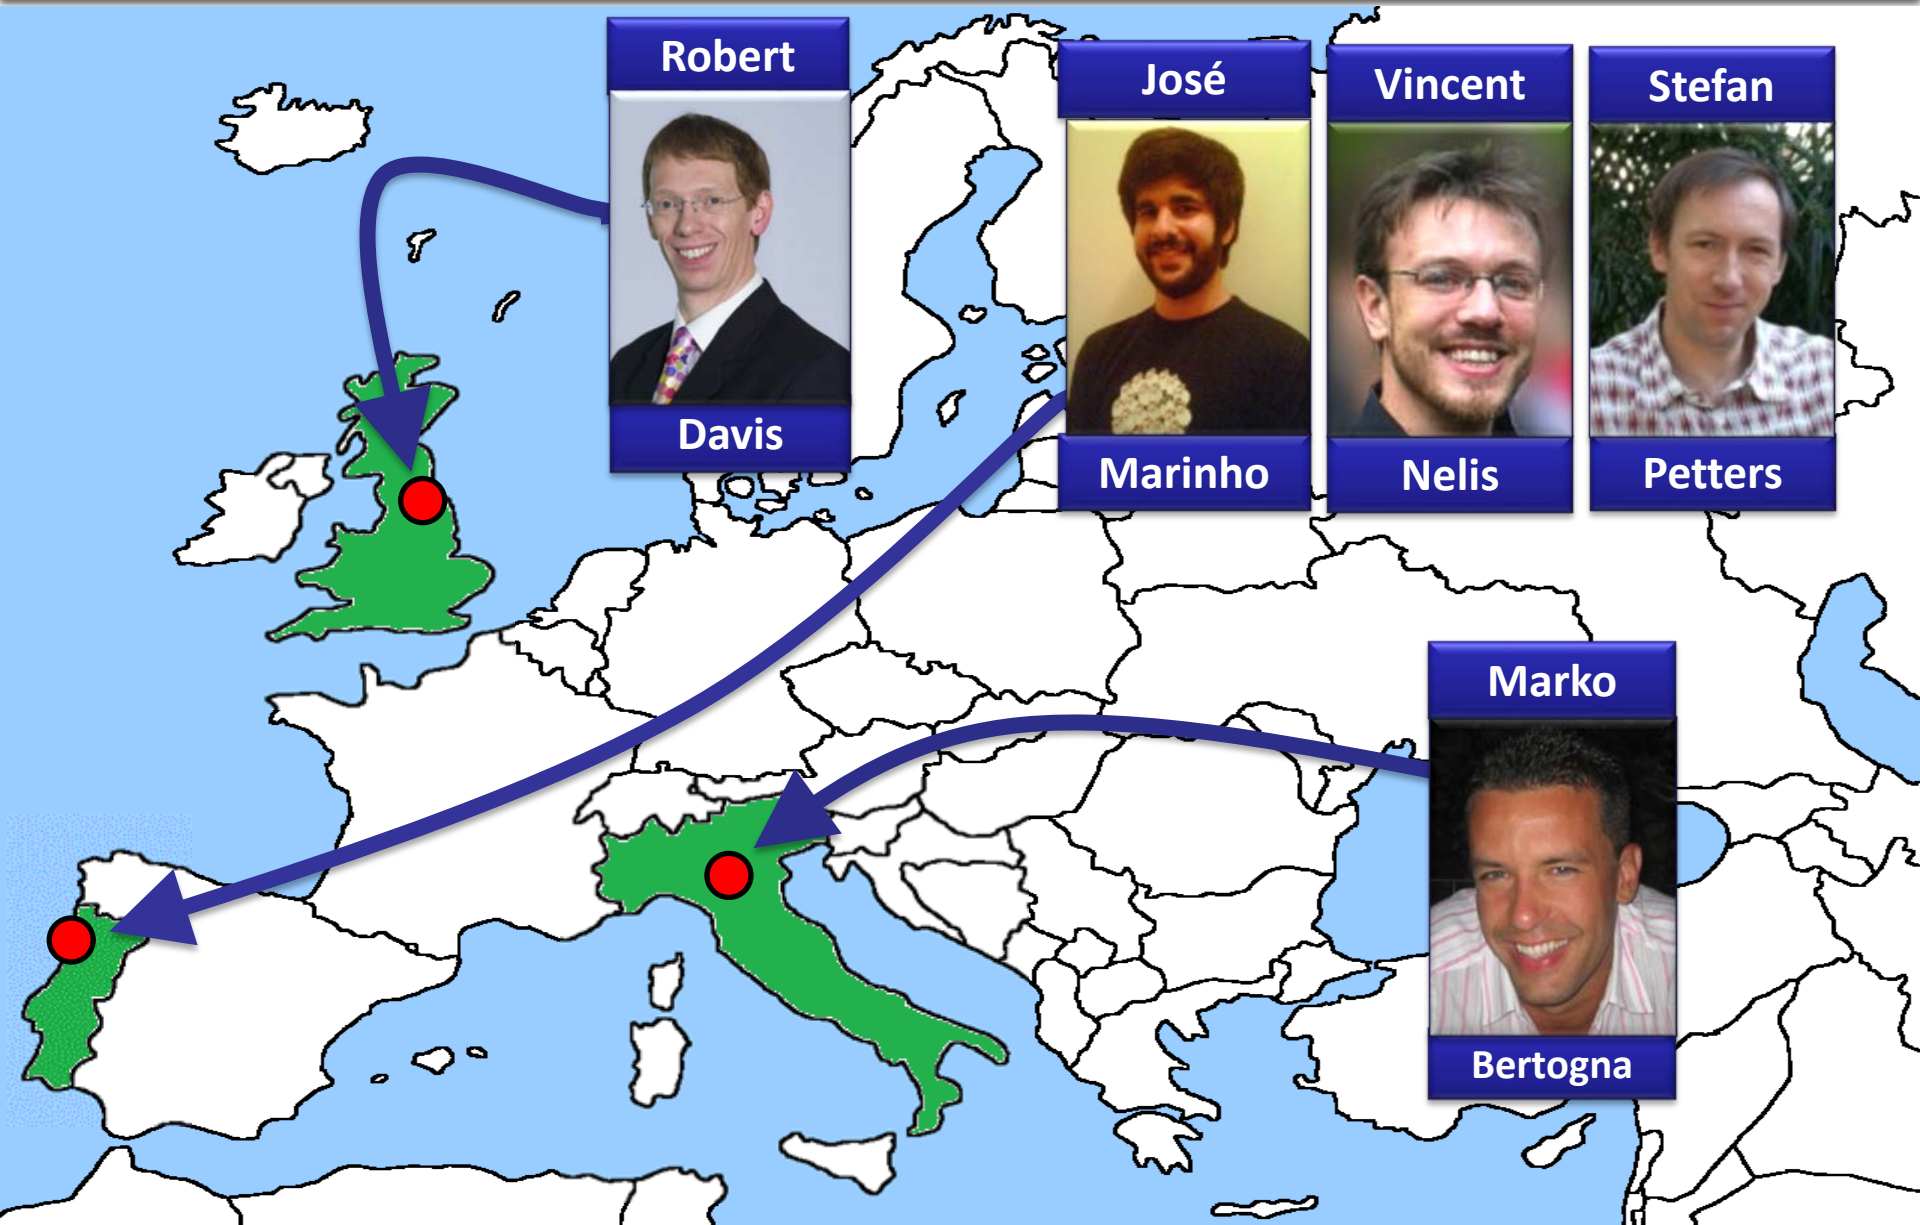




# Limited Pre-emptive Global Fixed Task Priority



Robert



Davis

José



Marinho

Vincent



Nelis

Stefan



Petters

Marko



Bertogna




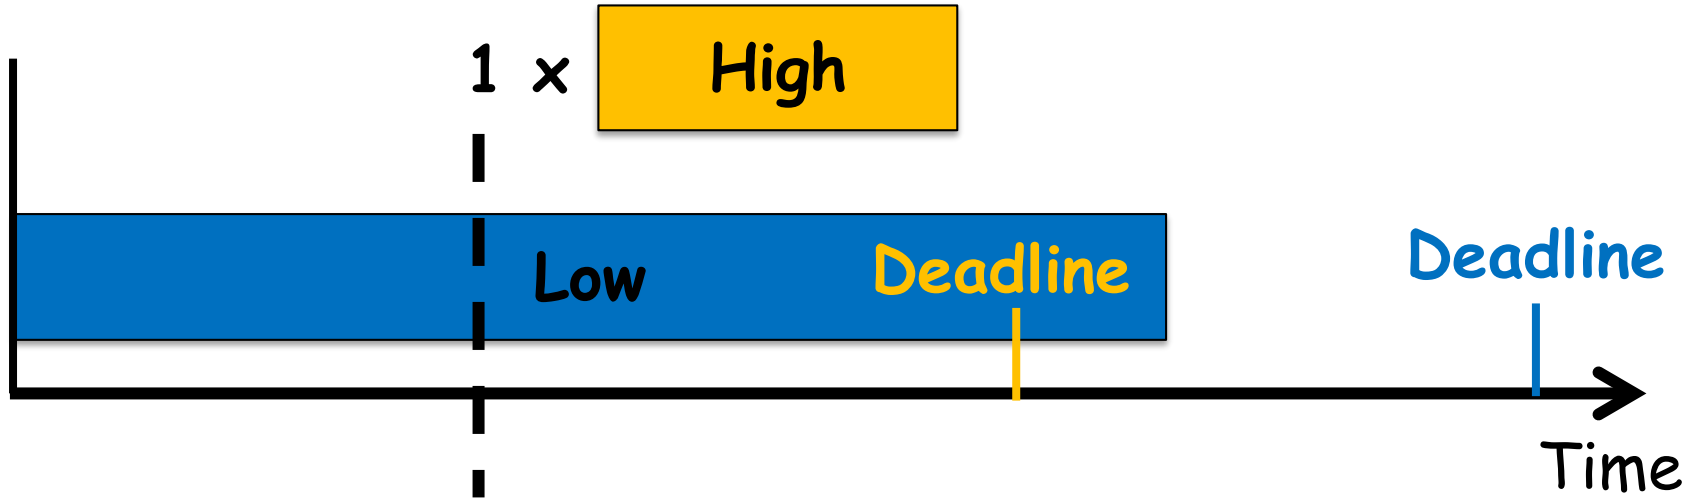
# In a nutshell...

We studied the preemptive scheduling problem on multicores

We focused on the **limited preemptive model**, which is somewhat between fully preemptive and non preemptive approaches


