

# Energy-Efficient Parallel Real-Time Scheduling on Clustered Multi-Core

Ashikahmed Bhuiyan<sup>1</sup>, Di Liu<sup>2</sup>, Aamir Khan, Abusayeed Saifullah<sup>3</sup>, Nan Guan<sup>4</sup>, and Zhishan Guo<sup>5</sup>

**Abstract**—Energy-efficiency is a critical requirement for computation-intensive real-time applications on multi-core embedded systems. Multi-core processors enable intra-task parallelism, and in this work, we study energy-efficient real-time scheduling of constrained deadline sporadic parallel tasks, where each task is represented as a directed acyclic graph (DAG). We consider a clustered multi-core platform where processors within the same cluster run at the same speed at any given time. A new concept named *speed-profile* is proposed to model per-task and per-cluster energy-consumption variations during run-time to minimize the expected long-term energy consumption. To our knowledge, no existing work considers energy-aware real-time scheduling of DAG tasks with *constrained deadlines*, nor on a *clustered multi-core platform*. The proposed energy-aware real-time scheduler is implemented upon an *ODROID XU-3* board to evaluate and demonstrate its feasibility and practicality. To complement our system experiments in large-scale, we have also conducted simulations that demonstrate a CPU energy saving of up to 67 percent through our proposed approach compared to existing methods.

**Index Terms**—Parallel task, real-time scheduling, energy minimization, cluster-based platform, heterogeneous platform

## 1 INTRODUCTION

MULTI-CORE processors appear as an enabling platform for embedded systems applications that require real-time guarantees, energy efficiency, and high performance. Intra-task parallelism (a task can be executed on multiple cores simultaneously) enables us to exploit the capability of the multi-core platform, and facilitates a balanced distribution of the tasks among the processors. Such a balanced distribution leads to energy efficiency [1]. Directed Acyclic Graph (DAG) task model [2] is one of the most generalized workload model for representing deterministic intra-task parallelism. Recently, quite some effort has been spent on developing real-time scheduling strategies and schedulability analysis of DAG tasks, few to mention [2], [3], [4], [5], [6], [7], [8].

There are several real-world application that uses the DAG model. For example, the work in [3] studies problems related to scheduling parallel real-time tasks, modeled as DAG, on multiprocessor architectures. In a homogeneous computing environment, a low-complexity compile-time algorithm for scheduling DAG tasks is proposed in [9].

- A. Bhuiyan and Z. Guo are with the Department of Electrical and Computer Engineering, University of Central Florida, Orlando, FL 32816. E-mail: abon2@mst.edu, zsguo@ucf.edu.
- D. Liu is with Yunnan University, Kunming 650106, China. E-mail: d.liu@liacs.leidenuniv.nl.
- A. Khan is with Brainco Inc., Somerville, MA 02143. E-mail: aamir.khan@brainco.tech.
- A. Saifullah is with the Department of Computer Science, Wayne State University, Detroit, MI 48202. E-mail: saifullah@wayne.edu.
- N. Guan is with the Department of Computing, Hong Kong Polytechnic University, Hong Kong. E-mail: csguan@comp.polyu.edu.hk.

Manuscript received 24 July 2019; revised 6 Mar. 2020; accepted 1 Apr. 2020. Date of publication 14 Apr. 2020; date of current version 24 Apr. 2020.

(Corresponding author: Ashikahmed Bhuiyan.)

Recommended for acceptance by Q. Zheng.

Digital Object Identifier no. 10.1109/TPDS.2020.2985701

Another example would be systems that control asynchronous devices, such as the local-area network adapters that implement real-time communication protocols.

Since many of those applications are battery-powered, considering energy-efficient approaches for designing such a platform is crucial. Thanks to the fact that modern generation processors support dynamic voltage and frequency scaling (DVFS), where each processor can adjust the voltage and frequency at runtime to minimize power consumption, per-core energy minimization becomes possible during runtime. Despite the hardness of the problem [10], a significant amount of work has considered power minimization for non-parallel tasks on a multi-core platform (refer to [11] for a survey). Regarding parallel tasks, Guo *et al.* studied energy-efficient real-time scheduling for DAG tasks as an early research effort [12]. They adopted the federated scheduling and task decomposition framework [2] for minimizing system energy consumption via per-core speed modulation. As the only step (that we are aware of) towards energy-aware scheduling of real-time DAG tasks, they targeted an exciting problem and laid some of the foundations of this work. However, the attention of [12] is restricted to implicit deadline tasks with a system model of per-core DVFS.

Unfortunately, per-core DVFS becomes inefficient as it increases the hardware cost [13]. For balancing the energy efficiency and the hardware cost, there is an ongoing trend to group processors into islands, where processors in the same island execute at the same speed. For example, a big LITTLE platform (e.g., ODROID XU-3 [14]) consists of high performance (but power-hungry) cores integrated into ‘big’ clusters and low-power cores into ‘LITTLE’ clusters. Such a platform executes several real-life applications with heavy computational demands (e.g., video streaming [15]) in an energy-efficient manner. Apart from the energy consumption issue, a multi-core platform enables task execution with

high-performance demand and tight deadlines, essential for computation-intensive real-time systems, e.g., autonomous vehicles [16].

Despite the urgent need, to our knowledge, no work has been done that considered the energy-efficient scheduling of DAG tasks in clustered multi-core platforms, where cores form a group of frequency/voltage clusters. Such kind of system balances the energy efficiency and hardware cost compared to the traditional (with individual frequency scaling feature) multi-core models. The scheduling problem becomes highly challenging on such platforms because:

- (i) The relationship between the execution time, frequency, and the energy consumption is nonlinear, making it highly challenging to minimize energy consumption while guaranteeing real-time correctness, i.e., none of the tasks miss their deadline.
- (ii) Existing solution (e.g., [12]) relies on the assumption that each processor can freely adjust its speed. That solution performs poorly as the assumption is no longer valid under a more realistic platform model considered in this paper.
- (iii) The speed of a cluster becomes unpredictable when shared by multiple tasks with sporadic release patterns.

*Contribution.* In this paper, we propose a novel technique for energy-efficient scheduling of constrained deadline DAG tasks in a cluster-based multi-core system. To the best of our knowledge, no work has investigated the energy-efficient scheduling of DAG tasks on such a cluster-based platform. It is also the first work that has addressed the power awareness issue considering constrained deadline DAG tasks. Specifically, we make the following contributions:

- We consider a more practical *cluster-based system model* where the cores must execute at the same speed at any time instant within each cluster.
- To better handle constrained deadlines, one need to capture the gaps between deadlines and upcoming releases, as well as handling sporadic releases. Considering a continuous frequency scheme, we first propose a novel concept of *speed-profile* to present the energy-consumption behavior for each task as well as each cluster, such that they could guide task partitioning in an energy-efficient manner. An efficient *greedy* algorithm is proposed to partition DAG tasks according to the constructed speed-profiles.
- We propose an approach to creating the speed-profile to adapt to the discrete frequency scheme. Also, we extend our approach to apply to the heterogeneous platform.
- To evaluate the effectiveness of our proposed technique, we implement it on the ODROID XU-3 board, a representative multi-core platform for embedded systems [14]. The experiments report that our approach can save energy consumption by 18 percent compared to a reference approach. For larger-scale evaluation, we perform simulations using synthetic workloads and compare our technique with two existing baselines [12], [17]. The simulation results demonstrate that our method can reduce energy consumption by up to 66 percent compared to the existing ones under the cluster-based platform setting.

*Organization.* The rest of the paper is organized as follows. Section 2 presents the workload, power, and platform models, and the problem statement. Section 3, describes the importance of creating a speed-profile for an individual task and the whole cluster. Section 4 discusses the approaches to create the speed-profile (considering both the continuous and discrete frequency mode) for each task. In this section, we also propose a greedy algorithm to allocate multiple tasks in the same cluster. Sections 5 and 6 presents the experimental and simulation results. Section 7 discusses related work including a detailed comparison with our existing work [12], [18]. Section 8 concludes this paper.

## 2 SYSTEM MODEL, PROBLEM STATEMENT, AND BACKGROUND

### 2.1 System Model and Problem Statement

*Workload Model.* We consider a set of sporadic parallel task denoted by  $\tau = \{\tau_1, \dots, \tau_n\}$ , where each  $\tau_i \in \tau$  ( $1 \leq i \leq n$ ) is represented as a DAG with a minimum inter-arrival separation (i.e., period) of  $T_i$  time units, and a relative deadline of  $D_i (\leq T_i)$  time units. An implicit deadline task has the same relative deadline and period, i.e.,  $D_i = T_i$ . As a DAG task, the execution part of task  $\tau_i$  contains a total of  $N_i$  nodes, each denoted by  $\mathcal{N}_i^j$  ( $1 \leq j \leq N_i$ ). A directed edge from  $\mathcal{N}_i^j$  to  $\mathcal{N}_i^k$  ( $\mathcal{N}_i^j \rightarrow \mathcal{N}_i^k$ ) implies that execution of  $\mathcal{N}_i^k$  can start if  $\mathcal{N}_i^j$  finishes for every instance (precedence constraints). In this case,  $\mathcal{N}_i^j$  is called a parent of  $\mathcal{N}_i^k$  ( $\mathcal{N}_i^k$  is a child of  $\mathcal{N}_i^j$ ). A node may have multiple parents or children. The *degree of parallelism*,  $M_i$ , of  $\tau_i$  is the number of nodes that can be simultaneously executed.  $c_i^j$  denotes the execution requirement of node  $\mathcal{N}_i^j$ .  $C_i := \sum_{j=1}^{N_i} c_i^j$  denotes the worst case execution requirement (WCET) of  $\tau_i$ .

A *critical path* is a directed path with the maximum total execution requirements among all other paths in a DAG.  $L_i$  is the sum of the execution requirements of all the nodes that lie on a critical path. It is the minimum make-span of  $\tau_i$ , i.e., in order to make  $\tau_i$  schedulable, at least  $L_i$  time units are required even when number of cores is unlimited. Since at least  $L_i$  time units are required for  $\tau_i$ , the condition  $T_i \geq L_i$  (implicit deadline tasks) and  $D_i \geq L_i$  (constrained deadline tasks) must hold for  $\tau_i$  to be schedulable. A schedule is said to be feasible if upon satisfying the precedence constraints, all the sub-tasks (nodes) receive enough execution from their arrival times, i.e.,  $C_i$  within  $T_i$  (implicit deadline) or  $D_i$  (constrained deadline) time units. These terms are illustrated in Fig. 2a.

