

Scheduling independent stochastic tasks under deadline and budget constraints

Louis-Claude Canon^{*†}, Aurélie Kong Win Chang^{*}, Yves Robert^{*‡}, Frédéric Vivien^{*}

^{*}Univ Lyon, CNRS, ENS de Lyon, Inria, Université Claude-Bernard Lyon 1, LIP UMR5668, F-69342, LYON Cedex 07, France

[†]FEMTO-ST, Université de Bourgogne Franche-Comté, France

[‡]University of Tennessee, Knoxville, TN, USA

{louis-claude.canon|aurelie.kong-win-chang|yves.robert|frederic.vivien}@ens-lyon.fr

Abstract—This paper discusses scheduling strategies for the problem of maximizing the expected number of tasks that can be executed on a cloud platform within a given budget and under a deadline constraint. The execution times of tasks follow IID probability laws. The main questions are how many processors to enroll and whether and when to interrupt tasks that have been executing for some time. We provide complexity results and an asymptotically optimal strategy for the problem instance with discrete probability distributions and without deadline. We extend the latter strategy for the general case with continuous distributions and a deadline and we design an efficient heuristic which is shown to outperform standard approaches when running simulations for a variety of useful distribution laws.

Index Terms—independent tasks, stochastic cost, scheduling, budget, deadline, cloud platform

I. INTRODUCTION

This paper deals with the following problem: given an infinite bag of stochastic tasks, and an infinite set of available Virtual Machines (VMs, or processors¹), how to successfully execute as many tasks as possible in expectation, under both a budget and a deadline constraint? The execution times of the tasks are IID (independent and identically distributed) random variables that follow a common probability distribution. The amount of budget spent during the execution of a given task is proportional to the length of its execution. At each instant, the scheduler can decide whether to continue the execution (until success) or to interrupt the task and start a new one. Intuitively, the dilemma is the following: (i) continuing execution means spending more budget, and taking the risk of waiting very long until completion, but it capitalizes on the budget already spent for the task; (ii) interrupting the task wastes the budget already spent for the task, but enables starting afresh with a new, hopefully shorter task. Of course there is a big risk here, since the new task could turn out to have an even longer execution than the interrupted one.

In addition to deciding which tasks to interrupt and when, the scheduler must also decide how many processors to enroll (this is the resource provisioning problem). There is again a trade-off here. On the one hand, enrolling many processors is mandatory when the deadline is small and the budget is large, and it allows us to make better scheduling decisions,

because we can dynamically observe many events taking place in parallel². On the other hand, enrolling too many processors increases the risk of having many unfinished tasks when budget runs out and/or when deadline strikes.

This difficult scheduling problem naturally arises with many applications in the context of cloud computing and data mining (see Section II for a detailed discussion). Informally, the goal is to extract as much information as possible from some big data set, by launching analysis tasks whose execution time strongly depends upon the nature of the data sample being processed. Not all data sample must be processed, but the larger the number of data samples successfully processed, the more accurate the analysis.

The main contribution of this work are the following:

- We provide a comprehensive set of theoretical results for the problem instance with discrete distributions and no deadline. These results show the difficulty of the general scheduling problem under study, and lay the foundations for its analysis;
- We design an asymptotically optimal scheduling strategy for the above problem instance (discrete distribution, no deadline);
- We design an efficient heuristic, OPTRATIO, for the general problem. This heuristic extends the asymptotically optimal scheduling strategy for discrete distributions to continuous ones, and accounts for the deadline constraint by enrolling the adequate number of processors. The heuristic computes a threshold at which tasks should be interrupted, which we compute for a variety of standard probability distributions (exponential, uniform, beta, gamma, inverse-gamma, Weibull, half-normal, and lognormal);
- We report a set of simulation results for three widely used probability distributions (exponential, uniform, and lognormal) that demonstrate both the superiority of the OPTRATIO heuristic over other approaches, and its good performance with short deadlines.

II. RELATED WORK

Due to lack of space, we only quote a few references here, and refer to [6] for a detailed survey of related work.

¹Throughout the text, we use both terms VM and processor indifferently.

²See the examples of Section IV-A for an illustration.

Cloud computing: See the surveys [3], [22], [23]. Resource provisioning and scheduling are key steps to the efficient execution of workflows on cloud platforms. The multi-objective scheduling problem that consists in meeting deadlines and either respecting a budget or minimizing the cost (or energy) has been extensively studied for deterministic workflows [2], [16], [25], but has received much less attention in a stochastic context. Indeed, most of the studies assume a *clairvoyant* setting: the resource provisioning and task scheduling mechanisms know in advance, and accurately, the execution time of all tasks. A handful of additional studies also consider that tasks may fail [15], [21]. Among these articles, Poola et al. [21] differ as they assume that tasks have uncertain execution times. However, they assume they know these execution times with a rather good accuracy (the standard deviation of the uncertainty is 10% of the expected execution time). They are thus dealing with uncertainties rather than a true non-clairvoyant setting. The work in [5] targets stochastic tasks but is limited to taking static decisions (no task interruption). Some works are limited to a particular type of application like MapReduce [12], [24].

