# Reservation and Checkpointing Strategies for Stochastic Jobs (Extended Version) 

Ana Gainaru, Brice Goglin, Valentin Honoré, Guillaume Pallez, Padma Raghavan, Yves Robert, Hongyang Sun

## RESEARCH

# Reservation and Checkpointing Strategies for Stochastic Jobs (Extended Version) 

Ana Gainaru* ${ }^{*}$, Brice Goglin ${ }^{\dagger \dagger}$, Valentin Honoré ${ }^{\dagger}$, Guillaume Pallez ${ }^{\dagger}$, Padma Raghavan*, Yves Rober地, Hongyang Sun*<br>Project-Teams TADaaM and ROMA<br>Research Report n 9294 - October 2019 - 38 pages


#### Abstract

In this work, we are interested in scheduling and checkpointing stochastic jobs on a reservation-based platform. We assume that jobs can be interrupted at any time to take a checkpoint, and that job execution times follow a known probability distribution. The user has to determine a sequence of fixed-length reservation requests, and to decide whether to checkpoint the state of the execution, or not, at the end of each request. The execution of the job is successful only when it terminates within a request, otherwise it must be resubmitted, using the next request in the reservation sequence, and restarting execution from the last checkpointed state. The cost of each reservation depends on both its duration and on the actual utilization of the platform during that request, which includes a restart if some previous reservation was terminated with a checkpoint, and possibly a checkpoint at the end of the current request. The cost of a job is then the cumulated cost of all the reservations that were needed until its completion. Overall, the objective is to find a reservation sequence that minimizes the total expected cost to execute a job. We provide an optimal strategy for discrete probability distributions of job execution times, and we design fully polynomial-time approximation strategies for continuous distributions with bounded support. We experimentally evaluate these strategies for jobs following a wide range of usual probability distributions, as well as one distribution obtained from traces of a neuroscience application. We compare our strategies with standard approaches that use periodic-length reservations (the next reservation is longer than the previous one by a constant amount of time) and simple checkpointing strategies (either checkpoint all reservations, or none).


Key-words: scheduling, checkpointing, stochastic cost, computing platform, sequence of requests, neuroscience applications

[^0]```
RESEARCH CENTRE
BORDEAUX - SUD-OUEST
200 avenue de la Vieille Tour
33405 Talence Cedex
```


## Stratégies de réservation avec points de sauvegarde pour l'ordonnancement de tâches stochastiques

## Résumé :

Dans ce rapport, nous nous intéressons à l'ordonnancement de tâches stochastiques exécutées sur une plateforme à réservations, où l'utilisateur réalise des requêtes successives de temps de calcul. Le temps d'exécution des tâches considérées n'est pas connu à l'avance. Ce temps est représenté par une loi de probabilité, décrite sous la forme d'une densité de probabilité. Nous nous intéressons à ordonnancer une instance d'une telle tâche, c'est à dire que nous ne connaissons pas son temps d'exécution qui reste constant tout au long de l'ordonnancement. Dans ce cas, le coût de l'ordonnancement dépend à la fois de la durée des requêtes et du temps d'exécution de la tâche considérée.

Nous supposons de plus que la tâche peut être interrompue à tout instant (tâche divisible) pour un point de sauvegarde. Nous avons donc la possibilité de prendre un point de sauvegarde à la fin de certaines réservations bien choisies. L'avantage est de pouvoir repartir de l'état sauvegardé à la prochaine réservation au lieu de repartir du début, si la tâche ne termine pas pendant la réservation courante. Le prix à payer est le temps de sauvegarde, et de redémarrage.

L'objectif de ce travail est de déterminer une stratégie de réservation optimale qui minimise le coût total de l'ordonnancement. Une stratégie de réservations est une séquence de requêtes croissantes, qui sont payées les unes à la suite des autres en ordre croissant jusqu'à complétion de la tâche. Pour chaque réservation, nous indiquons si un point de sauvegarde doit être pris ou non.

Nous décrivons une telle solution optimale pour des distributions de probabilité discrètes, et nous donnons un schéma d'approximation polynomial pour les distributions continues à support compact. Nous comparons expérimentalement ces stratégies aux stratégies usuelles qui incrémentent la longueur de chaque réservation par une valeur constante, et qui décident de sauvegarder soit toutes les réservations, soit aucune. Nous utilisons un grand nombre de lois de probabilité usuelles (i.e. Uniforme, Exponentielle, Log-Normale, Weibull, Beta etc), ainsi qu'une distribution basée sur l'interpolation de traces d'applications de neurosciences exécutées sur une plateforme HPC.
Mots-clés : ordonnancement, coût stochastique, plateformes de calcul, séquence de réservations, point de sauvegarde, applications de neurosciences

## Contents

1 Introduction ..... 4
2 Framework ..... 7
2.1 Stochastic jobs ..... 7
2.2 Cost model ..... 7
2.3 Expected cost ..... 9
2.4 Optimization Problem ..... 10
3 Algorithms ..... 10
3.1 Expected cost ..... 10
3.2 Dynamic programming for discrete distributions ..... 12
3.3 Approximation algorithm for continuous distributions ..... 14
3.4 Extensions. ..... 18
3.5 Periodic Heuristics ..... 19
3.5.1 All-ChECKPOINT-PERIODIC for Exponential distributions ..... 19
3.5.2 All-Checkpoint-Periodic for Uniform distributions ..... 20
4 Performance evaluation ..... 21
4.1 Evaluation methodology ..... 21
4.2 Results for Scenario 1 ..... 23
4.3 Results for Scenario 2 ..... 28
5 Experiments ..... 29
5.1 Experimental setup ..... 32
5.2 Experimental results ..... 32
6 Related Work ..... 34
6.1 Reservation-based scheduling ..... 35
6.2 Stochastic scheduling and checkpointing ..... 35
7 Conclusion and Future Work ..... 35

## 1 Introduction

In this report, we revisit our recent work on reservation strategies for stochastic jobs [3]. Stochastic jobs originate from Big Data or Machine Learning workloads, whose performance is widely dependent on characteristics of input data. Figure 1 shows an example of a Neuroscience job. Reservation strategies provide a sequence of fixed length reservations to execute a stochastic job. If the reservation is too short for the job, it is restarted in a longer reservation. We extend the approach to include the possibility of checkpointing at the end of some (well-chosen) reservations. The idea of checkpointing is very natural and widely used in practice, in particular for long jobs lasting several hours, but it dramatically complicates the design of scheduling strategies. To the best of our knowledge, existing approaches either checkpoint at the end of all reservations, or never checkpoint. For large-scale applications, checkpointing to save intermediate results at the end of each reservation is the de facto standard approach.


Figure 1: Execution times from 2017 for a Structural identification of orbital anatomy application, and its fitted distribution (in red).

We use an example to help understand the challenges of the problem under study. Consider the jobs depicted in Figure 1. We model their execution time with $\mathcal{D}$, a truncated LogNormal probability distribution on the domain $[a, b]=[0,80 \mathrm{~h}]$ (mean $\mu=21 \mathrm{~h}$, standard deviation $\sigma=20 \mathrm{~h})$. The exact execution time $X$ of the next job to be scheduled is not known until that job has successfully completed, but instead is randomly and uniformly sampled from the target probability distribution $\mathcal{D}$. The objective is to minimize the expected cost of scheduling this job. To do so, we have to derive a sequence of reservations. Then we compute the cost of the job given that sequence, and aim at minimizing the expected value. To determine the cost of a reservation, we use the generic model from our previous work [3]. This model has been shown to encompass a variety of scenarios, ranging from the Reserved Instances of Cloud Computing where one pays (for a cheaper cost) only the reserved time [2] , to High-Performance Computing (HPC) platforms where one pays the total execution time (wait time and runtime).

Specifically, for a reservation of length $W_{1}$ and an actual execution duration of length $X$, the cost is expressed as:

$$
\begin{equation*}
\alpha W_{1}+\beta \min \left(W_{1}, X\right)+\gamma \tag{1}
\end{equation*}
$$

where $\alpha, \beta$ and $\gamma$ are constant parameters that depend on the platform and the cost model. The first component $\alpha W_{1}$ is proportional to the reservation length (pay for what you ask). The second component $\beta \min \left(W_{1}, X\right)$ is proportional to the actual execution time (pay for what you use). Finally, the third and last component is a start-up time possibly associated with the first and/or second components.

To illustrate the contribution of this work, we use $\alpha=1, \beta=\gamma=0$ in the example. In Figure 2, we depict three strategies, and their expected costs in hours: (i) $S_{1}$ (Standard), which reserves the upper bound of $\mathcal{D}, W_{1}=b=80$; (ii) $S_{2}$ (No Checkpoint), which introduces a first reservation of size $W_{1}=20$ before the second reservation $W_{2}=80$; (iii) $S_{3}$ (With Checkpoint), which introduces a first checkpointed reservation of size $W_{1}=20+7$ (20 to cover jobs shorter than 20 , and 7 (red box) is the cost to checkpoint), then a second non-checkpointed reservation of size $W_{2}=7+20$ ( 7 (green box) is the cost to restart, 20 to cover jobs larger than 20 and smaller than 40 ), and a third reservation of size $W_{3}=7+60$ ( 7 is the cost to restart, 60 to cover jobs of size up to $b$ ). Now, we show how to compute the expected cost of the different strategies. For $S_{1}$, there is a unique reservation that represents the total cost:

$$
\mathbb{E}\left(S_{1}\right)=80
$$

For $S_{2}$, the expected cost is decomposed as follows. Either one reservation is sufficient to success, or the two reservations are needed. In the first case, the reservation $W_{1}$ of size 20 is performed and is weighted by $\mathbb{P}(X \leq 20)$, the probability of job success with this single reservation. In the second case, the first reservation has failed but has been paid, and one has to perform the second reservation $W_{2}=80$, for a total cost of $80+20=100$. This scenario occurs only if $20<X \leq 80$. Hence, we have

$$
\begin{aligned}
\mathbb{E}\left(S_{2}\right) & =20 \cdot \mathbb{P}(X \leq 20)+(80+20) \cdot \mathbb{P}(20<X \leq 80) \\
& =20 \times 0.66+100 \times 0.34 \\
& =47.2
\end{aligned}
$$

Regarding the expected cost of $S_{3}$, we apply the same principle as in the calcultion of $S_{2}$. However, we leverage the wastage of failed reservations by using checkpointing (inducing an extra-cost in the reservation) and updating the probability associated to each reservation:

$$
\begin{aligned}
\mathbb{E}\left(S_{3}\right) & =27 \cdot \mathbb{P}(X \leq 20)+(27+27) \cdot \mathbb{P}(20<X \leq 40)+(54+67) \cdot \mathbb{P}(40<X) \\
& =27 \times 0.66+54 \times 0.26+121 \times 0.08 \\
& =41.54
\end{aligned}
$$

Note that $\tilde{S}_{3}$, the variant of $S_{3}$ where the second reservation is also checkpointed, would have a larger expected cost due to this second checkpoint: $\mathbb{E}\left(\tilde{S}_{3}\right)=27 \times 0.66+61 \times 0.26+128 \times 0.08=$ 43.92. Similarly one can verify that not performing the second reservation at all would also have increased the expected cost. This example shows that checkpointing does help for some scenarios but has too much overhead for others, and suggests that finding the best trade-off is difficult.

Indeed, in the general case, one has to decide which reservations should be checkpointed, depending on application profile and platform parameters. Moreover, determining the expected cost of a given reservation sequence together with scheduling decisions gets quite complicated. Section 2 gives a detailed formula for the expected cost, and Theorem 1 in Section 3.1 provides a simplified version. In our previous work without checkpoints [3], we have been able to analytically characterize the optimal sequence of reservations for any smooth probability distribution (except the length of the first reservation which had to be found numerically). The problem with


Figure 2: Illustration of different reservation strategies. The checkpoint (red) and restart (green) costs are equal to 7 .
checkpoints is dramatically more difficult, but we provide a holistic approach: we show how to compute the optimal solution for any discrete probability distribution, using a sophisticated dynamic programming algorithm. Then we show how to approximate the optimal solution for any continuous probability distribution with bounded support, by providing a reservation sequence (and its checkpointing decisions) whose expected cost is arbitrarily close to the optimal one. In practice, the restriction to bounded support is not a limitation. Given, say, a Lognormal or Weibull probability distribution defined on $[a, \infty)$, it is very natural to truncate it on a bounded interval $[a, b]$ where $b$ corresponds to the quantile $Q(1-\varepsilon)$ for a small value of $\varepsilon$. This amount to discarding job execution times that are unreasonably too long, and never encountered in practice.

The main contributions of this work are the following:

- The characterization of an optimal reservation sequence, together with its checkpointing decisions, for any discrete probability distribution, using a sophisticated dynamic programming algorithm.
- An approximation of the optimal solution for any continuous probability distribution with bounded support, by providing an algorithm to compute a reservation sequence (and its checkpointing decisions) whose expected cost is arbitrarily close to the optimal one.
- An extensive set of simulation results based on execution times from nine probability distributions and from neuroscience application traces.
- Experimentation with neuroscience applications on a manycore platform, showing the efficiency of our strategies in a HPC environment.

The rest of the paper is organized as follows. Section 2 introduces the framework and main notations, and provides a detailed formula for the expected cost of a reservation sequence ant its checkpointing decisions. Section 3 describes our key algorithmic contributions. Section 4 is devoted to experimental evaluation and comparison with existing approaches through simulation process. Section 5 establishes performance evaluation of real applications in an HPC framework. Section 6 presents the related work. Finally, we provide concluding remarks and hints for future work in Section 7

## 2 Framework

In this section, we introduce some notations and formally define the optimization problem under study.

### 2.1 Stochastic jobs

We consider stochastic jobs whose execution times are unknown but (i) deterministic with respect to input data, so that two successive executions of the same job will have the same duration; and (ii) randomly and uniformly sampled from a given probability distribution law $\mathcal{D}$, whose density function (PDF) is $f$ and cumulative distribution function (CDF) is $F$. The probability distribution is assumed to be nonnegative, since we model execution times, and it is defined either on a finite support $[a, b]$, where $0 \leq a<b$, or on an infinite support $[a, \infty)$ where $a \geq 0$.

Hence, the execution time of a job is a random variable $X$, and

$$
\mathbb{P}(X \leq T)=F(T)=\int_{a}^{T} f(t) d t
$$

For notational convenience, we sometimes extend the domain of $f$ outside the support of $\mathcal{D}$ by letting $f(t)=0$ for $t \in[0, a] \cup[b, \infty)$.

In addition, we assume that we can interrupt the jobs at any time (divisible load application) to take a checkpoint: this will save the current progress of the execution, and enable to restart from that point on. Divisible load applications can be found, for example, in biological computations, image and video processing [20]. We assume that the cost of checkpoint and of recovery is constant throughout the execution: let $C$ be the cost to checkpoint the data at the end of an execution, and $R$ the cost to read the data to restart a computation.

Remark 1. By setting either $C$ or $R$ to $\infty$, then it is never useful to checkpoint, hence this problem can be reduced to the problem without checkpointing.

### 2.2 Cost model

We use the cost model motivated in our previous work [3]. For a reservation of length $W$ and an actual execution duration $w$ for the job, the cost is $\alpha W+\beta \min (W, w)+\gamma$, where $\alpha>0, \beta \geq 0$ and $\gamma \geq 0$. If the job does not complete within $W$ seconds, then another reservation should be paid for.

However, we take checkpoints into account in this work. If the job did not complete its execution during the last reservation, but was checkpointed during the last $C$ seconds of that reservation, then in the current reservation, the job can restart from that checkpoint during the first $R$ seconds, and then continue execution from its saved state. On the contrary, if no checkpoint was taken during the last reservation, the work done during that reservation is lost, and the execution must restart from the last checkpoint (or from the very beginning if no checkpoint was taken yet).

