

Mixed-Criticality Scheduling upon Permitted Failure Probability and Dynamic Priority

Zhishan Guo, *Senior Member, IEEE*, Sudharsan Vaidhun, *Student Member, IEEE*, Luca Satinelli, *Member, IEEE*, Samsil Arefin, Jun Wang, *Senior Member, IEEE*, and Kecheng Yang *Member, IEEE*,

Abstract—Many safety-critical real-time systems are considered certified when they meet failure probability requirements with respect to the maximum permitted incidences of failure per hour. In this paper, the mixed-criticality task model with multiple worst-case execution time (WCET) estimations is extended to incorporate such system-level certification restrictions. A new parameter is added to each task, characterizing the distribution of the WCET estimations – the likelihood of all jobs of a task finishing their executions within the less pessimistic WCET estimates. Efficient algorithms are derived for scheduling mixed-criticality systems represented using this model for both uniprocessor and multiprocessor platforms for independent tasks. Furthermore, a 0/1 covariance matrix is introduced to represent the failure-dependency between tasks. An efficient algorithm is proposed to schedule such failure-dependent tasks. Experimental analyses show our new model and algorithm outperform current state-of-the-art mixed-criticality scheduling algorithms.

Index Terms—Multicore real-time scheduling, Mixed-Criticality, Failure probability, Dependency.

I. INTRODUCTION

Safety-critical systems are as failure prone as any other system, and today's system certification approaches recognize this and specify *permitted system failure probabilities*. The underlying idea is to certify considering more realistic system models which account for any possible behavior, including faulty conditions, and the probability of these behaviors occurring. The gap that still exists is between such enhanced models and the current conservative deterministic analyses which tend to be pessimistic.

The *worst-case execution time* (WCET) abstraction plays a central role in the analysis of real-time systems. The WCET of a given piece of code upon a specified platform represents an upper bound to the duration of time needed to finish execution. Unfortunately, even when severe restrictions are placed upon the structure of the code (e.g., known loop bounds), it is still extremely difficult to determine the absolute WCET. An illustrative example is provided in [2], which demonstrates how the simple operation “ $a = b + c$ ” on integer variables

Manuscript received May 30, 2020; revised Sept. 15, 2020 and Dec. 2, 2020; accepted Jan. 16, 2021. A preliminary version of this paper with early results on uniprocessor was published in RTCSA 2015 [1]. This work was supported in part by National Science Foundation CNS-1850851, SPX-2028481.

Z. Guo, S. Vaidhun, and J. Wang are with the Department of Electrical and Computer Engineering, University of Central Florida, Orlando, FL 32816, USA. E-mail: zsguo@ucf.edu.

L. Satinelli is with ONERA, Toulouse, France and Airbus, Germany.

S. Arefin is with Microsoft New England Research & Development Center (NERD), Cambridge, MA 02142, USA.

K. Yang is with the Department of Computer Science, Texas State University, San Marcos, TX 78666, USA.

could take anywhere between 3 and 321 cycles upon a widely-used modern CPU. The number of execution cycles highly depends upon factors such as the state of the cache when the operation occurs. WCET analysis has always been a very active and thriving area of research, and sophisticated timing analysis tools have been developed (see [3] for an excellent survey).

Traditional rigorous WCET analysis may lead to a result of much pessimism, and the occurrence of such WCET is extremely unlikely, unless under highly pathological circumstances. For instance, although a conservative tool would assign the “ $a = b + c$ ” operation a WCET bound of 321 cycles, a less conservative tool may assign it a much smaller WCET (e.g., 30) with the understanding that the bound may be violated on rare occasions under certain (presumably highly unlikely to occur) pathological conditions.

Mixed-Criticality Systems. The gap between the actual running time and the WCET may be significantly large. Instead of completely wasting the processor capacities within the gap, recent researches focused on implementing functionalities of different degrees of importance, or *criticalities*, upon a common platform, so that the less important tasks that may execute in these gaps under normal circumstances, may be dropped in occasional situations where jobs of higher importance level execute beyond their estimated common case running time. Much prior research on mixed-criticality scheduling (see [4] for a review) has focused upon the phenomenon that different tools for determining WCET bounds may be more or less conservative than one another, which results in multiple WCET estimations for each individual task (piece of code). Typically in the two-criticality-level case, each task is designated as being of either higher (HI) or lower (LO) criticality, and two WCETs are specified for each HI-criticality task: a LO-WCET determined by a less pessimistic tool, and a larger HI-WCET determined by a more conservative one, which is sometimes larger than the LO-WCET by several orders of magnitude. The scheduling objective is to determine a run-time scheduling strategy which ensures that (i) all jobs of all tasks complete by their deadlines if each job completes upon executing for no more than its LO-WCET; and (ii) all jobs of tasks designated as being of HI criticality continue to complete by their deadlines (although the LO-criticality jobs may not) if any job requires execution for more than its LO-WCET (but no larger than its HI-WCET) to complete.

Under the current mixed-criticality model, it is assumed that *all* HI-criticality jobs may require executions up to their HI-WCETs in HI mode simultaneously. However, since WCET

tools are normally quite pessimistic, LO-WCET are not very likely to be exceeded during run-time.

Example I.1. Consider a system comprised of two independent¹ HI-criticality tasks τ_1 and τ_2 , where each task is denoted by two utilization estimations $u_{\text{LO}} \leq u_{\text{HI}}$. The two tasks $\tau_1 = \{0.4, 0.6\}$, $\tau_2 = \{0.3, 0.5\}$, represented by utilizations in different modes, are to be scheduled on a preemptive unit-speed uniprocessor. It is evident that this system cannot be scheduled correctly under the traditional model, since the HI-criticality utilization, at $(0.6 + 0.5)$, is greater than the processor capacity which is 1.

However, suppose that: (i) absolute certainty of correctness is not required; instead it is specified that the system failure probability should not exceed 10^{-6} per hour; and (ii) it is known that the timing analysis tools used to determine LO-criticality WCETs ensure that the likelihood of any job of a task exceeding its LO-WCET is no larger than 10^{-4} per hour.

Based on the task independence assumption, the probability of jobs from both tasks exceeding their LO-WCETs is $10^{-4} \times 10^{-4} = 10^{-8}$ per hour. Thus, we know that it is safe to ignore the case that both tasks simultaneously exceed their LO-WCETs. Hence, the system is probabilistically feasible, since the total remaining utilization will not exceed:

$$\max\{0.4 + 0.3, 0.4 + 0.5, 0.6 + 0.3\} = 0.9 \leq 1.$$

Example I.1 gives us an intuition that with the help of probabilistic analysis, we may be able to ignore some extremely unlikely cases, and come up with some *less pessimistic* schedulability analysis – if we have the prior knowledge that there will be at most a fixed number of HI-criticality tasks with execution exceptions per hour, then dropping of less important jobs may not be necessary at all.

In traditional MC models, each HI-criticality task is characterized by two WCETs, c_{LO} and c_{HI} , which could be derived with different timing analysis tools. By the level of pessimism and/or other properties in the timing analysis, such a tool usually provides a confidence for its resulting WCET estimates. *Confidence level* of the execution budget estimates is the likelihood that the estimate is true, given the information available. However, very few work on MC analysis has leveraged any information from the confidence of the provisioned WCET.

Existing MC analysis usually makes the most pessimistic assumption that *every* HI-criticality task may execute beyond its LO-WCET and reach its HI-WCET *simultaneously*. In real applications, the industry standards usually only require the expected probability of missing deadlines within a specified duration to be below some specified small value, as the deadline miss can be seen as a faulty condition. Instead, our work aims at leveraging probabilistic information from the timing analysis tools (i.e. confidence) to rule out the too pessimistic scenarios and to improve schedulability of the whole system under a probabilistic standard.

Our work also differs from most prior work on WCET analysis as follows: Existing timing analysis works usually

¹Two events are independent if the occurrence of one event does not have any impact on the other.

analyze the WCET for a task on a per-job basis; i.e., by focusing on the distribution of WCETs of jobs of a certain task[3]. When it comes to analyzing a series of consecutive jobs generated from the same task, the distribution is directly applied. It is usually assumed that i) all jobs WCET of a certain task obey the same distribution (identically distributed), and ii) the WCET of a job is probabilistically drawn from the distribution with no dependence on other jobs of the same task (independence). While the independence assumption holds for the WCET, as we will see in the next section, it may not hold for the task execution time. For example, in many applications such as video frames processing, the execution times of processing consecutive frames of a certain video are usually dependent. However, the event that a certain task has ever over run its provisioned execution time in time intervals of a certain adequate large length (e.g., an hour) is independent from the scenario in other such intervals, and the probability of such event should be derived from the confidence of corresponding timing analysis tools only.

Contributions. In addition to the existing mixed-criticality task model, this work introduces a new parameter to each task that represents the distribution information about its WCET. This work aims to provide schedulability analysis to instances with this additional probability information, with respect to the given safety certification requirement of the whole system, which is the permitted system failure probability per hour.

We consider the scheduling of dual-criticality task systems upon both preemptive uniprocessor and multiprocessor platforms. As stated above, dual-criticality tasks are traditionally characterized with two WCET estimations – a LO-WCET and a larger HI-WCET. Our contributions are as follows:

(i) We propose a supplement to current MC task models: an additional parameter for each HI-criticality task, denoting the *probability* of no job of this task exceeding its LO-WCET *within an hour* of execution.

(ii) We further generalize our notion of system behavior by allowing for the specification of a *permitted system failure probability per hour*, denoting an upper bound on the probability that the system may fail to meet its timing constraints during any hour of running.

(iii) We derive a novel scheduling algorithm (and an associated sufficient schedulability test) for a given MC task set and an allowed system failure probability on uniprocessor platforms. We seek to schedule the system such that the probability of failing to meet timing constraints during runtime is guaranteed to be no larger than the specified allowed system failure probability.

(iv) We further extend our uniprocessor scheduling technique to multiprocessor platforms by combining with the partitioned scheduling techniques.

(v) While our initial scheduling technique focuses on independent task sets, we further introduce a covariance matrix in the system model to represent the failure-dependencies between tasks in a task set, and we propose an efficient scheduling technique to schedule task sets with given failure-dependencies.

We emphasize that our algorithm, in the two criticality level case, requires just one probabilistic parameter per task – the

probability that the actual execution requirement will exceed the specified LO-WCET in an hour. We believe our scheduling algorithm is novel in that it is, to our knowledge, the first MC scheduling algorithm that makes scheduling decisions (e.g., when to trigger a mode switch) based not only on release times, deadlines, and WCETs, but also on the probabilities drawn from PTA tools (see, e.g., [5] [6] [7]).

