

Mixed-Criticality Real-Time Scheduling of Gang Task Systems

Ashikahmed Bhuiyan · Kecheng Yang ·
Samsil Arefin · Abusayeed Saifullah ·
Nan Guan · Zhishan Guo*

Received: date / Accepted: date

Abstract Mixed-criticality (MC) scheduling of sequential tasks (with no intra-task parallelism) has been well-explored by the real-time systems community. However, till date, there has been little progress on MC scheduling of parallel tasks. MC scheduling of parallel tasks is highly challenging due to the requirement of various assurances under different criticality levels. In this work, we address the MC scheduling of parallel tasks of gang model that allows workloads to execute on multiple cores simultaneously, as well as the change to degree of parallelism of a task upon a mode switch. It represents an efficient mode-based parallel processing scheme with many potential applications. To schedule such task sets, we propose a new technique GEDF-VD, which integrates Global Earliest Deadline First (GEDF) and Earliest Deadline First with Virtual Deadline (EDF-VD). We prove the correctness of GEDF-VD and

Ashikahmed Bhuiyan
University of Central Florida
E-mail: ashik@knights.ucf.edu

Kecheng Yang
Texas State University
E-mail: yangk@txstate.edu

Samsil Arefin
Microsoft New England Research and Development Center
E-mail: Samsil.Arefin@microsoft.com

Abusayeed Saifullah
Wayne State University
E-mail: saifullah@wayne.edu

Nan Guan
The Hong Kong Polytechnic University
E-mail: nan.guan@polyu.edu.hk

Zhishan Guo
University of Central Florida
E-mail: zsguo@ucf.edu

* = Corresponding author

provide a detailed quantitative evaluation in terms of speedup bound in both the MC and the non-MC cases. Specifically, we show that GEDF provides a speedup bound of 2 for non-MC gang tasks, while the speedup for GEDF-VD considering MC gang tasks is $\sqrt{5}+1$. Experiments on randomly generated gang task sets are conducted to validate our theoretical findings and to demonstrate the effectiveness of the proposed approach.

1 Introduction

Due to size, weight, and power considerations, there is a trend that multiple tasks with different criticality levels (that are subject to varying degrees of assurance/verification) share a computing platform [49]. This type of system is commonly known as a *mixed-criticality (MC) system*, where each task can be associated with various execution budgets. During normal operation, all tasks are scheduled according to their typical execution budget. However, some critical tasks may exceed their typical budget and need more resources to finish their execution. Suppose the available resources are not sufficient in these scenarios. In that case, the less critical task will be sacrificed to free up the resources for accommodating the additional computational requirements requested by the more critical ones.

Take an avionics software standard as an example, where the ground control subsystems are more safety-critical than ground communication and light controls. During the incident of emergency (e.g., an accident), it is more important to execute the safety-critical components rather than the other components. On the other hand, in normal condition, all these components are required to perform smoothly (for more details, refer to the Table 1.1 of [35], which demonstrates the RTCA DO-178B avionics software standard). MC scheduling has received considerable attention (refer to [17] for a thorough and updated survey) as it brings significant improvements in resource efficiency.

Note that safety-critical MC systems have tight correctness requirements. These requirements can be verified by two related but orthogonal perspectives: *a priori verification* and *run-time robustness* [5]. Before run-time, a priori verification determines whether a system will behave correctly (or not) during execution, while run-time robustness deals with unexpected system behavior at run-time. There are some debates regarding MC's applicability to run-time robustness [25, 26]. Although an MC system has imitation from the perspective of a priori verification (which is our work scope), these criticisms usually do not hinder the applicability of an MC system for its designed scope of a priori verification [5].

Parallel Computing Workloads. Recent advances in parallel computing allow executing a single piece of code simultaneously on multiple computing units or a set of threads to execute on multiple processors concurrently. Such design provides a much better capability of exploiting the benefits provided by modern platforms. As a result, there is an urgent need in handling workload models that allow intra-task parallelism (i.e., parallel tasks). Parallel com-

puting systems perform a large number of computations and often need to interact with their surroundings under real-time constraints, e.g., arms system (RADAR). In these types of applications, a lot of processors co-operate with each other, and these communications are timing critical. It is necessary for a system to have both high performance and predictability; i.e., efficient control that minimizes the introduced overhead, while responding to external events (coming through sensors) in real-time. In this work, we consider the gang scheduling, where all threads of a task are grouped into a gang, and during execution, the whole group is concurrently scheduled on distinct cores. The gang task model is a practical, widely used, and representative workload model for intra-task parallelism [1, 21, 30, 33]. Also, the gang task model is supported by some widely used parallel computing programming standards (e.g., OpenACC [42]), which is commonly used in graphics processing unit (GPU).

Existing Work. The real-time systems and parallel computing communities have given considerable attention towards these two directions: *MC scheduling* and *scheduling of parallel tasks*. These two emerging trends bring in some critical and exciting problems, and there is an emerging need in integrating those two trends. There has been extensive research on the (a) MC scheduling of sequential (i.e., non-parallel) tasks (refer to the recent survey in [18, 17]) and (b) scheduling of parallel tasks with a single-criticality level [33, 30, 21, 2, 16, 38]. Till date, very few efforts [40, 4, 39] have been made towards the combined problem of MC scheduling of parallel tasks. To our knowledge, none of these efforts has considered mixed-criticality gang task scheduling on multi-core platforms.

Motivation Behind This Work. Multi-core platform enables applications that require better energy efficiency, higher performance, and real-time guarantees. The notion of MC systems with the intra-task parallelism stems from many current trends. For example, the number of cores fabricated on a chip is increasing rapidly. Besides, the computational demand for an individual task (with stringent timing requirements) is rising, which makes it essential to consider the intra-task parallelism. Furthermore, when safety-critical and non-safety-critical tasks share a common computational platform, there is an increasing demand to integrate functionality with different criticality levels. Such demand promotes the idea of MC scheduling, i.e., combining various functionalities of varying criticality levels onto the same computing platform.

In this work, we study the mixed-criticality gang task scheduling, and the fault-tolerant systems would be a promising application for such study. A fault-tolerant system often follows the mixed-criticality model [17]. If a fault is identified in such a system, it is recovered via various recovery techniques, e.g., exception handling, recovery blocks, and task replication. Some extra work has to be undertaken upon identifying a fault, which leads to the abandonment or delay of some less critical works. The impact of a fault in such a system ranges from no visible effect to an entire system crash. To overcome these faults, the *ASTEROID* project is proposed (a cross-layer fault-tolerance solution for the mixed-criticality platform) that detects errors and recovers the system in different software and hardware layers [20, 24]. Considering such a cross-

layer platform, a recent work in [44] proposed a replica-aware co-scheduling (with Strict Priority Preemptive protocols) for a mixed-criticality system that improves the system performance. They have considered the replicas as a gang that is activated concurrently on multiple cores.

Challenges. In gang task model, a task cannot start execution until the number of available cores is no less than what is required by it (i.e., a task's *degree of parallelism*). This simple constraint adds a huge restriction on real-time schedulability and makes the problem highly challenging. We are aware of only one known correct schedulability analysis [21] under Global Earliest Deadline First (GEDF) for gang tasks. Besides, integrating MC in gang scheduling scheme adds additional challenges due to the dual notion of correctness. In the normal mode, a task may have a utilization less than 1, while in the critical mode, the utilization could be much higher than one [21]. Such a change in the utilization adds significant complexity in speedup bound analysis [6]. For example, in the speedup bound analysis for MC scheduling of ordinary sporadic tasks, an individual task's utilization is at most the processor's speed is a straight forward and necessary feasibility condition. At the same time, it no longer holds for the gang tasks. Besides, the scheduler does not know the exact behavior of each task before run-time (non-clairvoyant). Hence, the scheduler must be able to detect the critical condition early enough to allocate more resources to the more critical tasks to handle this drastic change and still be able to meet the deadlines.

This research. In this paper, we study the real-time scheduling of MC gang tasks on identical multi-core platforms. We propose the first scheduling algorithm GEDF-VD (GEDF with Virtual Deadline) for MC gang tasks. Our approach leverages the synthesis of uniprocessor scheduling techniques such as EDF-VD [7] as well as *GEDF* [21] that was designed for non-MC gang tasks. To our knowledge, this is the first work that studies the MC scheduling of the gang task model. Specifically, we make the following contributions:

- We generalize the gang task model to the MC context by incorporating extensions on both the execution time and the degree of parallelism dimensions, and propose GEDF-VD for the generalized model. We also conduct a utilization based schedulability test and prove its correctness formally.
- We prove that the speedup bound [6] for GEDF to gang tasks in a non-MC platform is at most $(2 - 1/(M + 1 - \min_i\{m_i\}))$, where M denotes the total number of processor cores and m_i denotes the degree of parallelism of task τ_i . To our knowledge, this is the first speedup bound result for GEDF scheduling of gang tasks.
- With the result from the previous step, we then derive a speedup bound of $\sqrt{5} + 1$ for GEDF-VD considering MC gang tasks.
- Extensive simulations under randomly generated task sets are conducted to demonstrate the real-time performance and effectiveness of the proposed algorithm in terms of acceptance ratio, which is defined as the ratio of the number of schedulable task sets over the total number of task sets.

Organization. The remainder of this paper is organized as follows. Section 2 describes the task model, notations, and preliminaries. Section 3 provides a detailed description of our scheduling algorithm and prove its correctness. Section 4 derives the speedup bounds for the non-MC and MC platform, under GEDF and GEDF-VD scheduling algorithms, respectively. Simulation results are presented in Section 5. Section 6 discusses related prior work. Section 7 concludes this paper and points out future research directions.

2 Dual-Criticality Gang Task Model

In this work, we consider the problem of scheduling a task set $\tau = \{\tau_1, \dots, \tau_n\}$ of n independent *implicit deadline* (i.e., the period of a task is equal to its deadline) sporadic MC gang tasks on M identical cores. In this model, each task generates an infinite number of MC gang jobs (the j^{th} job of task τ_i is denoted as $\tau_{i,j}$). To describe the dual-criticality gang task model, first, we provide details on traditional *non-MC gang task model* and *MC sporadic sequential task model*. Then, by leveraging these two models, we generalize the gang task model to the MC context. In this work, we restrict our attention to dual-criticality model, and consider it is an important step towards the multi-criticality systems.

