Efficient Partitioning of Sporadic Real-Time Tasks with Shared Resources and Spin Locks

Alexander Wieder Björn B. Brandenburg

Max Planck Institute for Software Systems (MPI-SWS)



 Max
 8th IEEE International Symposium on Industrial Embedded Systems

 Institute
 Porto, Portugal

 for
 19.06.2013

Motivation



- space, weight and power constraints
- tasks with real-time requirements
- multi-core architectures
- shared resources protected by spin locks

Motivation



Our Focus: Processor Efficiency

Make best use of multi-core architectures despite shared resources protected by spin locks

Example: Autosar



multicore architectures:





partitioned fixed-priority scheduling

Example: Autosar



multicore architectures:





partitioned fixed-priority scheduling

shared resources:

- sensors
- communication bus
- kernel objects



global resources: non-preemptable spin locks

local resources: Priority Ceiling Protocol (PCP)

T₁ T₂ T₃ T₄ l₁ l₂ Assignment Algorithm

core 1 core 2 T_1 T_3 T_4 T_2 T_2 T_4 core 4 l_1 l_2

sporadic real-time tasks



shared resources protected by **spin locks**





Challenge: Task sets using spin locks



Challenge: Task sets using spin locks

Suppose T_4 assigned to core 3 instead...



Challenge: Task sets using spin locks



Challenge: Task sets using spin locks

Remote task can miss deadline!



How efficient are prior heuristics...?

This Paper

Observation

Contribution

Part I

Optimality matters! with shared resources, potential wasted by prior heuristics

This Paper Contribution Observation Part II Part I **Optimality matters! Optimal** ILP-based partitioning scheme for with shared resources, task sets with shared potential wasted by prior heuristics resources

This Paper Contribution Observation Part II Part I **Optimality matters! Optimal** ILP-based partitioning scheme for with shared resources, task sets with shared potential wasted by prior heuristics resources

Prior sharing-aware heuristics are complicated and brittle

This Paper **Contribution** Observation Part I Part II **Optimal** ILP-based **Optimality matters!** with shared resources, partitioning scheme for potential wasted by task sets with shared prior heuristics resources

Part III

Prior sharing-aware heuristics are complicated and brittle Greedy Slacker: simple and robust heuristic

Task Model

- sporadic tasks: T_i : (e_i, d_i, p_i)
- constrained deadlines: $d_i \leq p_i$
- shared resources accessed in mutual exclusion
- coordinating resource access:
 global: non-preemptable FIFO spinlocks
 local: SRP (blocking equivalent to PCP)

Task Model

- sporadic tasks: T_i : (e_i, d_i, p_i)
- constrained deadlines: $d_i \leq p_i$
- shared resources accessed in mutual exclusion
- coordinating resource access:
 - global: non-preemptable FIFO spinlocks
 local: SRP (blocking equivalent to PCP)

Multiprocessor Stack Resource Policy [1]

19

[1] P. Gai, G. Lipari, and M. D. Natale, "Minimizing memory utilization of real-time task sets in single and multi-processor systems-on-a-chip," in Proc. RTSS, 2001.

Part I How efficient are prior heuristics?

Task Assignment with Heuristics



Schedulability Experiments



Exploring Wasted Potential



akshmanan, D. de Niz, and R. Raikumar, "Coordinated task schedul

[1] K. Lakshmanan, D. de Niz, and R. Rajkumar, "Coordinated task scheduling, allocation and synchronization on multiprocessors," in Proc. RTSS, 2009.



schedulable



8 processors, 16 resources, critical section lengths in [1us,100us], periods in [3ms,33ms], 10% average task utilization



Part II An Optimal ILP-based Partitioning Scheme

Optimal Task Assignment



Optimality

What is an optimal partitioning scheme?

If a valid partitioning <u>under a given analysis</u> exists, a valid partitioning can be found.

Optimality

What is an optimal partitioning scheme?



Optimality

What is an optimal partitioning scheme?

If a valid partitioning <u>under a given analysis</u> exists, a valid partitioning can be found.

We use the MSRP blocking analysis from Gai, Lipari, Di Natale (2001).

Basic ILP Model

Integer Linear Programming model encodes for a **fixed number of processors**:

- task assignment
- priority assignment
- constraints to enforce valid assignments

Basic ILP Model

Integer Linear Programming model encodes for a **fixed number of processors**:

- task assignment
- priority assignment
- constraints to enforce valid assignments

We need to encode the MSRP blocking analysis into the ILP!

Encoding Blocking in ILP



Classic fixed-priority response-time analysis [1]

[1] S. Baruah and E. Bini, "Partitioned scheduling of sporadic task systems: An ILP based approach," in Proc. DASIP, 2008.

