


# DUAL PRIORITY SCHEDULING IS NOT OPTIMAL

PONTUS EKBERG  
UPPSALA UNIVERSITY

ECRTS 2019  
STUTT GART, GERMANY



# DUAL PRIORITY SCHEDULING?

# DUAL PRIORITY SCHEDULING?

## DUAL PRIORITY ASSIGNMENT: A Practical Method for Increasing Processor Utilisation

A. Burns and A.J. Wellings

Department of Computer Science  
University of York, UK

### Abstract

*Static priority schemes have the disadvantage that processor utilisations less than 100% must be tolerated if a system is to be guaranteed off-line. By comparison earliest deadline scheduling can theoretically utilise all of a processors capacity, although in practice processor overheads are increased.*

$$U = \sum_{i=1}^n \frac{C_i}{T_i}$$

Test (1) converges, approximately, on a utilisation value of 0.69 for large  $n$ . For

(Euromicro Workshop on Real-Time Systems, 1993)

# HOW DOES IT WORK?

## Assumptions

- Implicit deadline periodic tasks
- Single preemptive processor

# HOW DOES IT WORK?

## Assumptions

- Implicit deadline periodic tasks
- Single preemptive processor

## Configurations

For each task  $\tau_i = (e_i, p_i)$ , assign

- a phase change point  $\delta_i \in \{0, \dots, p_i\}$
- two priority levels:  $\pi_i^1$  and  $\pi_i^2$

# HOW DOES IT WORK?


## Assumptions

- Implicit deadline periodic tasks
- Single preemptive processor

## Configurations

For each task  $\tau_i = (e_i, p_i)$ , assign

- a phase change point  $\delta_i \in \{0, \dots, p_i\}$
- two priority levels:  $\pi_i^1$  and  $\pi_i^2$



$\leftarrow p_i \rightarrow$

# HOW DOES IT WORK?


## Assumptions

- Implicit deadline periodic tasks
- Single preemptive processor

## Configurations

For each task  $\tau_i = (e_i, p_i)$ , assign

- a phase change point  $\delta_i \in \{0, \dots, p_i\}$
- two priority levels:  $\pi_i^1$  and  $\pi_i^2$



# HOW DOES IT WORK?


## Assumptions

- Implicit deadline periodic tasks
- Single preemptive processor

## Configurations




For each task  $\tau_i = (e_i, p_i)$ , assign

- a phase change point  $\delta_i \in \{0, \dots, p_i\}$
- two priority levels:  $\pi_i^1$  and  $\pi_i^2$









# COMPARING WITH FP AND EDF



# COMPARING WITH FP AND EDF



# COMPARING WITH FP AND EDF



# IS DUAL PRIORITY SCHEDULING OPTIMAL?

...  
priorities be found:

The answer to the first takes the form of a conjecture:

## Conjecture C1

For any task set with total utilisation less than or equal to 100% there exists a dual priority assignment that will meet all deadlines.

Analysis of this conjecture is given in section 4.

$\tau_3$				
$\tau_4$	20	6		
Total Utilisation	100%			
LCM	240			

Table 2: Task Set E2

With no task having a second phase, the system is not

## Conjecture 1 (Burns and Wellings, 1993)

Dual priority scheduling is *optimal* for implicit deadline periodic tasks.

(Sadly not)

WHY WAS THIS A DIFFICULT CONJECTURE?

## WHY WAS THIS A DIFFICULT CONJECTURE?

### Disproving the conjecture

Find a task set that is

- feasible ( $U \leq 1$ ), and
- not dual priority schedulable.

# WHY WAS THIS A DIFFICULT CONJECTURE?

## Disproving the conjecture

Find a task set that is

- feasible ( $U \leq 1$ ), and
- not dual priority schedulable.

Easy!

## WHY WAS THIS A DIFFICULT CONJECTURE?

### Disproving the conjecture

Find a task set that is

- feasible ( $U \leq 1$ ), and
- not dual priority schedulable.

Easy!

### But in practice...

- Almost all task sets are schedulable
- Evaluating the schedulability can be *very* costly



## WHY WAS THIS A DIFFICULT CONJECTURE?

### Disproving the conjecture

Find a task set that is

- feasible ( $U \leq 1$ ), and
- not dual priority schedulable.

Easy!

### But in practice...

- Almost all task sets are schedulable
- Evaluating the schedulability can be *very* costly

### Schedulability test

For every configuration (setting of  $\pi_i^1, \pi_i^2, \delta_i$ ), simulate the hyper-period until a deadline miss.

# HOW MANY CONFIGURATIONS ARE THERE?

## Assumptions

- All priority levels are unique  
⇒ A total of  $(2n)!$  permutations
- Phase change points  $(\delta_i)$  are integer  
⇒ A total of  $\prod_{i=1}^n (p_i + 1)$  combinations

# HOW MANY CONFIGURATIONS ARE THERE?

## Assumptions

- All priority levels are unique  
⇒ A total of  $(2n)!$  permutations
- Phase change points  $(\delta_i)$  are integer  
⇒ A total of  $\prod_{i=1}^n (p_i + 1)$  combinations

A task set  $\mathcal{T} = \{\tau_1, \dots, \tau_n\}$  has

$$(2n)! \times \prod_{i=1}^n (p_i + 1)$$

distinct configurations.

