Probabilities and Mixed-Criticalities: the Probabilistic C-Space

Luca Santinelli[1] and Laurent George[2]
[1]ONERA - Toulouse, luca.santinelli@onera.fr
[2] University of Paris Est / LIGM - ESIEE Paris, Laurent.George@univ-mlv.fr

***Abstract*—Probability distributions bring flexibility as well as accuracy in modeling and analyzing real-time systems. On the other end, the adding of probabilities increases the complexity of the scheduling problem, especially in case of mixed-criticalities where tasks of different criticalities have to be taken into account on the same computing platform. In this work we explore the flexibility of probabilistic distributions applied to mixed-critical task sets for defining the probabilistic space of Worst Case Execution Time and evaluating the effects of changes on the task execution conditions. Finally, we start formalizing and making use of probabilistic sensitivity analysis for evaluating mixed-critical scheduling performance.**

I. Introduction

During the last decade, real-time systems designers are facing the arrival of new COTS architectures and new functionalities which may result into important variability of the execution time of programs. Worst-case reasoning may have reached its limits since it considers worst-case values without taking into account that the probability of occurrence of such values may be vanishingly small [1].

In [2], the authors characterize the space of feasible Worst-Case Execution Times (WCETs) for the Earliest Deadline First (EDF) scheduling paradigm, denoted C-space. The C-space is a convex polytop of n dimensions (the number of tasks), such that for any vector of execution times inside the polytop, the task set is always feasible with EDF scheduling (deterministic approach). The deterministic approach supposes that all the jobs of the tasks can experience their WCETs at run time which is very unlikely.

Approaches that take into account tasks probabilistic worst case execution time (distributions of values instead of single values) may lead to important reduction of computing capability over-provisioning.

First papers on probabilistic approaches describe execution times of tasks by random variables, according to discrete [3], [4] or continuous [5] models. Since Edgar and Burns [6], several papers have worked on obtaining safe and reliable probabilistic Worst-Case Execution Time (pWCET) estimates [7], [8], [9].

The probabilistic worst-case reasoning leads to schedulability analysis with probabilities. Diaz et al. [10] developed the first analysis for systems with execution times described by random variables. Recent works have extended schedulability to probabilistic arrival times, using an average arrival model as in [11], or to probabilistic minimal inter-arrival times model as in [12]. Other approaches lately have revised the notion of bounding curves with probabilities, i.e. arrival curves, demand bound function and workload bound function, [13], for probabilistic guarantees of the timing constraints.

The Mixed-Criticality (MC) problem comes from the need for using the platform resources efficiently. This is facilitated by noting that the task parameters depend on their criticality level, in particular the WCET estimate will be derived by a process dictated by the criticality level. The higher the criticality level, the more conservative the verification process and hence the greater will be the WCET, [14]. In [15], the confidence on the WCET estimations has been leveraged to develop MC scheduling algorithms. In this work we intend to continue researching in that direction formalizing the relationship between pWCETs and mixed-criticalities.

Contribution: This paper formalizes the scheduling problem with tasks of different criticalities through probabilistic models. The pWCETs are applied to construct the probabilistic version of the C-space. Such fine grained probabilistic representations (pWCETs and probabilistic C-Space) are applied to leverage probabilistic information for the MC scheduling problem. This work intends to provide an initial evaluation to the flexibility brought by the probabilistic models and the probabilistic scheduling to the mixed-criticality problem. In it, the sensitivity analysis is enhanced with probabilities for the first time and it is applied to the probabilistic C-space for some initial thoughts and possible strategies for resource allocation with mixed-critical tasks.

II. Modeling with Probabilities

Jobs of tasks can exhibit multiple durations at run time due to interferences from the system elements and the environment. It is then reasonable to describe task execution time with random processes[1].

The probabilistic worst-case execution time random variable $\mathcal{C}_i$ of a task τ_i generalizes the deterministic WCET. It is defined as the worst-case distribution that upper-bounds any possible task execution time the task can exhibit, [16]. Hence, worst-case execution time distributions represents a way to account for the system variabilities as the worst-case model to all of them. In its abstract interpretation, $\mathcal{C}_i$ would includes multiple WCET values, each with the probability of being the worst-case[2]. For example, given a trace of task execution time which would be an empirical distribution due to the task execution time variability, the worst-case execution time distribution could be the distribution made out of the maximum of blocks of execution times, each block representing a specific task execution condition.

[1]A random process is a sequence of random variables describing a process whose outcomes do not follow a deterministic pattern, but follow probability distributions.

[2]Calligraphic letters are used to represent distributions while non calligraphic letters are for scalars.

The $\text{pdf}_{\mathcal{C}_i}$ is the probability distribution function (pdf) representation of the random variable $\mathcal{C}_i$. Without loss of generality, we could consider discrete pWCET distributions, that is:

$$\text{pdf}_{\mathcal{C}_i} = \begin{pmatrix} c_{i,1} & \dots & c_{i,k_i} \\ P(\mathcal{C}_i = c_{i,1}) & \dots & P(\mathcal{C}_i = c_{i,k_i}) \end{pmatrix}, \quad (1)$$

with $\text{pdf}_{\mathcal{C}_i}(c_{i,r}) = P(\mathcal{C}_i = c_{i,r})$, $\sum_{r=1}^{k_i} \text{pdf}_{\mathcal{C}_i}(c_{i,r}) = 1$, and k_i is the number of worst-case execution time values in the pWCET distribution of τ_i. It is $P(\mathcal{C}_i \leq C_i) = 1$, and the other values $c_{i,k}$, $1 \leq k \leq k_i$ are such that $C_{i,k} \leq C_i$.

