Analytical Approximations in Probabilistic Analysis of Real-Time Systems

Filip Marković[1,2], Thomas Nolte[1], and Alessandro Vittorio Papadopoulos[1]
[1]*Mälardalen University (MDU)*, Sweden
[2]*Max Planck Institute for Software Systems (MPI-SWS)*, Germany

***Abstract*—Probabilistic timing and schedulability analysis of real-time systems is constrained by the problem of often intractable exact computations. The intractability problem is present whenever there is a large number of entities to be analysed, e.g., jobs, tasks, etc. In the last few years, the analytical approximations for deadline-miss probability emerged as an important solution in the above problem domain.**

In this paper, we explore analytical solutions for two major problems that are present in the probabilistic analysis of real-time systems. First, for a safe approximation of the entire probability distributions (e.g., of the accumulated execution workloads) we show how the Berry-Esseen theorem can be used. Second, we propose an approximation built on the Berry-Esseen theorem for efficient computation of the quantile functions of probability execution distributions. We also show the asymptotic bounds on the execution distribution of the fixed-priority preemptive tasks.

In the evaluation, we investigate the complexity and accuracy of the proposed methods as the number of analysed jobs and tasks increases. The methods are compared with the circular convolution approach. We also investigate the memory footprint comparison between the proposed Berry-Esseen-based solutions and the circular convolution. . The contributions and results presented in this paper complement the state-of-the-art in accurate and efficient probabilistic analysis of real-time systems.

***Index Terms*—probabilistic timing analysis, probabilistic schedulability analysis, analytical bounds, Berry-Esseen theorem**

I. Introduction

The analysis of hard real-time systems has been built on the foundations of various mathematical concepts such as analytical bounds, fixed-point recursions, Linear Programming, etc. Among the most important concepts being used, there are the linear and non-linear bounds which allow for efficient and accurate analysis of different aspects of real-time systems, e.g., feasibility, schedulability, resource bandwidth, etc.

When hard real-time systems are considered, the bounds must be deterministic. The evolution of bound-based analysis started from the seminal paper by Liu and Layland [28], resulting in many analytical bounds for various model assumptions, e.g. [2], [7], [9], [25].

However, the majority of real-time systems exhibit an execution time that is typically lower than the estimated worst-case, which often leads to the corresponding resource provisioning being pessimistic. Diverse research efforts have been devoted to overcoming such pessimism while providing tools for analysing relevant real-time properties. More specifically, in recent years, analytical bounds on deadline miss probability have been proposed to solve the following (generalised) problem.

Problem 1. *How to efficiently and accurately derive an upper bound on the probability that a distribution (e.g., of an execution workload) exceeds a given value (e.g., an arbitrary time point, or a deadline)?*

To solve the problem, several probabilistic inequalities were formally adjusted to be used in real-time systems, e.g., the Hoeffding [26] and Bernstein bound [3], as shown by von der Brüggen et al. [47], [48], and the Chernoff bound [16], as shown by Chen et al. [12], [14], [15].

As discussed in the survey by Davis and Cucu-Grosjean [18] there are also problems that directly benefit from the computation of the entire probabilistic response time distributions, task workloads and their cumulative distribution functions (CDFs). To quote the survey by Davis and Cucu-Grosjean [18](Section 3.4, p.23) *"Two key problems that remain with probabilistic response time analysis are the tractability of the analysis for task sets of practical sizes"*. In the same survey [18] the authors highlighted the significance of the probabilistic response time analysis that considers multiple hyper-periods (Section 3.2, p.18): *"In contrast to classical task models, task sets containing a number of tasks with execution times described by random variables can usefully have a total worst-case processor utilisation that exceeds 1. This means that there is a backlog, meaning outstanding task execution with a finite probability of occurrence, at the end of each hyperperiod. This backlog makes the analysis of probabilistic response times for each job in the hyperperiod much more complex"*. Therefore, we formulate Problem 2, the *first problem* addressed in this paper.

Problem 2. *How to efficiently, accurately, and safely, approximate a probability distribution (e.g., of an execution workload or a response-time) whose exact computation is intractable?*

This problem is also relevant in areas such as probabilistic cache and WCET analysis [17], [36] (see [19] for a more comprehensive list), and for this reason, we stated the problem in a more general form. In Problem 2, the term *intractability* considers the computation demands in terms of space (memory) and time, which cannot be met by computing the exact distributions (e.g., using the linear or circular convolution-based approaches). This is a common problem, as identified by Davis and Cucu-Grosjean [18], and remains one of the important unsolved problems in the probabilistic analysis of real-time systems. The *second problem* addressed in this paper is stated below.

Problem 3. *How to efficiently and accurately derive the least value x (e.g., time point within some interval), such that the probability that a distribution (e.g., execution workload) precedes x is greater than or equal to some probability threshold p.*

This problem can be particularly relevant for the analysis and control of real-time systems and is the inverse of Problem 1. In general, quantiles can provide important information on where certain workloads should be placed in the timeline so that they minimise the probabilistic interference upon the other workloads they might affect. The main issue with utilising this approach is that it is often too expensive to compute the exact quantile, especially in the domain of probabilistic real-time system analysis, where possibly many hyper-periods dictate the final distribution lengths, as pointed out by Davis and Cucu-Grosjean [18]. The three problems are discussed in more detail in the context of the state-of-the-art in Section II.

A. Contributions of this paper

In this paper, we propose and investigate the effectiveness of solutions (see Table I) to Problems 1, 2, and 3, taking advantage of the properties of the Lyapunov Central Limit Theorem (CLT). The regular font in the table indicates bounds that have been adjusted and improved in the literature for the probabilistic analysis of real-time systems. The bold font indicates bounds adjusted, improved, or proposed in this paper.

TABLE I
LIST OF INEQUALITIES AND BOUNDS THAT ARE ADJUSTED AND PROPOSED (**BOLD**) FOR PROBABILISTIC ANALYSIS OF REAL-TIME SYSTEMS.

Problem 1	**Problem 2**	**Problem 3**
• Chernoff bound [12] • Bernstein bound [48] • Hoeffding bound [48] • **Lyapunov CLT**	• **Berry-Esseen bound**	• Short quantile [42] • **Berry-Esseen quantile**

Regarding the bounds in this paper, the contributions are:

- We prove the safe bounds on the asymptotic behaviour of the probability distribution of the preemptive independent tasks, thus contributing also to the asymptotic behaviour with respect to the Problems 1 to 3, using Lyapunov CLT.
- We adjust the Berry-Esseen inequality [4], [22] in order to safely approximate the entire distributions in the probabilistic analysis of real-time systems, addressing Problem 2.
- We use the Berry-Esseen inequality and Lyapunov's CLT in order to address Problem 3.

The evaluation shows that the proposed use of the Berry-Esseen theorem accurately, safely, and efficiently approximates the accumulated distributions of computation-demanding tasks and job workloads as the problem size increases. The improvement in computation time and space efficiency is significantly greater compared to the circular convolution approach [32], which is the most efficient and accurate method found in the state-of-the-art for Problems 2 and 3. The conclusion of the evaluation is that the two methods complement each other, where the circular convolution is very efficient when the number of analysed entities is small, whereas, by increasing the number of analysed entities, the Berry-Esseen formulation is increasingly more accurate and efficient, which is a direct implication of the central limit theorem that is integrated within the proposed method.

B. Paper outline

Section II reviews the related work. Section III presents the mathematical background. Section IV establishes the safety and applicability of the approximations proposed in the paper, while in Section V we analyse the asymptotic behaviour of the approximation. In Section VI we describe the proposed solution to Problem 2. Section VII describes the proposed solution to Problem 3. Evaluation of the solutions is presented in Section VIII. The paper is concluded in Section IX.

II. Related Work

Analytical bounds emerged as efficient solutions for the intractability problems of probabilistic timing and schedulability analysis, as shown by Davis and Cucu-Grosjean [18], [19].

Considering Problem 1, the seminal work on those bounds was contributed by Chen et al. [12] where they adapted the use of Chernoff bound for the probabilistic analysis of real-time systems. Also, Chen et al. [14] proposed the use of the golden-section search to address the optimisation problem that naturally occurs in Chernoff bounds. Taking into account closed-form solutions, von der Brüggen et al. [47], [48] proposed the use of Bernstein and Hoeffding bounds that are more efficient than the Chernoff bound-based solutions, but also less accurate. Lastly, Chen et al. [13] revisited the critical instant for probabilistic real-time systems and proposed corrections. One of them is used in this paper (Theorem 2).

In the domain of non-analytical contributions to Problems 1 and 2, Milutinović et al. [36] proposed several improvements for computing convolution between random variables in the context of probabilistic cache analysis, although applicable in the general problem space of real-time system analysis. Recently, Marković et al. [32] proposed the use of circular convolution to efficiently compute the sum of random variables that represent the execution time modes. There is also a rich area of downsampling methods which addressed the problems of time and space complexity of performing convolutions, e.g. [21], [27], [34], [35], [37]. Bozhko et al. [10] proposed the Monte Carlo simulation to calculate deadline-miss probability.

Considering Problem 2, Section 3 in the survey from Davis and Cucu-Grosjean [18] points to many non-analytical methods to calculate probabilistic response time, execution workload, or backlog distributions. However, as quoted from [18] in Section I, such methods suffer from concerns about efficiency and intractability, especially when a job backlog is possible. Various non-analytical methods exist for the periodic task model [1], [11], [20], [21], [23], [27], [30], [44]–[46], [50], and a few for the sporadic task model [33], [49]. The majority of them rely on convolutions with the noticeable exception of [11] which uses a Stochastic Timed Petri Nets.

Considering Problem 3, its formulation is very similar to the quantile function, and this concept was used several times in the analysis and control-optimisation problems of real-time systems. For example, Short and Proenza [42] proposed a stochastic error model in a real-time system, and one of the problems in that work was the need for efficient and accurate computation of the upper-tail quantiles, due to quick control decisions that were based on the result of the computation. This work later motivated the mathematical contribution [41] by Short, where the problem of the efficient approximation of the upper-tail quantile functions for Poisson and Binomial Distribution was proposed. Following a similar line of utilising quantiles for control decisions, Bertini et al. [5] showed that such an approach may improve the predictability of workloads thus directly propagating in more accurate quality of service control. In a different real-time setup, Marković et al. [31] proposed the use of quantile functions for controlling the distribution of the preemption overheads thus minimising the deadline-miss probability of tasks in a system.

In this work, we aim at providing an efficient and accurate approximation, addressing Problem 2, and also providing the efficient and accurate way for computing quantiles, given a general random variable, not constraining the problem only to Poisson and Binomial distributions (since those were addressed by Short and Proenza).

III. Mathematical notation, system assumptions and terminology

In this section, we first describe mathematical notation and terminology, and then provide a system model that can straightforwardly benefit from the equations and theorems used in this paper. Throughout the main content of the paper, the final equations consider a real-time task model, described in Section III-B, while some equations are mathematically generalised so that they can be adapted to broader assumptions.