Altogether, the user needs to schedule a (possibly infinite) sequence of reservations $\mathcal{W}=$ $\left(W_{1}, W_{2}, \ldots, W_{i}, W_{i+1}, \ldots\right)$ to execute any job whose execution time follows the distribution $\mathcal{D}$, and to launch these reservations one after the other, until the job successfully terminates within the duration of some reservation. In addition, the user should decide whether to take a checkpoint or not at the end of each reservation.


Figure 3: Graphical representation of elapsed time for the reservation sequence $\mathcal{S}=$ $\left\{\left(W_{1}, 1\right),\left(W_{2}, 0\right),\left(W_{3}, 1\right),\left(W_{4}, 0\right)\right\}$.


Figure 4: Graphical representation of job progress (and showing $t_{k}$ versus $T_{k}$ ) for the reservation sequence $\mathcal{S}=\left\{\left(W_{1}, 1\right),\left(W_{2}, 0\right),\left(W_{3}, 1\right),\left(W_{4}, 0\right)\right\}$.

## Reservation-based strategy

Definition 1 (Reservation sequence for $\mathcal{D}$ ). Given a probability distribution $\mathcal{D}$, a reservation sequence $\mathcal{S}=\left\{\left(W_{1}, \delta_{1}\right),\left(W_{2}, \delta_{2}\right), \ldots\right\}$, is defined as a sequence of reservation lengths $W_{k}$ and a sequence of checkpointing decisions $\delta_{k} \in\{0,1\}: \delta_{k}=1$ means the $k^{\text {th }}$ reservation ends with a checkpoint, and $\delta_{k}=0$ means it does not.

Then, the $k^{\text {th }}$ reservation can be decomposed into:

$$
\begin{equation*}
W_{k}=R_{k}+T_{k}+C_{k} \tag{2}
\end{equation*}
$$

where $R_{k}$ is the time spent for restart, $T_{k}$ for actual job execution, and $C_{k}$ for checkpoint. We have $C_{k}=\delta_{k} C$ by definition. There is a restart if and only if there has been a checkpoint at some point before, hence

$$
R_{k}=\left(1-\prod_{i=1}^{k-1}\left(1-\delta_{i}\right)\right) R
$$

(assuming $R_{1}=0$ for the first reservation).
It is hard to keep track of actual job progress when using only the $\left(W_{k}, \delta_{k}\right)$ values. Consider for instance the following sequence $\mathcal{S}=\left\{\left(W_{1}, 1\right),\left(W_{2}, 0\right),\left(W_{3}, 1\right),\left(W_{4}, 0\right)\right\}$, which is depicted in Figure 3. If the actual job duration is $X=t$, during which reservation will the job complete its execution? We introduce another view of the reservation sequence $\mathcal{S}$ by introducing the milestones $\left(t_{k}\right.$ 's $)$ as shown in Figure 4. A milestone $t_{k}$ represents the amount of work that has actually been executed at the end of the $k^{\text {th }}$ reservation. Then, the last reservation for the job of length $t$ is $W_{k}$, where $t_{k-1} \leq t \leq t_{k}$. Of course, we need that $t \leq t_{4}$ for all values of $\mathcal{D}$ (equivalently, the upper bound of the support of $\mathcal{D}$ is $b \leq t_{4}$ ) for all jobs to complete successfully with the four reservations of $\mathcal{S}$.

The relationship between the milestone $t_{k}$ (actual work progress) and the value of $T_{k}$ (time spent computing during reservation $W_{k}$; see Equation (2)) is the following:

$$
\begin{equation*}
t_{k}=T_{k}+\sum_{i=1}^{k-1} \delta_{i} T_{i} \tag{3}
\end{equation*}
$$

Indeed, the work actually progresses only from the last checkpoint, while the work executed during the previous non-checkpointed reservations is lost whenever these non-checkpointed reservations do not allow for the full completion of the job. Another way to express the relationship between $t_{k}$ and $T_{k}$ is the following:

$$
\begin{equation*}
t_{k}=T_{k}+\max \left\{t_{i} \mid 1 \leq i \leq k-1 \text { and } \delta_{i}=1\right\} \tag{4}
\end{equation*}
$$

Indeed, Equation (4) gives a recursive way to compute $t_{k}$ from its definition. We recapitulate the relations between all notations introduced in Figures 3 and 4 :

$$
\begin{align*}
W_{k} & =R_{k}+T_{k}+C_{k}  \tag{5}\\
R_{k} & =\left(1-\Pi_{i<k}\left(1-\delta_{i}\right)\right) R  \tag{6}\\
T_{k} & =t_{k}-\sum_{i<k} \delta_{i} T_{i} \\
& =t_{k}-\max \left\{t_{i} \mid 1 \leq i \leq k-1 \text { and } \delta_{i}=1\right\}  \tag{7}\\
C_{k} & =\delta_{k} C \tag{8}
\end{align*}
$$

In the following, we use milestones $t_{k}$ rather than reservation lengths $W_{k}$ to characterize a reservation sequence, and we write

$$
\mathcal{S}=\left\{\left(t_{1}, \delta_{1}\right),\left(t_{2}, \delta_{2}\right), \ldots\right\}
$$

instead of

$$
\mathcal{S}=\left\{\left(W_{1}, \delta_{1}\right),\left(W_{2}, \delta_{2}\right), \ldots\right\}
$$

because it is easier to use milestones when computing the expected cost of a sequence, as shown below. For notational convenience, we define $t_{0}=0$ as the first milestone of each sequence $\mathcal{S}$. Note also that we can restrict to sequences where $t_{k-1}<t_{k}$, because otherwise (if $t_{k-1}=t_{k}$ ), the execution does not progress during the $k^{\text {th }}$ reservation.

### 2.3 Expected cost

Given a reservation sequence $\mathcal{S}=\left\{\left(t_{i}, \delta_{i}\right)\right\}_{i}$ and a job with execution time $t$ such that $t_{k-1}<$ $t \leq t_{k}$, the cost of the sequence for that job is given by:

$$
\begin{equation*}
C_{\mathcal{S}}(k, t)=\sum_{i=1}^{k-1}\left(\alpha W_{i}+\beta W_{i}+\gamma\right)+\alpha W_{k}+\beta\left(R_{k}+t-\left(t_{k}-T_{k}\right)\right)+\gamma \tag{9}
\end{equation*}
$$

where the first part is the total cost from the $k-1$ first reservations that did not allow the job to complete, and the second part is the cost of the $k^{\text {th }}$ reservation. The actual execution time during the $k^{\text {th }}$ reservation is $t-\left(t_{k}-T_{k}\right)$, because $t_{k}-T_{k}$ is the amount of work done up to the beginning of that reservation; we add the restart time $\left(R_{k}\right)$ but do not need to checkpoint (if $\delta_{k}=1$ ) because the job successfully completes before it is taken.

We let $k(t)=k$ for a job of length $t$ such that $t_{k-1}<t \leq t_{k}$. Now, the expected cost of the reservation sequence $\mathcal{S}$ over a job whose execution time is a random variable $X$ is

$$
\begin{equation*}
\mathbb{E}(\mathcal{S}(X))=\int_{0}^{\infty} C_{\mathcal{S}}(k(t), t) f(t) d t=\sum_{k=1}^{\infty} \int_{t_{k-1}}^{t_{k}} C_{\mathcal{S}}(k, t) f(t) d t \tag{10}
\end{equation*}
$$

### 2.4 Optimization Problem

We are now ready to state the optimization problem:

Definition 2 (Stochastic). Given a random variable $X$ following the distribution $\mathcal{D}$ (with PDF $f$ and CDF $F$ ) for the execution times of stochastic jobs, and given a cost function given by Equation (9) (with parameters $\alpha>0, \beta \geq 0$ and $\gamma \geq 0$ ), find a reservation strategy $\mathcal{S}$ with minimal expected cost $\mathbb{E}(\mathcal{S}(X))$ as given in Equation 10 .

We further define ReservationOnly to be the instance of Stochastic where the cost is a linear function of the reservation length only, i.e., when $\beta=\gamma=0$. For ReservationOnly, we can further consider $\alpha=1$ without loss of generality. For instance, such costs are incurred when making reservations of resources to schedule jobs on some cloud platforms, with hourly or daily rates. Throughout the paper, we focus on the usual probability distributions, hence we assume that the density function $f$ and the CDF $F$ of $\mathcal{D}$ are smooth (infinitely differentiable), and that $\mathcal{D}$ has finite expectation.

## 3 Algorithms

In this section, we establish some key properties of an optimal solution in the general setting.

### 3.1 Expected cost

We start by establishing a simpler expression for the expected cost function of Stochastic.
Theorem 1. Given a random variable $X$ and a reservation sequence $\mathcal{S}=\left\{\left(t_{1}, \delta_{1}\right),\left(t_{2}, \delta_{2}\right), \ldots\right\}$, the expected cost $\mathbb{E}(\mathcal{S}(X))$ of a strategy $\mathcal{S}$ given by Equation 10 , with parameters $\alpha, \beta$ and $\gamma$, can be rewritten as:

$$
\begin{align*}
\mathbb{E}(\mathcal{S}(X))= & \beta \cdot \mathbb{E}[X]+\alpha\left(t_{1}+\delta_{1} C\right)+\gamma \\
& +\sum_{i=2}^{\infty}\left(\alpha W_{i}+\beta\left(R_{i}+\left(1-\delta_{i-1}\right) T_{i-1}+C_{i-1}\right)+\gamma\right) \cdot \mathbb{P}\left(X>t_{i-1}\right) \tag{11}
\end{align*}
$$

For simplicity, when there is no ambiguity on the random variable $X$, we denote $\mathbb{E}(\mathcal{S}(X))=\mathbb{E}(\mathcal{S})$.
Proof. Firstly we rewrite Equation as follows:

$$
\begin{equation*}
\mathbb{E}(\mathcal{S})=\beta \cdot \mathbb{E}[X]+\sum_{i=1}^{\infty}\left(\alpha W_{i}+\beta\left(R_{i}+\left(1-\delta_{i-1}\right) T_{i-1}+C_{i-1}\right)+\gamma\right) \cdot \mathbb{P}\left(X>t_{i-1}\right) \tag{12}
\end{equation*}
$$

with initialization $\delta_{0}=W_{0}=R_{1}=0$ and $t_{0}=0$.
From Equations (9) and (10), we have

$$
\begin{equation*}
\mathbb{E}(\mathcal{S})=\mathbb{E}_{1}+\mathbb{E}_{2}+\mathbb{E}_{3} \tag{13}
\end{equation*}
$$

where

$$
\begin{aligned}
& \mathbb{E}_{1}=\sum_{k=1}^{\infty} \int_{t_{k-1}}^{t_{k}}\left(\sum_{i=1}^{k}\left(\alpha W_{i}+\gamma\right)\right) f(t) d t \\
& \mathbb{E}_{2}=\sum_{k=1}^{\infty} \int_{t_{k-1}}^{t_{k}}\left(\sum_{i=1}^{k-1} \beta W_{i}\right) f(t) d t \\
& \mathbb{E}_{3}=\sum_{k=1}^{\infty} \int_{t_{k-1}}^{t_{k}} \beta\left(t+R_{k}+T_{k}-t_{k}\right) f(t) d t
\end{aligned}
$$

Using $t_{0}=0$, we compute the first term as

$$
\begin{aligned}
\mathbb{E}_{1} & =\sum_{k=1}^{\infty} \sum_{i=1}^{k}\left(\alpha W_{i}+\gamma\right) \int_{t_{k-1}}^{t_{k}} f(t) d t \\
& =\sum_{k=1}^{\infty} \sum_{i=1}^{k}\left(\alpha W_{i}+\gamma\right) \cdot \mathbb{P}\left(t_{k-1}<X \leq t_{k}\right) \\
& =\sum_{i=1}^{\infty}\left(\alpha W_{i}+\gamma\right) \sum_{k=i}^{\infty} \mathbb{P}\left(t_{k-1}<X \leq t_{k}\right) \\
& =\sum_{i=1}^{\infty}\left(\alpha W_{i}+\gamma\right) \cdot \mathbb{P}\left(X>t_{i-1}\right)
\end{aligned}
$$

Similarly, using $W_{0}=0$, we express the second term as:

$$
\mathbb{E}_{2}=\sum_{i=1}^{\infty} \beta W_{i} \cdot \mathbb{P}\left(X>t_{i}\right)
$$

Finally, we derive the third term as:

$$
\begin{aligned}
\mathbb{E}_{3} & =\int_{t_{0}}^{\infty} \beta t f(t) d t+\sum_{k=1}^{\infty} \int_{t_{k-1}}^{t_{k}} \beta\left(R_{k}+T_{k}-t_{k}\right) f(t) d t \\
& =\beta \cdot \mathbb{E}[X]+\sum_{i=1}^{\infty} \beta\left(R_{i}+T_{i}-t_{i}\right) \cdot \mathbb{P}\left(t_{i-1}<X \leq t_{i}\right)
\end{aligned}
$$

Plugging these three terms back into Equation (13), we get:

$$
\begin{align*}
\mathbb{E}(\mathcal{S}) & =\beta \cdot \mathbb{E}[X]+\sum_{i=1}^{\infty}\left(\alpha W_{i}+\beta W_{i-1}+\gamma\right) \cdot \mathbb{P}\left(X>t_{i-1}\right)+\sum_{i=1}^{\infty} \beta\left(R_{i}+T_{i}-t_{i}\right) \cdot \mathbb{P}\left(t_{i-1}<X \leq t_{i}\right) \\
& =\beta \cdot \mathbb{E}[X]+\sum_{i=1}^{\infty}\left(\alpha W_{i}+\beta\left(R_{i}+T_{i-1}+C_{i-1}\right)+\gamma\right) \cdot \mathbb{P}\left(X>t_{i-1}\right)-\sum_{i=1}^{\infty} \beta\left(t_{i}-T_{i}\right) \cdot \mathbb{P}\left(t_{i-1}<X \leq t_{i}\right) \tag{14}
\end{align*}
$$

For the last derivation, we used

$$
\sum_{i=1}^{\infty}\left(R_{i-1} \cdot \mathbb{P}\left(X>t_{i-1}\right)+R_{i} \cdot \mathbb{P}\left(t_{i-1}<X \leq t_{i}\right)\right)=\sum_{i=1}^{\infty}\left(R_{i} \cdot \mathbb{P}\left(X>t_{i-1}\right)\right.
$$

Now, we study the second part of Equation (14) above. For all $j \leq 1$, let $\phi^{\circ}(j)$ denote the index of the $j^{\text {th }}$ checkpointed reservation of $\mathcal{S}$. For instance in the example of Figures 3 and 4 , $\phi^{o}(1)=1$ and $\phi^{o}(2)=3$. Then, using Equation (7), we have

$$
\begin{aligned}
\sum_{i=1}^{\infty}\left(t_{i}-T_{i}\right) \cdot \mathbb{P}\left(t_{i-1}<X \leq t_{i}\right) & =\sum_{j=1}^{\infty} t_{\phi(j)} \cdot \mathbb{P}\left(t_{\phi(j)}<X \leq t_{\phi(j+1)}\right) \\
& =\sum_{j=1}^{\infty} T_{\phi(j)} \cdot \mathbb{P}\left(X>t_{\phi(j)}\right) \\
& =\sum_{i=1}^{\infty} \delta_{i} T_{i} \cdot \mathbb{P}\left(X>t_{i}\right)
\end{aligned}
$$

Plugging the above back into Equation (14), we get the desired result shown in Equation (12).