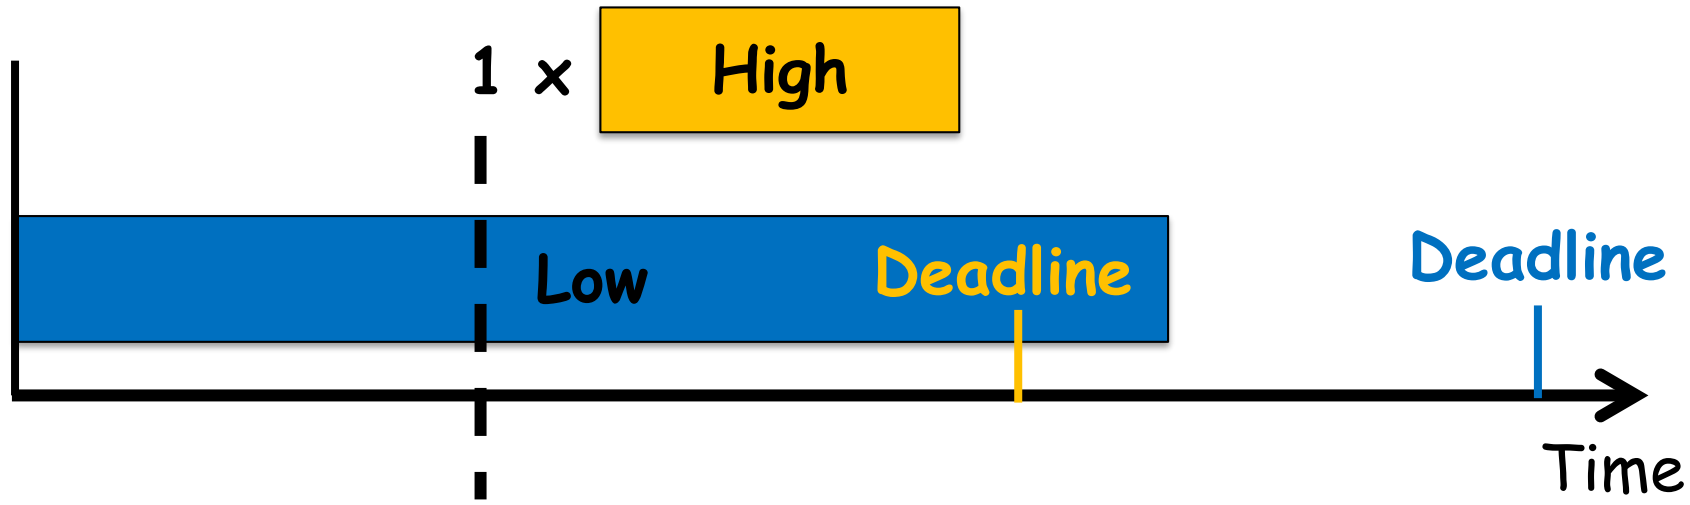


# Preemption or no preemption?





# Preemption or no preemption?





# In a nutshell...

We studied the preemptive scheduling problem on multicores

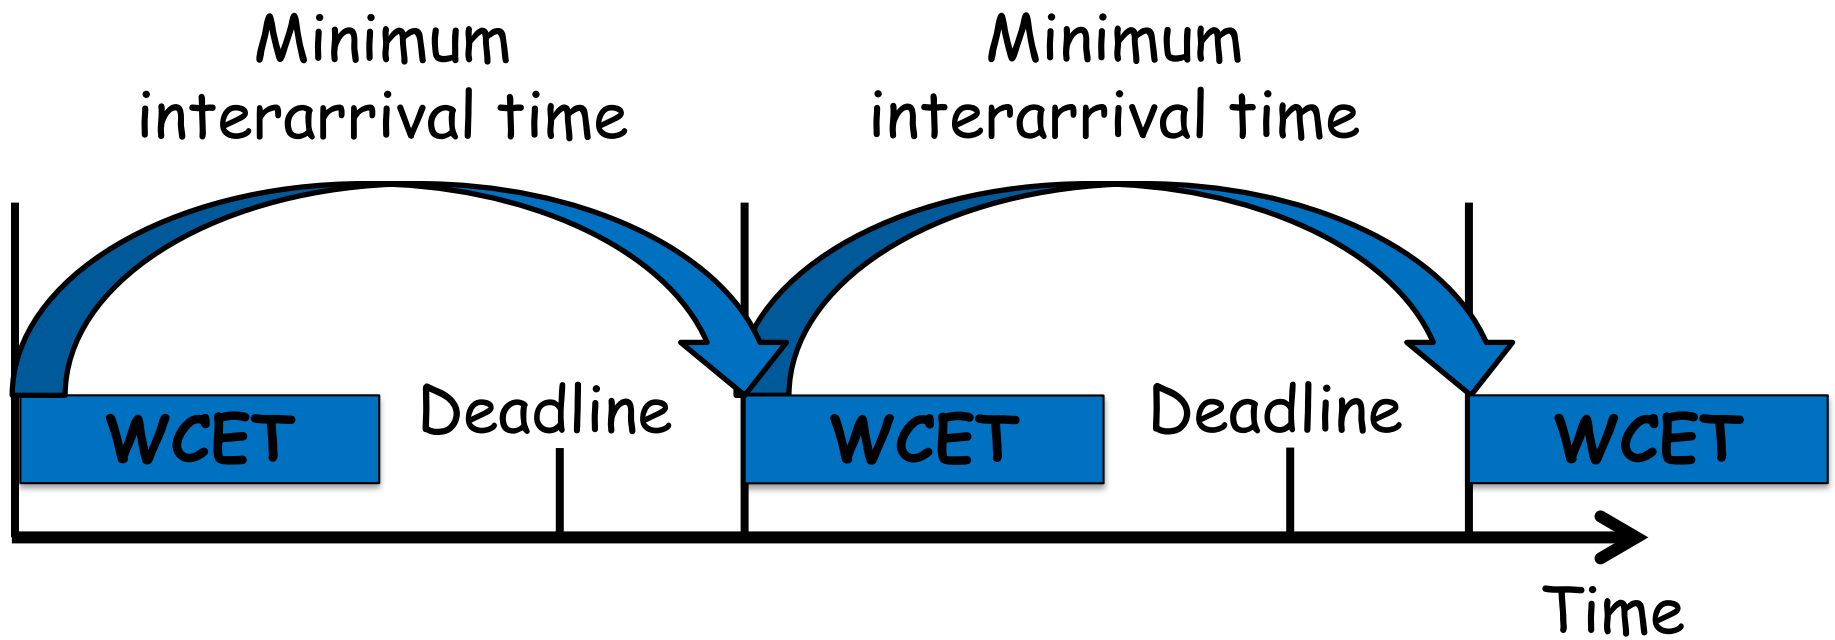
We focused on the **limited preemptive model**, which is somewhat between fully preemptive and non preemptive approaches