Bags of tasks: A bag of tasks is an application comprising a set of independent tasks sharing some common characteristics: either all tasks have the same execution time or they are instances coming from a same distribution. Several works devoted to bag-of-tasks processing explicitly target cloud computing [10], [20]. Some of them consider the classical clairvoyant model [10]. A group of authors including Oprescu and Kielmann have published several studies focusing on budget-constrained makespan minimization in a non clairvoyant settings [18]–[20]. They do not assume they know the distribution of execution times but try to learn it on the fly [18], [19]. This work differs from ours as these authors do not consider deadlines. For instance, in [20], the objective is to try to complete all tasks, possibly using replication on faster machines, and, in case the proposed solution fails to achieve this goal, to complete as many tasks as possible. The implied assumption is that all tasks can be completed within the budget. We implicitly assume the opposite: there are too many tasks to complete all of them by the deadline, and therefore we attempt to complete as many as possible; we avoid replication, which would be a waste of resources.

Task model: Our task model assumes that some tasks may not be executed. This model is very closely related to *imprecise computations* [1], particularly in the context of real-time computations. In imprecise computations, it is not necessary for all tasks to be completely processed to obtain a meaningful result. Most often, tasks in imprecise computations are divided into a mandatory and an optional part: our work then perfectly corresponds to the optimization of the processing of the optional parts [9], [11], [14], [17]. Our task model also corresponds to the overload case of [4] where jobs can be *skipped* or *aborted*. Another, related model, is that of *anytime tasks* [13] where a task can be interrupted at any time, with the assumption that the longer the running, the higher the quality of its output. Such a model requires a function relating the time spent to a notion of reward.

Altogether, the present study appears to be unique because it is non-clairvoyant and assumes an overall deadline in addition to a budget constraint.

III. PROBLEM DEFINITION

This section details the framework and scheduling objective.

a) Tasks: We aim at scheduling a set of independent tasks whose execution times are IID (independent and identically distributed) random variables. The common probability distribution of the execution time is denoted as \mathcal{D} . We consider both discrete and continuous distributions in this work. Discrete distributions are used to better understand the problem. Continuous distributions are those typically used in the literature, namely exponential, uniform, and lognormal.

b) Platform: The execution platform is composed of identical VMs, or processors. Without loss of generality, we assume unit speed and unit cost for each VM, and we scale the task execution times when we aim at changing granularity. Execution time and budget are expressed in seconds. There is an unlimited number of VMs that can be launched by the user.

c) Constraints and optimization objective: The user has a limited budget b and an execution deadline d . The optimization problem is to maximize the expected number of tasks that can be completed until: (i) the deadline is reached; and (ii) the totality of the budget is spent. More precisely:

- The scheduler decides how many VMs to launch and which VMs to stop at each second;
- Each VM executes a task as soon as it is started;
- Each VM is interrupted as soon as the deadline or the budget is exceeded, whichever comes first;
- Each task can be deleted by the scheduler at any second before completion;
- The execution of each task is non-preemptive, except in Section IV-B that summarizes complexity results. In a non-preemptive execution, interrupted tasks cannot be relaunched, and the time/budget spent computing until interruption is completely lost. On the contrary, in a preemptive execution, a task can be interrupted temporarily (e.g., for the execution of another task, or until some event on another VM) and resumed later on.

IV. DISCRETE DISTRIBUTIONS

This section provides theoretical results when execution times follow a discrete probability distribution $\mathcal{D} = \{(p_i, w_i)\}_{1 \leq i \leq k}$. There are k possible execution times $w_1 < w_2 < \dots < w_k$ (expressed in seconds) and a task has an execution time w_i with probability p_i , where $\sum_{i=1}^k p_i = 1$. The w_i are also called *thresholds*, because they represent instants at which we should take decisions: if the current task did not complete successfully, then either we continue its execution (if the remaining budget allows for it), or we interrupt the task and start a new one. Of course the discrete distribution of the thresholds is somewhat artificial: in practice, we have continuous distributions for the execution times of the tasks. With continuous distributions, at any instant, we do not know for sure that the task will continue executing until

some fixed delay. On the contrary with discrete distributions, we know that the execution will continue (at least) until the next threshold. However, any continuous distribution can be approximated by a discrete distribution, and the more threshold values, the more accurate the approximation. In Section V, we use the results obtained for discrete distributions to design efficient strategies for continuous distributions.

In this section, we further assume that there is no scheduling deadline d , or equivalently, that the deadline is equal to the budget: $d = b$. We re-introduce deadlines when dealing with continuous distributions in Section V. To help the reader apprehend the difficulty of the problem, we start with an example in Section IV-A. We discuss problem complexity without deadline in Section IV-B, providing pseudo-polynomial optimal algorithms and comparing three scenarios: sequential, sequential with preemption, and parallel. Then in Section IV-C, we focus on cases where the budget is large and design an asymptotically optimal strategy. This strategy determines the optimal threshold at which to interrupt all yet unsuccessful tasks. This result is key to the design of an efficient heuristic for continuous distributions in Section V-A.