Organization. Sec. II introduces the model and shows its advantage by a motivating example. Sec. III formally defines probabilistic schedulability and related concepts. In Sec. IV, we propose a clustering-based scheduling strategy, and the corresponding schedulability test, while Sec. V presents the multiprocessor scheduling algorithm and the corresponding schedulability test. In Sec. VI, we introduce the covariance matrix to represent the failure-dependencies between the tasks and present the scheduling technique for such task sets, and finally, Sec. VII performs their experimental evaluations and comparisons. Sec. VIII elaborates the existing results and Sec. IX concludes and suggests future work.

II. MODEL

We start out considering a workload model consisting of *independent implicit-deadline sporadic tasks*, where the deadline and the period of a task share the same value. In Sec. VI, we extend this model to allow for pairwise dependencies between tasks. Throughout this paper, an integer model of time is assumed — all task periods are assumed to be non-negative integers, and all job arrivals are assumed to occur at integer instants in time.

Before detailing our task model, a few statistical notions need to be introduced in order to clarify previous and next observations. Given a task τ_i , its pWCET estimate comes from a random variable (the worst-case execution time distribution), notably continuous distributions² denoted by \mathcal{C}_i . Equivalent representations for distributions are the probabilistic density functions (pdfs), $f_{\mathcal{C}_i}$, the Cumulative Distribution Functions (CDFs) $F_{\mathcal{C}_i}$, and the Complementary Cumulative Distribution Functions (CCDFs), $F'_{\mathcal{C}_i}$. In the following, calligraphic upper-case letters are used to refer to probabilistic distributions, while non calligraphic letters are used for single value parameters.

The CCDF representation relates *confidence* to *probabilities*; indeed, from $F'_{\mathcal{C}_i}(c_{(LO)})$ we have the probability of exceeding $c_{(LO)}$. *The confidence is then for $c_{(LO)}$ being an upper-bound to task execution time.* The WCET threshold, simply named pWCET or WCET in the rest of the paper, is a tuple $\langle c_{(LO)}, p_{(LO)} \rangle$, where the probability $p_{(LO)}$ sets the confidence (at the job level) of exceeding $c_{(LO)}$, $p_{(LO)} = F'_{\mathcal{C}_i}(c_{(LO)}) = P(\mathcal{C}_i > c_{(LO)})$. By decreasing the probability threshold $p_{(LO)}$, the confidence on the upper-bounding worst-case, $c_{(LO)}$, increases.

Given the event A that a job exceeds its threshold and its probability of happening $p^A = P(\mathcal{C}_i > c_{(LO)})$; given B the event that another job exceeds its threshold (in a different execution interval) with $p^B = P(\mathcal{C}_i > c_{(LO)})$ its probability of

²The timing analysis that makes use of the EVT, by definition provides continuous distributions as pWCET estimates [5]; they are then discretized, to ease their representation, by assigning them a discrete support.

happening. With separate jobs as well as separate execution intervals, and considering WCETs, the conditional probability $P(A|B)$ is equal to $P(A)$, thus the joint probability is

$$P(A, B) = P(A|B) \times P(B) = P(A) \times P(B), \quad (1)$$

due to the independence between WCETs. Projecting the per job probability threshold $p_{(LO)} = F'_{\mathcal{C}_i}(c_{(LO)})$ to one hour task execution interval, we make use of the joint probability of all the exceeding threshold events within the one hour interval. The joint probability is

$$1 - P(\mathcal{C}_i \leq c_{(LO)}, \mathcal{C}_i \leq c_{(LO)}, \dots, \mathcal{C}_i \leq c_{(LO)}), \quad (2)$$

as the probability of at least one task job exceeding its thresholds $c_{(LO)}$. With full independence, the probability of exceeding threshold in one hour would be at most $1 - F_{\mathcal{C}_i}(c_{(LO)}) \times [T_i/3,600,000]$, with the task τ_i period T_i expressed in msec.

III. PROBABILISTIC SCHEDULABILITY

System failure probability F_s . In our model, an *allowed system failure probability* F_S is specified. It describes the permitted probability of the system failing to meet timing constraints during one hour of execution³. F_S may be very close to zero (e.g., 10^{-12} for some safety critical avionics functionalities)⁴.

Failure probability. A *failure probability* parameter f_i is added to HI-criticality task τ_i , denoting the probability that the actual execution requirement of *any* job of the task exceeding $c_i(LO)$ (but still below $c_i(HI)$) in one hour (i.e., the adequate long time interval we assumed in this paper). f_i depends on a failure distribution $F_i(t)$ that describes the task τ_i probability of failure up to and including time t . Since $F_i(t)$ would refer to time (interval) and to task execution, it is going to be the one we computed for one hour interval or any another interval, Eq. (2). Thus, f_i can be directly derived from $F_{\mathcal{C}_i}$ ⁵.

A HI-criticality task is represented by: $\tau_i = ([c_i(LO), c_i(HI)], f_i, T_i, \chi_i)$, where T_i is the task period and $\chi_i \in \{LO, HI\}$ is the criticality level of the task; LO-criticality tasks continue to be represented with three parameters as before. This enhanced model is essentially asserting, for each HI-criticality task τ_i , within a time interval of one hour, no job of τ_i has an execution greater than $c_i(HI)$ and the probability of *any* job of τ_i having an execution greater than $c_i(LO)$ is f_i — we would expect f_i to be a very small positive value. In our work we assume $c_i(HI)$ the deterministic WCET, $\langle c_i(HI), 0 \rangle$, while $\langle c_i(LO), f_i \rangle > 0$ the probabilistic WCET with $c_i(LO) \leq c_i(HI)$. Normally we do not guarantee higher assurance for LO-criticality tasks (than HI-criticality ones), and thus only $c_i(LO)$ is adopted for them. In traditional MC model, the $c_i(LO)$ and $c_i(HI)$

³Failure probability are easily referable to failure rate, being careful at considering the failure rate as a probability.

⁴From DO-178B/C at the highest DAL - Level A, the acceptable failure rate is below $10^{-9}/h$.[8]

⁵It is possible to apply existing timing analysis tools to determine f_i — by monitoring executions of a piece of code for enough length, one may derive a stable pWCET, or may need to adapt EVT in case there are significant changes of execution time (to guarantee the safety of pWCET).

values are chosen from the worst-case execution time distribution. However, the corresponding probabilities, f_i and 0 respectively, are omitted from the task model. The proposed task model retains the probability information associated with the WCET estimates.

Definition 1 (MC Task Instance). *A MC task instance I is composed of a MC task set $\tau = \{\tau_1, \tau_2, \dots, \tau_n\}$ and a system failure requirement $F_S \in (0, 1)$. (Although F_S may be arbitrarily close to 0, $F_S = 0$ is not an acceptable value — “nothing is impossible.”)*

Let $n_{\text{HI}} \leq n$ denote the number of HI-criticality tasks in τ . We assume that the tasks are indexed such that the HI-criticality ones have lower indices; i.e., the HI-criticality tasks are indexed $1, 2, \dots, n_{\text{HI}}$.

We seek to determine the **probabilistic schedulability** of any given MC task instance:

Definition 2 (probabilistic schedulability). *A MC task set is strongly probabilistic schedulable by a scheduling strategy if it possesses the property that upon execution, the probability of missing any deadline is less than F_S . It is weakly probabilistic schedulable if the probability of missing any HI-criticality deadline is less than F_S . (In either case, all deadlines are met during system runs where no job exceeds its LO-WCET.)*

That is, if a schedulability test returns strongly schedulable, then all jobs meet their deadlines with a probability no less than $1 - F_S$, while weakly schedulable only guarantees (with probability no less than $1 - F_S$) that HI-criticality jobs meet their deadlines. Moreover, similar to all MC works, for either strongly or weakly probabilistic schedulable, all deadlines are met when all jobs finish upon executing their LO-WCETs. Again, F_S comes from the natural need of some system certifications, while f_i is the additional information for each task that we need to derive from WCET estimations to achieve such probabilistic certification levels.

A. On the WCET Dependencies

In our model, the failure probability per hour of each task f_i represents the probability of *any* job of the task τ_i exceeding its LO-WCET. Thus dependences between tasks and task executions could have a strong impact on f_i . We hereby detail how we intend to cope with statistical dependence.

In [9] it has been shown that *neither* probabilistic dependence among random variables *nor* statistical dependence of data implies the loss of independence between tasks’ pWCETs or WCET estimates. The WCET is an upper-bound to any execution time, which makes the important consequence on the independence between WCETs: jobs and tasks modeled with WCETs are independent because WCETs already embed dependence effects. Although both execution bounds (LO-WCET, HI-WCET) are so far called worst-case execution time estimations, the LO-WCET may also serve as an execution time upper-bound, where dependence between tasks and within tasks needs to be more carefully accounted for (see [10] for the original definition of MC task model).

Each MC task may generate an unbounded number of jobs. Since jobs generated from the same task set typically represent

execution of the same piece of code, and consecutive such jobs could experience similar circumstances, in the definition of the failure probability f_i (of a task τ_i), we naturally assume dependence among jobs of the *same* task; i.e., it represents the likelihood that the required execution time of *any* job generated within an hour by τ_i will exceed $c_i(\text{LO})$. In [11], [12] it has been shown that real safety-critical embedded systems have natural variability on the task execution time, thus it is reasonable to assume independence or extremal independence between jobs.

Concerning task dependencies, we can cope with the dependence by specifying the task pairwise dependence model. Assuming we are given a list of pairs (τ_i, τ_j) indicating that (WC)ETs of these two tasks may be dependent on each other. It means that the probability of them both exceeding their LO-WCET is no longer the product of their individual probabilities. By knowing $P(C_i > c_i(\text{LO}), C_j > c_j(\text{LO}))$ we are able to model (τ_i, τ_j) dependence including execution time task dependencies in our framework, Section IV-A. For many real-world systems, it is reasonable to assume that many (or most) task pairs do not have such dependencies to each other (although at the execution time level), since the limited impact of one task to another in a mixed-critical partitioned system. In Sec. IV and V, we consider inter-task independence (no pairwise failure-dependencies) to schedule the MC task set on uniprocessor and multiprocessor platforms respectively. Furthermore, in Sec. VI we introduce a covariance matrix to represent the pairwise failure-dependencies and present the corresponding scheduling technique.

To summarize, intra-task as well as inter-task job dependences are covered by our model.

B. Utilization Costs

The notion of **additional utilization cost**, defined below, helps quantify the capacity that must be provisioned under HI-criticality mode.