We came up with **two different scheduling techniques** and derived **schedulability analyses** for each of them

We played around with the parameters to further **reduce the number of preemption** and **increase the schedulability**

## Application model:

Sporadic constrained-deadline tasks






## Scheduler model:

Global Fixed Task Priority Scheduling  
Fixed Non-Preemptive Region

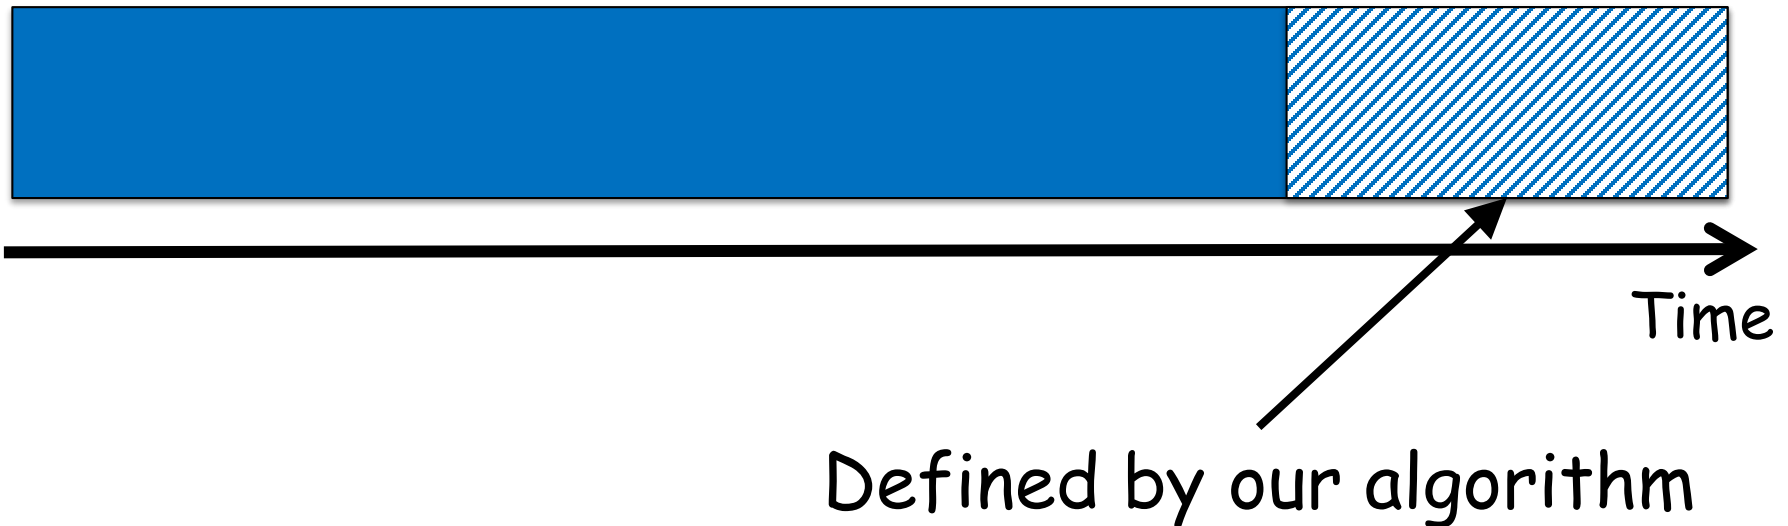
Non-preemptive Regions






## Scheduler model:

Global Fixed Task Priority Scheduling  
Fixed Non-Preemptive Region







## Platform model:

Identical multiprocessor platform  
m processors




## Assumptions:

No migration cost,  
No preemption cost  
Independent tasks

Focus on the two new scheduling techniques that we propose and the impact on the schedulability

**A bit of history...**






# Background on fully-preemptive schedulability analysis

Arrival

Deadline

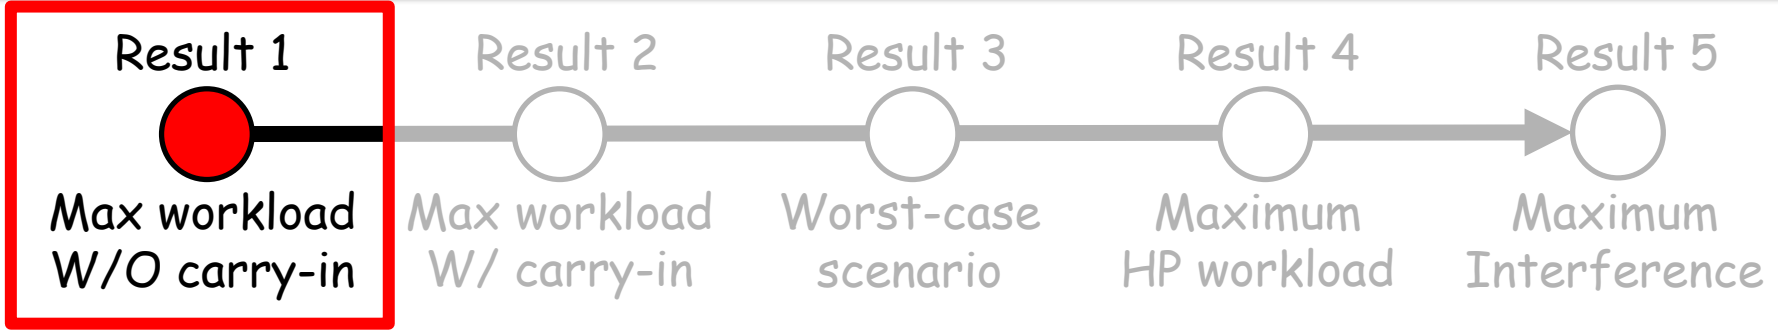


←—————→

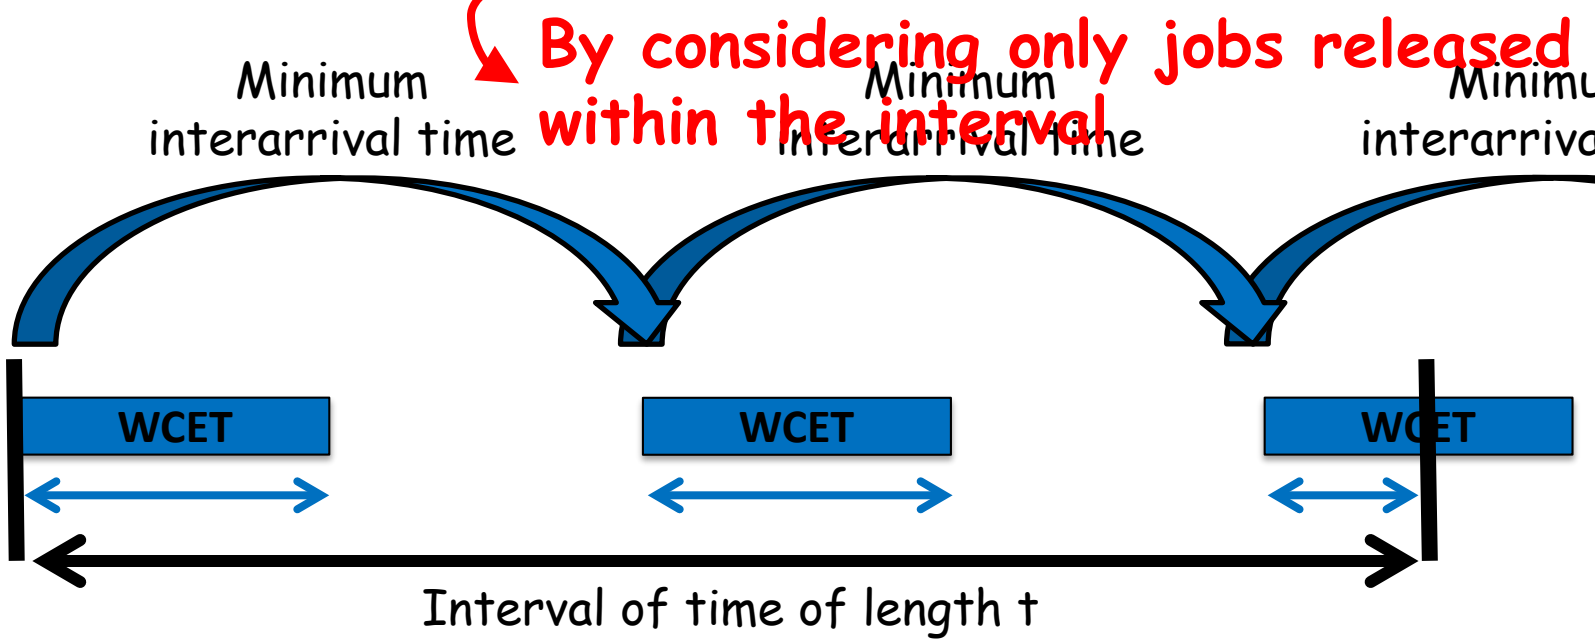
Maximum  
interference  
from HP workload  
in  $t$  time units