A. Example

We consider the following example with $k = 3$ thresholds: $\mathcal{D} = \{(0.4, 2), (0.15, 3), (0.45, 7)\}$. In other words, with a probability of 40% the execution time of a task is 2 seconds, with a probability of 15% it is 3 seconds, and with a probability of 45% it is 7 seconds. We assume that we have a total budget $b = 6$ (and recall that there is no deadline, or equivalently $d = 6$). Because $b = 6 < w_3 = 7$, no task will ever be executed up to its third threshold. We first define and evaluate the optimal policy with a single processor. Then, we exhibit a policy for two processors that achieves a better performance.

a) With a single processor: Let $E(b)$ denote the optimal expected number of completed tasks when the total budget is equal to b . To define the optimal policy for a budget of 6, we first compute $E(b)$ for the lower values of b that will appear recursively in the expression of $E(6)$:

- $E(1) = 0$, because $w_1 = 2$.
- $E(2) = p_1 \times 1 + (p_2 + p_3) \times 0 = 0.4$: when the budget is equal to 2, the only thing we can do is run the task for two units of time and check whether it completed, which happens with probability p_1 . Otherwise, no task is completed.
- $E(3) = (p_1 + p_2) \times 1 + p_3 \times 0 = 0.55$. Once again, we execute the task for two units of time. If it has not succeeded, it would be pointless to kill it because the remaining budget is 1 and $E(1) = 0$ (and if it has succeeded, we cannot take advantage of the remaining budget). Hence, if the task has not completed after two units of time, we continue its computation for the remaining unit of time and check whether it has succeeded.
- $E(4) = \max\{p_1 + E(2), p_1(1 + E(2)) + p_2(1 + E(1)) + p_3(0 + E(1))\} = 2p_1 = 0.8$. Here, two policies can be envisioned. Either, we decide to kill the first task if it has not completed by time 2 or, if it has not completed, we let it continue up to time 3 where we kill it if it has not completed (we do not

have the budget to let it run up to w_3). In the second case, we distinguish two sub-cases depending on the actual task duration. The reasoning will be the same for $E(6)$.

- $E(6) = \max\{p_1 + E(4), p_1(1 + E(4)) + p_2(1 + E(3))\} = 3p_1 = 1.2$. Once again, two policies can be envisioned. Either, we decide to kill the first task if it has not completed by time 2 or, if it has not completed, we let it pursue up to time 3 where we kill it if it has not completed (we do not have the budget to let it run up to w_3).

Therefore, the optimal expectation with a single processor is to complete 1.2 tasks.

b) With two processors: We consider the following policy: (i) we start two tasks in parallel; (ii) if none of them completes by time 2, we let them run up to time 3; (iii) otherwise, we kill at time 2 any not-yet completed task and start a new task instead. The following case analysis displays the expected number of completed tasks for each case of execution time of the two tasks initially started:

	w_1	w_2	w_3
w_1	$2 + p_1$	$1 + p_1$	$1 + p_1$
w_2	$1 + p_1$	2	1
w_3	$1 + p_1$	1	0

For instance, the square at the intersection of the column w_1 and the row w_2 corresponds to the case where the task on the first processor completes in two units of time, where the task on the second processor would have needed 3 units of time. Because of our policy, this second task is killed and at time 2 and we have completed a single task. There remain 2 units of time and we start a third task, which will complete in this budget with probability p_1 . Therefore, the total expected number of completed task in this configuration is $1 + p_1$, and this configuration happens with probability $p_1 p_2$.

The total expected number of completed tasks is:

$$E' = p_1^2(2 + p_1) + 2p_1(p_2 + p_3)(1 + p_1) + 2p_2^2 + 2p_2p_3 = 1.236.$$

Therefore, this two-processor policy is more efficient than the optimal single processor policy! Even in the absence of deadline parallelism may help to achieve better performance.

This example helps comprehend the difficulty of the scheduling problem under study.

B. Complexity results

Due to lack of space, we only state results, and refer to the extended version [6] for proofs and algorithms. This section is the only one in the paper where we allow preemption. We compare the performance of sequential scheduling, without or with preemption, to that of parallel scheduling, for the problem instance without deadline. More precisely, in [6]:

- We provide three pseudo-polynomial dynamic programming algorithms to compute the optimal expected number of tasks that can be completed for a given budget b : (i) Without preemption on a single processor, the complexity is $O(kb)$; (ii) With preemption on a single processor, it is $O\left(\prod_{s=1}^{k-1} \left(1 + \frac{b}{w_s}\right)\right)$; and (iii) Without preemption on parallel processors it is $O((b+k)b^3w_k^b)$.