Definition 3 (additional utilization cost). *The additional utilization cost of HI-criticality task τ_i is given by*

$$\delta_i = (c_i(\text{HI}) - c_i(\text{LO}))/T_i. \quad (3)$$

Since we consider EDF schedulability instead of fixed-priority, we would like to know whether, and *how likely* system utilization may exceed 1: (i) if it is extremely unlikely that the total HI-criticality utilization exceeds 1 (weakly probabilistic schedulable), we could assert a system that is infeasible in traditional MC model to be probabilistic feasible; (ii) if it is extremely unlikely that total system utilization exceeds 1 (strongly probabilistic schedulable), we could decide not to drop any LO-criticality task even if some HI-criticality tasks accidentally suffer from failures (that they require more execution time than expected).

Example I.1 has shown an infeasible task set (under traditional MC scheduling) being weakly probabilistic schedulable under our model. As seen from the definitions, existing mixed-criticality systems are often analyzed under two modes – the HI mode and the LO mode, and mode switch is triggered when any HI-criticality job exceeds its LO-WCET without signaling

finishing. Upon such a mode switch, deadlines of all LO-criticality jobs will no longer be guaranteed. A natural question arises – is such sacrifice (dropping *all* LO-criticality jobs) necessary whenever a HI-criticality job requires execution for more than its LO-WCET? The following example illustrates the potential benefits in terms of enhanced schedulability of the proposed probabilistic MC model.

Example III.1. Consider a system composed of the three independent MC tasks that $\tau_1 = \{[2, 3], 0.1, 5, \text{HI}\}$, $\tau_2 = \{[3, 4], 0.05, 10, \text{HI}\}$, and $\tau_3 = \{[1, 1], 10, \text{LO}\}$, to be scheduled on a preemptive uniprocessor, with desired system failure probability threshold of $F_S = 0.01$.

Since HI-utilization of the system is $u_{\text{HI}} = 3/5 + 4/10 = 1$, any deterministic MC scheduling algorithm will prioritize τ_1 and τ_2 over the LO-criticality task τ_3 , and drop τ_3 if any HI-criticality job exceeds its LO-WCET.

With the additional probability information provided in our richer model, however, more sophisticated scheduling and analysis can be done. Recall from the definition of f_i , τ_1 has a probability of no larger than 0.1 to exceed a 2-unit execution within an hour, while the probability of any job in τ_2 exceeding a 3-unit execution within an hour is 0.05. Under the task-level independence assumption, the probability of jobs from both HI-criticality tasks requiring more than their LO-WCETs in an hour ($P(x_1 = x_2 = 1) = P(x_1 = 1) \times P(x_2 = 1) = 0.1 \times 0.05 = 0.005$) is smaller than F_S ⁶.

Hence, in the schedulability test, we do not need to consider the case that both HI-criticality tasks exceed their LO-WCETs simultaneously. Moreover, either one of them exceeding its LO-WCET will not result in an over-utilized system – a HI-criticality “server” with budget 0.2 and period 1 can be added to provide the additional capacity (over and above the LO-WCET amount). This server will be scheduled and executed as a virtual task, and both HI-criticality tasks may run on the server. The additional budget of 0.2 is sufficient to handle either $\delta_1 = 0.2$ or $\delta_2 = 0.1$, which is necessary when one of the HI tasks exceed their LO-WCET. The server needs not provide $\delta_1 + \delta_2 = 0.3$ budget, since the probability of such an event is less than F_S .

The total system utilization thus provisioned for the HI-criticality tasks is $2/5 + 3/10 + 0.2/1 = 0.9$; upon provisioning an additional utilization of $1/10 = 0.1$ for the LO-criticality task τ_3 , the total utilization becomes 1. Thus under any optimal uniprocessor scheduling strategy, e.g., EDF, the failure (any deadline miss) rate of the system in any hour will be no greater than F_S , and the MC instance is strongly probabilistic schedulable under this scheduling strategy (EDF plus the HI-criticality server) for the specified threshold F_S . \square

IV. SCHEDULING STRATEGY

A. The LFF-Clustering Algorithm

In this subsection, we present our strategy for scheduling independent preemptive MC task instances, by combining

⁶In general, we cannot simply ignore an event when its failure probability is below F_S . Instead, we do not need to consider a set of events only when the sum of their failure probability is below F_S . More details on this can be found in Section III.

HI-criticality tasks into clusters intelligently, and provide a sufficient schedulability test for it. Consider what we have done in Example III.1 above. We essentially: (i) conceptually combined the HI-criticality tasks τ_1 and τ_2 into a single cluster, provisioning an additional server into the system to accommodate their possible occasional HI-mode behaviors (execution beyond their LO-WCETs); and (ii) performed two EDF schedulability tests: one considering only HI-criticality tasks (with LO-WCETs) and the server, and the other also considering the LO-criticality task (τ_3). Since both tests succeed, we declare *strongly probabilistic schedulable* for the given instance; we would have declared *weakly probabilistic schedulable* if the second schedulability test had failed while the first one succeeded.

The technique that was illustrated in Example III.1 forms the basis of the scheduling strategy that we derive in this section. To obtain a good upper bound to HI-criticality utilization of the system, we combine tasks into *clusters* – suppose that the n_{HI} HI-criticality tasks have been partitioned into M clusters G_1, G_2, \dots, G_M , and let $y_i \in \{1, 2, \dots, M\}$ denote to which cluster (number) task τ_i is assigned.

Definition 4. (Failure probability of a cluster) Failure of a cluster G_m is defined as jobs generated by more than one tasks in a single cluster exceeding their LO-WCETs within an hour. The probability of a failure occurring in cluster m is denoted as g_m and is given by

$$g_m \stackrel{\text{def}}{=} 1 - \prod_{i|y_i=m} (1 - f_i) - \sum_{j|y_j=m} f_j \frac{\prod_{i|y_i=m} (1 - f_i)}{1 - f_j}, \quad (4)$$

where the second term of right hand side is the probability of no task (in the cluster) exceeding its LO-WCET, and the last term represents the probability of exactly one of the tasks exceeding its LO-WCET in an hour.

Lemma 1. If $g_m < F_S/M$ holds for every cluster G_m , then the probability of having no failure in all clusters is greater than $(1 - F_S)$.

Proof. Since clusters do not overlap with each other (each HI-criticality task belongs to a single cluster) and thus are independent of each other, the probability of having no failure in all clusters is given by the product of each cluster being failure-free, which is: $\prod_{m=1}^M (1 - g_m) > \prod_{m=1}^M (1 - F_S/M) = (1 - F_S/M)^M \geq 1 - F_S$ (From Binomial Theorem). \square

Lemma 1 provides a safe failure threshold F_S/M for each cluster; i.e., the rule for forming *clusters* is $g_m < F_S/M$, where M is the current number of clusters.

The additional utilization cost of a cluster G_m is defined to be equal to the additional utilization cost (δ_i) of the task within the cluster with the largest δ_i value; i.e.,

$$\Delta_m \stackrel{\text{def}}{=} \max_{i|\tau_i \in G_m} \delta_i. \quad (5)$$

The total system additional utilization cost is given by the sum of additional utilization cost of all M clusters;

$$\Delta \stackrel{\text{def}}{=} \sum_{m=1}^M \Delta_m. \quad (6)$$

A critical observation is that, if a task τ_i with additional utilization cost δ_i has been assigned to a cluster, assigning any other task τ_j with $\delta_j \leq \delta_i$ to the cluster will not increase the additional utilization cost. To minimize the total additional utilization cost of the entire task set, we therefore greedily expand existing clusters with tasks of larger additional utilization cost while ensuring that the relationship $g_m < F_S/M$ continues to hold, which leads to the Largest Fit First (LFF)-Clustering algorithm.

Algorithm 1: Algorithm LFF-Clustering

Input: $F_S, \{f_i\}_{i=1}^{n_{\text{HI}}}, \{\delta_i\}_{i=1}^{n_{\text{HI}}}$
Output: maximum total additional utilization cost Δ

begin

- Sort the tasks in non-increasing order of δ_i ;
- $m \leftarrow 1, M \leftarrow n_{\text{HI}}, y_i \leftarrow 0$ for $i = 1, \dots, n$;
- while** $\prod_{i=1}^{n_{\text{HI}}} y_i = 0$ (*an unassigned task exists*) **do**

 - $\Delta_m \leftarrow 0$ (*additional utilization of each cluster*);
 - for** $i \leftarrow 1$ to n_{HI} **do**

 - if** $y_i > 0$: **continue**;
 - $y_i \leftarrow m, M \leftarrow M - 1$;
 - if** $g_m \geq F_S/(M+m)$: $y_i \leftarrow 0, M \leftarrow M + 1$;
 - end**
 - $\Delta_m \leftarrow \max_{i|y_i=m} \delta_i; m \leftarrow m + 1, M \leftarrow M + 1$;

 - end**
 - return** $\sum_{m=1}^M \Delta_m$;

end

This algorithm greedily expands each existing cluster with unassigned tasks while the condition $g_m < F_S/M$ holds; while a new cluster is created only if it is not possible to assign a task to any current cluster without violating the condition ($g_m < F_S/M$).

Remark 1. Similar to what has been done in [13] and [14], we may achieve a precise distribution to the total utilization of all tasks by applying the *convolution* operation ‘ \otimes ’, which results in an exponential ($O(2^{n_{\text{HI}}})$, to be precise) running time (see Appendix B). The sufficient schedulability test based on the LFF-Clustering algorithm runs in $O(n_{\text{HI}}^2)$ time, where n_{HI} is the number of HI-criticality tasks.

Remark 2. In the case that *all tasks share the same f_i value*, the schedulability test based on LFF-Clustering becomes *necessary and sufficient*.

Run-Time Strategy. During execution, a HI-criticality *server* τ_s with utilization Δ and a period of 1 tick is added to the task system. The server is represented as $\tau_s = \{\Delta, 1, \text{HI}\}$.

We need the server period as 1 tick because the mechanism and the analysis will not work if there is release or deadline within a server period. At any time instant that the server is executing, if there are active⁷ HI-criticality jobs, they are executed following earliest deadline first policy; if not, then

⁷A job is *active* if it is released and incomplete at that time instant.

the current server job is dropped⁸. All jobs including the server are scheduled and executed in EDF order, and a job is dropped at its deadline if it is not completed by then.

Although we introduce a server task with period of 1, preemption does not necessarily happen that often. The goal of the sever task with utilization Δ is to preserve a “bandwidth” of at least Δ for HI-criticality jobs if the HI-criticality ready queue is not empty. There are three situations to be considered:

Situation 1: The job with the earliest deadline is a HI-criticality job. In this situation, we execute the HI-criticality job with 100% processor share, and no more preemption is incurred by the server.