Non-MC gang task model. In traditional non-MC gang task model, each task τ_i is represented with a 4-tuple (m_i, c_i, T_i, D_i) , where m_i is the degree of parallelism and each job of task τ_i requires access to m_i cores for at most c_i time units to complete its execution, T_i is the task period, and D_i is the relative deadline. In gang scheduling, each task consists of multiple threads (referred to as a gang), and each thread of the same task occupies a processor for the same time quantum [33,21]. Hence, in the time-space, the execution requirement of any job $\tau_{i,j} \in \tau_i$ can be represented as an $m_i \times c_i$ rectangle. The relative deadline D_i specifies that for each of the released jobs $\tau_{i,j}$ (of task τ_i), its deadline $d_{i,j} = r_{i,j} + D_i$, where $r_{i,j}$ denotes the release time of $\tau_{i,j}$ [21]. The *utilization* u_i of each task $\tau_i \in \tau$ is given by $u_i = (m_i c_i) / T_i$, and the overall system utilization is: $U_{sum} = \sum_{\tau_i \in \tau} u_i$. Note that, it is possible that the value of u_i is larger than one, which is different from the traditional sequential task model. Based on the scheduling flexibility, a gang task τ_i can be categorized into three groups. A task τ_i is said to be:

- *rigid*, if m_i is fixed a priori and does not change throughout the execution,
- *moldable*, if m_i is fixed during its activation and does not change throughout the execution,
- *malleable*, if m_i is not fixed and can be changed during its execution by the scheduler.

We currently restrict our attention to the *rigid* task model. This model suits various applications that use parallelism, some of which are implemented using the message-passing approach and tools like MPI.

MC sporadic task model. In a dual-criticality systems, the criticality level of τ_i is represented by $\chi_i = \{\text{LO}, \text{HI}\}$. The worst case execution time (WCET)

estimations of each task is also represented by a tuple $(c_i^{\text{LO}}, c_i^{\text{HI}})$ where c_i^{LO} and c_i^{HI} represent the LO and HI-criticality WCETs respectively. c_i^{HI} is measured by a more pessimistic tool by considering all possible scenarios, while c_i^{LO} is calculated using a less pessimistic yet realistic tool. Collection of all LO- and HI-criticality tasks in τ are denoted by τ_{LO} and τ_{HI} respectively. u_i^{LO} and u_i^{HI} denotes the utilization of τ_i in LO- and HI-criticality mode respectively, where $u_i^{\text{LO}} = c_i^{\text{LO}}/T_i$ and $u_i^{\text{HI}} = c_i^{\text{HI}}/T_i$.

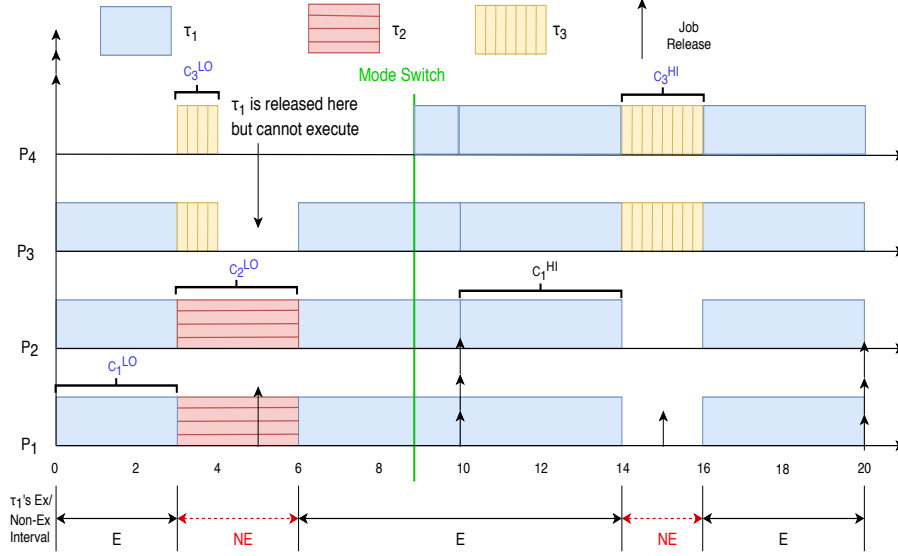


Fig. 1: A GEDF scheduling of the MC gang task-set from Table 1 and the executing (E)/non-executing (NE) intervals of τ_1 .

MC gang task model. By leveraging the above two models, in our work, we consider a workload model of MC gang tasks, where each task τ_i is represented by a 7-tuple $(m_i^{\text{LO}}, m_i^{\text{HI}}, \chi_i, c_i^{\text{LO}}, c_i^{\text{HI}}, T_i, D_i)$, where

- $m_i^{\text{LO}}(m_i^{\text{HI}})$ = degree of parallelism of task τ_i in LO(HI) criticality mode, and $m_i^{\text{LO}} \leq m_i^{\text{HI}}$.
- χ_i = criticality level of each task τ_i and $\chi_i \in \{\text{LO}, \text{HI}\}$.
- $c_i^{\text{LO}}(c_i^{\text{HI}})$ = τ_i 's WCET in LO(HI)-criticality mode, and $c_i^{\text{LO}} \leq c_i^{\text{HI}}$.
- T_i = minimum inter-arrival time between jobs.
- D_i = relative deadline.

In this model, a task τ_i occupies $m_i^{\text{LO}}(m_i^{\text{HI}})$ processors for $c_i^{\text{LO}}(c_i^{\text{HI}})$ time quantum at LO(HI)-criticality mode. Note that, if $\forall \tau_i, m_i^{\text{LO}} = m_i^{\text{HI}} = 1$, i.e., degree of parallelism for each gang task is 1, our analysis (Section 3 and 4) will reduce to the existing MC scheduling method designed for the sporadic

task model. We believe this is common for a restricted special case of a more complex and expressive model. For example, the directed acyclic graph (DAG) task model [16, 38] is popular to represent intra-task parallelism. Many of the existing schedulability analysis considering the DAG model would also reduce to prior study for ordinary sporadic tasks if the number of nodes of each DAG task is equal to 1.

Now, we generalize the utilization concepts to suit the MC gang task model, which are analogous to the above-mentioned concepts. Refer to the Example 1 for details.

$$U_{\text{LO}}^{\text{LO}} \stackrel{\text{def}}{=} \sum_{\tau_i \in \tau_{\text{LO}}} m_i^{\text{LO}} \times c_i^{\text{LO}} / T_i,$$

$$U_{\text{HI}}^{\text{LO}} \stackrel{\text{def}}{=} \sum_{\tau_i \in \tau_{\text{HI}}} m_i^{\text{LO}} \times c_i^{\text{LO}} / T_i,$$

$$U_{\text{HI}}^{\text{HI}} \stackrel{\text{def}}{=} \sum_{\tau_i \in \tau_{\text{HI}}} m_i^{\text{HI}} \times c_i^{\text{HI}} / T_i$$

Table 1: An MC gang task set with GEDF schedule shown in Figure 1.

Task ID	c_i^{LO}	c_i^{HI}	T_i	χ_i	m_i^{LO}	m_i^{HI}
τ_1	3	4	5	HI	3	4
τ_2	3	3	10	LO	2	2
τ_3	1	2	10	HI	2	2

Example 1 Consider the task-set $\tau = (\tau_1, \tau_2, \tau_3)$ in Table 1. For this task-set we derive the utilization as follows:

$$U_{\text{LO}}^{\text{LO}} = c_2^{\text{LO}} \times m_2^{\text{LO}} / T_2 = 0.6, U_{\text{HI}}^{\text{LO}} = c_1^{\text{LO}} \times m_1^{\text{LO}} / T_1 + c_3^{\text{LO}} \times m_3^{\text{LO}} / T_3 = 2, \text{ and } U_{\text{HI}}^{\text{HI}} = c_1^{\text{HI}} \times m_1^{\text{HI}} / T_1 + c_3^{\text{HI}} \times m_3^{\text{HI}} / T_3 = 3.6.$$

Example 2 Consider the MC gang task set in Table 1 to be scheduled in four cores. A GEDF schedule for this task set is shown in Figure 1. The system starts at LO-criticality mode, and all the tasks (τ_1, τ_2, τ_3) will execute up-to C_i^{LO} . Recall that, m_i^{LO} is the degree of parallelism of τ_i in LO-criticality mode. Hence, τ_1 cannot execute at $t = 5$ as it needs three cores to execute while only two cores (P_3 and P_4) are idle. At a mode switch ($t = 9$), all LO-criticality tasks (τ_2) are dropped, and all HI-criticality tasks (τ_1, τ_3) will execute up-to C_i^{HI} (at m_i^{HI} cores). After a mode switch, all HI-criticality jobs (including the ones which are currently executing) will execute up-to their HI-criticality WCET.

Motivations behind this model. Some commonly used parallel computing programming standards (e.g., OpenACC [42]) support the gang task model. OpenACC is one of the parallel computing programming standards used for GPU architecture. Recently, there has been extensive research on GPU architecture (few to mention [19, 23, 50, 51]). GPU architecture is popular because

of the features like (1) highly threaded but low context switch latency architecture, (2) high parallelism and (3) minimal dependency between data elements, etc. Previous work on GPU scheduling considered limited or no preemption policy [23,50]. However, this work is motivated by some recent attempts to incorporate the *preemptive* support in GPUs. For example, a prototype has been implemented and tested with preemptive support (at the pixel level and the thread level) in a virtualized environment in a recent work [19]. Its prototype is EDF based, and enhanced with a bandwidth isolation mechanism (e.g., constant/total bandwidth servers [45]) for the graphics and computing workloads. Also, the prototype is tested on a recent NVIDIA Tegra-based system on a chip (SoCs) [41]. Since some recent work study the preemptive support in the GPU architecture, there is a need for a comprehensive study of gang task scheduling using GEDF.

Now, we introduce some definitions and preliminaries which will be frequently used in later sections of this paper.

Definition 1 (MC-correct schedule): Scheduling strategy must ensure an *MC-correct* schedule, as defined below [39].

- If the system stays in normal condition (i.e., each task in the system finishes execution within its LO-criticality WCET), all tasks must meet their deadlines.
- If the system transits into a critical condition (i.e., there exists a HI-criticality task executing beyond its LO-criticality WCET), all HI-criticality tasks must meet their deadlines, while LO-criticality tasks need not so.

Definition 2 (Executing/Non-Executing interval) An interval $[t_1, t_2)$ (where $t_1 < t_2$) is an **executing** interval for a task τ_i if m_i out of M cores are executing the current active job released by τ_i throughout this interval. Otherwise, $[t_1, t_2)$ is a **non-executing** interval for τ_i . An illustrative example is shown in Figure 1 by pointing the executing and non-executing intervals for task τ_1 .

Definition 3 (Active/pending task) If there exists a task $\tau_i \in \tau$, such that it has a job $\tau_{i,j}$ where $r_{i,j} \leq t < d_{i,j}$. Here, $r_{i,j}$ and $d_{i,j}$ respectively denotes the release time and deadline of $\tau_{i,j}$, then τ_i is considered as an active task at time t . A job is *pending* if it is released but not finished [21].

Definition 4 Maximum possible number of idle cores Δ_i^{LO} (Δ_i^{HI}) for a task τ_i refers to the maximum number of available cores (that are not executing any job) at any time in LO (HI-criticality mode) during τ_i 's non-executing intervals in which it has a pending job [21].