$\text{cdf}_{\mathcal{C}_i}$ denotes the cumulative distribution function (cdf) representation of $\mathcal{C}_i$ such that $\text{cdf}_{\mathcal{C}_i}(c) = P(\mathcal{C}_i \leq c) = \sum_{x=1}^{c} \text{pdf}_{\mathcal{C}_i}(x)$, with discrete distributions. The inverse cumulative distribution function (icdf) $\text{icdf}_{\mathcal{C}_i}(c)$ outlines the exceedence thresholds, $\text{icdf}_{\mathcal{C}_i}(c) = P(\mathcal{C}_i \geq c)$ as the probability of having worst-case execution time larger than c. With discrete random variable, it is $\text{icdf}_{\mathcal{C}_i}(c) = 1 - \sum_{x=1}^{c} \text{pdf}_{\mathcal{C}_i}(x)$.

We also assume that the pWCET are finite distributions, with finite support, such that the safe[3] worst-case (the worst-case such that its cumulative probability is 1) is finite and is the deterministic WCET C_i; $C_i \in \mathcal{C}$ and $\text{cdf}_{\mathcal{C}_i}(C_i) = 1$. The finite support assumption allows to have the deterministic WCET belonging to the pWCET distribution, thus it is possible to do hard real-time analysis. Recent works have investigated how to derive continuous pWCETs estimates from execution time measurements in different system conditions, [17]. Discrete and finite pWCETs can always be derived as approximations at relatively low probabilities of such continuous pWCET estimates, [15].

A task τ_i is also characterized by a period T_i and a relative deadline $D_i \leq T_i$; thus the task model $\tau_i = (\mathcal{C}_i, T_i, D_i)$. In this paper we consider a set of n periodic tasks $\Gamma = \{\tau_1, \dots, \tau_n\}$, with the hyperperiod H being the least common multiple of all task periods, $H = \text{lcm}(T_1, \dots, T_n)$. Γ is scheduled with EDF.

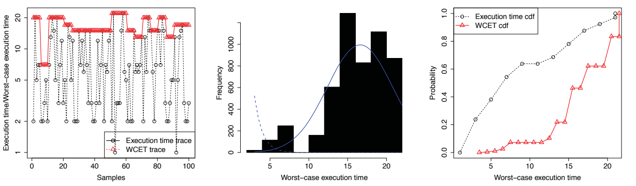

(a) Execution time and worst-case execution time trace (b) Worst-case execution time histogram representation (c) Execution time and worst-case execution time cdfs

Fig. 1. Representation of discrete task execution time and worst-case execution time distributions. Execution times are in time units.

Figure 1 shows an example of discrete pWCET. Figure 1(a) depicts a possible meaning for the worst-case distribution from a trace of execution time: a worst-case value could come from the maximum of a block of execution times, and due to system variabilities and different execution conditions, the maximum could change. Thus the maximum could end up in a probabilistic model. Figure 1(b) is for the histogram representation of the pWCET, which in that case resembles to a normal distribution. Figure 1(c) outlines the differences between execution time distribution and pWCET with the cdf representation.

[3]Safety comes from the overestimation of any possible measured behavior.

From $\mathcal{C}_i$ it is possible to define *WCET thresholds* $\langle c_{i,k}, p_{i,k} \rangle$, $1 \leq k \leq k_i$. The worst-case value $c_{i,k}$ is associated to the probability $p_{i,k}$ of being the WCET for task τ_i. $p_{i,k} \stackrel{def}{=} \text{cdf}_{\mathcal{C}_i}(c_{i,k})$ quantifies the accuracy of the WCET $c_{i,k}$, equivalently the *confidence* on $c_{i,k}$ of being the WCET. $1 - p_{i,k}$ is the probability of passing $c_{i,k}$. Depending on the granularity of the pWCET representation it would be possible to define WCET thresholds at 10^{-3}, 10^{-6}, 10^{-9}, and beyond.

A. Tasks Relationship

Most of the algebra in probability theory relies on the degree of dependence between random variables, the so called *statistical dependence*. Therefore, the task relationship can be evaluated with the degree of statistical dependence between pWCET distributions.

The joint probability, which expresses the composition of random variables, is affected by the degree of dependence among the random variables. For a couple of worst-case execution time distributions $\mathcal{C}_i$ and $\mathcal{C}_j$, respectively for τ_i and τ_j, the joint probability $P(\mathcal{C}_i = c_{i,r}, \mathcal{C}_j = c_{j,s})$ defines the probability of the worst-case execution times $\mathcal{C}_i = c_{i,r}$ and $\mathcal{C}_j = c_{j,s}$; both $\mathcal{C}_i$ and $\mathcal{C}_j$ represent events at the same time, thus concurrently executing tasks. It is $P(\mathcal{C}_i = c_{i,r}, \mathcal{C}_j = c_{j,s}) = P(\mathcal{C}_j = c_{j,s} | \mathcal{C}_i = c_{i,r}) \times P(\mathcal{C}_i = c_{i,r}) = P(\mathcal{C}_i = c_{i,r} | \mathcal{C}_j = c_{j,s}) \times P(\mathcal{C}_j = c_{j,s})$, with $P(\mathcal{C}_i = c_{i,r} | \mathcal{C}_j = c_{i,s})$ is the conditional probability defined as the probability of having $\mathcal{C}_i = c_{i,r}$ once also τ_j is executing and having $\mathcal{C}_j = c_{j,s}$. Equivalently, with the pdfs and the cdfs it is respectively:

$$\text{pdf}_{\mathcal{C}_i, \mathcal{C}_j} = \text{pdf}_{\mathcal{C}_i | \mathcal{C}_j} \otimes \text{pdf}_{\mathcal{C}_j} = \text{pdf}_{\mathcal{C}_j | \mathcal{C}_i} \otimes \text{pdf}_{\mathcal{C}_i} \quad (2)$$

$$\text{cdf}_{\mathcal{C}_i, \mathcal{C}_j} = \text{cdf}_{\mathcal{C}_i | \mathcal{C}_j} \times \text{cdf}_{\mathcal{C}_j} = \text{cdf}_{\mathcal{C}_j | \mathcal{C}_i} \times \text{cdf}_{\mathcal{C}_i}; \quad (3)$$

$\otimes$ being the convolution operator between random variables and $\text{pdf}_{\mathcal{C}_j | \mathcal{C}_i}$ is the conditional pWCET of τ_i concurrently executing with τ_j.

Two tasks τ_i and τ_j are independent, $\tau_i \overline{\rhd} \tau_j$ (equivalently $\tau_j \overline{\rhd} \tau_i$), if the execution of one task does not have any impact on the execution of the other task. Whenever the two tasks worst-case execution times are random variables $\mathcal{C}_i$ and $\mathcal{C}_j$ respectively, the independence states that the execution of one task does not affect the distribution of the other task (statistical independence): $\text{pdf}_{\mathcal{C}_i | \mathcal{C}_j} = \text{pdf}_{\mathcal{C}_i}$ and $\text{pdf}_{\mathcal{C}_j | \mathcal{C}_i} = \text{pdf}_{\mathcal{C}_j}$. Two tasks τ_i and τ_j are dependent, $\tau_i \rhd \tau_j$ (equivalently $\tau_j \rhd \tau_i$), if the execution of one task does have impact on the execution of the other task. With $\mathcal{C}_i$ and $\mathcal{C}_j$ the pWCETs of respectively τ_i and τ_j, with dependence the execution of one task does affect the distribution of the other task (statistical dependence): $\text{pdf}_{\mathcal{C}_i | \mathcal{C}_j} \neq \text{pdf}_{\mathcal{C}_i}$ and $\text{pdf}_{\mathcal{C}_j | \mathcal{C}_i} \neq \text{pdf}_{\mathcal{C}_j}$.

1) Worst-Case Independence: Interferences to task execution from concurrently executing tasks or concurrent access to shared resources have the effect of increasing the task execution time. Assuming $\mathcal{C}_i$ to be the probabilistic worst-case distribution of τ_i, it means that the distribution has already accounted for all the possible interferences, including those from other tasks. It implies also that every time there is an interference (a concurrent task or an other system element acting at the same time as τ_i), $\mathcal{C}_i$ as already embedded its effects [18]. The distribution does not change anymore in the presence of such effects being already the pWCET: $\text{pdf}_{\mathcal{C}_i | \mathcal{C}_j} = \text{pdf}_{\mathcal{C}_i}$. We can say that the $\mathcal{C}_i$, with respect to the empiric execution time distributions (from the measurements of the system actual behavior), quantifies the effect of dependence to the task executions.

III. Probabilistic Mixed-Criticality

In this paper, we consider the two-criticality-level case (high and low) of the MC problem, each task is designated as being either of high (HI) or low (LO) criticality. With the deterministic model two WCET values are specified for each task: a LO-WCET $c(\text{LO})$ determined by a less pessimistic timing analysis tools, and a larger HI-WCET $c(\text{HI})$ determined by more conservative timing analysis tools, which is sometimes larger than the LO-WCET by several orders of magnitude in COTS platforms. For τ_i, the WCET of a task is a vector $C_i = (c_i(\text{LO}), c_i(\text{HI}))$.

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. The expected probability of deadline miss depends on the criticality level $crit$, e.g. $crit \in \{\text{LO}, \text{HI}\}$, under which the system is running.

The pWCET distribution effectively defines different WCET threshold estimates for the same task, for different criticality levels depending on the different requirements confidence, e.g. on the maximum tolerable failure rate as the pWCET estimates embeds the effect of faults on the task executions. That translates into a MC two-criticality task model such that $\tau_i = ((\langle c_i(\text{LO}), p_i(\text{LO})\rangle, (\langle c_i(\text{HI}), p_i(\text{HI})\rangle), T_i, D_i)$, where the WCET values have a confidence of being worst-cases.

The cdf representation of the pWCET relates *probability* to *confidence* of the criticality levels. $p_i(\text{LO}) = P(\mathcal{C}_i \leq c_i(\text{LO})) \equiv \mathsf{cdf}_{\mathcal{C}_i}(c_i(\text{LO}))$ expresses the confidence of $c_i(\text{LO})$ of being the upper-bound of the task worst-case execution time in its LO-criticality. Similarly, $p_i(\text{HI}) = P(\mathcal{C}_i \leq c_i(\text{HI})) \equiv \mathsf{cdf}_{\mathcal{C}_i}(c_i(\text{HI}))$ the confidence of $c_i(\text{HI})$ of being the upper-bound of the task worst-case execution time in its HI-criticality We call $p_i(crit)$, $crit \in \{\text{LO}, \text{HI}\}$ the confidence on the criticality level $crit$.