*Platform Model.* We consider a clustered multi-core platform, where processors within the same cluster run at the same speed (frequency and supply voltage) at any given time. Such additional restriction comparing to traditional multi-core platform makes the model more realistic in many scenarios. For example, our experiment is conducted on the ODROID XU-3 platform with one ‘LITTLE’ cluster of four energy-efficient ARM Cortex-A7 and one ‘big’ cluster of four performance-efficient ARM Cortex-A15. Note that we do not restrict the hardware-dependent energy parameters (e.g.,  $\alpha$ ,  $\beta$  and  $\gamma$  in the power model discussed below) to be identical across different clusters—these parameters can be derived using any curve-fitting method, e.g., [19].




Fig. 1. Comparison of the power model (Equation (1)) with experimental results in [24]. Here,  $\alpha = 1.76 \text{ Watts/GHz}^3$ ,  $\gamma = 3$ , and  $\beta = 0.5 \text{ Watts}$ . This figure is adopted from [20].

**Energy Model.** Assuming frequency (speed) of a processor at a specific instant  $t$  is  $s(t)$  (in short, denoted as  $s$ ), then its power consumption  $P(s)$  can be calculated as

$$P(s) = P_s + P_d(s) = \beta + \alpha s^\gamma. \quad (1)$$

Here,  $P_s$  and  $P_d(s)$  respectively denote the static and dynamic power consumption. Whenever a processor remains on, it introduces  $P_s$  in the system (due to leakage current). Switching activities introduce  $P_d(s)$  which is frequency dependent and represented as  $\alpha s^\gamma$ . Here, the  $\alpha > 0$  depends on the effective switching capacitance [20];  $\gamma \in [2, 3]$  is a fixed parameter determined by the hardware;  $\beta > 0$  represents the static part of power consumption. From this model, the energy consumption over any given period  $[b, f]$  is calculated as  $E(b, f) = \int_b^f P(s(t)) dt$ .

Our motivation behind selecting this power model comes from the fact that it complies with many existing works in the community, few to mention [10], [12], [18], [20], [21], [22], [23]. Beside this, recently this model was shown to be highly realistic by showing its similarity with actual power consumption [21]. Fig. 1 shows comparison between the original power consumption results from [24] and the power model in Equation (1).

**Assumptions.** In this paper, we make the following assumptions: (i) we focus on CPU power consumption, and (ii) Dynamic power management (DPM) is not considered. Appendix B, which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TPDS.2020.2985701>, provides the details behind these assumptions, their impacts, and some hints to overcome the drawbacks.

**Problem Statement.** Considering a constrained deadline sporadic DAG task-set on a clustered multi-core platform, we focus on finding a correct scheduling strategy, while the CPU power consumption is minimized.

## 2.2 Background and Existing Concepts

In this section, we describe some existing concepts and techniques for handling real-time parallel task scheduling, and that constitute an initial step for our proposed work.

**Task Decomposition.** The well-known task decomposition technique [2] transforms a parallel task  $\tau_i$  into a set of sequential tasks as demonstrated in Fig. 2b. Upon task decomposition, each node  $\mathcal{N}_i^l \in \tau_i$  is converted into an individual sub-task with its scheduling window (defined by its own release time and deadline) and execution requirement ( $c_i^l$ ). The allocation of release time and deadline respect all the dependencies (represented by edges in the DAG). Considering that a task is allowed to execute on an unlimited number of cores, starting from the beginning, a vertical line




Fig. 2. (a) A DAG,  $\tau_i$ , (b) transformed DAG  $\tau_i$  after applying task decomposition. Both of them are adopted from [12].

is drawn at every time instant where a node  $\mathcal{N}_i^l$  starts or ends. So the DAG is partitioned into several segments which may contain single/multiple thread(s). Threads assigned to the same segment share equal amount of execution length; e.g.,  $\mathcal{N}_i^3$ ,  $\mathcal{N}_i^4$ , and  $\mathcal{N}_i^5$  all have 2-time units assigned to the 3rd segment, as demonstrated in Fig. 2b.

**Segment Extension.** The deadline for each node via task decomposition may be unnecessarily restrictive, e.g., the decomposition of the DAG in Fig. 2a will restrict  $\mathcal{N}_i^3$  within the 2nd and 3rd segment. To eliminate such unnecessary restriction and allow  $\mathcal{N}_i^3$  to execute in the 4th segment, segment extension should be applied, e.g., the green rectangle for  $\mathcal{N}_i^3$  in the 4th segment in Fig. 2b.

**Intra-Task Processor Merging.** After applying task decomposition and segment extension upon a DAG task  $\tau_i$ , some of these cores (where  $\tau_i$  is allocated) can be very lightly loaded. Those core cause massive leakage power consumption in the long run and should be avoided when necessary. Intra-task merging [12] seeks to merge those cores to gain overall energy efficiency by reducing the total number of active cores. For example, in Fig. 2b, the third core (executing  $\mathcal{N}_i^5$ ) is lightly loaded, and thus it is better to merge all the execution into the second core and shut it off completely. Such a reduction on the number of active cores minimizes leakage power consumption (see Equation (1) and Fig. 2 in [12]) as well as the total number of clusters.

## 3 SPEED-PROFILE FOR TASK AND CLUSTER

This section discusses how different tasks share a cluster where all processors in a cluster execute at the same speed. When multiple tasks share a cluster, they may not align well due to sporadic releases and different periods. In a cluster-based platform, the processor having the maximum speed dominates the others in the same cluster. Hence, existing energy-saving techniques may perform poorly in a cluster-based platform. To tackle this problem, we propose a new concept called *speed-profile*. We provide the definition of speed-profile and its motivation in Section 3.1. Section 3.2 describes how speed-profiles are handled when two tasks are partitioned into the same cluster.

### 3.1 Speed-Profile for Each DAG

Interesting energy-saving techniques (e.g., segment extension) have been proposed in [12] for the implicit deadline tasks. For the constrained deadline tasks, this technique becomes incompetent because of the non-negligible idle gaps between the task deadline and its next release. For example, consider the task  $\tau_i$  in Fig. 2b with  $D_i = 10$  and  $T_i = 12$ . Segment extension can stretch  $\mathcal{N}_i^3$  to the end of the 4th segment but cannot utilize the idle time of 2 units. Besides, the sub-



optimal solution provided in [12] becomes *non-convex* (in a convex function, we can find the global maximum or minimum, for some variables of this function, which does not hold for a non-convex function) in a cluster-based platform (see Lemma 1).

**Lemma 1.** *In a cluster-based platform, the convex optimization problem constructed in [12] becomes non-convex.*

**Proof.** The following set of constraints ensure the real-time correctness for each node  $\mathcal{N}_i^l \in \tau_i$ , i.e.,  $\mathcal{N}_i^l$  receives enough time to finish execution within its scheduling window

$$\forall l : \mathcal{N}_i^l \in \tau_i :: \sum_{j=b_i^l}^{d_i^l} t_j^c s_{i,j} \geq c_i^{\mathcal{N}_i^l}. \quad (2)$$

We introduce the following inequalities to bound the total length for all segments in task  $\tau_i$

$$\sum_{j=1}^Z t_j^c \leq T_i. \quad (3)$$

Any value of execution speed and segment length ensures real-time correctness if Equations (2) and (3) are respected. However, the work in [12] considered that the execution speed of a node,  $\mathcal{N}_i^l$ , is *constant* within its scheduling window (from  $b_i^l$  to  $d_i^l$ ), and can be represented by a function of nodes execution requirement and its scheduling window. Also, the work in [12] considered that a single DAG executes at a time, and, hence the execution speed of a node is not affected by the execution speed of other nodes (of other tasks). In this work, we consider the cluster-based platform, and the execution speed of a node depends on the execution speed of other nodes (of other tasks) in the same cluster. As a result, we cannot express the execution speed of a node as a function of its execution requirement, resulting in quadratic inequality constraints (Equation (2)). This makes the optimization problem *non-convex*.  $\square$

Due to the characteristics of a clustered platform, at each instant, all cores in a cluster must execute at the speed of the fastest one. If these tasks are not well aligned (concerning their execution speed), the cluster as a whole will perform poorly (w.r.t. energy efficiency). Assigning tasks with similar speed shape on the same cluster may not be an energy efficient option (due to their sporadic releases pattern). Fig. 3 and Example 1 demonstrates one such scenario.

**Example 1.** In this example, we describe how the sporadic arrival pattern of a task influences the energy efficiency of the whole cluster. Consider two tasks  $\tau_1$  and  $\tau_2$  with the predefined necessary speed of execution on two processors each, to be partitioned on to the same cluster (of four processors). Fig. 3a shows the synchronous release case, where the whole cluster could run at 0 speed between [3,4) and [7,8). While Fig. 3b shows the scenario when  $\tau_1$ 's initial release is delayed by one-time unit, where the whole cluster will need to run at a higher speed (of 0.8) most (75 percent) of the time and thus consumes more energy.

In this example, from  $\tau_2$ 's perspective, direct energy reduction with existing per-task WCET based techniques




Fig. 3. When two tasks  $\tau_1$  and  $\tau_2$  with fixed speed patterns each are partitioned on to the same cluster, the resultant speed pattern ( $\tau_{12}$ ) of the cluster may vary for their ( $\tau_1$  and  $\tau_2$ ) different release offsets. In order to satisfy platform model restrictions while guaranteeing the correctness, the processors (of the same cluster) must run at the maximum/larger of the two individual speeds at each instant.

may not help much, as it may be another task dominating the speed of the whole cluster most of the time. The critical observation is that, due to the extra restriction of the more realistic platform model, the speed of a cluster is determined by the heavier DAG running on it, as well as how synchronous are the releases, which could be entirely random. Moreover, even a task finishes its execution early (say,  $\tau_2$  requires no execution over [5,7)), we may not be able to reduce the cluster speed at all.

To address this issue, we propose a novel concept of speed-profile to capture the energy consumption behavior of all possible alignment scenarios.

**Definition 1.** *The Speed-profile of a task describes the percentage/likelihood of all possible speeds that the task may execute at over a period. It is a random variable  $S$  with an associated probability function (PF)  $f_S(s) = \mathbb{P}(S = s)$ , where  $s$  is a speed from the finite set of possible speeds, and  $f_S(s)$  represents the portions of the time (likelihoods) when it is running at speed  $s$ .*

**Example 2.** Let us consider a task  $\tau_i$  with  $T_i = 15$  executing at a speed of 0.6 for 5 time units (not necessarily to be continual), and at a speed of 0.5 for the rest of the time. The speed-profile of the task is thus  $\mathcal{S}_i = \begin{pmatrix} 0.6 & 0.5 \\ 5/15 & 10/15 \end{pmatrix} = \begin{pmatrix} 0.6 & 0.5 \\ 0.33 & 0.67 \end{pmatrix}$ . At any specific time,  $t$ , there is about 33 percent probability that the cores are running at the speed of 0.6 unit and about 67 percent probability that the cores are running at the speed of 0.5 unit.