Example 3 Let us consider a 4-core platform and task set $\tau = \{\tau_1, \tau_2, \tau_3\}$ from Table 1. At LO-criticality mode, the degree of parallelism for these tasks are given as: $m_1^{LO} = 3$, $m_2^{LO} = m_3^{LO} = 2$. For this task set, the maximum possible number of idle cores at LO-criticality mode is: $\Delta_1^{LO} = 2$, $\Delta_2^{LO} = \Delta_3^{LO} = 1$. This is because τ_1 cannot execute at time t (even when it has a pending job) if τ_2 or τ_3 is executing at t . The degree of parallelism for τ_2 (or τ_3) is 2. So, the maximum number of idle cores for τ_1 is Δ_1 , where $\Delta_1 = 4 - 2 = 2$. We can

calculate the value of Δ_2 and Δ_3 in the same approach (refer to Algorithm 1 in [21]). At HI-criticality mode, degree of parallelism for these tasks becomes: $m_1^{\text{HI}} = 4$, $m_2^{\text{HI}} = m_3^{\text{HI}} = 2$, and the maximum possible number of idle cores will be: $\Delta_1^{\text{HI}} = 2$, $\Delta_2^{\text{HI}} = \Delta_3^{\text{HI}} = 0$.

System behavior and scope of this work. It is expected that an MC system starts execution in normal mode. The system-wide mode transition is triggered if a HI-criticality task τ_i has received cumulative execution length beyond its LO-criticality WCET and did not signal its finishing. Like the Vestal model [49], after a mode switch, no LO-criticality tasks get any service guarantee. After mode transition (from LO-criticality to HI-criticality), at the first idle instant, the system switches back to the LO-criticality mode again. All other scenarios (e.g., a HI-criticality task runs for more than its HI-criticality WCET) are considered as *erroneous*, where no guarantees will be made and hence is not considered.

3 GEDF-VD for Dual-Criticality System

Now we describe our algorithm for the MC task systems considering the GEDF-VD algorithm. In this work, we consider implicit deadline (So, we use the terms *deadline* and *period* interchangeably) sporadic task systems on identical multi-core platforms. We integrate an uniprocessor MC scheduling technique (*EDF-VD* [7]) with a multiprocessor gang task scheduling technique (*GEDF* [21]) and derive a new algorithm named GEDF-VD (Subsections 3.1 and 3.2). In our approach, we determine a *scaling factor*, which scales the deadline of all HI-criticality tasks at LO-criticality mode. This factor will be calculated in such a way that the correctness of the system can be guaranteed at both LO- and HI-criticality modes (Subsections 3.3 and 3.4).

3.1 EDF-VD and GEDF-VD: An Overview

EDF-VD. In case of a mode switch (LO to HI), to generate an *MC-correct schedule* (Definition 1), a scheduler must ensure that all HI-criticality tasks meet their deadlines (while LO-criticality tasks can be sacrificed). To guarantee this criterion, a specific amount of CPU time must be reserved for those HI-criticality tasks even if the system is running at LO-criticality mode. This reservation of time can be achieved by shortening the deadlines of HI-criticality tasks under normal mode—those are virtual deadlines.

In EDF-VD, deadlines of all HI-criticality tasks are shortened by multiplying them with a *scaling factor*, and this updated deadline is called the *virtual deadline*. During run-time (at LO-criticality mode), all HI-criticality (LO-criticality) tasks are executed according to their virtual (actual) deadline, according to the EDF order. Upon a mode switch, only the HI-criticality tasks are executed in the EDF order w.r.t their actual/original deadlines.

In the case of a LO- to HI-criticality mode-switch, a HI-criticality task demands additional computational requirements. Setting a virtual deadline for the HI-criticality tasks leaves enough time to finish the extra workload within their actual deadlines. If the virtual deadline is too short, it increases the system density at normal (i.e., LO-criticality) mode. In contrast, a large virtual deadline threatens the schedulability of the system after a LO- to HI-criticality mode switch. The trick is to determine a balanced scaling factor x , such that the correctness under both execution modes can be guaranteed. [7] showed the steps to calculate the minimum x that guarantees the schedulability of all tasks in the system. They also proved the improvement in system schedulability by reducing the deadline for HI-criticality tasks at LO-criticality mode.

Remark 1 In this work, we consider a completely different Gang task workload model in a multi-core platform. As a result, the approach proposed by Baruah et al. [7] to calculate the scaling factor x , as well as the schedulability test, are no longer applicable for our case. We propose a novel approach to calculate a feasible scaling factor x in this section.

GEDF-VD. Now, we provide an overview of our algorithm (GEDF-VD) considering an implicit-deadline sporadic MC gang task system τ to be scheduled on M identical cores. The GEDF-VD algorithm starts by checking whether GEDF can successfully schedule the *regular task system*. A regular task system denotes that, all LO-criticality tasks will execute up-to their LO-criticality WCET and all HI-criticality tasks will execute up-to their HI-criticality WCET. It returns *SUCCESS* immediately if the regular task system is schedulable. Otherwise, all HI-criticality tasks can execute up-to their LO-criticality WCETs and their deadline is shortened (i.e., virtual deadline) and set to $\hat{T}_i = xT_i$, while all LO-criticality tasks execute up-to their LO-criticality WCETs with their original deadline. If any of the currently executing job (of a HI-criticality task) executed beyond its LO-criticality WCET and did not signal its completion by \hat{T}_i , the scheduler immediately discards all currently active LO-criticality jobs. Also, the deadline for all HI-criticality jobs is changed to their release time plus their actual deadline. Subsection 3.2 provides a detailed description of GEDF-VD algorithm.

3.2 GEDF-VD: A Detailed Description

In this subsection, we describe the GEDF-VD scheduling approach in a two-phase process. First, we describe what happens prior to run-time (denoted as a *pre-processing phase*). In this phase, GEDF-VD determines whether (or not) it is required to set a virtual deadline for the HI-criticality tasks. A lower and an upper bound of the virtual deadline is also calculated in this phase. Then, we discuss how the jobs are scheduled at run-time (denoted as *handling the dispatched jobs at run-time*). We present the pseudo-code for (the run-time part of) GEDF-VD in Algorithm 1.

Pre-processing phase. In this phase, we perform a schedulability test for

Algorithm 1: GEDF-VD (online part)

Input: A dual-criticality task-set $\tau = \{\tau_1, \tau_2, \dots, \tau_n\}$ and a feasible x .

```

1 /* Handling tasks at run-time */
2 Whenever a job is released by tasks  $\tau_i$  at time instant  $t$ 
3 if  $\tau_i \in \tau_{\text{HI}}$  then
4   |  $d_{i,j} = t + xT_i$ ;
5 end
6 if  $\tau_i \in \tau_{\text{LO}}$  then
7   |  $d_{i,j} = t + T_i$ ;
8 end
9 Schedule all active jobs by GEDF according to  $d_{i,j}$ 's.
10 if  $\exists \tau_{i,j} \in \tau_{\text{HI}}$  that is not finished by  $d_{i,j}$  at time  $t$  then
11   | /* Mode Switch */
12   | for  $\forall \tau_i \in \tau_{\text{HI}}$  do
13     |  $d_{i,j} = d_{i,j} + (1-x)T_i$ 
14   | end
15   | Discard all  $\tau_i \in \tau_{\text{LO}}$ 
16   | Schedule  $\tau' = \{\tau_{\text{HI}}\}$  by GEDF.
17 end

```

ordinary (*non-MC*) GEDF to determine whether (or not) it can successfully schedule: (i) all $\tau_i \in \tau_{\text{LO}}$ up-to their LO-criticality WCET (c_i^{LO}), and (ii) all $\tau_i \in \tau_{\text{HI}}$ up-to their HI-criticality WCET (c_i^{HI}). If the GEDF test fails, then, for each HI-criticality task $\tau_i \in \tau_{\text{HI}}$, a virtual deadline \hat{T}_i is computed (Step-2), and they execute up-to their LO-criticality WCET (c_i^{LO}).

Step 1. We start by checking whether the task-set can be successfully scheduled by GEDF. If so, then GEDF directly schedules the system. Else, we modify the task deadlines (Step 2).

Step 2: An additional *virtual deadline* parameter \hat{T}_i is calculated for each HI-criticality task τ_i , where $\hat{T}_i = xT_i$. A schedulability test for GEDF-VD is provided next. Furthermore, when the schedulability test is passed, x can be arbitrarily chosen from the range $[A, B]$ while GEDF-VD is guaranteed to generate an MC-correct schedule, where A and B are defined and can be easily calculated for any given system by the following equations:

$$A = \max\{A_1, A_2\}; \quad (1)$$

$$A_1 = \max_{i:\tau_i \in \tau_{\text{LO}}} \left\{ \frac{U_{\text{HI}}^{\text{LO}}}{M - \Delta_i^{\text{LO}} - U_{\text{LO}}^{\text{LO}}} \right\}; \quad (2)$$

$$A_2 = \max_{i:\tau_i \in \tau} \left\{ \frac{m_i^{\text{LO}} U_{\text{HI}}^{\text{LO}} + u_i^{\text{LO}} (M - \Delta_i^{\text{LO}} - m_i^{\text{LO}})}{m_i^{\text{LO}} (M - \Delta_i^{\text{LO}} - U_{\text{LO}}^{\text{LO}})} \right\}; \quad (3)$$

$$B = \min_{i:\tau_i \in \tau_{\text{HI}}} \left\{ 1 - \frac{m_i^{\text{HI}} U_{\text{HI}}^{\text{HI}} + u_i^{\text{HI}} (M - \Delta_i^{\text{HI}} - m_i^{\text{HI}})}{m_i^{\text{HI}} \times (M - \Delta_i^{\text{HI}})} \right\}. \quad (4)$$

A schedulability test for GEDF-VD. The following theorem provides a sufficient schedulability test for GEDF-VD.

Theorem 1 *An MC gang task system is schedulable under GEDF-VD upon M identical unit-speed processors if both conditions hold:*

$$U_{\text{LO}}^{\text{LO}} < M - \max_i \{\Delta_i^{\text{LO}}\}, \quad (5)$$

$$A \leq B. \quad (6)$$

We will prove this theorem later by proving Lemmas 1 and 2 in Subsections 3.3 and 3.4.