### 3.2 Dynamic programming for discrete distributions

We study the problem for a finite discrete distribution: $Y \sim\left(v_{i}, f_{i}\right)_{1 \leq i \leq n}$, where $v_{i}<v_{i+1}$ for all $1 \leq i \leq n-1$ and $f_{i}=\mathbb{P}\left(Y=v_{i}\right)$. We assume that $f_{n} \neq 0$ and $\sum_{i=1}^{n} f_{i}=1$. Consider a strategy $\mathcal{S}=\left\{\left(t_{1}, \delta_{1}\right),\left(t_{2}, \delta_{2}\right), \ldots,\left(t_{|\mathcal{S}|}, \delta_{|\mathcal{S}|}\right)\right\}$, where $t_{i}=v_{\pi(i)}$ and $t_{i}<t_{i+1}$ for all $1 \leq i \leq|\mathcal{S}|-1$. Also, the last reservation is necessarily $t_{|\mathcal{S}|}=v_{n}$ to ensure that the expected cost of the strategy is finite. By convention, we let $t_{0}=v_{0}=a$, hence $\mathbb{P}\left(Y>t_{0}\right)=1$. Note that we can safely restrict to strategies where each milestone $t_{i}$ is equal to some threshold $v_{j}$ of the discrete distribution: otherwise, replacing $t_{i}$ by the largest $v_{j}$ such that $v_{j} \leq t_{i}$ leads to a smaller cost.

Rewriting Equation (11) with $W_{i}=R_{i}+T_{i}+C_{i}$, and since $W_{0}=0$, the expected cost of strategy $\mathcal{S}$ can be expressed as:

$$
\begin{align*}
\mathbb{E}(\mathcal{S})= & \beta \cdot \mathbb{E}[Y] \\
& +\sum_{i=1}^{|\mathcal{S}|}\left(\alpha\left(R_{i}+T_{i}+C_{i}\right)+\beta R_{i}+\gamma\right) \cdot \mathbb{P}\left(Y>t_{i-1}\right) \\
& +\sum_{i=1}^{|\mathcal{S}|-1} \beta\left(\left(1-\delta_{i}\right) T_{i}+C_{i}\right) \cdot \mathbb{P}\left(Y>t_{i}\right) \tag{15}
\end{align*}
$$

Based on Equation (15), and using Equations (6) to (8), we construct a dynamic programming algorithm to compute the optimal reservation sequence:

Theorem 2. For a discrete distribution $Y \sim\left(v_{i}, f_{i}\right)_{1 \leq i \leq n}$, the optimal expected cost is returned
by $\mathbb{E}_{\text {ckpt }}(0,0)$, where, for $0 \leq i_{c} \leq i_{l} \leq n, \mathbb{E}_{\text {ckpt }}\left(i_{c}, i_{l}\right)$ is:

$$
\begin{aligned}
& =\beta \cdot \mathbb{E}[Y], \quad \text { if } i_{l}=n \\
& =\min _{\substack{i_{l}+1 \leq j \leq n, \Delta_{j} \in\{0,1\}}}\left(\mathbb{E}_{\text {ckpt }}\left(\Delta_{j} j, j\right)+\left(\alpha\left(v_{j}+\Delta_{j} C\right)+\gamma\right) \cdot \sum_{k=i_{l}+1}^{n} f_{k}+\beta\left(\left(1-\Delta_{j}\right) v_{j}+\Delta_{j} C\right) \cdot \sum_{k=j+1}^{n} f_{k}\right) \\
& \text { if } i_{c}=0 \\
& =\min _{\substack{i_{l}+1 \leq j \leq n, \Delta_{j} \in\{0,1\}}}\left(\mathbb{E}_{\text {ckpt }}\left(\left(1-\Delta_{j}\right) i_{c}+\Delta_{j} j, j\right)+\left(\alpha\left(R+\left(v_{j}-v_{i_{c}}\right)+\Delta_{j} C\right)+\beta R+\gamma\right) \cdot \sum_{k=i_{l}+1}^{n} f_{k}\right. \\
& \left.+\beta\left(\left(1-\Delta_{j}\right)\left(v_{j}-v_{i_{c}}\right)+\Delta_{j} C\right) \cdot \sum_{k=j+1}^{n} f_{k}\right), \quad \quad \text { otherwise }
\end{aligned}
$$

The optimal solution can be computed in $O\left(n^{3}\right)$ time.
Intuitively, $i_{c}$ denotes the index of the last checkpointed value, while $i_{l}$ denotes the index of the last value that was tried before we try the next one with index $j$. Here, $\Delta_{j}$ indicates whether the value $v_{j}$ will be checkpointed or not.

Proof. To prove the optimality, consider $\mathbb{E}(\mathcal{S})$ given in Equation for any reservation sequence:

$$
\begin{aligned}
\mathcal{S} & =\left\{\left(t_{1}, \delta_{1}\right), \ldots,\left(t_{|\mathcal{S}|}, \delta_{|\mathcal{S}|}\right)\right\} \\
& =\left\{\left(v_{\pi(1)}, \Delta_{\pi(1)}\right), \ldots,\left(v_{\pi(|\mathcal{S}|)}, \Delta_{\pi(|\mathcal{S}|)}\right)\right\}
\end{aligned}
$$

and define $\mathbb{E}_{\ell}$ as the following partial sum:

$$
\begin{align*}
\mathbb{E}_{\ell}= & \beta \cdot \mathbb{E}[Y] \\
& +\sum_{i=\ell+1}^{|\mathcal{S}|}\left(\alpha\left(R_{i}+T_{i}+C_{i}\right)+\beta R_{i}+\gamma\right) \cdot \mathbb{P}\left(Y>t_{i-1}\right) \\
& +\sum_{i=\ell+1}^{|\mathcal{S}|-1} \beta\left(\left(1-\delta_{i}\right) T_{i}+C_{i}\right) \cdot \mathbb{P}\left(Y>t_{i}\right) \tag{16}
\end{align*}
$$

Note that $\mathbb{E}_{0}=\mathbb{E}(\mathcal{S})$. We show by induction that the following invariant is true for all $\ell=|\mathcal{S}|,|\mathcal{S}|-1, \ldots, 0$ :

$$
\begin{equation*}
\mathbb{E}_{\ell} \geq \mathbb{E}_{\mathrm{ckpt}}(\pi(\bar{\ell}), \pi(\ell)) \tag{17}
\end{equation*}
$$

where $\pi(\bar{\ell})$ is the index of the last reservation not larger than $\pi(\ell)$ and such that $\Delta_{\pi(\bar{\ell})}=1$. We denote the corresponding reservation by $t_{\bar{\ell}}=v_{\pi(\bar{\ell})}$.

For the base case with $\ell=|\mathcal{S}|$, we have $\pi(|\mathcal{S}|)=n$, and $\mathbb{E}_{|\mathcal{S}|}=\beta \cdot \mathbb{E}[Y]=\mathbb{E}_{\text {ckpt }}(\pi(\overline{\mathcal{S} \mid}), n)$. Now, suppose $\mathbb{E}_{\ell+1} \geq \mathbb{E}_{\text {ckpt }}(\pi(\overline{\ell+1}), \pi(\ell+1))$ for $\ell+1 \leq|\mathcal{S}|$. Here, we note that $\pi(\overline{\ell+1})=$ $\pi(\ell+1)$ if $\Delta_{\pi(\ell+1)}=1$ (i.e., if $v_{\pi(\ell+1)}$ is checkpointed). Otherwise, we have $\pi(\overline{\ell+1})=\pi(\bar{\ell})$.

Then, from Equation (16), we derive:

$$
\begin{aligned}
\mathbb{E}_{\ell}= & \mathbb{E}_{\ell+1}+\left(\alpha\left(R_{\ell+1}+T_{\ell+1}+C_{\ell+1}\right)+\beta R_{\ell+1}+\gamma\right) \cdot \mathbb{P}\left(Y>t_{\ell}\right) \\
& +\beta\left(\left(1-\delta_{\ell+1}\right) T_{\ell+1}+C_{\ell+1}\right) \cdot \mathbb{P}\left(Y>t_{\ell+1}\right) \\
\geq & \mathbb{E}_{\mathrm{ckpt}}(\pi(\overline{\ell+1}), \pi(\ell+1))+\left(\alpha\left(R_{\ell+1}+\left(t_{\ell+1}-t_{\bar{\ell}}\right)+C_{\ell+1}\right)+\beta R_{\ell+1}+\gamma\right) \cdot \mathbb{P}\left(Y>v_{\pi(\ell)}\right) \\
& \left.+\beta\left(\left(1-\delta_{\ell+1}\right)\left(t_{\ell+1}-t_{\bar{\ell}}\right)+C_{\ell+1}\right) \cdot \mathbb{P}\left(Y>v_{\pi(\ell+1}\right)\right) \\
= & \mathbb{E}_{\mathrm{ckpt}}\left(\left(1-\Delta_{\pi(\ell+1)}\right) \pi(\bar{\ell})+\Delta_{\pi(\ell+1)} \pi(\ell+1), \pi(\ell+1)\right) \\
& +\left(\alpha\left(\mathbb{1}_{\pi(\bar{\ell}) \neq 0} R+\left(v_{\pi(\ell+1)}-v_{\pi(\bar{\ell})}\right)+\Delta_{\pi(\ell+1)} C\right)+\beta \mathbb{1}_{\pi(\bar{\ell}) \neq 0} R+\gamma\right) \cdot \sum_{k=\pi(\ell)+1}^{n} f_{k} \\
& +\beta\left(\left(1-\Delta_{\pi(\ell+1)}\right)\left(v_{\pi(\ell+1)}-v_{\pi(\bar{\ell})}\right)+\Delta_{\pi(\ell+1)} C\right) \cdot \sum_{k=\pi(\ell+1)+1}^{n} f_{k} \\
\geq & \mathbb{E}_{\mathrm{ckpt}}(\pi(\bar{\ell}), \pi(\ell))
\end{aligned}
$$

In the derivation above, the first inequality is due to the inductive hypothesis, and using $T_{\ell+1}=t_{\ell+1}-t_{\bar{\ell}}$ from Equation (7). The last inequality is due to the definition of $\mathbb{E}_{\mathrm{ckpt}}$. s Thus, by induction, we get $\mathbb{E}_{\text {ckpt }}(0,0)=\mathbb{E}_{\text {ckpt }}(\pi(\overline{0}), \pi(0)) \leq \mathbb{E}_{0}=\mathbb{E}(\mathcal{S})$. This shows that $\mathbb{E}_{\text {ckpt }}(0,0)$ is not greater than the expected cost of any reservation sequence $\mathcal{S}$, thus returning the optimal solution.

Finally, one can pre-compute values of $\sum_{k=\ell+1}^{n} f_{k}$ for all $0 \leq \ell<n$ (in linear time) and store them. Then, computing $\mathbb{E}_{\text {ckpt }}\left(i_{c}, i_{l}\right)$ depends on $2\left(n-i_{l}\right)$ other $\mathbb{E}_{\text {ckpt }}$ values, thus takes $O\left(n-i_{l}\right)$ time. The overall complexity is therefore $O\left(n^{3}\right)$.

### 3.3 Approximation algorithm for continuous distributions

In this section, we provide an approximation algorithm of the optimal strategy for continuous distributions with bounded support $[a, b]$, where $a \geq 0$ and $b$ is finite. Because we model job execution times, it is natural to truncate continuous distributions whose support is $[0, \infty[$ such as an Exponential or Lognormal distribution, say, to a bounded support $[a, b]$.

The result for continuous distribution is particularly important: we have shown in recent work $[10$ that continuous distributions gave strategies that allowed using small data samples to find an efficient strategy. Here, it returns an arbitrarily good quality solution with low complexity.

More precisely, let $X$ be a continuous random variable defined on $[a, b]$ modeling the probability distribution $\mathcal{D}$, where $0 \leq a<b$, with CDF $F$ and PDF $f$. Theorem 3 shows that Algorithm 1 computes a close-to-optimal strategy for Stochastic.

We start by showing the following Lemma:
Lemma 1. Given a random variable $X$ and a strategy $\mathcal{S}=\left\{\left(t_{1}, \delta_{1}\right), \ldots,\left(t_{|S|}, \delta_{|S|}\right)\right\}$, if there exists an index $i_{0}>1$ such that $t_{1}<\cdots<t_{i_{0}-1} \leq \min (R, \varepsilon \mathbb{E}[X])<t_{i_{0}}<\cdots<t_{|S|}$, then the strategy $\tilde{\mathcal{S}}=\left\{(\min (R, \varepsilon \mathbb{E}[X]), 0),\left(t_{i_{0}}, \delta_{i_{0}}\right), \ldots,\left(t_{|S|}, \delta_{|S|}\right)\right\}$ satisfies:

$$
\mathbb{E}(\tilde{\mathcal{S}}(X)) \leq(1+\varepsilon) \cdot \mathbb{E}(\mathcal{S}(X))
$$

Intuitively, this lemma states that restricting to strategies such that the first reservation is at least $\min (R, \varepsilon \mathbb{E}[X])$ only increases the cost by at most a factor of $1+\varepsilon$.

```
Algorithm 1 Dyn-Prog-Count \((X, \varepsilon)\)
    Let \([a, b]\) be the domain of \(X\), with \(0 \leq a<b\)
    \(c_{0}=3(b-a) \min \left(\frac{1}{\min (\max (a, \varepsilon \mathbb{E}[X] / 3), R, C)}, \frac{\alpha+\beta}{\gamma}\right)\)
    \(n \leftarrow\left\lceil c_{0} / \varepsilon\right\rceil\)
    Define the discrete distribution \(Y_{n} \sim\left(v_{i}, f_{i}\right)_{i=1 \ldots n}\) s.t.
