

Chronos: A Unifying Optimization Framework for Speculative Execution of Deadline-critical MapReduce Jobs

Maotong Xu, Sultan Alamro, Tian Lan and Suresh Subramaniam
Department of ECE, the George Washington University
{htfy8927, alamro, tlan, suresh}@gwu.edu

Abstract—Meeting desired application deadlines in cloud processing systems such as MapReduce is crucial as the nature of cloud applications is becoming increasingly mission-critical and deadline-sensitive. It has been shown that the execution times of MapReduce jobs are often adversely impacted by a few slow tasks, known as stragglers, which result in high latency and deadline violations. While a number of strategies have been developed in existing work to mitigate stragglers by launching speculative or clone task attempts, none of them provides a quantitative framework that optimizes the speculative execution for offering guaranteed Service Level Agreements (SLAs) to meet application deadlines. In this paper, we bring several speculative scheduling strategies together under a unifying optimization framework, called *Chronos*, which defines a new metric, Probability of Completion before Deadlines (PoCD), to measure the probability that MapReduce jobs meet their desired deadlines. We systematically analyze PoCD for popular strategies including Clone, Speculative-Restart, and Speculative-Resume, and quantify their PoCD in closed-form. The result illuminates an important tradeoff between PoCD and the cost of speculative execution, measured by the total (virtual) machine time required under different strategies. We propose an optimization problem to jointly optimize PoCD and execution cost in different strategies, and develop an algorithmic solution that is guaranteed to be optimal. Chronos is prototyped on Hadoop MapReduce and evaluated against three baseline strategies using both experiments and trace-driven simulations, achieving 50% net utility increase with up to 80% PoCD and 88% cost improvements.

I. INTRODUCTION

Distributed cloud computing frameworks, such as MapReduce, have been widely employed by social networks, financial operations and big data analytics due to their ability to process massive amounts of data by splitting large jobs into smaller, parallel tasks. Such frameworks are known to be susceptible to heavy tails in response time. The challenge arises from the fact that execution times of MapReduce jobs are often adversely impacted by a few slow tasks, causing the jobs to possibly suffer high latency and miss their application deadlines.

Prior work has reported that these slow tasks could run up to 8 times slower than the median task [1–4]. These slow tasks, known as *stragglers*, could significantly impact the overall performance of deadline-sensitive cloud applications and result in the violation of Service Level Agreements (SLAs).

Stragglers are inevitable in cloud environments due to a number of reasons. First, data center computing nodes are typically composed of commodity components, which could be

heterogeneous in nature and thus cause various nodes to perform differently. Second, various sources of hardware/software errors exist in large-scale data centers and could lead to node failures interrupting task execution on the failed nodes. Finally, with virtualization and resource sharing, co-scheduled tasks running simultaneously on the same physical machines could create an environment with resource contention and network congestion, which are shown to contribute significantly towards the occurrence of stragglers [2–4]. It has been observed that links in data centers could remain under congestion for up to several hundreds of seconds [5].

Hadoop has a speculation mode available (which we call Hadoop-S) to mitigate the straggler effect. In this mode, Hadoop can launch extra attempts for map (reduce) tasks after at least one map (reduce) task from the same job has finished. Periodically, Hadoop calculates the difference between the estimated completion time of each running task and the average completion time of finished tasks. It launches one extra attempt for the task that has the largest difference. Thus, if task sizes have large variation, Hadoop-S will launch extra attempts for a large number of tasks, and waste cluster resources. On the other hand, if task sizes are uniform, extra attempts are rarely launched; and so stragglers continue to straggle.

Several better approaches, both proactive and reactive, have been proposed recently to deal with the straggler problem [2–4, 6–10]. In particular, reactive approaches typically detect stragglers by monitoring the progress of tasks and then launch speculative copies of slow-running tasks. For example, Dryad [11] employs a heuristic scheme to detect tasks that are running slower than others, and schedules duplicate attempts. LATE [2] proposes a scheduling algorithm to launch speculative copies based on progress score. The progress score is the fraction of data processed. Later, Mantri [3] presents techniques to detect stragglers and act based on their causes. Dolly [4] is a proactive cloning approach. It launches replica task clones along with the originals and before straggler occurs, to avoid waiting and speculation altogether. Wrangler [8] applies a statistical learning technique and predicts stragglers before they occur and aborts their launch. Of the existing approaches, Mantri achieves a large amount of reduction in job completion time and resource usage compared with LATE [2]

and Dryad [11]. In Mantri, if there is an available container and there is no task waiting for a container, it keeps launching new attempts for a task whose remaining execution time is 30 sec larger than the average task execution time, until the number of extra attempts of the task is larger than 3. Mantri also periodically checks the progress of each task, and leaves one attempt with the best progress running. However, all these existing works only focus on mitigating stragglers without considering deadlines, and therefore cannot provide any guarantee to meet individual application deadlines which can vary significantly in practice.

Meeting desired deadlines is crucial as the nature of cloud applications is becoming increasingly mission-critical and deadline-sensitive [12, 13]. In this paper, we bring various scheduling strategies together under a unifying optimization framework, called *Chronos*, which is able to provide probabilistic guarantees for deadline-critical MapReduce jobs. In particular, we define a new metric, Probability of Completion before Deadlines (PoCD), to measure the probability that MapReduce jobs meet their desired deadlines. Assuming that a single task execution time follows a Pareto distribution [6, 7, 14, 15], we analyze three popular classes of (proactive and reactive) strategies, including Clone, Speculative-Restart, and Speculative-Resume, and quantify their PoCD in closed-form under the same framework. We note that our analysis also applies to MapReduce jobs, whose PoCD for map and reduce stages can be optimized separately. The result allows us to analytically compare the achievable PoCD of various existing strategies with different context and system configurations. In particular, we show that for the same number of speculative/clone attempts, Clone and Speculative-Resume always outperform Speculative-Restart, and we derive the sufficient conditions under which these strategies attain their highest PoCD. The result also illuminates an important tradeoff between PoCD and the cost of speculative execution, which is measured by the total (virtual) machine time that is required to cover the speculative execution in these strategies. The optimal tradeoff frontier that is characterized in this paper can be employed to determine user’s budget for desired PoCD performance, and vice versa. In particular, for a given target PoCD (e.g., as specified in the SLAs), users can select the corresponding scheduling strategy and optimize its parameters, to optimize the required budget/cost to achieve the PoCD target.

The main contributions of this paper are summarized as follows:

- We propose a novel framework, Chronos, which unifies Clone, Speculative-Restart, and Speculative-Resume strategies and enable the optimization of speculative execution of deadline-sensitive MapReduce jobs.
- We define PoCD to quantify the probability that a MapReduce job meets its individual application deadline, and analyze PoCD and execution (VM time) cost for the three strategies in closed-form.
- A joint optimization framework is developed to balance PoCD and cost through a maximization of net utility.

We develop an efficient algorithm that solves the non-convex optimization and is guaranteed to find an optimal solution.

- Chronos is prototyped on Hadoop MapReduce and evaluated against three baseline strategies using both experiments and trace-driven simulations. Chronos on average outperforms existing solutions by 50% in net utility increase, with up to 80% PoCD and 88% cost improvements.

II. RELATED WORK

There have been several research efforts to improve the execution time of MapReduce-like systems to guarantee meeting Quality of Service (QoS) [12, 16–40]. Some focus on static resource provisioning to meet a given deadline in MapReduce, while others propose resource scaling in response to resource demand and cluster utilization in order to minimize the overall cost. Moreover, [41–43] proposed frameworks to improve MapReduce job performance. These papers are similar to our proposed work in the sense that resources are optimized to minimize energy consumption and reduce operating cost. However, the above mentioned works do not optimize job execution times in the presence of stragglers, as we do in this work.

Efficient task scheduling is critical to reduce the execution time of MapReduce jobs. A large body of research exists on task scheduling in MapReduce with deadlines. These works range from deadline-aware scheduling [12, 13, 44–49] to energy- and network-aware scheduling [50–54]. However, these works do not consider dealing with stragglers which might severely prolong a job’s execution time and violate the QoS [2–4, 8].

The complexity of cloud computing continues to grow as it is being increasingly employed for a wide range of domains such as e-commerce and scientific research. Thus, many researchers have shown interest in mitigating stragglers and improving the default Hadoop speculation mechanism. They proposed new mechanisms to detect stragglers reactively and proactively and launch speculative tasks accordingly [2–4, 6–10]. This is important to ensure providing high reliability to satisfy a given QoS, as it can be at risk when stragglers exist or when failures occur. Different from these works, we jointly maximize the probability of meeting job deadlines and minimize the cost resulting from speculative/duplicate task execution and find the optimal number of speculative copies for each task. Note that in contrast to launching one copy, we launch r extra attempts. Moreover, our work includes introducing an accurate way to estimate the task finishing time by taking the JVM launching time into account, which in turn reduces the number of false positive decisions in straggler detection.

In addition, a number of works have considered avoiding re-executing the work done by the original tasks (stragglers) [55–58]. Basically, the key idea is to checkpoint running tasks and make the speculative task start from the checkpoint. This idea is similar to our Speculative-Resume strategy. However, in our strategy, we detect stragglers based on jobs’ deadlines and

launch the optimal number of extra attempts for each straggler in order to jointly optimize PoCD and cost. To save cost, we check the progress of all attempts and keep the fastest one when attempts' progress can be relatively accurately estimated. To save overhead, we only check progress of tasks two times, once for detecting stragglers, and another time to kill slower extra attempts.

III. BACKGROUND AND SYSTEM MODEL

Consider M MapReduce jobs that are submitted to a datacenter, where job i is associated with a deadline D_i and consists of N_i tasks for $i = 1, 2, \dots, M$. Job i meets the desired deadline if all its N_i tasks are processed by the datacenter before time D_i .¹ Tasks whose completion time exceed D are considered as stragglers. To mitigate stragglers, we launch multiple parallel attempts for each task belonging to job i , including one original attempt and r_i speculative/extra attempts. A task is completed as soon as one of its $r_i + 1$ attempts is successfully executed. Let $T_{i,j,k}$ denote the (random) execution time of attempt k of job i 's task j . Thus, we find job completion time T_i and task completion time $T_{i,j}$ by:

$$T_i = \max_{j=1, \dots, N_i} T_{i,j}, \text{ where } T_{i,j} = \min_{k=1, \dots, r_i+1} T_{i,j,k}, \forall j. \quad (1)$$

The Pareto distribution is proposed to model the execution times of tasks in [59], and is used in [6, 7, 14] to analyze the straggler problem. Following these papers, we assume the execution time $T_{i,j,k}$ of each attempt follows a Pareto distribution with parameters t_{\min} and β . The probability density function of $T_{i,j,k}$ is

$$f_{T_{i,j,k}}(t) = \begin{cases} \frac{\beta \cdot t_{\min}^\beta}{t^{\beta+1}} & t \geq t_{\min}, \\ 0 & t < t_{\min}, \end{cases} \quad (2)$$

where t_{\min} is the minimum execution time and β is the tail index, while different attempts are assumed to have independent execution times.

As introduced earlier, Chronos consists of three scheduling strategies to mitigate stragglers, namely Clone, Speculative-Restart, and Speculative-Resume. Clone is a proactive approach, wherein $r + 1$ copies of a task (i.e., one original attempt and r extra attempts) are launched simultaneously. Here, r is a variable that is optimized to balance the PoCD with the cost of execution. Speculative-Restart launches r extra copies for each straggler detected. Speculative-Resume is a work-preserving strategy, where the r extra copies of a detected straggler start processing right after the last byte offset processed by the original task. To model the cost of executing job i under each strategy, we consider an on-spot price of γ_i dollars per unit time for each active virtual machine (VM) running attempts/tasks of job i . The price γ_i depends on the VM-type subscribed by job i , and is assumed to be known when optimizing job scheduling strategies. Our goal is to jointly maximize the probability of meeting job deadlines and minimize the cost resulting from speculative/extra task/attempt

¹When speaking in the context of a single task or job, we drop the subscript(s) for clarity.

scheduling. In Chronos, we use progress score to determine if extra attempts are needed. The progress score is defined as the percentage of workload processed at a given time t .