## A COUNTEREXAMPLE

Not dual priority schedulable

	$e_i$	$p_i$
$\tau_1$	8	19
$\tau_2$	13	29
$\tau_3$	9	151
$\tau_4$	14	197

hyper-period: 16 390 597  
utilization:  $\sim 0.9999971$

## A COUNTEREXAMPLE

Not dual priority schedulable

	$e_i$	$p_i$
$\tau_1$	8	19
$\tau_2$	13	29
$\tau_3$	9	151
$\tau_4$	14	197

hyper-period: 16 390 597  
utilization:  $\sim 0.9999971$

$$\text{\#configurations} = (2n)! \times \prod_{i=1}^n (p_i + 1) = 728\,082\,432\,000$$

## A COUNTEREXAMPLE

Not dual priority schedulable

	$e_i$	$p_i$
$\tau_1$	8	19
$\tau_2$	13	29
$\tau_3$	9	151
$\tau_4$	14	197

hyper-period: 16 390 597  
utilization:  $\sim 0.9999971$

$$\text{\#configurations} = (2n)! \times \prod_{i=1}^n (p_i + 1) = 728\,082\,432\,000$$

Simulating the full hyper-period for all configurations would take *hundreds of years* on my computer.

## THE SAVING GRACE

Most configurations lead to a deadline miss very early.

## THE SAVING GRACE

Most configurations lead to a deadline miss very early.

For the previous counterexample, less than

0.00019 %

of the hyper-period is simulated on average.



## THE SAVING GRACE

Most configurations lead to a deadline miss very early.

For the previous counterexample, less than

0.00019 %

of the hyper-period is simulated on average.

Simulation time is  $\sim 2.5$  days on an office computer.

## RECOGNIZING THE NEEDLE: RM+RM

$$\text{\#configurations} = (2n)! \times \prod_{i=1}^n (p_i + 1)$$

## RECOGNIZING THE NEEDLE: RM+RM

$$\text{\#configurations} = (2 \times n)! \times \prod_{i=1}^n (p_i + 1)$$

## RECOGNIZING THE NEEDLE: RM+RM

$$\text{\#configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n (p_i + 1)$$

### Definition: RM+RM

A dual priority configuration is called *RM+RM* if

- 1 phase 1 priorities  $(\pi_i^1)$  are RM
- 2 phase 2 priorities  $(\pi_i^2)$  are RM
- 3  $\max_i \{\pi_i^2\} \leq \min_i \{\pi_i^1\}$

## RECOGNIZING THE NEEDLE: RM+RM

$$\# \text{configurations} = (2 \times n!) \times \prod_{i=1}^n (p_i + 1)$$

### Definition: RM+RM

A dual priority configuration is called *RM+RM* if

- 1 phase 1 priorities  $(\pi_i^1)$  are RM
- 2 phase 2 priorities  $(\pi_i^2)$  are RM
- 3  $\max_i \{\pi_i^2\} \leq \min_i \{\pi_i^1\}$

### Conjecture 2 (George et al., 2014)

RM+RM is an optimal choice of priorities.

(Sadly not, even considering only point **1** above)

## RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)


$$\text{\#configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n (p_i + 1)$$

## RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\text{\#configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(i+1)}$$

# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)


$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(i+1)}$$






# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(i+1)}$$




# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(d_i + 1)}$$




# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(d_i + 1)}$$




# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(i+1)}$$




# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(d_i + 1)}$$




# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(i+1)}$$




# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(i+1)}$$



# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)


$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(i+1)}$$






# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(i+1)}$$




# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(i+1)}$$




# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(i+1)}$$




# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(q_i + 1)}$$




# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(q_i + 1)}$$




# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(i+1)}$$




# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(i+1)}$$



# RECOGNIZING THE NEEDLE: FDMS (FAUTREL ET AL., 2018)

$$\# \text{configurations} = \cancel{(2^n)!} \times \prod_{i=1}^n \cancel{(i+1)}$$



Conjecture 3 (Fautrel et al., 2018)

FDMS always finds optimal phase change points.

(Sadly not)



## RECOGNIZING THE NEEDLE

### A simple schedulability analysis strategy

- 1 Try *RM+RM* priorities with *FDMS*
- 2 If not successful, check configurations exhaustively

Doing 1 is *much* faster than 2, and works most of the time.

WHERE IS MY HAYSTACK?

# WHERE IS MY HAYSTACK?

## An ad hoc search space

- 4 tasks
- Utilization  $\in [0.99999, 1]$
- Periods chosen from the first 100 primes
- $\prod_{i=1}^n p_i \leq 35\,000\,000$

# WHERE IS MY HAYSTACK?

## An ad hoc search space

- 4 tasks
- Utilization  $\in [0.99999, 1]$
- Periods chosen from the first 100 primes
- $\prod_{i=1}^n p_i \leq 35\,000\,000$

Random task sets were tested from this search space until an unschedulable one was found. This is the breakdown:

	# task sets	% of explored search space
Schedulable with RM+RM using FDMS	129 823	$\sim 99.67\%$
Schedulable with other configurations	431	$\sim 0.33\%$
Unschedulable	1	$\sim 0.0008\%$

SO DUAL PRIORITY SCHEDULING IS NO GOOD?