Recall that, $\Delta_i^{\text{LO}} < M$ and $\Delta_i^{\text{HI}} < M$ for all i . Therefore, $m_i U_{\text{HI}}^{\text{LO}} + u_i^{\text{LO}}(M - \Delta_i^{\text{LO}} - m_i^{\text{LO}}) = m_i^{\text{LO}}(U_{\text{HI}}^{\text{LO}} - u_i^{\text{LO}}) + u_i^{\text{LO}}(M - \Delta_i^{\text{LO}}) > 0$, which with (5) together implies $A > 0$; and also $m_i^{\text{HI}} U_{\text{HI}}^{\text{HI}} + u_i^{\text{HI}}(M - \Delta_i^{\text{HI}} - m_i^{\text{HI}}) = m_i^{\text{HI}}(U_{\text{HI}}^{\text{HI}} - u_i^{\text{HI}}) + u_i^{\text{HI}}(M - \Delta_i^{\text{HI}}) > 0$, which implies $B < 1$. Thus, both (5) and (6) being true implies that $0 < A \leq B < 1$, which guarantees that any x chosen from $[A, B]$ must be a valid scaling factor such that $0 < x < 1$.

Run-time dispatch. Similar to GEDF, at any specific time instant, a task with the earliest deadline gets the highest priority. In case of ties, task with a smaller index is favored. Let a binary variable ξ indicate the system-criticality level, then consider the following two possible cases:

Case 1. System is in the LO-criticality mode ($\xi = 0$), j^{th} job of task τ_i arrives at time t :

- (i) If τ_i is a LO-criticality task, set the deadline as $d_{i,j} = t + T_i$, else set $d_{i,j} = t + \hat{T}_i$, where $\hat{T}_i = xT_i$.
- (ii) If any of the currently executing jobs executes for more than c_i^{LO} and does not signal completion, then the system switches to the HI-criticality mode (Case 2).

Case 2. While the system is in the HI-criticality mode ($\xi = 1$):

- (i) Discard all LO-criticality tasks (or use background scheduling).
- (ii) Update the deadline for the currently active HI-criticality jobs into release time (of these jobs) plus their actual relative deadline.
- (iii) For any future HI-criticality task τ_i that releases a job at time t , the deadline is set to $t + T_i$.
- (iv) When there is an idle instant, switch to the LO-criticality mode (Case 1)¹.

3.3 Proof of Correctness in the LO-Criticality Mode

In this subsection, we show that GEDF-VD and its schedulability test given by Theorem 1 are able to guarantee MC correctness at LO-criticality mode.

Lemma 1 *If both (5) and (6) are true, GEDF-VD guarantees that all LO-criticality tasks meet their deadlines and all HI-criticality tasks meet their virtual deadlines during LO-criticality mode.*

¹ Note that HI-criticality mode exists for certification purposes. Such both directions of mode switch should be unlikely events during run time. Please also refer to the discussions about apriori verification and run-time robustness in Section 1.

Proof Dong et al. [21] have proved that, given any real-time implicit deadline sporadic gang task system τ , GEDF can schedule it successfully if

$$\begin{aligned} U_{sum} &\leq (M - \Delta_i) \times \left(1 - \frac{u_i}{m_i}\right) + u_i \\ \iff U_{sum} &\leq M - \Delta_i + u_i \left(1 - \frac{M - \Delta_i}{m_i}\right) \end{aligned} \quad (7)$$

holds for all $\tau_i \in \tau$ (refer to Theorem 2). The virtual deadline increases the utilization of these HI-criticality tasks (and hence the whole system). Note that, in the LO-criticality mode, each HI-criticality task is scheduled by its virtual relative deadline xT_i while each LO-criticality task is scheduled by its actual deadline T_i . Therefore, it is sufficient to view each LO-criticality task as a sporadic task with utilization u_i^{LO} and view each HI-criticality task as a sporadic task with utilization u_i^{LO}/x , in order to meet every LO-criticality deadline and every HI-criticality virtual deadline in LO-criticality mode. Then, for every i such that $\tau_i \in \tau$, we discuss the two cases for $M - \Delta_i^{\text{LO}} - m_i^{\text{LO}}$. Therefore, it suffice to evaluate (7) under such utilizations for every task τ_i . We show this by two cases: 1) $\tau_i \in \tau_{\text{HI}}$, and 2) $\tau_i \in \tau_{\text{LO}}$.

Case 1: $\tau_i \in \tau_{\text{HI}}$. In this case, using (7) as a result from [21], we just need the following inequality to hold for any $\tau_i \in \tau_{\text{HI}}$.

$$\begin{aligned} U_{\text{LO}}^{\text{LO}} + \frac{U_{\text{HI}}^{\text{LO}}}{x} &\leq M - \Delta_i^{\text{LO}} + \frac{u_i^{\text{LO}}}{x} \left(1 - \frac{M - \Delta_i^{\text{LO}}}{m_i^{\text{LO}}}\right) \\ \iff \frac{U_{\text{HI}}^{\text{LO}}}{x} + \frac{u_i^{\text{LO}}}{x} \left(\frac{M - \Delta_i^{\text{LO}}}{m_i^{\text{LO}}} - 1\right) &\leq M - \Delta_i^{\text{LO}} - U_{\text{LO}}^{\text{LO}} \\ \iff \frac{m_i^{\text{LO}} U_{\text{HI}}^{\text{LO}} + u_i^{\text{LO}} (M - \Delta_i^{\text{LO}} - m_i^{\text{LO}})}{m_i^{\text{LO}} \cdot x} &\leq M - \Delta_i^{\text{LO}} - U_{\text{LO}}^{\text{LO}} \end{aligned} \quad (8)$$

Notice that (5) implies

$$M - \Delta_i^{\text{LO}} - U_{\text{LO}}^{\text{LO}} > 0 \text{ for all } i \text{ such that } \tau_i \in \tau, \quad (9)$$

and (6) allows $x \in [A, B]$ can be chosen so that $x \geq A$, which, by (1) and (3), implies

$$x \geq \frac{m_i^{\text{LO}} U_{\text{HI}}^{\text{LO}} + u_i^{\text{LO}} (M - \Delta_i^{\text{LO}} - m_i^{\text{LO}})}{m_i^{\text{LO}} (M - \Delta_i^{\text{LO}} - U_{\text{LO}}^{\text{LO}})} \text{ for all } i \text{ such that } \tau_i \in \tau. \quad (10)$$

It is clear that (9) and (10) imply (8).

Case 2: $\tau_i \in \tau_{\text{LO}}$. In this case, using (7) as a result from [21], we just need the following condition to hold for any $\tau_i \in \tau_{\text{LO}}$.

$$U_{\text{LO}}^{\text{LO}} + \frac{U_{\text{HI}}^{\text{LO}}}{x} \leq M - \Delta_i^{\text{LO}} + u_i^{\text{LO}} \left(1 - \frac{M - \Delta_i^{\text{LO}}}{m_i^{\text{LO}}}\right) \quad (11)$$

Subcase 2.1: $M - \Delta_i^{\text{LO}} - m_i^{\text{LO}} \leq 0$. In this case, $M - \Delta_i^{\text{LO}} \leq m_i^{\text{LO}} \implies 1 - \frac{M - \Delta_i^{\text{LO}}}{m_i^{\text{LO}}} \geq 0$. Therefore, the following inequality implies (11):

$$U_{\text{LO}}^{\text{LO}} + \frac{U_{\text{HI}}^{\text{LO}}}{x} \leq M - \Delta_i^{\text{LO}}. \quad (12)$$

Notice that (5) implies

$$M - \Delta_i^{\text{LO}} - U_{\text{LO}}^{\text{LO}} > 0 \text{ for all } i \text{ such that } \tau_i \in \tau, \quad (13)$$

and (6) allows $x \in [A, B]$ can be chosen so that $x \geq A$, which, by (1) and (2), implies

$$x \geq \frac{U_{\text{HI}}^{\text{LO}}}{M - \Delta_i^{\text{LO}} - U_{\text{LO}}^{\text{LO}}} \text{ for all } i \text{ such that } \tau_i \in \tau_{\text{LO}}. \quad (14)$$

It is clear that (13) and (14) imply (12)

Subcase 2.2: $M - \Delta_i^{\text{LO}} - m_i^{\text{LO}} > 0$. In this case, $M - \Delta_i^{\text{LO}} > m_i^{\text{LO}} \implies 1 - \frac{M - \Delta_i^{\text{LO}}}{m_i^{\text{LO}}} < 0$. So, $\frac{u_i^{\text{LO}}}{x} (1 - \frac{M - \Delta_i^{\text{LO}}}{m_i^{\text{LO}}}) < u_i^{\text{LO}} (1 - \frac{M - \Delta_i^{\text{LO}}}{m_i^{\text{LO}}})$, as $0 < x < 1$. Therefore, the following inequality implies (11).

$$U_{\text{LO}}^{\text{LO}} + \frac{U_{\text{HI}}^{\text{LO}}}{x} \leq M - \Delta_i^{\text{LO}} + \frac{u_i^{\text{LO}}}{x} (1 - \frac{M - \Delta_i^{\text{LO}}}{m_i^{\text{LO}}}) \quad (15)$$

By the same reasoning as that for Case 1, (15) always holds because (9) and (10) are “for any $\tau_i \in \tau$ ” and both HI- and LO-criticality tasks are included in the set τ . That is, (11) is also true in Case 2.2 here.

Combining Cases 1 and 2 (the latter includes Sub-cases 2.1 and 2.2), the lemma follows.

3.4 Proof of Correctness in the HI-Criticality Mode

In this subsection, we show that GEDF-VD and its schedulability test given by Theorem 1 are able to guarantee MC correctness at HI-criticality mode.

Lemma 2 *If both (5) and (6) are true, GEDF-VD guarantees that all HI-criticality tasks meet their deadlines during HI-criticality mode.*

Proof At the mode switch point from the lo- to HI-criticality mode, a job from any task $\tau_i \in \tau_{\text{HI}}$ must be either completed or has a deadline at least $(1-x)T_i$ after this mode-switch point; otherwise, an earlier time instant would have been the mode switch point.

Afterwards, any job from any task $\tau_i \in \tau_{\text{HI}}$ has at least T_i time units (which is more than $(1-x)T_i$ as $0 < x < 1$) from their releases in the HI-criticality mode to their corresponding deadlines.

Therefore, viewing each task $\tau_i \in \tau_{\text{HI}}$ in the HI-criticality mode as a sporadic task with utilization $\frac{u_i^{\text{HI}}}{(1-x)}$ and using (7) as a result from [21], the following inequality is sufficient to ensure that all HI-criticality tasks meet their actual deadlines during HI-criticality mode. For all i such that $\tau_i \in \tau_{\text{HI}}$,

$$m_i^{\text{HI}} \times \frac{U_{\text{HI}}^{\text{HI}}}{(1-x)} \leq m_i^{\text{HI}} \times (M - \Delta_i^{\text{HI}}) - \frac{u_i^{\text{HI}}}{1-x} \times (M - \Delta_i^{\text{HI}} - m_i^{\text{HI}}) \quad (16)$$




Fig. 2: Any value of the scaling factor x , where $A \leq x \leq B$, guarantees an MC-correct schedule.