Clone Strategy. Under this strategy, $r + 1$ attempts of each task are launched at the beginning. The progress scores of the $r + 1$ attempts are checked at time τ_{kill} , and the attempt with the best progress is left for processing data, while the other r attempts are killed to save machine running time. Figure 1(a) illustrates the Clone strategy for a single task.

Speculative-Restart Strategy. Under this strategy, one attempt (original) of each task is launched at the beginning. At time τ_{est} , the task attempt completion time is estimated, and if it exceeds D , r extra attempts are launched that start processing data from the beginning.² At time τ_{kill} , the progress scores of all $r + 1$ attempts are checked, and the attempt with the smallest estimated completion time is left running, while the other r attempts are killed. Figure 1(b) illustrates the Speculative-Restart strategy for a task when the execution time of the original attempt exceeds D .

Speculative-Resume Strategy. This strategy is similar to the Speculative-Restart strategy in its straggler detection. The difference is that the detected straggler is killed and $r + 1$ attempts are launched for the straggling task. These attempts, however, do not reprocess the data that has already been processed by the original attempt, and start processing the data after the last byte offset when the straggler is detected. At time τ_{kill} , the progress scores of all attempts are checked, and the attempt with the smallest estimated completion time is left running while the other r attempts are killed. Figure 1(c) illustrates the Speculative-Resume strategy for a task for the case when the execution time of the original attempt exceeds D . Here, the processed byte offset of the original attempt at τ_{est} is b . Extra attempts launched at τ_{est} start to process data from byte offset b .

IV. ANALYSIS OF PoCD AND MACHINE RUNNING TIME

We now formally define PoCD, and analyze PoCD for each strategy. PoCD expressions for the three strategies, under the assumption that task attempt execution times are iid Pareto, are presented in Theorems 1, 3, and 5. We also analyze the machine running time for each strategy, and present expressions for these in Theorems 2, 3, and 5.

Definition PoCD is the probability that an arriving job completes before its deadline.

We use the notation R_{strategy} to denote the PoCD of a particular strategy, which can be Clone, Speculative-Restart, or Speculative-Resume.

A. Clone

We start by analyzing Clone, and derive PoCD and machine running time expressions in *Theorem 1* and *Theorem 2*, respectively.

²No extra attempts are launched if the estimated finishing time is not greater than D .

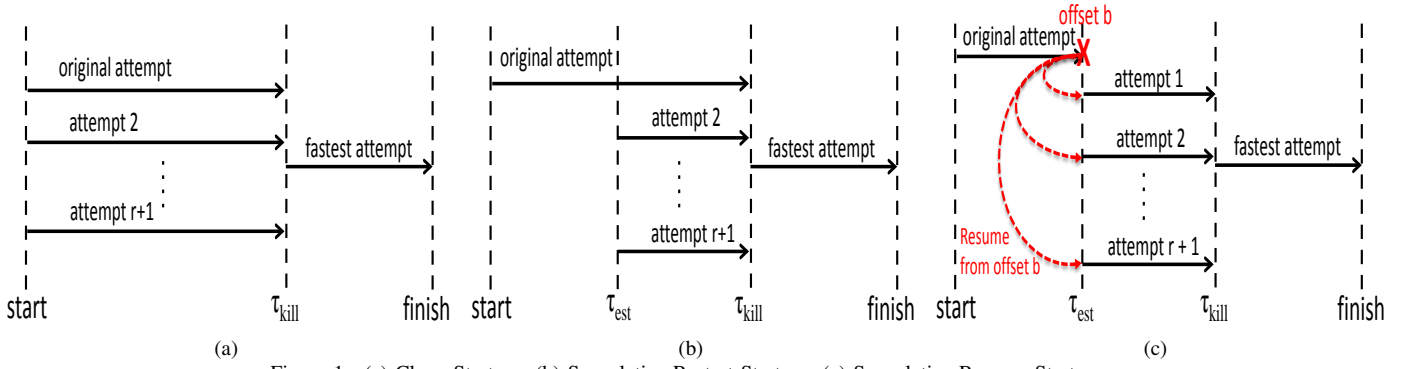


Figure 1. (a) Clone Strategy, (b) Speculative-Restart Strategy, (c) Speculative-Resume Strategy.

Theorem 1. Under Clone strategy, the PoCD

$$R_{\text{Clone}} = \left[1 - \left(\frac{t_{\min}}{D} \right)^{\beta \cdot (r+1)} \right]^N. \quad (3)$$

Proof. We first derive the probability that a single task completes before the deadline, and then derive the PoCD by considering all N tasks of the job.

Let us denote the probability that a single attempt of a task fails to finish before deadline D by P_{Clone} . Then, we have

$$P_{\text{Clone}} = \int_D^\infty \frac{\beta t_{\min}^\beta}{t^{\beta+1}} dt = \left(\frac{t_{\min}}{D} \right)^\beta. \quad (4)$$

The task fails to finish before D when all $r+1$ attempts fail to finish before D . Thus, the probability that a task finishes before D is $1 - (P_{\text{Clone}})^{r+1}$. The job finishes before deadline D when all its N tasks finish before D . Thus, the PoCD is given by

$$R_{\text{Clone}} = [1 - (P_{\text{Clone}})^{r+1}]^N = \left[1 - \left(\frac{t_{\min}}{D} \right)^{\beta \cdot (r+1)} \right]^N. \quad (5)$$

Before moving to *Theorem 2*, we first introduce and prove *Lemma 1* for computing the expected execution time of $\min(T_{j,1}, \dots, T_{j,n})$, where n is a positive integer.

Lemma 1. Let $W = \min(T_{j,1}, \dots, T_{j,n})$, where $T_{j,a}, a = 1, 2, \dots, n$ follows the Pareto distribution with parameters t_{\min} and β . Then, $E(W)$ can be computed as:

$$E(W) = \frac{t_{\min} \cdot n \cdot \beta}{n \cdot \beta - 1}. \quad (6)$$

Proof. We introduce

$$W - t_{\min} = \int_{t_{\min}}^W dt = \int_{t_{\min}}^\infty I(t) dt, \quad (7)$$

where $I(t)$ is the indicator function for event $\{W > t\}$:

$$I(t) = I\{W > t\} \stackrel{\text{def}}{=} \begin{cases} 1, & W > t; \\ 0, & W \leq t. \end{cases} \quad (8)$$

Taking expectations we have that

$$E(W) = \int_{t_{\min}}^\infty E(I(t)) dt + t_{\min} = \int_{t_{\min}}^\infty P(W > t) dt + t_{\min}. \quad (9)$$

Since $P(W > t) = [P(T_{j,a} > t)]^n$, we have:

$$E(W) = \int_{t_{\min}}^\infty \left(\frac{t_{\min}}{t} \right)^{n \cdot \beta} dt + t_{\min} = \frac{t_{\min} \cdot n \cdot \beta}{n \cdot \beta - 1} \quad (10)$$

□

Theorem 2. Under Clone strategy, the expected execution time of a job

$$E_{\text{Clone}}(T) = N \cdot \left[r \cdot \tau_{\text{kill}} + t_{\min} + \frac{t_{\min}}{\beta \cdot (r+1) - 1} \right], \quad (11)$$

where T denotes execution time of a job.

Proof. We first derive the machine running time of a task by adding the machine running time of attempts killed at τ_{kill} and the machine running time of the attempt that successfully completes, and then get the job's machine running time by adding the machine running time of all N tasks.

$E_{\text{Clone}}(T)$ equals the expectation of the machine running time of all N tasks, i.e., $E_{\text{Clone}}(T) = N \cdot E(T_j)$, where T_j is the machine running time of task j . Also, $E(T_j)$ equals the sum of the machine running times of the r killed attempts and the execution time of the attempt with the best progress score at τ_{kill} . Denote $\min\{T_{j,1}, \dots, T_{j,r+1}\}$ as W_j^{all} , where $T_{j,a}$ is the execution time of attempt a belonging to task j . Then,

$$E(T_j) = r \cdot \tau_{\text{kill}} + E(W_j^{\text{all}}). \quad (12)$$

We can use the result from *Lemma 1* to compute $E(W_j^{\text{all}})$. The parameter n in *Lemma 1* equals $r+1$. Thus,

$$E(W_j^{\text{all}}) = \frac{t_{\min} \cdot \beta \cdot (r+1)}{\beta \cdot (r+1) - 1} = t_{\min} + \frac{t_{\min}}{\beta \cdot (r+1) - 1}. \quad (13)$$

□

B. Speculative-Restart

Here we present the PoCD and machine running time analysis for the Speculative-Restart strategy in *Theorem 3* and *Theorem 4*, respectively.

Theorem 3. Under Speculative-Restart, the PoCD

$$R_{\text{S-Restart}} = \left[1 - \frac{t_{\min}^{\beta \cdot (r+1)}}{D^\beta \cdot (D - \tau_{\text{est}})^{\beta \cdot r}} \right]^N. \quad (14)$$

Proof. As in the proof of *Theorem 1*, we first derive the probability that a task completes before the deadline. Then,

we can obtain the PoCD by considering that all N tasks of the job complete before the deadline. \square

Theorem 4. *Under Speculative-Restart, the expected execution time of a job, $E_{S\text{-Restart}}(T)$, equals*

$$E(T_j|T_{j,1} \leq D) \cdot P(T_{j,1} \leq D) + E(T_j|T_{j,1} > D) \cdot P(T_{j,1} > D), \quad (15)$$

where

$$\begin{aligned} P(T_{j,1} > D) &= 1 - P(T_{j,1} \leq D) = \left(\frac{t_{\min}}{D}\right)^\beta, \\ E(T_j|T_{j,1} \leq D) &= \frac{t_{\min} \cdot D \cdot \beta \cdot (t_{\min}^{\beta-1} - D^{\beta-1})}{(1-\beta) \cdot (D^\beta - t_{\min}^\beta)}, \\ E(T_j|T_{j,1} > D) &= \tau_{\text{est}} + r \cdot (\tau_{\text{kill}} - \tau_{\text{est}}) \\ &+ \frac{t_{\min}}{\beta \cdot r - 1} - \frac{t_{\min}^{\beta \cdot r}}{(\beta \cdot r - 1) \cdot (D - \tau_{\text{est}})^{\beta \cdot r - 1}} \\ &+ \int_{D - \tau_{\text{est}}}^{\infty} \left(\frac{D}{\omega + \tau_{\text{est}}}\right)^\beta \cdot \left(\frac{t_{\min}}{\omega}\right)^{\beta \cdot r} d\omega + t_{\min} \end{aligned} \quad (16)$$

Proof. We first derive the machine running time of a task by considering if the execution time of the original attempt is larger than D . If the execution time is no more than D (denoted by $E(T_j|T_{j,1} \leq D)$ in (16)), there is no extra attempt launched, and the machine running time is the execution time of the original attempt.

In the case that execution time is larger than D , the machine running time (denoted by $E(T_j|T_{j,1} > D)$ in (16)) consists of three parts, i.e., (i) execution of the original attempt from start to τ_{est} , (ii) machine time need to run $r+1$ attempts between τ_{est} and τ_{kill} , and (iii) execution of the fastest attempt from τ_{kill} until it finishes. Due to space limitation, we omit the proof details and refer to Section IX. \square

C. Speculative-Resume

Finally, we obtain PoCD and machine running time expressions for Speculative-Resume in *Theorems 5* and *6*, respectively. Let us denote the average progress of the original attempts at time τ_{est} as $\varphi_{j,\text{est}}$.