# Background on fully-preemptive schedulability analysis



Result 1: Maximum workload **without carry-in job**

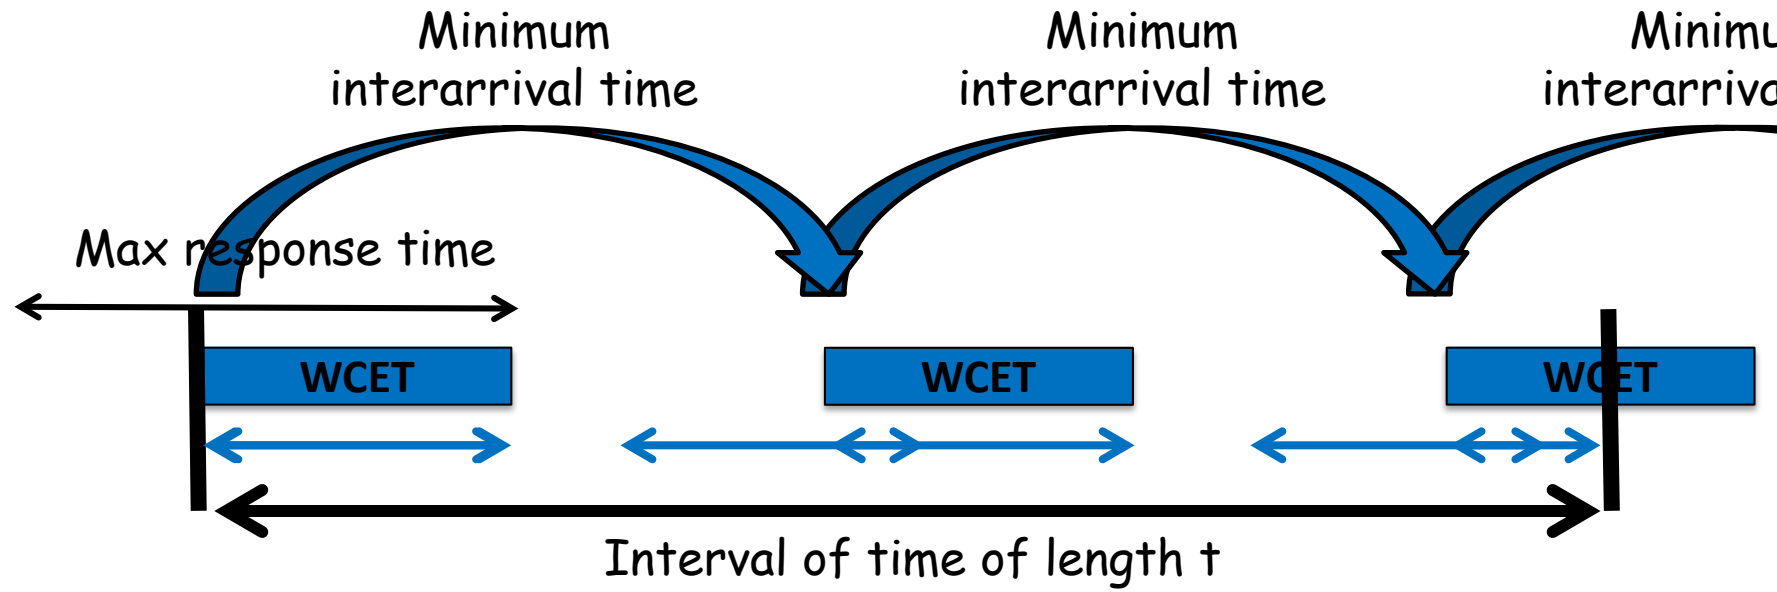




# Background on fully-preemptive schedulability analysis

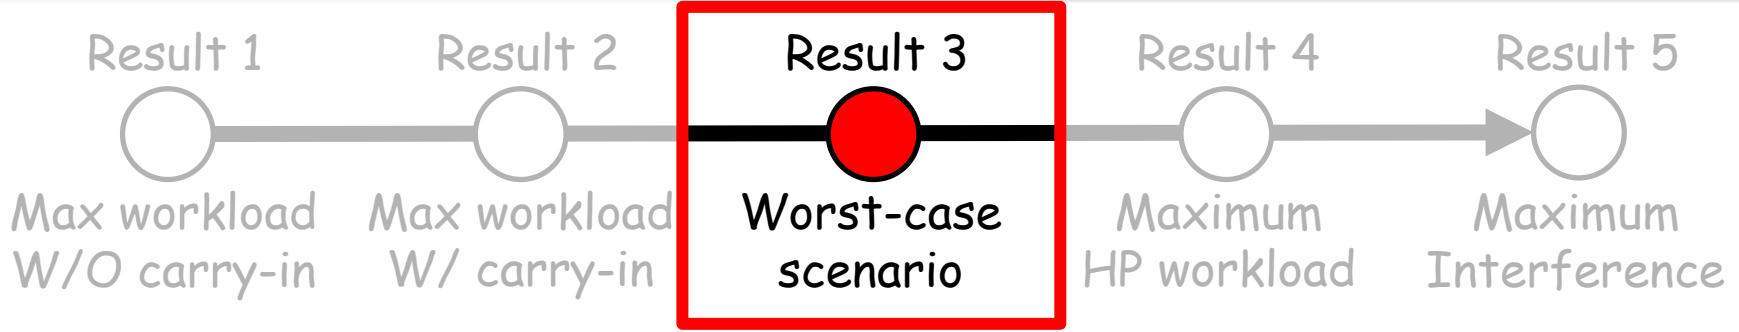


Result 2: Maximum workload **with** carry-in





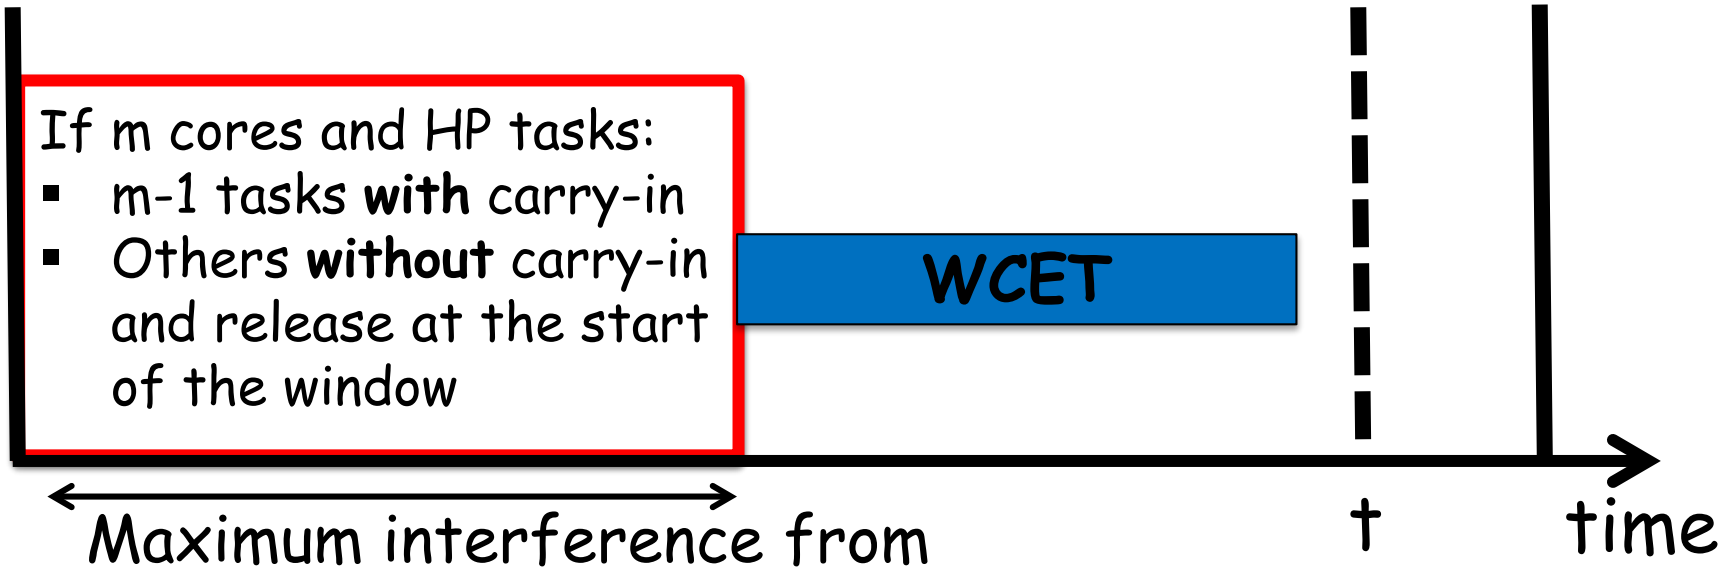
# Background on fully-preemptive schedulability analysis



Result 3: Worst-case interference scenario

Arrival

Deadline






# Background on fully-preemptive schedulability analysis



## Result 4: Maximum Higher Priority Workload

1. Worst-case interference scenario
2. Maximum task workload without carry-in job
3. Maximum task workload with carry-in job



Maximum workload that can be generated within the given time window




# Background on fully-preemptive schedulability analysis



Result 4: Maximum interference

Arrival **Max interference = max workload / m** Deadline