It is evident that from another task's point of view, the speed-profile provides probabilistic information on how the task of interest would restrict the lower bound to the speed of the cluster over time. As the alignment of releases between any two tasks is unknown, we assume in the future analysis that any phase difference is of equal chance over the long run.

**Remark 1.** The speed-profile  $\mathcal{S}_i$  of a given task  $\tau_i$  remains the same for an initial phase (release offset)  $\phi_i \geq 0$ . Regarding inter-task combinations, we assume uniform distribution for the phase of any task; i.e.,  $\phi_i \sim U[0, T_i)$ .

Section 4.1 details the calculation for task speed-profile. Here, we describe the calculation of the cluster speed-profile when two tasks are combined on to the same cluster.

### 3.2 Speed-Profile for the Cluster

As stated earlier, the property of the clustered platform and sporadic arrival pattern of a task makes the exact speed of the cluster unpredictable at a specific time instant (see Fig. 3 and Example 1). As a result, when two tasks  $\tau_i$  and  $\tau_j$  (with speed-profiles) are being considered allocating to the same cluster, we need to construct the merged speed-profile of the cluster (executing them both). To perform such calculation, we introduce a special  $\odot$  operator that takes the maximum of the two profiles on a probability basis.<sup>1</sup>

**Definition 2.** The special operator  $\odot$  operates on two (or more) random variables  $\mathcal{X}$  and  $\mathcal{Y}$ . During this operation, each entry  $\mathcal{X}_i \in \mathcal{X}$  is compared with each entry  $\mathcal{Y}_j \in \mathcal{Y}$  and the value  $\mathcal{Z}_{ij}$  is calculated as  $\mathcal{Z}_{ij} = \max(\mathcal{X}_i, \mathcal{Y}_j)$ , with a combined (multiplied) probability. If there are multiple occurrences of an entry, all of them are merged into a single entry, and their associated probability are summed together.

**Example 3.** Let  $\mathcal{S}_i = \begin{pmatrix} 6 & 5 \\ 0.4 & 0.6 \end{pmatrix}$ ,  $\mathcal{S}_j = \begin{pmatrix} 6 & 2 \\ 0.4 & 0.6 \end{pmatrix}$ . Then  $\mathcal{S}_i \odot \mathcal{S}_j = \begin{pmatrix} 6 & 6 & 6 & 5 \\ 0.16 & 0.24 & 0.24 & 0.36 \end{pmatrix} = \begin{pmatrix} 6 & 5 \\ 0.64 & 0.36 \end{pmatrix}$ .

Note that we allocate two different DAGs (with same/different periods) to the same cluster. The speed-profile indicates how long a DAG executes at different speeds within its deadline, i.e., the probability that a DAG executes at a specific speed. The task's period becomes irrelevant as speed-profile is a probability-based measurement. Once  $\tau_i$  and  $\tau_j$  are allocated to the same cluster, we use  $\mathcal{S}_{ij}$  to denote the speed-profile of the cluster (see Example 3).

In summary, energy minimization in a cluster-based platform is challenging because of sporadic release pattern and the idle gaps between a task deadline and its period. To tackle these problems, we have introduced the concept of speed-profile for both an individual task and a cluster where multiple tasks can be allocated.

## 4 TASK PARTITIONING ALGORITHM

The ultimate goal of the paper is to partition all DAGs into clusters, such that overall platform energy consumption is minimized. Recall that on a clustered multiprocessor platform, at a given instant, all processors in the same cluster must execute at the same speed. Due to this property of a cluster-based platform, if two tasks that are not well-aligned (in terms of execution speed) are allocated to the same cluster, it will result in reduced energy efficiency. So, we have proposed the concept of speed-profiles (refer to Section 3) which is a tool to measure the potential long-term energy saving of a cluster when partitioning any pair of DAGs into this cluster. So far we have discussed the importance of the concept of speed-profile but did not mention how to create them given a DAG task, which is the focus on Section 4.1. Then, Section 4.4 describes the task-to-cluster partitioning algorithm.

### 4.1 Creating the Speed-Profile of a Task

Given a DAG task  $\tau_i$ , we provide two approaches to create the speed-profile  $\mathcal{S}_i$ .

1. Although the appearance of the proposed operator is identical to [25], the calculation is quite different. This is due to the "larger value dominating" nature of the platform model considered in this paper.

*Approach A: Considering the Maximum Speed from all the Cores.* Upon applying the task decomposition, segment extension, and intra-task processor merging techniques (Section 2), some vital information (e.g., the speed of a core at a specific time and number of cores required) becomes available. This information plays a role in calculating the speed-profile  $\mathcal{S}_i$  of task  $\tau_i$ . At any time instant  $t$ , we consider the maximum speed from all the cores available. It ensures the sufficient speed so that even the heaviest node can finish execution within its scheduling window (defined after task decomposition). We consider constrained deadline (i.e.,  $D_i \leq T_i$ ), so the task must have to finish by  $D_i$  and rest of the time  $(T_i - D_i)$  it remains idle. For each segment  $j \in \tau_i$ , (summation of the length of these segments is equal to  $D_i$ ), we create a pair  $(s_{i,j}, p_{i,j})$ . For the  $j$ th segment,  $s_{i,j}$  and  $p_{i,j}$  respectively denote the maximum execution speed and the probability that the cluster will run at this speed. Let,  $M$  cores are allocated to  $\tau_i$ . At  $j$ th segment, we calculate  $s_{i,j}$  and  $p_{i,j}$  as follows:

$$s_{i,j} = \max_{k \leq M} \{s_{i,j,k}\}, p_{i,j} = \frac{t_j^c}{T_i}$$

Here,  $s_{i,j,k}$  denotes the speed of  $k$ th core at  $j$ th segment and  $t_j^c$  is the length of  $j$ th segment. The speed-profile  $\mathcal{S}_i$  will be

$$\mathcal{S}_i = \begin{pmatrix} s_{i,1} & s_{i,2} & \cdots & s_{i,z} & 0 \\ p_{i,1} & p_{i,2} & \cdots & p_{i,z} & (T_i - D_i)/T_i \end{pmatrix}.$$

The last pair reflects the fact that the core remains idle for the  $(T_i - D_i)$  time units at the end of each period.

**Example 4.** Consider a task  $\tau_i$  with  $T_i = 15$ ,  $D_i = 12$  and  $C_i = 6.5$ . Let, the task is partitioned into three segments of length 5, 7 and 3 time units respectively, where the processor is executing at a (maximum) speed of 0.6 in the first segment, speed of 0.5 in the second segment, and remain idle in the third segment. The speed-profile is

$$\mathcal{S}_i = \begin{pmatrix} 0.6 & 0.5 & 0 \\ 0.33 & 0.47 & 0.2 \end{pmatrix}.$$

Note that, if a cluster contains a single task  $\tau_i$ , then  $\mathcal{S}_i$  also represents the cluster speed-profile. If  $\tau_i$  and  $\tau_j$  (or more tasks) are executing on the same cluster, then the technique described in Section 3.2 needs to be applied before making any choices. The greedy choosing approach for task partitioning is detailed in Section 4.4.

*Approach B: A Single Speed Throughout.* Theorem 4 of [12] shares a valuable insight: *The total energy consumption (assuming processor remains on) is minimized in any scheduling window when execution speed remains uniform (the same) throughout the interval.* Motivated by it,<sup>2</sup> we propose another approach of selecting a single speed for a DAG task (job) during the whole duration from its release until its deadline.

2. Note that [12] considered that the speed remains constant within a scheduling slot for each processor. Also, they assumed per core speed scaling and calculated the speed within each scheduling slot through a convex optimization method. This paper considers the clustered platform where the objective function becomes *non-convex* (see Lemma 1) and thus the existing approach is inefficient.

In this approach, we consider the maximum workload (or the execution requirement) from all the cores available and determine the aggregated workload. Upon dividing the aggregated workload by the deadline, we get the desired single speed. Let  $M$  cores be allocated to task  $\tau_i$ . At  $j$ th segment, the execution requirement of the  $k$ th core is denoted by  $w_{i,j,k}$ , which is calculated as follows:

$$w_{i,j,k} = s_{i,j,k} \times t_j^c.$$

We determine the maximum execution requirement as follows:

$$w_{i,j} = \max_{k \leq M} \{w_{i,j,k}\}.$$

Let  $Z$  denotes the total number of segments in  $\tau_i$ . The maximum total workload  $w_i$  and the desired single speed  $s_i$  is calculated using the following equations:

$$w_i = \sum_{j=1}^Z w_{i,j}, \quad s_i = \frac{w_i}{D_i}. \quad (4)$$

Other than the idle pair  $(0, (T_i - D_i)/T_i)$ , we consider a single speed throughout the deadline so only a single pair  $(s_i, p_i)$  is required, where  $s_i = w_i/D_i$  and  $p_i = D_i/T_i$ .

**Example 5.** Consider the task described in Example 4 ( $T_i = 15$ ,  $D_i = 12$  and  $C_i = 6.5$ ). It must finish 6.5 unit of workloads within 12-time units. Using this approach its speed-profile is

$$S_i = \begin{pmatrix} 0.54 & 0 \\ 0.8 & 0.2 \end{pmatrix}.$$

**Lemma 2.** If a task  $\tau_i$  executes according to the speed-profile  $S_i$ , it guarantees real-time correctness.

**Proof.** It has been observed in [12] that the following constraint guarantees the real-time correctness

$$\forall l : \mathcal{N}_i^l \in \tau_i :: \sum_{k=b_i^l}^{d_i^l} t_k^c S_k^{M_i^l} \geq c_i \mathcal{N}_i^l. \quad (5)$$

Here,  $b_i^l$  and  $d_i^l$  denotes the release time and deadline of  $\mathcal{N}_i^l$ ,  $M_i^l$  denotes the node-to-processor mapping and  $S_k^{M_i^l}$  is the speed of the processor (where  $\mathcal{N}_i^l$  is allocated) at  $k$ th segment. Unlike to [12], at any time instant  $t$ , we choose either the maximum speed from all the cores running on the same cluster (Approach A) or a single speed that can guarantee the maximum execution requirement for the whole duration up to  $\tau_i$ 's deadline (Approach B). So, at any time instant, the cluster speed is larger or equals to the speed of any individual core. Considering Equations (2) and (5) we can deduce that

$$\forall l : \mathcal{N}_i^l \in \tau_i :: \sum_{k=b_i^l}^{d_i^l} t_k^c s_{i,k} \geq \sum_{k=b_i^l}^{d_i^l} t_k^c S_k^{M_i^l} \geq c_i \mathcal{N}_i^l.$$

So, we conclude that Executing a task with speed according to the speed-profile  $S_i$  guarantees real-time correctness.  $\square$

*An Efficient Approach for Implicit Deadline System.* By adopting simple modification in Equation (4), it is possible to apply the process mentioned above for the implicit deadline tasks also. The workload  $w_i$  should be divided by the period instead of the deadline. We consider the same speed through the task period, so only a single pair  $(s_i, p_i)$  is required, where  $s_i = w_i/T_i$  and  $p_i = 1$ .