Theorem 5. *Under Speculative-Resume, the PoCD*

$$R_{S\text{-Resume}} = \left[1 - \frac{(1 - \varphi_{j,\text{est}})^{\beta \cdot (r+1)} \cdot t_{\min}^{\beta \cdot (r+2)}}{D^\beta \cdot (D - \tau_{\text{est}})^{\beta \cdot (r+1)}} \right]^N. \quad (17)$$

Proof. The proof is similar to the proof of Theorem 1, and details are omitted. \square

Theorem 6. *Under Speculative-Resume, the expected execution time of a job $E_{S\text{-Resume}}(T)$ equals*

$$E(T_j|T_{j,1} \leq D) \cdot P(T_{j,1} \leq D) + E(T_j|T_{j,1} > D) \cdot P(T_{j,1} > D), \quad (18)$$

where

$$P(T_{j,1} > D) = 1 - P(T_{j,1} \leq D) = \left(\frac{t_{\min}}{D}\right)^\beta, \quad (19)$$

$$E(T_j|T_{j,1} \leq D) = \frac{t_{\min} \cdot D \cdot \beta \cdot (t_{\min}^{\beta-1} - D^{\beta-1})}{(1-\beta) \cdot (D^\beta - t_{\min}^\beta)}, \text{ and} \quad (20)$$

$$E(T_j|T_{j,1} > D) = \tau_{\text{est}} + r \cdot (\tau_{\text{kill}} - \tau_{\text{est}}) \quad (21)$$

$$+ \frac{t_{\min} \cdot (1 - \varphi_{j,\text{est}})^{\beta \cdot (r+1)}}{\beta \cdot (r+1) - 1} + t_{\min} \quad (22)$$

Proof. The proof is similar to the proof of Theorem 4, and is omitted. \square

We note that our analysis of PoCD and cost (including proof techniques of Theorems 1-6) actually works with other distributions as well (even though the exact equations/values may change in the derived results, depending on the distribution).

D. Comparing PoCD of different strategies

In this subsection, we compare the PoCDs of three strategies, and present results in Theorem 7. We denote $D - \tau_{\text{est}}$ as \bar{D} , and $1 - \varphi_{j,\text{est}}$ as $\bar{\varphi}_{j,\text{est}}$.

Theorem 7. *Given r , we can get three conclusions:*

- 1) $R_{\text{clone}} > R_{S\text{-Restart}}$,
- 2) $R_{S\text{-Resume}} > R_{S\text{-Restart}}$,
- 3) if $r > \frac{\beta \cdot \ln(\bar{\varphi}_{j,\text{est}} \cdot t_{\min}) - \ln \bar{D}}{\ln \bar{D} - \ln(\bar{\varphi}_{j,\text{est}} \cdot D)}$, $R_{\text{clone}} > R_{S\text{-Resume}}$;
Otherwise, $R_{\text{clone}} \leq R_{S\text{-Resume}}$.

We provide an outline of our proof. Given r , Clone strategy launches r extra attempts for each task from the beginning. For each straggler detected at τ_{est} , S-Restart strategy launches r extra new attempts (that restart from zero) and keeps the original attempt, while S-Resume strategy launches r extra attempts, which resume to process the remaining $\bar{\varphi}_{j,\text{est}}$ data, for each straggler. If r is not large, it is clearly better to kill the straggler, and launch extra $r+1$ attempts to process remaining workload, instead of leaving straggler running.

V. JOINT POCD AND COST OPTIMIZATION

Starting multiple speculative/clone tasks leads to higher PoCD, but also results in higher execution cost, which is proportional to the total VM time required for both original and speculative/clone tasks, e.g., when VMs are subscribed from a public cloud with a usage-based pricing mechanism. To exploit this tradeoff, we consider a joint optimization framework for the three proposed strategies, to maximize a ‘‘net utility’’ defined as PoCD minus execution cost. More precisely, we would like to:

$$\max U(r) = f(R(r) - R_{\min}) - \theta \cdot C \cdot \mathbb{E}(T), \quad (23)$$

$$\text{s.t. } r \geq 0, \quad (24)$$

$$\text{var. } r \in \mathbb{Z}. \quad (25)$$

Here, $R(r)$ is the PoCD that results by initiating r speculative/clone tasks in Clone, S-Restart, or S-Resume strategies. To guarantee a minimum required PoCD, we consider a utility function $f(R(r) - R_{\min})$, which is an increasing function and

drops to negative infinity if $R(r) < R_{\min}$. The usage-based VM price per unit time is C .

We use a tradeoff factor $\theta \geq 0$ to balance the PoCD objective $f(R(r) - R_{\min})$ and the execution cost $C \cdot \mathbb{E}(T)$. By varying θ , the proposed joint optimization can generate a wide range of solutions for diverse application scenarios, ranging from PoCD-critical optimization with small tradeoff factor θ and cost-sensitive optimization with large θ . Finally, while our joint optimization framework applies to any concave, increasing utility function f , in this paper, we focus on logarithmic utility functions, $f(R(r) - R_{\min}) = \lg(R(r) - R_{\min})$, which is known to achieve proportional fairness [60]. In the following, we will prove the concavity of the optimization objective $U(r)$ for different strategies and propose an efficient algorithm to find the optimal solution to the proposed joint PoCD and cost optimization.

A. Optimizing Chronos

To evaluate convexity of the optimization objective $U_{\text{strategy}}(r)$, we first present the following lemma on function compositions.

Lemma 2. [61] *Suppose $H(x)$ is a composite of $f(x)$ and $g(x)$, i.e., $H(x) = f(g(x))$. If $f(x)$ is increasing and both $f(x)$ and $g(x)$ are concave, then $H(x)$ is a concave function.*

Theorem 8. *The optimization objective $U_{\text{strategy}}(r)$ under Clone, S-Restart, and S-Resume, i.e.,*

$$U_{\text{strategy}}(r) = \lg(R_{\text{strategy}}(r) - R_{\min}) - \theta \cdot C \cdot E_{\text{strategy}}(T), \quad (26)$$

is a concave function of r , when $r > \Gamma_{\text{strategy}}$, where

$$\Gamma_{\text{Clone}} = -\beta^{-1} \cdot \log_{t_{\min}/D} N - 1. \quad (27)$$

$$\Gamma_{\text{S-Restart}} = \beta^{-1} \cdot \log_{t_{\min}/(D-\tau_{\text{est}})} \frac{D^\beta}{N \cdot t_{\min}^\beta}. \quad (28)$$

$$\Gamma_{\text{S-Resume}} = \beta^{-1} \cdot \log_{\frac{(1-\varphi_{j,\text{est}}) \cdot t_{\min}}{D-\tau_{\text{est}}}} \frac{D^\beta}{N \cdot t_{\min}^\beta} - 1. \quad (29)$$

Proof. $\lg(R_{\text{strategy}}(r) - R_{\min})$ is an increasing and concave function of $R_{\text{strategy}}(r)$. Also, when $r > \Gamma_{\text{strategy}}$, the second derivative of R_{strategy} is less than 0. So, R_{strategy} is a concave function of r , when $r > \Gamma_{\text{strategy}}$. Based on Lemma 2, we know that $\lg(R_{\text{Clone}} - R_{\min})$ is a concave function of r , when $r > \Gamma_{\text{strategy}}$.

The second derivative of $-E_{\text{strategy}}(T)$ is less than 0. Since the sum of two concave functions is also a concave function, $U_{\text{strategy}}(r)$ is a concave function when $r > \Gamma_{\text{strategy}}$. \square

We note that for non-deadline sensitive jobs, as job deadlines increase and become sufficiently large, the optimal r will approach zero. Thus, there is no need to launch speculative or clone attempts for these non-deadline sensitive jobs.

B. Our Proposed Optimization Algorithm

We present the algorithm for solving the optimal r to maximize the ‘‘PoCD minus cost’’ objective for the three strategies. First, our analysis shows that the optimization objective is guaranteed to be concave when r is above a certain threshold, i.e., $r > \Gamma_{\text{strategy}}$, for Clone, Speculative-Restart, and Speculative-Resume strategies. Thus, we can leverage the concavity property and employ a gradient descent method to find the optimal objective $U_{\text{strategy,opt}}(r_{\text{opt}})$ and solution r_{opt} , when the objective is concave, i.e., $r > \Gamma_{\text{strategy}}$ for different threshold Γ_{strategy} . Second, for the non-concave regime, $r \leq \Gamma_{\text{strategy}}$, which contains a limited number of integer points (typically, less than 4, as we will see later), we use a simple search to find the maximum objective value over all integers $r \leq \Gamma_{\text{strategy}}$. This hybrid algorithm can efficiently find the optimal solution to the joint PoCD and cost optimization, due to the optimality of convex optimization when $r > \Gamma_{\text{strategy}}$, as well as the limited complexity for searching $r \leq \Gamma_{\text{strategy}}$. The proposed algorithm is summarized in Algorithm 1. In Algorithm 1, η , α and ξ determine the accuracy and complexity of the algorithm, and their values depend on the values of N , θ , τ_{est} , τ_{kill} and C .

Theorem 9. *Algorithm 1 finds an optimal solution for the net utility optimization problem in (23).*

Proof. For $r \geq \lceil \Gamma_{\text{strategy}} \rceil$, $U_{\text{Strategy}}(r)$ is a concave function of r , and the Gradient-based Line Search can find the maximum value of $U_{\text{Strategy}}(r)$ in polynomial time, where *strategy* can be either Clone, S-Restart, or S-Resume. We can then compute the $U_{\text{Strategy}}(r)$, $\forall r \in 0, \dots, \lceil \Gamma_{\text{strategy}} \rceil - 1$, and compare the result from Gradient-based Line Search to get the maximum value of $U_{\text{Strategy}}(r)$, $\forall r \geq 0$. Since Γ_{strategy} is typically a small number, Algorithm 1 can find the optimal r for maximizing net utility. \square

VI. IMPLEMENTATION

We develop Chronos (including Clone, Speculative-Restart and Speculative-Resume) using Hadoop YARN, which consists of a central Resource Manager (RM), a per-application Application Master (AM) and a per-node Node Manager (NM). The AM negotiates resources from the RM and works with the NMs to execute and monitor an application’s tasks. Our optimization algorithm is implemented in the AM to calculate the optimal r upon job submission to be used by all strategies. Chronos uses the progress score which is provided by Hadoop to estimate the expected completion time of a given task. The progress score of the map phase represents the fraction of data processed. In Chronos, we modified the default Hadoop speculation method and implemented our new progress estimation mechanism, which significantly improves the estimation accuracy by taking into account the time overhead for launching MapReduce tasks, which is ignored by Hadoop.

A. Implementing Clone Strategy

Upon submission of a job, the AM creates tasks and queues them for scheduling. Then, the AM asks the RM for containers

Algorithm 1: Unifying Optimization Algorithm

```
strategy={Clone, S-Restart, S-Resume}
//Phase 1 : Gradient-based Line Search [61]
 $r = \lceil \Gamma_{\text{strategy}} \rceil$ 
while  $|\nabla U_{\text{Strategy}}(r)| > \eta$  do
   $\Delta r = -\nabla U_{\text{Strategy}}(r)$ 
   $\varepsilon = 1$ 
  while  $U_{\text{Strategy}}(r + \varepsilon \cdot \Delta r) >$ 
 $U_{\text{Strategy}}(r) + \alpha \cdot \varepsilon \cdot \nabla U_{\text{Strategy}}(r) \cdot \Delta r$  do
     $\varepsilon = \xi \cdot \varepsilon$ 
  end while
   $x = x + t \cdot \Delta r$ 
end while
 $U_{\text{strategy,opt}}(r_{\text{opt}}) = U_{\text{strategy}}(r)$ 
 $r_{\text{opt}} = r$ 
//Phase 2 : Finding optimal solution for  $r \geq 0$ 
for  $r = 0 : \lceil \Gamma_{\text{strategy}} \rceil - 1$  do
  Compute utility function  $U_{\text{strategy}}(r)$ 
  if  $U_{\text{strategy}}(r) > U_{\text{strategy,opt}}(r_{\text{opt}})$  then
     $U_{\text{strategy,opt}}(r_{\text{opt}}) = U_{\text{strategy}}(r)$ 
     $r_{\text{opt}} = r$ 
  end if
end for
```

to run its tasks. In Clone, the AM solves the proposed joint optimization problem to determine the optimal r before creating tasks for a given job. As discussed in Section V, r is the number of extra attempts for each task. Once r is found, whenever a map task belonging to the job is created, the AM also creates extra r copies of the same task, which start execution simultaneously with the original one. After launching the $r + 1$ attempts, the AM monitors all attempts' progress scores, and after τ_{kill} , the attempt with the highest progress score is kept alive and will continue running, while all other slower clone tasks are killed.

