HAL
open science

Checkpointing strategies for parallel jobs

Marin Bougeret, Henri Casanova, Mikael Rabie, Yves Robert, Frédéric Vivien

To cite this version:

Marin Bougeret, Henri Casanova, Mikael Rabie, Yves Robert, Frédéric Vivien. Checkpointing strategies for parallel jobs. [Research Report] RR-7520, INRIA. 2011, pp.45. inria-00560582v3

HAL Id: inria-00560582

https://inria.hal.science/inria-00560582v3

Submitted on 22 Apr 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Checkpointing strategies for parallel jobs

Marin Bougeret — Henri Casanova — Mikael Rabie — Yves Robert — Frédéric Vivien

N° 7520

April 2011

Distributed and High Performance Computing

Rapport de recherche

ISSN 0249-6399 ISRN INRIA/RR--7520--FR+ENG

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

INRIA

centre de recherche GRENOBLE - RHÔNE-ALPES

Checkpointing strategies for parallel jobs

Marin Bougeret[*] , Henri Casanova[†] , Mikael Rabie[*] , Yves Robert[*‡] , Frédéric Vivien[§]

Theme : Distributed and High Performance Computing
Équipe-Projet GRAAL

Rapport de recherche n° 7520 — April 2011 — 53 pages

Abstract: This work provides an analysis of checkpointing strategies for minimizing expected job execution times in an environment that is subject to processor failures. In the case of both sequential and parallel jobs, we give the optimal solution for exponentially distributed failure inter-arrival times, which, to the best of our knowledge, is the first rigorous proof that periodic checkpointing is optimal. For non-exponentially distributed failures, we develop a dynamic programming algorithm to maximize the amount of work completed before the next failure, which provides a good heuristic for minimizing the expected execution time. Our work considers various models of job parallelism and of parallel checkpointing overhead. We first perform extensive simulation experiments assuming that failures follow Exponential or Weibull distributions, the latter being more representative of real-world systems. The obtained results not only corroborate our theoretical findings, but also show that our dynamic programming algorithm significantly outperforms previously proposed solutions in the case of Weibull failures. We then discuss results from simulation experiments that use failure logs from production clusters. These results confirm that our dynamic programming algorithm significantly outperforms existing solutions for real-world clusters.

Key-words: Fault-tolerance, checkpointing, sequential job, parallel job, Weibull

[*] LIP, Ecole Normale Supérieure de Lyon
[†] Univ. of Hawai'i at Mānoa, Honolulu, USA
[‡] Yves Robert is with the Institut Universitaire de France. This work was supported in part by the ANR StochaGrid and RESCUE projects, and by the INRIA-Illinois Joint Laboratory for Petascale Computing.
[§] INRIA, Lyon, France

Centre de recherche INRIA Grenoble – Rhône-Alpes
655, avenue de l'Europe, 38334 Montbonnot Saint Ismier
Téléphone : +33 4 76 61 52 00 — Télécopie +33 4 76 61 52 52

Stratégies de checkpoint pour applications parallèles

Résumé : Nous présentons dans ce travail une analyse rigoureuse des stratégies de checkpoint, pour les applications séquentielles et parallèles. L'objectif est de minimiser l'espérance du temps de complétion d'une application s'exécutant sur des processeurs pouvant être de victimes de pannes. Dans le cas séquentiel, une résolution exacte est proposée lorsque les intervalles inter-pannes sont distribués selon une loi exponentielle. Il semble que ce résultat soit la première preuve rigoureuse de l'optimalité des stratégies périodiques de checkpoint dans ce cadre. Dans le cas général (c'est-à-dire pour des lois quelconques), nous fournissons une programmation dynamique permettant des calculs optimaux, à un quantum de temps fixé près. Dans le cas des applications parallèles, nous étendons le résultat exact (pour le cas exponentiel), et ce pour différents modèles d'applications et de coûts de checkpoints. Pour le cas général, nous proposons une deuxième programmation dynamique (plus rapide) dont l'objectif est de maximiser l'espérance du travail fait avant la prochaine panne, qui s'avère fournir de bonnes approximations pour le problème initial. Nous validons nos résultats grâce à de nombreuses simulations, réalisées pour des lois inter-pannes de distribution exponentielles et de Weibull (cette dernière distribution étant selon de nombreuses études plus appropriée pour modéliser des durées inter-pannes). Ces simulations confirment nos résultats théoriques. De plus, ils apparaît que notre algorithme de programmation dynamique fournit de bien meilleures performances que les solutions existantes pour le cas des lois de Weibull.

Mots-clés : Tolérance aux pannes, checkpoint, tâche séquentielle, tâche parallèle, Weibull

1 Introduction

Resilience is a key challenge for post-petascale high-performance computing (HPC) systems [9, 21] since failures are increasingly likely to occur during the execution of parallel jobs that enroll increasingly large numbers of processors. For instance, the 45,208-processor Jaguar platform is reported to experience on the order of 1 failure per day [18, 2]. Faults that cannot be automatically detected and corrected in hardware lead to failures. In this case, rollback recovery is used to resume job execution from a previously saved fault-free execution state, or *checkpoint*. Rollback recovery implies frequent (usually periodic) *checkpointing* events at which the job state is saved to resilient storage. More frequent checkpoints lead to higher overhead during fault-free execution, but less frequent checkpoints lead to a larger loss when a failure occurs. The design of efficient *checkpointing strategies*, which specify when checkpoints should be taken, is thus key to high performance.

We study the problem of finding a checkpointing strategy that minimizes the expectation of a job's execution time, or *expected makespan*. In this context, our novel contributions are as follows. For sequential jobs, we provide the optimal solution for exponential failures and an accurate dynamic programming algorithm for general failures. The optimal solution for exponential failures, i.e., periodic checkpointing, is widely known in the "folklore" but, to the best of our knowledge, we provide the first rigorous proof. Our dynamic programming algorithm provides the first accurate solution of the expected makespan minimization problem with Weibull failures, which are representative of the behavior of real-world platforms [10, 22, 17]. For parallel jobs, we consider a variety of execution scenarios with different models of job parallelism (embarrassingly parallel jobs, jobs that follow Amdahl's law, and typical numerical kernels such as matrix product or LU decomposition), and with different models of the overhead of checkpointing a parallel job (which may or may not depend on the total number of processors in use). In the case of Exponential failures we provide the optimal solution. In the case of general failures, since minimizing the expected makespan is computationally difficult, we instead provide a dynamic programming algorithm to maximize the amount of work successfully completed before the next failure. This approach turns out to provide a good heuristic solution to the expected makespan minimization problem. In particular, it significantly outperforms previously proposed solutions in the case of Weibull failures.

Sections 2 and 3 give theoretical results for sequential and parallel jobs, respectively. Section 4 presents our simulation methodology. Sections 5 and 6 discuss simulation results when using synthetic failure distributions and when using real-world failure data, respectively. Section 7 reviews related work. Finally, Section 8 concludes the paper with a summary of findings and a discussion of future directions.

2 Sequential jobs

2.1 Problem statement

We consider an application, or *job*, that executes on one *processor*. We use the term processor to indicate any individually scheduled compute resource (a core,

a multi-core processor, a cluster node) so that our work is agnostic to the granularity of the platform. The job must complete $\mathcal{W}$ units of (divisible) work, which can be split arbitrarily into separate *chunks*. The job state is checkpointed after the execution of every chunk. Defining the sequence of chunk sizes is therefore equivalent to defining the checkpointing dates. We use C to denote the time needed to perform a checkpoint. The processor is subject to *failures*, each causing a *downtime* period, of duration D, followed by a *recovery* period, of duration R. The downtime accounts for software rejuvenation (i.e., rebooting [13, 7]) or for the replacement of the failed processor by a spare. Regardless, we assume that after a downtime the processor is fault-free and begins a new lifetime at the beginning of the recovery period. This period corresponds to the time needed to restore the last checkpoint. Finally, we assume that failures can happen during recovery or checkpointing, but not during a downtime (otherwise, the downtime period could be considered part of the recovery period).

We study two optimization problems:

- MAKESPAN: Minimize the job's expected makespan;
- NEXTFAILURE: Maximize the expected amount of work completed before the next failure.

Solving MAKESPAN is our main goal. NEXTFAILURE amounts to optimizing the makespan on a "failure-by-failure" basis, selecting the next chunk size as if the next failure were to imply termination of the execution. Intuitively, solving NEXTFAILURE should lead to a good approximation of the solution to MAKESPAN, at least for large job sizes $\mathcal{W}$. Therefore, we use the solution of NEXTFAILURE in cases for which we are unable to solve MAKESPAN directly. We give formal definitions for both problems in the next section.

2.2 Formal problem definitions

We consider the processor from time t_0 onward. We do *not* assume that the failure stochastic process is memoryless. Failures occur at times $(t_n)_{n\geq 1}$, with $t_n = t_0 + \sum_{m=1}^{n} X_m$, where the random variables $(X_m)_{m\geq 1}$ are *iid* (independent and identically distributed). Given a current time $t > t_0$, we define $n(t) = \min\{n | t_n \geq t\}$, so that $X_{n(t)}$ corresponds to the inter-failure interval in which t falls. We use $P_{suc}(x|\tau)$ to denote the probability that the processor does not fail for the next x units of time, knowing that the last failure occurred τ units of time ago. In other words, if $X = X_{n(t)}$ denotes the current inter-arrival failure interval,

$$P_{suc}(x|\tau) = \mathbb{P}(X \geq \tau + x \mid X \geq \tau) \ .$$

For both problems stated in the previous section, a solution is fully defined by a function $f(\omega|\tau)$ that returns the size of the next chunk to execute given the amount of work ω that has not yet been executed successfully ($f(\omega|\tau) \leq \omega \leq \mathcal{W}$) and the amount of time τ elapsed since the last failure. f is invoked at each decision point, i.e., after each checkpoint or recovery. Our goal is to determine a function f that optimally solves the considered problem. Assuming a unit-speed processor without loss of generality, the time needed to execute a chunk of size ω is $\omega + C$ if no failure occurs.

Definition of MAKESPAN– For a given amount of work ω, a time elapsed since the last failure τ, and a function f, let $\omega_1 = f(\omega|\tau)$ denote the size of the first chunk, and let $T(\omega|\tau)$ be the random variable that quantifies the time needed

for successfully executing ω units of work. We can write the following recursion:

$$\begin{aligned} T(0|\tau) &= 0 \\ T(\omega|\tau) &= \\ &\begin{cases} \omega_1 + C + T(\omega - \omega_1 | \tau + \omega_1 + C) \\ \quad \text{if the processor does not fail during} \\ \quad \text{the next } \omega_1 + C \text{ units of time,} \\ T_{wasted}(\omega_1 + C|\tau) + T(\omega|R) \\ \quad \text{otherwise.} \end{cases} \end{aligned} \tag{1}$$

The two cases above are explained as follows:

- If the processor does not fail during the execution and checkpointing of the first chunk (i.e., for $\omega_1 + C$ time units), there remains to execute a work of size $\omega - \omega_1$ and the time since the last failure is $\tau + \omega_1 + C$;
- If the processor fails before successfully completing the first chunk and its checkpoint, then some additional delays are incurred, as captured by the variable $T_{wasted}(\omega_1 + C|\tau)$. The time wasted corresponds to the execution up to the failure, a downtime, and a recovery during which a failure may happen. We compute T_{wasted} in the next section. Regardless, once a successful recovery has been completed, there still remain ω units of work to execute, and the time since the last failure is simply R.

We define MAKESPAN formally as: find f that minimizes $\mathbb{E}(T(\mathcal{W}|\tau_0))$, where $\mathbb{E}(X)$ denotes the expectation of the random variable X, and τ_0 the time elapsed since the last failure before t_0.

Definition of NEXTFAILURE– For a given amount of work ω, a time elapsed since the last failure τ, and a function f, let $\omega_1 = f(\omega|\tau)$ denote the size of the first chunk, and let $W(\omega|\tau)$ be the random variable that quantifies the amount of work successfully executed before the next failure. We can write the following recursion:

$$\begin{aligned} W(0|\tau) &= 0 \\ W(\omega|\tau) &= \\ &\begin{cases} \omega_1 + W(\omega - \omega_1 | \tau + \omega_1 + C) \\ \quad \text{if the processor does not fail during} \\ \quad \text{the next } \omega_1 + C \text{ units of time,} \\ 0 \quad \text{otherwise.} \end{cases} \end{aligned} \tag{2}$$

This recursion is simpler than the one for MAKESPAN because a failure during the computation of the first chunk means that no work (i.e., no fraction of ω) will have been successfully executed before the next failure. We define NEXTFAILURE formally as: find f that maximizes $\mathbb{E}(W(\mathcal{W}|\tau_0))$.

2.3 Solving Makespan

A challenge for solving MAKESPAN is the computation of $T_{wasted}(\omega_1 + C|\tau)$. We rely on the following decomposition:

$$T_{wasted}(\omega_1 + C|\tau) = T_{lost}(\omega_1 + C|\tau) + T_{rec}\ , \qquad \text{where}$$

- $T_{lost}(x|\tau)$ is the amount of time spent computing before a failure, knowing that the next failure occurs within the next x units of time, and that the last failure has occurred τ units of time ago.

- T_{rec} is the amount of time needed by the system to recover from the failure (accounting for the fact that other failures may occur during recovery).