A. Mathematical notation and terminology

Definition 1 (Discrete Random Variable)**.** *A discrete random variable X on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is defined to be a measurable function $X : \Omega \to \mathbb{R}$ such that the image $X(\Omega)$ is a countable subset of $\mathbb{R}$, and $\{\omega \in \Omega : X(\omega) = x\} \in \mathcal{F}$ for $x \in \mathbb{R}$.*

In the above definition, Ω is a sample space, the set of all possible outcomes. $\mathcal{F}$ is an event space, where an event is a set of outcomes in the sample space. $\mathbb{P}$ represents a probability function, that assigns each event in the event space a probability.

The image of Ω under X is denoted with $\operatorname{Im} X$, and it is the set of values taken by X, with a strictly positive probability.

Given a random variable X, we define its probability mass function as $f_X(x) \triangleq \mathbb{P}(X = x)$, its cumulative distribution function as $F_X(x) \triangleq \mathbb{P}(X \leq x)$, and its expected value (also called expectation, mean, or the first moment) as $\mathbb{E}[X] \triangleq \sum_{x \in \operatorname{Im} X} x \cdot \mathbb{P}(X = x)$, while its variance is defined as $\mathbb{V}[X] \triangleq \sum_{x \in \operatorname{Im} X} \mathbb{P}(X = x) \cdot (x - \mathbb{E}[X])^2 = \mathbb{E}\left[(X - \mathbb{E}[X])^2\right]$.

Definition 2 (Usual Stochastic order [38])**.** *Two random variables X and Y, with cumulative distribution functions F_X and F_Y, are said to be in the usual stochastic order, denoted as $X \succeq Y$, if and only if $\forall x$, $F_X(x) \leq F_Y(x)$.*

In the literature of probabilistic timing analysis, X is considered a safe approximation (upper bound) of Y if $X \succeq Y$.

Definition 3 (Independence)**.** *Two (discrete) random variables X and Y are independent if the pair of events $\{X = x\}$ and $\{Y = y\}$ are independent for all $x, y \in \mathbb{R}$. Formally,*

$$\mathbb{P}(X = x, Y = y) = \mathbb{P}(X = x)\mathbb{P}(Y = y), \quad \forall\, x, y \in \mathbb{R}. \tag{1}$$

Definition 4 (Convolution or sum of random variables)**.** *If X and Y are independent discrete random variables on $(\Omega, \mathcal{F}, \mathbb{P})$, then $Z = X + Y$ has probability mass function*

$$\mathbb{P}(Z = z) = \sum_{x=-\infty}^{\infty} \mathbb{P}(X = x)\mathbb{P}(Y = z - x), \quad \forall\, z \in \mathbb{R}. \tag{2}$$

Definition 5 (The n-th moment of a random variable)**.** *The n-th moment of a random variable X is denoted as $\mathbb{E}[X^n]$ and it is defined as*

$$\mathbb{E}[X^n] \triangleq \sum_{x \,\in\, \operatorname{Im} X} x^n \cdot \mathbb{P}(X = x). \tag{3}$$

Property 1 (Linearity of expectation and variance)**.**

$$\mathbb{E}[a \cdot X + b \cdot Y] = a \cdot \mathbb{E}[X] + b \cdot \mathbb{E}[Y], \qquad \forall a, b \in \mathbb{R} \tag{4}$$

$$\mathbb{V}[a \cdot X + b \cdot Y] = a^2 \cdot \mathbb{V}[X] + b^2 \cdot \mathbb{V}[Y], \quad \forall a, b \in \mathbb{R} \tag{5}$$

Theorem 1 (Lyapunov Central Limit Theorem [6])**.** *Suppose $\{X_1, \ldots, X_n, \ldots\}$ is a sequence of independent random variables, each with finite expected value μ_i and variance σ_i^2. Let $S_n = \sum_{i=1}^{n} X_i$ and let $v_n^2 = \mathbb{V}[S_n]$. If for some $\delta > 0$, the Lyapunov condition*

$$\lim_{n \to \infty} \frac{1}{v_n^{2+\delta}} \sum_{i=1}^{n} \mathbb{E}\left(|X_i - \mu_i|^{2+\delta}\right) = 0 \tag{6}$$

is satisfied, then $\frac{S_n - \mathbb{E}[S_n]}{\sqrt{\mathbb{V}[S_n]}}$ converges in distribution to a standard normal random variable $\mathcal{N}(0,1)$, as n goes to infinity:

$$\frac{S_n - \mathbb{E}[S_n]}{\sqrt{\mathbb{V}[S_n]}} \xrightarrow{d} \mathcal{N}(0,1). \tag{7}$$

The above definitions are available in the book [24]. Additionally, we often consider normal distribution in this paper. The probability density function of a normally distributed random variable $X \sim \mathcal{N}(\mu, \sigma^2)$ is defined as

$$f_X\left(x \mid \mu, \sigma^2\right) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(x - \mu)^2}{2\sigma^2}\right] \tag{8}$$

TABLE II
LIST OF IMPORTANT SYMBOLS USED IN THE PAPER.

Symbol	Brief explanation
X, Y, Z	Discrete random variables.
$\operatorname{Im} X$	Image of random variable X.
$\mathbb{E}[X]$	Expected value of X.
$\mathbb{E}[X^n]$	The n-th moment of X.
$\mathbb{V}[X]$	Variance of X.
$\Phi(\cdot)$	Cumulative distribution function of normal distribution.
t	Time duration that is analysed.
τ_k	The k-th task in the system.
$\mathbf{T}_k$	Minimum inter-arrival time of τ_k.
$\mathbf{D}_k$	Relative deadline of τ_k.
$\alpha_{k,t}$	Upper bound on the number of releases of instances of τ_k, within an arbitrary time interval of length t.
C_k	Random variable that characterises the execution modes and their respective probabilities for τ_k.
$S_{k,t}$	Upper bound on the probabilistic workload accumulated over an interval of length t, for the jobs with a priority greater than or equal to the priority of τ_k.

and its Cumulative Distribution Function (CDF) is defined as

$$\Phi\left(\frac{x-\mu}{\sigma}\right) = \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{x-\mu}{\sigma\sqrt{2}}\right)\right] \tag{9}$$

where $\operatorname{erf}(\cdot)$ is the error function. In Table II, we provide the list of the most important and frequently used symbols in the paper.

B. Taskset model assumptions and notation

In this paper, we assume a taskset Γ consisting of g sporadic independent tasks such that $\Gamma \triangleq \{\tau_1, \tau_2, \dots, \tau_g\}$. For $1 \leq k \leq g$, each task τ_k releases an infinite number of task instances (jobs), and it is defined with a tuple $\langle C_k, \mathbf{T}_k, \mathbf{D}_k \rangle$, where C_k denotes the discrete random variable that characterises the execution time of the jobs of τ_k, $\mathbf{T}_k$ represents the minimum inter-arrival time between two consecutive instances of τ_k, and $\mathbf{D}_k$ represents the relative deadline of τ_k. C_k is a discrete random variable with a known distribution (same as the model in [10]). C_k can also represent the execution time modes according to the model proposed by von der Brüggen et al. [48] where τ_k has a finite set of execution time modes, and each job of τ_k is executed in one of those distinct modes. Each execution time mode is characterised by the worst-case execution time c such that $c \in \operatorname{Im} C_k$, and also the probability that an instance of τ_k is executed in the respective mode, denoted as $\mathbb{P}(C_k = c)$. The random variables $C_1, C_2, \dots, C_g$ are assumed to be independent. We consider a constrained deadline taskset model, i.e., $\mathbf{D}_k \leq \mathbf{T}_k$ for $1 \leq k \leq g$. It is assumed that the taskset is scheduled according to a preemptive fixed-priority scheduling policy such that the priority of each task instance (job) is identical to the priority of the corresponding task. Task priorities are assumed to be distinct and are represented by task indexes such that for two tasks τ_k and τ_l where $k < l$, τ_k has a higher priority than τ_l. Therefore, τ_1 is the task with the highest priority, while τ_g is the task with the lowest priority in Γ.

Definition 6. *The upper bound $\alpha_{k,t}$ on the number of jobs of τ_k that may affect the response times of lower-priority tasks within a time interval of length t, is defined as $\alpha_{k,t} \triangleq \left\lceil \frac{t+\mathbf{D}_k}{\mathbf{T}_k} \right\rceil$. Note that $\alpha_{k,t} \in \mathbb{N}$ for $t > 0$.*

Property 2. *The value of $\alpha_{k,t}$ can be bounded as*

$$\frac{t}{\mathbf{T}_k} \leq \alpha_{k,t} \leq \frac{t}{\mathbf{T}_k} + 2.$$

The upper bound $\frac{t}{\mathbf{T}_k} + 2$ follows from the system assumptions since we consider constrained deadlines. Since $\mathbf{D}_k \leq \mathbf{T}_k$ then $\frac{\mathbf{D}_k}{\mathbf{T}_k} \leq 1$, and since $\alpha_{k,t} \triangleq \left\lceil \frac{t+\mathbf{D}_k}{\mathbf{T}_k} \right\rceil = \left\lceil \frac{t}{\mathbf{T}_k} + \frac{\mathbf{D}_k}{\mathbf{T}_k} \right\rceil$ it follows that $\alpha_{k,t} \leq \frac{t}{\mathbf{T}_k} + 2$.

Theorem 2 (Corollary 12 in [13])**.** *Given a fully preemptive fixed-priority scheduler, a set of constrained-deadline sporadic tasks, and under the assumption that incomplete jobs are aborted at their deadline, let*

$$S^{\diamond}_{k,t} \triangleq C_k + \sum_{i=1}^{k-1} \sum_{j=1}^{\alpha_{i,t}} C_i. \tag{10}$$

Then, the following inequality holds,

$$\mathbf{DMP}_k \leq \min_{0<t\leq \mathbf{D}_k} \mathbb{P}\left(S^{\diamond}_{k,t} > t\right) \leq \mathbb{P}\left(S^{\diamond}_{k,t} > \mathbf{D}_k\right). \tag{11}$$

In the above inequality, $\mathbf{DMP}_k$ is the maximum possible deadline miss probability among all jobs of τ_k. The inequality states that $\mathbb{P}\left(S^{\diamond}_{k,t} > \mathbf{D}_k\right)$ is a safe upper bound on $\mathbf{DMP}_k$, as shown in [13].

Definition 7. *The random variable $S_{k,t}$ is the upper bound on the probabilistic workload accumulated over an interval of length t, for the jobs with a priority greater than or equal to the priority of τ_k, and it is defined as*

$$S_{k,t} \triangleq \sum_{i=1}^{k} \sum_{j=1}^{\alpha_{i,t}} C_i. \tag{12}$$

Eqs. (10) and (12) implicitly assume the synchronous busy period i.e., an arrival pattern in which τ_k and all tasks with a priority greater than or equal to the priority of τ_k release a job simultaneously at time 0 and continue to release jobs with minimal inter-arrival time [10]. The number of releases is safely upper bounded by $\alpha_{i,t}$ as shown by Chen et al. [13]. Note that $S_{k,t} \succeq S^{\diamond}_{k,t}$ due to the potentially additional sums of C_k, and thus $S_{k,t}$ upper-bounds $S^{\diamond}_{k,t}$ (see Definition 2).

IV. APPLICABILITY OF $S_{k,t}$

Let us now consider the applicability of the $S_{k,t}$ term for deriving the safe distribution approximation and consequently the deadline miss probabilities under the following two job-abortion policies: (*A*.) incomplete jobs aborted at their deadline while t is arbitrarily long (generalising above Theorem 2), and (*B*.) jobs run to completion despite the deadline miss. Note that synchronous busy period is implicitly assumed in both cases.