B. Implementing Speculation Strategies

Unlike Hadoop, Chronos launches r speculative attempts for each straggler task whose estimated completion time falls behind the target deadline. The optimal number of attempts r is calculated by the AM after a job is submitted, by solving the joint PoCD and cost optimization. The AM monitors all active tasks and launches speculative attempts for the tasks that have an estimated completion time (t_{ect}) larger than their target deadline D . In particular, we have $t_{\text{ect}} = t_{\text{lau}} + t_{\text{eet}}$, where t_{lau} and t_{eet} are the time at which the task is launched and the estimated execution time, respectively.³ In our experiments, we observed that Hadoop's default task completion time estimation is highly unreliable. In Hadoop, the task estimated execution time equals the difference between the current time and the time at which the task is launched, divided by the reported progress score.

³All times are relative to the start time of a job, which is assumed to be 0. Tasks are launched a little after time 0, namely t_{lau} .

The task completion time estimation error is mainly caused by Hadoop's assumption that a task starts running right after it is launched. However, this is not true especially in highly contended environments where JVM startup time is significant and cannot be ignored. We therefore take into consideration the time to launch a JVM when estimating task progress and completion time. Task execution time consists of two parts, i.e., the time for launching JVM, and the time for processing the workload. Chronos calculates the time for launching JVM by finding the difference between the time when the first progress report is received (t_{FP}) and t_{lau} . Therefore, the new estimated completion time is given by

$$t_{\text{ect}} = t_{\text{lau}} + (t_{\text{FP}} - t_{\text{lau}}) + \frac{t_{\text{now}} - t_{\text{FP}}}{CP - FP}, \quad (30)$$

where $\frac{t_{\text{now}} - t_{\text{FP}}}{CP - FP}$ is the time for processing the workload, and FP and CP are the first reported progress value and current reported progress value, respectively.

1) *Speculative-Restart*: For this strategy, the AM monitors all running tasks and estimates their t_{ect} . If the AM finds that a task's estimated completion time t_{ect} at τ_{est} exceeds deadline D , an optimal number r of extra attempts are created and launched. At τ_{kill} , all attempts are killed except for the one with the highest progress score. Unlike default Hadoop, Chronos launches speculative copies once it detects a straggler at τ_{est} .

2) *Speculative-Resume*: Our implementation of Speculative-Resume employs a work-preserving mechanism. The key idea is that the AM keeps track of an (original) attempt's data processing progress, maintains the latest (data) record's offset, and passes the offset to the new speculative attempts at τ_{est} if speculation is determined to be necessary. It ensures that the speculative attempts do not process the already processed data and can seamlessly "resume" executing data where the original attempt left off.

A challenge in implementing Speculative-Resume is that Chronos needs to consider the launching time of new speculative attempts, since JVM startup time postpones the actual start of execution of new attempts. Our solution estimates the number of bytes that will be processed (b_{extra}) by the original attempt during speculative attempts' launching time. Then, the speculative attempts will skip these bytes and start processing data from an anticipated offset. The speculative launching mechanism allows the original and speculative attempts to switch seamlessly and effectively avoids any JVM startup time overhead. More precisely, at τ_{est} , Chronos records the number of bytes processed by the original attempt (b_{est}), and then estimates the number of bytes to be processed by the original attempt during speculative attempt's startup time (b_{extra}) as follows:

$$\frac{b_{\text{est}}}{\tau_{\text{est}} - t_{\text{FP}}} \cdot (t_{\text{FP}} - t_{\text{lau}}). \quad (31)$$

Next, our new byte offset to be used by the speculative attempts to start processing data is calculated as $b_{\text{new}} = b_{\text{start}} + b_{\text{est}} + b_{\text{extra}}$, where b_{start} is the offset of the first byte of the split, which is the data to be processed by a task. Note

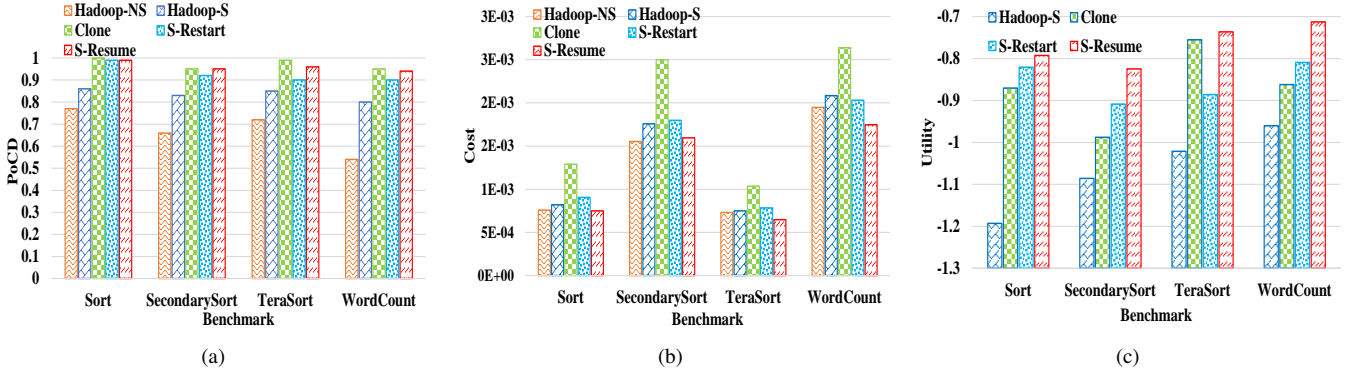


Figure 2. Comparisons of HNS, HS, Clone, S-Restart and S-Resume with respect to PoCD, Cost and Utilities with different benchmarks.

that in this strategy, Chronos launches $r + 1$ new copies for each straggler after killing the original one. This is because the original attempt was deemed to be a straggler (for reasons such as node malfunction or operating system disruption), and since all new attempts start from b_{new} , there is no data-processing delay for relaunching (and rejuvenating) the original attempt at a new location.

VII. EVALUATION

We evaluate Clone, Speculative Restart (S-Restart), and Speculative Resume (S-Resume) through both testbed experiments and large-scale simulations with real-world data center trace. In all three strategies, the optimal number r_{opt} of clone/speculative attempts are found by solving our proposed PoCD and cost optimization. We compare the three strategies with Default Hadoop without Speculation (Hadoop-NS), Default Hadoop with Speculation (Hadoop-S), and Mantri, which serve as baselines in our evaluation.

A. Experiment Results

1) *Setup*: We deploy our prototype of Clone, Speculative Restart, and Speculative Resume strategies on Amazon EC2 cloud testbed consisting of 40 nodes. Each node has 8 vCPUs and 2GB memory. The three strategies are evaluated by using the Map phases of four classic benchmarks, i.e., Sort, SecondarySort, TeraSort, and WordCount. Sort and SecondarySort are I/O bound applications, whereas WordCount and the map phase of TeraSort are CPU-bound. We generate workloads for Sort and TeraSort by using RandomWriter and TeraGen applications, respectively. Also, we generate random number pairs as workload for SecondarySort. The sizes of all workloads are 1.2GB. To emulate a realistic cloud infrastructure with resource contentions, we introduce background applications in the host OS of the physical servers, which repeatedly execute some computational tasks and inject background noise in our testbed. We observe that the task execution time measured on our testbed follows a Pareto distribution with an exponent $\beta < 2$. This is consistent with observations made in [6, 7, 14, 59]. As shown in Morpheus [62] and Jockey [63], job deadlines can be introduced by third parties and/or specified in Service Level Agreements (SLAs). In our experiments, we consider jobs with different deadline sensitivities by setting deadlines as different ratios of average job execution time. To evaluate the impact of background noise (in resource utilization), We use

the Stress utility to generate background applications. Stress is a tool to impose a configurable amount of CPU, memory, I/O, and disk stress on the system, and it is widely used to evaluate cloud and VM performance, such as in [64–67].

2) *Results*: In the following, we compare Hadoop-NS, Hadoop-S, Clone, S-Restart, and S-Resume with respect to PoCD, cost, and total net utility. We execute 100 MapReduce jobs using Chronos on our testbed. Each job has 10 tasks and a deadline equal to 100 sec or 150 sec (emulating jobs with different deadline sensitivities). We measure the PoCD by calculating the percentage of jobs that finish before their deadlines, and the cost by multiplying the machine running time, i.e., the average job running time (i.e., VM time required), and a fixed price per unit VM time that is obtained by Amazon EC2 average spot price (C). In all experiments, we set the tradeoff factor as $\theta = 0.0001$, i.e., 1% of PoCD utility improvement is equivalent to 100 units of VM cost reduction (since $1\% = \theta \cdot 100$). We then solve the corresponding joint PoCD and cost optimization and compare the performance of various algorithms.

Figure 2(a) and Figure 2(b) show PoCD and cost of the five strategies with different benchmarks, respectively. The deadline is 100 sec for Sort and TeraSort, and is 150 sec for SecondarySort and WordCount. For our three strategies, τ_{est} and τ_{kill} are set to 40 sec and 80 sec, respectively. Hadoop-NS has the lowest PoCD, and relatively large cost. Even though Hadoop-NS does not launch extra attempts, the large execution time of stragglers increases the cost. Hadoop-S can launch extra attempts for slow tasks, but only after at least one task finishes. Because Hadoop-S does not consider deadlines, it might launch extra attempts for tasks that may finish before the deadline. Clone launches extra attempts for each task, and it makes both PoCD and cost be the largest among the five strategies. S-Restart and S-Resume launch extra attempts for tasks whose estimated execution time is larger than the deadline. Due to the fact that S-Resume does not re-process data and kills the original attempt, S-Resume can achieve larger PoCD and smaller cost. Figure 2(c) compares the performance of the five strategies in terms of the overall net utility. The results show that our three strategies outperform Hadoop-NS and Hadoop-S by up to 33% on net utility value. In particular, the three strategies can improve PoCD by up to

80% and 19% over Hadoop-NS and Hadoop-S, respectively, while S-Resume introduces little additional cost compared with Hadoop-NS and Hadoop-S. This significant improvement comes from not only launching multiple (optimal number of) attempts for stragglers, but also maintaining only the fastest progress attempt at τ_{kill} , thereby introducing limited execution time overhead. For net utility, since we use the PoCD of Hadoop-NS as R_{min} , its utility is negative infinity.

Table I
PERFORMANCE COMPARISON WITH VARYING τ_{est} AND FIXED $\tau_{kill} - \tau_{est} = 0.5 \cdot t_{min}$.

	τ_{est}	τ_{kill}	PoCD	Cost	Utility
Clone	0	$0.5 \cdot t_{min}$	0.722	9373	-0.376
S-Restart	$0.1 \cdot t_{min}$	$0.6 \cdot t_{min}$	0.996	11458	-0.213
	$0.3 \cdot t_{min}$	$0.8 \cdot t_{min}$	0.988	9650	-0.199
	$0.5 \cdot t_{min}$	$1.0 \cdot t_{min}$	0.938	9486	-0.226
S-Resume	$0.1 \cdot t_{min}$	$0.6 \cdot t_{min}$	0.997	9121	-0.189
	$0.3 \cdot t_{min}$	$0.8 \cdot t_{min}$	0.992	8612	-0.187
	$0.5 \cdot t_{min}$	$1.0 \cdot t_{min}$	0.941	8712	-0.217

Table II
PERFORMANCE COMPARISON WITH VARYING τ_{kill} AND FIXED τ_{est} .