**Example 6.** Now we create the speed-profile for the task described in Examples 4 and 5 considering implicit deadline. So it has  $T_i = D_i = 15$  and  $C_i = 6.5$ . Let's assume that it is executed at a speed of 0.6 for 5-time units, at a speed of 0.35 for 10-time units. According to Approach A, the speed-profile is

$$S_i = \begin{pmatrix} 0.6 & 0.35 \\ 0.33 & 0.67 \end{pmatrix},$$

and according to Approach B, the speed-profile is

$$S_i = \begin{pmatrix} 0.43 \\ 1 \end{pmatrix}.$$

## 4.2 Discretization of the Speed-Profile

In Section 4.1, we have described two approaches to create the speed-profile for an individual task. While creating the speed-profiles, those approaches assume a continuous frequency scheme. From a practical point of view, discrete frequency mode should be preferred over the continuous frequency mode, because a real platform supports only a set of frequencies. Now, we describe the technique to discretize all the speeds available in a speed-profile (assuming that the speed-profile is already created).

Suppose, we execute a task  $\tau_i$  (and its speed-profile is  $S_i$ ) in a real-platform, and this platform supports only those speeds available on a speed-set  $\mathcal{Z}$ . Note that the content of  $\mathcal{Z}$  is dependent on the platform. For example, ODROID XU-3 supports a frequency range of 200-1400 MHz (LITTLE cluster) and 200-2000 MHz (big cluster) with scale steps of 100 MHz). Now, for each entry  $s_{i,j} \in S_i$ , we find the minimum speed  $\mathcal{Z}_k \in \mathcal{Z}$ , where  $\mathcal{Z}_k \geq s_{i,j}$ . Once, we find an appropriate  $\mathcal{Z}_k$ ; we set the value of  $s_{i,j}$  as  $s_{i,j} = \mathcal{Z}_k$ .

**Example 7.** Consider a task  $\tau_i$  with the same speed-profile from Example 4. Let us assume that we will execute  $\tau_i$  in a platform where  $\mathcal{Z} = \{0, 0.2, 0.4, 0.55, 0.75, \text{ and } 1\}$ , i.e., this platform supports only six discrete speeds, and all the speeds are normalized w.r.t. the maximum speed supported by this platform. Considering the speed-profile  $S_i$  (from Example 4) and the speed-set  $\mathcal{Z}$ , we find that:

- $s_{i,1} \leq \{\mathcal{Z}_5 \text{ and } \mathcal{Z}_6\}$
- $s_{i,2} \leq \{\mathcal{Z}_4, \mathcal{Z}_5 \text{ and } \mathcal{Z}_6\}$ , and
- $s_{i,3} \leq \{\mathcal{Z}_1, \mathcal{Z}_2, \mathcal{Z}_3, \mathcal{Z}_4, \mathcal{Z}_5 \text{ and } \mathcal{Z}_6\}$ .

Now, we choose the minimum  $\mathcal{Z}_k \in \mathcal{Z}$  such that  $\mathcal{Z}_k \geq s_{i,j}$ . So, we assign  $\mathcal{Z}_5$  to  $s_{i,1}$ ,  $\mathcal{Z}_4$  to  $s_{i,2}$ , and  $\mathcal{Z}_1$  to  $s_{i,3}$ . Now, the updated (i.e., discretized) speed-profile becomes

$$S_i = \begin{pmatrix} 0.75 & 0.55 & 0 \\ 0.33 & 0.47 & 0.2 \end{pmatrix}.$$

**Theorem 1.** When a task executes with its discretized speed-profile, it guarantees that the task will not miss the deadline.



TABLE 1  
Estimated Parameters for Different Cluster of  
an ODROID XU-3 Board

| Cluster Type | $\beta(W)$ | $\alpha(W/MHz^\gamma)$ | $\gamma$ |
|--------------|------------|------------------------|----------|
| big          | 0.155      | $3.03 \times 10^{-9}$  | 2.621    |
| LITTLE       | 0.028      | $2.62 \times 10^{-9}$  | 2.12     |

This table is adopted from [15].

**Proof.** We have shown in Lemma 2 that a task  $\tau_i$  will not miss the deadline if executed according to its speed-profile  $S_i$ . If we discretize  $\tau_i$ 's speed-profile and execute  $\tau_i$  according to this speed-profile, then the task still guarantees the real-time correctness. This is because any speed  $s_{i,j}$  of the discretized speed-profile is greater than or equal to  $s_{i,j}$  when it was continuous.  $\square$

### 4.3 Handling Platform Heterogeneity

In this section, we discuss a specific type of multi-core platform with diverse computing abilities: heterogeneous multi-core platform. We first discuss different types of heterogeneous platforms, and then explain how our proposed techniques can be extended to handle heterogeneity. In a heterogeneous platform, different cores have different computational capabilities. In terms of speed, Funk defined a widely-accepted classification of the heterogeneous platform [26] as follows, where the speed of the processor denotes the work completed (while executing a task) in a single-time unit by this processor.

- (i) *Identical multiprocessors*: On Identical multiprocessors, all tasks are executed at the same speed on any processor;
- (ii) *Uniform multiprocessors*: On Uniform multiprocessors, all the tasks execute at the same speed if allocated on the same processor, but at a different speed on different processors. So, the execution speed of a task depends on the processor where the task is allocated.
- (iii) *Unrelated multiprocessors*: On Unrelated multiprocessors, execution speeds of different tasks may vary on the same processor, i.e., a task's execution speed depends on both the task itself and the processor where it is allocated.

In a heterogeneous platform, each core is designed with a different computational capability, and an efficient task-to-core mapping improves the system resource efficiency. In the context of energy efficiency, two major directions have been mentioned in [27] for any heterogeneous platform:

- (i) Find an appropriate core/cluster for task mapping to reduce the overall power consumption of the whole platform.
- (ii) Deploy energy-aware scheduling techniques on each core/cluster to reduce power consumption.

Our proposed approach covers both directions. First, we use speed profile to identify efficient core/cluster to task mapping and then try to reduce the overall cluster speed as much as possible. It works for an identical heterogeneous platform (a.k.a., homogeneous multiprocessor) as task-to-core mapping does not impact energy consumption much.

TABLE 2  
The "Uncore" Power Consumption for Different Cluster  
of an ODROID XU-3 Board

| Freq(GHz)         | 2    | 1.8   | 1.6  | 1.4   | 1.2   | 1.0   |
|-------------------|------|-------|------|-------|-------|-------|
| big cluster(W)    | 0.8  | 0.528 | 0.39 | 0.309 | 0.244 | 0.182 |
| Freq(GHz)         | 1.4  | 1.2   | 1.0  | 0.8   | 0.6   | 0.4   |
| LITTLE cluster(W) | 0.04 | 0.04  | 0.04 | 0.04  | 0.04  | 0.04  |

This table is adopted from [15].

Now, we extend our approach to apply to the uniform heterogeneous platform by modifying the parameters in the power model in Equation (1), i.e., setting different  $\alpha$ ,  $\beta$ , and  $\gamma$  values for the 'big' and 'LITTLE' cluster. Under such consideration, different clusters no longer share the same power model, and the same task may have different execution requirements on different clusters. We report the estimated values of  $\alpha$ ,  $\beta$ , and  $\gamma$  in Table 1. These parameters are adopted from [15]. The work in [15] estimated these parameters for the ODROID XU-3 board using the real power measurements along with a curve fitting method. They have also assumed that there is another contributor to the total power consumption of a cluster, i.e., the "uncore" power consumption (reported in Table 2). The "uncore" power consumption introduced in the system from some components other than a processor, e.g., a shared cache. Similar to the dynamic power consumption, the "uncore" power consumption also depends on the processor frequency. However, unlike the dynamic power consumption, there is always some "uncore" power consumption as long as the cluster remains on (even if there is no workload on a processor).

Considering all the parameters from Tables 1 and 2, we bring the following modification in Equation (1)

$$P(s) = N_p \beta + \alpha s^\gamma + P_s(f). \quad (6)$$

Here,  $N_p$  denotes the number of cores per cluster, and  $P_s(f)$  denotes the "uncore" power consumption.

We have a different power model for the "big" and the "LITTLE" cluster, but we still don't know what the basis of assigning a task to a cluster is. Recall that, while creating the speed-profile, some vital information (e.g., the speed of a core at a specific time) were known to us (Section 4.1). If the execution speed of a task is greater than a certain threshold at any point from its release to its deadline, then we assign this task to the big cluster. Else, we assign this task to the LITTLE cluster. For the platform we consider (ODROID XU-3), we set the threshold to 0.7. It is the ratio of the maximum speed supported by the big cluster and the LITTLE cluster (see Table 2).

### 4.4 Task Partition: Greedy Merging With Speed-Profiles

We are now equipped with tools (speed-profiles) to measure the potential long-term energy saving of a cluster when partitioning any pair of DAG tasks into it. This subsection describes the scheme for selecting pair by pair so that the total number of clusters can be significantly smaller than the total number of tasks.

Let, we decide for each task whether it should be allocated on a LITTLE or a big cluster using the technique described in Section 4.3. To select a (task) pair that will share the same cluster, we greedily choose the pair that provides maximum power saving, as depicted in Algorithm 1. Note that we allow the pairing of two DAGs that are not merged previously. Also, if any task uses more cores than what is available in a cluster, that task cannot be merged with that cluster.