- We show that any algorithm designed to be executed on p processors with or without preemption can be simulated on a single processor with preemption with the same performance. Consider an algorithm A designed to be executed on p processors with or without preemption. We show how to build from A an algorithm B that executes on a single processor with preemption and such that, whatever the problem instance, B completes on a single processor at least as many tasks as A with p processors. This shows that the knowledge gained by attempting several executions in parallel cannot be used to successfully execute more tasks than in sequence (with preemption, and no deadline).
- We show that, without preemption, scheduling with parallel processors is never worse than scheduling with a single processor, and can achieve strictly better performance on some instances.
- We show that scheduling with preemption and with a single processor is never worse than scheduling without preemption and with parallel processors, and can achieve strictly better performance on some instances.

C. Asymptotic behavior

In this section, we derive an asymptotically optimal strategy when letting the budget tend to infinity (see [6] for all proofs). Because the scheduling strategy described below is applied independently on each processor, we can assume that $p = 1$ throughout this section without loss of generality. As stated earlier, recall that we assume that there is no deadline. Note that a fixed deadline would make no sense when $b \rightarrow +\infty$ and $p = 1$. We first describe the strategy in Section IV-C1 and show its asymptotic optimality in Section IV-C2. Throughout this section, we are given a discrete distribution $\mathcal{D} = \{(p_i, w_i)\}_{1 \leq i \leq k}$.

1) *Optimal fixed-threshold strategy*: For $1 \leq i \leq k$, the i -th fixed-threshold strategy, or FTS_i , interrupts every unsuccessful task at threshold w_i , i.e., when the task has been executing for w_i seconds without completing. There are k such strategies, one per threshold. Informally, our criterion to select the best one is to maximize the ratio

$$\mathcal{R} = \frac{\text{expected number of tasks completed}}{\text{budget}}.$$

Indeed, this ratio measures the success rate per time unit, or equivalently, per budget unit (since we have unit execution speed). Formally, we would like to compute

$$\mathcal{R}_i(b) = \frac{N_i(b)}{b} \quad (1)$$

where $N_i(b)$ is the expected number of tasks that are successfully completed when using strategy FTS_i that interrupts all unsuccessful tasks after w_i seconds, and proceeds until the budget b has been spent. It turns out that we can compute the limit \mathcal{R}_i of $\mathcal{R}_i(b)$ when the budget b tends to infinity:

Proposition 1.

$$\lim_{b \rightarrow \infty} \mathcal{R}_i(b) = \mathcal{R}_i \stackrel{\text{def}}{=} \frac{\sum_{j=1}^i p_j}{\sum_{j=1}^i p_j w_j + (1 - \sum_{j=1}^i p_j) w_i}$$

The optimal fixed-threshold strategy FTS_{opt} is defined as the strategy FTS_i whose ratio \mathcal{R}_i is maximal. If several strategies FTS_i achieve the maximal ratio \mathcal{R}_{opt} , we pick the one with smallest w_i (to improve success rate when the budget is limited and truncation must occur). Formally:

Definition 1. FTS_{opt} is the strategy FTS_{i_0} where $i_0 = \min_{1 \leq i \leq k} \{i \mid \mathcal{R}_i = \min_{1 \leq j \leq k} \mathcal{R}_j\}$.

To conclude this section, we work out a little example. Consider a distribution $\mathcal{D} = \{(p_i, w_i)\}_{1 \leq i \leq 3}$ with 3 thresholds. We have $\mathcal{R}_1 = \frac{p_1}{w_1}$, $\mathcal{R}_2 = \frac{p_1 + p_2}{p_1 w_1 + (1 - p_1) w_2}$, and $\mathcal{R}_3 = \frac{p_1 + p_2 + p_3}{p_1 w_1 + p_2 w_2 + (1 - p_1 - p_2) w_3} = \frac{1}{p_1 w_1 + p_2 w_2 + p_3 w_3}$. We pick the largest of these three values to derive FTS_{opt} .

2) *Asymptotic optimality of FTS_{opt}* : A scheduling strategy makes the following decisions for each task: when a new threshold is reached, and if the task is not successful at this point, decide whether either to continue execution until the next threshold, or to interrupt the task. In the most general case, these decisions may depend upon the remaining available budget. However, when the budget is large, it makes sense to restrict to strategies where such decisions are taken independently of the remaining budget, independently to past history, and either deterministically or non-deterministically but according to some fixed probabilities. We formally define such strategies as follows:

Definition 2. A mixed-threshold strategy $MTS(q_1, q_2, \dots, q_{k-1})$, where $0 \leq q_j \leq 1$ for $1 \leq j \leq k-1$ are fixed probabilities, makes the following decision when the execution of a task reaches threshold w_i , for $1 \leq i \leq k-1$, without success: it decides randomly to continue execution until the next threshold with probability q_i , and to interrupt the task otherwise, hence with probability $1 - q_i$.

Of course, the fixed-threshold strategy FTS_i coincides with $MTS(1, \dots, 1, 0, \dots, 0)$ where the last 1 is in position $i-1$: $q_j = 1$ for $j < i$ et $q_j = 0$ for $j \geq i$. In this section, we prove our main result for discrete distributions:

Theorem 1. FTS_{opt} is asymptotically optimal among all mixed-threshold strategies.