Proposition 1. *The* MAKESPAN *problem is equivalent to finding a function f minimizing the following quantity:*

$$\begin{aligned}&\mathbb{E}(T(\mathcal{W}|\tau)) = \\ &P_{suc}(\omega_1 + C|\tau)\Big(\omega_1 + C + \mathbb{E}(T(\mathcal{W} - \omega_1|\tau + \omega_1 + C))\Big) \\ &+(1 - P_{suc}(\omega_1 + C|\tau))\Big(\mathbb{E}(T_{lost}(\omega_1 + C|\tau)) \\ &\qquad + \mathbb{E}(T_{rec}) + \mathbb{E}(T(\mathcal{W}|R))\Big)\end{aligned} \tag{3}$$

where $\omega_1 = f(\mathcal{W}|\tau)$ and where $\mathbb{E}(T_{rec})$ is given by

$$\mathbb{E}(T_{rec}) = D + R + \frac{1 - P_{suc}(R|0)}{P_{suc}(R|0)}(D + \mathbb{E}(T_{lost}(R|0))).$$

Proof. Recall that $P_{suc}(x|\tau)$ denotes the probability of successfully computing during x time units, knowing that the last failure occurred τ units of time ago. Based on the recursion for T given in Equation 1, simply weighting the two cases by their probabilities of occurrence leads to Equation 3.

We can write the following recursion for T_{rec}:

$$T_{rec} = \begin{cases} D + R & \text{with probability } P_{suc}(R|0), \\ D + T_{lost}(R|0) + T_{rec} & \\ \qquad \text{with probability } 1 - P_{suc}(R, 0). \end{cases}$$

and compute $\mathbb{E}(T_{rec})$ in terms of $\mathbb{E}(T_{lost}(R|0))$ by weighting the two cases by their probabilities of occurrence. □

2.3.1 Results for the Exponential distribution

In this section we assume that the failure inter-arrival times follow an Exponential distribution with parameter λ, i.e., each $X_n = X$ has probability density $f_X(t) = \lambda e^{-\lambda t} dt$ and cumulative distribution $F_X(t) = 1 - e^{-\lambda t}$ for all $t \geq 0$. The advantage of the Exponential distribution, exploited time and again in the literature, is its "memoryless" property: the time at which the next failure occurs does not depend on the time elapsed since the last failure occurred. Therefore, in this section, we simply write $T(\omega)$, $T_{lost}(\omega)$, and $P_{suc}(\omega)$ instead of $T(\omega|\tau)$, $T_{lost}(\omega|\tau)$, and $P_{suc}(\omega|\tau)$.

Lemma 1. *With the Exponential distribution:*

$$\mathbb{E}(T_{lost}(\omega)) = \frac{1}{\lambda} - \frac{\omega}{e^{\lambda\omega} - 1} \qquad \textit{and}$$
$$\mathbb{E}(T_{rec}) = D + R + \frac{1 - e^{-\lambda R}}{e^{-\lambda R}}(D + \mathbb{E}(T_{lost}(R))).$$

Proof.

$$\begin{aligned}\mathbb{E}(T_{lost}(\omega)) &= \textstyle\int_0^\infty x\, \mathbb{P}(X = x|X < \omega)dx \\ &= \tfrac{1}{\mathbb{P}(X<\omega)} \textstyle\int_0^\omega x f_X(x)dx \\ &= \tfrac{1}{1-e^{-\lambda\omega}}([-xe^{-\lambda x}]_0^\omega + \textstyle\int_0^\omega e^{-\lambda x})\end{aligned}$$

The formula for $\mathbb{E}(T_{rec})$ is obtained directly from Proposition 1 by replacing $P_{suc}(R|0) = P_{suc}(R)$ by $e^{-\lambda R}$. □

The memoryless property makes it possible to solve the MAKESPAN problem analytically:

Theorem 1. *Let $\mathcal{W}$ be the amount of work to execute on a processor with failure inter-arrival times that follow an Exponential distribution with parameter λ. Let $K_0 = \frac{\lambda \mathcal{W}}{1+\mathbb{L}(-e^{-\lambda C-1})}$ where $\mathbb{L}$, the Lambert function, is defined as $\mathbb{L}(z)e^{\mathbb{L}(z)} = z$. Then the optimal strategy to minimize the expected makespan is to split $\mathcal{W}$ into $K^* = \max(1, \lfloor K_0 \rfloor)$ or $K^* = \lceil K_0 \rceil$ same-size chunks, whichever leads to the smaller value. The optimal expectation of the makespan is:*

$$\mathbb{E}(T^*(\mathcal{W})) = K^* \left(e^{\lambda R} \left(\frac{1}{\lambda} + D \right) \right) \left(e^{\lambda(\frac{\mathcal{W}}{K^*}+C)} - 1 \right).$$

Proof. We first observe that all possible executions for a given deterministic strategy f use the same sequence of chunk sizes. The only difference between two different executions is the number of times each chunk is tentatively executed before its successful completion. This is because the size of the first chunk, $\omega_1 = f(w|\tau) = f(\omega)$, does not depend upon τ (due to the memoryless property). Either its first execution is successful or it fails. If it fails, then the optimal solution consists in retrying a chunk of same size ω_1 since there remains ω units of work to be executed and ω_1 does not depend on τ. Once a chunk of size ω_1 has been successfully completed, the next attempted chunk is of size $\omega_2 = f(\mathcal{W} - \omega_1|\tau) = f(\mathcal{W} - \omega_1)$, which does not depend on τ, and for which the same reasoning holds.

Given the above, all executions can be described as sequences of the form $\omega_1^{(\ell_1)}\omega_2^{(\ell_2)} \dots \omega_k^{(\ell_k)} \dots$, where $\omega^{(\ell)}$ means that a chunk of size ω was tentatively executed ℓ times, the first $\ell - 1$ tentatives being unsuccessful and the last one being successful. Because each chunk size is successfully executed at least once, the time to execute it is always bounded below by C, the time to take a checkpoint. Say that the optimal strategy uses K successive chunk sizes, with K to be determined. Any execution following this optimal strategy will have a makespan at least as large as KC. Hence $\mathbb{E}(T^*(\mathcal{W}))$, the expectation of the makespan with the optimal strategy, is also greater than or equal to KC.

We first prove that, as might be expected, K is finite. Let us consider a simple, non-optimal, strategy, defined by $f(\omega) = \omega$. In other words, this strategy executes the whole work $\mathcal{W}$ as a single chunk, repeating its execution until it succeeds. Let $\mathbb{E}(T^{id}(\mathcal{W}))$ denote the expected makespan when this strategy is used. Using Lemma 1, we have:

$$\begin{aligned}
&\mathbb{E}(T^{id}(\mathcal{W})) = \\
&\quad P_{suc}(\mathcal{W}+C)(\mathcal{W}+C) \\
&\quad +(1 - P_{suc}(\mathcal{W}+C)) \left(\mathbb{E}(T_{lost}(\mathcal{W}+C)) + \mathbb{E}(T_{rec}) + \mathbb{E}(T^{id}(\mathcal{W})) \right).
\end{aligned}$$

Given that $P_{suc}(\mathcal{W}+C) = e^{-\lambda \mathcal{W}+C}$, we obtain:

$$\mathbb{E}(T^{id}(\mathcal{W})) = \mathcal{W} + C + \left(\mathbb{E}(T_{lost}(\mathcal{W}+C)) + \mathbb{E}(T_{rec}) \right) \frac{1 - e^{-\lambda \mathcal{W}+C}}{e^{-\lambda \mathcal{W}+C}}$$

This last expression shows that $\mathbb{E}(T^{id}(\mathcal{W}))$ is finite, implying that $\mathbb{E}(T^*(\mathcal{W}))$, the expected makespan for the optimal strategy, is also finite. Since it is bounded below by KC, we conclude that K is finite, meaning that the optimal solution uses a bounded number of chunk sizes. In the optimal solution, the i-th chunk size, $1 \le i \le K$, is $\omega_i = f^*(\mathcal{W} - \sum_{j=1}^{i-1} \omega_j)$. From Proposition 1, we derive that:

$$\begin{aligned}&\mathbb{E}(T^*(\mathcal{W})) = \\ &P_{suc}(\omega_1 + C)\,(\omega_1 + C + \mathbb{E}(T^*(\mathcal{W} - \omega_1)) \\ &+(1 - P_{suc}(\omega_1 + C))\,(\mathbb{E}(T_{lost}(\omega_1 + C)) + \mathbb{E}(T_{rec}) + \mathbb{E}(T^*(\mathcal{W}))\,.\end{aligned}$$

To shorten notations, let us define $\rho^* = \mathbb{E}(T^*(\mathcal{W}))$. Then:

$$\rho^* = \omega_1 + C + \mathbb{E}(T^*(\mathcal{W} - \omega_1)) + \frac{1 - P_{suc}(\omega_1 + C)}{P_{suc}(\omega_1 + C)}(\mathbb{E}(T_{lost}(\omega_1 + C)) + \mathbb{E}(T_{rec})).$$

We have $\frac{1-P_{suc}(\omega_1+C)}{P_{suc}(\omega_1+C)} = e^{\lambda(\omega_1+C)} - 1$ and, using Lemma 1, developing the expression for ρ^* leads to

$$\rho^* = \sum_{i=1}^{K}(\omega_i + C) + \sum_{i=1}^{K}\left(\left(\frac{1}{\lambda} - \frac{\omega_i + C}{e^{\lambda(\omega_i+C)} - 1} + \mathbb{E}(T_{rec})\right)(e^{\lambda(\omega_i+C)} - 1)\right)$$

Terms cancel out and the above simplifies into:

$$\rho^* = \left(\frac{1}{\lambda} + \mathbb{E}(T_{rec})\right)\sum_{i=1}^{K}(e^{\lambda(\omega_i+C)} - 1).$$

Since $e^{\lambda(\omega_i+C)}$ is a convex function of ω_i, ρ^* is minimized when all the chunks ω_i have the same size $\frac{\mathcal{W}}{K}$, in which case $\rho^* = K(\frac{1}{\lambda} + \mathbb{E}(T_{rec}))(e^{\lambda(\frac{\mathcal{W}}{K}+C)} - 1)$. We look for the value of K that minimizes $\psi(K) = K(e^{\lambda(\frac{\mathcal{W}}{K}+C)} - 1)$. We call K_0 this value and, differentiating, we must solve $\psi'(K_0) = 0$ where

$$\psi'(K_0) = e^{\lambda(\frac{\mathcal{W}}{K_0}+C)}\left(1 - \frac{\lambda\mathcal{W}}{K_0}\right) - 1. \tag{4}$$

This equation can be rewritten as $ye^y = -e^{-\lambda C-1}$, where $y = \frac{\lambda\mathcal{W}}{K_0} - 1$. The only solution is $y = \mathbb{L}(-e^{-\lambda C-1})$, where $\mathbb{L}$ is the Lambert function. Differentiating again, we easily see that ψ'' is always non-negative. The optimal value is thus obtained by one of the two integers surrounding the zero of ψ', which proves the theorem. □

Remark 1. *Although periodic checkpoints have been widely used in the literature, Theorem 1 is, to the best of our knowledge, the first proof that the optimal deterministic strategy uses a finite number of chunks and is periodic. In addition, as a consequence of Proposition 4.4.3 in [20], this strategy can be shown to be optimal among all deterministic and non-deterministic strategies.*

2.3.2 Results for arbitrary distributions

Solving the Makespan problem for arbitrary distributions is difficult because, unlike in the memoryless case, there is no reason for the optimal solution to use a single chunk size [23]. In fact, the optimal solution is very likely to use chunk sizes that depend on additional information that becomes available during the execution (i.e., failure occurrences to date). Using Proposition 1, we can write

$$\begin{aligned}&\mathbb{E}(T^*(\mathcal{W}|\tau)) = \\ &\min_{0<\omega_1\leq\mathcal{W}}\left(\begin{array}{l}P_{suc}(\omega_1 + C|\tau)\Big(\omega_1+C+\mathbb{E}(T^*(\mathcal{W}-\omega_1|\tau+\omega_1+C))\Big) \\ +(1 - P_{suc}(\omega_1 + C|\tau))\times \\ \qquad (\mathbb{E}(T_{lost}(\omega_1 + C|\tau))+\mathbb{E}(T_{rec})+\mathbb{E}(T^*(\mathcal{W}|R))\end{array}\right.\end{aligned}$$

Algorithm 1: DPMakespan (x,b,y,τ)

```
if x = 0 then
    return 0
if solution[x][b][y] = unknown then
    best ← ∞
    τ ← bτ + yu
    for i = 1 to x do
        exp_succ ← first(DPMakespan(x − i, b, y + i + C/u | τ))
        exp_fail ← first(DPMakespan(x, 0, R/u | τ))
        cur ← P_suc(iu + C|τ)(iu + C + exp_succ)
              +(1 − P_suc(iu + C|τ))( E(T_lost(iu + C|τ))
                                      + E(T_rec) + exp_fail )
        if cur < best then
            best ← cur;    chunksize ← i
    solution[x][b][y] ← (best, chunksize)
return solution[x][b][y]
```

which can be solved via dynamic programming. We introduce a time quantum u, meaning that all chunk sizes ω_i are integer multiples of u. This restricts the search for an optimal execution to a finite set of possible executions. The trade-off is that a smaller value of u leads to a more accurate solution, but also to a higher number of states in the algorithm, hence to a higher compute time.