Situation 2: The job with the earliest deadline is a LO-criticality job *and* the HI-criticality ready queue is empty. In this situation, we execute the LO-criticality job with 100% processor share, and hence there is no additional preemption in this situation either.

Situation 3: The job with the earliest deadline is a LO-criticality job *and* the HI-criticality ready queue is *not* empty. In this situation, we want to preserve a processor share of Δ for HI-criticality jobs and to execute the LO-criticality ones with the rest $1 - \Delta$ of the processor capacity. Therefore, the server creates preemptions every time unit.

That is, only in Situation 3, our algorithm “introduces” extra preemptions due to the server scheme, and normal EDF scheduling is applied in other cases. One may claim that such server allocation scheme may results in more preemptions than the approaches where the server capacity is only used for overruns. Actually this is because that the goal here is trying *not to drop* LO-criticality tasks even when a few HI-criticality ones exceed their LO-WCETs. Thus, in order to guarantee HI-deadline being always met, we have to make certain use of the server even when no HI-criticality behavior is detected – simply taking “precautions”. Alternative way such as assigning HI-criticality jobs virtual deadlines may lead to fewer preemptions, at a cost of losing the performance of schedulability ratio (see Sec. VII).

B. Schedulability Test

It is evident that for *strongly probabilistic schedulable* (i.e., to ensure that the probability of missing *any* deadline is no larger than the specified system failure probability F_S – see Definition 2), it is (necessary and) sufficient that $(\sum_{i=1}^n c_i(\text{LO})/T_i + \Delta)$ must be no larger than the capacity of the processor (which is 1).

For *weakly probabilistic schedulable* (i.e., to ensure that the probability of missing any HI-criticality deadline is no larger than F_S – again, see Definition 2), it is necessary that $(\sum_{i|_{\chi_i=\text{HI}}} c_i(\text{LO})/T_i + \Delta)$ must be no larger than 1 as well. The following theorem helps establish a sufficient condition for ensuring weakly probabilistic schedulable:

Theorem 1. If no job exceeds its LO-WCET, then no deadline is missed if

⁸Since an integer model of time is assumed (i.e., all task periods are integers and all job arrivals occur at integer instants in time), and the server has a period of 1, it is safe to drop the current job of the server if there are no active HI-criticality jobs since there can be no HI-criticality job releases in the current period of the server.

$$\Delta \cdot \left(1 - \sum_{i|\chi_i=\text{HI}} \frac{c_i(\text{LO})}{T_i}\right) + \sum_{i=1}^n \frac{c_i(\text{LO})}{T_i} \leq 1. \quad (7)$$

Proof. Please refer to the conference version [1]. \square

Theorem 1 yields the schedulability test pMC (Algorithm 2), while Theorem 2 below establishes its correctness.

Algorithm 2: Schedulability Test pMC

```

Input:  $\tau, F_S$ 
Output: schedulability

begin
    Calculate  $\delta_i$  values for all HI-criticality tasks in  $\tau$ ;
     $u_{\text{LO}} \leftarrow \sum_{i=1}^n c_i(\text{LO})/T_i$ ;
     $u'_{\text{LO}} \leftarrow \sum_{i|\chi_i=\text{HI}} c_i(\text{LO})/T_i$ ;
     $\Delta \leftarrow \text{LFF-Clustering}(F_S, \{f_i\}_{i=1}^{n_{\text{HI}}}, \{\delta_i\}_{i=1}^{n_{\text{HI}}})$ ;
    if  $u_{\text{LO}} + \Delta \leq 1$  then
        return strongly probabilistic schedulable;
    else if  $u'_{\text{LO}} + \Delta \leq 1, \Delta \cdot (1 - u'_{\text{LO}}) + u_{\text{LO}} \leq 1$  then
        return weakly probabilistic schedulable;
    return unknown;
end

```

Theorem 2. The schedulability test pMC is *sufficient* in the following sense: If it returns *strongly probabilistic schedulable*, the probability of any task missing its deadline is no greater than F_S ; and if it returns *weakly probabilistic schedulable*, the probability of any HI-criticality task missing its deadline is no greater than F_S , and no deadline is missed when all jobs finish upon execution of their LO-WCETs.

Proof. Please refer to the conference version [1]. \square

The schedulability test pMC returns *strongly probabilistic schedulable* if we are able to schedule the system such that the probability of missing any deadline is at most the specified threshold F_S , or *weakly probabilistic schedulable* if we are able to schedule the system such that the probability of missing any HI-criticality deadline is at most F_S . We will then use EDF to schedule and execute the task set with LO-WCETs and the additional server task $\tau_s = \{\Delta, 1, \text{HI}\}$.

In the case that the schedulability test pMC returns *unknown*, we are not able to schedule the system using the proposed probabilistic analysis technique. Normally it is because either the safety requirement of the system is too high; (i.e., the threshold F_S is too small), or the WCET estimations are not precise enough for HI-criticality tasks; (i.e., the f_i values are not small enough compared to F_S (and n_{HI}), and/or the $c_i(\text{LO})$ values are not differentiable enough against $c_i(\text{HI})$ values).

V. THE MULTI-PROCESSOR CASE

Multi-processor devices are becoming more and more popular, while it is becoming more efficient to schedule real-time tasks in multi-processor platforms to achieve better throughput. Considering the pragmatic application of applying probabilistic scheduling, implementing a similar technique in multi-processors will dissipate the pessimistic assumption of existing scheduling mechanism and improve the resource

efficiency. This section proposes a multi-processor scheduling technique of MC tasks considering the failure probability based on an partitioned-based approach.

The partitioned-based scheduling of implicit-deadline sporadic task system can be converted into a bin-packing problem [15]. Hence, each processor is modeled as a bin of capacity one, and each task τ_i has the capacity of size u_i (its utilization). As bin packing is NP-Hard [15], heuristics can be applied for solving the problem. As our system model considers MC tasks with failure probability, we need to modify the task sets to fit into a traditional bin packing heuristics. Here we briefly discuss three most-common heuristics (First-Fit, Best-Fit, and Worst-Fit) and then present an algorithm to schedule our system model in multi-processor environments.

Different Allocation Heuristics. For partitioned scheduling, most common heuristics are First-Fit (FF), Best-Fit (BF), and Worst-Fit (WF). Note that Reasonable Allocation Decreasing (RAD) algorithms are proven to provide optimal utilization bound [16]. Hence we adapt such approach and thus the three algorithms discussed above would have been called First-Fit Decreasing (FFD), Best-Fit Decreasing (BFD), and Worst-Fit Decreasing (WFD) partitioned algorithms.

Task Allocation. To schedule tasks in our proposed system model in multiprocessor platforms, we present a slightly different scheduling approach than the one for uniprocessor in Sec. IV. Initially, all the tasks are grouped into different clusters using the same LFF-Clustering algorithm presented before. Then we schedule the clusters on different processors by using different RAD partitioned heuristics. Note that, for every cluster, there is an additional utilization cost Δ_m which is also needed to be allocated in case any task of the cluster exceeds its LO-criticality WCET. For every processor, we need to allocate a server which has a utilization equal to the sum of additional utilization cost Δ_m of all the clusters allocated in that processor. For example, while considering scheduling three clusters on two processors, if the Clusters 1 and 2 are allocated on Processor 1 and the Cluster 3 is allocated on Processor 2, then we also need to allocate a server with utilization $(\Delta_1 + \Delta_2)$ to Processor 1, and another server with utilization Δ_3 to Processor 2. In short, while applying the partitioned heuristic, we need to accommodate the demand of server utilization as well. Specifically, the LO-criticality tasks are considered for allocation once all the HI-criticality clusters and their Δ s are allocated, using the same partitioning techniques used for the HI tasks.

A. Schedulability Analysis

The schedulability of the task set is determined in two different steps. First, we have to check whether the tasks can be properly allocated to the available processors. Upon successful allocation, we need to check whether the correctness of each uni-processor MC task system can be guaranteed in runtime even under worse conditions. In each step, there is a possibility that either only HI tasks or all the tasks are allocated/schedulable. Based on the allocation, the task sets can be strongly allocated or weakly allocated. If all the HI-task clusters, Δ s, and LO-tasks are allocated then we call it strongly allocated task set. If only the clusters and Δ s are allocated

properly but not the LO-tasks then we call it weakly allocated task set. If neither, then the task sets cannot be considered for further schedulability test. Upon successful allocation, we need to check the schedulability of the allocated task set on all the processors. The schedulability test is done by following the similar technique used for the uniprocessor scheduling.

Before going into the schedulability conditions, it is necessary to introduce the parameter α and β . α is the utilization factor of a task set, i.e., the maximum utilization among all tasks. β is the maximum number of tasks of α which fit into one processor under EDF scheduling. β can be expressed as a function of α

$$\beta = \lfloor 1/\alpha \rfloor \quad (8)$$

Lopez et al. [16] proved that for multiprocessor partitioned scheduling using EDF, FFD, BFD, and WFD algorithms provide the optimal upper bound, which is the following:

Lemma 2. [16] *Let $U(N, \alpha)$ denote the utilization bound to schedule N tasks using RAD algorithms on m processors and α represent the utilization factor for the task set, then $U(N, \alpha) = \frac{\beta N + 1}{\beta + 1}$ when $m > \beta N$.*

Task Allocation Conditions. For the HI-tasks we need to allocate each cluster with its Δ_m in the same processor as in the partitioned scheduling, a task is always needed to be executed on the same processor which it was initially allocated. Lets assume there are M clusters with utilization UC_1, UC_2, \dots, UC_M with the corresponding delta values as $\Delta_1, \Delta_2, \dots, \Delta_M$. As each cluster is to be allocated to a processor with its corresponding Δ , let's define HI mode utilization factor α_h and overall utilization factor α_s below:

$$\alpha_h = \max_{i \in 1, 2, \dots, M} (UC_i + \Delta_i); \quad \alpha_s = \max(\alpha_h, \alpha_l); \quad (9)$$

where α_l is the utilization factor for the LO-tasks.

Furthermore, corresponding β values can be defined:

$$\beta_h = \lfloor 1/\alpha_h \rfloor; \quad \beta_s = \lfloor 1/\alpha_s \rfloor \quad (10)$$

Theorem 3. *All the HI-criticality clusters along with their server allocation (Δ) and all the LO-criticality tasks can be allocated (i.e., strongly allocated task-set) to K processors if the following two condition holds,*

$$K > \beta_s(M + n); \quad U_s \leq \frac{\beta_s(M + n) + 1}{\beta_s + 1}; \quad (11)$$

where U_s is the sum of the utilization of all clusters, their Δ_s and LO-criticality tasks.