	τ_{est}	τ_{kill}	PoCD	Cost	Utility
Clone	0	$0.4 \cdot t_{min}$	0.718	9113	-0.376
	0	$0.6 \cdot t_{min}$	0.733	9713	-0.369
	0	$0.8 \cdot t_{min}$	0.731	10434	-0.378
S-Restart	$0.3 \cdot t_{min}$	$0.4 \cdot t_{min}$	0.981	8235	-0.189
	$0.3 \cdot t_{min}$	$0.6 \cdot t_{min}$	0.993	8848	-0.188
	$0.3 \cdot t_{min}$	$0.8 \cdot t_{min}$	0.988	9650	-0.199
S-Resume	$0.3 \cdot t_{min}$	$0.4 \cdot t_{min}$	0.987	7935	-0.183
	$0.3 \cdot t_{min}$	$0.6 \cdot t_{min}$	0.993	8220	-0.182
	$0.3 \cdot t_{min}$	$0.8 \cdot t_{min}$	0.992	8612	-0.187

B. Simulation Results

Next, to evaluate our proposed Chronos system at scale, we conduct a trace-driven simulation of a large-scale datacenter. We (i) leverage Google trace [68] to generate MapReduce jobs with different parameters and arrival times, and (ii) make use of spot instance price history from Amazon EC2 to calculate the cost of executing different strategies. We simulate 30 hours of job trace including a total of 2700 jobs with 1 million tasks from Google trace. For each job, we get the start time, the execution time distributions, the number of tasks, and job ID, and use Pareto distribution to generate the execution times of tasks that match the execution time distribution in the real-world trace.

To provide insights on determining the optimal parameters (such as τ_{est} and τ_{kill}) of Chronos, we first study the impact of choosing different parameters in Clone, S-Restart, and S-Resume. Table I shows the performance comparison with varying τ_{est} and fixed difference between τ_{est} and τ_{kill} , which implies fixed execution time $\tau_{kill} - \tau_{est}$ for all clone/speculative attempts.

Under Clone strategy, τ_{est} has only one value, which is 0. Under S-Restart and S-Resume, there exists an interesting tradeoff between estimation accuracy and timeliness of speculation, which is characterized by τ_{est} . Initially, as τ_{est} increases, both PoCD and cost decrease. This is because with

limited observation available for a small τ_{est} , Hadoop tends to overestimate the execution time of attempts at the beginning, while making more accurate estimations as more observations become available for large τ_{est} . More precisely, when τ_{est} is small, more tasks are considered as stragglers, and extra attempts are launched more aggressively. This leads to both high PoCD and large cost. On the other hand, when τ_{est} becomes too large, it might be too late to meet the deadline for stragglers. Based on the table, S-Restart and S-Resume can achieve the highest net utilities when $\tau_{est} = 0.3 \cdot t_{min}$.

Table II shows the performance comparisons for varying τ_{kill} and fixed τ_{est} . We can see that as τ_{kill} increasing, the execution time of clone/speculative attempts increases, resulting in higher cost, while more accurate estimation of progress and execution time can be obtained. Due to higher cost for executing clone/speculative attempts as τ_{kill} increases, the value of optimal r decreases, in order to re-balance PoCD and cost in the joint optimization. As a result, PoCDs are not monotonically increasing or decreasing. Based on the results, the best PoCD can be achieved when $\tau_{kill} = 0.3 \cdot t_{min}$.

In Figure 3(a), 3(b), and 3(c), we compare Mantri, Clone, S-Restart, and S-Resume in terms of PoCD, Cost, and Utility with different θ values to illustrate the tradeoff achieved by different strategies.

Figure 3(a) shows the effect of varying tradeoff factor θ on PoCD. As θ increases, the cost has a higher weight and becomes more critical in the joint optimization. Thus, fewer clone/speculative attempts are launched, and PoCD decreases, leading to decreased r to re-balance PoCD and cost in the optimal solution. It is seen that the PoCD of Clone decreases more than others since it incurs higher cost to execute clone attempts than speculative attempts. To validate this, we also present the histograms of the optimal value of r for Clone with $\theta = 1e-5$, Clone with $\theta = 1e-4$, S-Resume with $\theta = 1e-5$, and S-Resume with $\theta = 1e-4$ in Figure 5. As θ increases from $1e-5$ to $1e-4$, the optimal r for the vast majority of jobs decreases from $r = 2$ to $r = 1$ under Clone strategy, whereas the optimal majority r decreases from $r = 4$ to $r = 3$ under S-Resume strategy. Recall that Mantri aggressively launches and kills extra attempts. Thus, with a larger number of extra attempts running along with original tasks, Mantri can achieve high PoCD, but it also runs up significant cost to execute the speculative attempts. Our results show that S-Resume has the highest PoCD and outperforms Clone strategy by 56%.

Figure 3(b) shows the effect of varying θ on Cost. As θ increases, fewer clone/speculative attempts are launched, and the costs of Clone, S-Restart, and S-Resume all decrease. With a larger number of clone/speculative attempts running along with original tasks, the cost of Mantri is much higher than others. The results show that Mantri introduces 50%, 67%, and 88% higher cost compared with Clone, S-Restart, and S-Resume, respectively, making it non-optimal in terms of net utility optimization.

Figure 3(c) compares the net utility for varying θ . The utility of Mantri decreases the most as θ increases, since it introduces the largest cost compared with other strategies. S-

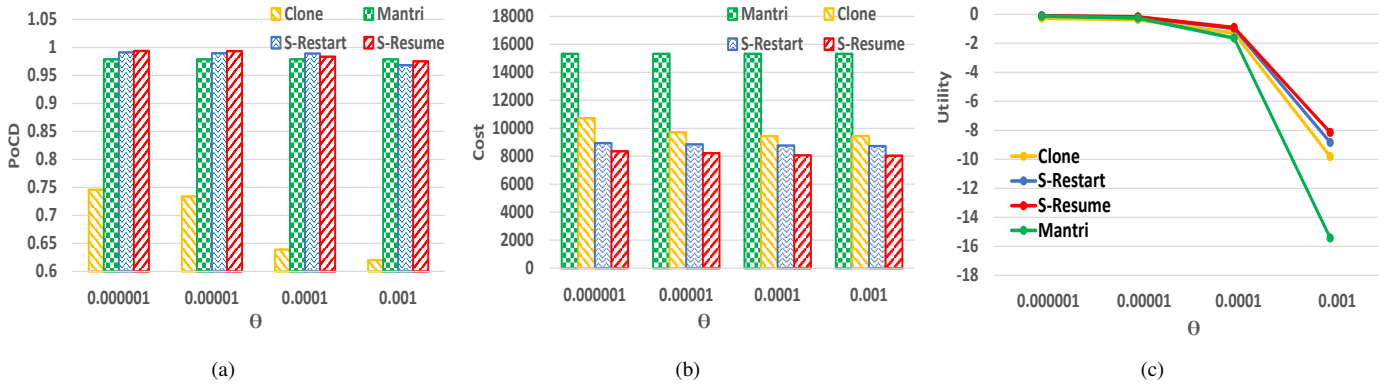


Figure 3. Comparisons of Mantri, Clone, S-Restart and S-Resume on PoCD, Cost, and Utility with different tradeoff factors θ .




Figure 4. Comparisons of HNS, HS, Clone, S-Restart and S-Resume in terms of PoCD, MRT and net utilities with different β s.




Figure 5. Histogram of the optimal r values for different strategies

Resume attains the best net utility, since it has the highest PoCD and lowest cost. Clone, S-Restart, and S-Resume can outperform Mantri by 31%, 43%, and 50%, respectively.

Figure 4(a) and 4(b) compares the performance of different strategies by varying β , the exponent in Pareto distribution. A smaller β corresponds to a smaller decaying in task execution time distribution, implying a heavier tail. In this simulation, we set deadline as 2 times the average task execution time. Because average execution time is $t_{\min} + \frac{t_{\min}}{\beta-1}$ under Pareto distribution, as β increases, the average execution time decreases, which results in decreasing costs. Also, as β increasing, the value of optimal r decreases due to a smaller tail in task execution time.

Figure 4(c) shows net utility achieved by different strategies

under various β . We observe that Clone, S-Restart and S-Resume outperform Hadoop-NS and Hadoop-S when β varies from 1.1 to 1.9.

In summary, as our analysis and experiment results show, if the unit cost of machine running time decreases, Clone becomes more preferable than S-Resume due to smaller up-front cost that is resulted from running the clone attempts (starting from $t = 0$). Further, S-Resume always outperforms S-Restart, since it is able to capitalize on the partially computed results and pick up the task execution from the breakpoint. But we also note that S-Resume may not be possible in certain (extreme) scenarios such as system breakdown or VM crash, where only S-Restart is feasible.

VIII. CONCLUSION

In this paper, we present Chronos, a unifying optimization framework to provide probabilistic guarantees for jobs with deadlines. Our framework aims to maximize the probability of meeting job deadlines and minimize the execution cost. We present and optimize three different strategies to mitigate stragglers proactively and reactively. Our solution includes an algorithm to search for the optimal number of speculative copies needed. Moreover, we present an improved technique to estimate the expected task completion time. Our results show that Chronos can outperform by up to 80% in terms of PoCD over Hadoop-NS, and by 88% in terms of cost over Mantri. Multi-wave executions will be considered in our future work.