Proposition 2. *Using a time quantum u, and for any failure inter-arrival time distribution,* DPMakespan *(Algorithm 1) computes an optimal solution to* Makespan *in time $\mathcal{O}(\frac{\mathcal{W}}{u}^3(1+\frac{C}{u})a)$, where a is an upper bound on the time needed to compute $E(T_{lost}(\omega|t))$, for any ω and t.*

Proof. Our goal is to compute $f(\omega|\tau)$, for any ω and τ values that can occur during the execution of $\mathcal{W}$ work units. A first attempt would be to design an algorithm A such that $A(x,y)$ computes an optimal solution assuming that the remaining work to process is $\omega = xu$ and the time since the last failure is $\tau = yu$, with $x \in [|0, \frac{\mathcal{W}}{u}|]$ and $y \in [|0, \delta|]$. To bound δ, we observe that the maximum possible elapsed time without failure occurs when (successfully) executing $\frac{\mathcal{W}}{u}$ chunks of size u, leading to $\delta = \frac{\frac{\mathcal{W}}{u}(u+C)+\tau}{u}$. To avoid using an arbitrarily large value τ, we instead introduce a boolean b which is equal to 1 only if a failure has never occurred since the initial state $(\omega|\tau)$. We now define an optimal solution as DPMakespan(x,b,y,τ), where the remaining work is $\omega = xu$ and the last failure occurred at time $\tau = b\tau + yu$, with $x \in [|0, \frac{\mathcal{W}}{u}|]$ and $y \in [|0, \frac{\mathcal{W}}{u}(1+\frac{C}{u})|]$. Note that DPMakespan returns a couple formed by the optimal expectation, and the corresponding optimal chunk size. All elements of array *solution* are initialized to *unknown* and T_{rec} can be computed using Proposition 1. Thus, the size of the chunk $f(\omega|\tau)$ is obtained by computing DPMakespan$(\frac{\omega}{u}, 1, 0, \tau)$. The complexity result is immediate. □

Algorithm 1 provides an approximation of the optimal solution to the Makespan problem. We evaluate this approximation experimentally in Section 5, including a direct comparison with the optimal solution in the case of Exponential failures (in which case the optimal can be computed via Theorem 1).

2.4 Solving NextFailure

Weighting the two cases in Equation 2 by their probabilities of occurrence, we obtain the expected amount of work successfully computed before the next failure:

$$\mathbb{E}(W(\omega|\tau)) = P_{suc}(\omega_1 + C|\tau)(\omega_1 + \mathbb{E}(W(\omega - \omega_1|\tau + \omega_1 + C))).$$

Here, unlike for MAKESPAN, the objective function to be maximized can easily be written as a closed form, even for arbitrary distributions. Developing the expression above leads to the following result:

Proposition 3. *The* NEXTFAILURE *problem is equivalent to maximizing the following quantity:*

$$\mathbb{E}(W(\mathcal{W}|\tau_0)) = \sum_{i=1}^{K} \omega_i \times \prod_{j=1}^{i} P_{suc}(\omega_j + C|t_j) \ ,$$

where $t_j = \tau_0 + \sum_{\ell=1}^{j-1}(\omega_\ell + C)$ *is the total time elapsed (without failure) before the start of the execution of chunk* ω_j*, and* K *is the (unknown) target number of chunks.*

Unfortunately, there does not seem to be an exact solution to this problem. However, just as for the MAKESPAN problem, the recursive definition of $\mathbb{E}(W(\mathcal{W}|\tau))$ lends itself naturally to a dynamic programming algorithm. The dynamic programming scheme is simpler because the size of the i-th chunk is only needed when no failure has occurred during the execution of the first $i-1$ chunks, regardless of the value of the τ parameter. More formally:

Proposition 4. *Using a time quantum* u*, and for any failure inter-arrival time distribution,* DPNEXTFAILURE *(Algorithm 2) computes an optimal solution to* NEXTFAILURE *in time* $\mathcal{O}(\frac{\mathcal{W}}{u}^3)$.

Proof. The goal is to compute $f(\omega|\tau)$, for any ω and τ that may occur during the execution of the job. We define DPNEXTFAILURE(x, n, τ) as an optimal solution for time quantum u, where $\omega = xu$ is the remaining work, n is the number of chunks already computed successfully, and τ is the amount of time elapsed since the last failure. Like DPMAKESPAN, DPNEXTFAILURE returns a couple. Note that x and n are in $[|0, \frac{\mathcal{W}}{u}|]$. Given x and n, the last failure necessarily occurred $\tau + (\mathcal{W} - xu) + nC$ units of time ago. Finally, we suppose that all elements of array *solution* are initialized to *unknown*. The size of the first chunk $f(\omega|\tau)$ is obtained by computing DPNEXTFAILURE$(\frac{\omega}{u}, 0, \tau)$, and the complexity result is immediate. □

3 Parallel jobs

3.1 Problem statement

We now turn to parallel jobs that can execute on any number of processors, p. We consider the following relevant scenarios for checkpointing/recovery overheads and for parallel execution times.

Algorithm 2: DPNextFailure (x,n,τ)

```
if x = 0 then
    return 0
if solution[x][n] = unknown then
    best ← ∞
    τ ← τ + (W − xu) + nC
    for i = 1 to x do
        work = first(DPNextFailure(x − i, n + 1|τ))
        cur ← P_suc(iu + C|τ) × (iu + work)
        if cur < best then
            best ← cur;    chunksize ← i
    solution[x][n] ← (best, chunksize)
return solution[x][n]
```

Checkpointing/recovery overheads – Checkpoints are synchronized over all processors. We use $C(p)$ and $R(p)$ to denote the time for saving a checkpoint and for recovering from a checkpoint on p processors, respectively (we assume that the downtime D does not depend on p). Assuming that the application's memory footprint is V bytes, with each processor holding V/p bytes, we consider two scenarios:

- Proportional overhead: $C(p) = R(p) = \alpha V/p$ for some constant α. This is representative of cases in which the bandwidth of the outgoing communication link of each processor is the I/O bottleneck.
- Constant overhead: $C(p) = R(p) = \alpha V$, which is representative of cases in which the incoming bandwidth of the resilient storage system is the I/O bottleneck.

Parallel work – Let $\mathcal{W}(p)$ be the time required for a failure-free execution on p processors. We use three models:

- Embarrassingly parallel jobs: $\mathcal{W}(p) = \mathcal{W}/p$.
- Amdahl parallel jobs: $\mathcal{W}(p) = \mathcal{W}/p + \gamma\mathcal{W}$. As in Amdahl's law [1], $\gamma < 1$ is the fraction of the work that is inherently sequential.
- Numerical kernels: $\mathcal{W}(p) = \mathcal{W}/p + \gamma\mathcal{W}^{2/3}/\sqrt{p}$. This is representative of a matrix product or a LU/QR factorization of size N on a 2D-processor grid, where $\mathcal{W} = O(N^3)$. In the algorithm in [3], $p = q^2$ and each processor receives $2q$ blocks of size N^2/q^2. Here γ is the communication-to-computation ratio of the platform.

We assume that the parallel job is tightly coupled, meaning that all p processors operate synchronously throughout the job execution. These processors execute the same amount of work $\mathcal{W}(p)$ in parallel, chunk by chunk. The total time (on one processor) to execute a chunk of size ω, and then checkpointing it, is $\omega + C(p)$. For the Makespan and NextFailure problems, we aim at computing a function f such that $f(\omega|\tau_1, \ldots, \tau_p)$ is the size of the next chunk that should be executed on every processor given a remaining amount of work $\omega \leq \mathcal{W}(p)$ and given a system state $(\tau_1, \ldots, \tau_p)$, where τ_i denotes the time elapsed since the last failure of the i-th processor. We assume that failure arrivals at all processors are *iid*.

An important remark on rejuvenation Two options are possible for recovering after a failure. Assume that a processor, say P_1, fails at time t. A first

option found in the literature [5, 24] is to rejuvenate *all* processors together with P_1, from time t to $t + D$ (e.g., via rebooting in case of software failure). Then all processors are available at time $t + D$ at which point they start executing the recovery simultaneously. In the second option, *only* P_1 is rejuvenated and the other processors are kept idle from time t to $t + D$. With this option any processor other than P_1 may fail between t and $t + D$ and thus may be itself in the process of being rejuvenated at time $t + D$.

Let us consider a platform with p processors that experience *iid* failures according to a Weibull distribution with scale parameter λ and shape parameter k, i.e., with cumulative distribution $F(t) = 1 - e^{-\frac{t^k}{\lambda^k}}$, and mean is $\mu = \lambda\Gamma(1 + \frac{1}{k})$. Define a *platform* failure as the occurrence of a failure at any of the processors. When rejuvenating all processors after each failure, platform failures are distributed according to a Weibull distribution with scale parameter $\frac{\lambda}{p^{1/k}}$ and shape parameter k. The MTBF for the platform is thus $D + \frac{\mu}{p^{1/k}}$ (note that the processor-level MTBF is $D + \mu$). When rejuvenating only the processor that failed, the platform MTBF is simply $\frac{D+\mu}{p}$. If $k = 1$, which corresponds to an Exponential distribution, rejuvenating all processors leads to a higher platform MTBF and is beneficial. However, if $k < 1$, rejuvenating all processors leads to a lower platform MTBF than rejuvenating only the processor that failed because $D \ll \frac{\mu}{p}$ in practical settings. This is shown on an example in Figure 1, which plots the platform MTBF vs. the number of processors. This behavior is easily explained: for a Weibull distribution with shape parameter $k < 1$, the probability $\mathbb{P}(X > t + x | X > t)$ strictly increases with t. In other words, a processor is less likely to fail the longer it remains in a fault-free state. It turns out that failure inter-arrival times for real-life systems have been modeled well by Weibull distributions whose shape parameter are strictly lower than 1 (either 0.7 or 0.78 in [10], 0.50944 in [17], between 0.33 and 0.49 in [22]). The overall conclusion is then that rejuvenating all processors after a failure, albeit commonly used in the literature, is likely not appropriate for large-scale platforms. Furthermore, even for Exponential distributions, rejuvenating all processors is not meaningful for hardware failures. Therefore, in the rest of this paper we assume that after a failure only the failed processor is rejuvenated[1].

3.2 Solving Makespan

In the case of the Exponential distribution, due to the memoryless property, the p processors used for a job can be conceptually aggregated into a virtual "macro-processor" with the following characteristics:

- Failure inter-arrival times follow an Exponential distribution of parameter $\lambda' = p\lambda$;
- The checkpoint and recovery overheads are $C(p)$ and $R(p)$, respectively.

A direct application of Theorem 1 yields the optimal solution of the MAKESPAN problem for parallel jobs:

Proposition 5. *Let $\mathcal{W}(p)$ be the amount of work to execute on p processors whose failure inter-arrival times follow iid Exponential distributions with parameter λ. Let $K_0 = \frac{p\lambda\mathcal{W}(p)}{1+\mathbb{L}(-e^{-p\lambda C(p)-1})}$. Then the optimal strategy to minimize the*

[1] For the sake of completeness, we considered both rejuvenation options for Exponential failures in the simulations (see Appendix B). We observe similar results for both options.

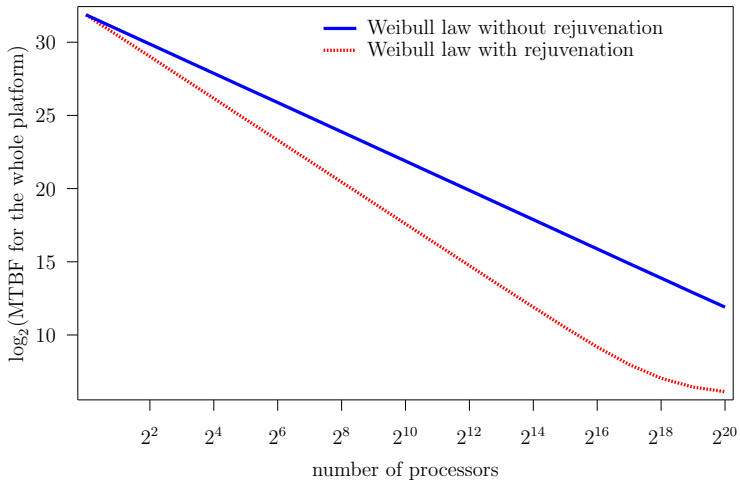

Figure 1: Impact of the two rejunevation options on the platform MTBF for a Weibull distribution with shape parameter 0.70, a processor-level MTBF of 125 years, and a downtime of 60 seconds.

expected makespan time is to split $\mathcal{W}(p)$ *into* $K^* = \max(1, \lfloor K_0 \rfloor)$ *or* $K^* = \lceil K_0 \rceil$ *same-size chunks, whichever minimizes* $\psi(K^*) = K^*(e^{p\lambda(\frac{\mathcal{W}(p)}{K^*}+C(p))} - 1)$.

Interestingly, although we know the optimal solution with p processors, we are not able to compute the optimal expected makespan analytically. Indeed, $\mathbb{E}(T_{rec})$, for which we had a closed form in the case of sequential jobs, becomes quite intricate in the case of parallel jobs. This is because during the downtime of a given processor another processor may fail. During the downtime of that processor, yet another processor may fail, and so on. We would need to compute the expected duration of these "cascading" failures until all processors are simultaneously available.

For arbitrary distributions, i.e., distributions without the memoryless property, we cannot (tractably) extend the dynamic programming algorithm DPMakespan. This is because one would have to memorize the evolution of the time elapsed since the last failure for all possible failure scenarios for each processor, leading to a number of states exponential in p. Fortunately, the dynamic programming approach for solving NextFailure can be extended to the case of a parallel job, as seen in Section 3.3. This was our motivation for studying NextFailure in the first place, and in the case of non-exponential failures, we use the solution to NextFailure as a heuristic solution for Makespan.