Proof. Here, each cluster and LO-criticality tasks can be seen as a single entity with specific utilization demand. The total number of entity here is $(M + n)$. Thus according to Lemma 2, the maximum utilization bound can be $\frac{\beta_s(M + n) + 1}{\beta_s + 1}$. So for a successful allocation, the system utilization U_s must be no greater than the utilization bound. \square

Theorem 4. *All the HI-criticality clusters along with their server allocation (Δ) can be allocated (i.e., weakly allocated*

task-set) to K processor if the following two condition holds,

$$K > \beta_h M; \quad U = \frac{\beta_h M + 1}{\beta_h + 1}; \quad (12)$$

Proof. The proof is very similar to Theorem 3 (omitted). \square

Partition Approach and Its Correctness. Upon successful allocation (either strongly or weakly allocated), the schedulability of the task set can be determined by running the pMCMP (probabilistic Mixed-Criticality on MultiProcessor) algorithm (presented in Algorithm 3). pMCMP algorithm basically use pMC algorithm presented in Sec. IV on all the processors to check the schedulability. The result of pMCMP can be strongly schedulable, weakly schedulable, or non-determined. If a task set is only weakly allocated on the available processors, the task set can never be considered as strongly schedulable.

Algorithm 3: pMCMP algorithm

Data: Allocation of HI-criticality Δ_i s and LO-criticality task-set τ_i on each processor κ_i
Result: The schedulability of the task set
if $\forall \kappa_i, pMC$ returns strongly-schedulable **then**
| return strongly-schedulable;
else if $\forall \kappa_i, pMC$ returns weakly-schedulable **then**
| return weakly-schedulable;
else
| return non-determined;
end

Theorem 5. *Algorithm pMCMP is correct. I.e., if pMCMP returns strongly probabilistic schedulable, the probability of any task missing its deadline is no greater than F_S ; while if pMCMP returns weakly probabilistic schedulable, the probability of any HI-criticality task missing its deadline is no greater than F_S , and no deadline is missed when all jobs finish upon execution of their LO-WCETs.*

Proof. After a successful allocation, each processor has a specific set of Δ 's and LO-criticality tasks assigned for scheduling. Let $\overline{UC}_1, \overline{UC}_2, \dots, \overline{UC}_M$ denote the number of clusters assigned to K processors (note that there are total M number of clusters). Hence, each processor assignment can be seen as a subset problem of uniprocessor scheduling presented in [1].

For an arbitrary processor i , similar to the proof of Eqn. 5.2 [1], the failure probability of processor i is no greater than $(\overline{UC}_i \times F_S)/M$. As the tasks are independent, the total failure probability of the system is no greater than $\sum_{i=1}^M (\overline{UC}_i \times F_S)/M = F_S$. \square

VI. CONSIDERING FAILURE DEPENDENCY

In previous sections, we have considered only independent tasks, i.e., the failure of a task τ_i is independent to whether another task fails or not. In other words, whether a HI task fails to complete within its LO-WCET budget does not affect the completion of other HI tasks. On the contrary, most existing work in MC scheduling assumed that all HI tasks may exceed their LO-WCET budgets at the same time. That means failure probabilities of all HI tasks are dependent on

each other, which is rather a pessimistic assumption in certain scenarios. However, in practical systems, not every failure probability is independent. It is possible that while most of the HI tasks are independent, a certain amount of task is directly dependent to each other with respect to their probability of exceeding LO-WCET, i.e., once a HI task exceeds the LO-WCET, other dependent tasks may also exceed their LO-WCET budgets, *regardless of their own failure probability*. Techniques such as fault tree analysis [17], [18] can be used to analyse tasks for failure dependency. With the dependency information provided as input, we propose a covariance matrix representation and propose a graph-coloring based partitioning approach to handle such a scenario.

A. Covariance Matrix

We introduce a covariance matrix to represent the dependencies between each pair of tasks regarding their potential failures to complete before LO-WCETs. In Table I, a sample covariance matrix is shown for eight HI-criticality tasks. Here the covariance matrix is binary, where for each item, 0 represents the independence between two tasks while 1 represents that there is dependency on failure probability between two tasks. Note that, 1 does not mean fully dependant, while is a safe (although pessimistic) measurement in terms of schedulability guarantees. As a result, if there is a dependency between two tasks, we assume that upon exceeding the LO-WCET budget of one task, the other dependent task/s will also exceed their LO-WCET budget simultaneously.

TABLE I: Covariance matrix of a set of eight tasks.

	τ_2	τ_3	τ_4	τ_5	τ_6	τ_7	τ_8
τ_1	0	0	0	1	0	0	0
τ_2	-	1	0	0	1	0	0
τ_3	-	-	0	0	1	1	0
τ_4	-	-	-	0	0	0	0
τ_5	-	-	-	-	0	0	1
τ_6	-	-	-	-	-	0	0
τ_7	-	-	-	-	-	-	0

Remark 2. Failure dependency is not statistical dependence.

Definition 5. (Failure-dependent tasks) *When scheduling a pair of tasks on the same processor, if one of those tasks exceeds its LO-WCET and, the other one's failure likelihood becomes larger than the given value, or the other way around, then two tasks are considered failure-dependent.*

In other words, it is safe to assume that both of the tasks exceed their LO-WCET simultaneously. Take the task set shown in Table I as an example, τ_2 and τ_3 are failure dependent but τ_1 and τ_2 are not. Note that two tasks are failure-dependent only if those are scheduled on the same processor. Upon scheduling on different processors, those tasks can be executed independently since we conduct partitioned scheduling and assume isolation between processors.

B. Task Isolation using Graph Model

Similar to the clustering problem without covariance, finding the optimal clustering for the new problem is also NP-Hard. Thus, we propose a heuristic to find an efficient solution

to the problem. From the covariance matrix, we can visualize a task set as a graph problem. We can assume the tasks as a node of the graph and the failure dependencies as edges between the nodes, i.e., if there is a 1 between two tasks, we can consider an edge between those two nodes (tasks). Note that, if the graph is fully connected, the problem will become a traditional MC scheduling problem as in our strategy, we will need to create a cluster for each task. In practical scenarios, the graph is assumed to be a disconnected graph as there can be both independent and failure dependent tasks in a system. Hence there will be multiple islands in the graph which consist of interdependent tasks.

Definition 6. (Transitive failure-dependency.) *If two tasks are not directly dependent but have a common failure-dependent task, then we call the failure-dependency between the first two tasks transitive failure-dependency.*

For example, τ_1 and τ_8 in Table I are not directly failure-dependent, but they both are dependent with τ_5 . Hence, if we partition these three tasks into a same processor, they may exceed their LO-WCET budgets simultaneously. However, if we schedule τ_5 in a separate processor, τ_1 and τ_8 will have no failure-dependency and can act as independent task pairs.

Upon transforming the covariance matrix into a graph, we apply our clustering heuristics. In the aforementioned LFF-Clustering algorithm we sorted all the tasks in descending order based on their additional utilization δ_i value and we keep adding the tasks one by one to clusters until the Lemma 1 is not violated. While scheduling the task set including failure-dependent tasks, we isolate the tasks of each island into m (number of processors) number of groups in a way such that there are no failure-dependent tasks in any group. By doing that, we ensure that no two inter-dependent tasks are grouped into a single processor.

C. Isolating Tasks of Each Island

Once we convert the task set into a graph with different islands, m number of groups of independent tasks are created. One can visualize this problem as a m -coloring graph problem, where the nodes of a graph are colored with at most m number of colors with no two nodes of same color adjacent to each other (i.e., there is no edge between them). If we can properly color the graph with m or less number of colors then we can easily allocate nodes with the same color into the same processor. Unfortunately, the m -coloring problem is an NP-Complete problem [19]. As a result, we will need an approximate algorithm to solve such problems. Furthermore, it may not always be possible to color the subgraph in each island with m colors by using the approximate algorithm (even by using an optimal algorithm). Hence, in this subsection we propose a modified approximate algorithm to color each subgraph with m colors. To accomplish this goal, we adapt a modified greedy coloring algorithm. As our base greedy coloring algorithm, we use the Welsh Powell Algorithm [20] (also known as Largest-First (LF) coloring algorithm).

In LF coloring, node with the highest degree⁹ will be colored first, and then the adjacent nodes will be colored and

⁹degree of a node is the number of edges connected to the node.

corresponding edges are deleted. By doing this, we remove the dependencies of other nodes with the colored nodes. At any step, if the sum of the utilizations exceeds one, we stop coloring for that particular color and leave that node for the next coloring iteration. However, when we already use $m - 1$ colors and only one color is left to use, we need to contract (merge) the remaining failure-dependent tasks. When at the last processor, there are still some nodes which are connected (failure-dependent), we contract those nodes in a single node and use the u_i^{LO} and δ_i of the node as the sum of the corresponding values of all the connected nodes. The steps of the coloring technique are shown in Algorithm 4 and further demonstrated in Example VI.1.

Algorithm 4: m-coloring algorithm

```

Data:  $G = \{V, E\}$ ,  $m$ 
Result:  $m$  - colored graph
Sort all nodes  $V_i$ s in non-increasing order of their degrees;
for  $i \leftarrow 1$  to  $m$  do
    /* All nodes are colored */ 
    if  $\forall V_i$  is colored then
        | return;
    end
    /* If at last processor, there exists
       some connected nodes */ 
    if ( $i == m$ ) and ( $\exists E_i$ ) then
        forall Connected subgraph do
            |  $V_{\text{new}} \leftarrow \{\text{connected\_nodes}\}$ ;
            |  $V_{\text{new}}.u_i^{\text{LO}} \leftarrow \{\text{connected\_nodes}\}.u_i^{\text{LO}}$ ;
            |  $V_{\text{new}}.\delta_i \leftarrow \{\text{connected\_nodes}\}.\delta_i$ ;
        end
    else
        | Color the nodes following LF [20] Algorithm;
    end
end

```

Example VI.1. In Figure 1, the tasks of one island is shown. Here, we need to allocate the tasks on two processors, i.e., we have to color the graph with 2 colors. To do this, we first take the node with the highest degree (τ_4 with degree 4) and color it with green. Then we color the non-adjacent nodes of τ_4 (i.e., τ_1 , τ_2 , and τ_8 with green and remove the edges of those nodes. We can no longer use green in this graph. Now we have one processor (color) left but there are still two tasks τ_6 and τ_7 , which are failure-dependent to each other. So, we merge these two tasks and all the nodes become independent. Finally, we color all the nodes with blue (τ_3 , τ_5 , and τ_{6+7}).