A. Incomplete Jobs Aborted (JA) at their deadline

Corollary 1. *$\mathbb{P}(S_{k,t} > t)$ is a safe upper bound on the probability that the accumulated workload of τ_k within an arbitrary t-long time interval exceeds t time units.*

Proof. Let us assume an additional virtual task τ_*, with $C_* = 0$ and $\mathbf{D}_* = \mathbf{T}_* = t$, whose priority is lower than the priority of τ_k and greater than the priority of τ_{k+1}. $\mathbb{P}\left(S^{\diamond}_{*,t} > t\right)$ is a safe upper bound, as follows from Theorem 2. Note that $S^{\diamond}_{*,t} = S_{k,t}$ as $C_* = 0$ thus $\mathbb{P}(S_{k,t} > t)$ is also a safe upper bound, and the corollary holds. □

The above implies that $\mathbb{P}(S_{k,t} \leq t)$ is also a safe bound on the CDF of the cumulative distribution of τ_k within an arbitrary t-long time interval. This follows from the fact that $\mathbb{P}(S_{k,t} \leq t) = 1 - \mathbb{P}(S_{k,t} > t)$ and since $\mathbb{P}(S_{k,t} > t)$ is a safe upper bound on the cumulative distribution of τ_k exceeding t, $\mathbb{P}(S_{k,t} \leq t)$ is a safe lower bound on the cumulative distribution of τ_k being less than or equal to t (see Definition 2).

B. Jobs run to Completion (JC)

Given a taskset Γ_{JC} (with the JC policy), a job of $\tau_k \in \Gamma_{\text{JC}}$ may be delayed due to the previous job of τ_k that is executing over its assigned deadline. For this reason, in the context of the JC policy, we need to restrict the assumptions compared to the JA policy for $S_{k,t}$ to be a safe bound. We first observe the following property of the JC policy.

Property 3. *In the taskset Γ_{JC}, deadlines do not affect scheduling decisions at any time point, since the jobs are not aborted.*

We prove the following theorem:

Theorem 3. *If no task has an accumulated workload at time 0 (e.g., at system startup), then $\mathbb{P}(S_{k,t} > t)$ is a safe upper bound on the probability that the accumulated workload of $\tau_k \in \Gamma_{\text{JC}}$ exceeds t time units.*

Proof. Taking into account the time interval $[0, t]$, we can transform any taskset Γ_{JC} (with the JC policy) into a taskset Γ_{JA} (with the JA policy) that has the same scheduling properties. Such a mapping is defined by the map $m : \Gamma_{\text{JC}} \times \mathbb{N}_0 \to \Gamma_{\text{JA}}$, that maps jobs of Γ_{JC} into distinctive tasks of Γ_{JA}. More specifically, let us assume that a job $\tau_{i,j}$ of the task $\tau_i \in \Gamma_{\text{JC}}$ that can be released within $[0, t)$ is a distinctive task $\tau_{m(i,j)}$ in Γ_{JA}, where $j \in [1, \ldots, \alpha_{i,t}]$.

The mapping is constructed such that $\forall \tau_{m(i,j)} \in \Gamma_{\text{JA}}$, $\mathbf{T}_{m(i,j)} = t$, which preserves the property of $\tau_{i,j} \in \Gamma_{\text{JC}}$ being released at most once within $[0, t]$. Then, we also construct that $\mathbf{D}_{m(i,j)} = t$ which implies that none of these tasks can be aborted due to a deadline miss. Next, to preserve the arrival properties of the jobs in Γ_{JC}, we assume that the release time $r_{m(i,j+1)}$ of each task $\tau_{m(i,j+1)} \in \Gamma_{\text{JA}}$ is conditioned on the release time $r_{m(i,j)}$ of $\tau_{m(i,j)} \in \Gamma_{\text{JA}}$ such that $r_{m(i,j+1)} \geq r_{m(i,j)} + \mathbf{T}_i$. This defines the mapping between Γ_{JC} and Γ_{JA}.

With this mapping, in Γ_{JA} we preserved the same workloads, with the same execution distributions, priorities and the minimum inter-arrival times as in Γ_{JC}. Due to Property 3 that holds for Γ_{JC} and the fact that tasks in Γ_{JA} cannot be aborted due to a deadline miss within t time units, the two tasksets enjoy the same scheduling properties within the interval $[0, t]$.

If we assume a virtual task $\tau_* \in \Gamma_{\text{JA}}$ whose priority is lower than the priority of $\tau_k \in \Gamma_{\text{JC}}$ and higher than the priority of $\tau_{k+1} \in \Gamma_{\text{JC}}$, while $C_* = 0$ and $\mathbf{T}_* = t$, then due to the same scheduling properties it holds that $S_{k,t} = S_{*,t}$.

From Corollary 1, for $\tau_* \in \Gamma_{\text{JA}}$ it is true that $\mathbb{P}(S_{*,t} > t)$ is a safe upper bound, then due to the same scheduling properties the same applies to $\mathbb{P}(S_{k,t} > t)$ in Γ_{JC}. □

Corollary 2. *Under the assumptions of Theorem 3, $S_{k,t}$ is a safe bound on the accumulated workload of τ_k within t time units, under the JC policy.*

Proof. From Theorem 3 it follows that $\mathbb{P}(S_{k,t} > t)$ is an upper bound on the accumulated workload of τ_k being greater than t. Since $\mathbb{P}(S_{k,t} \leq t) = 1 - \mathbb{P}(S_{k,t} > t)$ it follows that $\mathbb{P}(S_{k,t} \leq t)$ is a safe lower bound on the accumulated workload of τ_k being less than or equal to t. Due to Definition 2, the corollary holds. □

Note that Theorem 2 does not hold for the JC policy, so any consecutive job of τ_k may experience the maximum deadline-miss probability, which may be computationally intractable. In the following section, we consider the asymptotic behaviour of $S_{k,t}$ while in Section VI we propose an efficient and safe analytical approximation of $S_{k,t}$ that enables analysis within multiple hyperperiods.

V. Asymptotic behaviour of $S_{k,t}$

In this section, we prove the asymptotic behaviour of the cumulative probability distribution of τ_k as $t \to \infty$.

Let us first notice that $S_{k,t}$ (Eq. (12)) is a summation of multiple random variables (C_i) that increases in the number of addends as t grows. Recalling Theorem 1, as the number of addends grows, a certain transformation of the sum converges in distribution to a standard normal distribution.

In Fig. 1, we show the simple example in which the random variable X has the opposite properties compared to the normal distribution, i.e., large values in the tails and small values around the mean. However, already after twenty additions,

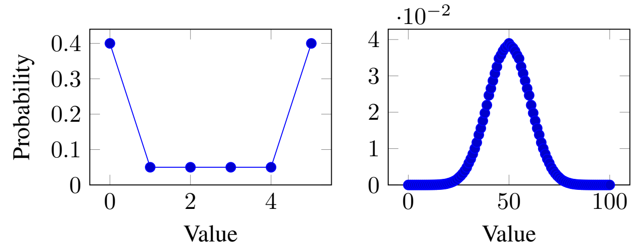

(a) X (b) $\sum_{k=1}^{20} X$

Fig. 1. (From [32]) Consecutive additions of the same random variable. The points are connected in order to show the resemblance (or lack of it) to the continuous normal distribution.

the sum resembles the normal distribution. This implication of the central limit theorem is important due to its potential to be utilised in the real-time system analysis since it is often the case that such analysis accounts for the iteratively enlarged probabilistic workload (e.g., accounting job releases in $S_{k,t}$).

Since the distributions of execution times are not necessarily identical between different tasks, we need to show that the Lyapunov CLT holds for the assumed task model. We start by proving the following lemmas.

Lemma 1. *The third centred moment* $\mathbb{E}\left[(C_i - \mathbb{E}[C_i])^3\right]$ *of the execution time distribution* C_i *(*$1 \leq i \leq q$*) is a finite value.*

Proof. It follows from the system model assumption since each C_i is defined as a discrete random variable with the minimum possible value of 0 and the maximum possible value of $c \in \operatorname{Im} C_i$. Under these assumptions and according to Definition 5 and the definition of expected value, the lemma is valid. □

Lemma 2. *The variance* $\mathbb{V}[C_i]$ *is a finite value.*

Proof. Similarly to the proof of Lemma 1, it follows from the system model assumptions and the definition of variance. □

Now, we prove that Lyapunov's condition holds for $S_{k,t}$ in the following adjusted form for $t \to \infty$, since the value of t dictates the number of summed execution time distributions in $S_{k,t}$.

Proposition 1. *For* $\delta = 1$*, and* $v_{k,t}^2 = \sum_{i=1}^{k} \sum_{j=1}^{\alpha_{i,t}} \mathbb{V}[C_i]$*,*

$$\lim_{t\to\infty} \frac{1}{v_{k,t}^{2+1}} \sum_{i=1}^{k} \sum_{j=1}^{\alpha_{i,t}} \mathbb{E}\left[|C_i - \mathbb{E}[C_i]|^{2+1}\right] = 0. \tag{13}$$

Proof. Let $Q \triangleq \max_{i=1}^{k} \mathbb{E}\left[|C_i - \mathbb{E}[C_i]|^3\right]$, and $W \triangleq \min_{i=1}^{k} \mathbb{V}[C_i]$. Then

$$\lim_{t\to\infty} \frac{\sum_{i=1}^{k} \sum_{j=1}^{\alpha_{i,t}} \mathbb{E}\left[|C_i - \mathbb{E}[C_i]|^{2+1}\right]}{\left(\sum_{i=1}^{k} \sum_{j=1}^{\alpha_{i,t}} \mathbb{V}[C_i]\right)^{3/2}} \leq$$

$$\lim_{t\to\infty} \frac{\sum_{i=1}^{k} \sum_{j=1}^{\alpha_{i,t}} Q}{\left(\sum_{i=1}^{k} \sum_{j=1}^{\alpha_{i,t}} W\right)^{3/2}} \leq \lim_{t\to\infty} \frac{\sum_{i=1}^{k} \left(\frac{t}{\mathbf{T}_i} + 2\right) \cdot Q}{\left(\sum_{i=1}^{k} \frac{t}{\mathbf{T}_i} \cdot W\right)^{3/2}} \leq$$

$$\lim_{t\to\infty} \frac{t \cdot \sum_{i=1}^{k} \frac{\left(Q + \frac{2 \cdot \mathbf{T}_i \cdot Q}{t}\right)}{\mathbf{T}_i}}{t^{3/2} \cdot \left(\sum_{i=1}^{k} \frac{W}{\mathbf{T}_i}\right)^{3/2}} \leq \lim_{t\to\infty} \frac{t \cdot Q^*}{t^{3/2} \cdot (W^*)^{3/2}} = 0$$

where $Q^* \triangleq \sum_{i=1}^{k} \frac{Q}{\mathbf{T}_i}$, and $W^* \triangleq \sum_{i=1}^{k} \frac{W}{\mathbf{T}_i}$, and Property 2 was used. This proves that the Lyapunov condition holds. □

Now we prove the following statement about the CDF of $S_{k,t}$.