V. CONTINUOUS DISTRIBUTIONS

In this section, we build upon the previous results and deal with continuous distributions. We do assume we have a fixed budget and a deadline. Thus, in contrast to Section IV, the distribution \mathcal{D} is now continuous and has expected value $\mu_{\mathcal{D}}$ and variance $\sigma_{\mathcal{D}}^2$. Let $F(x)$ be its cumulative distribution function and $f(x)$ its probability density function. The objective remains to execute as many tasks as possible given a budget b , a deadline d and a potentially unlimited number of processors.

We start by designing several heuristics in Section V-A and then we assess their efficiency through experiments in Section V-B. The code and scripts used for the simulations and the data analysis are publicly available online [7].




Figure 1. Efficiency (ratio \mathcal{R} of number of tasks successfully executed per budget unit) for different probability distributions. Some distributions have an optimal finite cutting threshold depicted with a vertical red line.

A. Heuristics

We present below different heuristics, among which an extension of the asymptotically optimal greedy strategy of Section IV-C to the continuous case. In all cases, we enroll $\lceil \frac{b}{d} \rceil$ machines. The rationale for this choice is that this is the maximum number of machines that can work in parallel and continuously, up to the deadline. We have three main classes of heuristics:

- **MEANVARIANCE**(x) is the family of heuristics that kill a task as soon as its execution time reaches $\mu_D + x\sigma_D$, where x is some positive or negative constant.
- **QUANTILE**(x) is the family of heuristics that kill a task when its execution time reaches the x -quantile of the distribution \mathcal{D} with $0 \leq x \leq 1$.
- **OPTRATIO** is the heuristic inspired by the asymptotically optimal strategy for discrete distributions. OPTRATIO interrupts all (unsuccessful) tasks at time $l = \arg \max_l \mathcal{R}(l)$ where

$$\mathcal{R}(l) = \frac{F(l)}{\int_0^l x f(x) dx + l(1 - F(l))}.$$

The idea behind OPTRATIO is that it maximizes the ratio of the probability of success (namely $F(l)$) to the expected amount of budget spent for a single task when the task is interrupted at time l (i.e., $\int_0^l x f(x) dx$ for the cases when the task terminates sooner than l and $\int_l^\infty l f(x) dx = l(1 - F(l))$ otherwise). This is a continuous extension of the approach proposed in Section IV-C, and we expect OPTRATIO to perform well for large budgets.

We now analyze OPTRATIO with some classical probability distributions defined on nonnegative values (task execution times need to be nonnegative). For the exponential distribution, which is memoryless, $\mathcal{R}(l) = \lambda$ where λ is the rate of the

Table I
PROBABILITY DISTRIBUTIONS WITH THEIR PROBABILITY DISTRIBUTION FUNCTION (PDF) AND DENSITY GRAPH. SUPPORTS ARE $[0, \infty)$ FOR ALL DISTRIBUTIONS EXCEPT FOR UNIFORM, WHERE IT IS $[a, b]$ AND BETA, WHERE IT IS $[0, 1]$. NOTE THAT $B(\alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}$.

Name	PDF	Density
Uniform	$\frac{1}{b-a}$	
Exponential	$\lambda e^{-\lambda x}$	
Half-normal	$\frac{\sqrt{2}}{\theta\sqrt{\pi}} e^{-\frac{x^2}{2\theta^2}}$	
Lognormal	$\frac{1}{x\beta\sqrt{2\pi}} e^{-\frac{(\log(x)-\alpha)^2}{2\beta^2}}$	
Beta	$\frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha,\beta)}$	
Gamma	$\frac{1}{\Gamma(k)\theta^k} x^{k-1} e^{-\frac{x}{\theta}}$	
Weibull	$\frac{k}{\theta^k} x^{k-1} e^{-(\frac{x}{\theta})^k}$	
Inverse-gamma	$\frac{\theta^k}{\Gamma(k)} x^{-k-1} e^{-\frac{\theta}{x}}$	

distribution. In this case, any l can be chosen and the tasks may be interrupted at any moment with OPTRATIO without modifying the performance. For the uniform distribution (between a and b), $\mathcal{R}(l) = 2 \frac{l-a}{-l^2+2bl-a^2}$, which takes its maximum value for $l = b$ ($\mathcal{R}(b) = \frac{2}{a+b}$). In this case, tasks should never be interrupted to maximize performance. We established these results for exponential and uniform distributions through simple algebraic manipulations.

In addition to the exponential and uniform distributions,

Table I presents other standard distributions. For these distributions, we provide some code [7] to numerically compute the optimal time l at which tasks should be interrupted. Note that there exist many relations between probability distributions. For instance, the beta distribution with both shape parameters equal to one is the same as the uniform distribution, whereas it has a U-shape with both equal to 0.5, and a bell-shape with both equal to 2. Also, the exponential distribution is a special case of the gamma and Weibull distributions when their shape parameter is one.