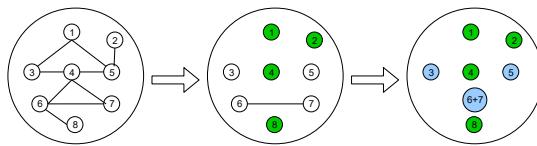


Fig. 1: Graph coloring with $m = 2$.

D. Task Allocation and Scheduling

Given a task set, we first run Depth-First-Search (DFS) over the covariance matrix and convert them to different number of islands. Then we use Algorithm 4 to color the nodes of

each island with at most m number of distinct colors. Then the nodes in each island are sorted in descending order with respect to their δ_i values and the islands themselves are then sorted in descending order based on $\max(\delta_i)$ values. HI-criticality tasks are allocated by following rules:

(i) Assign tasks of each color of each island to a distinct processor with the WF heuristics on processor capacity.

(ii) While assigning a task to a processor, we create a cluster following the LFF-Clustering algorithm. First, we keep adding the same colored tasks in an island to the existing task/s assigned to that processor. If we cannot assign the new task to an existing cluster, we create a new cluster. Once all nodes of the same color are allocated, we allocate the next color tasks to a different processor.

(iii) Every time we add a new task to a processor, we update the Δ value of the processor.

(iv) Once all islands with multiple tasks are assigned, we assign the remaining single-task islands by following the WF partitioning heuristic.

If all the tasks are allocated, the task set becomes strongly allocated, while only the allocation of HI-criticality tasks results in a weakly allocation. Similar to the multiprocessor case with no failure dependent tasks, successful allocation does not imply schedulability of the allocated task sets. If a task set is not strongly allocated, then the task set cannot be strongly schedulable, because of the lack of guarantees to the LO-tasks. Upon successful allocation, either strong or weak, we run the pMCMP algorithm to check the schedulability of the task set. Algorithm 5 details the task allocation procedure.

Algorithm 5: Task Allocation Algorithm with Covariance

```

Data:  $F_S$ ,  $\{f_i\}_{i=1}^{n_{\text{HI}}}$ ,  $\{u_i^{\text{LO}}\}_{i=1}^n$ ,  $\{\delta_i\}_{i=1}^n$ , covariance matrix
Result: Task allocation result
Run DFS on covariance matrix and get islands  $\{I_j\}_{j=1}^l$ ;
Color the nodes of each island using Algorithm 4;
Sort  $\forall \tau_i \in \forall I_j$  in descending order w.r.to  $\delta_i$ ;
Sort  $\forall I_i$  based on  $\max(\delta_i) \in I_i$ ;
 $\mathcal{K} = \{\kappa_1, \kappa_2, \dots, \kappa_K\}$ ;
forall multi-task islands  $I_i$  do
    allocated =  $\phi$ ;
    forall  $\forall \tau_j \in I_i$  do
        | allocate  $\tau_j$  into  $\kappa_k \in \{\mathcal{K} - \text{allocated}\}$  following WF
        | and update  $\Delta_k$ ;
        | allocated = allocated  $\cup \kappa_k$ 
    end
end
if all LO tasks are allocated using WF then
    | return strongly - allocated;
end
return weakly - allocated;

```

VII. SCHEDULABILITY EXPERIMENTS

In this section, we present extensive experimental evaluations to show the performance of our proposed algorithms. We present our results in two different categories. First, we present the schedulability result for uniprocessor platforms to show the efficiency of our algorithm. We performed simulation for different constraints and also compared with other state-of-the-art MC scheduler. For the next part, we present the

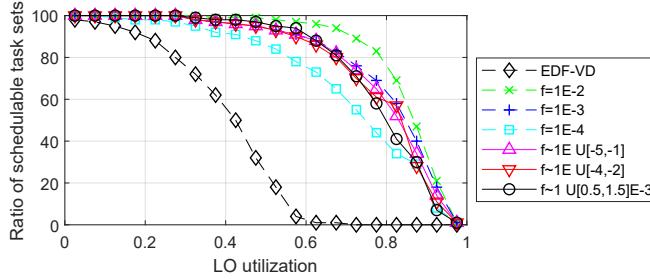


Fig. 2: Schedulability ratio comparison of EDF-VD and pMC, where HI utilization varies from 0.9 to 1 uniformly.

schedulability performance on multiprocessor platforms. We have performed a number of experiments by varying different important factors to observe the efficiency of our algorithm.

A. Results for UniProcessor Platforms

We have conducted schedulability tests on randomly-generated task systems, comparing our proposed method with existing one. The objective was to demonstrate the benefits of our model: by adding a probability estimation f_i to each task, our algorithm may successfully schedule (return *probabilistically correct* or *partial probabilistically correct*) many task sets that are unschedulable according to existing MC-scheduling algorithms; e.g., the EDF-VD algorithm [21].

Since this is the first work that combines pWCET and schedulability with mixed-criticality, it is hard to find a fair base line to compare with. The reason EDF-VD is selected here since (i) it is a widely accepted MC scheduling strategy; (ii) it is the most general algorithm in the whole VD family; and (iii) HI-criticality tasks are treated as a whole in both algorithms – EDF-VD sets virtual deadline according to a common factor, while we make use of a HI-criticality server. We need to point out that EDF-VD assumes unknown f_i for each task (not simply 0 or 1), and thus our algorithm has privilege naturally.

We use the algorithm *UUniFast* [22] to generate task sets for various values of cumulative LO utilization ($u(\text{LO}) = \sum_{i=1}^n c_i(\text{LO})/T_i$) and HI utilization ($u(\text{HI}) = \sum_{i|\chi_i=\text{HI}} c_i(\text{HI})/T_i$). The parameter $u(\text{LO})$ is ranged from 0 to 1, while $u(\text{HI})$ is ranged from 0 to 1.5, each with step 0.01.

Each task set contains 20 tasks, each of which is assigned LO or HI criticality with equal probability. LO-criticality utilizations are assigned according to *UUniFast*; given an expected HI utilization $u(\text{HI})$, we inflate the LO-criticality utilizations of the HI-criticality tasks using random factors chosen to ensure that the cumulative HI utilization of the task set equals the desired value with high probability.

Among the 626,200 valid task sets that we generated, EDF-VD succeeds to schedule 306,299 (48.9%) of them, and the proposed pMC reports probabilistic schedulable for a total of 438787 sets (70.1%), and only 121,426 sets (19.4%) are reported unknown. Even when focused only upon systems for which HI-criticality utilization is less than 1, EDF-VD fails to schedule 18.0%, while pMC returns unknown for only 8.4% of the sets. Although EDF-VD and pMC do not dominate each other, pMC generally significantly outperforms EDF-VD, particularly upon task-sets with large HI-utilization. Due to

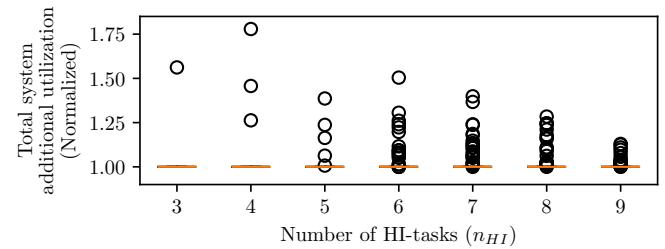


Fig. 3: Comparison of additional utilization Δ calculated by LFF-clustering normalized by the precise value

space constraints, the detailed schedulability comparison with EDF-VD algorithm is omitted in this paper, but is available in the conference version of this work.[1]

To show the robustness of our algorithm with respect to different f_i distributions, we focus on task sets with HI utilization between 0.9 and 1. Fig. 2 reports the ratios of schedulable (i.e., weakly probabilistic schedulable) sets over different LO utilizations. With the additional probability information, the schedulable ratio is significantly improved for heavy tasks comparing to EDF-VD [21]. The introduced parameter f_i is assigned to tasks in different ways; i.e., all sharing a same value, following uniform or log-uniform distribution ($f_i = 10^x$, where x is uniformly chosen). Generally speaking, smaller average f leads to higher ratio of acceptance, and there is no significant difference between different distributions of f_i with the same average, which indicates that our algorithm is *robust* to different combinations of output measurement probabilities from probabilistic timing analysis tools.

The boxplot in Figure 3 shows the total additional utilization resulting from LFF-Clustering normalized by the precise value calculated using exhaustive search. Due to high complexity for calculating precise values, 100 task sets under each n_{HI} setting is considered. The mean (orange line) is close to 1 under each setting, meaning over 75% of the task sets were precisely clustered. However, the outliers get closer to 1 as n_{HI} increases—LFF-Clustering is particularly beneficial when there are large number of HI-tasks and calculating precise clustering becomes prohibitively expensive.

B. Results for Multiprocessor Platforms

Workload Generation. To conduct the experiments, we have generated MC tasks based on the following parameters.

- m : The number of processor cores.
- U_a : The average utilization for the task set. The average is calculated by averaging the LO and HI-criticality utilization of the task set.
- $P_{\text{HI}} = 0.5$: Probability of a task to be a HI-criticality one.
- $R = 4$: Denotes the maximum ratio of u_i^{HI} to u_i^{LO} . u_i^{HI} is generated uniformly from $[u_i^{\text{LO}}, R \times u_i^{\text{LO}}]$.
- F_S : The system-wide permitted failure probability, default set as 10^{-6} for our experiments.

We performed the simulation for average utilization ranging from $0.05m$ to $2m$ with increasing at step size $0.05m$. For every average utilization, we generate 100 task sets which consist of 20 tasks each. Note that, for most of the experiments, we have measured the performance with respect to

average utilization as we wanted to show the improved quality of service for generated MC task sets.

At first, for a specific average utilization, we use *UUniFast algorithm* [23] to generate a lognormal distribution of U_a for all the tasks in a task set. 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 $[u_i^{\text{LO}}, R \times u_i^{\text{LO}}]$.

Evaluation Results. We execute a set of MC tasks under our proposed algorithm by varying different parameters. Simulation results for various scenarios are presented in Figures 4 and 5. We perform the following simulations:

The schedulability performance of pMCMP is shown in Fig. 4. Fig. 4(a) shows the acceptance ratio (ratio of successfully scheduled task sets over total number of task sets) with respect to average utilization U_a , while the Fig. 4(b) shows the acceptance ratio with respect to u_i^{LO} . In both figures, we show the acceptance ratio for both strongly schedulable and weakly schedulable task sets. In the Fig. 4(b), we can see that a good number of task sets is schedulable when the average utilization of the task set is one or even higher as the average utilization is calculated based on both u_i^{LO} and u_i^{HI} but pMCMP doesn't need to allocate full HI-WCET budget. To understand the schedulability with respect to u_i^{LO} , we further performed the experiment presented in Fig. 4(b) where the u_i^{LO} is generated following the lognormal distribution using the *UUniFast algorithm* [23] algorithm.