---

**Algorithm 1.** Greedy Merging
 

---

```

1: Input: Task-set  $\tau$ , with speed-profile  $\mathcal{S}_i$  (computed using
   approach A or approach B) for each task
2: Output: Speed-profile  $\tilde{\mathcal{S}}$  (with processor power saving).
3:  $\bar{\mathcal{S}}, \tilde{\mathcal{S}} \leftarrow \emptyset \triangleright$  All the possible/selected speed-profiles
4: for  $i = 1$  to  $n$  do
5:   for  $j = i + 1$  to  $n$  do
6:      $\mathcal{S}_{ij} \leftarrow \mathcal{S}_i \odot \mathcal{S}_j; \bar{\mathcal{S}} \leftarrow \bar{\mathcal{S}} \cup \mathcal{S}_{ij};$ 
7:   end for
8: end for
9: while  $\exists \mathcal{S}_{xy} \in \bar{\mathcal{S}}$  and  $\mathcal{S}_{xy}$  provides non-zero power saving do
10:   $\mathcal{S}_{xy} \leftarrow$  the pair from  $\bar{\mathcal{S}}$  with maximum power saving
11:   $\tilde{\mathcal{S}} \leftarrow \tilde{\mathcal{S}} \cup \mathcal{S}_{xy}$ 
12:  for  $k = 1$  to  $n$  do
13:     $\bar{\mathcal{S}} \leftarrow \bar{\mathcal{S}} - \mathcal{S}_{kx} - \mathcal{S}_{xk} - \mathcal{S}_{ky} - \mathcal{S}_{yk}$ 
14:  end for
15: end while
16: return  $\tilde{\mathcal{S}}$ .

```

---

Algorithm 1 creates two empty lists  $\bar{\mathcal{S}}$  and  $\tilde{\mathcal{S}}$  that will contain all the possible and selected speed-profiles (Line 3). Lines 4–8, calculate all the possible speed-profiles and insert them into  $\bar{\mathcal{S}}$ . We greedily select a pair of DAGs that provide the maximum power saving (calculated using Equations (6) and (10) from [12]) and update the list  $\bar{\mathcal{S}}$  by removing the pair from any further merging (Lines 9–15). The list  $\tilde{\mathcal{S}}$  is also updated by adding the selected pair (Line 11). We conclude by returning the updated list  $\tilde{\mathcal{S}}$  (Line 16).

**Theorem 2.** *Executing a task with a speed according to the cluster speed-profile guarantees real-time correctness.*

**Proof.** We have shown in Lemma 2 that a task  $\tau_i$  will not miss the deadline if executed according to its speed-profile  $\mathcal{S}_i$ . If  $\tau_i$  share a cluster with another task  $\tau_j$  and executes according to the merged (i.e., cluster) speed-profile  $\mathcal{S}_{ij}$ , then it still guarantees the real-time correctness, because  $\mathcal{S}_{ij} \geq \mathcal{S}_i$  holds at any time instant.  $\square$

**Remark 2.** For  $n$  tasks, the time complexity to generate all possible speed-profiles,  $\bar{\mathcal{S}}$ , is  $O(n^2Z)$ , where  $Z$  is the maximum number of segments of all DAGs in the set after decomposition (related to the structure and size of the DAGs). Algorithm 1 greedily choose a speed-profile by iterating through  $\bar{\mathcal{S}}$  and then update, which takes  $O(n^2)$  time as well. Thus the total complexity of the proposed method is  $O(n^2)$ .

In summary, we have proposed two methods (Section 4.1) to create the speed-profile for a constrained-deadline DAG. We also show that if a task executes according to the speed-profile, it ensures real-time correctness. According to the techniques provided in Section 3, we could evaluate and compare all potential pairs of the combination by calculating

the cluster speed-profile after merging. Finally, Section 4.4 discussed how to use these speed-profiles to find suitable partners to share a cluster greedily.

## 5 SYSTEM EXPERIMENTS

In this section, we present experimental results conducted on an ODROID XU-3 board. The platform runs on Ubuntu 16.04 LTS with Linux kernel 3.10.105. It is fabricated with Samsung Exynos5422 Octa-core SoC, consisting of two quad-core clusters, one ‘LITTLE’ cluster with four energy-efficient ARM Cortex-A7 and one ‘big’ cluster with four performance-efficient ARM Cortex-A15. Four TI INA231 power sensors are integrated onto the board to provide real-time power monitoring for the A-7 and A-15 clusters, GPU, and DRAM. An energy monitoring script, `emoxu3` [28], is used to log energy consumption of the workloads.

*DAG Generation.* In this experiment, we generate two task sets each with 300 DAGs, and use the widely used Erdos-Renyi [29] method to generate a DAG. We tune a parameter  $p$ , that denotes the probability of having an edge between two nodes. In this experiment, we set  $p$  to 0.25 generate DAGs with an uncomplicated structure. If a disconnected DAG is generated, we add the fewest number of edges to make it connected. For experimentation, we have considered arbitrary task periods, and it is determined using Gamma distribution [30]. We set the periods with  $T_i = L_i + 2(C_i/m)(1 + \Gamma(2, 1)/4)$  [2]. Here,  $L_i$  is the critical path length of  $\tau_i$ , calculated according to the definition of  $L_i$  (refer to Section 2).

After generating the topology of each DAG of a set, we partition them into two subsets according to the proposed approach, one to the “big” and the other one to the “LITTLE” cluster, and measure the energy consumption over the hyper-period of all DAGs. We use `rt-app` [31] to emulate the workload for each node. `rt-app` simulates a real-time periodic load and utilizes the POSIX threads model to call and execute threads. For each thread, an execution time needs to be assigned. In this experiment, for each node, we randomly select an execution time ranged between [300 ms, 700 ms]. `rt-app` itself has a latency that varies randomly between 13 – 150 ms per thread. Therefore, we add the maximum latency of `rt-app`, i.e., 150 ms, to the execution time of each thread from an analytical point of view.

*DAG Scheduling.* We use the Linux built-in real-time scheduler `sched_FIFO` to schedule the DAGs. Compared to the other system tasks, DAGs are assigned with higher priorities so that they can execute without interference. Our approach is also applicable to other preemptive schedulers which feature the work-conserving property.

*Frequency Scaling.* According to the frequency/speed-profile (Section 4), we use `cpufreq-set` program (from `cpufrequtils` package) to change the system’s frequency online. We use the ODROID XU-3 board, where scaling-down (up) the frequency of the big cluster takes at most 60 (40) ms, respectively. On the LITTLE cluster, both the operation takes at most 15 ms. Due to this delay, the hyper-period of all DAGs becomes large (230s, in this experiment). We detail the reasons behind this delay in Appendix B.2, available in the online supplemental material.

*The Reference Approach.* Since no work has studied the same problem considered in this paper, we do not have a



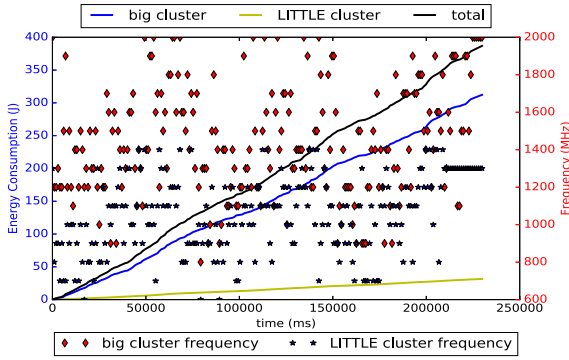


Fig. 4. The energy consumption and the frequency variation of our proposed approach on ODROID XU-3.

direct baseline for comparison. So, we propose a reference approach based on the studies for energy-efficient scheduling of sequential tasks [32]. They assigned an operational frequency to each task, and at run-time, schedule them according to their frequency. In this reference approach, we compute an operational frequency for each DAG. This frequency stretches out execution length of these DAGs as much as possible without violating their deadlines. As stated earlier, the reference approach executes the DAGs with the same partition, but without the merging techniques proposed in Section 4.

*Results.* The experimental results are plotted in Figs. 4 and 5. In these figures, we show (i) the energy consumption over the hyper-period (230s), where the three lines show the energy consumption of the big and LITTLE cluster, and the total system; and (ii) frequency variation during the runtime, where the diamond and star marks denote the operational frequency of the big and the LITTLE cluster at a specific time instant, respectively. Note that the GPU and DRAM also contribute the energy consumption of the total system. Hence, the total energy consumption is a bit higher than the summation of the contribution of the big and the LITTLE cluster, but it is observed that there is a negligible difference for the energy consumption of GPU and DRAM between the two approaches. Besides, it is worth noticing that this energy consumption also accounts for energy consumption of the operating system.

Table 3 summarizes the comparison of the experimental results, where the energy consumption of the two clusters and the total system is presented, and the energy saving from our approach is given. As can be seen, our approach




Fig. 5. The energy consumption and the frequency variation of the reference approach on ODROID XU-3.

TABLE 3  
Summary of Experimental Results

|                | Ours ( $J$ ) | Ref ( $J$ ) | Energy Saving (%) |
|----------------|--------------|-------------|-------------------|
| big cluster    | 312          | 389         | 20                |
| LITTLE cluster | 32           | 38          | 16                |
| Total          | 387          | 472         | 18                |

consumes 312 $J$  and 32 $J$  on the big and the LITTLE cluster, respectively. Comparing to the reference approach, we save energy consumption by 20 and 16 percent. In total, our approach saves energy consumption by 18 percent.

The result can be justified as the reference approach changes the frequency for each DAG, while ours have a fine-grained frequency adjustment at each segment (Section 4.1), and could scale down the frequency if required. Fig. 6 presents the frequency occurrence probability of two clusters which is recorded per second by emoxu3. We observe that within the same time interval the reference approach has a higher probability to execute at a higher frequency, while our approach is more likely to execute at the lower frequencies, thus reducing the energy consumption.

**Remark 3.** Each heavy DAG ( $C_i > T_i$ ) needs two or more cores while executing and the ODROID XU-3 board contains four cores per cluster. So, in this experimental setup, we can not execute more than four heavy DAGs at a time. Such a restriction is not applicable to the light DAGs ( $C_i \leq T_i$ ). We also consider that a heavy DAG cannot be allocated in multiple clusters.

## 6 SIMULATIONS

For large-scale evaluation, we perform simulations and compare the results with existing baselines. We generate DAGs using the Erdos-Renyi method (Section 5). We consider two types of task periods; (a) *harmonic periods*, where the task period  $T_i$  is enforced to be an integral power of 2. We define  $T_i$  as  $T_i = 2^\alpha$ , where  $\alpha$  is the minimum value such that  $2^\alpha \geq L_i$ , where  $L_i$  is the critical path length of  $\tau_i$ ; (b) *arbitrary periods*,  $T_i$  is determined using Gamma distribution (see Section 5).