Notice that (6) allows $x \in [A, B]$ can be chosen so that $x \leq B$, which, by (4), implies the following equation holds for all i such that $\tau_i \in \tau_{\text{HI}}$:

$$x \leq 1 - \frac{m_i^{\text{HI}} U_{\text{HI}}^{\text{HI}} + u_i^{\text{HI}} (M - \Delta_i^{\text{HI}} - m_i^{\text{HI}})}{m_i^{\text{HI}} \times (M - \Delta_i^{\text{HI}})} \quad (17)$$

Furthermore, Equation (17) is equivalent to Equation (16), as $0 < x < 1$ and $\Delta_i^{\text{HI}} < M$. Thus, the lemma follows.

Finishing up. We establish Theorem 1 by combining Lemma 1 and 2, and it serves as a sufficient schedulability test for GEDF-VD to schedule MC gang task sets on M identical processors. In addition, Figure 2 gives a high-level intuition for validating Theorem 1, given that Lemma 1 and 2 have been proven. Note that we did leverage some insights (in our analysis) from prior works on MC scheduling and gang scheduling. However, considering both of these directions brought increased complexity in our system model. The existing analysis of MC scheduling or gang scheduling is not directly applicable to our work. For example, in the speedup bound analysis for MC scheduling of ordinary sporadic tasks, *an individual task's utilization is at most the speed of a processor* is a straightforward and necessary feasibility condition, while it no longer holds for the gang tasks.

Remark 2 If we consider that the degree of parallelism of a task τ_i remains unchanged even after a mode-switch, i.e., $m_i^{\text{LO}} = m_i^{\text{HI}} = m_i$, then the correctness proofs become identical to the correctness proofs established in [14].

4 Speedup Bound Analysis

In this section, we evaluate the effectiveness of our algorithm GEDF-VD based on *speedup bound* metric, which is a widely accepted tool for evaluating the effectiveness of multiprocessor scheduling algorithms [6]. We will first provide the related definition and some existing results, and then (in Subsection 4.1) will derive the speedup bound for gang tasks under GEDF algorithm considering the *non-MC* systems. This is the first speedup bound for (non-MC) gang tasks under GEDF scheduling policy, which lays the foundation for deriving a speedup bound for MC gang tasks. Finally, in Subsection 4.2, considering

the MC sporadic gang tasks, we prove a speedup bound for our proposed algorithm GEDF-VD. We derive the speedup bound for non-MC gang tasks (Subsection 4.1) and the MC gang tasks (Subsection 4.2) assuming that the degree of parallelism of a task τ_i does not change after a mode-switch, i.e., $m_i^{\text{LO}} = m_i^{\text{HI}} = m_i$.

Definition 5 (Speedup factor and speedup bound) For a scheduler \mathcal{S} , a speedup factor \mathcal{V} ($\mathcal{V} \geq 1$) (also known as resource augmentation factor) means that any task set that is schedulable by an optimal scheduler on a platform of speed-1 cores will be schedulable by \mathcal{S} on a platform of speed- \mathcal{V} cores.

For a scheduler \mathcal{S} , a speedup bound refers to the lower bound of the speedup factor \mathcal{V} achievable by it. A speedup bound for a scheduler \mathcal{S} provides an estimation of how far the performance of \mathcal{S} is from an optimal scheduler, and the lower the better.

Limitations. Our speedup factors result in this section rely on the following assumption that

$$m_i \leq \frac{M+1}{2} \text{ for all } \tau_i \in \tau, \quad (18)$$

The speedup factors results in this section apply only to systems that satisfy the condition (18). Nonetheless, condition (18) was not required for the schedulability test and analysis in the last section. Therefore those schedulability results apply to a broader range of MC gang task systems. In practice, condition (18) is often satisfied because of the number of cores in modern platforms increases.

Note that gang tasks cannot be scheduled on uniprocessor platforms due to their natures of the *mandatory* parallel processor access request. Therefore, in order to compare with a potential optimal scheduler on a uniprocessor, we propose a *De-gang* transformation between a multiprocessor gang task set and a corresponding Liu-and-Layland (LL) task set:

- **De-gang**: Given a gang task set $\tau = \{\tau_1, \dots, \tau_n\}$, for each task $\tau_i = \{m_i, c_i, T_i\}$, construct m_i LL tasks $\{\tau_i^{(1)}, \dots, \tau_i^{(m_i)}\}$, each with the same execution length and period, i.e., $\tau_i^{(j)} = \{c_i, T_i\}$ for any $j = 1, \dots, m_i$. For mapping of the other way around, any deganged LL task set can be clustered into n groups, where there are m_i tasks from the i^{th} group sharing the same execution time c_i and the same period T_i , resulting in a gang task $\tau_i = \{m_i, c_i, T_i\}$ of the same “total” utilization. The extension to MC task set is trivial—treat c_i as a vector and maintain the values during the transfer.

A moment thought should convince the reader that it suffices to restrict our attention to the de-ganged LL task set when deriving the speedup bound, as the de-ganged LL task set being schedulable is *necessary* for the corresponding gang task set to be schedulable. This transformation does not change the overall set utilization and thus does not change the utilization-based necessary

schedulability conditions (i.e., basis of the speedup proofs). Throughout the proofs in this section, the following Greek letters will be used frequently:

$$\begin{aligned}\psi &= M/(2 - \frac{1}{M}); \\ \phi &= \frac{\sqrt{5} + 1}{2} \quad (\text{i.e., golden ratio}); \\ \Phi &= \frac{\sqrt{5} - 1}{2}\end{aligned}\tag{19}$$

4.1 Speedup Bound for Gang Tasks under GEDF

In this subsection, we derive the speedup bound (shown in Theorem 2) for the algorithm GEDF, considering the *non-MC gang task set* τ , executing on \mathcal{V} -speed cores. This is the first speedup bound result for gang task under GEDF scheduling. This analysis lays the basis for deriving the speedup bound for the proposed MC gang task scheduler.

Theorem 2 *Given any de-ganged task set that is schedulable on a speed- M uni-processor, the corresponding gang task set will pass the schedulability test of GEDF upon a M -core system, each of speed $\mathcal{V} = 2 - 1/(M + 1 - \min_i\{m_i\})$.*

Proof Because $1 \leq m_i \leq M$ for any i , we know that for all $\tau_i \in \tau$,

$$\begin{aligned}\mathcal{V} &= 2 - \frac{1}{M + 1 - \min_i\{m_i\}} \\ &\geq 2 - \frac{1}{M + 1 - m_i} \\ &= \frac{2M + 1 - 2m_i}{M + 1 - m_i};\end{aligned}\tag{20}$$

From feasibility of the LL task set on a speed- M uniprocessor, we have $U_{sum} \leq M$. So,

$$\begin{aligned}(20) \implies \mathcal{V} &\geq \frac{U_{sum} + M + 1 - 2m_i}{M + 1 - m_i} \\ \iff \frac{U_{sum}}{\mathcal{V}} &\leq M - (m_i - 1) + \frac{(2m_i - M - 1)}{\mathcal{V}} \\ \iff m_i \frac{U_{sum}}{\mathcal{V}} &\leq m_i M - m_i(m_i - 1) + \frac{m_i}{\mathcal{V}}(2m_i - M - 1) \\ \implies m_i \frac{U_{sum}}{\mathcal{V}} &\leq m_i M - m_i(m_i - 1) + \frac{u_i}{\mathcal{V}}(2m_i - M - 1) \\ &[u_i \leq m_i, 2m_i - M - 1 \leq 0]\end{aligned}\tag{21}$$

The condition $2m_i - M - 1 \leq 0$ is equivalent to (18); while $u_i \leq m_i$ is true for any gang task because the utilization of each gang task τ_i is $u_i = m_i(c_i/T_i)$, where $c_i \leq T_i$. Note that u_i can also be viewed as the total utilization of the

m_i de-ganged LL tasks that correspond to the gang task τ_i . Again, de-ganging preserves the utilization of the set. From Equation (21):

$$\begin{aligned} m_i \frac{U_{sum}}{\mathcal{V}} &\leq m_i M - (m_i - 1)m_i + (m_i - 1) \frac{u_i}{\mathcal{V}} + \frac{u_i}{\mathcal{V}}(m_i - M) \\ \implies m_i \frac{U_{sum}}{\mathcal{V}} &\leq m_i M - (m_i - 1)(m_i - \frac{u_i}{\mathcal{V}}) - \frac{u_i}{\mathcal{V}}(M - m_i) \\ \implies m_i \frac{U_{sum}}{\mathcal{V}} &\leq m_i M - \Delta_i(m_i - \frac{u_i}{\mathcal{V}}) - \frac{u_i}{\mathcal{V}}(M - m_i) \end{aligned}$$

[From Definition 4: $0 \leq \Delta_i \leq m_i - 1$]

$$\begin{aligned} &= m_i(M - \Delta_i) - (M - \Delta_i - m_i) \frac{u_i}{\mathcal{V}} \quad [\text{re-arrange}] \\ &= (M - \Delta_i)(m_i - \frac{u_i}{\mathcal{V}}) + m_i \frac{u_i}{\mathcal{V}} \quad [\text{re-arrange}] \\ \implies \frac{U_{sum}}{\mathcal{V}} &\leq (M - \Delta_i)(1 - \frac{u_i}{m_i \mathcal{V}}) + \frac{u_i}{\mathcal{V}} \\ &[\text{divide } m_i \text{ on both sides, re-arrange}], \text{ for all } i \end{aligned} \tag{22}$$

The equation above implies that the corresponding gang set is GEDF schedulable on M speed- \mathcal{V} processors (Theorem 2 in [21]). Note that the last step is true because under speed of \mathcal{V} , all utilizations in the test should be treated as the ones under speed 1 divided by \mathcal{V} , in order to apply the original schedulability test under a speed-1 platform.

Theorem 2 indicates that the speedup factor of the GEDF schedulability test in Theorem 2 of [21] (for gang task set) is no greater than $\mathcal{V} = 2 - 1/(M + 1 - \min_i\{m_i\})$. Because $1 \leq m_i \leq M$ for any i , $\mathcal{V} \leq 2 - \frac{1}{M}$. Therefore, the following corollary follows directly from Theorem 2.

Corollary 1 *Given any de-ganged task set that is schedulable on a speed- M uni-processor, the corresponding gang task set will pass the schedulability test of GEDF upon a M -core system, each of speed $(2 - \frac{1}{M})$.*

Furthermore, scaling all speeds by a factor of $1/(2 - \frac{1}{M})$ lead to the following corollary.