## Definition of fully preemptive scheduling on multicores


~~"policy where at any time,~~ the  $m$  highest priority tasks execute on the  $m$  processors of the platform."



In **limited** preemptive scheduling


## Definition of fully preemptive scheduling on multicores

~~"policy where at any time,~~ the  $m$  highest priority tasks execute on the  $m$  processors of the platform."



## Definition of fully preemptive scheduling on multicores


~~"policy where at any time, the m highest priority tasks execute on the m processors of the system."~~



**Regular Deferred Scheduling (RDS):** New arriving higher priority workload preempts the first preemptible lower priority workload.

## Definition of fully preemptive scheduling on multicores

~~"policy where at any time,~~ the  $m$  highest priority tasks execute on the  $m$  processors of the platform."

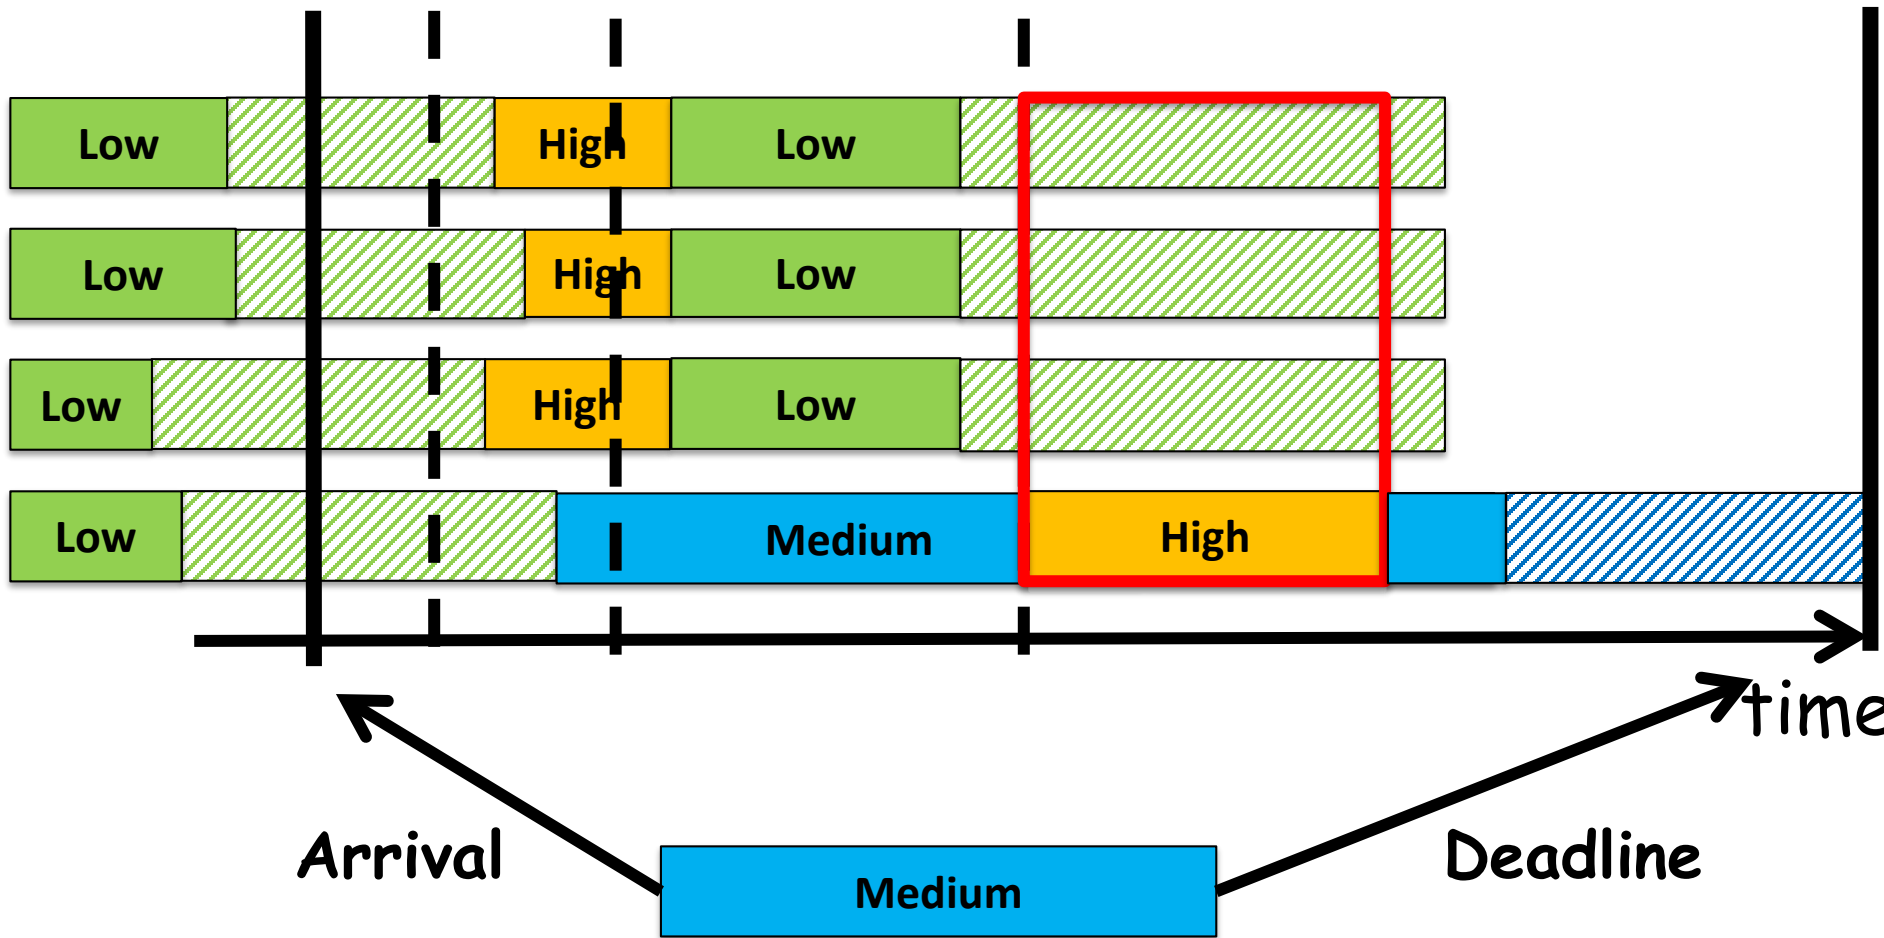


**Adapted Deferred Scheduling (ADS):** New arriving higher priority workload preempts the lowest priority job once it gets preemptible.