3.3 Solving NextFailure

For solving NextFailure using dynamic programming, there is no need to keep for each processor the time elapsed since its last failure as parameter of the recursive calls. This is because the τ variables of all processors evolve identically: recursive calls only correspond to cases in which no failure has occurred. Formally, the goal is to find a function $f(\omega|\tau) = \omega_1$ maximizing $\mathbb{E}(W(\omega|\tau_1, \ldots, \tau_p))$, where $\mathbb{E}(W(0|\tau_1, \ldots, \tau_p)) = 0$ and

$$\mathbb{E}(W(\omega|\tau_1,\ldots,\tau_p)) = \begin{cases} \omega_1 + \mathbb{E}(W(\omega-\omega_1|\tau_1+\omega_1+C(p),\ldots,\tau_p+\omega_1+C(p)) \\ \quad \text{if no processor fails during the next } \omega_1+C(p) \\ \quad \text{units of time} \\ 0 \quad \text{otherwise.} \end{cases}$$

Using a straightforward adaptation of DPNextFailure, which computes the probability of success

$$P_{suc}(x|\tau_1,\ldots,\tau_p) = \prod_{i=1}^{p} \mathbb{P}(X \geq x + \tau_i | X \geq \tau_i),$$

we obtain:

Proposition 6. *Using a time quantum u, for any failure inter-arrival time distribution,* DPNextFailure *computes an optimal solution to* NextFailure *with p processors in time $O(p\frac{\mathcal{W}}{u}^3)$.*

Even if a linear dependency in p, due to the computation of P_{suc}, seems a small price to pay, the above computational complexity is not tractable. Typical platforms in the scope of this paper (Jaguar [4], Exascale platforms) consist of tens of thousands of processors. The DPNextFailure algorithm is thus unusable, especially since it must be invoked after each failure. In what follows we propose a method to reduce its computational complexity.

Rather than working with the set of all p τ_i values, we approximate this set. With distributions such as the Weibull distribution, the smallest τ_i's have the highest impact on the overall probability of success. Therefore, we keep in the set the exact n_{exact} smallest τ_i values. Then we approximate the $p-n_{\text{exact}}$ remaining τ_i values using only n_{approx} "reference" values $\tau_1^{\approx}$, ..., $\tau_{n_{\text{approx}}}^{\approx}$. To each processor P_i whose τ_i value is not one of the n_{exact} smallest τ_i values, we associate one of the reference values. We can then simply keep track of how many processors are associated to each reference value, thereby vastly reducing computational complexity. We pick the reference values as follows. $\tau_1^{\approx}$ is the smallest of the remaining $p - n_{\text{exact}}$ exact τ_i values, while $\tau_{n_{\text{approx}}}^{\approx}$ is the largest. (Note that if processor P_i has never failed to date then $\tau_i = \tau_{n_{\text{approx}}}^{\approx}$.) The remaining $n_{\text{approx}} - 2$ reference values are chosen based on the distribution of the (*iid*) failure inter-arrival times. Assuming that X is a random variable distributed according to this distribution, then, for $i \in [2, n_{\text{approx}} - 1]$, we compute $\tau_i^{\approx}$ as

$$\tau_i^{\approx} = \text{quantile}\left(X, \frac{n_{\text{approx}} - i}{n_{\text{approx}} - 1}\,\mathbb{P}(X \geq \tau_1^{\approx}) + \frac{i-1}{n_{\text{approx}} - 1}\,\mathbb{P}(X \geq \tau_{n_{\text{approx}}}^{\approx})\right).$$

We have implemented DPNextFailure with $n_{\text{exact}} = 10$ and $n_{\text{approx}} = 100$. For the simulation scenario detailed in Section 5.2.2, we have studied the precision of this approximation by evaluating the relative error incurred when computing the probability using the approximated state rather than the exact one, for chunks of size 2^{-i} times the MTBF of the platform, with $i \in \{0..6\}$ and failure inter-arrival times following a Weibull distribution. It turns out that the larger the chunk size, the less accurate the approximation. Over the whole execution of a job in the settings of Section 5.2.2 (i.e., for 45,208 processors),

the worst relative error is lower than 0.2% for a chunk of duration equal to the MTBF of the platform. In practice, the chunks used by DPNextFailure are far smaller, and the approximation of their probabilities of success is thus far more accurate.

The running time of DPNextFailure is proportional to the work size $\mathcal{W}$. If $\mathcal{W}$ is significantly larger than the platform MTBF, which is very likely in practice, then with high probability a failure occurs before the last chunks of the solution produced by DPNextFailure are even considered for execution. In other words, a significant portion of the solution produced by DPNextFailure is unused, and can thus be discarded without a significant impact on the quality of the end result. In order to further boost the execution time of DPNextFailure, rather than invoking it on the size of the remaining work ω, we invoke it for a work size equal to $\min(\omega, 2 \times MTBF/p)$, where $MTBF$ is the processor-level mean time between failure ($MTBF/p$ is thus the platform mean time between failure). We use only the first half of the chunks in the solution produced by DPNextFailure so as to avoid any side effects due the truncated ω.

With all these optimizations, DPNextFailure runs in a few seconds even on the largest platforms. In all the application execution times reported in Sections 5 and 6 the execution time of DPNextFailure is taken into account.

4 Simulation framework

In this section we detail our simulation methodology. We use both synthetic and log-based failure distributions. The source code and all simulation results are publicly available at: http://graal.ens-lyon.fr/~fvivien/checkpoint.

4.1 Heuristics

Our simulator implements the following eight checkpointing policies (recall that $MTBF/p$ is the mean time between failures of the whole platform):

- Young is the periodic checkpointing policy of period $\sqrt{2 \times C(p) \times \frac{MTBF}{p}}$ given in [26].
- DalyLow is the first order approximation given in [8]. This is a periodic policy of period: $\sqrt{2 \times C(p) \times (\frac{MTBF}{p} + D + R(p))}$.
- DalyHigh is the periodic policy (high order approximation) given in [8].
- Bouguerra is the periodic policy given in [5].
- Liu is the non-periodic policy given in [17].
- OptExp is the periodic policy whose period is given in Proposition 5.
- DPNextFailure is the dynamic programming algorithm that maximizes the expectation of the amount of work completed before the next failure occurs.
- DPMakespan is the dynamic programming algorithm that minimizes the expectation of the makespan. For parallel jobs, DPMakespan makes the false assumption that all processors are rejuvenated after each failure (without this assumption this heuristic cannot be used).

	p_{total}	D	C,R	$MTBF$	$\mathcal{W}$
1-proc	1	60 s	600 s	1 h, 1 d, 1 w	20 d
Peta	$45,208$	60 s	600 s	125 y, 500 y	$1,000$ y
Exa	2^{20}	60 s	600 s	1250 y	$10,000$ y

Table 1: Parameters used in the simulations (C, R and D chosen according to [11, 6]). The first line corresponds to one-processor platforms, the second to Petascale platforms, and the third to Exascale platforms.

Our simulator also implements LowerBound, an omniscient algorithm that knows when the next failure will happen and checkpoints just in time, i.e., $C(p)$ time units before the failure. The makespan of LowerBound is thus an absolute lower bound on the makespan achievable by any policy. Note that LowerBound is unattainable in practice. Along the same line, the simulator implements PeriodLB, which implements a numerical search for the optimal period by evaluating each target period on 1,000 randomly generated scenarios (which would have a prohibitive computational cost in practice). The period computed by OptExp is multiplied and divided by $1 + 0.05 \times i$ with $i \in \{1, ..., 180\}$, and by 1.1^j with $j \in \{1, ..., 60\}$. PeriodLB corresponds to the periodic policy that uses the best period found by the search.

We point out that DalyLow, DalyHigh, and OptExp compute the checkpointing period based solely on the MTBF, which comes from the implicit assumption that failures are exponentially distributed. For the sake of completeness we nevertheless include them in all our simulations, simply using the MTBF value even when failures follow a Weibull distribution.

Performance evaluation. We compare heuristics using average makespan degradation, defined as follows. Given an experimental scenario (i.e., parameter values for failure distribution and platform configuration), we generate a set $X = \{tr_1, \ldots, tr_{600}\}$ of 600 traces. For each trace tr_i and each of the heuristics $heur_j$, we compute the achieved makespan, $res_{(i,j)}$. The makespan degradation for heuristic $heur_j$ on trace tr_i is defined as $v_{(i,j)} = res_{(i,j)}/min_{j \neq 0}\{res_{(i,j)}\}$ (where $heur_0$ is LowerBound). Finally, we compute the average degradation for heuristic $heur_j$ as $\sum_{i=1}^{600} v_{(i,j)}/600$. Standard deviations are small and thus not plotted on figures (see Table 2, Table 3 and Table 4 where standard deviations are reported).

4.2 Platforms

We target two types of platforms: Petascale and Exascale. For Petascale we choose as reference the Jaguar supercomputer [4], which contains $p_{total} = 45,208$ processors. We consider jobs that use between 1,024 and 45,208 processors. We then corroborate the Petascale results by running simulations of Exascale platforms with $p_{total} = 2^{20}$ processors. For both platform types, we determine the job size $\mathcal{W}$ so that a job using the whole platform would use it for a significant amount of time in the absence of failures, namely ≈ 8 days for Petascale platforms and ≈ 3.5 days for Exascale platforms. The parameters are listed in Table 1.

4.3 Generation of failure scenarios

Synthetic failure distributions

To choose failure distribution parameters that are representative of realistic systems, we use failure statistics from the Jaguar platform. Jaguar is said to experience on the order of 1 failure per day [18, 2]. Assuming a 1-day platform MTBF gives us a processor MTBF equal to $\frac{p_{total}}{365} \approx 125$ years, where $p_{total} = 45,208$ is the number of processors of the Jaguar platform. To verify that our results are consistent over a range of processor MTBF values, we also consider a processor MTBF of 500 years. We then compute the parameters of Exponential and Weibull distributions so that they lead to this MTBF value (recall that $MTBF = \mu + D \approx \mu$, where μ is the mean of the underlying distribution). Namely, for the Exponential distribution we set $\lambda = \frac{1}{MTBF}$ and for the Weibull distribution, which requires two parameters k and λ, we set $\lambda = MTBF/\Gamma(1 + 1/k)$. We first fix $k = 0.7$ based on the results of [22], and then vary it between 0.1 and 1.0.

Log-based failure distributions

We also consider failure distributions based on failure logs from production clusters. We used logs from Los Alamos National Laboratory [22], i.e., the logs for the largest clusters among the preprocessed logs in the *Failure trace archive* [14]. In these logs, each failure is tagged by the node —and not just the processor— on which the failure occurred. Among the 26 possible clusters, we opted for the logs of the only two clusters with more than 1,000 nodes. The motivation is that we need a sample history sufficiently large to simulate platforms with more than ten thousand nodes. The two chosen logs are for clusters 18 and 19 in the archive (referred to as 7 and 8 in [22]). For each log, we record the set$\mathcal{S}$ of availability intervals. The discrete failure distribution for the simulation is generated as follows: the conditional probability $\mathbb{P}(X \geq t \mid X \geq \tau)$ that a node stays up for a duration t, knowing that it had been up for a duration τ, is set equal to the ratio of the number of availability durations in $\mathcal{S}$ greater than or equal to t, over the number of availability durations in $\mathcal{S}_x$ greater than or equal to τ.

Scenario generation

Given a p-processor job, a failure trace is a set of failure dates for each processor over a fixed time horizon h. In the one-processor case, h is set to 1 year. In all the other cases, h is set to 11 years and the job start time, t_0, is assumed to be one-year to avoid side-effects related to the synchronous initialization of all nodes/processors. Given the distribution of inter-arrival times at a processor, for each processor we generate a trace via independent sampling until the target time horizon is reached. Finally, for simulations where the only varying parameter is the number of processors $a \leq p \leq b$, we first generate traces for b processors. For experiments with p processors we then simply select the first p traces. This ensures that simulation results are coherent when varying p.

The two clusters used for computing our log-based failure distribution consist of 4-processor nodes. Hence, to simulate a 45,208-processor platform we generate 11,302 failure traces, one for each four-processor node.

5 Simulations with synthetic failures

5.1 Single processor jobs

For a single processor, we cannot use a 125-year MTBF, as a job would have to run for centuries in order to need a few checkpoints. Hence we study scenarios with smaller values of the MTBF, from one hour to one week. This study, while unrealistic, allows us to compare the performance of DPNextFailure with that of DPMakespan.

5.1.1 Exponential failures

Table 2 shows the average makespan degradation for the eight heuristics and the two lower bounds, in the case of exponentially distributed failure inter-arrival times. Unsurprisingly, LowerBound is significantly better than all heuristics, especially for a small MTBF. At first glance it may seem surprising that PeriodLB achieves results close but not equal to 1. This is because although the expected optimal solution is periodic, checkpointing with the optimal period is not always the best strategy for a given random scenario.

A first interesting observation is that the performance by the well-known Young, DalyLow and DalyHigh heuristics is indeed close to optimal. While this result seems widely accepted, we are not aware of previously published simulation studies that have demonstrated it. Looking more closely at the results (see Appendix A) we find that, in the neighborhood of the optimal period, the performance of periodic policies is almost independent of the period. This explains while the Young, DalyLow and DalyHigh heuristics have near optimal performance even if their periods differ.