Theorem 4. *As* $t \to \infty$, $\mathbb{P}(S_{k,t} \leq x) \to \Phi\left(\frac{x-\mu}{\sqrt{\sigma^2}}\right)$ *where* $\mu = \sum_{i=1}^{k} \alpha_{i,t} \cdot \mathbb{E}[C_i]$ *and* $\sigma^2 = \sum_{i=1}^{k} \alpha_{i,t}^2 \cdot \mathbb{V}[C_i]$.

Proof. From Proposition 1, the Lyapunov CLT holds

$$\frac{S_{k,t} - \mathbb{E}[S_{k,t}]}{\sqrt{\mathbb{V}[S_{k,t}]}} \xrightarrow{d} \mathcal{N}(0,1) \text{ which implies}$$

$$\mathbb{P}\left(\frac{S_{k,t} - \mathbb{E}[S_{k,t}]}{\sqrt{\mathbb{V}[S_{k,t}]}} \leq x\right) \to \Phi\left(\frac{x-0}{1}\right)$$

$$\mathbb{P}\left(S_{k,t} \leq x \cdot \sqrt{\mathbb{V}[S_{k,t}]} + \mathbb{E}[S_{k,t}]\right) \to \Phi(x).$$

Let $a = x \cdot \sqrt{\mathbb{V}[S_{k,t}]} + \mathbb{E}[S_{k,t}]$, then $x = \frac{a - \mathbb{E}[S_{k,t}]}{\sqrt{\mathbb{V}[S_{k,t}]}}$, and

$$\mathbb{P}(S_{k,t} \leq a) \to \Phi\left(\frac{a - \mathbb{E}[S_{k,t}]}{\sqrt{\mathbb{V}[S_{k,t}]}}\right) = \Phi\left(\frac{a-\mu}{\sqrt{\sigma^2}}\right),$$

where $\mathbb{E}[S_{k,t}] = \mu$ and $\mathbb{V}[S_{k,t}] = \sigma^2$ are true statements due to Property 1. □

Similarly, we can derive the asymptotic behaviour of the probability that $S_{k,t}$ exceeds the time point x as $t \to \infty$.

Corollary 3. *As* $t \to \infty$, $\mathbb{P}(S_{k,t} > x) \to 1 - \Phi\left(\frac{x-\mu}{\sqrt{\sigma^2}}\right)$ *where* $\mu = \sum_{i=1}^{k} \alpha_{i,t} \cdot \mathbb{E}[C_i]$ *and* $\sigma^2 = \sum_{i=1}^{k} \alpha_{i,t}^2 \cdot \mathbb{V}[C_i]$.

Proof. Follows from Theorem 4, as $\mathbb{P}(S_{k,t} > x) = 1 - \mathbb{P}(S_{k,t} \leq x)$, analogously holding for normal distribution. □

Using the corollary, we conclude this section by formulating the deadline-miss probability approximation that shows the convergence of the deadline-miss probability of a task as $t \to \infty$. This can be very important for tasksets with JC policy, since, as pointed out by Davis and Cucu-Grosjean [18] a task running on a JC policy may experience a backlog and thus its deadline-miss probability needs to be analysed over multiple, possibly infinitely many hyperperiods.

Corollary 4 (Asymptotic deadline miss probability)**.**

$$\lim_{t\to\infty} \mathbb{P}(S_{k,t} > t) \leq 1 - \Phi\left(\frac{1 - \sum_{i=1}^{k} \frac{\mathbb{E}[C_i]}{\mathbf{T}_i}}{\sqrt{\sum_{i=1}^{k} \frac{\mathbb{V}[C_i]}{(\mathbf{T}_i)^2}}}\right) \tag{14}$$

Proof. From Corollary 3, for $x = t$ it holds that

$$\lim_{t\to\infty} \mathbb{P}(S_{k,t} > t) = 1 - \Phi\left(\frac{t-\mu}{\sqrt{\sigma^2}}\right)$$

with $\mu = \sum_{i=1}^{k} \alpha_{i,t} \cdot \mathbb{E}[C_i]$ and $\sigma^2 = \sum_{i=1}^{k} \alpha_{i,t}^2 \cdot \mathbb{V}[C_i]$. If the argument of $\Phi(\cdot)$ converges to a finite value, then the asymptotic deadline miss probability will be upper bounded by a probability strictly less than 1. Therefore, we can analyze the convergence of the argument of $\Phi(\cdot)$:

$$\lim_{t\to\infty} \frac{t - \sum_{i=1}^{k} \alpha_{i,t} \cdot \mathbb{E}[C_i]}{\sqrt{\sum_{i=1}^{k} \alpha_{i,t}^2 \cdot \mathbb{V}[C_i]}} \leq \lim_{t\to\infty} \frac{t - \sum_{i=1}^{k} \left(\frac{t}{\mathbf{T}_i} + 2\right) \cdot \mathbb{E}[C_i]}{\sqrt{\sum_{i=1}^{k} \left(\frac{t}{\mathbf{T}_i}\right)^2 \cdot \mathbb{V}[C_i]}}$$

$$= \lim_{t\to\infty} \frac{t \cdot \left(1 - \sum_{i=1}^{k} \left(\frac{\mathbb{E}[C_i]}{\mathbf{T}_i}\right)\right)}{t \cdot \sqrt{\sum_{i=1}^{k} \frac{\mathbb{V}[C_i]}{(\mathbf{T}_i)^2}}} = \frac{1 - \sum_{i=1}^{k} \frac{\mathbb{E}[C_i]}{\mathbf{T}_i}}{\sqrt{\sum_{i=1}^{k} \frac{\mathbb{V}[C_i]}{(\mathbf{T}_i)^2}}},$$

where Property 2 was used to compute the upper bound. This concludes the proof. □

VI. Approximating the Entire Cumulative Distribution

In this section, we solve Problem 2: *How to efficiently, accurately, and safely, approximate a probability distribution (e.g., of an execution workload or a response-time) whose exact computation is intractable?*

To solve the above problem, we build on the results of the previous section, knowing that the sum of random variables will converge to a normal distribution. After Aleksandr Lyapunov proved Theorem 1, mathematicians were interested in quantifying the rate of convergence given an arbitrarily long sum, and the mathematical contributions to that problem are used in this section.

Exemplification: Let us provide an intuitive example of what is stated above. In Fig. 2, we illustrate a CDF of $S_{k,t}$, denoted with $\mathbb{P}(S_{k,t} \le x)$. For a large value of t and (or) complex task parameters, computing the exact value of $S_{k,t}$ is very costly or even not possible on commodity platforms. Along such CDF we illustrate the CDF of the normal distribution $\Phi\left(\frac{x-\mu}{\sigma}\right)$ to which the CDF of $S_{k,t}$ converges, as t increases. This property is proved in Theorem 4, and for a larger value of t than the one depicted in the figure, the resemblance between the two functions would be even more obvious. However, the relevant question is, if we know the value of t, can we say something more about the rate of convergence of the normal distribution? Do we know the accuracy of such approximation?

The answer is yes, and to approximate the cumulative distribution of the sum of random variables, we use the Berry-Esseen theorem, which states the following inequality.

Theorem 5 (Berry-Esseen inequality, proposed in [4], [22])**.** *Let $X_1, \ldots, X_n$ be independent random variables such that*

$$1 \le i \le n, \quad \mathbb{E}[X_i] = 0, \quad \mathbb{E}\left[X_i^2\right] > 0, \quad \mathbb{E}\left[|X_i|^3\right] < \infty,$$

and let $\sigma_i^2 = \mathbb{E}\left[X_i^2\right]$, $\rho_i = \mathbb{E}\left[|X_i|^3\right]$, and

$$S_n \triangleq \frac{X_1 + X_2 + \cdots + X_n}{\sqrt{\sigma_1^2 + \sigma_2^2 + \cdots + \sigma_n^2}},$$

where S_n is the normalized n-th partial sum. Let F_n be the CDF of S_n, and Φ be the CDF of the standard normal distribution. Then, for any $n \in \mathbb{N}$,

$$\sup_{x \in \mathbb{R}} |F_n(x) - \Phi(x)| \le A \cdot \psi,$$

where A is a positive constant, and $\psi \triangleq \sum_{i=1}^{n} \rho_i \cdot \left(\sum_{i=1}^{n} \sigma_i^2\right)^{-3/2}$.

The Berry-Esseen theorem essentially states that for any $x \in \mathbb{R}$, the supremum of the difference between the cumulative distribution $F_n(x)$, of the n-th partial normalised sum S_n, and the standard normal distribution $\Phi(x)$, never exceeds $A \cdot \psi$. Intuitively, in Fig. 2 we are aware of the convergence property, but we cannot quantify it. However, using the Berry-Esseen inequality, in the following proofs we will get to the point where the rate of convergence is well known (see Fig. 3). In Fig. 3, we are able to provide the lower and the upper bound on the possible values of the CDF of $S_{k,t}$, and because of that we do not need to compute $S_{k,t}$ with costly convolutions. Instead, we can use the knowledge about how closely $S_{k,t}$ resembles the CDF of the corresponding normal distribution, i.e. $\Phi\left(\frac{x-\mu}{\sigma}\right)$. Comparing the two figures, one may notice that adding and subtracting the term $A \cdot \psi$ plays a major role in quantifying the bounds. The tight value of the absolute positive constant A was an important topic in the mathematics community, and it is even up to this day. To the best knowledge of the authors, the tightest bound at the moment of writing is the one proved by Irina Shevtsova [39], [40], and Theorem 5 holds safely for the value $A = 0.5583$. **Note that in the following logical statements and proofs** we assume the values of A and ψ as defined in Theorem 5.

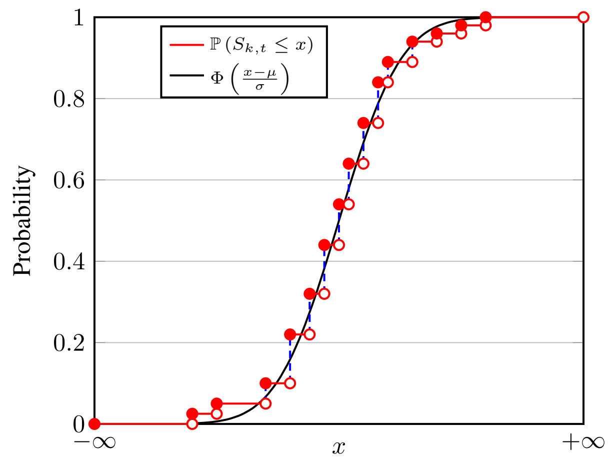

Fig. 2. Exact CDF of $S_{k,t}$ (red, step function) compared to the normal distribution to which the CDF converges (black, continuous) as t increases.

Note that in Theorem 5, there is an assumption that $\mathbb{E}[X_i] = 0$. We prove that in a more general case when $\mathbb{E}[X_i]$ is not necessarily equal to zero, the following holds.

