

## Putting RUN into practice

### *Implementation and evaluation*

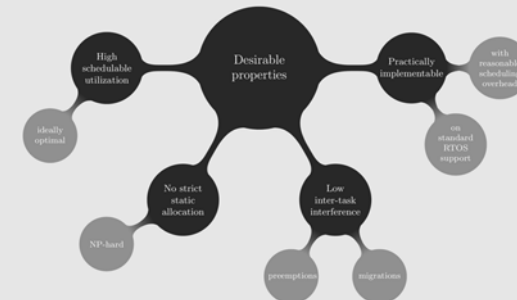
Davide Compagnin, Enrico Mezzetti and Tullio Vardanega  
University of Padua, Italy

26<sup>th</sup> EUROMICRO Conference on Real-time Systems (ECRTS)  
Madrid, 9 July 2014

This project and the research leading to these results has received funding from the European Community's Seventh Framework Programme [FP7 / 2007-2013] under grant agreement 611085

[www.proxima-project.eu](http://www.proxima-project.eu)

## Multiprocessor scheduling requisites



- ❑ Balancing good theoretical properties and viability requirements
  - **Low interference** and **high system utilization**
  - Standard **RTOS** support and reasonable **scheduling overheads**

3

09/07/2014

PROXIMA

## Outline

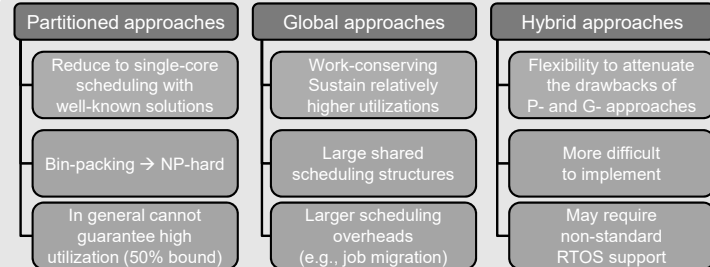
- ❑ Motivation
- ❑ Brief recap of Reduction to UNiprocessor
- ❑ RUN implementation and evaluation
- ❑ Conclusions and future work

2

09/07/2014

PROXIMA

## Multiprocessor scheduling state-of-the-art



- ❑ Reduction to UNiprocessor (RUN)
  - **Optimal** for *implicit-deadline periodic independent* tasks
  - Low interference with **few job migrations**
  - Reduces to P-EDF when a perfect partitioning exists

4

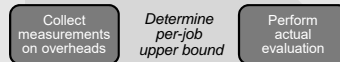
09/07/2014

PROXIMA



## Experimental setup

- ❑ LITMUS<sup>RT</sup> on an 8-core AMD Opteron™ 2356
- ❑ Collected measurements for RUN, P-EDF, G-EDF
  - Hundreds of automatically generated task sets
  - Harmonic and non-harmonic, with global utilization @ 50%-100%
  - Representative of small up to large tasks
- ❑ Two-step process
  - Preliminary empirical determination of overheads

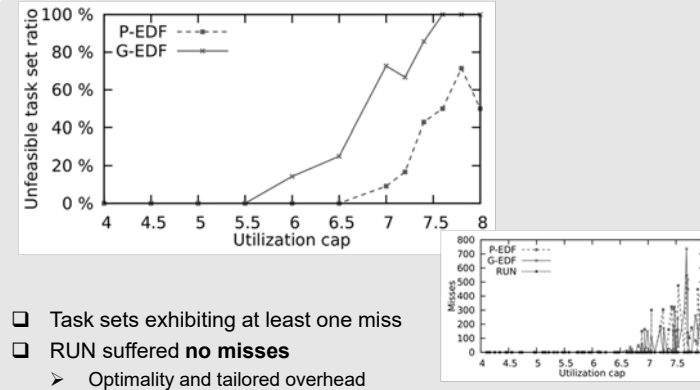


9

09/07/2014

PROXIMA

## Empirical schedulability




- ❑ Task sets exhibiting at least one miss
- ❑ RUN suffered **no misses**
  - Optimality and tailored overhead