REFERENCES

- [1] J. Dean and S. Ghemawat, "Mapreduce: simplified data processing on large clusters," *Communications of the ACM*, vol. 51, no. 1, pp. 107–113, 2008.
- [2] M. Zaharia, A. Konwinski, A. D. Joseph, R. H. Katz, and I. Stoica, "Improving mapreduce performance in heterogeneous environments," in *OSDI*, vol. 8, no. 4, 2008, p. 7.
- [3] G. Ananthanarayanan, S. Kandula, A. G. Greenberg, I. Stoica, Y. Lu, B. Saha, and E. Harris, "Reining in the outliers in map-reduce clusters using mantri," in *OSDI*, vol. 10, no. 1, 2010, p. 24.
- [4] G. Ananthanarayanan, A. Ghodsi, S. Shenker, and I. Stoica, "Effective straggler mitigation: Attack of the clones," in *Presented as part of the 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13)*, 2013, pp. 185–198.
- [5] S. Kandula, S. Sengupta, A. Greenberg, P. Patel, and R. Chaiken, "The nature of data center traffic: measurements & analysis," in *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*. ACM, 2009, pp. 202–208.
- [6] H. Xu and W. C. Lau, "Optimization for speculative execution in big data processing clusters," *IEEE Transactions on Parallel and Distributed Systems*, vol. 28, no. 2, pp. 530–545, 2017.
- [7] D. Wang, G. Joshi, and G. Wornell, "Using straggler replication to reduce latency in large-scale parallel computing," *ACM SIGMETRICS Performance Evaluation Review*, vol. 43, no. 3, pp. 7–11, 2015.
- [8] N. J. Yadwadkar, G. Ananthanarayanan, and R. Katz, "Wrangler: Predictable and faster jobs using fewer resources," in *Proceedings of the ACM Symposium on Cloud Computing*. ACM, 2014, pp. 1–14.
- [9] S. Ibrahim, H. Jin, L. Lu, B. He, G. Antoniu, and S. Wu, "Maestro: Replica-aware map scheduling for mapreduce," in *Cluster, Cloud and Grid Computing (CCGrid), 2012 12th IEEE/ACM International Symposium on*. IEEE, 2012, pp. 435–442.
- [10] J. Rosen and B. Zhao, "Fine-grained micro-tasks for mapreduce skew-handling," *White Paper, University of Berkeley*, 2012.
- [11] M. Isard, M. Budiu, Y. Yu, A. Birrell, and D. Fetterly, "Dryad: Distributed data-parallel programs from sequential building blocks," ser. EuroSys '07. New York, NY, USA: ACM, 2007, pp. 59–72.
- [12] D. Cheng, J. Rao, C. Jiang, and X. Zhou, "Resource and deadline-aware job scheduling in dynamic hadoop clusters," in *Parallel and Distributed Processing Symposium (IPDPS), 2015 IEEE International*. IEEE, 2015, pp. 956–965.
- [13] D. Li, C. Chen, J. Guan, Y. Zhang, J. Zhu, and R. Yu, "Dcloud: deadline-aware resource allocation for cloud computing jobs," *IEEE Transactions on Parallel and Distributed Systems*, 2016.
- [14] X. Ren, G. Ananthanarayanan, A. Wierman, and M. Yu, "Hopper: Decentralized speculation-aware cluster scheduling at scale," *ACM SIGCOMM Computer Communication Review*, vol. 45, no. 4, pp. 379–392, 2015.
- [15] M. Xu, S. Alamro, T. Lan, and S. Subramaniam, "Laser: A deep learning approach for speculative execution and replication of deadline-critical jobs in cloud," in *Computer Communication and Networks (ICCCN), 2017 26th International Conference on*. IEEE, 2017, pp. 1–8.
- [16] X. Xu, M. Tang, and Y.-C. Tian, "Theoretical results of qos-guaranteed resource scaling for cloud-based mapreduce," *IEEE Transactions on Cloud Computing*, vol. PP, no. 99, pp. 1–1, 2016.
- [17] A. Verma, L. Cherkasova, and R. H. Campbell, "Aria: automatic resource inference and allocation for mapreduce environments," in *Proceedings of the 8th ACM international conference on Autonomic computing*. ACM, 2011, pp. 235–244.
- [18] P. Lama and X. Zhou, "Aroma: Automated resource allocation and configuration of mapreduce environment in the cloud," in *Proceedings of the 9th international conference on Autonomic computing*. ACM, 2012, pp. 63–72.
- [19] K. Chen, J. Powers, S. Guo, and F. Tian, "Cresp: Towards optimal resource provisioning for mapreduce computing in public clouds," *IEEE Transactions on Parallel and Distributed Systems*, vol. 25, no. 6, pp. 1403–1412, 2014.
- [20] M. Malekimajd, A. M. Rizzi, D. Ardagna, M. Ciavotta, M. Passacantando, and A. Movaghar, "Optimal capacity allocation for executing mapreduce jobs in cloud systems," in *Symbolic and Numeric Algorithms for Scientific Computing (SYNASC), 2014 16th International Symposium on*. IEEE, 2014, pp. 385–392.
- [21] H. Herodotou, F. Dong, and S. Babu, "No one (cluster) size fits all: automatic cluster sizing for data-intensive analytics," in *Proceedings of the 2nd ACM Symposium on Cloud Computing*. ACM, 2011, p. 18.
- [22] E. Hwang and K. H. Kim, "Minimizing cost of virtual machines for deadline-constrained mapreduce applications in the cloud," in *2012 ACM/IEEE 13th International Conference on Grid Computing*. IEEE, 2012, pp. 130–138.
- [23] K. Kc and K. Anyanwu, "Scheduling hadoop jobs to meet deadlines," in *Cloud Computing Technology and Science (CloudCom), 2010 IEEE Second International Conference on*. IEEE, 2010, pp. 388–392.
- [24] J. Polo, D. Carrera, Y. Becerra, M. Steinder, and I. Whalley, "Performance-driven task co-scheduling for mapreduce environments," in *2010 IEEE Network Operations and Management Symposium-NOMS 2010*. IEEE, 2010, pp. 373–380.
- [25] W. Shi and B. Hong, "Clotho: an elastic mapreduce workload/runtime co-design," in *Proceedings of the 12th International Workshop on Adaptive and Reflective Middleware*. ACM, 2013, p. 5.
- [26] W. Zhang, S. Rajasekaran, T. Wood, and M. Zhu, "Mimp: deadline and interference aware scheduling of hadoop virtual machines," in *Cluster, Cloud and Grid Computing (CCGrid), 2014 14th IEEE/ACM International Symposium on*. IEEE, 2014, pp. 394–403.
- [27] M. AbdelBaky, H. Kim, I. Rodero, and M. Parashar, "Accelerating mapreduce analytics using cometcloud," in *Cloud Computing (CLOUD), 2012 IEEE 5th International Conference on*. IEEE, 2012, pp. 447–454.
- [28] M. Mattess, R. N. Calheiros, and R. Buyya, "Scaling mapreduce applications across hybrid clouds to meet soft deadlines," in *Advanced Information Networking and Applications (AINA), 2013 IEEE 27th International Conference on*. IEEE, 2013, pp. 629–636.
- [29] M. Cardosa, P. Narang, A. Chandra, H. Pucha, and A. Singh, "Steaming: Driving mapreduce provisioning in the cloud," in *2011 18th International Conference on High Performance Computing*. IEEE, 2011, pp. 1–10.
- [30] B. T. Rao and L. Reddy, "Scheduling data intensive workloads through virtualization on mapreduce based clouds," *International Journal of Computer Science and Network Security (IJCSNS)*, vol. 13, no. 6, p. 105, 2013.
- [31] J. Polo, C. Castillo, D. Carrera, Y. Becerra, I. Whalley, M. Steinder, J. Torres, and E. Agyuadé, "Resource-aware adaptive scheduling for mapreduce clusters," in *ACM/IFIP/USENIX International Conference on Distributed Systems Platforms and Open Distributed Processing*. Springer, 2011, pp. 187–207.
- [32] H. Li, X. Wei, Q. Fu, and Y. Luo, "Mapreduce delay scheduling with deadline constraint," *Concurrency and Computation: Practice and Experience*, vol. 26, no. 3, pp. 766–778, 2014.
- [33] S. Li, S. Hu, S. Wang, L. Su, T. Abdelzaher, I. Gupta, and R. Pace, "Woha: deadline-aware map-reduce workflow scheduling framework over hadoop clusters," in *Distributed Computing Systems (ICDCS), 2014 IEEE 34th International Conference on*. IEEE, 2014, pp. 93–103.
- [34] Z. Tang, J. Zhou, K. Li, and R. Li, "A mapreduce task scheduling algorithm for deadline constraints," *Cluster computing*, vol. 16, no. 4, pp. 651–662, 2013.
- [35] J. Wang, Q. Li, and Y. Shi, "Slo-driven task scheduling in mapreduce environments," in *Web Information System and Application Conference (WISA), 2013 10th*. IEEE, 2013, pp. 308–313.
- [36] C.-H. Chen, J.-W. Lin, and S.-Y. Kuo, "Mapreduce scheduling for deadline-constrained jobs in heterogeneous cloud computing systems."
- [37] Z. Liu, Q. Zhang, M. F. Zhani, R. Boutaba, Y. Liu, and Z. Gong, "Dreams: Dynamic resource allocation for mapreduce with data skew," in *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*. IEEE, 2015, pp. 18–26.
- [38] B. Palanisamy, A. Singh, and L. Liu, "Cost-effective resource provisioning for mapreduce in a cloud," *IEEE Transactions on Parallel and Distributed Systems*, vol. 26, no. 5, pp. 1265–1279, 2015.
- [39] N. Lim, S. Majumdar, and P. Ashwood-Smith, "Mrep-rm: a technique for resource allocation and scheduling of mapreduce jobs with deadlines," *IEEE Transactions on Parallel and Distributed Systems*, 2016.
- [40] L. Gu, D. Zeng, P. Li, and S. Guo, "Cost minimization for big data processing in geo-distributed data centers," *IEEE Transactions on Emerging Topics in Computing*, vol. 2, no. 3, pp. 314–323, 2014.
- [41] F. Ahmad, S. T. Chakradhar, A. Raghunathan, and T. Vijaykumar, "Tarazu: optimizing mapreduce on heterogeneous clusters," in *ACM SIGARCH Computer Architecture News*, vol. 40, no. 1. ACM, 2012, pp. 61–74.

[42] B. Heintz, A. Chandra, and J. Weissman, "Cross-phase optimization in mapreduce," in *Cloud Computing for Data-Intensive Applications*. Springer, 2014, pp. 277–302.

[43] Y. Ying, R. Birke, C. Wang, L. Y. Chen, and N. Gautam, "Optimizing energy, locality and priority in a mapreduce cluster," in *Autonomic Computing (ICAC), 2015 IEEE International Conference on*. IEEE, 2015, pp. 21–30.

[44] M. Xu, S. Alamro, T. Lan, and S. Subramaniam, "Cred: Cloud right-sizing with execution deadlines and data locality," *IEEE Transactions on Parallel and Distributed Systems*, vol. 28, no. 12, pp. 3389–3400, 2017.

[45] M. L. Della Vedova, D. Tessera, and M. C. Calzarossa, "Probabilistic provisioning and scheduling in uncertain cloud environments," in *Computers and Communication (ISCC), 2016 IEEE Symposium on*. IEEE, 2016, pp. 797–803.

[46] F. Teng, H. Yang, T. Li, F. Magoulès, and X. Fan, "Mus: a novel deadline-constrained scheduling algorithm for hadoop," *International Journal of Computational Science and Engineering*, vol. 11, no. 4, pp. 360–367, 2015.

[47] Z. Tang, L. Jiang, J. Zhou, K. Li, and K. Li, "A self-adaptive scheduling algorithm for reduce start time," *Future Generation Computer Systems*, vol. 43, pp. 51–60, 2015.

[48] S. Li, S. Hu, and T. Abdelzaher, "The packing server for real-time scheduling of mapreduce workflows," in *21st IEEE Real-Time and Embedded Technology and Applications Symposium*. IEEE, 2015, pp. 51–62.

[49] K. Bok, J. Hwang, J. Lim, Y. Kim, and J. Yoo, "An efficient mapreduce scheduling scheme for processing large multimedia data," *Multimedia Tools and Applications*, pp. 1–24, 2016.

[50] P. Li, L. Ju, Z. Jia, and Z. Sun, "Sla-aware energy-efficient scheduling scheme for hadoop yarn," in *High Performance Computing and Communications (HPCC), 2015 IEEE 7th International Symposium on CyberSpace Safety and Security (CSS), 2015 IEEE 12th International Conference on Embedded Software and Systems (ICES), 2015 IEEE 17th International Conference on*. IEEE, 2015, pp. 623–628.

[51] J. Wang, X. Li, and J. Yang, "Energy-aware task scheduling of mapreduce cluster," in *2015 International Conference on Service Science (ICSS)*. IEEE, 2015, pp. 187–194.

[52] A. Gregory and S. Majumdar, "A constraint programming based energy aware resource management middleware for clouds processing mapreduce jobs with deadlines," in *Companion Publication for ACM/SPEC on International Conference on Performance Engineering*. ACM, 2016, pp. 15–20.

[53] L. Mashayekhy, M. M. Nejad, D. Grosu, Q. Zhang, and W. Shi, "Energy-aware scheduling of mapreduce jobs for big data applications," *IEEE Transactions on Parallel and Distributed Systems*, vol. 26, no. 10, pp. 2720–2733, 2015.

[54] J. Wang and X. Li, "Task scheduling for mapreduce in heterogeneous networks," *Cluster Computing*, vol. 19, no. 1, pp. 197–210, 2016.

[55] Y. Wang, W. Lu, R. Lou, and B. Wei, "Improving mapreduce performance with partial speculative execution," *Journal of Grid Computing*, vol. 13, no. 4, pp. 587–604, 2015.

[56] H. Wang, H. Chen, Z. Du, and F. Hu, "Betl: Mapreduce checkpoint tactics beneath the task level," *IEEE Transactions on Services Computing*, vol. 9, no. 1, pp. 84–95, 2016.