```

$$
\left\{\begin{aligned}
v_{i} & =a+i \cdot \frac{b-a}{n} & & \text { for } 0 \leq i \leq n \\
f_{i} & =\mathbb{P}\left(Y_{n}=v_{i}\right)=\mathbb{P}\left(v_{i-1}<X \leq v_{i}\right) & & \text { for } 1 \leq i \leq n
\end{aligned}\right.
$$

$\mathcal{S}_{n}^{\mathrm{dp}} \leftarrow$ Optimal strategy for $Y_{n}$ (Theorem 2 )
return $\mathcal{S}_{n}^{\text {dp }}$

Proof. Consider a strategy $\mathcal{S}=\left\{\left(t_{1}, \delta_{1}\right), \ldots,\left(t_{|S|}, \delta_{|S|}\right)\right\}$ for a random variable $X$, such that there exists an index $i_{0}>1$ with

$$
t_{1}<\cdots<t_{i_{0}-1} \leq \min (R, \varepsilon \mathbb{E}[X])<t_{i_{0}}<\cdots<t_{|S|}
$$

For simplicity, we denote by $\tilde{a}=\min (R, \varepsilon \mathbb{E}[X])$, and define strategy $\tilde{\mathcal{S}}=\left\{(\tilde{a}, 0),\left(t_{i_{0}}, \delta_{i_{0}}\right), \ldots,\left(t_{|S|}, \delta_{|S|}\right)\right\}$. From Equation we have:

$$
\begin{aligned}
& \mathbb{E}(\mathcal{S}(X)) \geq \beta \mathbb{E}[X]+\alpha\left(t_{1}+C_{1}\right)+\gamma+\left(\alpha W_{i_{0}}+\beta\left(R_{i_{0}}+\left(1-\delta_{i_{0}-1}\right) T_{i_{0}-1}+C_{i_{0}-1}\right)+\gamma\right) \cdot \mathbb{P}\left(X>t_{i_{0}-1}\right) \\
&+\sum_{i=i_{0}+1}^{|\mathcal{S}|}\left(\alpha W_{i}+\beta\left(R_{i}+\left(1-\delta_{i-1}\right) T_{i-1}+C_{i-1}\right)+\gamma\right) \cdot \mathbb{P}\left(X>t_{i-1}\right) \\
& \mathbb{E}(\tilde{\mathcal{S}}(X))=\beta \mathbb{E}[X]+(\alpha \tilde{a}+\gamma)+\left(\alpha \tilde{W}_{i_{0}}+\beta \tilde{a}+\gamma\right) \cdot \mathbb{P}(X>\tilde{a}) \\
&+\sum_{i=i_{0}+1}^{|\mathcal{S}|}\left(\alpha \tilde{W}_{i}+\beta\left(\tilde{R}_{i}+\left(1-\delta_{i-1}\right) \tilde{T}_{i-1}+\tilde{C}_{i-1}\right)+\gamma\right) \cdot \mathbb{P}\left(X>t_{i-1}\right)
\end{aligned}
$$

We obviously have $C_{i}=\tilde{C}_{i}, \forall i \geq i_{0}$. We now show the following property: $\forall i \geq i_{0}, W_{i} \geq \tilde{W}_{i}$. To show this, we consider two cases: (i) the last checkpoint before $t_{i}$ was done during $t_{j}$ with $j \geq i_{0}$ or there was no checkpoint before $t_{i}$. In this case, we obviously have $W_{i}=\tilde{W}_{i}$; (ii) the last checkpoint before $t_{i}$ was done during $t_{j}$ with $j<i_{0}$ in $\mathcal{S}$, and there was no checkpoint done in $\tilde{\mathcal{S}}$ before $t_{i}$. In this case, we have $W_{i}=R+\left(t_{i}-t_{j}\right)+\delta_{i} C$ and $\tilde{W}_{i}=t_{i}+\delta_{i} C$. Since $t_{j} \leq t_{i_{0}-1} \leq R$, we get $W_{i} \geq \tilde{W}_{i}$. Similarly, we can show that, $\forall i \geq i_{0}, R_{i} \geq \tilde{R}_{i}$.

Further, since $\mathbb{P}(X>\tilde{a}) \leq \mathbb{P}\left(X>t_{i_{0}-1}\right) \leq 1$, we can now derive that:

$$
\begin{aligned}
\mathbb{E}(\tilde{\mathcal{S}}(X))-\mathbb{E}(\mathcal{S}(X)) \leq & \alpha\left(\tilde{a}-t_{1}-C_{1}\right) \\
& +\beta\left(\tilde{a}-R_{i_{0}}-\left(1-\delta_{i_{0}-1}\right) T_{i_{0}-1}-C_{i_{0}-1}\right) \cdot \mathbb{P}\left(X>t_{i_{0}-1}\right) \\
\leq & (\alpha+\beta) \tilde{a} \\
\leq & \varepsilon(\alpha+\beta) \mathbb{E}[X]
\end{aligned}
$$

Finally, note that we immediately have $\mathbb{E}(\mathcal{S}(X)) \geq(\alpha+\beta) \mathbb{E}[X]+\gamma$, because this is the cost of an omniscient strategy that makes a single reservation of exactly the right size for each job. Therefore, we get the result:

$$
\mathbb{E}(\tilde{\mathcal{S}}(X))-\mathbb{E}(\mathcal{S}(X)) \leq \varepsilon \cdot \mathbb{E}(\mathcal{S}(X))
$$

which completes the proof of Lemma 1.

Theorem 3. Given a continuous random variable $X$ on the domain $[a, b]$, where $0 \leq a<b$, and given a constant $\varepsilon>0$, Dyn-Prog-Count $(X, \varepsilon)$ is a $1+\varepsilon$-approximation algorithm for Stochastic and executes in time $\mathcal{O}\left(\frac{1}{\varepsilon^{3}}\right)$.
Proof. Given a continuous random variable $X$ of support $[a, b]$, we define the discrete random variable $Y_{n} \sim\left(v_{i}, f_{i}\right)_{i=1 \ldots n}$ as stated in Algorithm 1;

$$
\left\{\begin{array}{rlr}
v_{i}=a+i \cdot \frac{b-a}{n} & & \text { for } 0 \leq i \leq n  \tag{18}\\
f_{i}=\mathbb{P}\left(Y_{n}=v_{i}\right)=\mathbb{P}\left(v_{i-1}<X \leq v_{i}\right) & & \text { for } 1 \leq i \leq n
\end{array}\right.
$$

Let $\mathcal{S}^{\text {opt }}=\left\{\left(\tilde{t}_{i}^{o}, \tilde{\delta}_{i}^{o}\right)\right\}_{1 \leq i \leq\left|\mathcal{S}^{\circ p \mathrm{p}}\right|}$ denote the optimal solution for $X$, and let $\mathcal{S}_{n}^{\text {dp }}$ denote the optimal solution for $Y_{n}$ returned by Theorem 2. We want to show that

$$
\mathbb{E}\left(\mathcal{S}_{n}^{\mathrm{dp}}(X)\right) \leq(1+\varepsilon) \cdot \mathbb{E}\left(\mathcal{S}^{\mathrm{opt}}(X)\right)
$$

In order to do that, we construct two intermediate strategies $\mathcal{S}_{\varepsilon / 3}^{\text {opt }}$ and $\mathcal{S}^{\text {algo }}$ as follows.
First, $\mathcal{S}_{\varepsilon / 3}^{\text {opt }}=\left(\left(t_{i}^{o}, \delta_{i}^{o}\right)\right)_{i}$ is constructed in such a way that if $\tilde{t}_{1}^{o} \geq \min \left(R, \frac{\varepsilon \mathbb{E}[X]}{3}\right)$, then $\mathcal{S}_{\varepsilon / 3}^{\text {opt }}=$ $\mathcal{S}^{\text {opt }}$, otherwise we construct $\mathcal{S}_{\varepsilon / 3}^{\text {opt }}$ from $\mathcal{S}^{\text {opt }}$ by following Lemma 1 (with the value $\frac{\varepsilon}{3}$ ). Then, according to Lemma 1 we have:

$$
\begin{equation*}
\mathbb{E}\left(\mathcal{S}_{\varepsilon / 3}^{\mathrm{opt}}(X)\right) \leq\left(1+\frac{\varepsilon}{3}\right) \cdot \mathbb{E}\left(\mathcal{S}^{\mathrm{opt}}(X)\right) \tag{19}
\end{equation*}
$$

Second, $\mathcal{S}^{\text {algo }}=\left(\left(t_{i}^{a}, \delta_{i}^{a}\right)\right)_{1 \leq i \leq\left|\mathcal{S}_{\varepsilon / 3}^{\text {opt }}\right|}$ (hence $\left.\left|\mathcal{S}^{\text {algo }}\right|=\left|\mathcal{S}_{\varepsilon / 3}^{\text {opt }}\right|\right)$, is such that for $1 \leq i \leq\left|\mathcal{S}_{\varepsilon / 3}^{\text {opt }}\right|$, we let $\left(t_{i}^{a}, \delta_{i}^{a}\right)=\left(v_{\pi^{o}(i)}, \delta_{i}^{o}\right)$. Here, we use the sequence $\left(v_{i}\right)_{i=0 \ldots n}$ from Equation 18), and the function $\pi^{o}$ defined below:

$$
\begin{equation*}
v_{\pi^{o}(i)-1}<t_{i}^{o} \leq v_{\pi^{o}(i)} \tag{20}
\end{equation*}
$$

In other words, for each reservation, $\mathcal{S}^{\text {algo }}$ chooses the first discrete value larger than or equal to the corresponding one chosen by $\mathcal{S}_{\varepsilon / 3}^{\text {opt }}$, and makes the same checkpointing decision.

Lemma 2. $\mathbb{E}\left(\mathcal{S}^{\text {algo }}(X)\right) \leq\left(1+\frac{\varepsilon}{3}\right) \cdot \mathbb{E}\left(\mathcal{S}_{\varepsilon / 3}^{\text {opt }}(X)\right)$.
Proof. We use the notations $T_{i}^{o}, R_{i}^{o}, C_{i}^{o}, W_{i}^{o}$ for the parameters of $\mathcal{S}_{\varepsilon / 3}^{\text {opt }}$, and $T_{i}^{a}, R_{i}^{a}, C_{i}^{a}, W_{i}^{a}$ for the parameters of $\mathcal{S}^{\text {algo }}$. From Equations (5) to (8), we see that, for $1 \leq i \leq\left|\mathcal{S}_{\varepsilon / 3}^{\text {opt }}\right|$, we have:

1. $\delta_{i}^{o}=\delta_{i}^{a}$;
2. $R_{i}^{o}=R_{i}^{a}$;
3. $C_{i}^{o}=C_{i}^{a}$; and
4. $W_{i}^{a}-W_{i}^{o}=T_{i}^{a}-T_{i}^{o}$.

In addition, if $\sigma^{o}(i)$ (resp. $\sigma^{a}(i)$ ) is the index of the last checkpoint before $t_{i}^{o}$ (resp. $t_{i}^{a}$ ), then $\sigma^{o}(i)=\sigma^{a}(i)$, and,

$$
\begin{aligned}
\left|T_{i}^{a}-T_{i}^{o}\right| & =\left|\left(t_{i}^{a}-t_{\sigma^{a}(i)}^{a}\right)-\left(t_{i}^{o}-t_{\sigma^{o}(i)}^{o}\right)\right| \\
& =\left|\left(v_{\pi^{o}(i)}-v_{\pi^{o}\left(\sigma^{o}(i)\right)}\right)-\left(t_{i}^{o}-t_{\sigma^{o}(i)}^{o}\right)\right| \\
& =\left|\left(v_{\pi^{o}(i)}-t_{i}^{o}\right)-\left(v_{\pi\left(\sigma^{o}(i)\right)}-t_{\sigma^{o}(i)}^{o}\right)\right| \\
& \leq \max \left(v_{\pi^{o}(i)}-t_{i}^{o}, v_{\pi^{o}\left(\sigma^{o}(i)\right)}-t_{\sigma^{o}(i)}^{o}\right) \leq \frac{b-a}{n}
\end{aligned}
$$

From Equation (11) we have:

$$
\begin{aligned}
& \mathbb{E}\left(\mathcal{S}_{\varepsilon / 3}^{\mathrm{opt}}(X)\right)=\beta \mathbb{E}[X]+\sum_{i=1}^{\left|\mathcal{S}_{\varepsilon / 3}^{\mathrm{opt}}\right|}\left(\alpha W_{i}^{o}+\beta\left(R_{i}^{o}+\left(1-\delta_{i-1}^{o}\right) T_{i-1}^{o}+C_{i-1}^{o}\right)+\gamma\right) \cdot \mathbb{P}\left(X>t_{i-1}^{o}\right) \\
& \mathbb{E}\left(\mathcal{S}^{\mathrm{algo}}(X)\right)=\beta \mathbb{E}[X]+\sum_{i=1}^{\left|\mathcal{S}_{\varepsilon / 3}^{\mathrm{ot}}\right|}\left(\alpha W_{i}^{a}+\beta\left(R_{i}^{a}+\left(1-\delta_{i-1}^{a}\right) T_{i-1}^{a}+C_{i-1}^{a}\right)+\gamma\right) \cdot \mathbb{P}\left(X>t_{i-1}^{a}\right)
\end{aligned}
$$

We first observe that $\mathbb{P}\left(X>t_{i-1}^{a}\right) \leq \mathbb{P}\left(X>t_{i-1}^{o}\right)$ because $t_{i-1}^{a} \geq t_{i-1}^{o}$. We can derive that:

$$
\begin{aligned}
\mathbb{E}\left(\mathcal{S}^{\mathrm{algo}}(X)\right)-\mathbb{E}\left(\mathcal{S}_{\varepsilon / 3}^{\mathrm{opt}}(X)\right) & \leq \sum_{i=1}^{\left|\mathcal{S}_{\varepsilon / 3}^{\mathrm{opt}}\right|}\left(\alpha\left|T_{i}^{a}-T_{i}^{o}\right|+\beta\left(1-\delta_{i-1}^{o}\right)\left|T_{i-1}^{a}-T_{i-1}^{o}\right|\right) \cdot \mathbb{P}\left(X>t_{i-1}^{o}\right) \\
& \leq \alpha \frac{b-a}{n}+\sum_{i=1}^{\left|\mathcal{S}_{\varepsilon / 3}^{\mathrm{opt} \mid}\right|-1}\left(\left(\alpha+\beta\left(1-\delta_{i}^{o}\right)\right) \frac{b-a}{n}\right) \cdot \mathbb{P}\left(X>t_{i}^{o}\right) \\
& \leq \frac{b-a}{n}\left(\alpha+(\alpha+\beta) \sum_{i=1}^{\left|\mathcal{S}_{\varepsilon / 3}^{\mathrm{opt}}\right|-1} \mathbb{P}\left(X>t_{i}^{o}\right)\right)
\end{aligned}
$$

We also observe that:

$$
\mathbb{E}\left(\mathcal{S}_{\varepsilon / 3}^{\mathrm{opt}}(X)\right) \geq \gamma+\sum_{i=1}^{\mid \mathcal{S}_{\varepsilon / 3}^{\mathrm{opt} \mid-1}} \gamma \cdot \mathbb{P}\left(X>t_{i}^{o}\right)
$$

Further, for $1 \leq i \leq\left|\mathcal{S}_{\varepsilon / 3}^{\text {opt }}\right|$, we have $W_{i}^{o} \geq R_{i}^{o}+T_{i}^{o} \geq \min (R, \tilde{a})$, where $\tilde{a}=\max (a, \min (R, \varepsilon \mathbb{E}[X] / 3))$. This is because either $T_{i}^{o} \geq \tilde{a}$ according to Lemma 1 (when there was no checkpoint before $t_{i}^{o}$ ), or $R_{i}^{o}=R$ (when there was a checkpoint before $t_{i}^{o}$ ). Therefore, we can derive:

$$
\mathbb{E}\left(\mathcal{S}_{\varepsilon / 3}^{\mathrm{opt}}(X)\right) \geq \min (\max (a, \varepsilon \mathbb{E}[X] / 3), R, C)\left(\alpha+(\alpha+\beta) \sum_{i=1}^{\left|\mathcal{S}_{\varepsilon / 3}^{\mathrm{opt}}\right|-1} \mathbb{P}\left(X>t_{i}^{o}\right)\right)
$$

Note that $\min (R, \max (a, \min (R, \varepsilon \mathbb{E}[X] / 3)))=\min (\max (a, \varepsilon \mathbb{E}[X] / 3), R)$.
Using the definition of $c_{0}=3(b-a) \min \left(\frac{1}{\min (\max (a, \varepsilon \mathbb{E}[X] / 3), R, C)}, \frac{\alpha+\beta}{\gamma}\right)$ in Algorithm 1 , we obtain:

$$
\begin{aligned}
\mathbb{E}\left(\mathcal{S}^{\text {algo }}(X)\right)-\mathbb{E}\left(\mathcal{S}_{\varepsilon / 3}^{\mathrm{opt}}(X)\right) & \leq \frac{c_{0}}{n} \cdot \mathbb{E}\left(\mathcal{S}_{\varepsilon / 3}^{\mathrm{opt}}(X)\right) \\
& \leq \frac{\varepsilon}{3} \cdot \mathbb{E}\left(\mathcal{S}_{\varepsilon / 3}^{\mathrm{opt}}(X)\right)
\end{aligned}
$$

which concludes the proof of Lemma 2.