Corollary 5. *Let $Y_1, \ldots, Y_n$ be random variables such that*

$$1 \le i \le n, \quad \mathbb{E}[Y_i] \in \mathbb{R}, \quad \mathbb{E}\left[Y_i^2\right] > 0, \quad \mathbb{E}\left[|Y_i|^3\right] < \infty.$$

Then, for any $x \in \mathbb{R}$, the supremum of the difference between the CDF of the sum of random variables $Y_1, \ldots, Y_n$ at x and the CDF of the general normal distribution defined as $\mathcal{N}(\sum_{i=1}^n \mathbb{E}[Y_i], \sqrt{\sum_{i=1}^n \sigma_i^2})$ is never greater than $A \cdot \psi$. Formally defined, it holds that

$$\sup_{x \in \mathbb{R}} \left| \mathbb{P}\left(\sum_{i=1}^{n} Y_i \le x\right) - \Phi\left(\frac{x - \sum_{i=1}^n \mathbb{E}[Y_i]}{\sqrt{\sum_{i=1}^n \sigma_i^2}}\right) \right| \le A \cdot \psi, \tag{15}$$

where $\sigma_i^2 = \mathbb{E}\left[(Y_i - \mathbb{E}[Y_i])^2\right] = \mathbb{V}[Y_i]$.

Proof. Let $X_i = Y_i - \mathbb{E}[Y_i]$. Then, the following holds due to linearity of expectation and the initial assumptions,

$$\mathbb{E}[X_i] = 0, \quad \mathbb{E}\left[X_i^2\right] > 0, \quad \mathbb{E}\left[|X_i|^3\right] < \infty.$$

Furthermore, the Berry-Esseen theorem holds as follows

$$\sup_{x\in\mathbb{R}}\left|\mathbb{P}\left(\frac{\sum_{i=1}^{n}X_i}{\sqrt{\sum_{i=1}^{n}\sigma_i^2}}\leq x\right)-\Phi(x)\right|\leq A\cdot\psi, \tag{16}$$

Substituting X_i with $Y_i-\mathbb{E}[Y_i]$ yields

$$\sup_{x\in\mathbb{R}}\left|\mathbb{P}\left(\frac{\sum_{i=1}^{n}(Y_i-\mathbb{E}[Y_i])}{\sqrt{\sum_{i=1}^{n}\sigma_i^2}}\leq x\right)-\Phi(x)\right|\leq A\cdot\psi, \tag{17}$$

Considering the inner inequality of the CDF function, we can multiply both sides with $\sqrt{\sum_{i=1}^{n}\sigma_i^2}$ thus deriving

$$\sup_{x\in\mathbb{R}}\left|\mathbb{P}\left(\sum_{i=1}^{n}(Y_i-\mathbb{E}[Y_i]))\leq x\cdot\sqrt{\sum_{i=1}^{n}\sigma_i^2}\right)-\Phi(x)\right|\leq A\cdot\psi. \tag{18}$$

Then, by adding $\sum_{i=1}^{n}\mathbb{E}[Y_i]$ to both sides of the inequality, the following inequality is obtained

$$\sup_{x\in\mathbb{R}}\left|\mathbb{P}\left(\sum_{i=1}^{n}Y_i\leq x\cdot\sqrt{\sum_{i=1}^{n}\sigma_i^2}+\sum_{i=1}^{n}\mathbb{E}[Y_i]\right)-\Phi(x)\right|\leq A\cdot\psi. \tag{19}$$

Let $a=x\cdot\sqrt{\sum_{i=1}^{n}\sigma_i^2}+\sum_{i=1}^{n}\mathbb{E}[Y_i]$, then,

$x=\dfrac{a-\sum_{i=1}^{n}\mathbb{E}[Y_i]}{\sqrt{\sum_{i=1}^{n}\sigma_i^2}}$, from which follows that

$$\sup_{a\in\mathbb{R}}\left|\mathbb{P}\left(\sum_{i=1}^{n}Y_i\leq a\right)-\Phi\left(\frac{a-\sum_{i=1}^{n}\mathbb{E}[Y_i]}{\sqrt{\sum_{i=1}^{n}\sigma_i^2}}\right)\right|\leq A\cdot\psi. \tag{20}$$

□

Next, we approximate only the cumulative distribution of the sum $S_n=\sum_{i=1}^{n}Y_i$ of random variables, and from Corollary 5 it follows that for any $x\in\mathbb{R}$,

$$\Phi\left(\frac{x-\mu}{\sigma}\right)-A\cdot\psi\leq\mathbb{P}(S_n\leq x)\leq\Phi\left(\frac{x-\mu}{\sigma}\right)+A\cdot\psi,$$
$$\text{where}\quad \mu=\sum_{i=1}^{n}\mathbb{E}[Y_i]\quad\text{and}\quad\sigma=\sqrt{\sum_{i=1}^{n}\mathbb{E}\left[(Y_i-\mathbb{E}[Y_i])^2\right]}. \tag{21}$$

Considering the real-time system model assumption, now we can apply the Berry-Esseen theorem to bound the cumulative distribution function of $S_{k,t}$ as follows.

Theorem 6 (Berry-Esseen theorem applied to $S_{k,t}$)**.**

$$\Phi\left(\frac{x-\mu}{\sigma}\right)-A\cdot\psi\leq\mathbb{P}(S_{k,t}\leq x)\leq\Phi\left(\frac{x-\mu}{\sigma}\right)+A\cdot\psi,$$
$$\textit{where }\mu=\sum_{i=1}^{k}\alpha_{i,t}\cdot\mathbb{E}[C_i],\quad\sigma=\sqrt{\sum_{i=1}^{k}\alpha_{i,t}^2\cdot\mathbb{V}[C_i]}$$
$$\rho_i=|\alpha_{i,t}\cdot C_i-\mathbb{E}[\alpha_{i,t}\cdot C_i]|^3=|\alpha_{i,t}\cdot(C_i-\mathbb{E}[C_i])|^3$$
$$\psi=\left(\sum_{i=1}^{k}\alpha_{i,t}^2\cdot\mathbb{V}[C_i]\right)^{-3/2}\cdot\sum_{i=1}^{k}\rho_i. \tag{22}$$

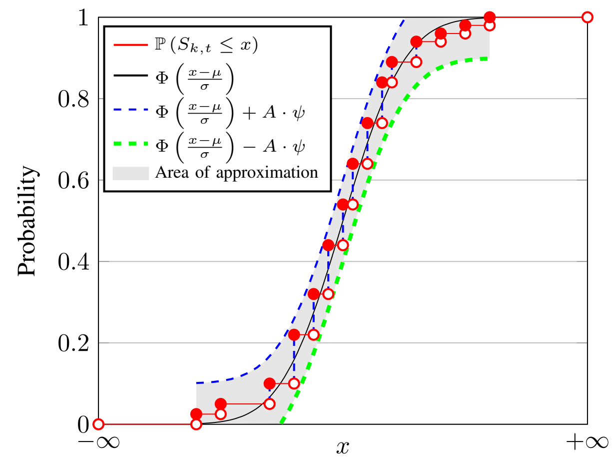

Fig. 3. CDF of the exact $S_{k,t}$ compared to the CDF bounds derived from the normal distribution in Theorem 6. The upper bound is blue, while the lower bound is light green. The red step function is the actual CDF of $S_{k,t}$.

Proof. Follows directly from Corollary 5 and Property 1. □

The benefit of Theorem 6 is that it approximates the entire cumulative distribution function of $S_{k,t}$ from both sides (see Fig. 3) using the CDF of the generalised normal distribution to which it converges. In Fig. 3 we illustrate the most important terms of the theorem. Red step function represents the potential CDF of $S_{k,t}$. The black line is the normal CDF approximation that has the mean and standard deviation identical to the one of $S_{k,t}$, while the blue and green lines represent the upper and lower bound on potential CDF values for $S_{k,t}$. Gray area represents all the potential values for CDF of $S_{k,t}$ (assuming that one cannot compute the actual $F_{S_{k,t}}$). Considering the normal CDF approximation, it can be efficiently computed using Eq. (9).

Theorem 7 (Safe approximation of $S_{k,t}$)**.** *To safely approximate $S_{k,t}$ we define random variable Z,*

$$Z(x)=\begin{cases}1 & x\geq\sum_{i=1}^{k}\alpha_{i,t}\cdot\max(\operatorname{Im}C_i)\\ 0 & \Phi\left(\frac{x-\mu}{\sigma}\right)-A\cdot\psi<0\\ \Phi\left(\frac{x-\mu}{\sigma}\right)-A\cdot\psi & \textit{otherwise}\end{cases}$$

for which holds that

$$Z\succeq S_{k,t}$$

Proof. Directly follows from Theorem 6 and the fact that

$$\mathbb{P}\left(S_{k,t}\leq\sum_{i=1}^{k}\alpha_{i,t}\cdot\min(\operatorname{Im}C_i)\right)=0$$

and

$$\mathbb{P}\left(S_{k,t}\leq\sum_{i=1}^{k}\alpha_{i,t}\cdot\max(\operatorname{Im}C_i)\right)=1.$$

□

VII. Approximating the Value with a Predefined Exceedance Probability

In this section, we solve Problem 3: *How to efficiently and accurately derive the least value x (e.g., time point within some interval), such that the probability that a random variable X (e.g., of execution workload) precedes x is greater than or equal to some probability threshold p.*

For this purpose, we consider the quantile function $Q_X(p)$, defined as

$$Q_X(p) \triangleq \inf\left\{x \in \operatorname{Im} X : p \leq F_X(x)\right\}, \quad p \in (0,1). \quad (23)$$

In general terms, finding the result of $Q_X(p)$ for the given distribution X and the predefined threshold p can, for example, be relevant to controlling the deadline-miss probability when assigning the releases of the new workloads. This is the case due to the fact that the result x of $Q_X(p)$ can also represent the time point for which its exceedance probability $\mathbb{P}\left(X \geq x\right) \leq 1-p$, as follows from Eq. (23). This is then analogous to the case where the result x of $Q_X(p)$ represents the earliest time point for which the deadline-miss probability of X is less than or equal to $1-p$.

Exemplification: Let us provide a bit more concrete and easy-to-follow example of the lengthy mathematical description above. Similarly as before, in Fig. 4 we consider the accumulated execution workload $S_{k,t}$ of τ_k. Consider the following problem. The task τ_k is not required to always finish its execution, but there is a requirement that the probability that it finishes its execution is at least 0.2. Additionally, we want to add another workload to the timeline, but without invalidating the above requirement. How to determine the time point t for which the above requirement is satisfied? First of all, there are multiple such points. In the figure, $\forall x > x'$ (green thick line) the defined property will hold. However, if we delay adding the new workload, then we may push the other jobs that may come for execution and possibly increase their deadline-miss probabilities, or even the deadline-miss probability of the newly added task. Thus, it is of interest to add such workload as early as possible, i.e. as soon as the requirement is met.

The earliest time point which satisfies the requirement can be derived exactly by using Eq. (23) and substituting X for $S_{k,t}$. The value of $Q_{S_{k,t}}$ would be x' but it can only be derived after computing the exact CDF of $S_{k,t}$, which can be intractable. Therefore, once more we will try to derive an efficient approximation of $Q_{S_{k,t}}$ for which we know by construction that it is safe. It means that the result of approximation must be larger than x', otherwise, the requirement that the probability of τ_k finishing its execution being at least 0.2 will not be satisfied.