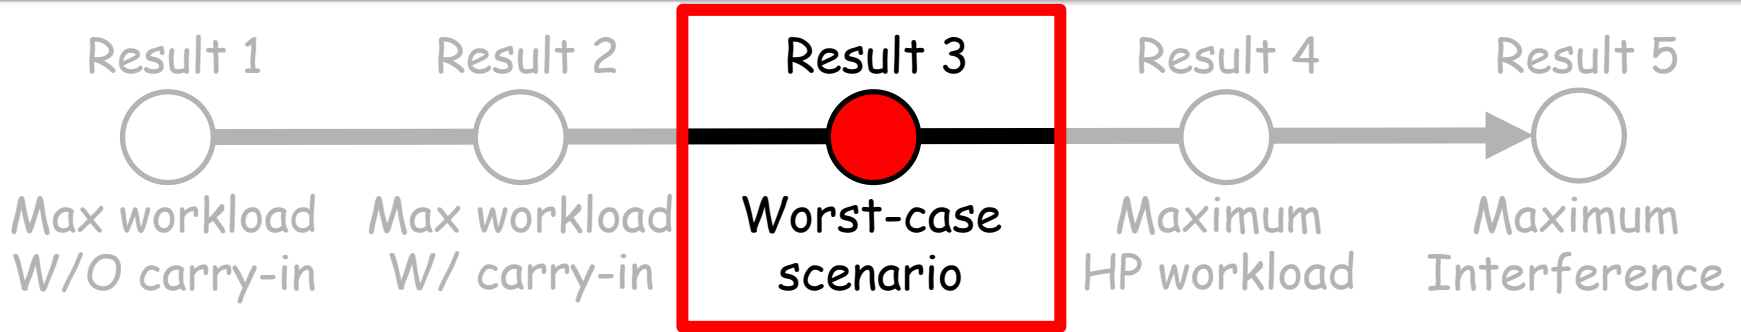


# What's wrong with RDS?

3 x High      3 x Low      1 x High



# What's wrong with RDS



Result 3: Worst-case

**With RDS, all the tasks, including HP and LP, can interfere with the task under analysis**

If  $m$  cores and  $n$  HP tasks:

- $m-1$  tasks with carry-in
- Others without carry-in and release at the start of the window

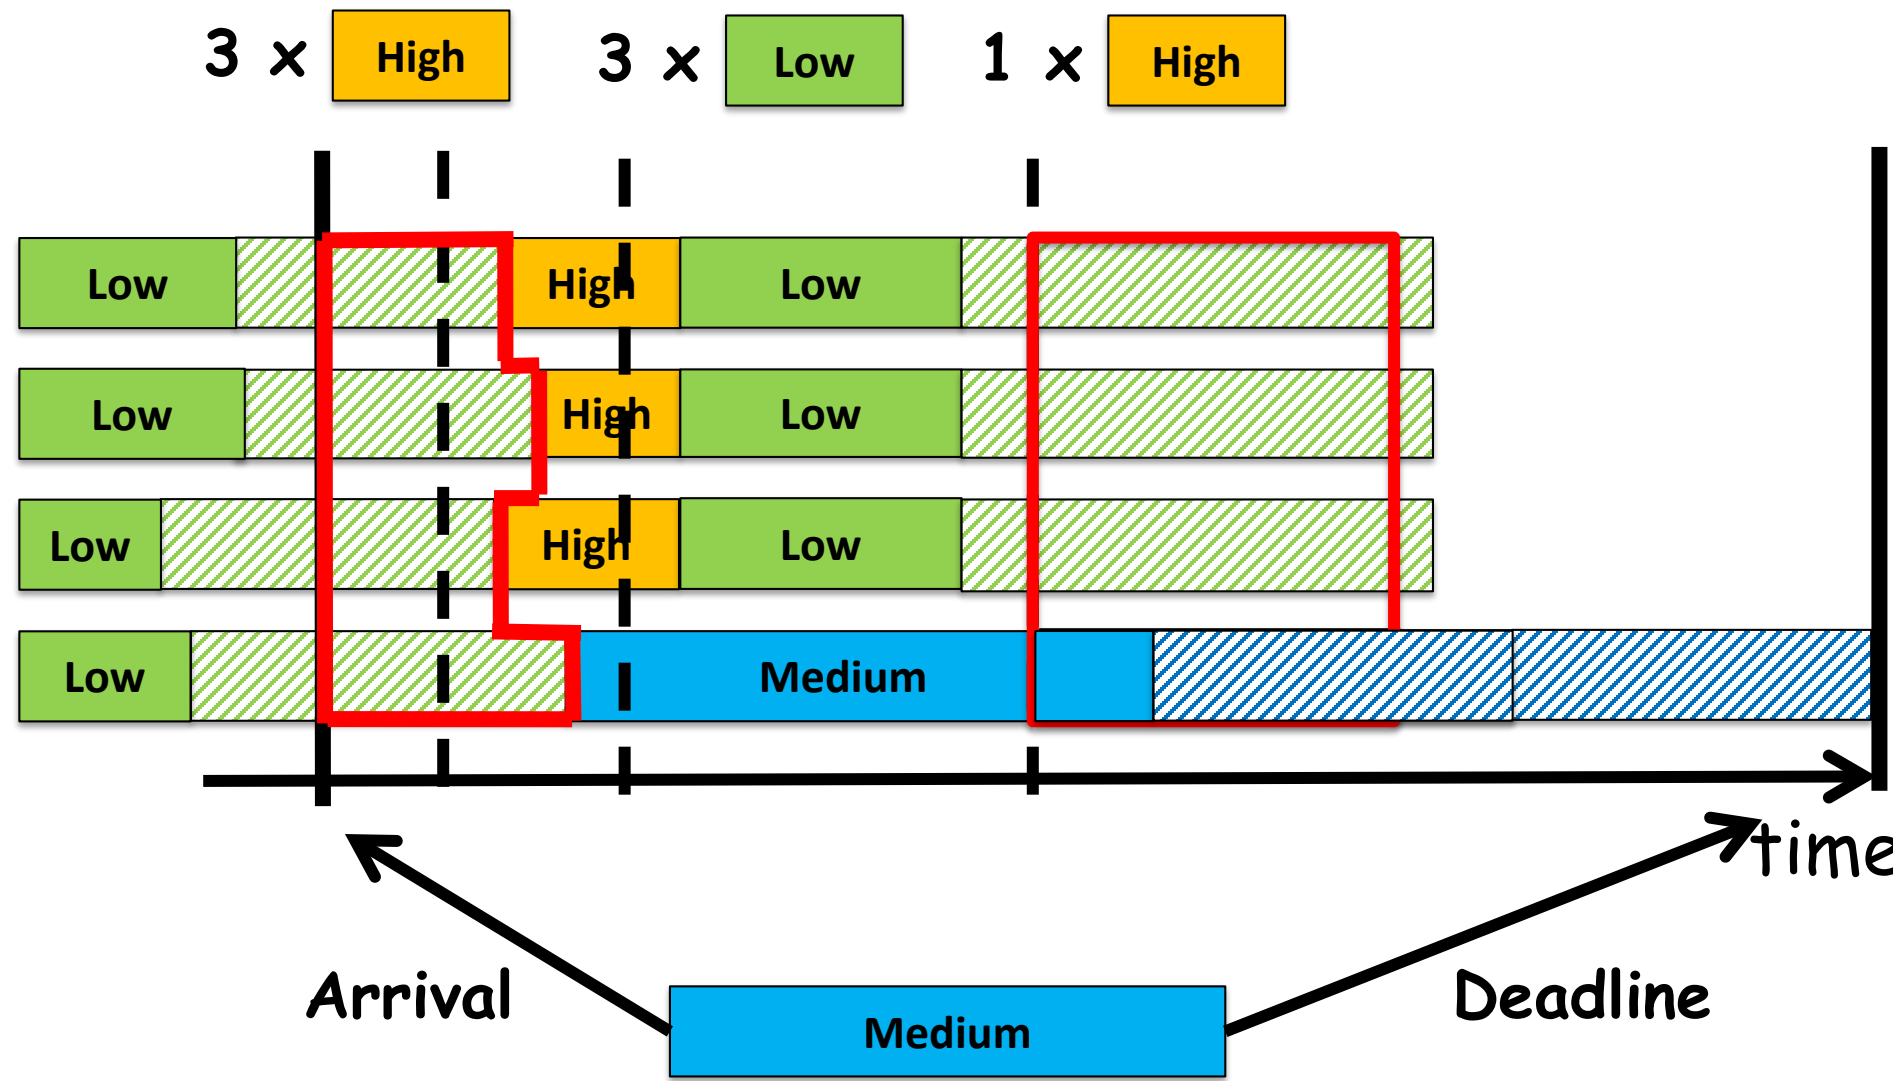
0 ← Maximum interference from HP workload in  $t$  time units