Lemma 3. $\mathbb{E}\left(\mathcal{S}_{n}^{\operatorname{dp}}(X)\right) \leq \mathbb{E}\left(\mathcal{S}^{\text {algo }}(X)\right)$
Proof. Given any reservation strategy $\mathcal{S}=\left\{\left(t_{i}, \delta_{i}\right)\right\}_{1 \leq i \leq|\mathcal{S}|}$ such that $\forall i, t_{i} \in\left\{v_{1}, \ldots, v_{n}\right\}$, we show that:

$$
\mathbb{E}\left(\mathcal{S}\left(Y_{n}\right)\right)-\mathbb{E}(\mathcal{S}(X))=\beta\left(\mathbb{E}\left[Y_{n}\right]-\mathbb{E}[X]\right)
$$

Indeed, for the two distributions $Y_{n}$ and $X$, the only differences in the cost function are: (i) the expectations $\mathbb{E}\left[Y_{n}\right]$ and $\mathbb{E}[X]$; and (ii) the probability values $\mathbb{P}\left(Y_{n}>t_{i}\right)$ and $\mathbb{P}\left(X>t_{i}\right), \forall i$. But if $t_{i} \in\left\{v_{1}, \ldots, v_{n}\right\}$, we have:

$$
\begin{aligned}
\mathbb{P}\left(Y_{n}>t_{i}\right) & =\mathbb{P}\left(Y_{n}>v_{k}\right) \\
& =\mathbb{P}\left(Y_{n} \in \cup_{j=k+1}^{n}\left\{v_{j}\right\}\right)=\sum_{j=k+1}^{n} \mathbb{P}\left(Y_{n}=v_{j}\right) \\
& \left.\left.\left.\left.=\sum_{j=k+1}^{n} \mathbb{P}(X \in] v_{j-1}, v_{j}\right]\right)=\mathbb{P}(X \in] v_{k}, v_{n}\right]\right) \\
& =\mathbb{P}\left(X>v_{k}\right)=\mathbb{P}\left(X>t_{i}\right)
\end{aligned}
$$

We obtain that:

$$
\mathbb{E}\left(\mathcal{S}\left(Y_{n}\right)\right)-\mathbb{E}(\mathcal{S}(X))=\beta\left(\mathbb{E}\left[Y_{n}\right]-\mathbb{E}[X]\right)
$$

We apply this result to both $\mathcal{S}_{n}^{\mathrm{dp}}$ and $\mathcal{S}^{\text {algo }}$ and derive that:

$$
\mathbb{E}\left(\mathcal{S}_{n}^{\mathrm{dp}}\left(Y_{n}\right)\right)-\mathbb{E}\left(\mathcal{S}_{n}^{\mathrm{dp}}(X)\right)=\mathbb{E}\left(\mathcal{S}^{\text {algo }}\left(Y_{n}\right)\right)-\mathbb{E}\left(\mathcal{S}^{\text {algo }}(X)\right)
$$

or equivalently,

$$
\mathbb{E}\left(\mathcal{S}_{n}^{\mathrm{dp}}\left(Y_{n}\right)\right)-\mathbb{E}\left(\mathcal{S}^{\mathrm{algo}}\left(Y_{n}\right)\right)=\mathbb{E}\left(\mathcal{S}_{n}^{\mathrm{dp}}(X)\right)-\mathbb{E}\left(\mathcal{S}^{\mathrm{algo}}(X)\right)
$$

But $\mathcal{S}_{n}^{\mathrm{dp}}$ is optimal for $Y_{n}$, hence

$$
\mathbb{E}\left(\mathcal{S}_{n}^{\mathrm{dp}}\left(Y_{n}\right)\right)-\mathbb{E}\left(\mathcal{S}^{\mathrm{algo}}\left(Y_{n}\right)\right) \leq 0
$$

Therefore,

$$
\mathbb{E}\left(\mathcal{S}_{n}^{\mathrm{dp}}(X)\right)-\mathbb{E}\left(\mathcal{S}^{\text {algo }}(X)\right) \leq 0
$$

This concludes the proof of Lemma 3 .
Now, combining Lemma 2, Lemma 3 and Equation (19), we get:

$$
\begin{aligned}
\mathbb{E}\left(\mathcal{S}_{n}^{\mathrm{dp}}(X)\right) & \leq \mathbb{E}\left(\mathcal{S}^{\text {algo }}(X)\right) \\
& \leq\left(1+\frac{\varepsilon}{3}\right) \cdot \mathbb{E}\left(\mathcal{S}_{\varepsilon / 3}^{\mathrm{opt}}(X)\right) \\
& \leq\left(1+\frac{\varepsilon}{3}\right)\left(1+\frac{\varepsilon}{3}\right) \cdot \mathbb{E}\left(\mathcal{S}^{\mathrm{opt}}(X)\right) \\
& \leq(1+\varepsilon) \cdot \mathbb{E}\left(\mathcal{S}^{\mathrm{opt}}(X)\right)
\end{aligned}
$$

which concludes the proof of Theorem 3.

### 3.4 Extensions

All the results presented in Sections 3.1 to 3.3 namely the cost model (Theorem 1), the optimal algorithm for discrete distributions (Theorem [2], and the approximation algorithm for continuous distributions with bounded support (Theorem 3), can be extended to some variants of the problem where the checkpoint strategy is determined a priori.

Indeed, there are two important and natural variants to consider: strategies where no reservation is checkpointed, and strategies where all reservations are checkpointed. The former variant (called No-Checkpoint) was studied in our previous work [3], where we derived an optimal algorithm for discrete distributions with reduced time complexity $O\left(n^{2}\right)$ instead of $O\left(n^{3}\right)$ as in Theorem 2. The latter variant (called All-Checkpoint) also admits an optimal dynamic programming algorithm of reduced time complexity $O\left(n^{2}\right)$ :

Theorem 4. For a discrete distribution $Y \sim\left(v_{i}, f_{i}\right)_{1 \leq i \leq n}$, the optimal expected cost for All-Checkpoint (when all reservations are checkpointed) is returned by $\mathbb{E}_{\mathrm{AllCkpt}}(0)$, where $v_{0}=0$ and:
$\mathbb{E}_{\text {AllCkpt }}(n)=\beta \cdot \mathbb{E}[Y]$
$\mathbb{E}_{\text {AllCkpt }}(i)=\min _{i+1 \leq j \leq n}\left(\mathbb{E}_{\text {AllCkpt }}(j)+\beta C \cdot \sum_{k=j+1}^{n} f_{k}+\left(\alpha\left(\mathbb{1}_{i \neq 0} R+\left(v_{j}-v_{i}\right)+\mathbb{1}_{j \neq n} C\right)+\beta \mathbb{1}_{i \neq 0} R+\gamma\right) \cdot \sum_{k=i+1}^{n} f_{k}\right)$
The optimal solution can be computed in $O\left(n^{2}\right)$ time.

### 3.5 Periodic Heuristics

In addition to the algorithms presented in Section 3, we propose a periodic heuristic for the case of bounded distributions. This strategy, described in Algorithm 2, is a natural policy, where successive reservations differ in length by a constant amount of time $T$, called the period. A checkpoint is performed at the end of each period. Hence, the value of $W_{i}$ associated with each $t_{i}$ is constant in this strategy. The algorithm specifies the number of chunks $\tau$ in the domain $[a, b]$ of the bounded distribution, thus the period can be computed as $T=\frac{b-a}{\tau}$.

```
Algorithm 2 All-Checkpoint-Periodic \((X, \tau)\)
    Let \([a, b]\) be the domain of \(X\), and let \(T=\frac{b-a}{\tau}\)
    \(\left(t_{i}, \delta_{i}\right)= \begin{cases}(a+i \cdot T, 1) & \text { for } i=1,2, \ldots, \tau-1 \\ (b, 0) & \text { for } i=\tau\end{cases}\)
    3: return \(\mathcal{S}_{\tau}^{\text {period }} \leftarrow\left(\left(t_{i}, \delta_{i}\right)\right)_{1 \leq i \leq \tau}\)
```

For this policy, one can derive optimal strategies for some distributions such as Uniform distributions. One can also prove that All-Checkpoint and its periodic counterpart are identical. The next subsections are dedicated to proving those assertions.

### 3.5.1 All-Checkpoint-Periodic for Exponential distributions

In this section, we are interested in proving that All-Checkpoint and All-Checkpoint-Periodic are similar for Exponential distribution.
Theorem 5. When the execution time of a job follows a distribution $\mathcal{D} \sim$ Exponential $(\lambda)$, the solution of All-Checkpoint and All-Checkpoint-Periodic are identical.

Proof. For this specific result, we express the solutions under the form: $\left\{T_{1}, T_{2}, T_{3}, \cdots\right\}$.
We have the following properties:

- Given $\left(t_{1}, t_{2}, \ldots,\right)$ a solution to All-Checkpoint, then the associated $\left\{T_{1}, T_{2}, T_{3}, \cdots\right\}$ satisfy: forall $i, t_{i}=\sum_{j \leq i} T_{j}$. This is a direct corollary of Equations (3) and (4).
- For $\mathcal{D} \sim$ Exponential $(\lambda)$,

$$
\begin{equation*}
\mathbb{P}\left(X>V_{1}+V_{2}\right)=\mathbb{P}\left(X>V_{1}\right) \cdot \mathbb{P}\left(X>V_{2}\right) \tag{21}
\end{equation*}
$$

We define $\left(u_{1}, u_{2}, \cdots\right)$ the solution that minimizes

$$
\begin{equation*}
\sum_{i=1}^{\infty}\left(\alpha\left(R+u_{i}+C\right)+\beta(R+C)+\gamma\right) \cdot \mathbb{P}\left(X>\sum_{j<i} u_{j}\right) \tag{22}
\end{equation*}
$$

RR n ${ }^{\circ} 9294$

To study $\mathcal{S}^{\text {opt }}=\left(T_{1}^{o}, T_{2}^{o}, \ldots\right)$ the optimal solution given by All-Checkpoint for $\mathcal{D}$, we rewrite the expected cost of a solution $\mathcal{S}=\left(T_{1}, T_{2}, \ldots\right)$ for All-Checkpoint:

$$
\begin{aligned}
\mathbb{E}(\mathcal{S}) & =\beta \cdot \mathbb{E}[X]+\sum_{i=1}^{\infty}\left(\alpha W_{i}+\beta\left(R_{i}+\left(1-\delta_{i-1}\right) T_{i-1}+C_{i-1}\right)+\gamma\right) \cdot \mathbb{P}\left(X>t_{i-1}\right) \\
& =\beta \cdot \mathbb{E}[X]-C+\sum_{i=1}^{\infty}\left[\alpha\left(R+T_{i}+C\right)+\beta(R+C)+\gamma\right] \cdot \mathbb{P}\left(X>\sum_{j<i} T_{j}\right)
\end{aligned}
$$

From this formulation, we can see that $T_{1}^{o}=u_{1}, T_{2}^{o}=u_{2}, T_{3}^{o}=u_{3}, \ldots$ as they satisfy the same equations.

By rewriting and using Eq 21,

$$
\begin{align*}
\mathbb{E}(\mathcal{S})= & \beta \cdot \mathbb{E}[X]+\left(\alpha W_{1}+\beta R+\gamma\right)+\sum_{i=2}^{\infty}\left(\alpha W_{i}+\beta(R+C)+\gamma\right) \cdot \mathbb{P}\left(X>\sum_{j<i} T_{j}\right) \\
= & \beta \cdot \mathbb{E}[X]+\alpha\left(W_{1}+\beta C+\gamma\right) \\
& +\mathbb{P}\left(X>T_{1}\right) \cdot \sum_{i=2}^{\infty}\left(\alpha\left(R+T_{i}+C\right)+\beta(R+C)+\gamma\right) \cdot \mathbb{P}\left(X>\sum_{2 \leq j<i} T_{j}\right) \\
=\beta \cdot & \mathbb{E}[X]+\alpha\left(W_{1}+\beta C+\gamma\right) \\
& +\mathbb{P}\left(X>T_{1}\right) \cdot \sum_{i=1}^{\infty}\left(\alpha\left(R+T_{i-1}+C\right)+\beta(R+C)+\gamma\right) \cdot \mathbb{P}\left(X>\sum_{j<i-1} T_{j+1}\right) \tag{23}
\end{align*}
$$

From this last formulation, we can see that the sequence $\left(T_{2}^{o}, T_{3}^{o}, \ldots\right)$ that minimizes Eq. (23) can be optimized independently of $T_{1}^{o}$, and that it is the solution that minimizes Eq. 22). Hence we obtain, $T_{2}^{o}=u_{1}, T_{3}^{o}=u_{2}, \ldots$ Iterating the process, we obtain the result: $T_{1}^{o}=u_{1}=T_{2}^{o}=$ $u_{2}=T_{3}^{o}=\ldots$, and the solution to All-Checkpoint is a periodic solution.

### 3.5.2 All-Checkpoint-Periodic for Uniform distributions

In this section, we consider the All-Checkpoint approach for Uniform distributions in the ReservationOnly scenario. Specifically, we are able to characterize the best periodic approach: for a Uniform distribution $\mathcal{D}$ over $[a, b]$ (where $0 \leq a<b$ ), consider a reservation sequence $\mathcal{S}^{(n)}=\left(W_{i}, 1\right)_{1 \leq i \leq n}$ where $n \geq 2, W_{1}=a+\frac{b-a}{n}+C, W_{i}=R+\frac{b-a}{n}+C$ for $1<i<n$ and $W_{n}=R+\frac{b-a}{n}$. In other words, we have $n-1$ evenly distributed checkpoints in the interval $[a, b]$. The first reservation has a checkpoint but no restart, the last reservation (whenever used) has a restart but no checkpoint, and all intermediate reservations have both a restart and a checkpoint. Finally, let $\mathcal{S}^{(1)}=(b, 0)$ be the sequence with a unique reservation of length (and cost) $b$ (no need to checkpoint in this particular case). The following proposition provides the optimal value of $n$ :

Proposition 1. With the above notations, $\mathbb{E}\left(\mathcal{S}^{(n)}\right)$ is minimized either for $n=\max \left(1,\left\lfloor n^{\text {opt }}\right\rfloor\right)$ or $n=\left\lceil n^{\mathrm{opt}}\right\rceil$ where $n^{\mathrm{opt}}=\sqrt{\frac{b-a-2 C}{C+R}}$ if $b-a \geq 2 C$ and $n^{\mathrm{opt}}=1$ otherwise.

Proof. Because the distribution is uniform, the probability that $i$ reservations are need for a given job is always equal to $\frac{1}{n}$, hence the expected cost of $\mathcal{S}^{(n)}$ for $n \geq 2$ is

$$
\mathbb{E}\left(\mathcal{S}^{(n)}\right)=\frac{1}{n} \sum_{i=1}^{n} \sum_{j=1}^{i} W_{j}
$$

where $W_{i}$ is the cost of the $i$-th reservation. After several algebraic manipulations, we derive that

$$
\mathbb{E}\left(\mathcal{S}^{(n)}\right)=\frac{n-1}{2 n} a+\frac{n+1}{2 n} b+\frac{n^{2}+n-2}{2 n} C+\frac{n-1}{2} R
$$

Differentiating, the derivative gets zeroed for $n=n^{\text {opt }}$ when $b-a \geq 2 C$, and otherwise it stays positive, hence the result.