An approximation of quantile $Q_{S_{k,t}}(p)$ can be derived

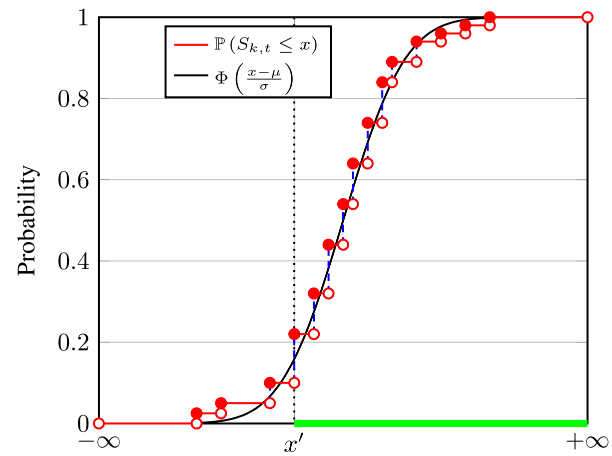

Fig. 4. Exact quantile $Q_{S_{k,t}}(0.2) = x'$ of $S_{k,t}$ for probability equal to 0.2 (vertical line). A thick green line represents all values of $x \in \mathbb{R}$ for which the approximation on the exact quantile is safe.

according to Theorem 4, with the following equation

$$\begin{aligned} &Q^{\text{clt}}_{S_{k,t}}(p) \triangleq F^{-1}_{\mathcal{N}(\mu,\sigma)}(p) = \mu + \sigma\sqrt{2}\operatorname{erf}^{-1}(2p-1), \\ &\text{where } \mu = \sum_{i=1}^{k} \alpha_{i,t} \cdot \mathbb{E}\left[C_i\right] \\ &\text{and } \sigma = \sqrt{\sum_{i=1}^{k} \alpha^2_{i,t} \cdot \mathbb{V}\left[C_i\right]}. \end{aligned} \quad (24)$$

The above equation takes into account that the quantile $Q_{S_{k,t}}(p)$ of random variable $S_{k,t}$ for probability p is approximately equal to the corresponding quantile $F^{-1}_{\mathcal{N}(\mu,\sigma)}(p)$ of normal distribution $\mathcal{N}(\mu,\sigma)$ as $t \to \infty$. This is the case since $F_{S_{i,t}}(x)$ converges to $F_{\mathcal{N}(\mu,\sigma)}(x)$, as we showed in the previous sections. In Eq. (24), the quantile of the normal distribution $\mathcal{N}(\mu,\sigma)$ is expressed in the form of the inverse error function.

To safely approximate the quantile $Q_{S_{k,t}}(p)$ we need to find an approximation $Q^a_{S_{k,t}}(p)$ such that for any $p \in (0,1)$ $Q_{S_{k,t}}(p) \leq Q^a_{S_{k,t}}(p)$.

Theorem 8 (Safe approximation of quantile)**.**

$$Q_{S_{k,t}}(p) \leq F^{-1}_{\mathcal{N}(\mu,\sigma)}(p+A\cdot\psi) = \mu+\sigma\sqrt{2}\operatorname{erf}^{-1}(2(p+A\cdot\psi)-1) \quad (25)$$

Proof. From Theorem 6 it holds that

$$\Phi\left(\frac{x-\mu}{\sigma}\right) - A\cdot\psi \leq F_{S_{i,t}}(x), \quad x \in \operatorname{Im} X$$

Let $f(x) \triangleq \Phi\left(\frac{x-\mu}{\sigma}\right) - A\cdot\psi$. Since f is a monotonically increasing function, then f^{-1} is also monotonically increasing. Since $f(x) \leq F_{S_{i,t}}(x)$ and f^{-1} is monotonically increasing, then

$$f^{-1}(f(x)) \leq f^{-1}(F_{S_{i,t}}(x)) \quad \Rightarrow \quad x \leq f^{-1}(F_{S_{i,t}}(x))$$

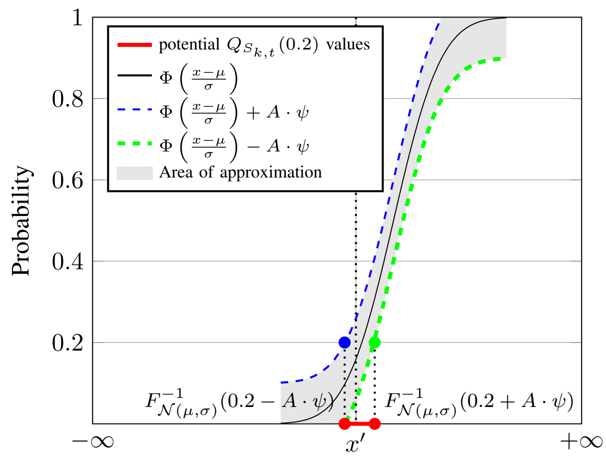

Fig. 5. The exact quantile of $Q_{S_{k,t}}(0.2)$ may be any of the points within the red line. The bounds on the range of the red line are depicted according to Remark 1, using the quantiles of the appropriate normal distributions.

Let $Q_{S_{i,t}}(p)$ be the quantile function of $F_{S_{i,t}}(x)$. Then it holds that $Q_{S_{i,t}}(F_{S_{i,t}}(x)) = x$, yielding

$$Q_{S_{i,t}}(F_{S_{i,t}}(x)) \leq f^{-1}(F_{S_{i,t}}(x)) \text{ which is the same as } Q_{S_{i,t}}(p) \leq f^{-1}(p).$$

Therefore the quantile function is upper bounded by $f^{-1}(p)$ that can be obtained by inverting the function $f(x)$:

$$p = f(x) = \Phi\left(\frac{x-\mu}{\sigma}\right) - A \cdot \psi$$

$$\Phi\left(\frac{x-\mu}{\sigma}\right) = p + A \cdot \psi$$

$$x = \sigma\Phi^{-1}(p + A \cdot \psi) + \mu = F^{-1}_{\mathcal{N}(\mu,\sigma)}(p + A \cdot \psi) = f^{-1}(p)$$

which concludes the proof. □

In general, following from Corollary 5, Theorem 6, and properties of quantile functions, the following remark holds.

Remark 1.

$$F^{-1}_{\mathcal{N}(\mu,\sigma)}(p - A \cdot \psi) \leq Q_{S_{i,t}}(p) \leq F^{-1}_{\mathcal{N}(\mu,\sigma)}(p + A \cdot \psi)$$

We depicted the essential terms of the remark in Fig. 5. The red line represents the potential range of the quantile for probability equal to 0.2. The projection of the green dot is the upper bound on the quantile, while the one of the blue dot is the lower bound on the quantile. Therefore, if one wants to more accurately approximate quantile $Q(S_{i,t}, p)$ it is enough to search it in the range given in the remark. One potential solution would be to use the methods of Problem 1, together with the binary search on the interval defined under the remark. However, this line of research remains for the future work.

VIII. Evaluation

The evaluation is organised into two parts, separated in the following sections; (A.) Approximation of the entire cumulative distribution of $S_{q,t}$, and (B.) Approximation of the quantile function of $S_{q,t}$.

Hardware and software configuration. We used a Macbook Pro with 2,6 GHz 6-Core Intel Core i7 CPU, and 16 GB of RAM memory. All equations are implemented in MATLAB, using Advanpix Multiprecision Computing Toolbox [29].

A. Approximation of the entire cumulative distribution of $S_{q,t}$

Goal of the evaluation and the evaluated entities. In this evaluation, we compared the approximation of the CDF of $S_{q,t}$ computed using Theorem 6 (labeled with *BE*), with the circular-convolution method (labeled with *CC*, from [32]) that computes the actual CDF of $S_{q,t}$. In the experiment, we compared the computation time and the memory footprint of the respective methods.

Experiment Setup. To approximate the CDF of $S_{q,t}$, we generated 1000 tasksets per each point in the graphs shown, considering the taskset sizes $5, 10, \ldots, 50$. Respective task utilisations were generated using UUniFast [8] with the utilisation of the taskset set to 0.7. The task periods were randomly generated from the log-uniform distribution with a range from 10 to 1000 ms. The worst-case execution time value $\max\{\operatorname{Im} C_i\}$ for each task was obtained with $m_i = \max\{\operatorname{Im} C_i\} = T_i \cdot U_i$, where U_i is the utilisation generated using UUnifast for the respective task. The discrete random variable C_i was generated randomly (up to m_i), using the MATLAB function *randi*, which allows random variability in the probabilistic distribution. The generation function was implemented using a library for random vectors, contributed by Roger Stafford [43]. To give a fair advantage to the circular convolution method, the generated timing values were supported by the discrete space where the minimum value is 50μs. Reducing the quanta of the support vector would reduce the efficiency of *CC*.

Experiment Results. In Fig. 6, 7, and 8, we report the average computation time, memory footprint, and the accuracy respectively, for computing the approximation and the exact result of $S_{q,t}$. $S_{q,t}$ is the maximum accumulated probabilistic workload over an interval of length t, for jobs with a priority greater than or equal to the priority of τ_q. τ_q is the task with the lowest priority in the generated taskset. There are three subfigures for the different values of t, which are $\mathbf{T}_q$, $50 \cdot \mathbf{T}_q$, and $100 \cdot \mathbf{T}_q$. An increase in the values of t has a purpose to increase the analysed workloads.

We observe that the proposed approximation *BE* is computationally efficient (see Fig. 8), since even for $t = 100 \cdot \mathbf{T}_q$, the computation time is below 6 ms. On the other hand, the exact calculation with *CC* takes significantly longer to compute, with each increase in the interval under consideration. For $t = 100 \cdot \mathbf{T}_q$ and the size of the task set 50, *CC* took on average 45 seconds to calculate the cumulative distribution function, while with *BE* the approximations were obtained in 5.4 ms at most. The efficiency of the *BE* method is due to the nature of its calculations and its use of the moments of random variables, and expectation-related values that are rather quickly computed using the linearity of expectation.

In Fig. 7, we also reported the average memory footprint, reported in megabytes (MB). We note that the average memory

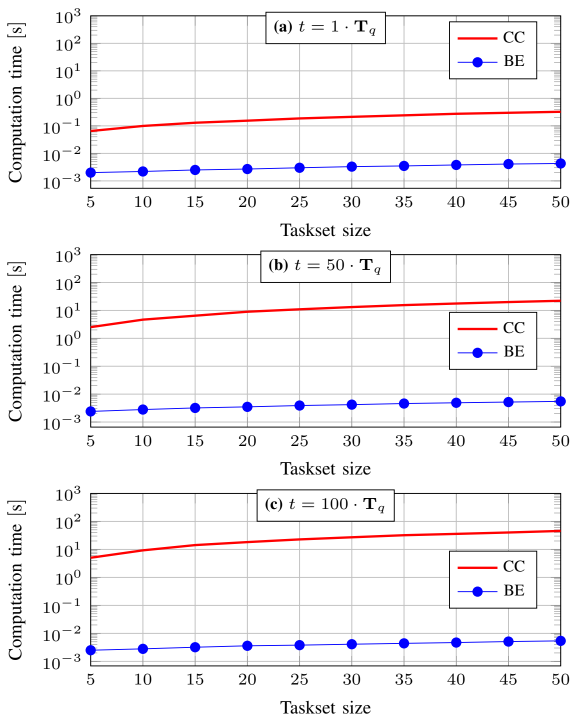

Fig. 6. Average computation time of the CDF approximation of $S_{q,t}$. Shown as a function of a taskset size, for different values of the analysed interval t.