A. Probability thresholding

It is possible to define a design parameter β as the *probability threshold* for the pWCET defining the level of confidence for a WCET limit imposed to a task. β comes from the quantile $q(p)$ as the probability threshold p, and $q(\mathcal{C}_i, \beta)$ is the WCET threshold such that $\beta \times 100\%$ of the worst-case execution time experienced by τ_i are below that threshold.

β offers another perspective to the task execution model. By fixing β it is possible to specify which is the limit WCET reachable, $c_i(\beta)$; β imposes a bound to the task WCET such that $c_i(\beta) = q(\mathcal{C}_i, \beta)$.

A trace $\mathcal{T}_{\mathcal{C}_i}$ reports the sequence of WCET values that τ_i has assumed from one execution instance to another. From $\mathcal{T}_{\mathcal{C}_i}$ it is then possible to infer the timing behavior of the task WCETs as well as identify $c_i(\beta) = q(\mathcal{T}_{\mathcal{C}_i}, \beta)$. Therefore, β can model the task (or the whole application) timing behavior and it could be applied as design parameter: by imposing $c_i(\beta)$ as the task WCET value the behavior of τ_i is limited to $c_i(\beta)$. With respect of the actual task behavior $\mathcal{T}_{\mathcal{C}_i}$ (which follow $\mathcal{C}_i$, β is the confidence that τ_i respects it WCET limit $c_i(\beta)$.

From β it is also possible to infer the criticality level $crit$ that would allow respecting $c_i(\beta)$:

$$max\{crit\} \text{ such that } c_i(crit) \leq c_i(\beta). \tag{4}$$

It is $\beta \neq p(crit)$, as $c(\beta) \neq c(crit)$, but there is a close relationship between the two thresholds $c(\beta)$ and $c(crit)$ which come from the probabilistic modeling of the task (the pWCET).

Example 1. *Given a task set* $\Gamma = \{\tau_1, \tau_2, \tau_3\}$*, with* $\tau_1 = (\mathcal{C}_1, 7, 5)$*,* $\tau_2 = (\mathcal{C}_2, 11, 7)$*,* $\tau_1 = (\mathcal{C}_3, 13, 10)$*, and the discrete worst-case execution time distributions are*

$$\mathsf{pdf}_{\mathcal{C}_1} = (\begin{pmatrix} 1 & 2 & 3 & 4 & 6 \\ 0.3 & 0.3 & 0.1 & 0.15 & 0.15 \end{pmatrix}$$

$$\mathsf{pdf}_{\mathcal{C}_2} = (\begin{pmatrix} 1 & 2 & 3 & 5 & 6 & 7 \\ 0.1 & 0.2 & 0.4 & 0.2 & 0.07 & 0.03 \end{pmatrix}$$

$$\mathsf{pdf}_{\mathcal{C}_3} = (\begin{pmatrix} 3 & 4 & 5 & 6 & 8 \\ 0.2 & 0.1 & 0.3 & 0.3 & 0.1 \end{pmatrix}.$$

Considering criticality levels such that for τ_1 *it is* $\{c(\text{LO}) = 1, c(\text{HI}) = 4\}$*, for* τ_2 *it is* $\{c(\text{LO}) = 1, c(\text{HI}) = 7\}$*, and for* τ_3 *it is* $\{c(\text{LO}) = 4, c(\text{HI}) = 8\}$*. Figure 2 represents the tasks pWCETs*

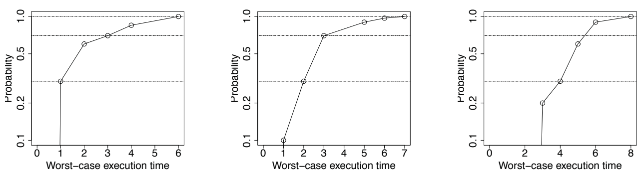

(a) τ_1 (b) τ_2 (c) τ_3

Fig. 2. A portion of the pWCET cdf for the three tasks. Three probability thresholds are outlined, respectively at $p = 0.3$ $p = 0.7$, and $p = 1$.

with a zoom of the cdfs on the largest cumulative probabilities. In Figure 3(a) an example of thresholds β *for* $\mathcal{C}_2$ *and their effects on the task [worst-case] executions. It is represented a trace of* 500 *worst-cases extracted from* $\mathcal{C}_2$ *by randomly picking values from the distributions law, and associated to the task execution. As already mention, the trace of worst-case execution times could represent the trace of maximum among different task execution conditions. For a* $\beta_3 = 1$ *the WCET limit is* 7*, the maximum allowed value. For* $\beta_2 = 0.8$ *the task WCET would be limited to* 5*, and for* $\beta_2 = 0.5$ *the task WCET would be limited to* 3*. For example, with* $c_2(\beta_2) = 5$ *there is a confidence of* 0.8 *of remaining below that WCET values while* τ_2 *executes. Hence, imposing* $C_2 = 5$ *as the max WCET for* τ_2 *it will be respected* 80% *of the case.*

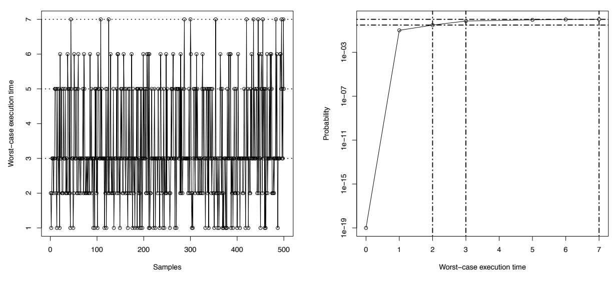

(a) βs from traces; $\beta_1 = 0.5$, $\beta_2 = 0.8$, and $\beta_3 = 1$

(b) cdf with criticality levels (horizontal lines - WCET thresholds) and βs (vertical lines - probability thresholds)

Fig. 3. Trace and cdf representations of the pWCET.