If $\mathcal{D} \sim \operatorname{Uniform}(a, b)$ with $[a, b]=[2,20]$ and $C=R=1$ we find $n^{\text {opt }}=\sqrt{8}$. One can compute that the cost for three reservations (ceil value for $\sqrt{8}$ ) is $\mathbb{E}\left(\mathcal{S}^{(3)}\right)=97 / 6 \approx 16.2$. Let us now define a two-reservation strategies at milestones 11 and 20 we can compute its associated cost by:

$$
\begin{aligned}
\mathbb{E}\left(\mathcal{S}^{(2)}\right) & =\int_{2}^{11}(\alpha 12+\beta t+\gamma) \mathbb{P}(X=t \mid X \leq 11) d t \\
& +\int_{11}^{20}[(\alpha 12+\beta 12+\gamma)+(\alpha 10+\beta(t-11)+\gamma)] \mathbb{P}(X=t \mid X \geq 11) d t
\end{aligned}
$$

This gives us $\mathbb{E}\left(\mathcal{S}^{(2)}\right)=17$ for distribution $\mathcal{D}$. Hence, it is more efficient to use three reservations than two.

## 4 Performance evaluation

In this section, we evaluate the performance of the different algorithms in simulation. In the following, performance stands for the expected cost of each algorithm under various job execution time distributions, $\mathrm{C} / \mathrm{R}$ overhead and cost functions. We use jobs that follow a wide range of usual probability distributions as well as a distribution obtained from traces of a real neuroscience application. The code for this section is publicly available on https: //gitlab.inria.fr/vhonore/ckpt-for-stochastic-scheduling.

### 4.1 Evaluation methodology

In this section, we evaluate five different algorithms from the following two sets of strategies:

- Dyn-Prog-Count: This set includes Algorithm 1 and its All-Checkpoint and No-Checkpoint variants described in Section 3.4,
- All-Checkpoint-Periodic: This set includes Algorithm 2 and its No-Checkpoint-Periodic counterpart where checkpointing is not allowed (i.e., $\delta_{i}=0, \forall i$ ).

The above algorithms are evaluated using two scenarios and two different cost functions.
The first cost function is ReservationOnly, presented in Section 2.4. This function is based on the Reserved Instance pricing scheme in AWS [2], where the user pays exactly what is requested. Hence, $\alpha=1, \beta=\gamma=0$. The second cost function, called HPC, add an additional

[^1]Table 1: Probability distributions and parameter instantiations.

| Distribution | PDF $f(t)$ | Instantiation | Support (in hours) |
| :---: | :---: | :---: | :---: |
| Distributions with infinite support |  |  |  |
| Exponential ( $\lambda$ ) | $\lambda e^{-\lambda t}$ | $\lambda=1.0 h^{-1}$ | $t \in[0, \infty)$ |
| $\operatorname{Weibull}(\lambda, \kappa)$ | $\frac{\kappa}{\lambda}\left(\frac{t}{\lambda}\right)^{\kappa-1} e^{-\left(\frac{t}{\lambda}\right)^{\kappa}}$ | $\begin{aligned} & \lambda=1.0 h \\ & \kappa=0.5 \end{aligned}$ | $t \in[0, \infty)$ |
| $\operatorname{Gamma}(\alpha, \beta)$ | $\frac{\beta^{\alpha}}{\Gamma(\alpha)} t^{\alpha-1} e^{-\beta t}$ | $\begin{aligned} & \alpha=2.0 \\ & \beta=2.0 h^{-1} \end{aligned}$ | $t \in[0, \infty)$ |
| Lognormal ( $\nu, \kappa$ ) | $\frac{1}{t \kappa \sqrt{2 \pi}} e^{-\frac{(\ln t-\nu)^{2}}{2 \kappa^{2}}}$ | $\begin{aligned} & \nu=3.0 h \\ & \kappa=0.5 \end{aligned}$ | $t \in(0, \infty)$ |
| Pareto $(\nu, \alpha)$ | $\frac{\alpha \nu^{\alpha}}{t^{\alpha+1}}$ | $\begin{aligned} & \nu=1.5 h \\ & \alpha=3.0 \end{aligned}$ | $t \in[\nu, \infty)$ |
| Distributions with finite support |  |  |  |
| Truncated $\operatorname{Normal}\left(\nu, \kappa^{2}, a, b\right)$ | $\frac{1}{\kappa} \sqrt{\frac{2}{\pi}} \cdot \frac{e^{-\frac{1}{2}\left(\frac{t-\nu}{\kappa}\right)^{2}}}{1-\operatorname{erf}\left(\frac{a-\nu}{\kappa \sqrt{2}}\right)}$ | $\begin{aligned} & \nu=8.0 h \\ & \kappa^{2}=2.0 h^{2} \\ & a=1.0 h \\ & b=20.0 h \end{aligned}$ | $t \in[a, b]$ |
| Uniform $(a, b)$ | $\frac{1}{b-a}$ | $\begin{aligned} & a=1.0 h \\ & b=20.0 h \end{aligned}$ | $t \in[a, b]$ |
| $\operatorname{Beta}(\alpha, \beta)$ | $\frac{t^{\alpha-1} \cdot(1-t)^{\beta-1}}{\mathrm{~B}(\alpha, \beta)}$ | $\begin{aligned} & \alpha=2.0 \\ & \beta=2.0 \end{aligned}$ | $t \in[0,1]$ |
| Bounded Pareto ( $L, H, \alpha$ ) | $\frac{\alpha L^{\alpha} t^{-\alpha-1}}{1-\left(\frac{L}{H}\right)^{\alpha}}$ | $\begin{aligned} & L=1.0 h \\ & H=20.0 h \\ & \alpha=2.1 \end{aligned}$ | $t \in[L, H]$ |

cost that is proportional to the actual execution time (pay for what you use). Thus, $\alpha=1, \beta=$ $1, \gamma=0$.

We now detail the two scenarios we use to evaluate the algorithms, using the two cost functions:

- Scenario 1 (Section 4.2): We consider nine usual probability distributions, five of which have infinite support (Exponential, Weibull, Gamma, Lognormal, Pareto) and four have finite support (Truncated Normal, Uniform, Beta, Bounded Pareto). Table 1 lists all distributions used in simulation with the instantiations of their parameters for evaluation. The first five distributions are truncated and fed as input to Algorithm 1. To do so, we set the upper bound of the infinite support to $b=Q(1-v)$, where $Q(x)=\inf \{t \mid F(t) \geq x\}$ is the quantile function and $v$ is a small constant. In our simulation, we set $v=10^{-7}$. During the discretization procedure in Algorithm 11, we then normalize the probabilities of all discrete values so that they sum to 1 . We set $C=R=360$ seconds ( 0.1 hour). This checkpointing cost is extracted from [18] and corresponds to an average checkpointing duration, where an optimistic one is 60 seconds and a pessimistic one is 600 seconds. We further discuss the impact of the checkpointing cost on the performance.
- Scenario 2 (Section 4.3): In this scenario, we consider the execution time traces of a real neuroscience application, and fit a Lognormal distribution to its execution times. To further evaluate the robustness of the algorithms, we perturb the parameters of the fitted distribution by varying its mean and standard deviation and show the impact on the performance.


### 4.2 Results for Scenario 1

We first evaluate the performance of Dyn-Prog-Count compared to the other strategies, when the values of $R$ and $C$ varies. Figure 5 presents the performance of these strategies normalized to that of Dyn-Prog-Count (black line for $y=1.0$ ) for all distributions of Table 1 using ReservationOnly cost function. We use $\varepsilon=0.1$ for Dyn-Prog-Count and its variants. Regarding periodic strategies, we choose the best value for the number of chunks $\tau$ in [1, 1000]. Not surprisingly, we can observe that when $C$ and $R$ are small, the best result is to use the All-Checkpoint strategy while when they are large, one should use the No-Checkpoint strategy. There exist thresholds on the sizes of $C$ and $R$ where Dyn-Prog-Count uses a mix of checkpointed and not checkpointed reservations. In that case, the gain of using Dyn-Prog-Count can be up to $10 \%$ in compared with its variants. An interesting future direction is to find properties on those threshold depending on the distribution. Finally, one should observe that the gain obtained with Dyn-Prog-Count compared to the best periodic solution is in general more important (for Truncated Normal, the performance of periodic solutions are worse than a factor 2 of Dyn-Prog-Count). For Exponential distribution, All-Checkpoint and its periodic counterpart are identical (proof can be found in Section 3.5.1), due to the memoryless property of the exponential distribution. Figure 6 shows the results for a similar setup with HPC cost function. We see that results are consistent between the two cost functions, with small variations in results but same general trends.

We then study the impact of $\varepsilon$ on the performance of Dyn-Prog-Count (DPC) for the two different cost functions, when $R=C=6 \mathrm{~min}, 30 \mathrm{~min}$ and 60 min . The idea is that when $\varepsilon=1$, this theoretically guarantees that the performance is at most twice $(=1+\varepsilon)$ that of the optimal, but in practice it can be a lot better. We study in Figure 7 the performance of Dyn-Prog-Count for various values of $\varepsilon$ for distributions of Table 1 with ReservationOnly cost function. All performance are normalized by Dyn-Prog-Count for $\varepsilon=0.1$. We can see that in practice, the convergence to the lower bound in performance is fast. Indeed, for $\varepsilon=1$ and $C=R=6 \mathrm{~min}$ (Figure 7a), almost all distributions already reach convergence, except for Weibull and Pareto (which have a much larger domain of definition and specific properties ${ }^{2}$. For those distributions, we see that they converged for $\varepsilon=0.1$. We observe similar trends in Figures 7 b and 7 c when $R$ and $C$ increases. Interestingly, one can note that Pareto distribution converges faster for $C=R=30$ or 60 min than for $C=R=6 \mathrm{~min}$, while Weibull distribution shows contrary behavior. Convergence is still achieved for $\varepsilon=0.1$. This shows the possible impact of application features on algorithm behavior.

Figure 8 shows the results of a similar setup for HPC cost function. One can note that all distributions converge even faster than for ReservationOnly cost function. Indeed, for $\varepsilon=0.4$, all distribution converge. Overall, the results are consistent in between the two cost functions. In both figures, the number of chunks $n$ in Dyn-Prog-Count varies between 50 to 1000 depending on the distribution and value of $\varepsilon$, showing the practicality of Dyn-Prog-Count for considered distributions.

Our final evaluation for this scenario is a study of the impact of the size of the period. Until now we have always chosen the period that minimized the objective functions. Table 2 shows the performance of both variants of the periodic algorithms, All-Checkpoint-Periodic and No-Checkpoint-Periodic, normalized by that of Dyn-Prog-Count $(\varepsilon=0.1)$, when $C=R=360$ s using ReservationOnly cost function. For each distribution: the second columns shows the best period found when $\tau$ varies from 1 to 1000 (with its associated cost normalized by that of Dyn-Prog-Count), and the other columns present results for specific

[^2]

Figure 5: Expected costs of the different strategies normalized to that of Dyn-Prog-Count $(X, 0.1)$ when $C=R$ vary from 60 to 3600 seconds, for all distributions in Table 1 with support considered in hours with ReservationOnly cost function. We indicate in brackets the mean $\mu$ of each distribution.


Figure 6: Expected costs of the different strategies normalized to that of Dyn-Prog-Count $(X, 0.1)$ when $C=R$ vary from 60 to 3600 seconds, for all distributions in Table 1 with support considered in hours with HPC cost function. We indicate in brackets the mean $\mu$ of each distribution.


Figure 7: Expected cost of Dyn-Prog-Count $(X, \varepsilon)$ as a function of $\varepsilon$ for different distributions for $X$ with ReservationOnly cost function. $C=R$ are set to 6,30 and 60 min .


Figure 8: Expected cost of Dyn-Prog-Count $(X, \varepsilon)$ as a function of $\varepsilon$ for different distributions for $X$ with HPC cost function. $C=R$ are set to 6,30 and 60 min .
values of $\tau$ in that interval. As was observed before, All-Checkpoint-Periodic is in general not able to match Dyn-Prog-Count (except for some distributions). We can also clearly see that No-Checkpoint-Periodic performs even worse than All-Checkpoint-Periodic. The reason is that the checkpointing cost is relatively low in this setup, so it is preferable to checkpoint more often than never. Hence, when $C=R=1 \mathrm{~h}$, Table 3 shows that No-Checkpoint-Periodic performs slightly better than All-Checkpoint-Periodic. Finally another observation is that a wrong period size can significantly deteriorate the performance of the periodic algorithms. Tables 4 and 5 show similar trends using the HPC cost function.

Table 2: Expected cost of All-Checkpoint-Periodic and No-Checkpoint-Periodic, normalized by Dyn-Prog-Count $(X, 0.1)$ for $C=R=360$ s, with ReservationOnly cost function.

| Distribution | All-Checkpoint-Periodic |  |  |  |  |  |  | No-Checkpoint-Periodic |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Best $\tau$ | $\tau=1$ | $\tau=200$ | $\tau=400$ | $\tau=600$ | $\tau=800$ | $\tau=1000$ | Best $\tau$ | $\tau=1$ | $\tau=200$ | $\tau=400$ | $\tau=600$ | $\tau=800$ | $\tau=1000$ |
| Exponential | 27 (1.00) | 9.82 | 2.31 | 4.02 | 5.75 | 7.47 | 9.19 | 14 (1.38) | 9.82 | 6.91 | 13.23 | 19.56 | 25.89 | 32.22 |
| Weibull | 380 (1.06) | 106.50 | 1.15 | 1.06 | 1.10 | 1.18 | 1.27 | 89 (2.54) | 106.50 | 3.26 | 5.32 | 7.52 | 9.74 | 11.98 |
| Gamma | 14 (1.02) | 6.02 | 3.68 | 6.76 | 9.85 | 12.93 | 16.02 | 8 (1.26) | 6.02 | 9.34 | 18.08 | 26.82 | 35.56 | 44.31 |
| Lognormal | 11 (1.11) | 3.60 | 3.92 | 7.05 | 10.18 | 13.32 | 16.45 | 4 (1.25) | 3.60 | 15.48 | 30.17 | 44.86 | 59.56 | 74.25 |
| Pareto | 1000 (1.01) | 228.77 | 1.75 | 1.23 | 1.09 | 1.03 | 1.01 | 562 (1.33) | 228.77 | 1.78 | 1.37 | 1.33 | 1.37 | 1.45 |
| TruncatedNormal | 2 (2.15) | 2.18 | 8.17 | 14.31 | 20.45 | 26.59 | 32.73 | 1 (2.18) | 2.18 | 197.54 | 394.01 | 590.47 | 786.93 | 983.39 |
| Uniform | 8 (1.01) | 1.57 | 3.17 | 5.51 | 7.86 | 10.20 | 12.54 | 1 (1.57) | 1.57 | 51.08 | 101.33 | 151.59 | 201.84 | 252.10 |
| Beta | 2 (1.06) | 1.11 | 30.78 | 61.00 | 91.23 | 121.45 | 151.67 | 1 (1.11) | 1.11 | 40.85 | 81.15 | 121.45 | 161.75 | 202.05 |
| BoundedPareto | 32 (1.01) | 7.53 | 1.73 | 2.71 | 3.70 | 4.70 | 5.69 | 14(1.44) | 7.53 | 6.51 | 12.28 | 18.06 | 23.83 | 29.61 |

Table 3: Expected cost of All-Checkpoint-Periodic and No-Checkpoint-Periodic, normalized by Dyn-Prog-Count $(X, 0.1)$ for $C=R=1 \mathrm{~h}$, with ReservationOnly cost function.