Corollary 2 *Given any de-ganged task set that is schedulable on a speed- ψ uni-processor, the corresponding gang task set will pass the schedulability test of GEDF upon M unit-speed processors, where $\psi = M/(2 - \frac{1}{M})$.*

4.2 Speedup Bound for Gang Tasks under GEDF-VD

The previous subsection proved the speedup bound for non-MC task under GEDF. We now brings MC and virtual deadlines into the picture, and derive the speedup bound for *MC gang task set* τ under GEDF-VD. From the definitions of ϕ and Φ in Equation (19), the following properties hold:

$$1 + \Phi = \phi = \frac{1}{\Phi} \quad (23)$$

$$\Phi + \Phi^2 = 1 \quad (24)$$

Theorem 3 *Given any de-ganged MC task set that is schedulable on a speed- $(\psi \cdot \Phi)$ uniprocessor, the corresponding MC gang task set will be schedulable under GEDF-VD upon M unit-speed processors, where $\psi = M/(2 - \frac{1}{M})$ and $\Phi = \frac{\sqrt{5}-1}{2}$.*

Proof The de-ganged MC task set being schedulable on a speed- $(\psi \cdot \Phi)$ uniprocessor implies

$$\max\{U_{\text{LO}}^{\text{LO}} + U_{\text{HI}}^{\text{LO}}, U_{\text{HI}}^{\text{HI}}\} \leq \psi \cdot \Phi. \quad (25)$$

We proceed the rest of this proof in two cases.

Case 1: $U_{\text{HI}}^{\text{LO}} \geq \Phi \cdot U_{\text{LO}}^{\text{LO}}$. By (25) and the condition of Case 1,

$$\begin{aligned} \psi \cdot \Phi &\geq U_{\text{LO}}^{\text{LO}} + U_{\text{HI}}^{\text{LO}} \geq (1 + \Phi)U_{\text{LO}}^{\text{LO}} \\ \implies U_{\text{LO}}^{\text{LO}} &\leq \frac{\Phi}{1 + \Phi} \cdot \psi = \Phi^2 \cdot \psi. \text{ [by (23)]} \end{aligned}$$

Then, by the above and (25),

$$U_{\text{LO}}^{\text{LO}} + U_{\text{HI}}^{\text{HI}} \leq \Phi^2 \cdot \psi + \Phi \cdot \psi = \psi. \text{ [by (24)]}$$

Thus, no virtual deadline needs to be set at all. Both HI- and LO-criticality tasks are scheduled by GEDF according to their actual deadlines on M unit-speed processors. By Corollary 2, no deadline will be missed.

Case 2: $U_{\text{HI}}^{\text{LO}} < \Phi \cdot U_{\text{LO}}^{\text{LO}}$. By (25) and the condition of Case 2,

$$\begin{aligned} \psi \cdot \Phi &\geq U_{\text{LO}}^{\text{LO}} + U_{\text{HI}}^{\text{LO}} > \left(\frac{1}{\Phi} + 1\right)U_{\text{HI}}^{\text{LO}} \\ &= \frac{1 + \Phi}{\Phi} \cdot U_{\text{HI}}^{\text{LO}} = \frac{1}{\Phi^2} \cdot U_{\text{HI}}^{\text{LO}}. \text{ [by (23)]} \end{aligned}$$

That is,

$$U_{\text{HI}}^{\text{LO}} < \Phi^3 \cdot \psi. \quad (26)$$

Then, we have

$$\begin{aligned} U_{\text{LO}}^{\text{LO}} + \frac{U_{\text{HI}}^{\text{LO}}}{1 - U_{\text{HI}}^{\text{HI}}/\psi} &= U_{\text{LO}}^{\text{LO}} + U_{\text{HI}}^{\text{LO}} + U_{\text{HI}}^{\text{LO}} \cdot \frac{U_{\text{HI}}^{\text{HI}}/\psi}{1 - U_{\text{HI}}^{\text{HI}}/\psi} \\ &\leq U_{\text{LO}}^{\text{LO}} + U_{\text{HI}}^{\text{LO}} + U_{\text{HI}}^{\text{LO}} \cdot \frac{\Phi}{1 - \Phi} \\ &\quad [U_{\text{HI}}^{\text{HI}} \leq \psi \cdot \Phi \text{ by (25)}] \\ &= U_{\text{LO}}^{\text{LO}} + U_{\text{HI}}^{\text{LO}} + U_{\text{HI}}^{\text{LO}} \cdot \frac{\Phi}{\Phi^2} \text{ [by (24)]} \\ &< \psi \cdot \Phi + \Phi^3 \cdot \psi \cdot \frac{\Phi}{\Phi^2} \text{ [by (25) and (26)]} \\ &= (\Phi + \Phi^2)\psi \text{ [rearrange]} = \psi, \text{ [by (24)]} \end{aligned}$$

which is concluded as

$$U_{LO}^{LO} + \frac{U_{HI}^{LO}}{1 - U_{HI}^{HI}/\psi} < \psi. \quad (27)$$

In this case, one could take $x = \frac{U_{HI}^{LO}}{\psi - U_{LO}^{LO}}$ as the scaling factor to set the virtual deadlines for HI-criticality tasks. Because the de-ganged task set is schedulable on a speed- $(\psi \cdot \Phi)$ uniprocessor, $U_{LO}^{LO} \leq \psi \cdot \Phi$, which implies $x > 0$, as $\Phi < 1$, $\psi > 0$, and $U_{HI}^{LO} > 0$. On the other hand, $U_{HI}^{LO} < \Phi \cdot U_{LO}^{LO}$ in Case 2, so

$$\begin{aligned} x &= \frac{U_{HI}^{LO}}{\psi - U_{LO}^{LO}} < \frac{\Phi \cdot U_{LO}^{LO}}{\psi - U_{LO}^{LO}} \\ &\leq \frac{\Phi \cdot \psi \cdot \Phi}{\psi - \psi \cdot \Phi} = \frac{\Phi^2}{1 - \Phi} = 1. \text{ [by (24)]} \end{aligned}$$

Thus, in this case, $0 < x < 1$ is guaranteed under this particular setting and therefore this x can always be used as the scaling factor to set the virtual deadlines for GEDF-VD.

Then, we first show that all LO-criticality tasks meet their actual deadlines and all HI-criticality tasks meet their *virtual* deadlines during the LO-criticality mode.

$$U_{LO}^{LO} + \frac{U_{HI}^{LO}}{x} = U_{LO}^{LO} + \psi - U_{LO}^{LO} = \psi.$$

By Corollary 2, the above equation implies that using GEDF-VD to schedule the gang task set on M unit-speed processors, all LO-tasks meet their actual deadlines and all HI-tasks meet their *virtual* deadlines during the LO-mode. Next, we show that all HI-criticality tasks, including any carryover (HI-) jobs across the mode-switch point, meet their *actual* deadlines during the HI-criticality mode.

Because the virtual deadlines are set as $x \cdot T_i$ for each HI-criticality task τ_i , every HI-criticality job including the one triggering the mode switch will have *at least* $(1 - x)T_i$ time units to finish its *at most* C_i^{HI} execution and to release its next job. It suffices to consider the schedulability when replacing each HI-criticality task in the HI-criticality mode by a implicit-deadline sporadic task with period $(1 - x)T_i$ and execution C_i^{HI} . It can be done by checking their total utilization

$$\sum_{\tau_i \in \tau_{HI}} \frac{C_i^{HI}}{(1 - x)T_i} = \frac{U_{HI}^{HI}}{1 - x}.$$

On the other hand, by (27), we have

$$\frac{U_{HI}^{LO}}{\psi - U_{LO}^{LO}} < 1 - U_{HI}^{HI}/\psi,$$

and $1 - x > U_{HI}^{HI}/\psi$ holds since we set $x = \frac{U_{HI}^{LO}}{\psi - U_{LO}^{LO}}$. Thus,

$$\frac{U_{HI}^{HI}}{1 - x} < \frac{U_{HI}^{HI}}{U_{HI}^{HI}/\psi} = \psi.$$

By Corollary 2, the above equation implies that using GEDF-VD to schedule the gang task set on M unit-speed processors, all HI-criticality tasks, including any carryover (HI-) jobs across the mode-switch point, meet their *actual* deadlines during the HI-criticality mode.

Finally, we can easily use Theorem 3 to derive a speedup bound for GEDF-VD to schedule MC gang task sets on identical processors, as stated in the following theorem.

Theorem 4 *If any potentially optimal algorithm can schedule an MC gang task set on M unit-speed processors, GEDF-VD is able to schedule the same MC gang task set on M speed- $(\sqrt{5} + 1)$ processors.*

Proof Theorem 3 directly implies that:

If any potentially optimal algorithm can schedule an MC gang task set on M speed- $(\psi \cdot \Phi/M)$ processors, GEDF-VD is able to schedule the same MC gang task set on M unit-speed processors.

This is because for a MC gang task set to be schedulable on M speed- $(\psi \cdot \Phi/M)$ processors, it is necessary for its corresponding de-ganged MC task set to be schedulable on a speed- $(\psi \cdot \Phi)$ uniprocessor. Note that, by definitions: $\psi = \frac{M}{2 - \frac{1}{M}}$ and $\Phi = \frac{\sqrt{5}-1}{2}$, the following statement is true:

If any potentially optimal algorithm can schedule a MC gang task set on M speed- $(\frac{1}{2 - \frac{1}{M}} \cdot \frac{\sqrt{5}-1}{2})$ processors, GEDF-VD is able to schedule the same MC gang task set on M unit-speed processors.

Scaling the speed unit up by $(2 - \frac{1}{M}) \frac{\sqrt{5}+1}{2}$ (please note that $\frac{\sqrt{5}-1}{2} \cdot \frac{\sqrt{5}+1}{2} = 1$), the above statement can be re-written as:

If any potentially optimal algorithm can schedule a MC gang task set on M unit-speed processors, GEDF-VD is able to schedule the same MC gang task set on M speed- $(2 - \frac{1}{M}) \frac{\sqrt{5}+1}{2}$ processors.

Since $(2 - \frac{1}{M}) \frac{\sqrt{5}+1}{2} < \sqrt{5} + 1$, the theorem follows.

5 Evaluation

In this section, we evaluate the performance of GEDF-VD through simulation results. While the simulation results provide some representation of the proposed scheduling's performance, they may not represent the exact behavior of our proposed approach in real systems for several reasons. For example, memory plays a vital role from an implementation point of view and needs to be available and allocated to parallel threads. Although a recent work [19] has implemented a preemptive EDF scheduler for GPU tasks providing bandwidth isolation, MC scheduling on the GPU platform (with preemptive EDF) and the memory partitioning technique to the gang tasks are yet to be explored.

In the future, we plant to explore implementation and experimentation on a real hardware platform. As our work is the first to propose MC gang task scheduling, there is no perfect baseline for comparison. We have performed many experiments by varying different factors to observe the efficiency of our algorithm.

5.1 Experimental Setup

Workload generation. We generate MC gang tasks based on the following parameters.