Fig. 4(c) presents the acceptance ratio for all three heuristics (FF, BF, and WF) discussed in Section V. The results match the discussion that all RAD algorithms share the same utilization bound while task partitioning. We use only strongly schedulable task set to calculate the acceptance ratio from this experiment as only the strongly schedulable task sets provide the graceful degradation to LO-criticality tasks.

Fig. 5 first presents the percentage of successful strongly schedulable task set under different number of processors (Fig. 5(a)) and different number of f_i values (Fig. 4(b)). As expected, the performance of schedulability decreases with the increase of the number of the processors by following the performance of the partition heuristics. On the other hand, with the lower f_i values, we get better acceptance ratio as the algorithm can create more clusters and thus needs to allocate a smaller Δ . We also evaluated the effect of the density of the covariance matrix on the acceptance ratio, calculated as the density of the undirected graph interpreted from the covariance matrix. The number of edges of the covariance graph (1 in covariance matrix) are randomly generated based on the density of edge. The simulation result is presented in Fig. 5(c). Under lower densities, our algorithms performed surprisingly well. With the increase of density, the possibility of merging also increases and the acceptance ratio decreases.

VIII. RELATED WORKS

There is a gap between the current conservative deterministic analysis and the richer models which include the probabilistic information about the WCET estimates. Besides the resource under-utilization issue due to over-pessimism, LO-critical tasks do not receive guarantees in HI-critical modes.

Graceful degradation techniques and imprecise computation have been proposed to provide guarantees to the LO-critical tasks as well. Specifically, reduced utilization budget [24] to scale budgets in HI-mode as well as precise computation techniques such as providing asymptotic rate guarantees [25], guaranteed completion rate [26] and QoS guarantees for LO-critical tasks [27]. While these approaches address the under-provision of resources to LO-critical tasks, they do not leverage the information from the WCET estimates.

Real-Time Models with Probabilities. In order to formally describe the uncertainty of the WCET estimations and overcome the over-pessimism, many attempts in introducing probability to real-time system model and analysis have been made. Edgar and Burns [28] made a major step forward in introducing the concept of *probabilistic confidence* to the task and the system model. Their work targets the estimation of probabilistic WCETs (pWCETs) from test data for individual tasks, while providing a suitable lower bound for the overall confidence level of a system. Since then, on one hand much work has been done to provide better WCET estimations and a predicted probability of any execution exceeding such estimation alongside the usage of extreme value theory (EVT), e.g., [6] [29] [7]. In static probabilistic timing analysis, random replacement caches are applied to compute exact probabilistic WCETs, and probabilistic WCET estimations with preemptions, [30]. More recently, researchers have initiated some pWCET estimation studies [31] [32] in the presence of permanent faults and disabling of hardware elements. On the other hand, there is only one piece of work which proposes probabilistic execution time (pET) estimation [33] based upon a tree-based technique. The pET of a task describes the probability that the execution time of the job is equal to a given value, while the pWCET of a task describes the probability that the WCET of that task does not exceed a given value.

Schedulability with Probabilities. Based upon the estimated pWCET and pET parameters (often as distributions with multiple values and associated probabilities), studies aim to provide estimations that the probability of missing a deadline of the given system is small enough for safety requirements; e.g., of the same order of magnitude as other dependability estimations. Tia et al. [34] focus on unbalanced heavy loaded system (with maximum utilization larger than 1 and much smaller average utilization) and provide two methods for probabilistic schedulability guarantees. Lehoczky [35] proposes the first schedulability analysis of task systems with probabilistic execution times. This work is further extended to specific schedulers, such as earliest deadline first (EDF, [36]) in [37] and under fixed-priority policy in [38]. [13] provides a very general analysis for probabilistic systems with pWCET estimations for tasks. In addition to WCET estimations, statistical guarantees are performed upon the minimum inter-arrival time (MIT) estimation as well [39] [14]. Schedulability analysis based on pETs (instead of pWCETs) is also done in [40] for limited priority level case (quantized EDF), and in [41] an associated schedulability analysis on multiprocessors is presented. Statistical response-time analysis, e.g., [42], can be further done to real-time systems based upon those probabilistic schedulability analysis. Unfortunately, to the best of our

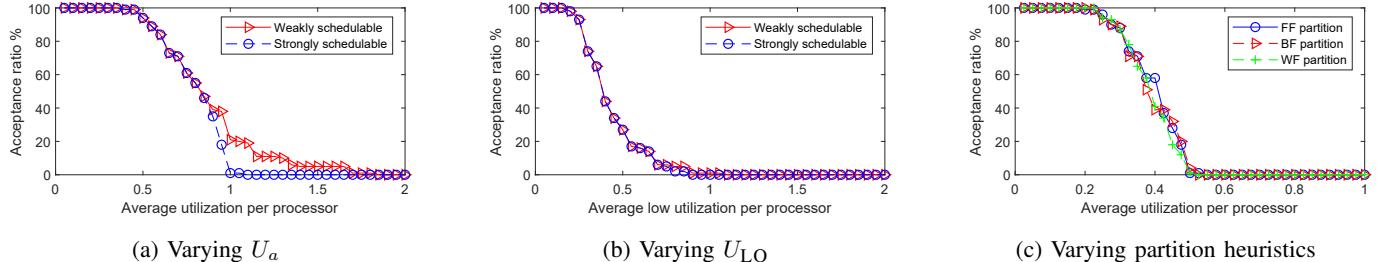


Fig. 4: Acceptance ratio for pMCMP in a 4-core platform under different utilizations and different partition heuristics.

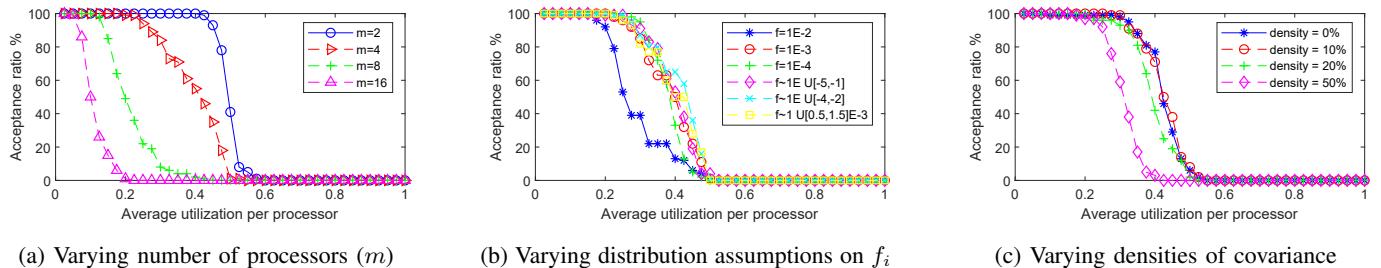


Fig. 5: Acceptance ratio for pMCMP upon a quad-core platform.

knowledge, most existing studies have only shown probabilistic schedulability analysis (e.g., estimating the likelihood for a system to miss any deadline) or probabilistic response time analysis to existing algorithms such as EDF and fixed-priority scheduling, instead of incorporating probabilistic information into the scheduling strategy. In other words, current research has not addressed the possibility of making *smarter scheduling decisions* with probabilistic models from existing powerful probabilistic timing analysis (PTA) tools (e.g., [43]) that provide WCET bounds and specified confidences¹⁰.

Mixed-Criticality Probabilistic Scheduling. Few recent works have applied probabilistic models to mixed-criticality scheduling. In [46], the probabilistic models are applied to LO-criticality modes and a scheduling algorithm is developed to leverage probabilities into schedulability analysis. In particular, it is quantified how LO-criticality jobs behave whenever HI-criticality jobs overrun their optimistic LO-criticality reservation. In [47], [48], the probabilistic models are applied into classical mixed-critical scheduling policies for fixed-priority task scheduling. These works have proven the advantages in resource utilization from the use of flexible probabilistic worst-case representations. [49] introduces a new probabilistic model for mixed-criticality systems; safety metrics are defined, and an analysis is developed to quantify safety levels according to the considered criticality level. [50] discusses the accuracy and the flexibility of probabilistic models as advantages for mixed-criticality schedulability analysis in efficient resource usage. [51] depicts a survey of recent probabilistic scheduling approaches, with a dedicated section of the above introduced mixed-criticality scheduling approaches.

¹⁰To our best knowledge, there is only one paper presenting scheduling algorithms for probabilistic WCETs of tasks described by random variables [44], which extends the optimality of Audsley's approach [45] in fixed-priority scheduling to the case WCETs are described by distribution functions.

IX. CONCLUSION

This paper presented some efforts into scheduling MC systems that account for probabilistic information. Existing MC task models are generalized with an additional parameter specifying the distribution information of the WCET. We require that it is *a priori* determined how likely jobs may exceed their LO-WCETs. We proposed a novel EDF-based scheduling algorithm, which exploits the probabilistic information to make mode-switching and LO-task-dropping decisions. Given a system failure probability threshold, the goal is to derive more precise schedulability analysis, which may deem a system that is infeasible under the traditional MC model as feasible, and will not drop any task unless it is probabilistically necessary. Experimental results show the advantages of the model and the proposed scheduling schemes.

The provided solution requires a server with period of 1—applying the idea of adaptive servers (with dynamic periods or budgets) may avoid many preemptions. So far we only considered systems with two criticality levels. Probabilistic correctness for multiple probability thresholds per task needs to be defined, and the scheduling problem is worth studying.

REFERENCES

- [1] Z. Guo, L. Santinelli, and K. Yang, "EDF schedulability analysis on mixed-criticality systems with permitted failure probability," in *Proceedings - IEEE 21st International Conference on Embedded and Real-Time Computing Systems and Applications, RTCSA 2015*, 2015.
- [2] J. Souyris, E. Pavec, G. Hibert, V. Jegu, and G. Borios, "Computing the worst case execution time of an avionics program by abstract interpretation," in *the 5th International Workshop on Worst-Case Execution Time Analysis (WCET'05)*, 2005.
- [3] R. Wilhelm, J. Engblom, A. Ermedahl, N. Holsti, S. Thesing, D. Whalley, G. Bernat, C. Ferdinand, R. Heckmann, T. Mitra *et al.*, "The worst-case execution-time problem: overview of methods and survey of tools," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 7, no. 3, pp. 1–53, 2008.
- [4] A. Burns and R. Davis, "Mixed-criticality systems: A review," 2018, <http://www-users.cs.york.ac.uk/burns/review.pdf>.