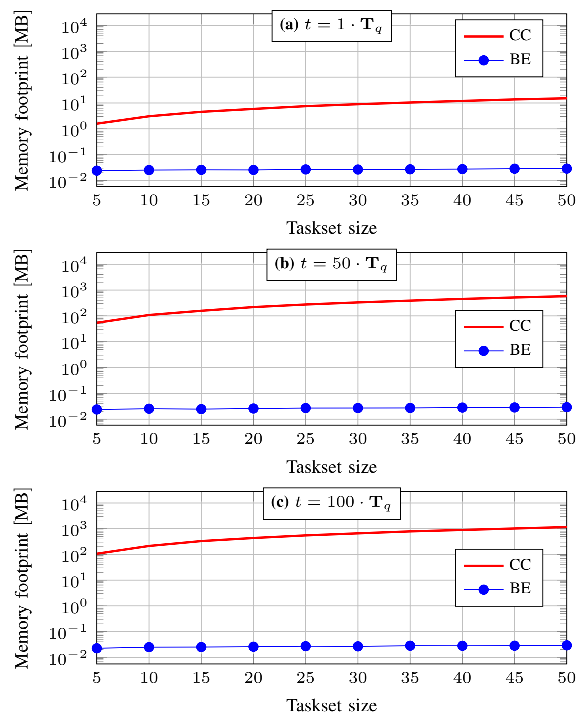

Fig. 7. Average memory footprint of the CDF approximation of $S_{q,t}$. Shown as a function of a taskset size, for different values of the analysed interval t.

footprint of *BE* is significantly better than that of *CC* for each analysed combination. For example (see Fig. 7), for $t = 100 \cdot \mathbf{T}_q$ and the taskset size 50, *CC* on average requires 1157 MB, while *BE* requires only 0.029 MB.

Next, in Fig. 8, we reported the accuracy of *BE*. The accuracy was measured with $\Delta = A \cdot \psi$ from Theorem 6, which is the highest possible deviation of *BE* compared to the exact result of *CC* for any value. Ideally, Δ would have a value of zero, which is not possible considering the finite number of computations. We observe that the accuracy of the approximation *BE* is benefiting from the increase of the taskset size (it is approaching zero). But, we can also observe that it benefits from the increase of t. This is expected due to the nature of the Berry-Esseen theorem and the Lyapunov CLT.

As observed by Davis and Cucu-Grosjean [18] computing the average deadline-miss probability of strictly periodic tasksets requires estimating the backlog distributions and deadline-miss probabilities at all periods within a hyperperiod (p.14 in [18]). In addition, Davis and Cucu-Grosjean also note that for the tasksets with JC policy, *"calculation becomes more complex if the task model permits a backlog of outstanding execution at the end of the hyperperiod"*((p.14 in [18])). Thus, to further investigate the efficiency and accuracy of *BE* we performed the experiment depicted in Fig. 9 (a), where t is equal to the Least Common Multiple (LCM) of the task periods. The average time of *BE* was below 6 milliseconds, while *CC* often resulted in memory allocation problems or did not provide a result within a few hours of observing the experiment progress, so those results are omitted.

The greatest improvement for *BE* in accuracy is observable in Fig. 9 (b), when the LCM was considered as an interval. For the taskset size of 50, the maximum deviation from the exact distribution was just 10^{-16}. This trend continues with each new increase in workload $S_{k,t}$, which is a very positive property when considering real-time systems that are often analysed with respect to iteratively increasing workloads.

B. Approximation of the quantile function of $S_{q,t}$

Goal of the evaluation and the evaluated entities. In this evaluation, we compared the accuracy of the quantile approximations for $Q_{S_{q,t}}(0.8)$ considering three different versions:

- *CLT-based* – A quantile approximation based on the Lyapunov CLT, that is increasingly more accurate as t grows, but possibly unsafe. It is given in Equation 24,
- *upper* – Upper-bound quantile approximation (Remark 1),
- *down* – Lower-bound quantile approximation (Remark 1).

We also compared the execution times needed for deriving the approximations (jointly labelled with *BE*) and the exact quantile result (labelled with *CC*).

Experiment Setup. In order to approximate $Q(S_{q,t}, 0.8)$, we use the same experiment setup as in the previous subsection, but the taskset size is fixed to 5, while we modify t values, given by the following set $\{1 \cdot \mathbf{T}_q, 5 \cdot \mathbf{T}_q, 10 \cdot \mathbf{T}_q, \ldots, 50 \cdot \mathbf{T}_q\}$. The results are presented as normalised average values over the actual value of the quantile (computed with *CC*). Thus, *CC* has a constant value of 100%, while if the approximation is higher, its value will be greater than 100%, and if it is lower, its value will be less than 100%.

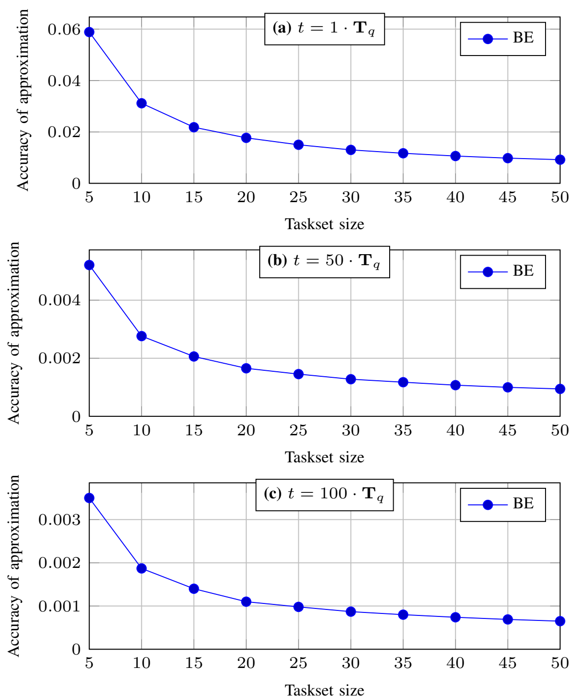

Fig. 8. Average accuracy of approximation Δ. $\Delta = 0$ theoretically means that the approximation is equivalent to the approximated exact CDF.

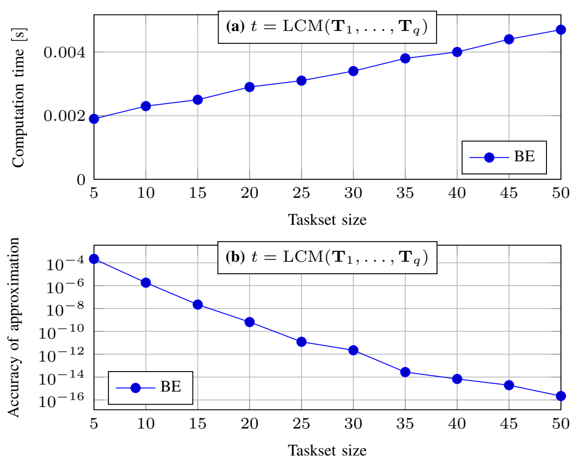

Fig. 9. Average computation time and the average accuracy of the approximation for t equal to the least common multiple of the periods of the analysed tasks.

Experiment Results. In Fig. 10 (a), we observe that the time needed to compute the exact quantile using *CC* is increasing rapidly (10 seconds for $t = 50 \cdot \mathbf{T}_q$), while it is less than 4 miliseconds for any t when *BE* is used.

In Fig. 10 (b), we can observe that the accuracy of each of the three quantile approximations is low for the values $1 \cdot \mathbf{T}_q$ and $5 \cdot \mathbf{T}_q$. This property of the *BE* bounds can also be seen with respect to the CDF approximation, in Fig. 8, for the small values of t and q. We can conclude that the proposed methods complement the state-of-the-art (especially *CC*), providing very efficient approximations (both in computation time and

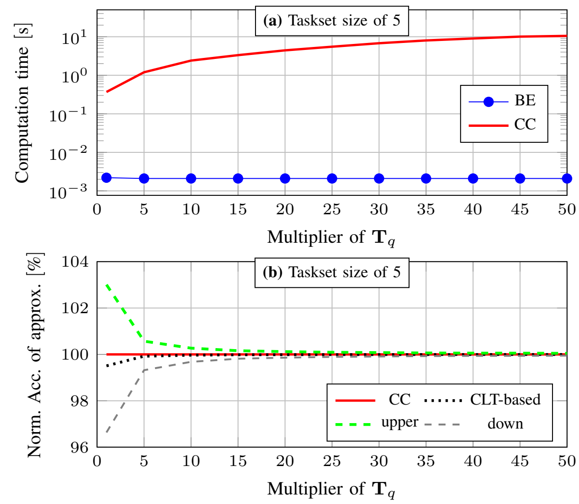

Fig. 10. Average computation time and the average accuracy of the quantile approximation for the taskset size of 5, and $t \in \{20\,\mathbf{T}_q, 40\,\mathbf{T}_q, 60\,\mathbf{T}_q, 80\,\mathbf{T}_q\}$.

space complexity) when the analysed tasksets and considered time intervals increase in size.

IX. Conclusion

In this paper, we addressed the following main problem in the probabilistic analysis of real-time systems. *How to efficiently and accurately derive an analytical approximation of the intractable probabilistic execution workloads*? To solve the above problem, we first applied and proved the Lyapunov central limit theorem considering the accumulated execution distribution of the fixed-priority fully-preemptive tasks, under the synchronous busy period. We considered the job-abortion and job-run-to-completion policies. This proof showed that as time goes to infinity, the accumulated execution distribution (e.g., probabilistic response time) can converge to a normal distribution of certain characteristics. We then proved the bound on the asymptotic behaviour of the deadline-miss probability as time goes to infinity.

To safely approximate the accumulated execution distribution of a task within an arbitrarily long finite time interval, we adjusted the Berry-Esseen theorem to the investigated task model and proposed two safe approximations: 1) for the cumulative distribution function and 2) for the quantile function.

In the evaluation, we generated synthetic tasksets to investigate the computational efficiency and the approximation accuracy compared to the actual results derived with the circular convolution. We showed that the proposed approximations exhibit an extremely low computational cost (below 4 ms for all experiment setups even in the case of analysing LCM intervals), while the accuracy of the approximations is increasingly better upon increasing the number of workloads whose probabilistic interaction is to be approximated. The evaluation suggests that the proposed methods complement the state-of-the-art of analytical approximations.

Acknowledgement

This work was supported by the Swedish Research Council (VR) via the projects "Practical Probabilistic Timing Analysis of Real-Time Systems (PARIS)" and "Pervasive Self-Optimizing Computing Infrastructures (PSI)", and by the Knowledge Foundation (KKS) via the project "Federated Choreography of an Integrated Embedded Systems Software Architecture (FIESTA)". Lastly, this project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 803111).

We thank the anonymous reviewers, our shepherd, Arvind Easwaran in his role as RTSS'22 Program Chair, and the authors of [13] for giving us the opportunity to revise our paper before publication in light of the findings in [13].

References

[1] Neil Audsley, Alan Burns, Mike Richardson, Ken Tindell, and Andy J Wellings. Applying new scheduling theory to static priority pre-emptive scheduling. *Software engineering journal*, 8(5):284–292, 1993.