- M : The number of processor cores.
- $m_{min}, m_{max}, m_{avg}$: The minimum, maximum, and average value for m (i.e., degree of parallelism), respectively. We generate the task set by varying these three parameters, where $m_{min}, m_{max} \in [1, M]$ and $m_{min} \leq m_{avg} \leq m_{max}$.
- U_{avg} : The average utilization for the task set. We have varied U_{avg} value from $0.05 \times M$ to $0.95 \times M$ with $0.05 \times M$ difference at each step.
- $P_{HI} = 0.5$: The probability of a task $\tau_i \in \tau_{HI}$.
- R : Denotes the maximum ratio of u_i^{HI} to u_i^{LO} , where $R \in [4, 8]$. We generate u_i^{HI} uniformly from $[u_i^{LO}, R \times u_i^{LO}]$.

At first, for a specific value of n (number of tasks per task set), we generate the m values for each task. m is uniformly generated from $[m_{min}, m_{max}]$ range in a way so that the average m for all tasks remains equal to m_{avg} . Next, for a specific value of average utilization U_{avg} , we calculate the average utilization u_i^a for each task by following the log-normal distribution. Note that, for n number of gang tasks, there are total $\sum_{i=1}^n m_i = m_{avg} \times n$ amount of single task instances in each task set. For the sake of a proper distribution, we extend the *UUniFast algorithm* [15] for Gang task. We use log-normal distribution over $\sum_{i=1}^n m_i$ task instances similarly as UUnifast, but for a single task, we take the average of all of its instances as the task's average utilization. The values of u_i^{LO} is uniformly generated from $[\frac{2 \times u_i^a}{R+1}, u_i^a]$ so that the value of u_i^{HI} is always in the range $[u_i^{LO}, R \times u_i^{LO}]$.

Simulation setup. We performed the simulation for average utilization ranging from $0.05M$ to $0.95M$ with a step size of $0.05M$. For each average, 100 task sets (each with 10 tasks) are generated.

5.2 Evaluation Results

We execute a set of gang tasks under our proposed algorithm by varying different parameters, and present the simulation results in Figure 3, 4, 5, 6, 7, 8, 9, and in Table 2. Note that, in Table 2, and in Figure 4-9, we consider that the degree of parallelism (of the HI-criticality tasks) does not change after a mode-switch.

Effect of changing the degree of parallelism after a mode switch. In this experiment, we set the value for M to 8, and vary (i.e., increase or un-

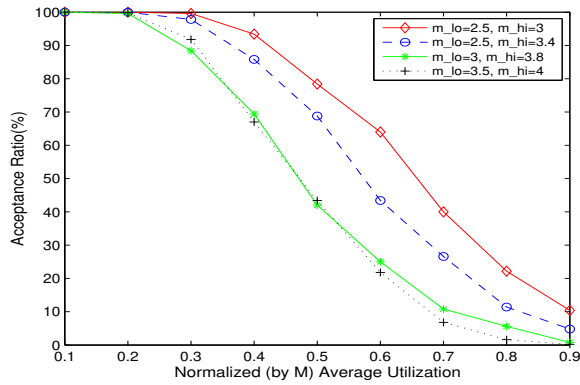


Fig. 3: Acceptance ratio for GEDF-VD with a different (after a mode-switch) average degrees of parallelism.

Table 2: Acceptance ratio for different amount of tasks generated under various average utilization and R value.

$U_{avg} \rightarrow$ # of tasks \downarrow	R	2	2.5	3	3.5	4	4.5	5
8	4	100	100	97	59	5	5	4
12	4	100	99	94	65	2	2	2
16	4	100	100	98	50	0	0	0
8	8	97	75	67	31	23	19	11
12	8	100	100	74	58	40	17	6
16	8	100	98	91	77	48	34	11

changed) the degree of parallelism of a HI-criticality task after a mode-switch. To incorporate the change in the task model, i.e., the degree of parallelism of a HI-criticality task can change after a mode-switch, we slightly modify the simulation setup described in Subsection 5.1. We report the acceptance ratio (in Figure 3) with a different average of m_i^{LO} (i.e., from 2.5 to 3.5) and m_i^{HI} (i.e., from 3 to 4) values. This figure reports that the *acceptance ratio* (i.e., the ratio of the number of schedulable task sets over the total number of task sets) decreases when the degree of parallelism increases, which can be explained by Equations (8) and (16). That is, a higher value of m_i^{LO} or m_i^{HI} inversely affects the acceptance ratio.

Effect of changing the degree of parallelism in a range with lower difference. In this set of experiments, for $M = 8$, we vary a task's degree of parallelism (m) in a different range, while the difference between the upper and lower bound in each range is fixed. The acceptance ratio under varying degree of parallelism (and different R values) is reported in Figure 4 and Figure 5. These figures indicate that in boundary cases (where the degree of parallelism is very low or very high) acceptance ratio changes proportionally with respect to the degree of parallelism. This behavior can be explained with

the help of Equations (8) and (16). When m increases or decreases by a large amount, acceptance ratio will increase or decrease respectively. However, for a small change of m , acceptance ratio may not change proportionally. This is because the schedulability conditions provided by Equations (8) and (16) are also effected by the maximum number of idle cores (Δ_i), which is dependent on m .




Fig. 4: Acceptance ratio for GEDF-VD in an 8-core platform with $R = 4$, and under same ranges of degrees of parallelism.




Fig. 5: Acceptance ratio for GEDF-VD in an 8-core platform with $R = 8$, and under same ranges of degrees of parallelism.

Effect of changing the total number of cores. In Figure 6 and Figure 7, we report the acceptance ratio of the task set by varying the number of cores




Fig. 6: Acceptance ratio for GEDF-VD in an M -core platform (with $R = 4$). In this experiment $M \in \{4, 8, 16, 32\}$.




Fig. 7: Acceptance ratio for GEDF-VD in an M -core platform (with $R = 8$). In this experiment $M \in \{4, 8, 16, 32\}$.

in the system, M . In this set of experiments, we set a value for m_{avg} , which is uniformly generated from a range of $[\frac{M}{2}, \frac{3M}{4}]$. Simulations are conducted for $M = 4, 8, 16$, and 32 and the average utilization is weighted with respect to the value of M . Figure 6 and Figure 7 shows that the acceptance ratio is not affected by different values of M and remains almost unchanged.

Effect of changing number of tasks per task set. In this set of experiments, we have randomly generated 100 task sets with different R values and 8, 12, and 16 tasks per task set (with U_{avg} changing from 2 to 5 with a step size of 0.5) and report the acceptance ratio in Table 2. From the reported data, it is clear that the acceptance ratio of the task set is not affected by the number of tasks per set. This result indicates the effectiveness of our proposed algorithm under a varying number of tasks in a task set.




Fig. 8: Acceptance ratio for GEDF-VD in an 8-core platform with $R = 4$ and a varying range of m_{avg} .

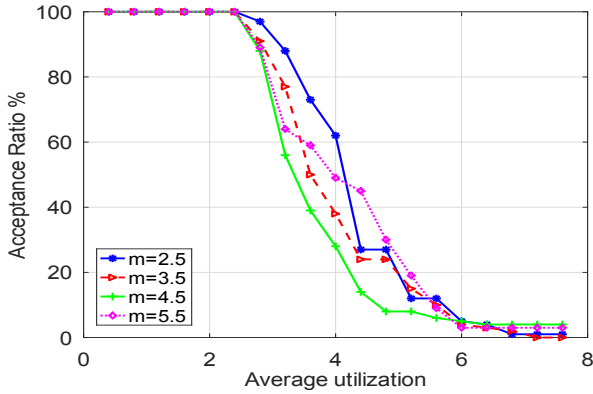


Fig. 9: Acceptance ratio for GEDF-VD in an 8-core platform with $R = 8$ and a varying range of m_{avg} .

Effect of changing m_{avg} value. In Figure 8 and Figure 9, we show the acceptance ratio by varying m_{avg} in an 8-core platform. The result does not demonstrate a direct relationship between m_{avg} and the acceptance ratio.

6 Related Work

Since Vestal's proposal [49] of MC workload model, much work has focused on scheduling MC tasks (refer to Burns et al. [17] for a survey). For uniprocessor platforms, many algorithms were proposed based on both fixed priority (e.g., Li et al. [36], Baruah et al. [8]) and dynamic priority scheduling (e.g., Easwaran et al. [22]). The work in [13,12] proposed the precise scheduling

policy, where all LO-criticality tasks receive a full-service guarantee even after a mode switch. On the other hand, numerous MC scheduling algorithms were proposed for multiprocessor platforms [34,9,47,3,48]. Considering the multiprocessor platforms, Lee et al. [34] and Baruah et al. [9] proposed fluid-based MC models, and a semi-partitioned based scheme is proposed by Awan et al. [3].

Considering different parallel tasks models (e.g., synchronous task model [2], DAG model [16,38,32,11,10,31] and gang models [33,30,21]) there have been a number of works that have provided the energy efficiency technique, schedulability analysis, and the speedup bound (i.e., resource augmentation bound) for various scheduling strategies. For synchronous tasks under GEDF scheduling, Andersson et al. [2] proved a resource augmentation bound of 2 with constrained deadlines tasks. Considering DAG tasks (with arbitrary deadlines) under GEDF, Li et al. [37] and Bonifaci et al. [16] simultaneously proved a resource augmentation bound of 2. Bonifaci et al. [16] also showed the bound to be 3 under global rate-monotonic scheduling. For implicit deadline DAG tasks under federated scheduling, a resource augmentation bound of 2 is showed by Li et al. [38]. Gang scheduling and Coscheduling was initially introduced to perform parallel processing with fine-grained interactions efficiently [27,43,28]. Both of these approaches allocate resources to the threads of the same task concurrently. However, gang scheduling imposes a strict requirement of executing all threads of the same task simultaneously. In contrast, in coscheduling, some threads may not execute concurrently with the rest of the threads in the same task. Some recent work used this concept to execute the parallel workload in cloud computing [46] and extended to incorporate hard real-time tasks [30]. The work in [30] also has proposed a DP-Fair based scheduling of periodic gang tasks and proved a speedup bound which is no larger than $(2 - 1/m)$. A recent work by Alahmad et al. [1] proposed the isochronous scheduling, which has some similarity to the traditional gang scheduling. Unlike the gang scheduling, the isochronous model assumed that the job versions might not be compatible with all the processors available. Although the work in [1] has some connections to the MC task model, they did not explicitly concern different criticality levels of a task. It aims to achieve higher design assurance levels by using adequate monitoring and improving mechanisms. In contrast, the MC scheduling that we propose focuses on providing service guarantee only to the high criticality jobs where computational resources are not adequate. Kato et al. [33] introduced gang task scheduling based on global EDF. Dong et al. [21] proposed a schedulability analysis based on lag-based reasoning. Few other related works, such as Goossens et al. [29], provided schedulability tests for fixed task-priority scheduling of real-time periodic gang tasks.