In Section 2.4, we claimed that DPNextFailure should provide a reasonable solution to the Makespan problem. We observe that, at least in the one-processor case, DPNextFailure does lead to solutions that are close to those computed by DPMakespan and to the optimal.

Heuristics	MTBF					
	1 hour		1 day		1 week	
	avg	std	avg	std	avg	std
LowerBound	0.62852	0.00761	0.90679	0.01243	0.97874	0.01430
PeriodLB	1.00739	0.00627	1.01600	0.00873	1.02285	0.00923
Young	1.01755	0.01051	1.01600	0.00928	1.02325	0.00937
DalyLow	1.02809	0.01198	1.01622	0.00943	1.02330	0.00937
DalyHigh	1.00732	0.00624	1.01596	0.00881	1.02339	0.00951
Liu	1.01755	0.0097	1.05519	0.03699	1.20965	0.16736
Bouguerra	1.02673	0.01175	1.02349	0.01314	1.02690	0.01290
OptExp	1.00739	0.00627	1.01604	0.00873	1.02285	0.00923
DPNextFailure	1.00787	0.00599	1.01705	0.00875	1.02830	0.01360
DPMakespan	1.00781	0.00639	1.01618	0.01015	1.03500	0.01976

Table 2: Degradation from best for a single processor with Exponential failures.

5.1.2 Weibull failures

Table 3 shows results when failure inter-arrival times follow a Weibull distribution (note that the LIU heuristic was specifically designed to handle Weibull distributions). Unlike in the exponential case, the optimal checkpoint policy may be non-periodic [23]. Results in the table show that all the heuristics lead to results that are close to the optimal, except LIU when the MTBF is not small. The implication is that, in the one-processor case, one can safely use YOUNG, DALYLOW, and DALYHIGH, which only require the failure MTBF, even for Weibull failures. In Section 5.2.2 we see that this result does not hold for multi-processor platforms. Note that, just like in the Exponential case, DPNEXTFAILURE leads to solutions that are close to those computed by DPMAKESPAN.

Heuristics	MTBF					
	1 hour		1 day		1 week	
	avg	std	avg	std	avg	std
LOWERBOUND	0.66351	0.00990	0.90994	0.01711	0.97598	0.01654
PERIODLB	1.00971	0.00669	1.01602	0.00949	1.02275	0.00955
YOUNG	1.00954	0.00689	1.01645	0.00920	1.02300	0.00969
DALYLOW	1.01159	0.00811	1.01654	0.00944	1.02304	0.00968
DALYHIGH	1.01726	0.00852	1.01606	0.00936	1.02304	0.00961
LIU	1.01024	0.00787	1.06873	0.05758	1.19364	0.17197
BOUGUERRA	1.02867	0.01274	1.01894	0.01041	1.02338	0.01064
OPTEXP	1.01731	0.00849	1.01659	0.00935	1.02284	0.00955
DPNEXTFAILURE	1.01353	0.00819	1.01686	0.00921	1.02727	0.01321
DPMAKESPAN	1.00728	0.00727	1.01570	0.01012	1.03528	0.02136

Table 3: Degradation from best for a single processor with Weibull failures.

5.2 Parallel jobs

Section 3.1 defines 3×2 combinations of parallelism and checkpointing overhead models. For our experiments we have instantiated these models as follows: $\mathcal{W}(p)$ is equal to either $\frac{\mathcal{W}}{p}$, $\frac{\mathcal{W}}{p} + \gamma\mathcal{W}$ with $\gamma \in \{10^{-4}, 10^{-6}\}$, or $\frac{\mathcal{W}}{p} + \gamma\frac{\mathcal{W}^{2/3}}{\sqrt{p}}$ with $\gamma \in \{0.1, 1, 10\}$; and $C(p) = R(p) = 600$ seconds or $C(p) = R(p) = 600 \times p_{total}/p$ seconds. Due to lack of space, in this paper we only report results for the embarrassingly parallel applications ($\mathcal{W}(p) = \mathcal{W}/p$) with constant checkpoint overhead ($C(p) = R(p) = 600$ seconds). Results for all other cases lead to the same conclusions regarding the relative performance of the various checkpointing strategies. We refer the reader to Appendix B, which contains the comprehensive set of results for all combinations of parallelism and checkpointing overhead models.

5.2.1 Exponential failures

Petascale platforms – Figure 2 shows results for Petascale platforms. The main observation is that, regardless of the number of processors p, the YOUNG, DALYLOW, and DALYHIGH heuristics compute an almost optimal solution (i.e.,

with degradation below 1.023) indistinguishable from that of OptExp and PeriodLB. By contrast, the degradation of Bouguerra is only slightly higher, and that of Liu[2] is ≈ 1.09. We see that DPNextFailure behaves satisfactorily: its degradation is less than 4.2‰ worse than that for OptExp for $p \geq 2^{13}$, and less than 1.85% worse overall. We also observe that DPNextFailure always performs better than DPMakespan. This is likely due to the false assumption in DPMakespan that all processors are rejuvenated after each failure. The same conclusions are reached when the MTBF per processor is 500 years instead of 125 years (see Appendix B).

Exascale platforms – Results for Exascale platforms, shown in Figure 3, corroborates the results obtained for Petascale platforms.

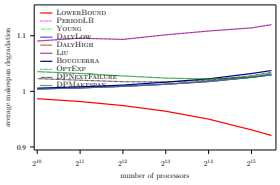

Figure 2: Evaluation of the different heuristics on a Petascale platform with Exponential failures.

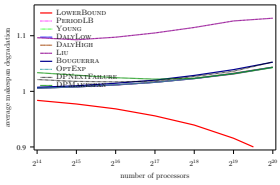

Figure 3: Evaluation of the different heuristics on an Exascale platform with Exponential failures.

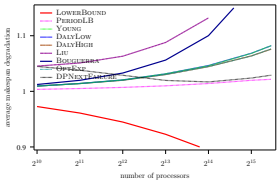

Figure 4: Evaluation of the different heuristics on a Petascale platform with Weibull failures.

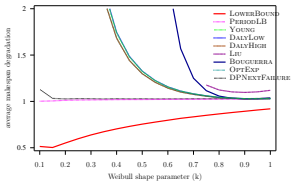

Figure 5: Varying the shape parameter k of the Weibull distribution for a Jaguar-like platform with $45,208$ processors.

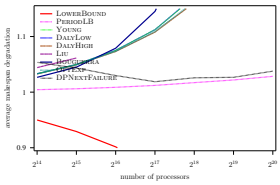

Figure 6: Evaluation of the different heuristics on an Exascale platform with Weibull failures.

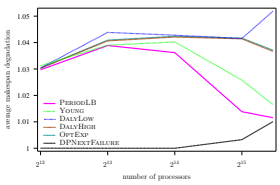

Figure 7: Evaluation of the different heuristics on a Petascale platform with failures based on the failure log of LANL cluster 19.

5.2.2 Weibull failures

Petascale platforms – A key contribution of this paper is the comparison between DPNextFailure and all previously proposed heuristics for the Makespan problem on platforms whose failure inter-arrival times follow a Weibull

[2]On most figures the curve for Liu is incomplete. Liu computes the *dates* at which the application should be checkpointed. In several cases the interval between two consecutive dates is smaller than the checkpoint duration, C, which is nonsensical. In such cases we do not report any result for Liu and speculate that there may be an error in [17].

distribution. Existing heuristics provide good solutions for sequential jobs (see Section 5.1). Figure 4 shows that this is no longer the case beyond $p = 1,024$ processors as demonstrated by growing gaps between heuristics and PeriodLB as p increases. For large platforms, only DPNextFailure is able to bridge this gap. For example with $45,208$ processors, Young, DalyLow, and DalyHigh are at least 4.3% worse than DPNextFailure, the latter being only 0.59% worse than PeriodLB. These poor results of previously proposed heuristics are partly due to the fact that the optimal solution is not periodic. For instance, throughout a complete execution with $45,208$ processors, DPNextFailure changes the size of inter-checkpoint intervals from $2,984$ seconds up to $6,108$ seconds. Liu, which is specifically designed to handle Weibull failures, has far worse performance than other heuristics, and fails to compute meaningful checkpoint dates for large platforms. Bouguerra is also supposed to handle Weibull failures but has poor performance because it relies on the assumption that all processors are rejuvenated after each failure. We conclude that our dynamic programming approach provides significant improvements over all previously proposed approaches for solving the Makespan problem in the case of large platforms. The same conclusions are reached when the MTBF per processor is 500 years instead of 125 years (see Appendix B.3).

Heuristics	Average degradation	Standard deviation
LowerBound	0.83366	0.01578
PeriodLB	1.02169	0.01345
Young	1.08226	0.03560
DalyLow	1.08211	0.03612
DalyHigh	1.07588	0.03333
Bouguerra	1.25020	0.02543
OptExp	1.07645	0.03348
DPNextFailure	1.02910	0.01647

Table 4: Average degradation from best and standard deviation for a 45208 processor platform with Weibull failures, embarrassingly parrallel job and fixed checkpoint/recovery costs

Number of spare processors necessary – In our simulations, for a job running around 10.5 days on a 45,208 processor platform, when using DPNextFailure, on average, 38.0 failures occur during a job execution, with a maximum of 66 failures. This provides some guidance regarding the number of spare processors necessary so as not to experience any job interruption, in this case circa 1‰.

Impact of the shape parameter k – We report results from experiments in which we vary the shape parameter k of the Weibull distribution in a view to assessing the sensitivity of each heuristic to this parameter. Figure 5 shows average makespan degradation vs. k. We see that, with small values of k, the degradation is small for DPNextFailure (below 1.033 for $k \geq 0.15$, and 1.130 for $k = 0.10$), while it is dramatically larger for all other heuristics. DPNextFailure achieves the best performance over all heuristics for the range of k values seen in practice as reported in the literature (between 0.33 and 0.78 [10, 17, 22]). Liu fails to compute a solution for $k \leq 0.70$. Bouguerra

leads to very poor solutions because it assumes processor rejuvenation, which is only acceptable for k close to 1 (i.e., for nearly Exponential failures) but becomes increasingly harmful as k becomes smaller.

Exascale platforms – Figure 6 presents results for Exascale platforms. The advantage of DPNextFailure over the other heuristics is even more pronounced than for Petascale platforms. The average degradation from best of DPNextFailure for platforms consisting of between 2^{16} and 2^{20} processors is less than 1.028, the reference being defined by the inaccessible performance of PeriodLB.

6 Simulations with log-based failures

To fully assess the performance of DPNextFailure, we also perform simulations using the failure logs of two production clusters, following the methodology explained in Section 4.3. We compare DPNextFailure to the Young, DalyLow, DalyHigh, and OptExp heuristics, adapting them by pretending that the underlying failure distribution is Exponential with the same MTBF at the empirical MTBF computed from the log. The same adaptation cannot be done for Liu, Bouguerra, and DPMakespan, which are thus not considered in this section.

Simulation results corresponding to one of the production cluster (LANL cluster 19, see Section 4.3) are shown in Figure 7. For the sake of readability, we do not display LowerBound as it leads to low values ranging from 0.80 to 0.56 as p increases (which underlines the intrinsic difficulty of the problem). As before, DalyHigh and OptExp achieve similar performance. But their performance, alongside that of DalyLow, is significantly worse than that of Young, especially for large p. The performance of all these heuristics is closer to the performance of PeriodLB than in the case of Weibull failures. The main difference with results for synthetic failures is that the performance of DPNextFailure is even better than that of PeriodLB. This is because, for these real-world failure distributions, periodic heuristics are inherently suboptimal. By contrast, DPNextFailure keeps adapting the size of the chunks that it attempts to execute. On a 45,208 processor platform, the processing time (or size) of the attempted chunks range from as (surprisingly) low as 60s up to 2280s. These values may seem extremely low, but the platform MTBF in this case is only 1,297s (while R+C=1,200s). This is thus a very difficult problem instance, but DPNextFailure solves it satisfactorily. More concretely, DPNextFailure saves more than 18,000 processor hours when using 45,208 processors, and more than 262,000 processor hours using 32,768 processors, compared to PeriodLB.

Simulation results based on the failure log of the other cluster (cluster 18) are similar, and even more in favor of DPNextFailure (see Appendix E).

7 Related work

In [8], Daly studies periodic checkpointing of applications executed on platforms where failures inter-arrival times are exponentially distributed. That study accounts for checkpointing and recovery overheads (but not for downtimes), and allows failures to happen during recoveries. Two estimates of the optimal period

are proposed. The lower order estimate is a generalization of Young's approximation [26], which takes recovery overheads into account. The higher order estimate is ill-formed as it relies on an equation that sums up non-independent probabilities (Equation (13) in [8]). That work was later extended in [12], which studies the impact of sub-optimal periods on application performance.

In [5], Bouguerra et al. study the design of an optimal checkpointing policy when failures can occur during checkpointing and recovery, with checkpointing and recovery overheads depending upon the application progress. They show that the optimal checkpointing policy is periodic when checkpointing and recovery overheads are constant, and when failure inter-arrival times follow either an Exponential or a Weibull distribution. They also give formulas to compute the optimal period in both cases. Their results, however, rely on the unstated assumption that all processors are rejuvenated after each failure and after each checkpoint. The work in [24] suffers from the same issue.