| Distribution | All-Checkpoint-Periodic |  |  |  |  |  |  | No-Checkpoint-Periodic |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Best $\tau$ | $\tau=1$ | $\tau=200$ | $\tau=400$ | $\tau=600$ | $\tau=800$ | $\tau=1000$ | Best $\tau$ | $\tau=1$ | $\tau=200$ | $\tau=400$ | $\tau=600$ | $\tau=800$ | $\tau=1000$ |
| Exponential | 12 (1.44) | 7.35 | 9.49 | 18.53 | 27.57 | 36.61 | 45.66 | 14 (1.03) | 7.35 | 5.17 | 9.90 | 14.64 | 19.38 | 24.12 |
| Weibull | 156 (1.26) | 70.94 | 1.29 | 1.65 | 2.11 | 2.60 | 3.11 | 89 (1.69) | 70.94 | 2.17 | 3.54 | 5.01 | 6.49 | 7.98 |
| Gamma | 7 (1.48) | 4.77 | 17.59 | 34.70 | 51.82 | 68.93 | 86.05 | 8 (1.00) | 4.77 | 7.40 | 14.32 | 21.25 | 28.18 | 35.11 |
| Lognormal | 4 (1.23) | 2.98 | 18.81 | 36.96 | 55.11 | 73.26 | 91.42 | 4 (1.04) | 2.98 | 12.82 | 24.98 | 37.14 | 49.30 | 61.47 |
| Pareto | 563 (1.20) | 175.44 | 1.58 | 1.24 | 1.21 | 1.24 | 1.31 | 562 (1.02) | 175.44 | 1.37 | 1.05 | 1.02 | 1.05 | 1.11 |
| TruncatedNormal | 1 (1.85) | 1.85 | 38.01 | 74.45 | 110.89 | 147.33 | 183.78 | 1 (1.85) | 1.85 | 167.50 | 334.08 | 500.67 | 667.25 | 833.83 |
| Uniform | 3 (1.01) | 1.23 | 13.46 | 26.26 | 39.07 | 51.88 | 64.69 | 1 (1.23) | 1.23 | 39.89 | 79.13 | 118.37 | 157.61 | 196.85 |
| Beta | 1 (1.08) | 1.08 | 207.32 | 414.09 | 620.87 | 827.64 | 1034.42 | 1 (1.08) | 1.08 | 39.93 | 79.31 | 118.70 | 158.08 | 197.47 |
| BoundedPareto | 14 (1.26) | 5.59 | 5.72 | 10.88 | 16.05 | 21.21 | 26.38 | 14 (1.07) | 5.59 | 4.82 | 9.11 | 13.39 | 17.68 | 21.96 |

Table 4: Performance of All-Checkpoint-Periodic and No-Checkpoint-Periodic, normalized by Dyn-Prog-Count $(X, 0.1)$ for $C=R=360$ s, with HPC cost function.

| Distribution | All-Checkpoint-Periodic |  |  |  |  |  |  | No-Checkpoint-Periodic |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Best $\tau$ | $\tau=1$ | $\tau=200$ | $\tau=400$ | $\tau=600$ | $\tau=800$ | $\tau=1000$ | Best $\tau$ | $\tau=1$ | $\tau=200$ | $\tau=400$ | $\tau=600$ | $\tau=800$ | $\tau=1000$ |
| Exponential | 21 (1.00) | 6.38 | 2.43 | 4.20 | 5.97 | 7.74 | 9.52 | 10 (1.35) | 6.38 | 8.14 | 15.59 | 23.04 | 30.50 | 37.95 |
| Weibull | 300 (1.04) | 62.32 | 1.07 | 1.05 | 1.12 | 1.21 | 1.31 | 64 (2.29) | 62.32 | 3.53 | 6.03 | 8.60 | 11.19 | 13.79 |
| Gamma | 12 (1.01) | 4.03 | 3.83 | 6.98 | 10.13 | 13.28 | 16.44 | 6 (1.24) | 4.03 | 10.80 | 20.89 | 30.99 | 41.08 | 51.18 |
| Lognormal | 4 (1.07) | 2.42 | 3.74 | 6.63 | 9.52 | 12.41 | 15.30 | 3 (1.17) | 2.42 | 16.55 | 32.24 | 47.93 | 63.62 | 79.32 |
| Pareto | 999 (1.00) | 128.29 | 1.39 | 1.11 | 1.03 | 1.01 | 1.00 | 415 (1.28) | 128.29 | 1.45 | 1.28 | 1.32 | 1.41 | 1.51 |
| TruncatedNormal | 2 (1.61) | 1.61 | 6.93 | 12.34 | 17.75 | 23.16 | 28.57 | 1 (1.61) | 1.61 | 207.16 | 413.80 | 620.45 | 827.09 | 1033.73 |
| Uniform | 6 (1.01) | 1.28 | 2.98 | 5.07 | 7.17 | 9.26 | 11.36 | 1 (1.28) | 1.28 | 54.39 | 107.91 | 161.42 | 214.93 | 268.44 |
| Beta | 1 (1.03) | 1.03 | 32.95 | 65.29 | 97.62 | 129.96 | 162.29 | 1 (1.03) | 1.03 | 45.95 | 91.22 | 136.49 | 181.76 | 227.03 |
| BoundedPareto | 26 (1.01) | 4.69 | 1.74 | 2.68 | 3.62 | 4.57 | 5.52 | 10 (1.38) | 4.69 | 7.22 | 13.62 | 20.03 | 26.44 | 32.85 |

### 4.3 Results for Scenario 2

We now present the simulation results for a probability distribution fitted to the execution time traces of a real neuroscience application (a code for structural identification of orbital anatomy) extracted from the Vanderbilt's medical imaging database [15]. Figure 1 shows the execution time traces of the application and its fitted Lognormal distribution. Figure 9 presents the performance of different algorithms for this fitted distribution using the ReservationOnly cost function, for

Table 5: Performance of All-Checkpoint-Periodic and No-Checkpoint-Periodic, normalized by Dyn-Prog-Count $(X, 0.1)$ for $C=R=1$ h, with HPC cost function.

| Distribution | All-Checkpoint-Periodic |  |  |  |  |  |  | No-Checkpoint-Periodic |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Best $\tau$ | $\tau=1$ | $\tau=200$ | $\tau=400$ | $\tau=600$ | $\tau=800$ | $\tau=1000$ | Best $\tau$ | $\tau=1$ | $\tau=200$ | $\tau=400$ | $\tau=600$ | $\tau=800$ | $\tau=1000$ |
| Exponential | 9 (1.25) | 4.81 | 11.24 | 22.19 | 33.15 | 44.11 | 55.07 | 10 (1.02) | 4.81 | 6.13 | 11.74 | 17.35 | 22.96 | 28.58 |
| Weibull | 116 (1.19) | 44.78 | 1.30 | 1.83 | 2.42 | 3.03 | 3.65 | 64 (1.65) | 44.78 | 2.54 | 4.33 | 6.18 | 8.04 | 9.91 |
| Gamma | 6 (1.29) | 3.26 | 21.21 | 42.11 | 63.02 | 83.93 | 104.84 | 6 (1.00) | 3.26 | 8.73 | 16.89 | 25.04 | 33.20 | 41.36 |
| Lognormal | 4 (1.13) | 2.12 | 21.39 | 42.15 | 62.91 | 83.67 | 104.43 | 3 (1.03) | 2.12 | 14.50 | 28.24 | 41.99 | 55.73 | 69.48 |
| Pareto | 438 (1.11) | 103.29 | 1.27 | 1.11 | 1.13 | 1.21 | 1.31 | 415 (1.03) | 103.29 | 1.17 | 1.03 | 1.06 | 1.13 | 1.21 |
| TruncatedNormal | 1 (1.45) | 1.45 | 41.13 | 81.05 | 120.98 | 160.91 | 200.83 | 1 (1.45) | 1.45 | 186.46 | 372.46 | 558.46 | 744.46 | 930.46 |
| Uniform | 2 (1.00) | 1.08 | 15.21 | 29.71 | 44.21 | 58.71 | 73.22 | 1 (1.08) | 1.08 | 45.90 | 91.05 | 136.21 | 181.36 | 226.52 |
| Beta | 1 (1.03) | 1.03 | 263.43 | 526.83 | 790.22 | 1053.62 | 1317.02 | 1 (1.03) | 1.03 | 45.65 | 90.62 | 135.59 | 180.56 | 225.53 |
| BoundedPareto | 11 (1.14) | 3.64 | 6.50 | 12.51 | 18.53 | 24.55 | 30.57 | 10 (1.07) | 3.64 | 5.59 | 10.56 | 15.53 | 20.50 | 25.46 |

different values of $C=R$. To evaluate the robustness of algorithms, we also vary the original mean $\mu_{o}$ (Figures 9a 9c 9e or standard deviation $\sigma_{o}$ (Figures 9b 9d 9f) of the distribution from their original values. For readability, all axis are in logscale. We fix $\varepsilon=1.0$ and test checkpoint/restart costs equals to $6 \mathrm{~min}, 1 \mathrm{~h}$ and 12 h . For periodic strategies, we use similar brute-force procedure as Scenario 1 to find the period that performs best. The expected costs of the algorithms are normalized by that of an omniscient scheduler (blue dashed line), which knows the execution time $t$ of a job a priori, and thus would pay the minimum possible cost by making a single reservation of length $t_{1}=t$. Averaging over all possible values of $t$ from the distribution, the omniscient scheduler has an expected cost

$$
\mathbb{E}^{o}=\int_{0}^{\infty}(\alpha t+\beta t+\gamma) f(t) d t=\alpha \cdot \mathbb{E}[X]+\gamma
$$

We can observe that Dyn-Prog-Count always gives the best performance. As previously observed, the checkpointing cost influences the performance of No-CHECKPoint and All-ChECKpoint with regard to Dyn-Prog-Count. When $C=R=600$ seconds or $C=R=3600$ s (Figure $9 \mathrm{a}-9 \mathrm{~d}$ ), the value is low enough to allow for checkpointing all reservations, the performance of Dyn-Prog-Count and All-Checkpoint are the same and outperforms No-Checkpoint by a wide margin. However, when the checkpoint/restart overhead increases to 12h (roughly $\frac{\mu_{o}}{2}$ ), we see that checkpointing all reservations become overcostly (No-CHECKPOINT is better than All-Checkpoint). In that case, Dyn-Prog-Count outperforms all other strategies. One can observe that when the ratio $\mu / \sigma$ is large (either by increasing the mean, or decreasing the standard deviation), the solutions converge to the omniscient scheduler. This could be expected, in this case the variability becomes negligible and the job behaves similarly to a deterministic job. As for the periodic algorithms, All-Checkpoint-Periodic has better performance than No-Checkpoint-Periodic. However, both algorithms have worse performance than Dyn-Prog-Count. The results demonstrate the robustness of Dyn-Prog-Count for a practical application with different distribution parameters.

Figure 10 shows similar trends using the same setup with the HPC cost function.

## 5 Experiments

In this section, we conduct real experiments on an HPC platform by using three stochastic neuroscience applications. The focus is to study the performance of different reservation and checkpointing strategies for scheduling multiple jobs in a shared HPC execution environment.


Figure 9: Normalized performance of algorithms with omniscient scheduler when $\mu$ or $\sigma$ vary, using ReservationOnly cost function ( $\alpha=1.0, \beta=\gamma=0$ ). Basis is the Lognormal distribution in Figure 1 ( $\left.\mu_{o}=21.4 \mathrm{~h}, \sigma_{0}=19.7 \mathrm{~h}\right) . C=R$ are set to 600,3600 and $43200 \mathrm{~s}(12 \mathrm{~h}), \varepsilon=1$.


Figure 10: Normalized performance of algorithms with omniscient scheduler when $\mu$ or $\sigma$ vary, using HPC cost function $(\alpha=\beta=1.0, \gamma=0)$. Basis is the Lognormal distribution in Figure 1 ( $\mu_{o}=21.4 \mathrm{~h}, \sigma_{0}=19.7 \mathrm{~h}$ ). $C=R$ are set to 600,3600 and $43200 \mathrm{~s}(12 \mathrm{~h}), \varepsilon=1$.

Table 6: Characteristics of the chosen neuroscience applications.

| Application Type | Walltime distribution | $C$ | $R$ |
| :---: | :---: | :---: | :---: |
| Diffusion model fitting (Qball) | Gamma $(k=1.18, \theta=34$, <br> $[a, b]=[146 s, 407 s])$ | 90 s | 40 s |
| Diffusion model fitting (SD) | Weibull $(k=1043811, \lambda=1174322466$, <br> $[a, b]=[46 \mathrm{~min}, 2.3 h])$ | 25 min | 10 min |
| Functional connectivity analysis (FCA) | Gamma $(k=3.6, \theta=72$, <br> $[a, b]=[165 s, 1003 s])$ | 150 s | 100 s |

### 5.1 Experimental setup

The chosen neuroscience applications are described in Table 6 along with their execution characteristics, which are extracted from the Vanderbilt's medical imaging database [15. In particular, the walltime distributions are obtained by fitting the execution time traces, and the checkpointing/restart costs are obtained by analyzing and averaging their memory footprints. Note that, for these applications, the restart costs $(R)$ are different from the checkpointing costs $(C)$.

Here, we focus on the evaluation of the following two different sets of strategies:

- An HPC-for-neuroscience strategy (called HPC in Section 5.2), which uses the average of the last 5 runs as the initial reservation length and then increases it by a factor of 1.5 for each subsequent reservation. This strategy is currently used by the MASI group 22 at Vanderbilt to handle stochastic neuroscience applications.
- Our proposed Dyn-Prog-Count strategy and its All-Checkpoint variant.

We ran the experiments on a 256 -thread Intel Processor (Xeon Phi $7230,1.30 \mathrm{GHz}$ ) while submitting jobs through the Slurm scheduler [34]. All three neuroscience applications are sequential (i.e., uses a single hardware thread) and performs some medical imaging analysis. The variation in the execution time is due to the different characteristics of the input data. However, as we do not have access to the raw input images, we used the information in the logs to simulate the characteristics of the input data, thereby forcing a job to run for a certain walltime and saving a specific amount of data for the checkpoints. In each experiment, we submitted 500 jobs from one of the three applications, and recorded the completion time of each job. We use the average job stretch (defined as the ratio between the total execution time of a job and its actual walltime) to show the individual job performance, and use the utilization (defined as the ratio between the sum of all jobs' walltimes and the total time required to execute them) to show the performance of the system for the whole job set. During the experiments, the scheduler has complete access to the entire platform, thus corresponding to the scenario with $\alpha=1, \beta=0, \gamma=0$.

### 5.2 Experimental results

By experimenting on a real system, we investigate the robustness of our strategy: 1) when multiple applications are running concurrently; 2) when the read/write times vary due to congestion while accessing I/O and/or due to application interference; 3) when the $\mathrm{C} / \mathrm{R}$ costs vary depending on when in the application the checkpoint/restart takes place (i.e., different values for different reservations). Figure 11 shows the performance of the three strategies when submitting 500 jobs from each application to the Slurm scheduler. We manually force the $\mathrm{C} / \mathrm{R}$ costs to be the same (as in Table 6) for each strategy so as to study the effects of application interference and the runtime system's performance variability on our model. The findings are consistent with
the simulation results (in Section 44, showing that Dyn-Prog-Count performs better than its All-Checkpoint variant in terms of both system utilization and average job stretch using all three applications. Moreover, the two algorithms outperform the simple HPC strategy.


Figure 11: Utilization and average job stretch for Dyn-Prog-Count, All-Checkpoint and the HPC strategies.

Depending on when the checkpoint is being taken, the checkpoint size and thus the time to save and restore the application can vary. Figure 12 shows the results when the $\mathrm{C} / \mathrm{R}$ costs could vary for different reservations. Based on the log traces of these three applications, we noticed that their memory footprints can vary by as much as $30 \%$ depending on when the checkpoint is taken (e.g., the checkpoint time can vary between 80 and 110 seconds for Qball). Our experiment generates random checkpoint sizes using a uniform distribution with the mean given by the average checkpoint size from the traces, and forces the application to read/write the corresponding amount at the beginning/end of the execution. In this experiment, we assume that the checkpointing time is included in the request time and is never responsible for applications exceeding their allocated time. While the Dyn-Prog-Count solution is computed using the average $\mathrm{C} / \mathrm{R}$ costs presented in Table 6 the experimental results show that its performance is robust up to $15-20 \%$ variability in the $\mathrm{C} / \mathrm{R}$ costs. Moreover, the average job stretch appears to be even more stable than the utilization, suggesting that most of the submitted jobs are not impacted by the fluctuation in the $\mathrm{C} / \mathrm{R}$ costs.