[2] Sanjoy K Baruah. Optimal utilization bounds for the fixed-priority scheduling of periodic task systems on identical multiprocessors. *IEEE Transactions on Computers*, 53(6):781–784, 2004.

[3] Sergei Bernstein. On a modification of chebyshev's inequality and of the error formula of laplace. *Ann. Sci. Inst. Sav. Ukraine, Sect. Math*, 1(4):38–49, 1924.

[4] Andrew C Berry. The accuracy of the gaussian approximation to the sum of independent variates. *Transactions of the american mathematical society*, 49(1):122–136, 1941.

[5] Luciano Bertini, Julius CB Leite, and Daniel Mossé. Generalized tardiness quantile metric: Distributed dvs for soft real-time web clusters. In *2009 21st Euromicro Conference on Real-Time Systems*, pages 227–236. IEEE, 2009.

[6] Patrick Billingsley. *Convergence of probability measures*. John Wiley & Sons, 2013.

[7] Enrico Bini, Giorgio Buttazzo, and Giuseppe Buttazzo. A hyperbolic bound for the rate monotonic algorithm. In *Proceedings 13th Euromicro Conference on Real-Time Systems*, pages 59–66. IEEE, 2001.

[8] Enrico Bini and Giorgio C Buttazzo. Measuring the performance of schedulability tests. *Real-Time Syst.*, 30(1-2):129–154, 2005.

[9] Enrico Bini, Giorgio C Buttazzo, and Giuseppe M Buttazzo. Rate monotonic analysis: the hyperbolic bound. *IEEE Transactions on Computers*, 52(7):933–942, 2003.

[10] Sergey Bozhko, Georg von der Brüggen, and Björn Brandenburg. Monte carlo response-time analysis. In *IEEE 42nd Real-Time Systems Symposium*, pages 342–355. IEEE, 2021.

[11] Laura Carnevali, Alessandra Melani, Luca Santinelli, and Giuseppe Lipari. Probabilistic deadline miss analysis of real-time systems using regenerative transient analysis. In *Proceedings of the 22nd International Conference on Real-Time Networks and Systems*, pages 299–308, 2014.

[12] Kuan-Hsun Chen and Jian-Jia Chen. Probabilistic schedulability tests for uniprocessor fixed-priority scheduling under soft errors. In *IEEE Int. Symp. Industrial Emb. Syst. (SIES)*, pages 1–8, 2017.

[13] Kuan-Hsun Chen, Mario Günzel, Georg von der Brüggen, and Jian-Jia Chen. Critical instant for probabilistic timing guarantees: Refuted and revisited. In *IEEE 43rd Real-Time Systems Symposium*. IEEE, 2022.

[14] Kuan-Hsun Chen, Niklas Ueter, Georg von der Brüggen, and Jian-Jia Chen. Efficient computation of deadline-miss probability and potential pitfalls. In *Design, Automation & Test in Europe Conf. & Exhibition (DATE)*, pages 896–901, 2019.

[15] Kuan-Hsun Chen, Georg von Der Brüggen, and Jian-Jia Chen. Analysis of deadline miss rates for uniprocessor fixed-priority scheduling. In *IEEE Int. Conf. Emb. and Real-Time Computing Syst. and Applications (RTCSA)*, pages 168–178, 2018.

[16] Herman Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, pages 493–507, 1952.

[17] Robert I Davis, Luca Santinelli, Sebastian Altmeyer, Claire Maiza, and Liliana Cucu-Grosjean. Analysis of probabilistic cache related pre-emption delays. In *2013 25th Euromicro Conference on Real-Time Systems*, pages 168–179. IEEE, 2013.

[18] Robert Ian Davis and Liliana Cucu-Grosjean. A survey of probabilistic schedulability analysis techniques for real-time systems. *LITES: Leibniz Trans. Emb. Syst.*, pages 1–53, 2019.

[19] Robert Ian Davis and Liliana Cucu-Grosjean. A survey of probabilistic timing analysis techniques for real-time systems. *LITES: Leibniz Transactions on Embedded Systems*, pages 1–60, 2019.

[20] José Luis Díaz, Daniel F García, Kanghee Kim, Chang-Gun Lee, L Lo Bello, José María López, Sang Lyul Min, and Orazio Mirabella. Stochastic analysis of periodic real-time systems. In *23rd IEEE Real-Time Systems Symposium, 2002. RTSS 2002.*, pages 289–300. IEEE, 2002.

[21] Jose Luis Diaz, Jose Maria Lopez, Manuel Garcia, Antonio M Campos, Kanghee Kim, and Lucia Lo Bello. Pessimism in the stochastic analysis of real-time systems: Concept and applications. In *IEEE Int. Real-Time Syst. Symp. (RTSS)*, pages 197–207, 2004.

[22] Carl-Gustav Esseen. On the liapunov limit error in the theory of probability. *Ark. Mat. Astr. Fys.*, 28:1–19, 1942.

[23] Mark K Gardner and Jane WS Liu. Analyzing stochastic fixed-priority real-time systems. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 44–58. Springer, 1999.

[24] Geoffrey Grimmett and Dominic Welsh. *Probability: an introduction*. Oxford U. Press, 2014.

[25] Nan Guan and Wang Yi. Fixed-priority multiprocessor scheduling: Critical instant, response time and utilization bound. In *2012 IEEE 26th International Parallel and Distributed Processing Symposium Workshops & PhD Forum*, pages 2470–2473. IEEE, 2012.

[26] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. In *The collected works of Wassily Hoeffding*, pages 409–426. Springer, 1994.

[27] Kanghee Kim, Jose Luis Diaz, Lucia Lo Bello, José María López, Chang-Gun Lee, and Sang Lyul Min. An exact stochastic analysis of priority-driven periodic real-time systems and its approximations. *IEEE Transactions on Computers*, 54(11):1460–1466, 2005.

[28] Chung Laung Liu and James W Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *Journal of the ACM (JACM)*, 20(1):46–61, 1973.

[29] Advanpix LLC. Multiprecision computing toolbox for MATLAB. URL: http://www.advanpix.com/.

[30] José María López, José Luis Díaz, Joaquín Entrialgo, and Daniel García. Stochastic analysis of real-time systems under preemptive priority-driven scheduling. *Real-Time Systems*, 40(2):180–207, 2008.

[31] Filip Marković, Jan Carlson, Radu Dobrin, Bjorn Lisper, and Abhilash Thekkilakattil. Probabilistic response time analysis for fixed preemption point selection. In *2018 IEEE 13th International Symposium on Industrial Embedded Systems (SIES)*, pages 1–10. IEEE, 2018.

[32] Filip Marković, Alessandro Vittorio Papadopoulos, and Thomas Nolte. On the convolution efficiency for probabilistic analysis of real-time systems. In *33rd Euromicro Conference on Real-Time Systems (ECRTS 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.

[33] Dorin Maxim and Liliana Cucu-Grosjean. Response time analysis for fixed-priority tasks with multiple probabilistic parameters. In *2013 IEEE 34th Real-Time Systems Symposium*, pages 224–235. IEEE, 2013.

[34] Dorin Maxim, Mike Houston, Luca Santinelli, Guillem Bernat, Robert I Davis, and Liliana Cucu-Grosjean. Re-sampling for statistical timing analysis of real-time systems. In *Int. Conf. Real-Time and Network Syst. (RTNS)*, pages 111–120, 2012.

[35] Dorin Maxim, Luca Santinelli, and Liliana Cucu-Grosjean. Improved sampling for statistical timing analysis of real-time systems. *Int. Conf. Real-Time and Network Syst. (RTNS)*, pages 17–20, 2010.

[36] Suzana Milutinović, Jaume Abella, Damien Hardy, Eduardo Quiñones, Isabelle Puaut, and Francisco J Cazorla. Speeding up static probabilistic timing analysis. In *Int. Conf. Architecture of Computing Syst. (ARCS)*, pages 236–247, 2015.

[37] Khaled S Refaat and Pierre-Emmanuel Hladik. Efficient stochastic analysis of real-time systems via random sampling. In *Euromicro Conf. Real-Time Syst. (ECRTS)*, pages 175–183, 2010.

[38] Moshe Shaked and J. George Shanthikumar, editors. *Univariate Stochastic Orders*, pages 3–79. Springer New York, New York, NY, 2007.

[39] Irina Shevtsova. On the absolute constants in the berry-esseen type inequalities for identically distributed summands. *arXiv preprint arXiv:1111.6554*, 2011.

[40] Irina Shevtsova. On the absolute constants in the Berry-Esseen inequality and its structural and nonuniform improvements. *Informatika i Ee Primeneniya [Informatics and its Applications]*, 7(1):124–125, 2013.

[41] Michael Short. Improved inequalities for the poisson and binomial distribution and upper tail quantile functions. *International Scholarly Research Notices*, 2013, 2013.

[42] Michael Short and Julián Proenza. Towards efficient probabilistic scheduling guarantees for real-time systems subject to random errors and random bursts of errors. In *2013 25th Euromicro Conference on Real-Time Systems*, pages 259–268. IEEE, 2013.

[43] Roger Stafford. Random Vectors with Fixed Sum (MATLAB Central File Exchange). https://www.mathworks.com/matlabcentral/fileexchange/9700-random-vectors-with-fixed-sum. Accessed: 2022-05.

[44] Bogdan Tanasa, Unmesh D Bordoloi, Petru Eles, and Zebo Peng. Probabilistic timing analysis for the dynamic segment of flexray. In *2013 25th Euromicro Conference on Real-Time Systems*, pages 135–144. IEEE, 2013.

[45] Bogdan Tanasa, Unmesh D Bordoloi, Petru Eles, and Zebo Peng. Probabilistic response time and joint analysis of periodic tasks. In *2015 27th Euromicro Conference on Real-Time Systems*, pages 235–246. IEEE, 2015.

[46] T-S Tia, Zhong Deng, Mallikarjun Shankar, Matthew Storch, Jun Sun, L-C Wu, and JW-S Liu. Probabilistic performance guarantee for real-time tasks with varying computation times. In *Proceedings Real-Time Technology and Applications Symposium*, pages 164–173. IEEE, 1995.

[47] Georg von der Brüggen. *Realistic Scheduling Models and Analyses for Advanced Real-Time Embedded Systems*. PhD thesis, TU Dortmund (Germany), 2019.

[48] Georg von der Brüggen, Nico Piatkowski, Kuan-Hsun Chen, Jian-Jia Chen, and Katharina Morik. Efficiently approximating the probability of deadline misses in real-time systems. In *Euromicro Conf. Real-Time Syst. (ECRTS)*, 2018.

[49] Georg von der Brüggen, Nico Piatkowski, Kuan-Hsun Chen, Jian-Jia Chen, Katharina Morik, and Björn B Brandenburg. Efficiently approximating the worst-case deadline failure probability under edf. In *2021 IEEE Real-Time Systems Symposium (RTSS)*, pages 214–226. IEEE, 2021.

[50] Michael H Woodbury and Kang G Shin. Evaluation of the probability of dynamic failure and processor utilization for real-time systems. In *Proceedings. Real-Time Systems Symposium*, pages 222–223. IEEE Computer Society, 1988.