

12-2020

## Review of CPU Energy-Aware Parallel Real-Time Scheduling

Sebastien Pierre

University of Central Florida, Spierre3@Knights.ucf.edu

Find similar works at: <https://stars.library.ucf.edu/realtimesystems-reports>

University of Central Florida Libraries <http://library.ucf.edu>

This Report is brought to you for free and open access by the Department of Electrical and Computer Engineering at STARS. It has been accepted for inclusion in Recent Advances in Real-Time Systems by an authorized administrator of STARS. For more information, please contact [STARS@ucf.edu](mailto:STARS@ucf.edu).

---

### STARS Citation

Pierre, Sebastien, "Review of CPU Energy-Aware Parallel Real-Time Scheduling" (2020). *Recent Advances in Real-Time Systems*. 6.

<https://stars.library.ucf.edu/realtimesystems-reports/6>

# Review of CPU Energy-Aware Parallel Real-Time Scheduling

Sebastien Pierre

Real Time Systems - 2020

University of Central Florida

Course: EEE 4775 - Real Time Systems

**Abstract**—This is a review of journals in the area of using Energy-Aware/Energy efficient real-time systems of various tasks on different types of multi-core platforms (in this case, specifically DAG). The journals have discussed the milestones that have been achieved in the area of embedded systems and how their is need for energy efficient solutions to handle this advancement. These papers also include some of the first scheduling algorithms that take prioritizes energy awareness when it comes to both global scheduling and federated sch Three articles have been reviewed titled as follows: -

- 1) CPU Energy-Aware Parallel Real-Time Scheduling by Zhishan Guo and others
- 2) Energy-Efficient Real-Time Scheduling of DAG Tasks by Zhishan Guo, Ashikahmed Bhuiyan and others

The following text describes a brief overview of what these articles have reviewed in this field.

## I. CPU ENERGY-AWARE PARALLEL REAL-TIME SCHEDULING

### A. Authors

(Zhishan Guo, Ashikahmed Bhuiyan, Sezana Fahmida, Venkata P. Modekurthy, Nathan Fisher).

### B. Overview

This paper reviews advanced hard real-time scheduling of parallel tasks while also trying to find ways to minimize the CPU energy consumption these schedules could have on on multi-core embedded systems. In this paper, tasks are represented in a directed acyclic graphs (DAG) with nodes that indicate various threads of execution and also with edges that indicated their dependencies. This paper claims to have developed a technique that will be able to determine the execution speeds of the nodes within their DAGs that both fulfill all of the task deadlines while

also having the lowest overall energy consumption while doing so.

The writers of this paper focused on this topic because they noticed that previous work on real-time parallel tasks, were mainly focused on real-time scheduling without making sure that there was efficient energy consumption. The following key areas formed the main text of this paper.

- Introduction/Background
- Energy-Aware Techniques
- Results

### C. Introduction/Background

1) *Introduction:* In this section of the paper, they present the case for why Energy-efficiency is an important factor in modern day embedded systems. With the rise of various technology such as mobile phones, robotics, cars, etc that rely on batteries or other energy sources, they argue that energy consumption is now more important than ever especially since many of the technological systems are so advanced that they require the use of parallel tasks. With this, they focused this paper on minimizing CPU energy consumption by means of dynamic voltage and frequency scaling (DVFS).

2) *Background:* This section of the paper briefly covers the previous related work that was done on this concept. They show that even though these papers were focused on the concept of "energy efficiency within multiprocessors", that they weren't the same since they focused on sequential tasks while this paper focuses on parallel tasks which are more challenging to make energy efficient than sequential due to sequential tasks only being able to run on one core at a time. The paper then goes on to give a brief educational description

on how parallel tasks work while also describing the model they followed while creating this paper. They also introduce the power/energy model they used to judge the energy efficiency of the various tasks and techniques they will use throughout the paper. Finally this section concludes with them briefly going over the two scheduling algorithms they seek to make energy efficient which are federated and global scheduling.

#### D. Energy-Aware techniques

1) *Federated Scheduling*: This section of the paper describes their first approach for minimizing energy consumption on multi-core processors with real-time parallel tasks in federated scheduling. Their game plan to making federated scheduling energy efficient is by trying to regulate the frequencies (speed) throughout the cores and various nodes of the same task. They believe this method will work because energy consumption of the cores is a function of the frequencies (speed). They designed their algorithm with the Worst case execution requirement (WCER) of the nodes in mind. Overall their approach involves an optimized engine and DVFS. The section then goes on into detail about how their algorithm works and also includes various conditions, definitions, and equations they used to test out their new algorithm.

2) *Global Scheduling*: In this section of the paper they described how they plan on minimizing energy consumption when it comes to global scheduling. Just like in their approach to federated scheduling, they plan on achieving minimized energy consumption by regulating the frequencies (speed) across the core and different nodes of the same task. Which again, leads them to involving an optimized engine and also DVFS into the basic real-time scheduling policies such as EDF and DM in hopes of making them more energy aware. The paper then goes on to describe how they plan on making both global edf and global DM more energy aware by describing different formulas and convexifications they will use to make them so.

#### E. Results

Finally the paper concludes by simulating their new found algorithms with various tasks in a custom built simulator in MATLAB. They test algorithms to make sure they achieved the results

that they set out for and document the results that they pulled from these simulations. In the end they concluded that their results showed a strong foundation for why future researchers should adopt their algorithms in hope of making embedded systems more energy efficient.