Although a good number of works studied MC scheduling and parallel tasks scheduling individually, very few works studied the scheduling of MC parallel tasks [40,4,39,44]. Rambo et al. [44] proposed a replica-aware co-scheduling approach (that is a combination of strict priority preemptive (SPP) policy and gang scheduling policy) for mixed-critical systems. Baruah et al. [4]

and Li et al. [39] proposed the MC scheduling of DAG models, while Liu et al. [40] proposed the MC scheduling of the synchronous task model. Unlike these works, we consider the gang task model, where a task cannot execute if the number of available cores is less than its degree of parallelism. This constraint makes the scheduling problem highly challenging.

7 Conclusion

Parallel computing with real-time constraints is gaining popularity due to its broad applicability and system efficiency [33, 2, 16, 38, 30, 21, 32, 10, 31]. WCET measurements are pessimistic due to increased uncertainty. So, there is an emerging need to introduce MC into parallel computation models and system designs. We leverage two existing algorithms (EDF-VD and GEDF) to schedule MC gang tasks efficiently in this work. We derive the first speedup bound for GEDF schedulability of (non-MC) gang tasks and further derived the bound for GEDF-VD of MC gang tasks.

Future Work. This work is an initial step of more substantial efforts in bringing richer system modeling and analysis into the emerging need in many applications for parallel computing and MC. In the future, by ensuring the a priori verification, we plan to consider run-time robustness, with moldable or malleable models for MC gang tasks. In this work, we have considered the preemptive scheduling, where a high priority task can interrupt a running task, and offers more flexibility in task scheduling. However, preemptive scheduling suffers from the switching overhead as a task from running state to ready state and vice-versa. In the future, we plan to study the non-preemptive scheduling (which does not suffer from task switching overhead) and provide a detailed comparison between these two directions. We also plan to evaluate our results (in this paper) by implementations on applicable hardware platforms.

Acknowledgements We thank the reviewers for their constructive feedback to improve this paper. This work is partially supported by NSF grant CNS-1850851.

References

1. Alahmad, B., Gopalakrishnan, S.: Isochronous execution models for high-assurance real-time systems. In: HASE. IEEE (2019)
2. Andersson, B., de Niz, D.: Analyzing global-EDF for multiprocessor scheduling of parallel tasks. In: OPODIS. Springer (2012)
3. Awan, M., Bletsas, K., Souto, P., Tovar, E.: Semi-partitioned mixed-criticality scheduling. In: ARCS. Springer (2017)
4. Baruah, S.: The federated scheduling of systems of mixed-criticality sporadic DAG tasks. In: RTSS. IEEE (2016)
5. Baruah, S.: Mixed-criticality scheduling theory: Scope, promise, and limitations. IEEE DESIGN AND TEST **35**(2), 31–37 (2018)
6. Baruah, S., Bertogna, M., Buttazzo, G.: Multiprocessor Scheduling for Real-Time Systems. Springer (2015)

7. Baruah, S., Bonifaci, V., D'Angelo, G., Li, H., Marchetti-Spaccamela, A., Van Der Ster, S., Stougie, L.: The preemptive uniprocessor scheduling of mixed-criticality implicit-deadline sporadic task systems. In: ECRTS. IEEE (2012)
8. Baruah, S., Bonifaci, V., D'angelo, G., Marchetti-Spaccamela, A., Van Der Ster, S., Stougie, L.: Mixed-criticality scheduling of sporadic task systems. In: ESA. Springer (2011)
9. Baruah, S., Eswaran, A., Guo, Z.: MC-Fluid: simplified and optimally quantified. In: RTSS. IEEE (2015)
10. Bhuiyan, A., Guo, Z., Saifullah, A., Guan, N., Xiong, H.: Energy-efficient real-time scheduling of DAG tasks. *ACM Transactions on Embedded Computing Systems (TECS)* **17**(5), 84 (2018)
11. Bhuiyan, A., Liu, D., Khan, A., Saifullah, A., Guan, N., Guo, Z.: Energy-efficient parallel real-time scheduling on clustered multi-core. *IEEE Transactions on Parallel and Distributed Systems* **31**(9), 2097–2111 (2020)
12. Bhuiyan, A., Reghenzani, F., Fornaciari, W., Guo, Z.: Optimizing energy in non-preemptive mixed-criticality scheduling by exploiting probabilistic information. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **39**(11), 3906–3917 (2020)
13. Bhuiyan, A., Sruti, S., Guo, Z., Yang, K.: Precise scheduling of mixed-criticality tasks by varying processor speed. In: RTNS (2019)
14. ahmed Bhuiyan, A., Yang, K., Arefin, S., Saifullah, A., Guan, N., Guo, Z.: Mixed-criticality multicore scheduling of real-time gang task systems. In: 2019 IEEE Real-Time Systems Symposium (RTSS), pp. 469–480. IEEE (2019)
15. Bolado, M., Posadas, H., Castillo, J., Huerta, P., Sanchez, P., Sánchez, C., Fouren, H., Blasco, F.: Platform based on open-source cores for industrial applications. In: Design, Automation and Test in Europe Conference and Exhibition, 2004. Proceedings, vol. 2, pp. 1014–1019. IEEE (2004)
16. Bonifaci, V., Marchetti-Spaccamela, A., Stiller, S., Wiese, A.: Feasibility analysis in the sporadic dag task model. In: ECRTS. IEEE (2013)
17. Burns, A., Davis, R.: Mixed criticality systems—a review. Department of Computer Science, University of York, Tech. Rep pp. 1–69 (2013)
18. Burns, A., Davis, R.: A survey of research into mixed criticality systems. *ACM Computing Surveys (CSUR)* **50**(6), 82 (2018)
19. Capodiceci, N., Cavicchioli, R., Bertogna, M., Paramakuru, A.: Deadline-based scheduling for GPU with preemption support. In: RTSS. IEEE (2018)
20. Döbel, B., Härtig, H., Engel, M.: Operating system support for redundant multithreading. In: Proceedings of the tenth ACM international conference on Embedded software, pp. 83–92 (2012)
21. Dong, Z., Liu, C.: Analysis techniques for supporting hard real-time sporadic gang task systems. In: RTSS. IEEE (2017)
22. Easwaran, A.: Demand-based scheduling of mixed-criticality sporadic tasks on one processor. In: RTSS. IEEE (2013)
23. Elliott, G.A., Ward, B.C., Anderson, J.H.: GPUSync: A framework for real-time GPU management. In: RTSS. IEEE (2013)
24. Engel, M., Döbel, B.: The reliable computing base—a paradigm for software-based reliability. *INFORMATIK 2012* (2012)
25. Ernst, R., Di Natale, M.: Mixed criticality systems—a history of misconceptions? *IEEE Design & Test* **33**(5), 65–74 (2016)
26. Esper, A., Nelissen, G., Nélis, V., Tovar, E.: How realistic is the mixed-criticality real-time system model? In: RTNS. ACM (2015)
27. Feitelson, D.G., Rudolph, L.: Gang scheduling performance benefits for fine-grain synchronization. *Journal of Parallel and distributed Computing* **16**(4), 306–318 (1992)
28. Gehringer, E.F., Siewiorek, D.P., Segall, Z.: *Parallel processing: the Cm* experience*. Digital Press (1987)
29. Goossens, J., Berten, V.: Gang FTP scheduling of periodic and parallel rigid real-time tasks. arXiv preprint arXiv:1006.2617 (2010)
30. Goossens, J., Richard, P.: Optimal scheduling of periodic gang tasks. *Leibniz transactions on embedded systems* **3**(1), 04–1 (2016)

31. Guo, Z., Bhuiyan, A., Liu, D., Khan, A., Saifullah, A., Guan, N.: Energy-efficient real-time scheduling of DAGs on clustered multi-core platforms. In: RTAS. IEEE (2019)
32. Guo, Z., Bhuiyan, A., Saifullah, A., Guan, N., Xiong, H.: Energy-efficient multi-core scheduling for real-time DAG tasks (2017)
33. Kato, S., Ishikawa, Y.: Gang EDF scheduling of parallel task systems. In: 30th IEEE Real-Time Systems Symposium, pp. 459–468. IEEE (2009)
34. Lee, J., Phan, K., Gu, X., Lee, J., Easwaran, A., Shin, I., Lee, I.: MC-Fluid: Fluid model-based mixed-criticality scheduling on multiprocessors. In: RTSS. IEEE (2014)
35. Li, H.: Scheduling mixed-criticality real-time systems. Ph.D. thesis, The University of North Carolina at Chapel Hill (2013)
36. Li, H., Baruah, S.: An algorithm for scheduling certifiable mixed-criticality sporadic task systems. In: RTSS. IEEE (2010)
37. Li, J., Agrawal, K., Lu, C., Gill, C.: Analysis of global EDF for parallel tasks. In: ECRTS. IEEE (2013)
38. Li, J., Chen, J.J., Agrawal, K., Lu, C., Gill, C., Saifullah, A.: Analysis of federated and global scheduling for parallel real-time tasks. In: ECRTS. IEEE (2014)
39. Li, J., Ferry, D., Ahuja, S., Agrawal, K., Gill, C., Lu, C.: Mixed-criticality federated scheduling for parallel real-time tasks. *Real-Time Systems* **53**(5), 760–811 (2017)
40. Liu, G., Lu, Y., Wang, S., Gu, Z.: Partitioned multiprocessor scheduling of mixed-criticality parallel jobs. In: RTCSA. IEEE (2014)
41. (2017). <http://www.nvidia.com/page/home.html>
42. (2017). <https://www.openacc.org/>
43. Ousterhout, J.K., et al.: Scheduling techniques for concurrent systems. In: ICDCS, vol. 82, pp. 22–30 (1982)
44. Rambo, E.A., Ernst, R.: Replica-aware co-scheduling for mixed-criticality. In: ECRTS 2017. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2017)
45. Spuri, M., Buttazzo, G.C.: Efficient aperiodic service under earliest deadline scheduling. In: RTSS, pp. 2–11 (1994)
46. Stavrinides, G.L., Karatza, H.D.: Scheduling real-time parallel applications in saas clouds in the presence of transient software failures. In: SPECTS. IEEE (2016)
47. Tobuschat, S., Ernst, R.: Efficient latency guarantees for mixed-criticality networks-on-chip. In: RTAS. IEEE (2017)
48. Trüb, R., Giannopoulou, G., Tretter, A., Thiele, L.: Implementation of partitioned mixed-criticality scheduling on a multi-core platform. *ACM Transactions on Embedded Computing Systems (TECS)* **16**(5s), 122 (2017)
49. Vestal, S.: Preemptive scheduling of multi-criticality systems with varying degrees of execution time assurance. In: RTSS. IEEE (2007)
50. Xiao, S., Feng, W.c.: Inter-block GPU communication via fast barrier synchronization. In: IPDPS. IEEE (2010)
51. Yang, M., Amert, T., Yang, K., Otterness, N., Anderson, J.H., Smith, F.D., Wang, S.: Making OpenVX really” real time”. In: RTSS. IEEE (2018)