We compare our approaches with some existing baselines studied in [12], [17], [33]. Total power consumption by our approach and by these baselines are calculated using




Fig. 6. Frequency occurrence probabilities.

Equation (6). As mentioned earlier, [12] considered per-core DVFS, i.e., each core individually is an island of the cluster-based platform. For a fair comparison, according to the scheduling policy of [12], when a task is allocated on some cores at any time instant  $t$ , we choose the maximum speed among all these cores. We consider [12] as a baseline because that work is closely related to ours. Although they have considered per-core DVFS and restrict their attention only to implicit deadline tasks, the task and the power model are same. Besides, although this work and [12] propose different approaches to power saving, the initial (preparation) steps of both approaches are based on commonly known techniques like task decomposition, task merging, etc.

The work in [17] studied a greedy slack stealing (GSS) scheduling approach considering inter-dependent sequential tasks. It considered the DAG model to represent dependencies among the tasks. In GSS, the *slack* (unused time in actual computation requirement of a task) is reclaimed by one task by shifting others towards the deadline. They did not consider repetitive tasks; hence it can be regarded as scheduling a single task. Besides this, the power and graph model used in [17] is different from ours. To ensure a fair comparison, we execute the GSS algorithm using the power model in Equation (1) and assume that once introduced in the system; a processor remains active. We also consider a minimum inter-arrival separation for a DAG. That work considered three different kinds of nodes: *AND*, *OR*, and *Computation* nodes (Section 2.1 in [17]). A computation node has both the maximum and average computation requirement. To comply with our work where the focus in energy reduction while guaranteeing worst-case temporal correctness, we execute the GSS algorithm considering only the computation nodes with their maximum computation requirement. We made all the changes in order to provide a fair comparison. Despite these differences, we chose [17] as a baseline because they studied a GSS approach for energy minimization. They considered the inter-dependent sequential tasks and their dependencies was represented by a DAG, which is similar to our task model.

We also consider [33] as a baseline because this work considered scheduling a set of independent periodic applications, where each application is modeled as a DAG. They proposed an approach for energy minimization combining the DVFS and the DPM policy. Similar to [12] and [17], the work in [33] considered per-core DVFS.

We compare power consumption by varying two parameters for each task: *task periods (utilization)* and *the number of nodes*. We randomly generate 25 sets of DAG tasks and compare the average power consumption.

*Notations of Referenced Approaches.* For the task partitioning step, either we randomly choose any two and allocate them to the same cluster, or greedily choose the ones with lowest speed as proposed. Regarding speed-profile calculation, there are also two options (Approaches A and B in Section 4.1). Combining these options in two steps lead to four baselines: *MaxSpeed\_Greedy*, *SingleSpeed\_Greedy*, *MaxSpeed\_Random*, *SingleSpeed\_Random*. Also, three baselines mentioned above are included for comparison:

- Federated scheduling with intra-task processor merging [12], denoted by *Fed\_Guo*;
- GSS algorithm [17], denoted by *GSS\_Zhu*.

- DVFS and DPM combination [33], denoted by *com\_Chen*.

## 6.1 Uniform Heterogeneous Platform With a Continuous Frequency Scheme

In this section, considering the uniform heterogeneous platform and a continuous frequency scheme, we report the power consumption comparison for different approaches mentioned earlier. Under such a platform, different clusters no longer share the same power model and we use the power model described in Equation (6). We present the power comparison results in an identical heterogeneous platform (from [34]) in Appendix A, available in the online supplemental material.

### 6.1.1 Constrained Deadline Task

Here, we report the power consumption under the scheme for constrained deadline tasks mentioned in Section 4. We evaluate the efficiency of our proposed method by changing two parameters; task period (utilization) and the number of nodes in the task.

*Effect of Varying Task Periods (Utilization).* Here we control the average task utilization through varying the task period. In order to make the task schedulable, the critical path length  $L_i$  of task  $\tau_i$  should not exceed its deadline  $D_i$ . We vary the period in a range ( $L_i \leq T_i \leq C_i$ ). The parameter  $L_i$  and  $C_i$  are measured once the DAG is generated according to the technique described in Section 5. We also use the following equation (according to [12]) to ensure that the value of  $T_i$  satisfies the range ( $L_i \leq T_i \leq C_i$ )

$$T_i = L_i + (1 - k)(C_i - L_i). \quad (7)$$

Here,  $k \in [0, 1]$  is task utilization. As we are considering the constrained deadline tasks,  $D_i$  is randomly picked from the range ( $L_i \leq D_i \leq T_i$ ). The results are presented in Fig. 7a. Note that when any parameter (e.g., number of nodes in a DAG, task utilization) changes, savings in energy randomly vary within a small range and we consider the minimum value among them. The results indicate a proportional relationship between the average power consumption and average task utilization. It happens because a higher task utilization imposes tighter real-time restrictions. It restricts (refer to Fig. 2b) the space for the segment length optimization. In this experiment, the number of nodes is fixed to 30. Fig. 7a shows that *SingleSpeed\_Greedy* approach performs better for a higher utilization value. On average, the *SingleSpeed\_Greedy* approach leads to a power saving of at least 30.23 and 60.2 percent compared to *Fed\_Guo* and *GSS\_Zhu* approaches, respectively. In *SingleSpeed\_Greedy* approach, a task executes with a single speed throughout the deadline. During the task partitioning step, a suitable partner (with similar speed-profile) leads to energy efficiency. However, for the other approaches task speed may vary throughout the deadline. In that case, evil alignment and a significant variation in the speed may reduce energy efficiency (see Fig. 3 and Example 1).

*Effect of Varying the Numbers of Nodes.* Now we vary the number of nodes (10 to 55) ( $T_i$  is fixed) and report the average power consumption. We report the average power consumption for harmonic deadline tasks in Fig. 7b and




Fig. 7. Power consumption comparison between different approaches for the *constrained* deadline tasks considering a *continuous* frequency scheme on the *uniform* heterogeneous platform.




Fig. 8. Power consumption comparison between different approaches for the *implicit* deadline tasks considering a *continuous* frequency scheme on the *uniform* heterogeneous platform.




Fig. 9. Power consumption comparison between different approaches for the *constrained* deadline tasks considering a *discrete* frequency scheme on the *uniform* heterogeneous platform.

arbitrary deadline tasks in Fig. 7c. We observe that the power consumption pattern does not change that much, i.e., SingleSpeed\_Greedy approach outperforms other approaches especially when the number of nodes (in each DAG) are high, 35 or higher. Specifically, under harmonic task periods, the SingleSpeed\_Greedy incurs 40.19 and 65.9 percent less power on average compared to Fed\_Guo and GSS\_Zhu; under arbitrary task periods, the savings potential are 33.43 and 61.96 percent, respectively.

### 6.1.2 Implicit Deadline Task

*Effect of Varying Task Periods (Utilization).* Using previous setup (Section 6.1.1), We observe that the average energy

consumption is directly proportional to the average task utilization.

Fig. 8a shows that *SingleSpeed\_Greedy* approach performs better for a higher utilization value and on average, saves at least 35.21 and 62.52 percent compared to *Fed\_Guo* and *GSS\_Zhu* approaches, respectively.

*Effect of Varying the Numbers of Nodes.* Figs. 8b and 8c report the average power consumption for the harmonic and arbitrary deadline tasks, respectively. We observe that the SingleSpeed\_Greedy approach outperforms other approaches when the number of nodes (in each DAG) are high. Under harmonic task periods, the SingleSpeed\_Greedy incurs 44.84 and 67.55 percent less power on average compared to






Fig. 10. Power consumption comparison between different approaches for the *implicit* deadline tasks considering a *discrete* frequency scheme on the *uniform* heterogeneous platform.

Fed\_Guo and GSS\_Zhu; under arbitrary task periods, the savings potential are 42.33 and 67.19 percent, respectively.

## 6.2 Uniform Heterogeneous Platform With a Discrete Frequency Scheme

In this section, we report the power consumption comparison for the (previously mentioned) approaches considering the uniform heterogeneous platform and a discrete frequency scheme. Under such a platform, we discretize the frequency using the technique described in Section 4.2.

### 6.2.1 Constrained Deadline Task

Here, we consider the constrained deadline tasks and report their average power consumption by changing two parameters: task period (or utilization) and the number of nodes.

*Effect of Varying Task Periods (Utilization).* Similar to the Figs. 7a, and 8a, we observe that the (i) average energy consumption is directly proportional to the average task utilization. (ii) *SingleSpeed\_Greedy* approach consumes less power than other approaches (see Fig. 9a).

*Effect of Varying the Numbers of Nodes.* We vary the number of nodes (10 to 55) and report the average power consumption for harmonic (arbitrary) deadline tasks in Fig. 9b (Fig. 9c). Similar to the Figs. 7 and 8, we observe that the (i) Performance of *SingleSpeed\_Random*, *SingleSpeed\_Greedy*, *MaxSpeed\_Greedy*, and *MaxSpeed\_Random* does not vary that much for a small number of nodes (typically 10 to 25) per DAG. (ii) *SingleSpeed\_Greedy* approach performs better (i.e., consume less power) than other approaches when the number of nodes per DAG is high.

### 6.2.2 Implicit Deadline Task

Now, we report the average power consumption using the same setup as described in Section 6.2.1, i.e., (a) for a fixed number of nodes (30) per task, change their utilization value, and (b) vary the number of nodes (10 to 55) per task, while keeping  $T_i$  fixed. We report the average power consumption in Fig. 10. From this figure, we observe that the (i) Performance of *SingleSpeed\_Random*, *SingleSpeed\_Greedy*, *MaxSpeed\_Greedy*, and *MaxSpeed\_Random* does not vary that much for a smaller task utilization or when the number of nodes per DAG is small (typically 10 to 35). (ii) *SingleSpeed\_Greedy* approach performs better (i.e., consume less

power) compared to the other approaches when the number of nodes per DAG is high.