[57] J.-A. Quiane-Ruiz, C. Pinkel, J. Schad, and J. Dittrich, "Rafting mapreduce: Fast recovery on the raft," in *2011 IEEE 27th International Conference on Data Engineering*. IEEE, 2011, pp. 589–600.

[58] Y. Wang, H. Fu, and W. Yu, "Cracking down mapreduce failure amplification through analytics logging and migration," in *Parallel and Distributed Processing Symposium (IPDPS), 2015 IEEE International*. IEEE, 2015, pp. 261–270.

[59] G. Ananthanarayanan, M. C.-C. Hung, X. Ren, I. Stoica, A. Wierman, and M. Yu, "Grass: trimming stragglers in approximation analytics," in *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*, 2014, pp. 289–302.

[60] Z. Huang, B. Balasubramanian, M. Wang, T. Lan, M. Chiang, and D. H. Tsang, "Need for speed: Cora scheduler for optimizing completion-times in the cloud," in *INFOCOM*, 2015.

[61] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.

[62] S. A. Jyothi, C. Curino, I. Menache, S. M. Narayanamurthy, A. Tumanov, J. Yaniv, R. Mavlyutov, I. Goiri, S. Krishnan, J. Kulkarni *et al.*,

"Morpheus: Towards automated slos for enterprise clusters." in *OSDI*, 2016, pp. 117–134.

[63] A. D. Ferguson, P. Bodik, S. Kandula, E. Boutin, and R. Fonseca, "Jockey: guaranteed job latency in data parallel clusters," in *Proceedings of the 7th ACM european conference on Computer Systems*. ACM, 2012, pp. 99–112.

[64] A. V. Katherine and K. Alagarsamy, "Software testing in cloud platform: A survey," *International Journal of computer applications*, vol. 46, no. 6, pp. 21–25, 2012.

[65] S. Nachiyappan and S. Justus, "Cloud testing tools and its challenges: A comparative study," *Procedia Computer Science*, vol. 50, pp. 482–489, 2015.

[66] G. Carrozza, L. Battaglia, V. Manetti, A. Marotta, R. Canonico, and S. Avallone, "On the evaluation of vm provisioning time in cloud platforms for mission-critical infrastructures," in *Cluster, Cloud and Grid Computing (CCGrid), 2014 14th IEEE/ACM International Symposium on*. IEEE, 2014, pp. 802–810.

[67] "Stress," <https://people.seas.harvard.edu/apw/stress/>.

[68] "Google trace," <https://github.com/google/cluster-data>, 2011.

IX. APPENDIX

In this section, we present details for proving Theorem 3, 4, 5, 6, 7, 8. To prove Theorem 4. To prove Theorem 4, we first present Lemma 3.

Lemma 3. *Let A and B follow Pareto distribution with (a_{\min}, β) and (b_{\min}, β) , respectively. Given $a_{\min} \leq b_{\min}$, then*

$$f_B(b) = f_A(b|A > b_{\min}) \quad (32)$$

PROOF OF THEOREM 3

Proof. In this proof, we first derive the probability of a task completes before the deadline, and then derive the PoCD by considering all N tasks completing before the deadline.

We first compute the probability that an original attempt fails to finish before D . We denote the probability that an original attempt fails to finish before D as $P_{\text{Restart},o}$, which equals

$$P_{\text{Restart},o} = P(T_{j,1} > D) = \left(\frac{t_{\min}}{D}\right)^\beta. \quad (33)$$

If the original attempt's execution time is larger than D , r extra attempts are launched at τ_{est} . If an extra attempt fails to finish before D , it means that the execution time is more than $D - \tau_{\text{est}}$. We denote the probability that an extra attempt fails to finish before D as $P_{\text{Restart},e}$, which equals

$$P_{\text{Restart},e} = P(T_{j,a} > D - \tau_{\text{est}}) = \int_{D - \tau_{\text{est}}}^{\infty} \frac{\beta t_{\min}^\beta}{t^{\beta+1}} dt = \left(\frac{t_{\min}}{D - \tau_{\text{est}}}\right)^\beta. \quad (34)$$

A task fails to finish before D when the original attempt and r extra attempts launched at τ_{est} fail to finish before D . Thus, the probability that a task can finish before D is $1 - P_{\text{Restart},o} \cdot P_{\text{Restart},e}^r$. Also, a job can finish before the deadline D when all N tasks can finish before D . Thus, PoCD $R_{S-\text{Restart}}$ equals

$$(1 - P_{\text{Restart},o} \cdot (P_{\text{Restart},e})^r)^N = \left[1 - \frac{t_{\min}^{\beta \cdot (r+1)}}{D^\beta \cdot (D - \tau_{\text{est}})^{\beta \cdot r}}\right]^N. \quad (35)$$

□

PROOF OF THEOREM 4

Proof. In this proof, we first derive machine running time of a task by considering if the execution time of the original attempt is larger than D . If execution time is no more than D , there is no extra attempt launched. The machine running time is execution time of original attempt. If the execution time is larger than D , task machine time is the summation of machine running time of extra attempts killed at τ_{kill} and machine running time of the attempt successfully completed, and then get job's machine running time by adding machine running time of N tasks together.

$E_{S\text{-Restart}}(T)$ equals the expectation machine running time of all N tasks, i.e., $E_{S\text{-Restart}}(T) = N \cdot E(T_j)$, where T_j is the machine running time of task j . For task j , if the execution time of the original attempt exceeds D ($T_{j,1} > D$), extra attempts are launched at τ_{est} . Thus, we consider $E(T_j)$ in whether the execution time of the original attempt exceeds D or not, i.e., $T_{j,1} \leq D$, or $T_{j,1} > D$, and $E(T_j)$ equals

$$E(T_j|T_{j,1} \leq D) \cdot P(T_{j,1} \leq D) + E(T_j|T_{j,1} > D) \cdot P(T_{j,1} > D), \quad (36)$$

where

$$P(T_{j,1} > D) = 1 - P(T_{j,1} \leq D) = \left(\frac{t_{min}}{D}\right)^\beta. \quad (37)$$

Case 1. $T_{j,1} \leq D$.

The execution time of the original attempt is no more than D , and there is no extra attempt launched at time τ_{est} . Because

$$E(T_j|T_{j,1} \leq D) = \int_{t_{min}}^{\infty} t \cdot f_{T_j}(t|T_{j,1} \leq D) dt, \quad (38)$$

we first compute $f_{T_j}(t|T_{j,1} \leq D)$, which is

$$f_{T_j}(t|T_{j,1} \leq D) = \begin{cases} \frac{\beta \cdot t_{min}^\beta \cdot D^\beta}{(D^\beta - t_{min}^\beta) \cdot t^{\beta+1}}, & D > t > t_{min}; \\ 0, & \text{otherwise,} \end{cases} \quad (39)$$

By deriving Equ. (38), we can get

$$E(T_j|T_{j,1} \leq D) = \frac{t_{min} \cdot D \cdot \beta \cdot (t_{min}^{\beta-1} - D^{\beta-1})}{(1 - \beta) \cdot (D^\beta - t_{min}^\beta)}. \quad (40)$$

Case 2. $T_{j,1} > D$.

The execution time of the original attempt is more than D , and r extra attempts are launched at time τ_{est} . At time τ_{kill} , a attempt with the smallest finish time is left for running, and other r attempts are killed.

$E(T_j|T_{j,1} > D)$ consists of three parts, i.e., machine running time of the original attempt before τ_{est} , machine running time of r killed attempts between τ_{est} and τ_{kill} , and execution time of the attempt with the smallest finish time after τ_{est} . Since the original attempt starts τ_{est} earlier than extra attempts, $T_{j,1} - \tau_{est}$ is the execution time of the original attempt after τ_{est} , and $\min(T_{j,1} - \tau_{est}, T_{j,2}, \dots, T_{j,r+1})$ is the execution time of the attempt with the smallest finish time after τ_{est} . We denote

$\min(T_{j,1} - \tau_{est}, T_{j,2}, \dots, T_{j,r+1})$ as W_j^{all} . So, $E(T_j|T_{j,1} > D)$ equals

$$\tau_{est} + r \cdot (\tau_{kill} - \tau_{est}) + E(W_j^{all}|T_{j,1} > D). \quad (41)$$

Based on Lemma 3, we can transform $E(W_j^{all}|T_{j,1} > D)$ to $E(\widehat{W}_j^{all})$, where $\widehat{W}_j^{all} = \min(\widehat{T}_{j,1} - \tau_{est}, T_{j,2}, \dots, T_{j,r+1})$, and the minimum values of $\widehat{T}_{j,1}$ is D . Also, we denote $\min(T_{j,2}, \dots, T_{j,r+1})$ as W_j^{extra} .

We compute $E(\widehat{W}_j^{all})$ by using Lemma 1, i.e.,

$$E(\widehat{W}_j^{all}) = \int_{t_{min}}^{\infty} P(\widehat{T}_{j,1} - \tau_{est} \geq \omega) \cdot P^r(T_{j,a} \geq \omega) d\omega + t_{min}, \quad (42)$$

where

$$P(\widehat{T}_{j,1} - \tau_{est} \geq \omega) = \begin{cases} \left(\frac{D}{\omega + \tau_{est}}\right)^\beta, & \omega \geq D - \tau_{est}; \\ 1, & \omega < D - \tau_{est}, \end{cases} \quad (43)$$

and

$$P(T_{j,a} \geq \omega) = \begin{cases} \left(\frac{t_{min}}{\omega}\right)^\beta, & \omega \geq t_{min}; \\ 1, & \omega < t_{min}. \end{cases} \quad (44)$$

Because $D - \tau_{est}$ should be no less than t_{min} , otherwise there is no reason for launching extra attempts, thus

$$E(\widehat{W}_j^{all}) = \frac{t_{min}}{\beta \cdot r - 1} - \frac{t_{min}^{\beta \cdot r}}{(\beta \cdot r - 1) \cdot (D - \tau_{est})^{\beta \cdot r - 1}} + \int_{D - \tau_{est}}^{\infty} \left(\frac{D}{\omega + \tau_{est}}\right)^\beta \cdot \left(\frac{t_{min}}{\omega}\right)^{\beta \cdot r} d\omega + t_{min} \quad (45)$$

□

PROOF OF THEOREM 5

Proof. In this proof, we first derive the probability that a task completes before the deadline, and then derive the PoCD by considering all N tasks complete before the deadline.

We first compute the probability of an original attempt failed to finish before D . We denote the probability that an original attempt fails to finish before D as $P_{Resume,o}$ that

$$P_{Resume,o} = P(T_{j,1} > D) = \left(\frac{t_{min}}{D}\right)^\beta. \quad (46)$$

If the original attempt's execution time is larger than D , the original attempt is killed, and $r+1$ extra attempts are launched at τ_{est} . Suppose the original attempt processed $\varphi_{j,est}$ fraction of data, extra attempts continue to process the remaining $1 - \varphi_{j,est}$ fraction of data. An extra attempt failed to finish before D means the execution time is more than $D - \tau_{est}$. We denote the probability that an extra attempt fails to finish before D as $P_{Resume,e}$ that

$$P_{Resume,e} = P((1 - \varphi_{j,est}) \cdot T_{j,a} > D - \tau_{est}) = \left[\frac{(1 - \varphi_{j,est}) \cdot t_{min}}{D - \tau_{est}}\right]^\beta. \quad (47)$$

A task fails to finish before D when the original attempt and the r extra attempts launched at τ_{est} fail to finish before D . Thus, the probability that a task finishes before D is $1 -$

$P_{\text{Resume},o} \cdot P_{\text{Resume},e}^{r+1}$. Also, the job finishes before the deadline D when all N tasks finish before D . Thus, PoCD $R_{S-\text{Resume}}$ equals

$$R_{S-\text{Resume}} = (1 - P_{\text{resume},o} \cdot (P_{\text{resume},e})^{r+1})^N = \left[1 - \frac{(1 - \varphi_{j,\text{est}})^{\beta \cdot (r+1)} \cdot t_{\min}^{\beta \cdot (r+2)}}{D^\beta \cdot (D - \tau_{\text{est}})^{\beta \cdot (r+1)}} \right]^N. \quad (48)$$

PROOF OF THEOREM 6

Proof. In this proof, we first derive machine running time of a task by considering if the execution time of the original attempt is larger than D . If execution time is no more than D , there is no extra attempt launched. The machine running time is execution time of original attempt. If the execution time is larger than D , task machine time is the summation of machine running time of extra attempts killed at τ_{kill} and machine running time of the attempt successfully completed, and then get job's machine running time by adding machine running time of N tasks together.