Deadline


t

time

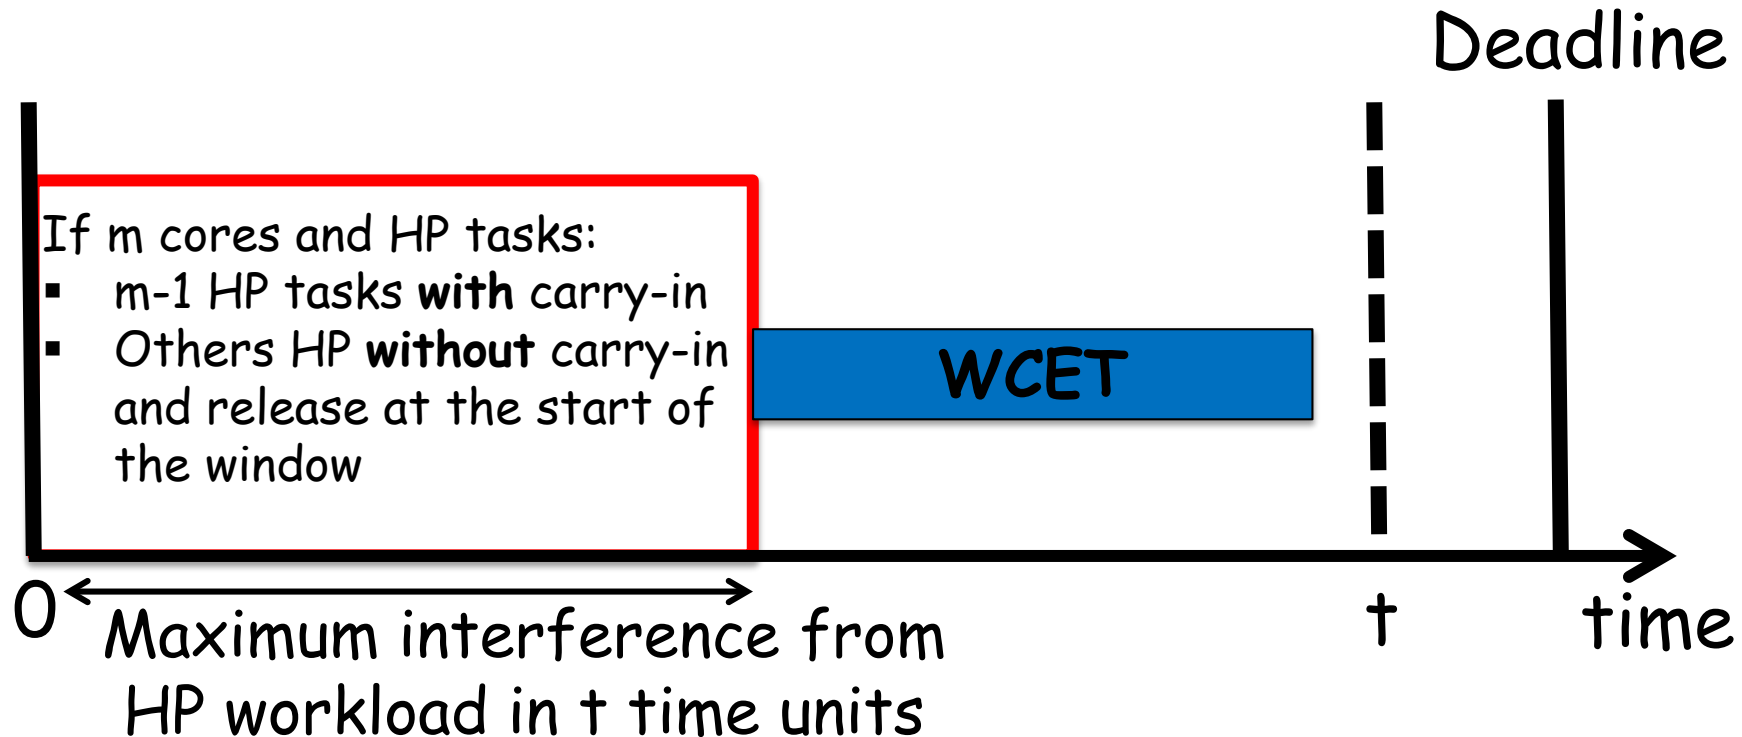
# And what about ADS?