## 7 RELATED WORK

Much work has been done aimed at energy-efficient scheduling of sequential tasks in a homogeneous multi-core platform (see [11] for a survey). Considering the mixed-criticality task model and varying-speed processors, the works on [23], [35], [36], [37], [38] proposed an approach to handle the energy minimization problem. The work in [27], [39], [40], [41], [42], [43] presented an energy-efficient approach for the heterogeneous platform. Considering the real-time tasks in clustered heterogeneous platforms, the work in [39] studied the partitioned EDF scheduling policy, while [42] proposed an optimal task-core mapping technique that is fully-migrative. Considering the heterogeneous multi-core platform, a two-phase algorithm was proposed by [27]. In the first phase, they proposed a tasks-core allocation approach with the aim of reducing the dynamic energy consumption, while the second phase seeks for a better sleep state to reduce the leakage power consumption. A low overhead, DVFS-cum-DPM enabled energy-aware approach, HEALERS, was proposed by [43]. However, none of them considered the intra-task parallelism. Considering a clustered heterogeneous MPSoC platform, a migrative cluster scheduling approach was proposed by [15]. In this approach, run-time migration (within different cores in the same cluster) for a task is allowed to improve resource utilization. The work in [44] studied the technique to utilize the parallelism in a hard real-time streaming application (represented as a Synchronous Data Flow (SDF) graph) in a clustered heterogeneous platform.

Till date, considering both the intra-task parallelization and power minimization has received less attention. A greedy slack stealing algorithm is proposed in [17] that deals with task represented by graphs but did not consider the periodic DAGs. Assuming per-core DVFS, [33] provided the technique to combine DVFS and DPM. Considering the real-time jobs (represented as a DAG) in cloud computing systems and in a heterogeneous multi-core platform, the work in [45], [46] studied a QoS-aware and energy-efficient scheduling strategy. They proposed a scheduling policy that utilizes per-core DVFS. With the aim of improving energy-efficiency in a heterogeneous real-time platform, [47] proposed a combined

approach considering the approximate computation and bin packing strategy. [48] investigated the energy awareness for cores that are grouped into blocks, and each block shares the same power supply scaled by DVFS. Benefits of (in terms of power saving) intra-task parallelism is proven theoretically in [1]. Considering the fork-join model, [49] reported an empirical evaluation of the power savings in a real test-bed. Based on level-packing, [50] proposed an energy efficient algorithm for implicit deadline tasks with same arrival time and deadline.

None of these works allows intra-task processor sharing considering the sporadic DAG task model. The recent work in [12], [18] is most related to ours. However, these works are significantly different from ours w.r.t the task model, platform, real-time constraints (deadlines), solution techniques, and the evaluation. *First*, the work in [12] considered a simplified model where only one DAG task executes at a time, while the work in [18] extends this work by allowing inter-task processor sharing. However, both of these works assumed that the number of cores are unlimited. *Second*, Both the works in [12], [18] assumed per-core speed scaling. However, many of the existing platforms (e.g., ODRROID XU-3) do not support such speed scaling—speeds of processors under the same cluster must execute at the same speed. As the number of cores fabricated on a chip increases, per-core speed scaling design is less likely to be supported due to the inefficiency on hardware levels [13]. *Third*, Both of these works have studied only the implicit deadline tasks and did not consider the *constrained* deadline tasks. Hence, the non-negligible idle gaps between the task deadline and its next release remain un-utilized. *Finally*, the evaluations in [12], [18] were done based on simulations without any implementation on a real platform.

## 8 CONCLUSION

In this paper, we have studied real-time scheduling of a set of implicit and constrained deadline sporadic DAG tasks. We schedule these tasks on the cluster-based multi-core platforms with the goal of minimizing the CPU power consumption. In a clustered multi-core platform, the cores within the same cluster run at the same speed at any given time. Such design better balances energy efficiency and hardware cost and appears in many systems. However, from the resource management point of view, this additional restriction leads to new challenges. By leveraging a new concept, i.e., *speed-profile*, which models energy consumption variations during run-time, we can conduct scheduling and task-to-cluster partitioning while minimizing the expected overall long-term CPU energy consumption. To our knowledge, this is the first work that has investigated energy-efficient scheduling of DAGs on clustered multi-core platform. Also, no work considered energy-aware real-time scheduling of constrained deadline DAG tasks.

We have implemented our result on an ODRROID XU-3 board to demonstrate its feasibility and practicality. We have also complemented our system experiments on a larger scale through realistic simulations that demonstrate an energy saving of up to 57 percent through our proposed approach compared to existing methods. In this work, we

have restricted our attention mainly to the CPU power consumption. In the future, we plan to consider other components that may affect the total power consumption, e.g., cache misses, context switches, I/O usage, etc. We also plan to study the effect of tasks sporadic release patterns (to the overall power consumption) and propose a task reallocation scheme.

## ACKNOWLEDGMENTS

This work was partially supported by the National Science Foundation CNS-1850851, CNS-1742985 and National Natural Science Foundation of China 61801418.

## REFERENCES

- [1] A. Paolillo, J. Goossens, P. M. Hettiarachchi, and N. Fisher, "Power minimization for parallel real-time systems with malleable jobs and homogeneous frequencies," in *Proc. IEEE 20th Int. Conf. Embedded Real-Time Comput. Syst. Appl.*, 2014, pp. 1–10.
- [2] A. Saifullah, D. Ferry, J. Li, K. Agrawal, C. Lu, and C. D. Gill, "Parallel real-time scheduling of DAGs," *IEEE Trans. Parallel Distrib. Syst.*, vol. 25, no. 12, pp. 3242–3252, Dec. 2014.
- [3] M. Qamhieh, F. Fauberteau, L. George, and S. Midonnet, "Global EDF scheduling of directed acyclic graphs on multiprocessor systems," in *Proc. 21st Int. Conf. Real-Time Netw. Syst.*, 2013, pp. 287–296.
- [4] S. Baruah, V. Bonifaci, A. Marchetti-Spaccamela, L. Stougie, and A. Wiese, "A generalized parallel task model for recurrent real-time processes," in *Proc. IEEE 33rd Real-Time Syst. Symp.*, 2012, pp. 63–72.
- [5] J. Li, K. Agrawal, C. Lu, and C. Gill, "Analysis of global EDF for parallel tasks," in *Proc. 25th Euromicro Conf. Real-Time Syst.*, 2013, pp. 3–13.
- [6] V. Bonifaci, A. Marchetti-Spaccamela, S. Stiller, and A. Wiese, "Feasibility analysis in the sporadic DAG task model," in *Proc. 25th Euromicro Conf. Real-Time Syst.*, 2013, pp. 225–233.
- [7] J. Li, J. J. Chen, K. Agrawal, C. Lu, C. Gill, and A. Saifullah, "Analysis of federated and global scheduling for parallel real-time tasks," in *Proc. 26th Euromicro Conf. Real-Time Syst.*, 2014, pp. 85–96.
- [8] S. Baruah, V. Bonifaci, and A. Marchetti-Spaccamela, "The global EDF scheduling of systems of conditional sporadic DAG tasks," in *Proc. 27th Euromicro Conf. Real-Time Syst.*, 2015, pp. 222–231.
- [9] T. Hagras and J. Janecek, "A high performance, low complexity algorithm for compile-time job scheduling in homogeneous computing environments," in *Proc. Int. Conf. Parallel Process. Workshops*, 2003, pp. 149–155.
- [10] H. Aydin and Q. Yang, "Energy-aware partitioning for multiprocessor real-time systems," in *Proc. Int. Parallel Distrib. Process. Symp.*, 2003, p. 9.
- [11] M. Bambagini, M. Marinoni, H. Aydin, and G. Buttazzo, "Energy-aware scheduling for real-time systems: A survey," *ACM Trans. Embedded Comput. Syst.*, vol. 15, no. 1, 2016, Art. no. 7.
- [12] Z. Guo, A. Bhuiyan, A. Saifullah, N. Guan, and H. Xiong, "Energy-efficient multi-core scheduling for real-time DAG tasks," in *Proc. 29th Euromicro Conf. Real-Time Syst.*, 2017, pp. 22:1–22:21.
- [13] S. Herbert and D. Marculescu, "Analysis of dynamic voltage/frequency scaling in chip-multiprocessors," in *Proc. Int. Symp. Low Power Electron. Des.*, 2007, pp. 38–43.
- [14] 2017. [Online]. Available: <http://www.hardkernel.com/>
- [15] D. Liu, J. Spasic, G. Chen, and T. Stefanov, "Energy-efficient mapping of real-time streaming applications on cluster heterogeneous MPSoCs," in *Proc. 13th IEEE Symp. Embedded Syst. Real-Time Multimedia*, 2015, pp. 1–10.
- [16] J. Kim, H. Kim, K. Lakshmanan, and R. R. Rajkumar, "Parallel scheduling for cyber-physical systems: Analysis and case study on a self-driving car," in *Proc. ACM/IEEE Int. Conf. Cyber-Physical Syst.*, 2013, pp. 31–40.
- [17] D. Zhu, D. Mosse, and R. Melhem, "Power-aware scheduling for AND/OR graphs in real-time systems," *IEEE Trans. Parallel Distrib. Syst.*, vol. 15, no. 9, pp. 849–864, Sep. 2004.
- [18] A. Bhuiyan, Z. Guo, A. Saifullah, N. Guan, and H. Xiong, "Energy-efficient real-time scheduling of DAG tasks," *ACM Trans. Embedded Comput. Syst.*, vol. 17, no. 5, 2018, Art. no. 84.