Encoding Blocking in ILP

Blocking analysis from Gai et al.:



Blocking under classic MSRP analysis from Gai et al.

Encoding Blocking in ILP

Blocking analysis from Gai et al.:


Blocking analysis from Gai et al.:



Blocking analysis from Gai et al.:





How can we express blocking in purely linear terms?

intorforonco

Transform to multiplication of variables with constants!

depends on *I*₁ priority/locality!

How long can T_1 's job spin?



spinning =



+

spinning = waiting for
$$l_1$$

waiting for l_2













Making ILP-based Partitioning Practical

In the real world, we also want to...



Making ILP-based Partitioning Practical In the real world, we also want to...



Making ILP-based Partitioning Practical In the real world, we also want to...



ILP Solving Overhead

4 processors average runtime [s] utilization 2.0 utilization 2.5 utilization 3.0

tasks

ILP Solving Overhead



ILP Solving Overhead



- task set size
- utilization
- resource contention

Only a one-time cost for exploiting wasted potential!

ILP Solving



solving time grows with:

- task set size
- utilization
- resource contention

Only a one-time cost for exploiting wasted potential!

Part III A Simple Sharing-Aware Partitioning Heuristic

Sharing-Aware Partitioning Heuristics

Prior Sharing-Aware Heuristics:

LNR-heuristic [1]

 Blocking-Aware Partitioning Algorithm (BPA) [2]

[1] K. Lakshmanan, D. de Niz, and R. Rajkumar, "Coordinated task scheduling, allocation and synchronization on multiprocessors," in Proc. RTSS, 2009.
[2] F. Nemati, T. Nolte, and M. Behnam, "Partitioning real-time systems on multiprocessors with shared resources," in Proc. OPODIS, 2010.



- identify connected components
- assign components
- if not possible, split
 - cost functions



- identify connected components
- assign components
- if not possible, split
 - cost functions



- identify connected components
- assign components
- if not possible, split
 cost functions

- identify connected components
- assign components
- if not possible, split
 - cost functions



- identify connected components
- assign components
- if not possible, split
 - cost functions



- identify connected components
- assign components
- if not possible, split
 - cost functions



- identify connected components
- assign components
- if not possible, split
 cost functions



- identify connected components
- assign components
- if not possible, splitcost functions



- identify connected components
- assign components
- if not possible, splitcost functions



- identify connected components
- assign components
- if not possible, split
 cost functions







Greedy Slacker

embarrassingly simple:

- disregard graph structure
- greedily try to maximize minimum slack

Greedy Slacker

embarrassingly simple:

disregard graph structure

greedily try to maximize minimum slack

> time until deadline is missed

Greedy Slacker

embarrassingly simple:

disregard graph structure

greedily try to maximize minimum slack

for each task T_i in order of increasing period: for each processor C_k : **compute slack when** T_i **assigned to** C_k if there is no C_r such that minimum slack \geq 0: **fail** else:

assign T_i to C_r s.t. minimum slack is maximized


No cost functions! Ignores graph structure! disregard graph structure

greedily try to maximize minimum slack

for each task T_i in order of increasing period: for each processor C_k : compute slack when T_i assigned to C_k if there is no C_r such that minimum slack > 0: fail else:

assign T_i to C_T s.t. minimum slack is maximized

roody Slacker Works with **any** blocking analysis.

No cost functions! Ignores graph structure! disregard graph structure

greedily try to maximize minimum slack

Can this possibly work?

else:

assign T_i to C_r s.t. minimum slack is maximized

Experimental Setup

Heuristics:

- sharing-oblivious (bin-packing)
- LNR-heuristic
- BPA
- Greedy Slacker

Experimental Setup

Heuristics:

- sharing-oblivious (bin-packing)
- LNR-heuristic
- BPA
- Greedy Slacker

Configuration:

- 8 processors
- 4 shared resources
- 10% average task utilization
- each resource accessed by 25% of tasks
- 100 samples

Resource Access Patterns

Unstructured



Resource Access Patterns

Unstructured



Structured

Resource Access Patterns

 T_4 I_3 L_2 Unstructured T_4 T_3 11 I_2 Structured Structured with T_4 T_3 T_2 T_1 global resources















Grouping of tasks and resources into functional components











Structured Resource Accesses with global resources



Grouping of tasks and resources into functional components, some resources access by all tasks

Structured Resource Accesses with global resources







Structured Resource Accesses with global resources



Summary

Optimal partitioning matters in the face of shared resources protected by spin locks.

Blocking due to spin locks in the MSRP can be expressed with purely linear expressions which allows using ILP techniques.

Summary

Fast and robust sharing-aware partitioning heuristic can be embarrassingly simple.

Thanks!