# Let's focus on ADS

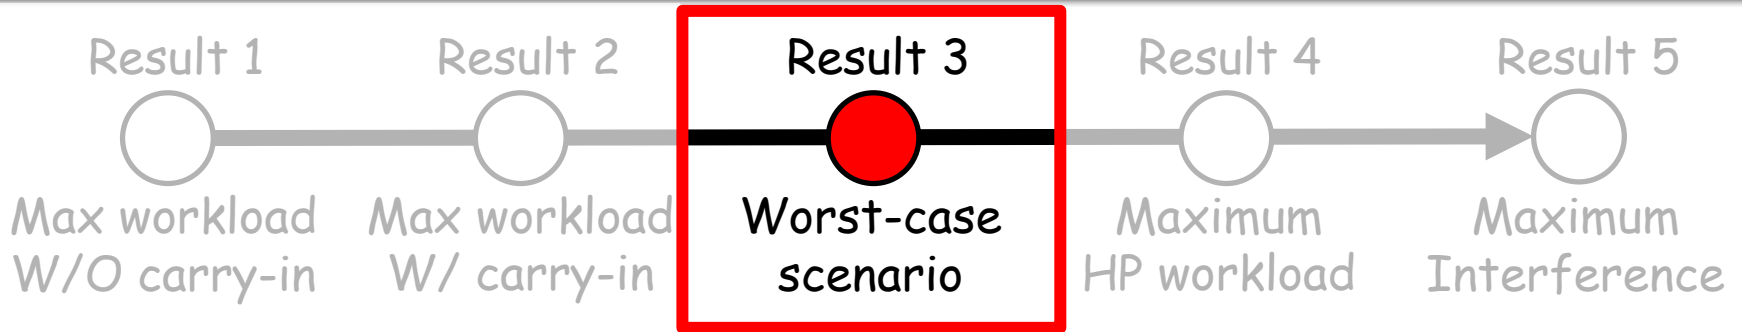


## Result 3: Worst-case interference scenario





# Let's focus on ADS

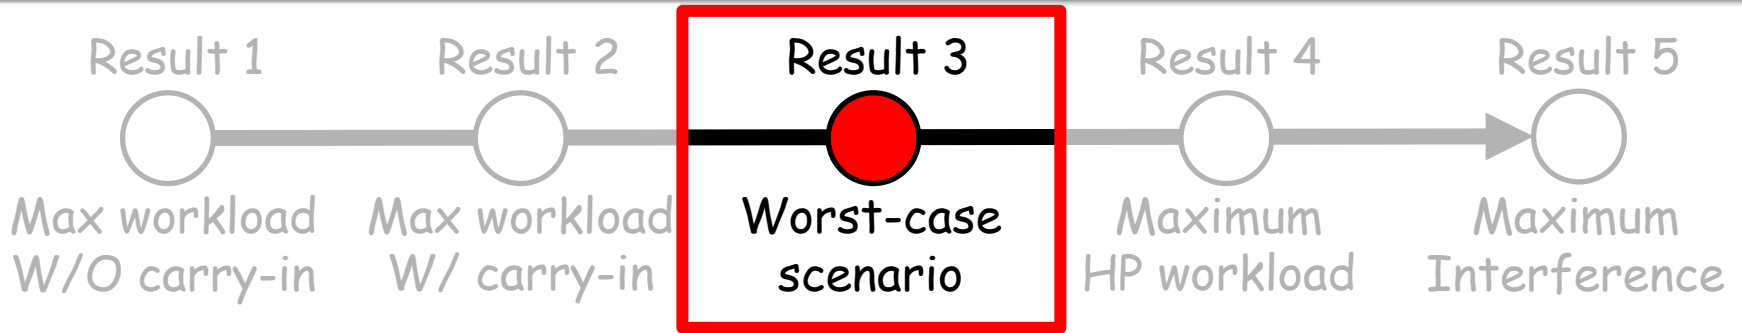


Result 3: Worst-case interference scenario

If  $m$  cores and HP tasks:

- $m-1$  HP tasks **with** carry-in

# Let's focus on ADS



Result 3: Worst-case interference scenario

<p>If <math>m</math> cores:</p> <ul style="list-style-type: none"> <li><math>m-1</math> HP</li> <li>1 LP</li> </ul>	<p>If <math>m</math> cores:</p> <ul style="list-style-type: none"> <li><math>m-2</math> HP</li> <li>2 LP</li> </ul>	<p>If <math>m</math> cores:</p> <ul style="list-style-type: none"> <li><math>m-3</math> HP</li> <li>3 LP</li> </ul>	<p>If <math>m</math> cores:</p> <ul style="list-style-type: none"> <li>0 HP tasks with carry-in</li> <li><math>m</math> LP carry-in jobs</li> </ul>
---	---	---	---


And so on...




# m - 4 HP carry-in tasks 4 LP carry-in jobs

Assume 4 carry-in jobs from LP tasks

The most compact the minimum workload generated by these 4 LP carry-in jobs




To compute the exact area requires to enumerate **all possible subsets** of 4 tasks out of the set of LP tasks



We propose three different techniques to upper-bound this area:

- Two of them are fairly straightforward
- The third one is more complex



Upper-bound on the area

 Maximum interfering workload from LP tasks

 Maximum interference from HP and LP tasks

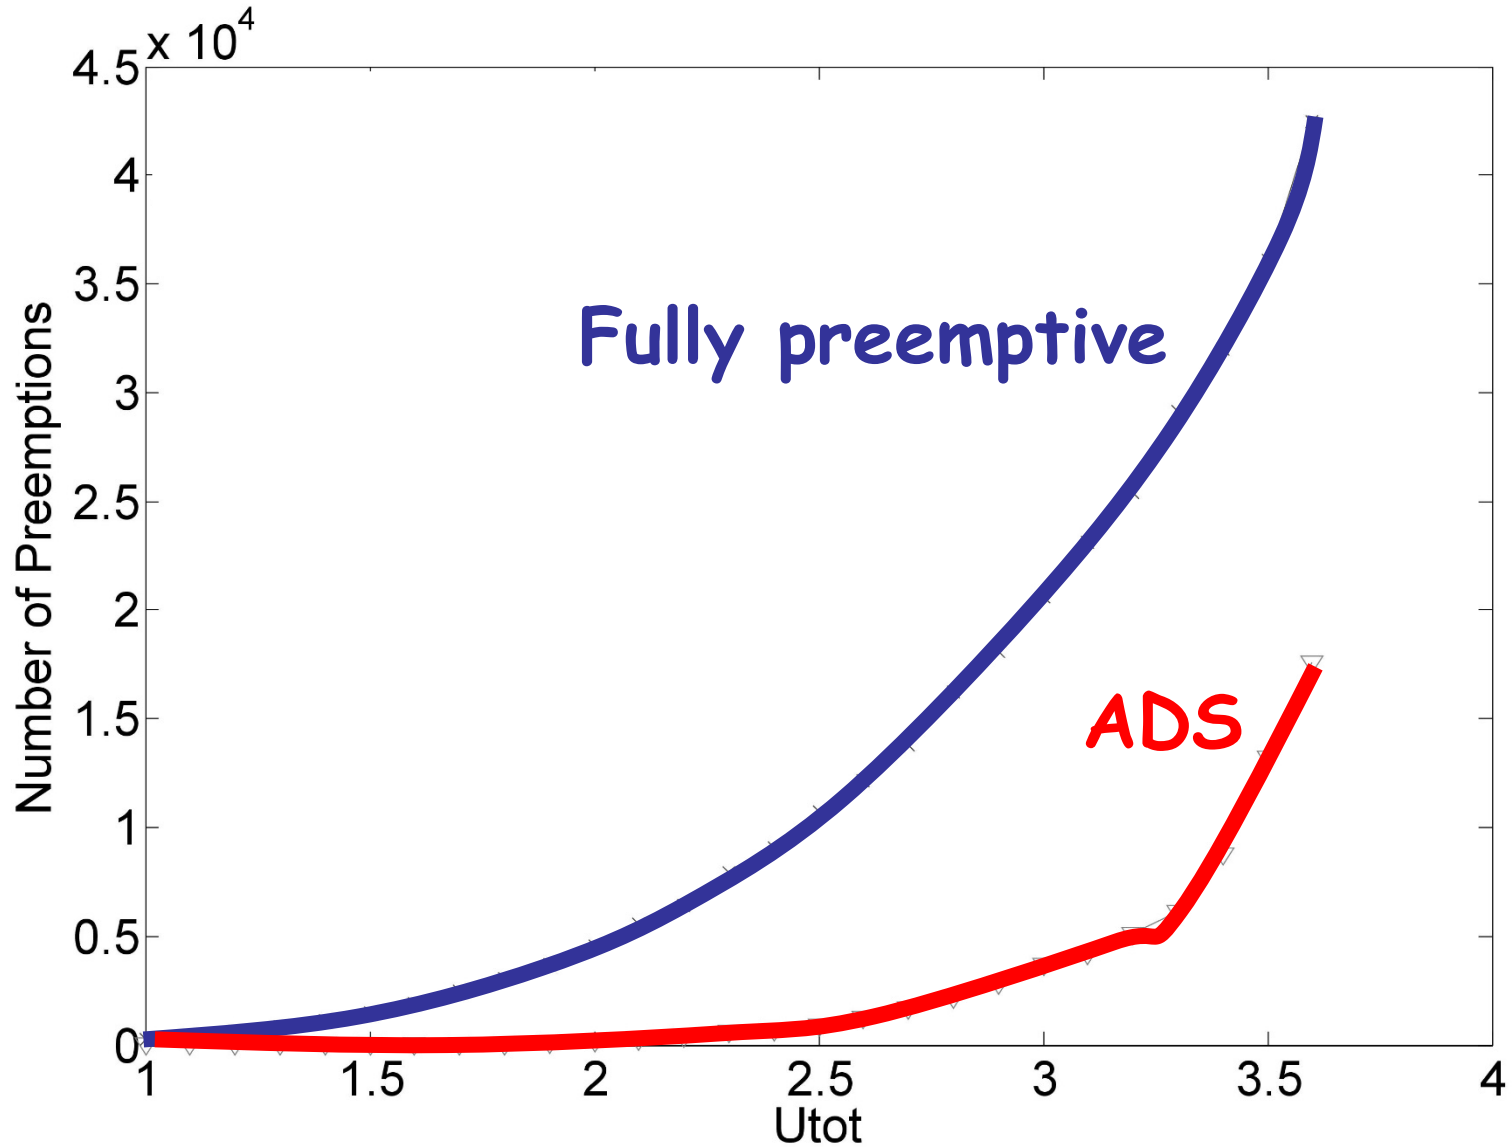
 Schedulability test

We designed a top-down approach (starting with the highest priority task) to assign to each task the maximum non-preemptive region length such that all the HP tasks are still schedulable.

 Reduce the number of preemption



# Some preliminary results



**Thank you for your attention 😊**