Figure 1 shows how $\mathcal{R}(l)$ varies as a function of the cutting threshold l , for the probability distributions shown in Table I. Recall that OPTRATIO will select the threshold l for which $\mathcal{R}(l)$ is maximum. For instance, this threshold is $l = 1$ for the uniform distribution, meaning that we should never interrupt any task. The threshold can be any value of l for the exponential distribution, and this is due to the memoryless property: we can interrupt a task at any moment, without any expected consequence. The threshold is $l = \infty$ for the half-normal distribution, meaning again that we should never interrupt any task, just as for uniform distributions. Note that the expected value of all distributions is not the same overall, because we use standard parameters in Figure 1, hence ratio values are not comparable across distributions.

We remark that the lognormal distribution, which presents a fast increase followed by a slow decrease with an heavy tail, exhibits an optimal cutting threshold during the execution of a task: on Figure 1, we see that the optimal threshold is $l \approx 1.73$ (we computed this value numerically) for the distribution $\text{Lognormal}(0, 1)$. We make a similar observation for the inverse-gamma distributions, where the optimal threshold is $l \approx 0.7$ for $\text{Inv-Gamma}(1.5, 0.5)$ and $l \approx 2.32$ for $\text{Inv-Gamma}(3, 2)$. These lognormal and inverse-gamma distributions share the following properties: the density is close to zero for small costs and has a steep increase. On the contrary, the bell-shape beta distribution $\text{Beta}(2, 2)$ has a small density for small costs but does not have a steep increase, and tasks should never be interrupted (in other words, the optimal cutting threshold is $l = 1$ for $\text{Beta}(2, 2)$).

Finally, we observe that three distributions are the most efficient when the cutting threshold tends to zero ($\text{Beta}(0.5, 0.5)$, $\text{Gamma}(0.5, 2)$ and $\text{Weibull}(0.5, 1/\Gamma(3))$). We point out that it is unlikely that such distributions would model actual execution times in practice.

B. Experiments

The following experiments make use of three standard distributions: exponential, uniform, and lognormal. The first two distributions are very simple and easy to use, while the latter has been advocated to model file sizes [8], and we assume that task costs could naturally obey this distribution too. Moreover, the lognormal distribution is positive, it has a tail that extends to infinity and the logarithm of the data values are normally distributed. Also, this distribution leads to a non-trivial cutting threshold, contrarily to exponential (interrupt anywhere) or uniform (never interrupt), thereby allowing for

a complete assessment of our approach. In all experiments, we submit tasks steadily until the budget and/or the deadline is exhausted.

Figure 2 shows the number of successfully executed tasks for each heuristic with three distributions (lognormal, uniform, exponential) of same expected value $\mu = 1$, with a budget and deadline $b = d = 100$. Note that to ensure a given expected value and standard deviation for the lognormal distribution, we set its parameters as follows: $\alpha = \log(\mu) - \log(\sigma^2/\mu^2 + 1)/2$ and $\beta = \sqrt{\log(\sigma^2/\mu^2 + 1)}$. Note also that using a standard deviation $\sigma = 3$ for the lognormal distribution corresponds to a high level of heterogeneity. To see this intuitively, take a discrete distribution with 11 equally probable costs, 10 of value 0.1 and 1 of value 10: its expected value is $\mu = 1$ while its standard deviation is $\sigma \approx 2.85$. Finally, we note that Figure 2 confirms that tasks with exponentially distributed costs can be interrupted at any time and that tasks with uniformly distributed costs should never be interrupted.

Next, we focus on the lognormal distribution. First, in Figure 3, we assess the impact of three important parameters: the standard deviation, the budget and the deadline, respectively. The expected value is always $\mu = 1$. By default, the standard deviation is $\sigma = 3$, and the budget and deadline are set to 100 ($b = d = 100$), which means that a single machine is enrolled. When we vary the standard deviation (first row in Figure 3), we keep $b = d = 100$. When we vary the budget (second row in Figure 3), we maintain the equality $b = d$. When we vary the deadline (third row in Figure 3), we keep $b = 100$, hence more VMs are enrolled (10 VMs when $d = 10$ and 100 VMs when $d = 1$). Each heuristic is run 100,000 times for each scenario. The error bars represent an interval from the mean of two standard deviations of the number of successes. For a normal distribution, this means that more than 95% of the values are in this interval. Note that the subfigures with $\sigma = 3$, $b = 100$ and $d = 100$ in Figure 3 are all the same as the subfigure with the lognormal distribution in Figure 2.

On Figure 3, we see that the higher the standard deviation, the larger the gain of every approach. With a low standard deviation, all approaches perform similarly. Increasing the budget tends to decrease the variability when running several times the same approach (the error bars are narrower with large budgets, which makes the approaches more predictable). This is a consequence of the law of large numbers. However, the expected efficiency (around 2.5 tasks per unit of time) remains similar even for a low budget of 30. Finally, decreasing significantly the deadline prevents some strategies from letting tasks run a long time. Long running tasks are then forced to be interrupted early, which is similar to the behavior of the more efficient approaches.

In all tested situations, the OPTRATIO algorithm with the optimal threshold achieved the best results.