Figure 3(b) outlines the relationship between β and the criticality levels.

IV. The Probabilistic C-space

Real-time systems with probabilistic models require schedulability conditions which involve probabilities. Given a random process $\mathcal{C}_i$ describing the evolution of τ_i worst-case execution time, we can state the notion of probabilistic demand bound function, [13], [19].

$\mathsf{dbf}_{i,j}(t)$ is the demand bound function in the interval $[0, t]$:

$$\mathsf{dbf}_{i,j}(t) \stackrel{def}{=} \lfloor \frac{t - D_i}{T_i} + 1 \rfloor \times c_{i,j}. \tag{5}$$

The bound is the result of a specific WCET threshold $c_{i,j}$, and by definition, it represents an upper-bound to any $\mathsf{dbf}_{i,k}$ obtained from $c_{i,k} \leq c_{i,j}$. Then, there exist an associate confidence of the bound, which is the exact confidence $p_{i,j}$ of the WCET $c_{i,j}$ applied. $c_{i,j}$ and Equation (5) define a probabilistic bound $\langle \mathsf{dbf}_{i,j}(t), p_{i,j} \rangle$ to the task resource demand; $\langle \mathsf{dbf}_{i,j}(t), p_{i,j} \rangle$ is such that $p_{i,j}$ is the probability that $\mathsf{dbf}_{i,j}(t)$ is an upper-bound to the τ_i resource demand in $[0, t]$. Equivalently to $\mathsf{dbf}_{i,j}(t)$ we can write $\mathsf{dbf}_i(t, c_{i,j})$.

As $\langle \mathsf{dbf}_i(t, c_{i,j}), p_{i,j} \rangle$ represents a single demand bound function with its associated confidence, there exist a $\mathsf{dbf}_i(t, c_{i,j})$ for each $c_{i,j} \in \mathcal{C}_i$. All together the $\langle \mathsf{dbf}_i(t, c_{i,j}), p_{i,j} \rangle$ form a distribution of demand bound functions, $\mathcal{DBF}_i(t) = (\mathsf{dbf}_i(t, \cdot), p_i(\cdot))$ which is the probabilistic demand bound function (probabilistic demand curve) of τ_i in $[0, t]$. $\mathcal{DBF}_i(t)$ collects the set of all demand bound functions $\mathsf{dbf}_i(t)$ and the set of all confidences p_i. In particular, the p_is forms the the cdf of $\mathcal{DBF}_i(t)$, $\mathsf{cdf}_{\mathcal{DBF}_i(t)}$, as cumulative probabilities. To note that the set of probabilities p_i does not change with the interval $[0, t]$, therefore form one interval to another is only the bounds $\mathsf{dbf}_i(t)$ to change, but not their confidence.

The application Γ probabilistic demand curve $\mathcal{DBF} = (\mathsf{dbf}(t, \cdot), p(\cdot))$ results from the combination of tasks demand bound functions $\mathcal{DBF}_i$:

$$\mathcal{DBF}(t) = \otimes_i \mathcal{DBF}_i(t), \tag{6}$$

with $\otimes$ the convolution of the distributions. $\mathsf{dbf}(t, \cdot)$ is the set of all the demand bound functions:

$$\mathsf{dbf}(t, \cdot) \stackrel{def}{=} +_i \mathsf{dbf}_i(t, \cdot), \tag{7}$$

with $+$ the combination (sum) of all the demand bound function. $p(\cdot)$ is the set of all the confidence probabilities:

$$p(\cdot) \stackrel{def}{=} \times_i p_i(\cdot), \tag{8}$$

with $\times$ the combination (product) of all the demand bound function probabilities.

The demand bound function $\mathsf{dbf}(t, \overline{c})$ is the application demand with $\overline{c} = (c_{1,j}, c_{2,k}, \ldots)$ the array of WCET thresholds used for achieving $\mathsf{dbf}(t, \overline{c})$; $p(\overline{c})$ is the confidence of $\mathsf{dbf}(t, \overline{c})$ such that:

$$p(\overline{c}) = p_1(c_{1,j}) \times p_2(c_{2,k}) \times \ldots. \tag{9}$$

The probability multiplication for the joint probability $p(\overline{c})$ is possible due to the worst-case distribution assumption. As $\mathcal{C}_i$ are pWCETs they are independent, the distributions $\mathcal{DBF}_i$ are independent among each other; consequently the joint probability $p(\cdot)$ could result from the probability multiplication, Equation (9).

A. Probabilistic C-space Representations

The schedulability under EDF states that

$$\forall t \in D, \quad \mathsf{dbf}(t) \leq t, \tag{10}$$

with D the set of Γ deadlines within the hyperperiod, according to [20], [21].