If application-level checkpoint is used, the application is usually aware of the checkpoint size, thus the checkpointing process can start before the reservation is over. The subsequent submissions can easily adapt to this deviation with the first checkpoints that are smaller than the one used to compute the sequence (this is the case for Figure 9). For system-level checkpoint, the application footprint usually remains similar throughout the execution of the application. In case the checkpointing time is causing the application to exceed the reserved time, the submission will fail and subsequent submissions can take this into account by adding the wasted time.

The limitation of our method is visible for applications with large variability in checkpointing size, which can be due to multiple factors, either within the application that presents different memory footprints throughout its execution, or by system-level causes, such as I/O congestion or failures. Such large variability in checkpointing size compared to what is used to compute the reservation sequence can result in worse performance when using our method, and the classic HPC model would be preferred in this case. We are currently investigating methods to incorporate variation of checkpointing size into the computation of the optimal reservation sequence, by either using historic information or adapting the subsequent request times based on the sizes of previous checkpoints. We plan to further analyze variable $C / R$ times in the future.


Figure 12: Utilization and average job stretch for the three applications (blue: Qball; Orange: SD; Green: FCA) when varying the C/R costs by different percentages ( 0 to $30 \%$ ) using the Dyn-Prog-Count strategy. Horizontal lines represent the results for the HPC strategy.

Table 7: Utilization and average job stretch for 10 different runs, each using 500 jobs from all three applications. The runs are ordered by the best improvement of Dyn-Prog-Count in utilization.

| DYN-PrOG-COUNT |  | HPC |  | Improvement |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| Utilization | Avg Stretch | Utilization | Avg Stretch | Utilization | Avg Stretch |
| 67 | 2.04 | 55 | 2.34 | $21 \%$ | $15 \%$ |
| 73 | 1.72 | 62 | 2.04 | $18 \%$ | $19 \%$ |
| 62 | 2.08 | 55 | 2.46 | $12 \%$ | $18 \%$ |
| 71 | 1.88 | 64 | 2.1 | $11 \%$ | $12 \%$ |
| 63 | 2.19 | 56 | 2.41 | $11 \%$ | $10 \%$ |
| 71 | 1.74 | 64 | 1.96 | $10 \%$ | $12 \%$ |
| 75 | 1.51 | 68 | 1.69 | $10 \%$ | $12 \%$ |
| 68 | 2.09 | 65 | 2.19 | $4 \%$ | $5 \%$ |
| 61 | 2.24 | 60 | 2.32 | $2 \%$ | $4 \%$ |
| 77 | 1.96 | 75 | 1.99 | $2 \%$ | $2 \%$ |

Finally, we conduct experiments in a more realistic scenario by running all three applications at the same time and investigating the impact on the strategies. In particular, we submitted a total of 500 jobs ( 100 from Qball, and 200 each from SD and FCA), and kept the C/R costs constant across different reservations. We recorded the utilization and average job stretch when using Dyn-Prog-Count compared to the HPC strategy for 10 different runs choosing different instances from the traces each time. The results are presented in Table 7 . We can see that Dyn-Prog-Count improves both utilization and average job stretch by $10 \%$ on average, and by up to $20 \%$ depending on the instances submitted. Overall, these results again illustrate the robustness of our algorithm and confirm its benefit for scheduling stochastic applications on reservation-based platforms, as long as checkpoint costs remain constant for each application

## 6 Related Work

We review some related work on reservation-based scheduling and checkpointing in HPC and cloud systems, as well as some prior work on dealing with stochastic applications.

### 6.1 Reservation-based scheduling

Batch schedulers are widely adopted by many resource managers in HPC systems, such as Slurm [34], Torque [29] and Moab [5]. Most batch schedulers use resource reservation in combination with backfilling [23, 25, 28], and rely on users to provide accurate estimates for the walltimes of the submitted jobs. While this works for applications with deterministic resource needs, it can cause resource over-estimation or under-estimation for stochastic jobs with large variations in the walltime, thus degrading system and/or application performance 12 . 31 .

Clusters of commodity servers that use big-data frameworks such as MapReduce [8] and Dryad [19] offer alternative solutions to running HPC workloads. Schedulers for these frameworks such as YARN [30] and Mesos [16] offer distinct features (e.g., fairness, resource negotiation) to manage the workloads, but they generally also require accurate information regarding the applications' resource demands.

Cloud computing platforms such as Amazon AWS [2] and Google GCP [14] have emerged as another option for executing HPC applications, with a variety of pricing and reservation schemes. Both on-demand and reservation models are available with the latter typically offering a lower price. Several works [1.6 6,33 have studied the pricing strategies for platform providers, as well as delay modeling and cost evaluation for the users.

### 6.2 Stochastic scheduling and checkpointing

Many prior works have considered stochastic scheduling for jobs with execution time uncertainty. Most research in this paradigm (e.g., [4, 13 21, 24 26, 27]) assumes that the execution time of a job follows a known probability distribution and aims at optimizing the expected response time or makespan for a set of jobs under various distributions. Most of them, however, do not consider the problem in the context of reservation-based scheduling. In our prior work [3], we have proposed near-optimal reservation strategies for a single job in both HPC and cloud systems. The work was later extended to scheduling a set of stochastic jobs, both sequential and parallel, using backfilling in a reservation-based scheduling environment 11, 12].

Another approach to coping with stochastic applications and/or platform unavailability is through checkpoint-restart [17, 32]. To ensure the robustness of the execution, the application's state is periodically checkpointed and in case of interrupt (due to either insufficient reservation or platform failure), the application can be recovered from the last saved checkpoint. In the context of fault tolerance, a lot of work (e.g., $[7,17,35]$ ) has been devoted to deriving the optimal checkpointing interval that minimizes the checkpointing overhead or resource waste.

In this paper, we present strategies that combine reservation and checkpointing for stochastic jobs with known execution time distributions. To the best of our knowledge, this is the first result to provide performance guarantee on the expected execution time while leveraging checkpointing in reservation-based scheduling environment.

## 7 Conclusion and Future Work

In this paper, we have studied the problem of performing checkpointing for stochastic applications running on a reservation-based platform. We presented an optimization framework through a scheduling problem that combines a sequence of reservations and associated checkpointing decisions, for a considered job. We provided an optimal solution via a dynamic programming algorithm in the case of discrete distributions, and showed the existence of a fully polynomialtime approximation scheme for bounded continuous distributions. Through an extensive set of simulations and experiments, we have demonstrated the effectiveness of those solutions in
comparison with classic strategies using both usual distributions and traces from real neuroscience applications.

For future work, we are interested in quantifying the critical checkpointing cost, below (or above) which the best strategy is to always (or never) checkpoint the reservations. This will help completely characterize the optimal solution for a given application. Another interesting direction is to incorporate non-constant checkpointing costs into the optimization problem, in order to design new reservation strategies that will be more robust than our current solutions. This would alleviate the limitation of our approach when checkpointing costs exhibit a large variability.

Through this work and our previous efforts [3, 11, 12] on scheduling stochastic applications, we would like to motivate users and systems administrators of HPC platforms to adopt our strategies in the case of stochastic job submissions. In particular, by modeling the execution time of a stochastic job with a distribution, we have demonstrated the benefits of requesting a sequence of reservations compared to a single reservation of the maximum execution time for the job. Thus, by leveraging such knowledge on the job profiles, we believe that our strategies will lead to significant improvements in terms of both system and application performance over the existing scheduling policies used in HPC platforms.

Acknowledgments Some of the simulations presented in this paper were carried out using the PlaFRIM experimental testbed, supported by Inria, CNRS (LABRI and IMB), Université de Bordeaux, Bordeaux INP and Conseil Régional d'Aquitaine (see https://www.plafrim.fr/ en/home/). The remaining simulation resources were provided by the computing facilities MCIA (Mésocentre de Calcul Intensif Aquitain) of the Université de Bordeaux and of the Université de Pau et des Pays de l'Adour.

## References

[1] M. Afanasyev and H. Mendelson. Service provider competition: Delay cost structure, segmentation, and cost advantage. Manufacturing $\mathcal{B}$ Service Operations Management, 12(2):213-235, 2010.
[2] Amazon. AWS pricing information. https://aws.amazon.com/ec2/pricing/. Accessed: 2018-10-11.
[3] G. Aupy, A. Gainaru, V. Honoré, P. Raghavan, Y. Robert, and H. Sun. Reservation Strategies for Stochastic Jobs. In IPDPS, 2019.
[4] J. Bruno, P. Downey, and G. N. Frederickson. Sequencing tasks with exponential service times to minimize the expected flow time or makespan. Journal of the ACM, 28(1):100-113, 1981.
[5] N. Capit, G. D. Costa, Y. Georgiou, G. Huard, C. Martin, G. Mounie, P. Neyron, and O. Richard. A batch scheduler with high level components. In CCGrid, pages 776-783, 2005.
[6] S. Chen, H. Lee, and K. Moinzadeh. Pricing schemes in cloud computing: Utilization-based versus reservation-based. Production and Operations Management, 2017.
[7] J. T. Daly. A higher order estimate of the optimum checkpoint interval for restart dumps. Future Generation Comp. Syst., 22(3):303-312, 2006.
[8] J. Dean and S. Ghemawat. MapReduce: Simplified data processing on large clusters. Commun. ACM, 51(1):107-113, Jan. 2008.
[9] L. Dierks and S. Seuken. Cloud pricing: the spot market strikes back. In The Workshop on Economics of Cloud Computing, 2016.
[10] A. Gainaru and G. Pallez. Making speculative scheduling robust to incomplete data. In Scala, 2019.
[11] A. Gainaru, G. Pallez, H. Sun, and P. Raghavan. Speculative scheduling for stochastic HPC applications. In ICPP, 2019.
[12] A. Gainaru, H. Sun, G. Aupy, Y. Huo, B. A. Landman, and P. Raghavan. On-the-fly scheduling versus reservation-based scheduling for unpredictable workflows. Int. J. High Perf. Computing Applications, 2019.
[13] A. Goel and P. Indyk. Stochastic load balancing and related problems. In FOCS, pages 579-586. ACM, 1999.
[14] Google. GCP pricing information. https://cloud.google.com/pricing/. Accessed: 2018-10-16.
[15] R. L. Harrigan, B. C. Yvernault, B. D. Boyd, S. M. Damon, K. D. Gibney, B. N. Conrad, N. S. Phillips, B. P. Rogers, Y. Gao, and B. A. Landman. Vanderbilt university institute of imaging science center for computational imaging XNAT: A multimodal data archive and processing environment. NeuroImage, 124:1097-1101, 2016.
[16] B. Hindman, A. Konwinski, M. Zaharia, A. Ghodsi, A. D. Joseph, R. Katz, S. Shenker, and I. Stoica. Mesos: A platform for fine-grained resource sharing in the data center. In $8 t h$ USENIX Conf. Networked Systems Design and Implementation, pages 295-308, 2011.
[17] T. Hérault and Y. Robert, editors. Fault-Tolerance Techniques for High-Performance Computing. Springer Verlag, 2015.
[18] Z. Hussain, T. Znati, and R. Melhem. Partial redundancy in hpc systems with non-uniform node reliabilities. In SC. IEEE Press, 2018.
[19] M. Isard, M. Budiu, Y. Yu, A. Birrell, and D. Fetterly. Dryad: Distributed data-parallel programs from sequential building blocks. In 2nd ACM SIGOPS/EuroSys European Conf. Computer Systems, 2007.
[20] L. Ismail and L. Khan. Implementation and performance evaluation of a scheduling algorithm for divisible load parallel applications in a cloud computing environment. Software: Practice and Experience, 45, 032014.
[21] J. Kleinberg, Y. Rabani, and E. Tardos. Allocating bandwidth for bursty connections. In STOC, pages 664-673, 1997.
[22] B. Landman. Medical-image Analysis and Statistical Interpretation (MASI) Lab. https: //my.vanderbilt.edu/masi/.
[23] D. A. Lifka. The ANL/IBM SP Scheduling System. In JSSPP, pages 295-303, 1995.
[24] R. H. Möhring, A. S. Schulz, and M. Uetz. Approximation in stochastic scheduling: The power of LP-based priority policies. Journal of the ACM, 46(6):924-942, 1999.
[25] A. W. Mu'alem and D. G. Feitelson. Utilization, Predictability, Workloads, and User Runtime Estimates in Scheduling the IBM SP2 with Backfilling. IEEE Trans. Parallel Distrib. Syst., 12(6):529-543, 2001.
[26] J. Niño Mora. Stochastic scheduling. Encyclopedia of Optimization, pages 3818-3824, 2009.
[27] M. L. Pinedo. Scheduling: Theory, Algorithms, and Systems. Springer-Verlag New York, Inc., Third edition, 2008.
[28] J. Skovira, W. Chan, H. Zhou, and D. A. Lifka. The EASY - LoadLeveler API Project. In $J S S P P$, pages 41-47, 1996.
[29] G. Staples. Torque resource manager. In Proceedings of the 2006 ACM/IEEE Conference on Supercomputing, SC '06, 2006.
[30] V. K. Vavilapalli, A. C. Murthy, C. Douglas, S. Agarwal, M. Konar, R. Evans, T. Graves, J. Lowe, H. Shah, S. Seth, B. Saha, C. Curino, O. O’Malley, S. Radia, B. Reed, and E. Baldeschwieler. Apache hadoop yarn: Yet another resource negotiator. In the 4 th Annual Symposium on Cloud Computing, pages 5:1-5:16, 2013.
[31] O. Weidner, M. Atkinson, A. Barker, and R. Filgueira Vicente. Rethinking high performance computing platforms: Challenges, opportunities and recommendations. In Proceedings of the ACM International Workshop on Data-Intensive Distributed Computing, pages 19-26, 2016.
[32] K. Wolter, editor. Stochastic Models for Fault Tolerance, Restart, Rejuvenation, and Checkpointing. Springer Verlag, 2010.
[33] H. Xu and B. Li. Dynamic cloud pricing for revenue maximization. IEEE Transactions on Cloud Computing, 1(2):158-171, July 2013.
[34] A. B. Yoo, M. A. Jette, and M. Grondona. Slurm: Simple linux utility for resource management. In $J S S P P$, pages 44-60, 2003.
[35] J. W. Young. A first order approximation to the optimum checkpoint interval. Comm. ACM, 17(9):530-531, 1974.

RESEARCH CENTRE
BORDEAUX - SUD-OUEST
200 avenue de la Vieille Tour 33405 Talence Cedex

Inria
Domaine de Voluceau - Rocquencourt
BP 105-78153 Le Chesnay Cedex
inria.fr


[^0]:    * Department of EECS, Vanderbilt University, Nashville, TN, USA
    $\dagger$ Inria, LaBRI, Univ. Bordeaux
    $\ddagger$ Laboratoire LIP, ENS Lyon \& University of Tennessee Knoxville, Lyon, France

[^1]:    ${ }^{1}$ The first reservation will be of length $11+1$ with the checkpointing overhead, and the second one of length $1+9$ with the restart cost.

[^2]:    ${ }^{2}$ For instance, Pareto is a long-tail distribution, meaning that it has a large number of occurrences that are far from the beginning and central part of its support. Formally, it means that $\frac{1-F(x+y)}{1-F(x)} \rightarrow 1$ when $x \rightarrow \infty, \forall y>0$.