Finally, Figure 4 depicts the efficiency of OPTRATIO with small deadlines. Even though our approach extends a strategy that is asymptotically optimal when both the budget and the deadline are large, it does perform well with small deadlines, as long as d is not lower than the cutting threshold. In the



Figure 2. Number of successfully executed tasks for each heuristic with three distributions (lognormal, uniform, exponential) of same expected value $\mu = 1$, with a budget and deadline $b = d = 100$ (which means that a single machine is enrolled). Each heuristic is run 100,000 times for each scenario. The error bars are computed with the mean plus/minus two standard deviations of the number of successes. The lognormal distribution has parameters $\alpha \approx -1.15$ and $\beta \approx 1.52$ to have an expected value $\mu = 1$ and a standard deviation $\sigma = 3$, and the optimal cutting threshold for OPTRATIO is $l \approx 0.1$). The exponential distribution has shape $\lambda = 1$ and the cutting threshold is arbitrarily set to $l = 2$. The uniform distribution has parameters $a = 0$ and $b = 2$, and the cutting threshold is $l = 2$.

settings of Figure 4, where the average execution time of a task is equal to 1, this means that as soon as the deadline is equal to 0.1, OPTRATIO achieves its asymptotic performance! (The reader can compare the performance of OPTRATIO for deadlines of 100 and 0.1 on Figures 2 and 4.) Finally note that on Figure 4, $b = 100$ and that, therefore, OPTRATIO uses 1,000 processors for a deadline $d = 0.1$. This confirms that neither the budget, nor the deadline need to be large for OPTRATIO to reach its best efficiency, and that this heuristic is extremely robust.

VI. CONCLUSION

This paper deals with scheduling strategies to successfully execute the maximum number of a bag of stochastic tasks on VMs (Virtual Machines) with a finite budget and under a deadline constraint. We first focused on the problem instance with discrete probability distributions and no deadline. We proposed three optimal dynamic programming algorithms for different scenarios, depending upon whether tasks may be preempted or not, and whether multiple VMs may be enrolled or only a single one. We also introduced an asymptotically optimal method that computes a cutting threshold that is independent of the remaining budget. Then, we extended this approach to the continuous case and with deadline. We designed OPTRATIO, an efficient heuristic which we validated through simulations with classical distributions such as exponential, uniform, and lognormal. Tests with several values of the deadline, leading to enroll different numbers of VMs, also confirm the relevance and robustness of our proposition.

Future work will be dedicated to considering heterogeneous tasks (still with stochastic costs), as well as heterogeneous VMs. Typically, cloud providers provide a few different categories of VM with different computer power and nominal cost, and it would be interesting (albeit challenging) to extend our study to such a framework. Another interesting direction would be to take into account start-up costs when launching a VM, thereby reducing the amount of parallelism, because fewer VMs will likely be deployed.

REFERENCES

- [1] M. Amirijoo, J. Hansson, and S. H. Son. Specification and management of qos in real-time databases supporting imprecise computations. *IEEE Trans. Computers*, 55(3):304–319, 2006.
- [2] V. Arabnejad, K. Bubendorfer, and B. Ng. Budget distribution strategies for scientific workflow scheduling in commercial clouds. In *12th IEEE International Conference on e-Science*, pages 137–146, Oct 2016.
- [3] M. U. Bokhari, Q. Makki, and Y. K. Tamandani. A survey on cloud computing. In D. M. V. Aggarwal, V. Bhatnagar, editor, *Big Data Analytics*, volume 654 of *Advances in Intelligent Systems and Computing*. Springer, 2018.
- [4] G. Buttazzo. Handling overload conditions in real-time systems. In S. M. Babamir, editor, *Real-Time Systems, Architecture, Scheduling, and Application*, chapter 7. InTech, Rijeka, 2012.
- [5] Y. Caniou, E. Caron, A. K. W. Chang, and Y. Robert. Budget-aware scheduling algorithms for scientific workflows with stochastic task weights on heterogeneous iaas cloud platforms. In *27th International Heterogeneity in Computing Workshop HCW 2013*. IEEE Computer Society Press, 2018.
- [6] L.-C. Canon, A. K. W. Chang, Y. Robert, and F. Vivien. Scheduling independent stochastic tasks under deadline and budget constraints. Research Report 9178, INRIA, June 2018.
- [7] L.-C. Canon, A. Kong Win Chang, F. Vivien, and Y. Robert. Code for scheduling independent stochastic tasks under deadline and budget constraints, June 2018. <https://doi.org/10.6084/m9.figshare.6463223.v2>.
- [8] D. Feitelson. Workload modeling for computer systems performance evaluation. *Version 1.0.3*, pages 1–607, 2014.
- [9] W. Feng and J. W. S. Liu. An extended imprecise computation model for time-constrained speech processing and generation. In *[1993] Proceedings of the IEEE Workshop on Real-Time Applications*, pages 76–80, May 1993.
- [10] A. Grekoti and N. V. Shakhlevich. Scheduling bag-of-tasks applications to optimize computation time and cost. In *PPAM 2013.*, volume 8385 of *Lecture Notes in Computer Science*. Springer, 2014.
- [11] H. Hassan, J. Simó, and A. Crespo. Flexible real-time mobile robotic architecture based on behavioural models. *Engineering Applications of Artificial Intelligence*, 14(5):685 – 702, 2001.
- [12] E. Hwang and K. H. Kim. Minimizing cost of virtual machines for deadline-constrained mapreduce applications in the cloud. In *Proceedings of the 2012 ACM/IEEE 13th International Conference on Grid Computing, GRID '12*, pages 130–138, Washington, DC, USA, 2012. IEEE Computer Society.
- [13] F. Jumel and F. Simonot-Lion. Management of anytime tasks in real time applications. In *XIV Workshop on Supervising and Diagnostics of Machining Systems*, Karpacz/Pologne, 2003. Colloque avec actes et comité de lecture. internationale.
- [14] H. Kobayashi and N. Yamasaki. Rt-frontier: a real-time operating system for practical imprecise computation. In *Proceedings. RTAS 2004. 10th IEEE Real-Time and Embedded Technology and Applications Symposium, 2004.*, pages 255–264, May 2004.