In [25], the authors claim to use an "optimal checkpoint restart model [for] Weibull's and Exponential distributions" that they have designed in another paper (referenced as [1] in [25]). However, this latter paper is not available, and we were unable to compare our work to their solution. However, as explained in [25] the "optimal" solution in [1] is found using the assumption that checkpoint is periodic (even for Weibull failures). In addition, the authors of [25] partially address the question of the optimal number of processors for parallel jobs, presenting experiments for four MPI applications, using a non-optimal policy, and for up to 35 processors. Our approach is radically different since we target large-scale platforms with up to tens of thousands of processors and rely on generic application models for deriving optimal solutions.

In this work, we solve the NextFailure problem to obtain heuristic solutions to the Makespan problem in the case of parallel jobs. The NextFailure problem has been studied by many authors in the literature, often for single-processor jobs. Maximizing the expected work successfully completed before the first failure is equivalent to minimizing the expected wasted time before the first failure, which is itself a classical problem. Some authors propose analytical resolution using a "checkpointing frequency function", for both infinite (see [16, 17]) and finite time horizons (see [19]). However, these works use approximations, for example assuming that the expected failure occurrence is exactly halfway between two checkpointing events, which does not hold for general failure distributions. Approaches that do not rely on a checkpointing frequency function are used in [23, 15], but only for infinite time horizons.

8 Conclusion

We have studied the problem of scheduling checkpoints for minimizing the makespan of sequential and parallel jobs on large-scale and failure-prone platforms, which we have called Makespan. An auxiliary problem, NextFailure, was introduced as an approximation of Makespan. Both problems are defined rigorously in general settings. For exponential distributions, we have provided a complete analytical solution of Makespan together with an assessment of the quality of the NextFailure approximation. We have also designed dynamic programming solutions for both problems, that can be applied for any failure distribution.

We have obtained a number of key results via simulation experiments. For Exponential failures, our approach allows us to determine the optimal checkpointing policy. For Weibull failures, we have demonstrated the importance of using the "single processor rejuvenation" model. With this model, we have shown that our dynamic programming algorithm leads to significantly more efficient executions than all previously proposed algorithms with an average decrease in the application makespan of at least 4.38% for our largest simulated Petascale platforms, and of at least 30.7% for our largest simulated Exascale platforms. We have also considered failures from empirical failure distributions extracted from failure logs of two production clusters. In this settings, once again our dynamic programming algorithm leads to significantly more efficient executions than all previously proposed algorithms. Given that our results also hold across our various application and checkpoint scenarios, we claim that our dynamic programming approach provides a key step for the effective use of next-generation large-scale platforms. Furthermore, our dynamic programming approach can be easily extended to settings in which the checkpoint and restart costs are not constants but depends on the progress of the application execution.

There are several promising avenues for future work. Interesting questions relate to computing the optimal number of processors for executing a parallel job. On a fault-free machine, and for all the scenarios considered in the paper (embarrassingly parallel, Amdahl law, numerical kernels), the execution time of the job decreases with the number of enrolled resources, and hence is minimal when the whole platform is used. In the presence of failures, this is no longer true, and the expected makespan may be smaller when using fewer processors than p_{total}. This leads to the idea of replicating the execution of a given job on say, both halves of the platform, i.e., with $p_{total}/2$ processors each. This could be done independently, or better, by synchronizing the execution after each checkpoint. The question of which is the optimal strategy is open. Another research direction is provided by recognizing that the (expected) makespan is not the only worthwhile or relevant objective. Because of the enormous energy cost incurred by large-scale platforms, along with environmental concerns, a crucial direction for future work is the design of checkpointing strategies that can trade off a longer execution time for a reduced energy consumption.

It is reasonable to expect that parallel jobs will be deployed successfully on exascale platforms only by using multiple techniques together (checkpointing, migration, replication, self-tolerant algorithms). While checkpointing is only part of the solution, it is an important part. This paper has evidenced the intrinsic difficulty of designing efficient checkpointing strategies, but it has also given promising results that greatly improve state-of-the art approaches.

References

[1] G. Amdahl. The validity of the single processor approach to achieving large scale computing capabilities. In *AFIPS Conference Proceedings*, volume 30, pages 483–485. AFIPS Press, 1967.

[2] L. Bautista Gomez, A. Nukada, N. Maruyama, F. Cappello, and S. Matsuoka. Transparent low-overhead checkpoint for GPU-accelerated clusters. https://wiki.ncsa.illinois.edu/download/attachments/17630761/

INRIA-UIUC-WS4-lbautista.pdf?version=1&modificationDate=1290470402000.

[3] L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley. *ScaLAPACK Users' Guide*. SIAM, 1997.

[4] A.S. Bland, R.A. Kendall, D.B. Kothe, J.H. Rogers, and G.M. Shipman. Jaguar: The World's Most Powerful Computer. In *GUC'2009*, 2009.

[5] Mohamed-Slim Bouguerra, Thierry Gautier, Denis Trystram, and Jean-Marc Vincent. A flexible checkpoint/restart model in distributed systems. In *PPAM*, volume 6067 of *LNCS*, pages 206–215, 2010.

[6] Franck Cappello, Henri Casanova, and Yves Robert. Checkpointing vs. migration for post-petascale supercomputers. In *ICPP'2010*. IEEE Computer Society Press, 2010.

[7] V. Castelli, R. E. Harper, P. Heidelberger, S. W. Hunter, K. S. Trivedi, K. Vaidyanathan, and W. P. Zeggert. Proactive management of software aging. *IBM J. Res. Dev.*, 45(2):311–332, 2001.

[8] J. T. Daly. A higher order estimate of the optimum checkpoint interval for restart dumps. *Future Generation Computer Systems*, 22(3):303–312, 2004.

[9] Jack Dongarra, Pete Beckman, Patrick Aerts, Frank Cappello, Thomas Lippert, Satoshi Matsuoka, Paul Messina, Terry Moore, Rick Stevens, Anne Trefethen, and Mateo Valero. The international exascale software project: a call to cooperative action by the global high-performance community. *Int. J. High Perform. Comput. Appl.*, 23(4):309–322, 2009.

[10] T. Heath, R. P. Martin, and T. D. Nguyen. Improving cluster availability using workstation validation. *SIGMETRICS Perf. Eval. Rev.*, 30(1):217–227, 2002.

[11] J.C.Y. Ho, C.L. Wang, and F.C.M. Lau. Scalable group-based checkpoint/restart for large-scale message-passing systems. In *Parallel and Distributed Processing, 2008. IPDPS 2008. IEEE International Symposium on*, pages 1–12. IEEE, 2008.

[12] W.M. Jones, J.T. Daly, and N. DeBardeleben. Impact of sub-optimal checkpoint intervals on application efficiency in computational clusters. In *HPDC'10*, pages 276–279. ACM, 2010.

[13] Nick Kolettis and N. Dudley Fulton. Software rejuvenation: Analysis, module and applications. In *FTCS '95*, page 381, Washington, DC, USA, 1995. IEEE CS.

[14] Derrick Kondo, Bahman Javadi, Alexandru Iosup, and Dick Epema. The failure trace archive: Enabling comparative analysis of failures in diverse distributed systems. *Cluster Computing and the Grid, IEEE International Symposium on*, 0:398–407, 2010.

[15] P. L'Ecuyer and J. Malenfant. Computing optimal checkpointing strategies for rollback and recovery systems. *IEEE Transactions on computers*, 37(4):491–496, 2002.

[16] Y. Ling, J. Mi, and X. Lin. A variational calculus approach to optimal checkpoint placement. *IEEE Transactions on computers*, pages 699–708, 2001.

[17] Y. Liu, R. Nassar, C. Leangsuksun, N. Naksinehaboon, M. Paun, and SL Scott. An optimal checkpoint/restart model for a large scale high performance computing system. In *IPDPS 2008*, pages 1–9. IEEE, 2008.

[18] E. Meneses. Clustering Parallel Applications to Enhance Message Logging Protocols. https://wiki.ncsa.illinois.edu/download/attachments/17630761/INRIA-UIUC-WS4-emenese.pdf?version=1&modificationDate=1290466786000.

[19] T. Ozaki, T. Dohi, H. Okamura, and N. Kaio. Distribution-free checkpoint placement algorithms based on min-max principle. *IEEE TDSC*, pages 130–140, 2006.

[20] Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 2005.

[21] Vivek Sarkar and others. Exascale software study: Software challenges in extreme scale systems, 2009. White paper available at: http://users.ece.gatech.edu/mrichard/ExascaleComputingStudyReports/ECSS%20report%20101909.pdf.

[22] B. Schroeder and G. A. Gibson. A large-scale study of failures in high-performance computing systems. In *Proc. of DSN*, pages 249–258, 2006.

[23] A.N. Tantawi and M. Ruschitzka. Performance analysis of checkpointing strategies. *ACM TOCS*, 2(2):123–144, 1984.

[24] S. Toueg and O. Babaoglu. On the optimum checkpoint selection problem. *SIAM J. Computing*, 13(3):630–649, 1984.

[25] K. Venkatesh. Analysis of Dependencies of Checkpoint Cost and Checkpoint Interval of Fault Tolerant MPI Applications. *Analysis*, 2(08):2690–2697, 2010.

[26] John W. Young. A first order approximation to the optimum checkpoint interval. *Communications of the ACM*, 17(9):530–531, 1974.

A Appendix : Results for one processor

A.1 Exponential law

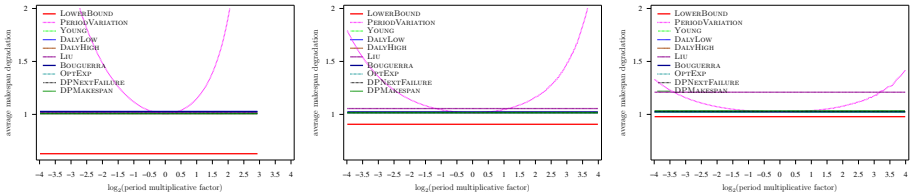

a) $MTBF = 1$ hour b) $MTBF = 1$ day c) $MTBF = 1$ week

Figure 8: Evaluation of the different heuristics on one processor with Exponential failures

A.2 Weibull law

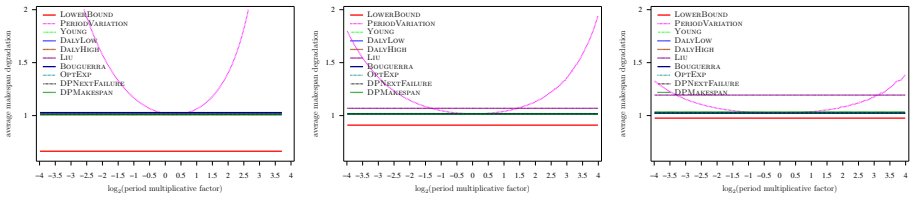

a) $MTBF = 1$ hour b) $MTBF = 1$ day c) $MTBF = 1$ week

Figure 9: Evaluation of the different heuristics on one processor with Weibull failures ($k = 0.7$)

B Appendix : Results for Petascale platforms

Notice that in some cases (typically when $MTBF = 500$ years, see Figure 11 b)), PeriodVariation leads to results that does not depend on the period multiplicative factor j. This can be explained by observing that for large j, $1.1^j \times$ OptExp becomes greater than $\mathcal{W}(p)$, and thus PeriodVariation simply tries to execute the whole work in one chunk, and this strategy no longer depends on j.

B.1 Exponential law, rejuvenating all processors

B.1.1 Fixed checkpoint and recovery cost

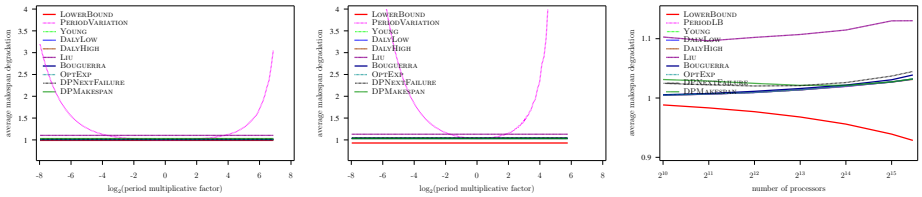

a) 1024 processors b) 45208 processors c)

Figure 10: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years) using using embarrassingly parallel job, and constant overhead model

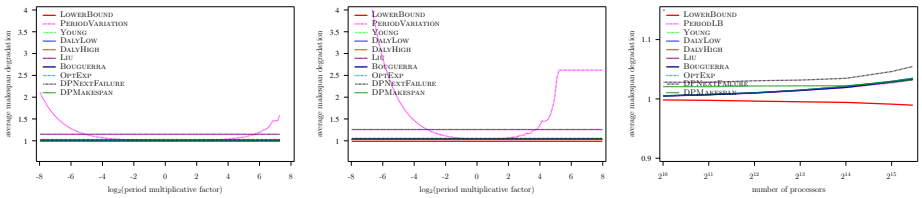

a) 1024 processors b) 45208 processors c)

Figure 11: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years) using using embarrassingly parallel job, and constant overhead model

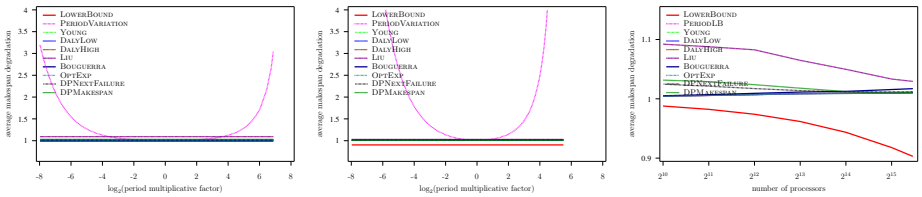

a) 1024 processors b) 45208 processors c)