With a probabilistic framework each condition $\mathsf{dbf}(t, \cdot) \leq t$ has a probability $p(t)$ associated, which is the confidence on the demand bound function $\mathsf{dbf}(t, \cdot)$. Being $p = p(\overline{c})$ the probability of not passing $\mathsf{dbf}(t, \overline{c})$, with the condition $\mathsf{dbf}(t) \leq t$ the probability p could be also interpreted as the probability of verifying the condition.

For all $t \in D$, there exist $\overline{c}^*$ such that $\mathsf{dbf}(t, \overline{c}^*) = max\{\mathsf{dbf}(t, \overline{c}) \mid \mathsf{dbf}(t, \overline{c}) \leq t\}$. $P(t) = p(\overline{c}^*)$ from Equation (9) is the probability for which $\mathsf{dbf}(t) \leq t$ is true. The overall schedulability probability P is given such that all the conditions are satisfied:

$$P = P(t_1) \times P(t_2) \times \ldots, \tag{11}$$

with $P(t_k)$ the schedulability probability of the k-th condition $\mathsf{dbf}(t_k) \leq t_k$, $t_k \in D$. The independence between conditions and the probability product as the joint probability, are guaranteed by the use of pWCET distributions. $1 - P$ is the probability that at least one condition is not respected, thus the probability of deadline miss.

From Condition (10) and Equation (11) it is possible to build the probabilistic version of the C-space (pC-space), [2]. The pC-space is the abstraction that applies the schedulability condition, Condition (10), to a vector of execution times $\overline{c} = \{c_1, c_2, \ldots\}$. Each point $\overline{c} = \{c_1, c_2, \ldots\}$ in the pC-Space is a combination of task WCET thresholds. Within the pC-Space, given the scheduling policy, it is possible to define the schedulability region where every point $\overline{c}$ within the region is a schedulable WCET thresholds configuration, and the points outside the region do not represent schedulable WCET thresholds configurations. [2] for the details on the definition of the deterministic C-space under EDF.

The pC-space maps also probabilities onto points. Each $\overline{c}$ within the space has a probability associated which is the probability of being the application set of worst-case execution times, Equation (9). Then, depending on where the point is with respect to the schedulability region, the probability could translate into schedulability probability. For the points at the feasibility region border, their ps, Equation (9), are exactly the schedulability probability, Equation (11). With the different probabilities P within the region and at the border it would be possible to classify portions of the regions with respect to the schedulability probability P.

B. pC-Space and Confidence

The probabilities within the pC-Space can be interpreted in various ways:

- as the confidence of not passing the WCET thresholds of $\overline{c}$. With the criticality levels there is also the probability of remaining at a certain criticality level $p(crit) = p_1(crit) \times p_2(crit) \times \ldots$. Consequently it is quantifiable the possibility of changing that level as $1 - p(crit)$;
- as the confidence on the system schedulability P, or the confidence per schedulability condition, P_k. The feasibility region is characterized by P and all the points inside

the region are schedulable but with a confidence of at least P, Equation (11). It translates into a per-condition schedulability probability of P_k;
- as the confidence β on the worst-case behavior of the tasks. β is the probability of passing the $\overline{c}(\beta)$; per-task it would be $c_j(\beta)$.

With different probability interpretations the pC-Space can be used for different purposes. At one end there is the modeling of the probabilistic applications; on the other end, it is possible to develop analysis on top of the pC-Space with probabilities.

Example 2. *Given the probabilistic task set of Example 1. The feasibility region of* Γ *does not depend on the input distributions (and their shape) but it describes the feasibility point according to the task period and deadline configuration,* (x, T_i, D_i)*. What is depending on the pWCETs instead, are the probabilities of each point. In Figure 4 all the possible WCET thresholds and the WCET thresholds combinations from the input pWCETs are plotted. To notice that the whole plane at* $\mathcal{C}_1 = 6$ *is an unfeasible plane, being outside the feasibility region. The confidences are presented together with the feasible points* $\overline{c}$ *(points with circles and in green). Points with crosses (in red) are WCET thresholds configurations unfeasible. Figure 5 gives a better insight with 2D representations of both the pC-Space and the feasibility regions. In both figures the* β*s are represented as constraints to the task execution behavior.*

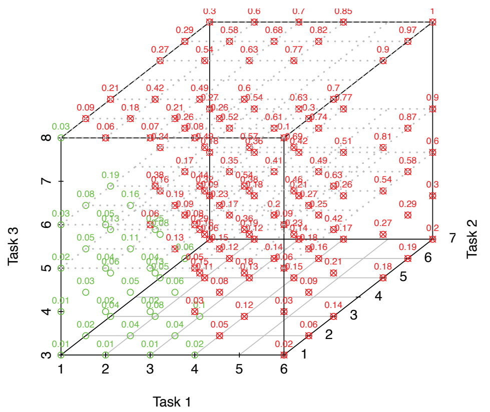

Fig. 4. Probabilistic C-space: feasibility and confidences (probabilities) within the pC-space. Circles for feasible $\overline{c}$, crosses for not feasible $\overline{c}$) and βs limitations ($\beta_1 = 0.5$,$\beta_2 = 0.8$, and $\beta_3 = 1$) for τ_3 are presented as horizontal 2D planes.