## II. ENERGY-EFFICIENT REAL-TIME SCHEDULING OF DAG TASKS

### A. Authors

(Ashikahmed Bhuiyan, Zhishan Guo, Abusayeed Saifullah, Nan Guan, Haoyi Xiong)

### B. Overview

This paper "studies energy-aware real-time scheduling of a set of sporadic Directed Acyclic Graph (DAG) tasks with implicit deadlines"[2]. They are trying to meet all real-time constraints while also trying to identify the best task allocation/execution patterns that will make the average power consumption of the whole platform be minimized overall. They believe that prior to this paper, no other body of work addresses the power consumption issue that has arisen when scheduling multiple DAG tasks on multi-cores while also allowing for intra-task processor sharing. This paper claims that they were able to achieve simulation results that allowed for 60 to 68 percent energy saving when compared to the current DAG task schedulers. Here are the 3 main sections of this paper:

- Introduction/Background
- Suboptimal/Optimal federated scheduling
- Simulation/Results

### C. Introduction/Background

1) *Introduction*: In this section just like the last paper, they present the case for why Energy-efficiency is an important factor in modern day embedded systems. They described how energy consumption is a cornerstone within the designs of embedded systems that rely on batter power. Due to this they say that finding ways to be energy efficient and power-aware has been a focal point within the embedded systems research community. Doing so will lead to reduced energy consumption and subsequently reduced power bills (which is by nature cost-effective). The rest of this section

than goes on to describe the structure of this paper and what it will present. The main contributions this paper presents are these: (i) *We propose a multi-processor scheduling algorithm along with the power consumption issues for sporadic DAG tasks with implicit deadlines.* (ii) *Under the federated scheduling and task decomposition framework, our table-driven scheduler is shown to be optimal in the sense of average power consumption (i.e., named sub-optimal due to extra constraints included).* (iii) *We propose an efficient processor merging technique that is widely applicable for energy-efficiency improvements to most of the existing work on federated DAG task scheduling. We formally prove the NP-completeness of the problem, propose an approximation algorithm, and prove the upper bound of its approximation ratio.* (iv) *Simulations are conducted to verify the theoretical results and demonstrate the effectiveness of our algorithm.* [2]

2) *Background:* In this section of the paper, they introduce the system/model they will base this paper off of. They will consider a multi-core platform where each task is represented as a DAG with minimum inter-arrival distance between time units (periods). They then briefly give a background of all of the tasks, models, and equations they will use throughout the rest of this paper and reference. They will use the Power/Energy model and these equations to describe the energy consumption

$$P(s) = P_s + P_d(s) = B + \alpha s$$

$$E(s) = (P + s)/s = B/s + \alpha$$

( $B = \beta$ ;  $\alpha = \alpha$ ;  $\gamma = \lambda$ )

#### D. Sub-optimal/Optimal federated scheduling

1) *Sub-optimal federated scheduling:* In this section, they introduce what they deem as sub-optimal federated scheduling in terms of energy efficiency for DAG tasks. Within this federated approach to multi-core scheduling, each task is restricted to a single processor or has exclusive rights to every processor it will execute under. They described the goal of this section as to minimize the energy consumption of a single DAG task since each processor is dedicated to one processor. They seek to do this by using a commonly used technique known as task decomposition which simplifies the scheduling analysis of parallel real-time tasks. They then go on to use different

examples, equations, algorithms to flesh out their sub-optimal federated scheduling. They introduced plenty of theorems also to back up their research also.

2) *Optimal federated scheduling:* Since task decomposition transforms the parallel task into a set of sequential tasks the process tries to maximize the degree of parallelism. However, some of the processors may have poor energy efficiency due to power leakage. Which is why they believe that the previous solution they introduced in the previous section was sub-optimal. They thus introduce another solution in this section that will reduce the amount of power leakage. They seek to improve the overall energy efficiency of the scheduler they described by merging the workloads assigned to different processors onto a single one. This section then begins to dive into the details of how they plan on merging these processors. They include various figures and equations that display how merging processors can be beneficial to their overall goal they seek to achieve in this section.

#### E. Simulation/Results

In this section they show off the experiments they used to test out the mechanisms they developed throughout the paper. These experiments were designed to see whether or not they were more energy efficient than classical algorithms for DAG tasks. They also wanted to calculate just how much better they were in comparison. This section then goes on to mention the methods they used to create these experiments so that others can repeat their results. They conducted experiments under single merging of processors and also multiple merging of processors to test out the differences. With these experiments, they were able to draw many conclusions with one of them being that the average energy consumption is directly proportional to the average task utilization.[2] In the end this paper sought to minimize the overall power consumption when using DAG task and studied various ways they could achieve that.

### III. CONCLUSIONS

As technology becomes more advanced and complex, the need for energy efficiency amongst embedded systems will exponentially increase as many of these systems rely upon unreliable battery

systems. To keep things optimal and cost effective, researchers amongst the real-time systems community have been to spend much of their time studying/introducing various ways they make these systems more energy-efficient (especially when comes to parallel tasks). These two papers both contributed heavily to this effort by introducing their own mechanisms and algorithms in hopes of furthering the future work in that field.

#### REFERENCES

- [1] Ashikahmed Bhuiyan, Zhishan Guo, Abusayeed Saifullah, Nan Guan, and Haoyi Xiong. 2018. Energy-Efficient Real-Time Scheduling of DAG Tasks. *ACM Trans. Embed. Comput. Syst.* 17, 5, Article 84 (November 2018), 25 pages. DOI:<https://doi.org/10.1145/3241049> [1]
- [2] Z. Guo, A. Bhuiyan, D. Liu, A. Khan, A. Saifullah and N. Guan, "Energy-Efficient Real-Time Scheduling of DAGs on Clustered Multi-Core Platforms," 2019 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS), Montreal, QC, Canada, 2019, pp. 156-168, doi: 10.1109/RTAS.2019.00021. [2]