11

09/07/2014

PROXIMA

## Primitive overheads and empirical bound



- ❑ Expectations confirmed
  - P-EDF needs lighter-weight scheduling primitives
- ❑ **Tree update (TUP)** triggered upon
  - Budget exhaustion event
  - Job release → REL includes TUP
- ❑ Empirical upper bound on RUN scheduling overhead
  - $OH_{RUN}^{Job} = REL + \widehat{SCHED} + CLK + k \times (TUP + \widehat{SCHED} + \max(PRE, MIG))$
  - $k = \lceil (3p + 1)/2 \rceil$  and  $\widehat{SCHED} = SCHED + CSW + LAT$ .


10

09/07/2014

PROXIMA

## Kernel interference

- ❑ Observing average preemptions and migrations




12

09/07/2014

PROXIMA

## Scheduling cost

### □ Average cost of core scheduling primitives



13

09/07/2014

PROXIMA

## Conclusions and future work

### □ Good news on RUN from this evaluation

- It can be **practically** and **efficiently** implemented
- It may exhibit **very modest kernel overhead**
  - Acceptable even on non-harmonic task sets
- It causes a tiny amount of **migrations**
  - Hence low inter-task interference

### □ Essential improvements

- Handle *sporadic* task sets
- Allow sharing of *logical* resources

### □ Further work

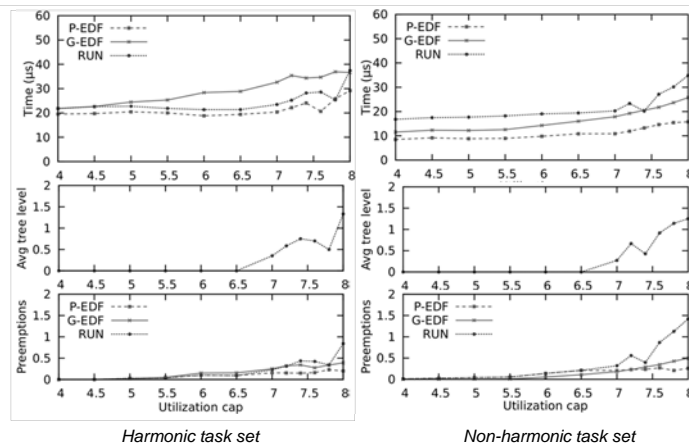
- Better understanding of the role of **packing policies**
  - Affecting the reduction tree, hence preemptions/migrations
- Further **comparisons** against other optimal solutions
  - High interest in *Quasi-Partitioned Scheduling* (QPS)

15

09/07/2014

PROXIMA

## Per-job scheduling overhead

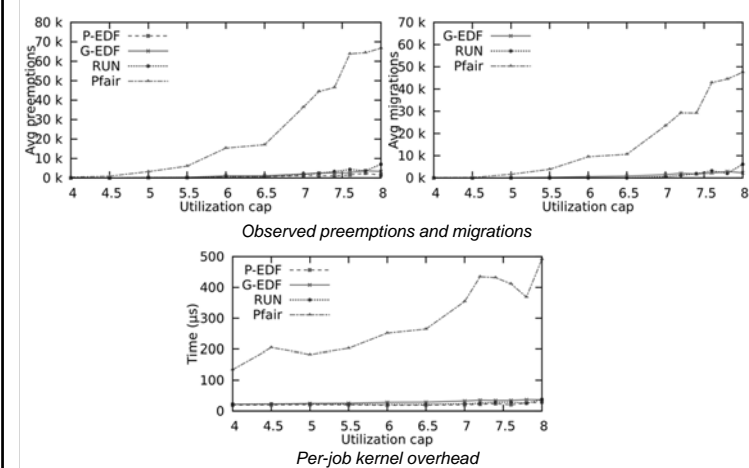


14

09/07/2014

PROXIMA

## Evaluation against S-PD<sup>2</sup>



16

09/07/2014

PROXIMA

## Reduction tree at run time