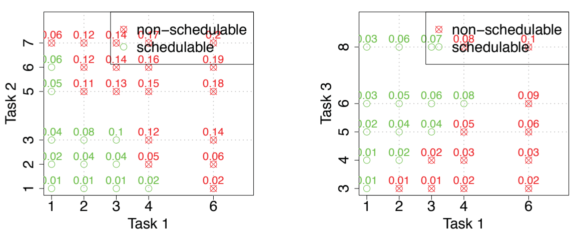

(a) 2D plane (τ_1, τ_2) with $\mathcal{C}_3 = 3$ (b) 2D plane (τ_1, τ_3) with $\mathcal{C}_2 = 1$

Fig. 5. Probabilistic C-space: 2D planes with bounding βs ($\beta_1 = 0.5$,$\beta_2 = 0.8$, and $\beta_3 = 1$).

V. Probabilistic Sensitivity Analysis

The probabilistic version of the sensitivity analysis [2] intends to combine the information from the probabilistic models (the pWCETs, β, and the confidences β and ps) and the pC-space representation.

We have seen that $\mathcal{C}_i$s discretize the pC-Space as they maps the points to only the possible WCET thresholds of the tasks. Out of that, the probabilistic sensitivity analysis can be applied to quantify the effects of changes in terms of schedulability, probabilities/confidences, and criticalities.

- What are the resource demand that can be accommodated? Hence, which task combinations can be accounted for a schedulable systems, the criticality levels that can be considered in order to make the system schedulable, etc.
- What can be done with β? By acting on β (limiting task WCETs to $c(\beta)$) it is possible to evaluate the effect on the execution of tasks. What are the effect of β on the tasks criticality levels? With the relationship $\beta \rightarrow crit$ is is possible to infer the criticality levels which subdue to the $\overline{c}(\beta)$ bounding.

Furthermore, with the probabilistic sensitivity analysis it is possible to evaluate the effect of changes on Γ. For example a change on β, from β to β' would result into a WCET threshold change $\overline{c}$ to $\overline{c}'$, such that $\overline{c} = (c_{1,j}, c_{2,k}, \dots)$ and $\overline{c}' = (c_{1,r}, c_{2,s}, \dots)$. The change of probabilities, from $p(\overline{c})$ to $p(\overline{c}')$, is an immediate consequence of the change of β. It would also be evident the effects of changes on the allowed criticality levels, from $\beta \rightarrow crit$.

While P does not change by moving the points toward the feasibility region (by limiting task execution behavior with β), it is possible to increase the confidence that the feasibility condition is respected.

Example 3. *The probabilistic sensitivity analysis, with respect to the previous example, could help replying to the following questions:*

- *With a certain mixed-criticality level, is the system schedulable? For example, in Figure 7(c), if* τ_3 *is in HI-criticality mode, then* τ_2 *can only be scheduled with* $\mathcal{C}_2 = 1$ *(LO-criticality) for guaranteeing schedulability. Another example from Figure 7(a), where only LO-criticality modes are schedulable for* τ_1 *and* τ_2*.*
- *What can be done to make the system schedulable? What are the costs of being schedulable?* β *and its limitation effects on the task WCETs can give answers to those questions. From Figure 7(b), only limiting WCETs with* $\beta < 0.5$ *would guarantee* Γ *schedulability. This translates into LO-criticality execution for both* τ_1 *and* τ_3*.*

Figure 6 representing the discretized feasibility region (to the possible WCET thresholds) and the feasible $\overline{c}$*s. Figure 7 for 2D representations.* β *limitation effects are evident to the tasks WCET thresholds and the criticality level allowed.*

VI. Conclusion

In this work we have begun combining the probabilities and the mixed-criticality problem with the help of the probabilistic C-space and of the probabilistic sensitivity analysis. Some of the observations on the probabilistic task sets are just initial ideas which could help designing and applying more effective MC scheduling.

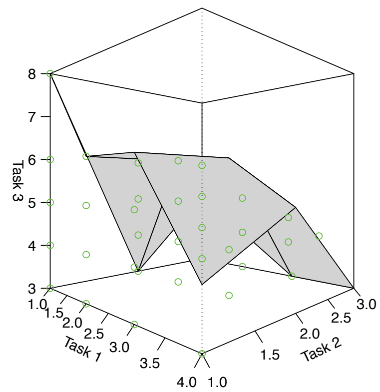

Fig. 6. Probabilistic sensitivity analysis from the feasibility region (points with circles) and the βs ($\beta_1 = 0.5$,$\beta_2 = 0.8$, and $\beta_3 = 1$).

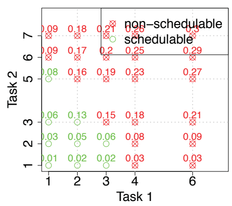

(a) 2D plane (τ_1, τ_2) with $\mathcal{C}_3 = 3$

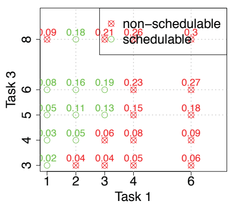

(b) 2D plane (τ_1, τ_3) with $\mathcal{C}_2 = 2$

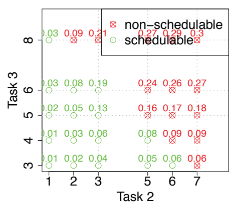

(c) 2D plane (τ_2, τ_3) with $\mathcal{C}_1 = 1$

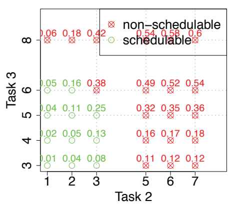

(d) 2D plane (τ_2, τ_3) with $\mathcal{C}_1 = 2$

Fig. 7. Probabilistic sensitivity analysis on the 2D planes.