Figure 3. Number of successfully executed tasks for each heuristic, with lognormal costs and expected value $\mu = 1$. Unless otherwise specified, the standard deviation is $\sigma = 3$, and the budget and deadline are $b = d = 100$. Each heuristic is run 100,000 times for each scenario. The error bars are computed with the mean plus/minus two standard deviations of the number of successes. The lognormal distribution has parameters $\alpha \approx -1.15$ and $\beta \approx 1.52$ by default (to have $\mu = 1$ and $\sigma = 3$) (the cutting threshold for OPTRATIO is $l \approx 0.1$). They are $\alpha \approx -0.35$ and $\beta \approx 0.83$ when $\sigma = 1$ ($l \approx 2.1$) and $\alpha \approx -0.8$ and $\beta \approx 1.27$ when $\sigma = 2$ ($l \approx 0.34$).




Figure 4. Number of successfully executed tasks for OPTRATIO with a budget $b = 100$ and optimal cutting threshold $l \approx 0.1$. OPTRATIO is run 100,000 times for each deadline. The error bars are computed with the mean plus/minus two standard deviations of the number of successes. The lognormal distribution has parameters $\alpha \approx -1.15$ and $\beta \approx 1.52$ to have an expected value $\mu = 1$ and a standard deviation $\sigma = 3$.

- [15] K. Liu, H. Jin, J. Chen, X. Liu, D. Yuan, and Y. Yang. A compromised-time-cost scheduling algorithm in swindow-c for instance-intensive cost-constrained workflows on a cloud computing platform. *Int. J. High Performance Computing Applications*, 24(4):445–456, 2010.
- [16] M. Malawski, G. Juve, E. Deelman, and J. Nabrzyski. Algorithms

- for cost- and deadline-constrained provisioning for scientific workflow ensembles in iaas clouds. *Future Gen. Comp. Syst.*, 48:1–18, 2015.
- [17] J. Meng, S. Chakradhar, and A. Raghunathan. Best-effort parallel execution framework for recognition and mining applications. In *2009 IEEE IPDPS*, pages 1–12, May 2009.
- [18] A. M. Oprescu and T. Kielmann. Bag-of-tasks scheduling under budget constraints. In *2010 IEEE Second International Conference on Cloud Computing Technology and Science*, pages 351–359, Nov. 2010.
- [19] A.-M. Oprescu, T. Kielmann, and H. Leahu. Budget estimation and control for bag-of-tasks scheduling in clouds. *Parallel Processing Letters*, 21(02):219–243, 2011.
- [20] A. M. Oprescu, T. Kielmann, and H. Leahu. Stochastic tail-phase optimization for bag-of-tasks execution in clouds. In *Fifth Int. Conf. on Utility and Cloud Computing*, pages 204–208. IEEE, Nov. 2012.
- [21] D. Poola, S. K. Garg, R. Buyya, Y. Yang, and K. Ramamohanarao. Robust scheduling of scientific workflows with deadline and budget constraints in clouds. In *AINA 2014*, pages 858–865, May 2014.
- [22] S. Singh and I. Chana. Cloud resource provisioning: survey, status and future research directions. *Knowledge and Information Systems*, 49(3):1005–1069, Dec. 2016.
- [23] S. Singh and I. Chana. A survey on resource scheduling in cloud computing: Issues and challenges. *J. Grid Comp.*, 14(2):217–264, 2016.
- [24] F. Tian and K. Chen. Towards optimal resource provisioning for running mapreduce programs in public clouds. In *2011 IEEE 4th International Conference on Cloud Computing*, pages 155–162. IEEE, July 2011.
- [25] C. Q. Wu, X. Lin, D. Yu, W. Xu, and L. Li. End-to-end delay minimization for scientific workflows in clouds under budget constraint. *IEEE Transactions on Cloud Computing*, 3(2):169–181, April 2015.