- [19] W. M. Kolb, "Curve fitting for programmable calculators," IMTEC, 1984.
- [20] S. Pagani and J.-J. Chen, "Energy efficient task partitioning based on the single frequency approximation scheme," in *Proc. IEEE 34th Real-Time Syst. Symp.*, 2013, pp. 308–318.
- [21] S. Pagani and J.-J. Chen, "Energy efficiency analysis for the single frequency approximation (SFA) scheme," *ACM Trans. Embedded Comput. Syst.*, vol. 13, no. 5s, 2014, Art. no. 158.
- [22] P. Huang, P. Kumar, G. Giannopoulou, and L. Thiele, "Energy efficient DVFS scheduling for mixed-criticality systems," in *Proc. Int. Conf. Embedded Softw.*, 2014, pp. 1–10.
- [23] S. Narayana, P. Huang, G. Giannopoulou, L. Thiele, and R. V. Prasad, "Exploring energy saving for mixed-criticality systems on multicores," in *Proc. IEEE Real-Time Embedded Technol. Appl. Symp.*, 2016, pp. 1–12.
- [24] J. Howard *et al.*, "A 48-core IA-32 processor in 45 nm CMOS using on-die message-passing and DVFS for performance and power scaling," *IEEE J. Solid-State Circuits*, vol. 46, no. 1, pp. 173–183, Jan. 2011.
- [25] D. Maxim and L. Cucu-Grosjean, "Response time analysis for fixed-priority tasks with multiple probabilistic parameters," in *Proc. IEEE 34th Real-Time Syst. Symp.*, 2013, pp. 224–235.
- [26] S. H. Funk, *EDF Scheduling on Heterogeneous Multiprocessors*. Chapel Hill, NC, USA: Univ. of North Carolina, 2004.
- [27] M. A. Awan, D. Masson, and E. Tovar, "Energy efficient mapping of mixed criticality applications on unrelated heterogeneous multicore platforms," in *Proc. 11th IEEE Symp. Ind. Embedded Syst.*, 2016, pp. 1–10.
- [28] 2017. [Online]. Available: <https://github.com/tuxamito/emoxu3>
- [29] D. Cordeiro, G. Mounié, S. Perarnau, D. Trystram, J.-M. Vincent, and F. Wagner, "Random graph generation for scheduling simulations," in *Proc. 3rd Int. ICST Conf. Simul. Tools Techn.*, 2010, Art. no. 60.
- [30] 2017. [Online]. Available: [http://en.wikipedia.org/wiki/Gamma\\_distribution](http://en.wikipedia.org/wiki/Gamma_distribution)
- [31] 2017. [Online]. Available: <https://github.com/scheduler-tools/rt-app/>
- [32] J.-J. Chen and C.-F. Kuo, "Energy-efficient scheduling for real-time systems on dynamic voltage scaling (DVS) platforms," in *Proc. 13th IEEE Int. Conf. Embedded Real-Time Comput. Syst. Appl.*, 2007, pp. 28–38.
- [33] G. Chen, K. Huang, and A. Knoll, "Energy optimization for real-time multiprocessor system-on-chip with optimal DVFS and DPM combination," *ACM Trans. Embedded Comput. Syst.*, vol. 13, no. 3s, 2014, Art. no. 111.
- [34] 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," in *Proc. IEEE Real-Time Embedded Technol. Appl. Symp.*, 2019, pp. 156–168.
- [35] S. Baruah and Z. Guo, "Mixed-criticality scheduling upon varying-speed processors," in *Proc. IEEE 34th Real-Time Syst. Symp.*, 2013, pp. 68–77.
- [36] S. Baruah and Z. Guo, "Scheduling mixed-criticality implicit-deadline sporadic task systems upon a varying-speed processor," in *Proc. IEEE Real-Time Syst. Symp.*, 2014, pp. 31–40.
- [37] Z. Guo and S. Baruah, "The concurrent consideration of uncertainty in WCETs and processor speeds in mixed-criticality systems," in *Proc. 23rd Int. Conf. Real-Time Netw. Syst.*, 2015, pp. 247–256.
- [38] A. Bhuiyan, S. Sruti, Z. Guo, and K. Yang, "Precise scheduling of mixed-criticality tasks by varying processor speed," in *Proc. 27th Int. Conf. Real-Time Netw. Syst.*, 2019, pp. 123–132.
- [39] A. Colin, A. Kandhalu, and R. Rajkumar, "Energy-efficient allocation of real-time applications onto heterogeneous processors," in *Proc. IEEE 20th Int. Conf. Embedded Real-Time Comput. Syst. Appl.*, 2014, pp. 1–10.
- [40] J.-J. Chen, A. Schranzhofer, and L. Thiele, "Energy minimization for periodic real-time tasks on heterogeneous processing units," in *Proc. IEEE Int. Symp. Parallel Distrib. Process.*, 2009, pp. 1–12.
- [41] C. Liu, J. Li, W. Huang, J. Rubio, E. Speight, and X. Lin, "Power-efficient time-sensitive mapping in heterogeneous systems," in *Proc. 21st Int. Conf. Parallel Archit. Compilation Techn.*, 2012, pp. 23–32.
- [42] H. S. Chwa, J. Seo, J. Lee, and I. Shin, "Optimal real-time scheduling on two-type heterogeneous multicore platforms," in *IEEE Real-Time Syst. Symp.*, 2015, pp. 119–129.
- [43] S. Moulik, R. Devaraj, and A. Sarkar, "HEALERS: A heterogeneous energy-aware low-overhead real-time scheduler," *IET Comput. Digit. Techn.*, vol. 13, no. 6, pp. 470–480, Nov. 2019.
- [44] J. Spasic, D. Liu, and T. Stefanov, "Energy-efficient mapping of real-time applications on heterogeneous MPSoCs using task replication," in *Proc. Int. Conf. Hardware/Softw. Codes. Syst. Synthesis*, 2016, pp. 1–10.
- [45] G. L. Stavrinides and H. D. Karatza, "Energy-aware scheduling of real-time workflow applications in clouds utilizing DVFS and approximate computations," in *Proc. IEEE 6th Int. Conf. Future Internet Things Cloud*, 2018, pp. 33–40.
- [46] G. L. Stavrinides and H. D. Karatza, "An energy-efficient, QoS-aware and cost-effective scheduling approach for real-time workflow applications in cloud computing systems utilizing DVFS and approximate computations," *Future Gener. Comput. Syst.*, vol. 96, pp. 216–226, 2019.
- [47] G. L. Stavrinides and H. D. Karatza, "Scheduling real-time DAGs in heterogeneous clusters by combining imprecise computations and bin packing techniques for the exploitation of schedule holes," *Future Gener. Comput. Syst.*, vol. 28, no. 7, pp. 977–988, 2012.
- [48] X. Qi and D.-K. Zhu, "Energy efficient block-partitioned multicore processors for parallel applications," *J. Comput. Sci. Technol.*, vol. 26, no. 3, 2011, Art. no. 418.
- [49] A. Paolillo, P. Rodriguez, N. Veshchikov, J. Goossens, and B. Rodriguez, "Quantifying energy consumption for practical fork-join parallelism on an embedded real-time operating system," in *Proc. 24th Int. Conf. Real-Time Netw. Syst.*, 2016, pp. 329–338.
- [50] H. Xu, F. Kong, and Q. Deng, "Energy minimizing for parallel real-time tasks based on level-packing," in *Proc. IEEE Int. Conf. Embedded Real-Time Comput. Syst. Appl.*, 2012, pp. 98–103.
- [51] A. Dunkels, F. Osterlind, N. Tsiftes, and Z. He, "Software-based on-line energy estimation for sensor nodes," in *Proc. 4th Workshop Embedded Netw. Sensors*, 2007, pp. 28–32.
- [52] H.-Y. Zhou, D.-Y. Luo, Y. Gao, and D.-C. Zuo, "Modeling of node energy consumption for wireless sensor networks," *Wireless Sensor Netw.*, vol. 3, no. 01, 2011, Art. no. 18.



**Ashikahmed Bhuiyan** received the bachelor of science degree in computer science and engineering from the Bangladesh University of Engineering and Technology (BUET), Dhaka, Bangladesh, in 2013. He is currently working toward the PhD degree in the Department of Electrical and Computer Engineering, University of Central Florida (UCF), Orlando, Florida, under the supervision of Zhishan Guo and Abusayeed Saifullah (from the Wayne State University). He is a member of the Real-Time & Intelligent Systems Lab, UCF. His research focuses on improving energy efficiency in real-time embedded systems, parallel computing, and mixed-criticality scheduling. He has received the Best Student Paper Award at the 40th IEEE Real-Time Systems Symposium (RTSS 2019).



**Di Liu** received the BEng and MEng degrees from Northwestern Polytechnical University, Xi'an, China, in 2007 and 2011, respectively, and the PhD degree from Leiden University, Leiden, The Netherlands, in 2017. He is currently an assistant professor with the School of Software, Yunnan University, China. His research interests include the fields of real-time systems, energy-efficient multicore/many core systems, and cyber-physical systems.



**Aamir Khan** received the MS degree from the Computer Engineering Department, Missouri University of Science and Technology, Rolla, Missouri, in 2018. He is currently an embedded systems engineer with BrainCo Tech in Boston. His areas of interest includes real-time systems and body-machine interface products.





**Abusayeed Saifullah** received the PhD degree in computer science and engineering with Turner Dissertation Award from Washington University in St. Louis, St. Louis, Missouri, in 2014. He is an assistant professor of the Computer Science Department, Wayne State University. His research primarily concerns Internet-of-Things, cyber-physical systems, real-time systems, embedded systems, and low-power wide-area networks. He received seven Best Paper Awards/Nominations in highly competitive conferences including ACM SenSys (2016 nomination), IEEE RTSS (2019, 2014, 2011), IEEE ICII (2018), and IEEE RTAS (2012 nomination). He also received multiple young investigator awards including the CAREER award (2019) and the CRII award (2016) of the National Science Foundation (NSF). He is serving as the program chair of IEEE ICSS 2020, served as a track chair of IEEE ICCCN 2019 and as a program committee member for various conferences including ACM SenSys, IEEE RTSS, ACM/IEEE IoTDI, IEEE RTAS, ACM/IEEE ICCPS, ACM MobiHoc, IEEE INFOCOM, EWSN, and ACM IWQoS. He also served as a guest editor of the *IEEE Transactions on Industrial Informatics*, and is currently an editor of the *Elsevier Pervasive and Mobile Computing Journal*.



**Nan Guan** received the PhD degree from Uppsala University, Uppsala, Sweden. He is currently an assistant professor with the Department of Computing, Hong Kong Polytechnic University. He worked with Northeastern University, China before joining The Hong Kong Polytechnic University. His research interests include the design and analysis of real-time systems, embedded systems, cyber-physical systems, and Internet-of-Things (IoT) systems. He received the EDAA Outstanding Dissertation Award in 2014, the CCF Outstanding Dissertation Award in 2013, the Best Paper Award of IEEE RTSS in 2009, the Best Paper Award of DATE in 2013, the Best Paper Award of ACM e-Energy 2018 and the Best Paper Award of IEEE ISORC 2019. He served as the TPC co-chair of EMSOFT 2015, ICSS 2017, SETTA 2019, the TPC track chair of RTAS 2018.



**Zhishan Guo** received the BE (with honor) degree in computer science and technology from Tsinghua University, Beijing, China, in 2009, the MPhil degree in mechanical and automation engineering from the Chinese University of Hong Kong, Hong Kong, in 2011, and the PhD degree in computer science from the University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, in 2016. He is an assistant professor with the Department of Electrical and Computer Engineering, University of Central Florida. His research and teaching interests include real-time scheduling, cyber-physical systems, and neural networks and their applications.

▷ **For more information on this or any other computing topic, please visit our Digital Library at [www.computer.org/csdl](http://www.computer.org/csdl).**