In the future we intend to enhance such observations and define the probabilistic sensitivity analysis in terms of change strategies and effect evaluation. We aim at leveraging the information from the probabilistic models (pWCET distributions and confidences) and provide system design feedbacks for an optimal (at least suboptimal) system resource utilization for different criticalities, thus different requirements.

References

[1] F. J. Cazorla, E. Quiñones, T. Vardanega, L. Cucu, B. Triquet, G. Bernat, E. D. Berger, J. Abella, F. Wartel, M. Houston, L. Santinelli, L. Kosmidis, C. Lo, and D. Maxim, "Proartis: Probabilistically analyzable real-time systems," *ACM Trans. Embedded Comput. Syst.*, vol. 12, no. 2s, p. 94, 2013.

[2] L. George and J. Hermant, "Characterization of the space of feasible worst-case execution times for earliest-deadline-first scheduling," *Journal of Aerospace Computing, Information and Communication (JACIC)*, vol. 6, no. 11, 2009.

[3] T. Tia, Z. Deng, M. Shankar, 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*, 1995.

[4] L. Abeni and Buttazzo, "QoS guarantee using probabilistic deadlines," in *IEEE Euromicro Conference on Real-Time Systems (ECRTS99)*, 1999.

[5] J. Lehoczky, "Real-time queueing theory," in *10th of the IEEE Real-Time Systems Symposium (RTSS96)*, 1996, pp. 186–195.

[6] S. Edgar and A. Burns, "Statistical analysis of WCET for scheduling," in *22nd of the IEEE Real-Time Systems Symposium*, 2001.

[7] J. Hansen, S. Hissam, , and G. Moreno, "Statistical- based wcet estimation and validation," in *the 9th International Workshop on Worst-Case Execution Time*, 2009.

[8] L. Cucu-Grosjean, L. Santinelli, M. Houston, C. Lo, T. Vardanega, L. Kosmidis, J. Abella, E. Mezzetti, E. Quiñones, and F. J. Cazorla, "Measurement-based probabilistic timing analysis for multi-path programs," in *ECRTS*, 2012, pp. 91–101.

[9] D. Hardy and I. Puaut, "Static probabilistic worst case execution time estimation for architectures with faulty instruction caches," in *International Conference on Real-Time Networks and Systems (RTNS)*, 2013.

[10] J. Díaz, D. Garcia, K. Kim, C. Lee, L. Bello, L. J.M., and O. Mirabella, "Stochastic analysis of periodic real-time systems," in *23rd of the IEEE Real-Time Systems Symposium (RTSS02)*, 2002, pp. 289–300.

[11] G. Kaczynski, L. Lo Bello, and T. Nolte, "Deriving exact stochastic response times of periodic tasks in hybrid priority-driven soft real-time systems," *12th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA'07), Greece*, 2007.

[12] D. Maxim and L. Cucu-Grosjean, "Response time analysis for fixed-priority tasks with multiple probabilistic parameters," in *IEEE Real-Time Systems Symposium*, 2013.

[13] L. Santinelli and L. Cucu-Grosjean, "A probabilistic calculus for probabilistic real-time systems," *ACM Trans. Embedded Comput. Syst.*, vol. 14, no. 3, p. 52, 2015.

[14] S. Vestal, "Preemptive scheduling of multi-criticality systems with varying degrees of execution time assurance," in *Proceedings of the 28th IEEE International Real-Time Systems Symposium*, ser. RTSS '07. IEEE Computer Society, 2007, pp. 239–243.

[15] Z. Guo, L. Santinelli, and K. Yang, "Edf schedulability analysis on mixed-criticality systems with permitted failure probability," in *21th IEEE International Conference on Embedded and Real-Time Computing System and Applications*, 2015.

[16] 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)*, 2013.

[17] L. Cucu-Grosjean, L. Santinelli, M. Houston, C. Lo, T. Vardanega, L. Kosmidis, J. Abella, E. Mezzeti, E. Quinones, and F. J. Cazorla, "Measurement-Based Probabilistic Timing Analysis for Multi-path Programs," in *23nd Euromicro Conference on Real-Time Systems (ECRTS)*. IEEE, 2012.

[18] L. Cucu-Grosjean, "Independence - a misunderstood property of and for (probabilistic) real-time systems," in *the 60th Anniversary of A. Burns, York*, 2013.

[19] L. Santinelli, P. Meumeu Yomsy, D. Maxim, and L. Cucu-Grosjean, "A component-based framework for modeling and analysing probabilistic real-time systems," in *16th IEEE International Conference on Emerging Technologies and Factory Automation*, 2011.

[20] E. Bini and G. C. Buttazzo, "The space of EDF feasible deadlines," in *19th Euromicro Conference on Real-Time Systems, ECRTS'07, 4-6 July 2007, Pisa, Italy, Proceedings*, 2007, pp. 19–28.

[21] S. K. Baruah, L. E. Rosier, and R. R. Howell, "Algorithms and complexity concerning the preemptive scheduling of periodic, real-time tasks on one processor," *Real-Time Systems*, vol. 2, no. 4, pp. 301–324, 1990.