$E_{S-\text{Resume}}(T)$ equals the expectation machine running time of all N tasks, i.e., $E_{S-\text{Resume}}(T) = N \cdot E(T_j)$, where T_j is the machine running time of task j . For task j , if the execution time of the original attempt exceeds D ($T_{j,1} > D$), extra attempts are launched at τ_{est} . Thus, we consider $E(T_j)$ in whether the execution time of the original attempt exceeds D or not, i.e., $T_{j,1} \leq D$, or $T_{j,1} > D$, and $E(T_j)$ equals

$$E(T_j | T_{j,1} \leq D) \cdot P(T_{j,1} \leq D) + E(T_j | T_{j,1} > D) \cdot P(T_{j,1} > D), \quad (49)$$

where

$$P(T_{j,1} > D) = 1 - P(T_{j,1} \leq D) = \left(\frac{t_{\min}}{D} \right)^\beta. \quad (50)$$

Case 1. $T_{j,1} \leq D$.

The execution time of the original attempt is no more than D , and there is no extra attempt launched at time τ_{est} . Because

$$E(T_j | T_{j,1} \leq D) = \int_{t_{\min}}^{\infty} t \cdot f_{T_j}(t | T_{j,1} \leq D) dt, \quad (51)$$

we first compute $f_{T_j}(t | T_{j,1} \leq D)$, which is

$$f_{T_j}(t | T_{j,1} \leq D) = \begin{cases} \frac{\beta \cdot t_{\min}^\beta \cdot D^\beta}{(D^\beta - t_{\min}^\beta) \cdot t^{\beta+1}}, & D > t > t_{\min}; \\ 0, & \text{otherwise,} \end{cases} \quad (52)$$

By deriving Equ. (51), we can get

$$E(T_j | T_{j,1} \leq D) = \frac{t_{\min} \cdot D \cdot \beta \cdot (t_{\min}^{\beta-1} - D^{\beta-1})}{(1 - \beta) \cdot (D^\beta - t_{\min}^\beta)}. \quad (53)$$

Case 2 $T_{j,1} > D$

The execution time of the original attempt is more than D , and the original attempt is killed at τ_{est} . $r + 1$ extra attempts are launched at τ_{est} , and continue to process remaining

$(1 - \varphi_{j,\text{est}})$ data. At τ_{kill} , r extra attempts are killed and leave an extra attempt with the smallest finishing time running. $E(T_j | T_{j,1} > D)$ consists of three parts, i.e., machine running time of the original attempt before τ_{est} , machine running time of killed extra attempts between τ_{est} and τ_{kill} , and execution time of the extra attempt with the smallest finish time. So,

$$E(T_j | T_{j,1} > D) = \tau_{\text{est}} + r \cdot (\tau_{\text{kill}} - \tau_{\text{est}}) + E(W_j^{\text{new}}), \quad (54)$$

□ where $W_j^{\text{new}} = \min(T_{j,2}, \dots, T_{j,r+1})$, and $T_{j,a} = (1 - \varphi_{j,\text{est}}) \cdot T_{j,1}$, $\forall a \in 2, \dots, r + 1$. Based on Lemma 1,

$$E(W_j^{\text{new}}) = \int_{t_{\min}}^{\infty} P^{r+1}((1 - \varphi_{j,\text{est}}) \cdot T_{j,1} > t) dt + t_{\min} \quad (55)$$

$$= \frac{t_{\min} \cdot (1 - \varphi_{j,\text{est}})^{\beta \cdot (r+1)}}{\beta \cdot (r + 1) - 1} + t_{\min}. \quad (56)$$

□

PROOF OF THEOREM 7

Proof. Suppose function $f(x)$ equals $(1 - x)^N$, where N is an integer. $f(x)$ is a monotonic non-increasing function, i.e., $f(x_1) \geq f(x_2)$, if and only if $x_1 \leq x_2$.

In following, we compare among $(1 - R_{\text{clone}})^{1/N}$, $(1 - R_{S-\text{Restart}})^{1/N}$, and $(1 - R_{S-\text{Resume}})^{1/N}$.

1) *PoCD comparison between Clone and Speculative-Restart:*

$$\frac{(1 - R_{\text{clone}})^{1/N}}{(1 - R_{S-\text{Restart}})^{1/N}} = \left(\frac{D - \tau_{\text{est}}}{D} \right)^{\beta \cdot r} < 1 \quad (57)$$

Thus, $R_{\text{clone}} > R_{S-\text{Restart}}$.

2) *PoCD comparison between Speculative-Restart and Speculative-Resume:*

$$\frac{(1 - R_{S-\text{Restart}})^{1/N}}{(1 - R_{S-\text{Resume}})^{1/N}} = \frac{(D - \tau_{\text{est}})^\beta}{t_{\min}^\beta \cdot (1 - \varphi_{j,\text{est}})^{\beta \cdot (r+1)}} \quad (58)$$

Given $D - \tau_{\text{est}} \geq t_{\min} \cdot (1 - \varphi_{j,\text{est}})$, we can get $R_{S-\text{Restart}} < R_{S-\text{Resume}}$.

3) *PoCD comparison between Clone and Speculative-Resume:*

$$\frac{(1 - R_{\text{clone}})^{1/N}}{(1 - R_{S-\text{Resume}})^{1/N}} = \frac{(D - \tau_{\text{est}})^{\beta \cdot (r+1)}}{(1 - \varphi_{j,\text{est}})^{\beta \cdot (r+1)} \cdot D^{\beta \cdot r} \cdot t_{\min}^\beta} \quad (59)$$

Given the original attempt misses the deadline, then $D - \tau_{\text{est}} < (1 - \varphi_{j,\text{est}}) \cdot D$. Thus, if

$$r > \log_{\frac{D - \tau_{\text{est}}}{(1 - \varphi_{j,\text{est}}) \cdot D}} \frac{(1 - \varphi_{j,\text{est}})^{\beta \cdot t_{\min}^\beta}}{D - \tau_{\text{est}}}, \quad (60)$$

then $R_{\text{clone}} > R_{S-\text{Resume}}$. Otherwise, $R_{\text{clone}} \leq R_{S-\text{Resume}}$. □

PROOF OF THEOREM 8

Proof. In following, we prove $\lg(R_{\text{Clone}}(r) - R_{\min})$ is a concave function when $r > \Gamma_{\text{Clone}} = -\beta^{-1} \cdot \log_{t_{\min}/D} N - 1$, and $-\theta \cdot C \cdot E_{\text{Clone}}(T)$ is a concave function for all values of r .

$\lg(R_{\text{Clone}}(r) - R_{\min})$ is an increasing and concave function of $R_{\text{Clone}}(r)$. Also, when $r > \Gamma_{\text{Clone}}$, the second order derivative of R_{Clone} is less than 0. So, R_{Clone} is a concave function of r , when $r > \Gamma_{\text{Clone}}$. Based on *Lemma 2*, we know that $\lg(R_{\text{Clone}} - R_{\min})$ is a concave function of r , when $r > \Gamma_{\text{Clone}}$.

The second order derivative of $-E_{\text{Clone}}(T)$ equals

$$-\frac{2 \cdot N \cdot t_{\min}}{[\beta \cdot (r + 1) - 1]^3}, \quad (61)$$

which is less than 0. Given the summation of two concave functions is also a concave function, thus $U_{\text{Clone}}(r)$ is a concave function when $r > \Gamma_{\text{Clone}}$. \square

Proof. In following, we prove $\lg(R_{\text{S-Restart}} - R_{\min})$ is a concave function when

$$r > \Gamma_{\text{S-Restart}} = \beta^{-1} \cdot \log_{t_{\min}/(D-\tau_{\text{est}})} \frac{D^\beta}{N \cdot t_{\min}}, \quad (62)$$

and $-\theta \cdot C \cdot E_{\text{S-Restart}}(T)$ is a concave function for all values of r .

$\lg(R_{\text{S-Restart}}(r) - R_{\min})$ is an increasing and concave function of $R_{\text{S-Restart}}(r)$. Also, when $r > \Gamma_{\text{S-Restart}}$, the second order derivative of $R_{\text{S-Restart}}$ is less than 0. So, $R_{\text{S-Restart}}$ is a concave function of r , when $r > \Gamma_{\text{S-Restart}}$. Based on *Lemma 2*, we know that $\lg(R_{\text{Clone}} - R_{\min})$ is a concave function of r , when $r > \Gamma_{\text{S-Restart}}$.

The second order derivative of $-E_{\text{S-Restart}}(T)$ is

$$\begin{aligned} & - \int_{D-\tau_{\text{est}}}^{t_{\min}} \left(\frac{t_{\min}}{D} \right)^\beta \left(\ln \frac{t_{\min}}{\omega} \right)^2 \cdot \beta^2 \cdot \left(\frac{t_{\min}}{\omega} \right)^{\beta \cdot r} d\omega \\ & - \int_{D-\tau_{\text{est}}}^{\infty} \left(\ln \frac{t_{\min}}{\omega} \right)^2 \cdot \beta^2 \cdot \left(\frac{t_{\min}}{\omega + \tau_{\text{est}}} \right)^\beta \left(\frac{t_{\min}}{\omega} \right)^{\beta \cdot r} d\omega, \end{aligned} \quad (63)$$

which is less than 0. Given the summation of two concave functions is also a concave function, thus $U_{\text{S-Restart}}(r)$ is a concave function when $r > \Gamma_{\text{S-Restart}}$. \square

Proof. In following, we prove that $\lg(R_{\text{S-Resume}}(r) - R_{\min})$ is concave when

$$r > \Gamma_{\text{S-Resume}} = \beta^{-1} \cdot \log_{\frac{\bar{t}_{\min}}{D-\tau_{\text{est}}}} \frac{D^\beta}{N \cdot t_{\min}^\beta} - 1, \quad (64)$$

and $-\theta \cdot C \cdot E_{\text{S-Resume}}(T)$ is concave for all values of r .

First, we note that $\lg(R_{\text{S-Resume}}(r) - R_{\min})$ is an increasing and concave function of $R_{\text{S-Resume}}(r)$. Also, when $r > \Gamma_{\text{S-Resume}}$, the second order derivative of $R_{\text{S-Resume}}$ is less than 0. So, $R_{\text{S-Resume}}$ is a concave function of r , when $r > \Gamma_{\text{Clone}}$. Based on *Lemma 2*, we know that $\lg(R_{\text{S-Resume}} - R_{\min})$ is a concave function of r , when $r > \Gamma_{\text{S-Resume}}$.

The second order derivative of $-E_{\text{S-Resume}}(T)$ equals

$$- \int_{t_{\min}}^{\infty} \left(\frac{t_{\min}}{D} \right)^\beta \left(\frac{\bar{t}_{\min}}{t} \right)^{\beta \cdot (r+1)} \cdot \left(\ln \frac{\bar{t}_{\min}}{t} \right)^2 \cdot \beta^2 dt \quad (65)$$

which is less than 0. Given the summation of two concave functions is also a concave function, thus $U_{\text{S-Resume}}(r)$ is a concave function when $r > \Gamma_{\text{S-Resume}}$. \square