Figure 12: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years) using Amdahl law with $\gamma = 10^{-4}$, and constant overhead model

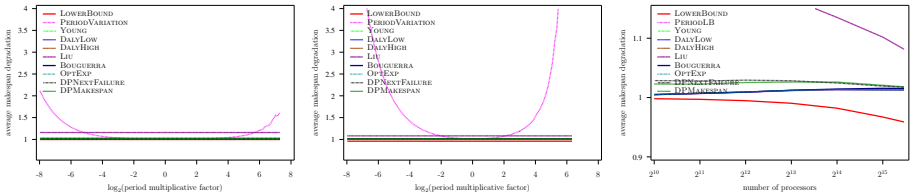

a) 1024 processors b) 45208 processors c)

Figure 13: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years) using Amdahl law with $\gamma = 10^{-4}$, and constant overhead model

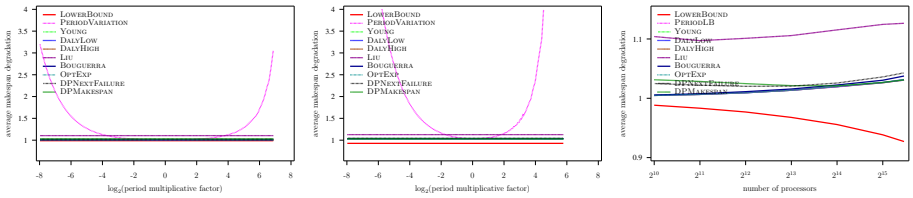

a) 1024 processors b) 45208 processors c)

Figure 14: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Amdahl law with $\gamma = 10^{-6}$, and constant overhead model

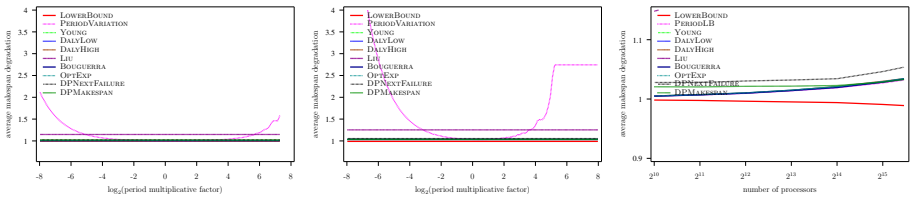

a) 1024 processors b) 45208 processors c)

Figure 15: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years), using Amdahl law with $\gamma = 10^{-6}$, and constant overhead model

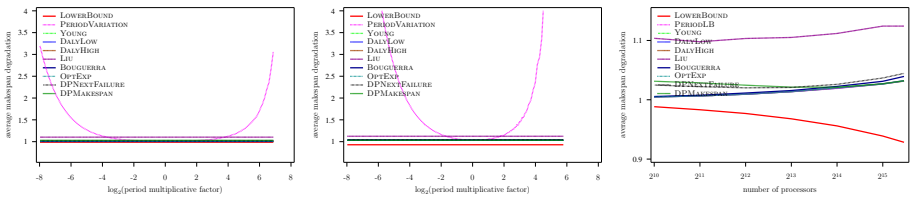

a) 1024 processors b) 45208 processors c)

Figure 16: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Numerical Kernel law with $\gamma = 0.1$, and constant overhead model

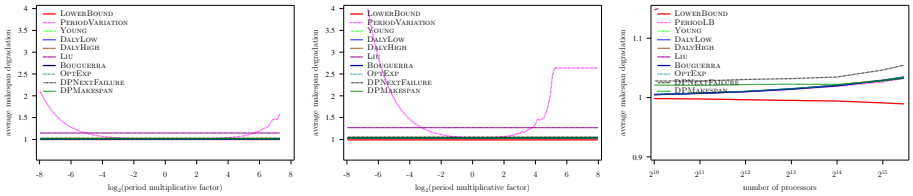

a) 1024 processors b) 45208 processors c)

Figure 17: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years), using Numerical Kernel law with $\gamma = 0.1$, and constant overhead model

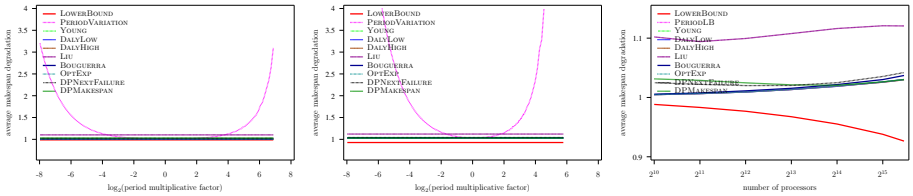

a) 1024 processors b) 45208 processors c)

Figure 18: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Numerical Kernel law with $\gamma = 1$, and constant overhead model

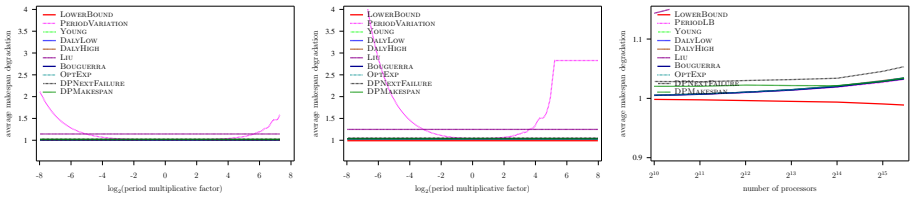

a) 1024 processors b) 45208 processors c)

Figure 19: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years), using Numerical Kernel law with $\gamma = 1$, and constant overhead model

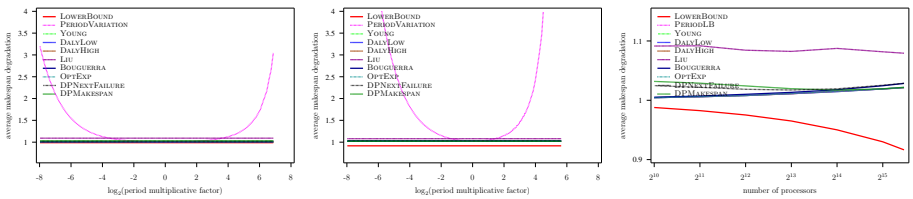

a) 1024 processors b) 45208 processors c)

Figure 20: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Numerical Kernel law with $\gamma = 10$, and constant overhead model

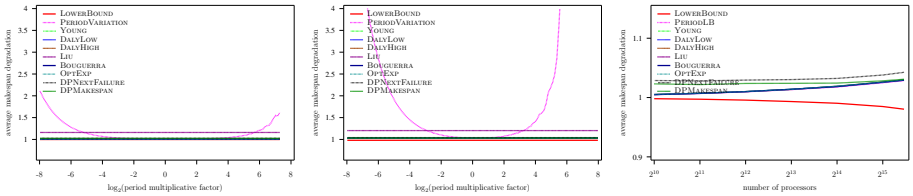

a) 1024 processors b) 45208 processors c)

Figure 21: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF = 500$ years), using Numerical Kernel law with $\gamma = 10$, and constant overhead model

B.1.2 Variable checkpoint and recovery cost

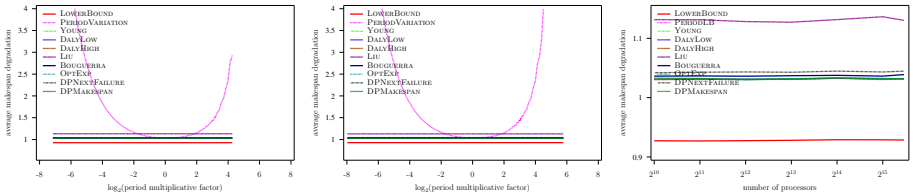

a) 1024 processors b) 45208 processors c)

Figure 22: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF = 125$ years), using embarrassingly parallel job, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

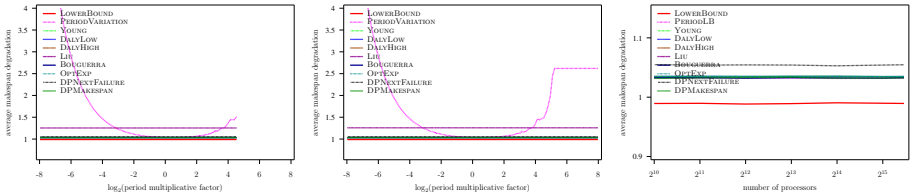

a) 1024 processors b) 45208 processors c)

Figure 23: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF = 500$ years), using embarrassingly parallel job, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

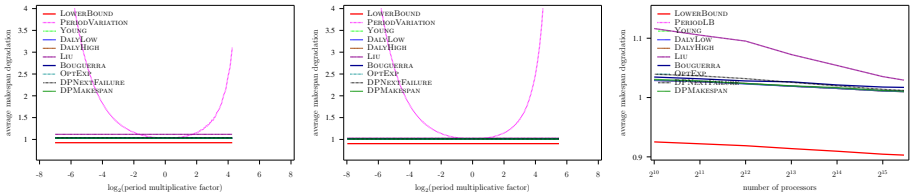

a) 1024 processors b) 45208 processors c)

Figure 24: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF = 125$ years), using Amdahl law with $\gamma = 10^{-4}$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

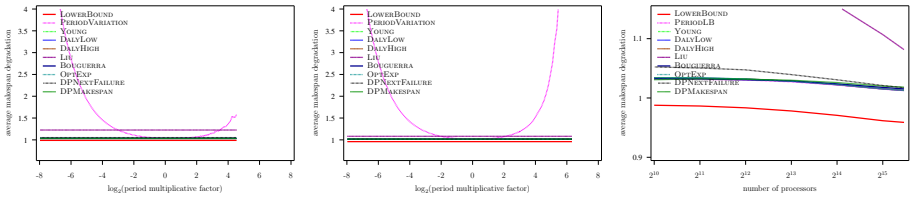

a) 1024 processors b) 45208 processors c)

Figure 25: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF = 500$ years), using Amdahl law with $\gamma = 10^{-4}$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

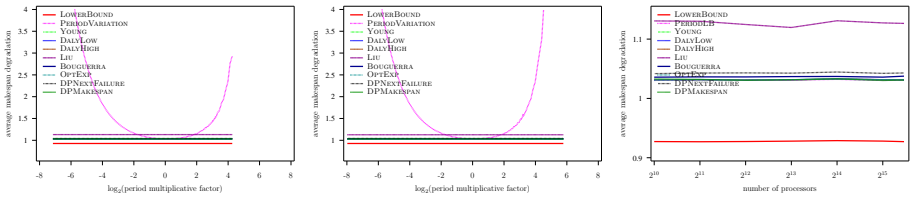

a) 1024 processors b) 45208 processors c)

Figure 26: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF = 125$ years), using Amdahl law with $\gamma = 10^{-6}$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

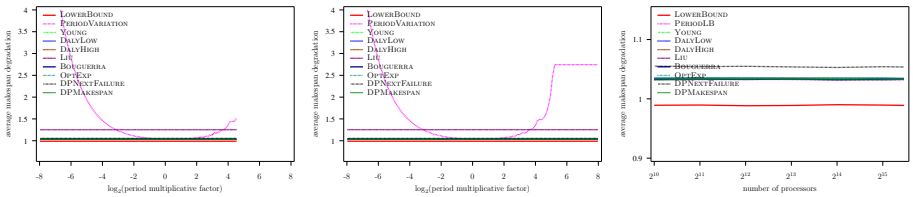

a) 1024 processors b) 45208 processors c)

Figure 27: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF = 500$ years), using Amdahl law with $\gamma = 10^{-6}$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

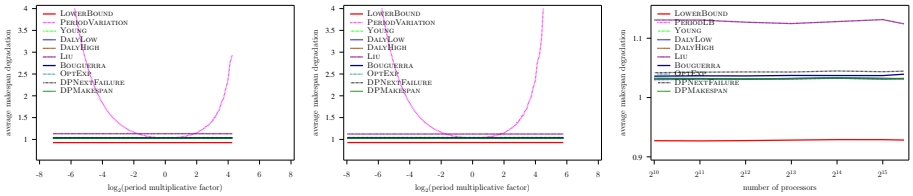

a) 1024 processors b) 45208 processors c)

Figure 28: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Numerical Kernel law with $\gamma = 0.1$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

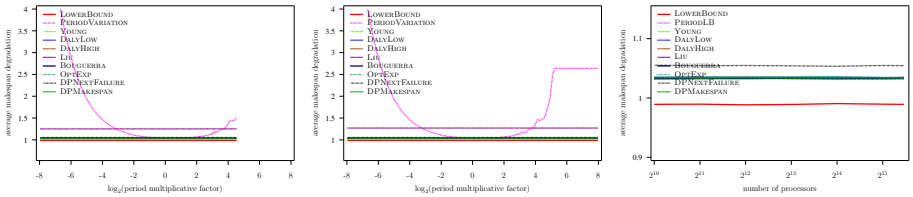

a) 1024 processors b) 45208 processors c)

Figure 29: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years), using Numerical Kernel law with $\gamma = 0.1$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

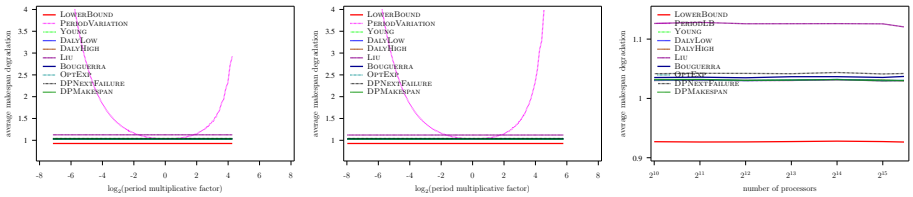

a) 1024 processors b) 45208 processors c)