## SO DUAL PRIORITY SCHEDULING IS NO GOOD?

### Fact

It took *26 years* and evaluating *millions of task sets* to find a single unschedulable one.

## SO DUAL PRIORITY SCHEDULING IS NO GOOD?

### Fact

It took *26 years* and evaluating *millions of task sets* to find a single unschedulable one.

But perhaps it behaves worse for larger task sets.

## SO DUAL PRIORITY SCHEDULING IS NO GOOD?

### Fact

It took *26 years* and evaluating *millions of task sets* to find a single unschedulable one.

But perhaps it behaves worse for larger task sets.

To fully exploit the apparent near-optimality we need efficient tests and methods for finding the right parameters.



## OPEN PROBLEMS

### Open problem 1

Can we efficiently determine if there exists a schedulable configuration?

## OPEN PROBLEMS

### Open problem 1

Can we efficiently determine if there exists a schedulable configuration?

In PSPACE, no lower bounds known

## OPEN PROBLEMS

### Open problem 1

Can we efficiently determine if there exists a schedulable configuration?

In PSPACE, no lower bounds known

### Open problem 2

If yes, can we efficiently find it?

## OPEN PROBLEMS

### Open problem 1

Can we efficiently determine if there exists a schedulable configuration?

In PSPACE, no lower bounds known

### Open problem 2

If yes, can we efficiently find it?

### Open problem 3

Can we efficiently evaluate a given configuration?

## OPEN PROBLEMS

### Open problem 1

Can we efficiently determine if there exists a schedulable configuration?

In PSPACE, no lower bounds known

### Open problem 2

If yes, can we efficiently find it?

### Open problem 3

Can we efficiently evaluate a given configuration?

In PSPACE, weakly NP-hard

## OPEN PROBLEMS

### Open problem 1

Can we efficiently determine if there exists a schedulable configuration?

In PSPACE, no lower bounds known

### Open problem 2

If yes, can we efficiently find it?

### Open problem 3

Can we efficiently evaluate a given configuration?

In PSPACE, weakly NP-hard

Simulation only works for periodic tasks (George et al., 2014)

## OPEN PROBLEMS

### Open problem 4

What is the utilization bound?

## OPEN PROBLEMS

### Open problem 4

What is the utilization bound?

Must be in the interval  $[\ln(2), 1)$



## OPEN PROBLEMS

### Open problem 4

What is the utilization bound?

Must be in the interval  $[\ln(2), 1)$

### Open problem 5

Is  $k$ -priority scheduling optimal for some constant  $k$ ?

## OPEN PROBLEMS

### Open problem 4

What is the utilization bound?

Must be in the interval  $[\ln(2), 1)$

### Open problem 5

Is  $k$ -priority scheduling optimal for some constant  $k$ ?

It is for  $n$ -priority scheduling, where  $n = |\mathcal{J}|$  (Pathan, 2015)

## OPEN PROBLEMS

### Open problem 4

What is the utilization bound?

Must be in the interval  $[\ln(2), 1)$

### Open problem 5

Is  $k$ -priority scheduling optimal for some constant  $k$ ?

It is for  $n$ -priority scheduling, where  $n = |\mathcal{T}|$  (Pathan, 2015)

### Open problem 6

What about rational phase change points?

## OPEN PROBLEMS

### Open problem 4

What is the utilization bound?

Must be in the interval  $[\ln(2), 1)$

### Open problem 5

Is  $k$ -priority scheduling optimal for some constant  $k$ ?

It is for  $n$ -priority scheduling, where  $n = |\mathcal{T}|$  (Pathan, 2015)

### Open problem 6

What about rational phase change points?

Can't be brute forced!

## WHAT IF NOT ALL PRIORITIES ARE UNIQUE?

### Precondition

We need to define the tie-breaking rule.

## WHAT IF NOT ALL PRIORITIES ARE UNIQUE?

### Precondition

We need to define the tie-breaking rule.

But most of them don't make sense!

## WHAT IF NOT ALL PRIORITIES ARE UNIQUE?

### Precondition

We need to define the tie-breaking rule.

But most of them don't make sense!

### Bad news

The number of priority orderings with ties are counted by the *Fubini numbers*, growing faster than the factorials.

## WHAT IF NOT ALL PRIORITIES ARE UNIQUE?

### Precondition

We need to define the tie-breaking rule.

But most of them don't make sense!

### Bad news

The number of priority orderings with ties are counted by the *Fubini numbers*, growing faster than the factorials.

### Result

The same counterexample remains un schedulable if priorities can be shared and FIFO or LIFO is used for ties.



## SPECIAL THANKS TO



Martina Maggio



Joël Goossens



Artifact evaluators

$\forall$  Thank you!



$\exists$  Questions?