- [5] F. Cazorla, E. Quinones, T. Vardanega, L. Cucu-Grosjean, B. Triquet, G. Bernat, E. Berger, J. Abella, F. Wartel, M. Houston, L. Santinelli, L. Kosmidis, C. Lo, and D. Maxim, "PROARTIS: Probabilistically analysable real-time systems," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 12, no. 2, 2013.
- [6] J. Hansen, S. Hissam, and G. Moreno, "Statistical-based WCET estimation and validation," in *the 9th International Workshop on Worst-Case Execution Time Analysis (WCET'09)*, 2009.
- [7] L. Cucu-Grosjean, L. Santinelli, M. Houston, C. Lo, T. Vardanega, L. Kosmidis, J. Abella, E. Mezzetti, E. Quinones, and F. Cazorla, "Measurement-based probabilistic timing analysis for multi-path programs," in *the 24th Euromicro Conference on Real-Time Systems (ECRTS'12)*, 2012.
- [8] R. F. S. 167, *Software considerations in airborne systems and equipment certification, RTCA document DO-178C*. RTCA, Incorporated, 1992.
- [9] L. Cucu-Grosjean, "Independence - a misunderstood property of and for probabilistic real-time systems," in *N. Audsley and S. Baruah, editors, Real-Time Systems: the past, the present and the future*, 2013.
- [10] S. Vestal, "Preemptive scheduling of multi-criticality systems with varying degrees of execution time assurance," in *the 28th IEEE Real-Time Systems Symposium (RTSS'07)*, 2007.
- [11] L. Santinelli, J. Morio, G. Dufour, and D. Jacquemart, "On the sustainability of the extreme value theory for WCET estimation," in *the 14th International Workshop on Worst-Case Execution Time Analysis (WCET'14)*, 2014.
- [12] A. Melani, E. Noulard, and L. Santinelli, "Learning from probabilities: Dependencies within real-time systems," in *the 8th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA'13)*, 2013.
- [13] J. Díaz, D. García, C. Lee, L. Bello, J. López, and O. Mirabella, "Stochastic analysis of periodic real-time systems," in *the 23rd IEEE Real-Time Systems Symposium (RTSS'02)*, 2002.
- [14] D. Maxim and L. Cucu-Grosjean, "Response time analysis for fixed-priority tasks with multiple probabilistic parameters," in *the 34th IEEE Real-Time Systems Symposium (RTSS'13)*, 2013.
- [15] B. Korte and J. Vygen, *Bin-Packing*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 471–488. [Online]. Available: https://doi.org/10.1007/978-3-642-24488-9_18
- [16] J. M. López, J. L. Díaz, and D. F. García, "Utilization bounds for edf scheduling on real-time multiprocessor systems," *Real-Time Systems*, vol. 28, no. 1, pp. 39–68, 2004.
- [17] I. E. Commission *et al.*, "Iec 61025: Fault tree analysis (fta)," *IEC Standards Online*, 2006.
- [18] I. . T. Committee *et al.*, "Analysis techniques for system reliability-procedure for failure mode and effects analysis (fmea)," *IEC 60812*, 2006.
- [19] R. W. Irving, "Np-completeness of a family of graph-colouring problems," *Discrete Applied Mathematics*, vol. 5, no. 1, pp. 111–117, 1983.
- [20] D. J. Welsh and M. B. Powell, "An upper bound for the chromatic number of a graph and its application to timetabling problems," *The Computer Journal*, vol. 10, no. 1, pp. 85–86, 1967.
- [21] S. Baruah, V. Bonifaci, G. D'Angelo, H. Li, A. Marchetti-Spaccamela, S. van der Ster, and L. Stougie, "The preemptive uniprocessor scheduling of mixed-criticality implicit-deadline sporadic task systems," in *the 24th Euromicro Conference on Real-Time Systems (ECRTS'12)*, 2012.
- [22] E. Bini and G. Buttazzo, "Measuring the performance of schedulability tests," *Real-Time Systems*, vol. 30, no. 1-2, pp. 129–154, 2005.
- [23] M. Bolado, H. Posadas, J. Castillo, P. Huerta, P. Sanchez, C. Sánchez, H. Fournet, and F. Blasco, "Platform based on open-source cores for industrial applications," in *DATE, 2004. Proceedings*, vol. 2. IEEE, 2004, pp. 1014–1019.
- [24] D. Liu, J. Spasic, N. Guan, G. Chen, S. Liu, T. Stefanov, and W. Yi, "Edf-vd scheduling of mixed-criticality systems with degraded quality guarantees," in *Proceedings of the 37th Real-Time Systems Symposium (RTSS), 2016 IEEE*. IEEE, 2016, pp. 35–46.
- [25] M. S. Branicky, S. M. Phillips, and W. Zhang, "Scheduling and feedback co-design for networked control systems," in *Proceedings of the 41st IEEE Conference on Decision and Control*, 2002., vol. 2. IEEE, 2002, pp. 1211–1217.
- [26] Z. Guo, K. Yang, S. Vaidhun, S. Arefin, S. K. Das, and H. Xiong, "Uniprocessor mixed-criticality scheduling with graceful degradation by completion rate," in *2018 IEEE Real-Time Systems Symposium (RTSS)*. IEEE, 2018, pp. 373–383.
- [27] D. Matschulat, C. A. Marcon, and F. Hessel, "Er-edf: A qos scheduler for real-time embedded systems," in *18th IEEE/IFIP International Workshop on Rapid System Prototyping (RSP'07)*. IEEE, 2007, pp. 181–188.
- [28] S. Edgar and A. Burns, "Statistical analysis of WCET for scheduling," in *the 22nd IEEE Real-Time Systems Symposium (RTSS'01)*, 2001.
- [29] D. Griffin and A. Burns, "Realism in statistical analysis of worst case execution times," in *the 10th International Workshop on Worst-Case Execution Time Analysis (WCET'10)*, 2010.
- [30] R. I. Davis, L. Santinelli, S. Altmeyer, C. Maiza, and L. Cucu-Grosjean, "Analysis of probabilistic cache related pre-emption delays," *Proceedings of the 25th IEEE Euromicro Conference on Real-Time Systems (ECRTS'13)*, 2013.
- [31] M. Slijepcevic, L. Kosmidis, J. Abella, E. Q. Nones, and F. J. Cazorla, "Dtm: Degraded test mode for fault-aware probabilistic timing analysis," in *the 25th Euromicro Conference on Real-Time Systems (ECRTS'13)*, 2013.
- [32] D. Hardy and I. Puaut, "Static probabilistic worst case execution time estimation for architectures with faulty instruction caches," in *the 21st International Conference on Real-Time and Networked Systems (RTNS'13)*, 2013.
- [33] L. David and I. Puaut, "Static determination of probabilistic execution times," in *the 16th Euromicro Conference on Real-Time Systems (ECRTS'04)*, 2004.
- [34] T. Tia, Z. Deng, M. Storch, J. Sun, L. Wu, and J. Liu, "Probabilistic performance guarantee for real-time tasks with varying computation times," in *IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS'95)*, 1995.
- [35] J. Lehoczky, "Real-time queueing theory," in *the 17th IEEE Real-Time Systems Symposium (RTSS'96)*, 1996.
- [36] C. Liu and J. Layland, "Scheduling algorithms for multiprogramming in a hard real-time environment," *Journal of the ACM*, vol. 20, no. 1, pp. 46–61, 1973.
- [37] H. Zhu, J. Hansen, J. Lehoczky, and R. Rajkumar, "Optimal partitioning for quantized EDF scheduling," in *the 23rd IEEE Real-Time Systems Symposium (RTSS'02)*, 2002.
- [38] M. Gardner and J. Liu, "Analyzing stochastic fixed-priority real-time systems," in *the 5th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'99)*, 1999.
- [39] L. Abeni and G. Buttazzo, "QoS guarantee using probabilistic deadlines," in *the 11th Euromicro Conference on Real-Time Systems (ECRTS'99)*, 1999.
- [40] J. Hansen, J. Lehoczky, H. Zhu, and R. Rajkumar, "Quantized EDF scheduling in a stochastic environment," in *the 16th IEEE International Parallel and Distributed Processing Symposium (IPDPS'02)*, 2002.
- [41] S. Manolache, P. Eles, and Z. Peng, "Schedulability analysis of applications with stochastic task execution times," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 3, no. 4, pp. 706–735, 2004.
- [42] Y. Lu, T. Nolte, I. Bate, and L. Cucu-Grosjean, "A statistical response-time analysis of real-time embedded systems," in *the 33rd IEEE Real-Time Systems Symposium (RTSS'12)*, 2012.
- [43] G. Bernat, A. Colin, and S. Petters, "pWCET: a tool for probabilistic worst-case execution time analysis of real-time systems," 2003, report - University of York Department of Computer Science YCS.
- [44] D. Maxim, O. Buffet, L. Santinelli, L. Cucu-Grosjean, and R. Davis, "Optimal priority assignment algorithms for probabilistic real-time systems," in *the 19th International Conference on Real-Time and Networked Systems (RTNS'11)*, 2011.
- [45] N. Audsley, "On priority assignment in fixed priority scheduling," *Information Processing Letters*, vol. 79, no. 1, pp. 39–44, 2001.
- [46] M. Kuttler, M. Roitzsch, Claude-Joachim, and M. Volp, "Probabilistic analysis of low-criticality execution," *Proc. WMC, RTSS*, 2017.
- [47] Y. Abdeddaïm and D. Maxim, "Probabilistic schedulability analysis for fixed priority mixed criticality real-time systems," in *Design, Automation & Test in Europe Conference & Exhibition, DATE 2017, Lausanne, Switzerland, March 27-31, 2017*, 2017, pp. 596–601.
- [48] D. Maxim, R. I. Davis, L. Cucu-Grosjean, and A. Easwaran, "Probabilistic analysis for mixed criticality systems using fixed priority preemptive scheduling," in *Proceedings of the 25th International Conference on Real-Time Networks and Systems, RTNS 2017, Grenoble, France, October 04 - 06, 2017*, 2017, pp. 237–246.
- [49] S. Draskovic, P. Huang, and L. Thiele, "On the safety of mixed-criticality scheduling," *Proc. WMC, RTSS*, 2017.
- [50] L. Santinelli and Z. Guo, "On the criticality of probabilistic worst-case execution time models," in *Dependable Software Engineering. Theories, Tools, and Applications - Third International Symposium, SETTA 2017, Changsha, China, October 23-25, 2017, Proceedings*, 2017, pp. 59–74.
- [51] R. I. Davis and L. Cucu-Grosjean, "A survey of probabilistic schedulability analysis techniques for real-time systems," *LITES*, vol. 6, no. 1, pp. 04:1–04:53, 2019.