Figure 30: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Numerical Kernel law with $\gamma = 1$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

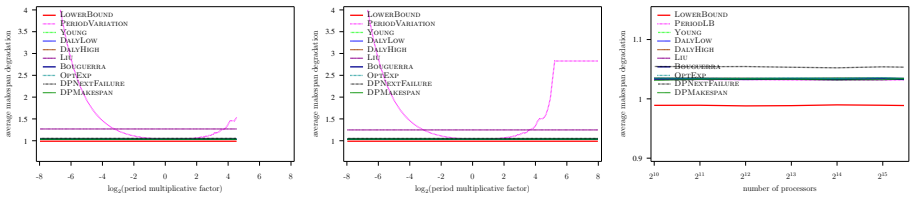

a) 1024 processors b) 45208 processors c)

Figure 31: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years), using Numerical Kernel law with $\gamma = 1$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

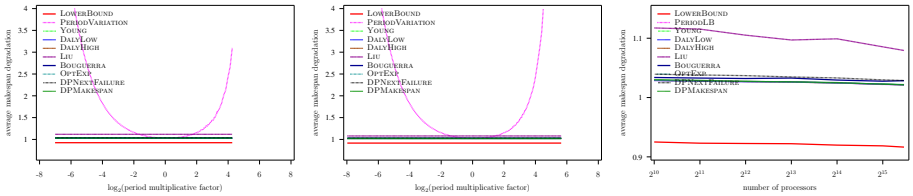

a) 1024 processors b) 45208 processors c)

Figure 32: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Numerical Kernel law with $\gamma = 10$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

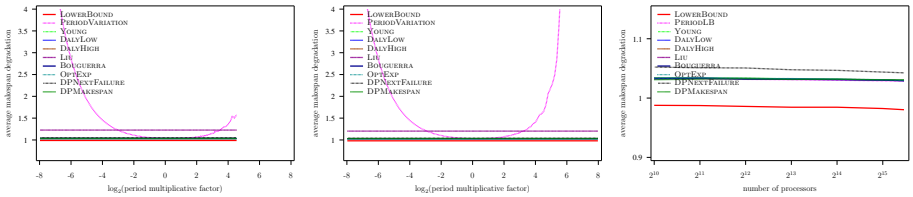

a) 1024 processors b) 45208 processors c)

Figure 33: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years), using Numerical Kernel law with $\gamma = 10$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

B.2 Exponential law, rejuvenating only the faulty processor(s)

B.2.1 Fixed checkpoint and recovery cost

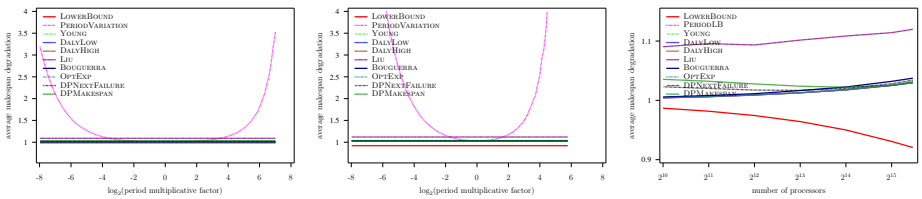

a) 1024 processors b) 45208 processors c)

Figure 34: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years) using embarrassingly parallel job, and constant overhead model

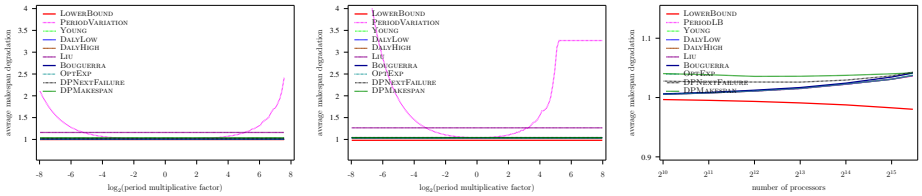

a) 1024 processors b) 45208 processors c)

Figure 35: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years) using embarrassingly parallel job, and constant overhead model

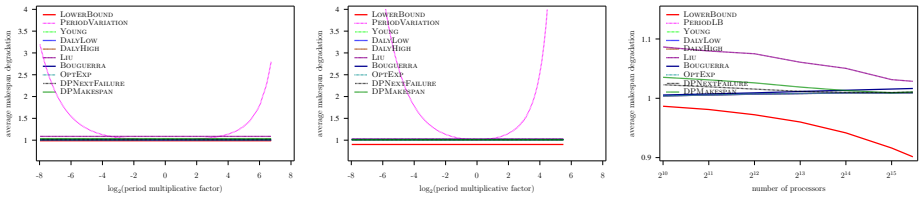

a) 1024 processors b) 45208 processors c)

Figure 36: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years) using Amdahl law with $\gamma = 10^{-4}$, and constant overhead model

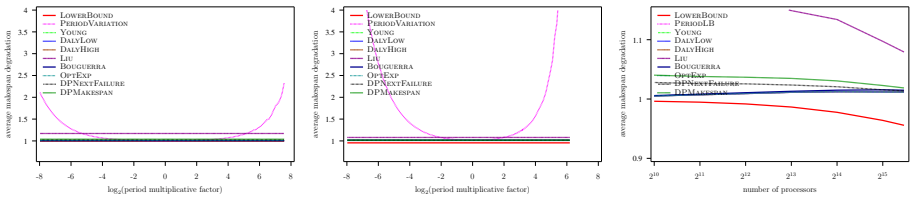

a) 1024 processors b) 45208 processors c)

Figure 37: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years) using Amdahl law with $\gamma = 10^{-4}$, and constant overhead model

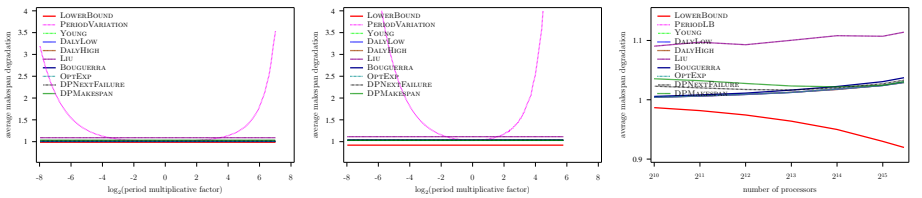

a) 1024 processors b) 45208 processors c)

Figure 38: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Amdahl law with $\gamma = 10^{-6}$, and constant overhead model

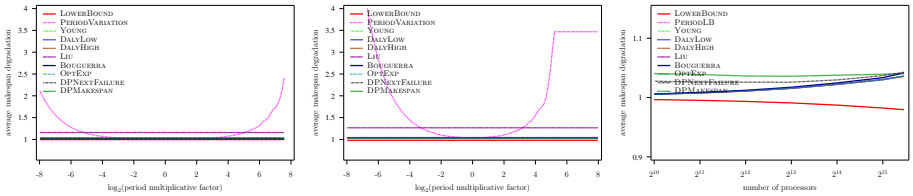

a) 1024 processors b) 45208 processors c)

Figure 39: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years), using Amdahl law with $\gamma = 10^{-6}$, and constant overhead model

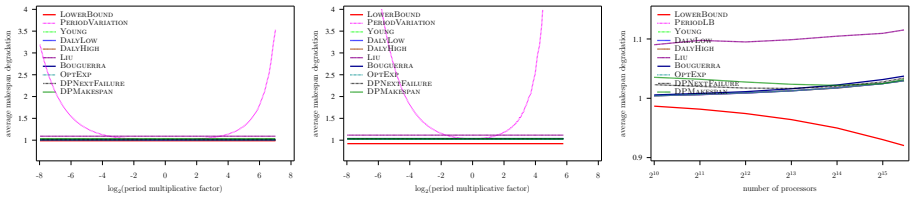

a) 1024 processors b) 45208 processors c)

Figure 40: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Numerical Kernel law with $\gamma = 0.1$, and constant overhead model

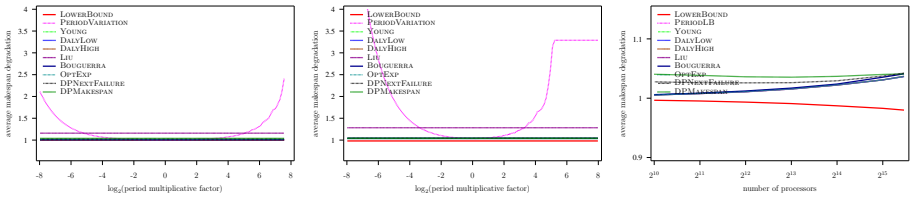

a) 1024 processors b) 45208 processors c)

Figure 41: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years), using Numerical Kernel law with $\gamma = 0.1$, and constant overhead model

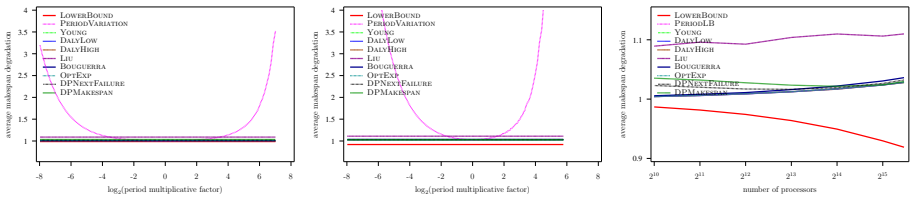

a) 1024 processors b) 45208 processors c)

Figure 42: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Numerical Kernel law with $\gamma = 1$, and constant overhead model

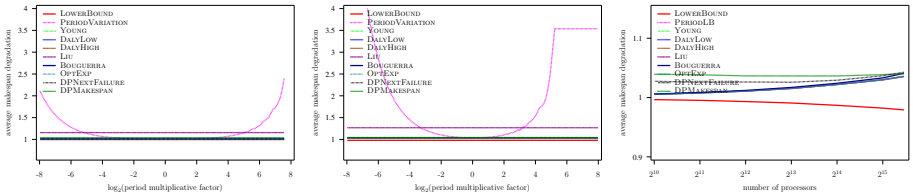

a) 1024 processors b) 45208 processors c)

Figure 43: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years), using Numerical Kernel law with $\gamma = 1$, and constant overhead model

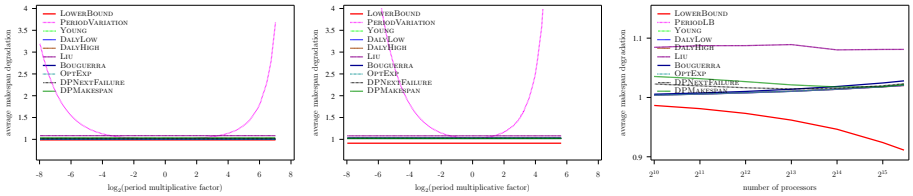

a) 1024 processors b) 45208 processors c)

Figure 44: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Numerical Kernel law with $\gamma = 10$, and constant overhead model

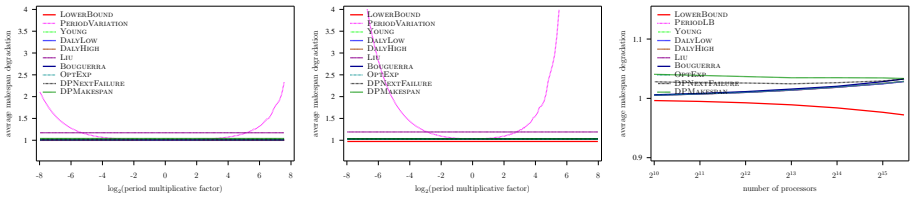

a) 1024 processors b) 45208 processors c)

Figure 45: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years), using Numerical Kernel law with $\gamma = 10$, and constant overhead model

B.2.2 Variable checkpoint and recovery cost

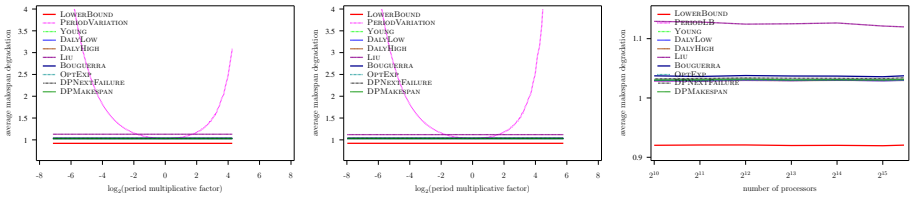

a) 1024 processors b) 45208 processors c)

Figure 46: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF = 125$ years), using embarrassingly parallel job, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

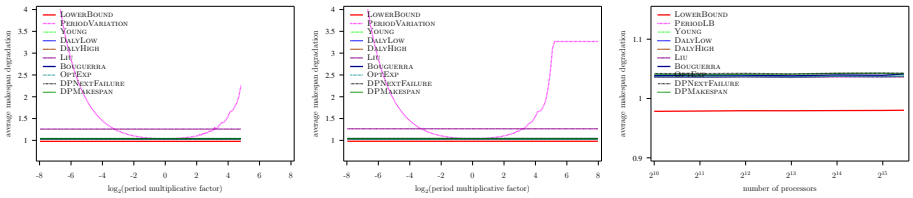

a) 1024 processors b) 45208 processors c)

Figure 47: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF = 500$ years), using embarrassingly parallel job, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

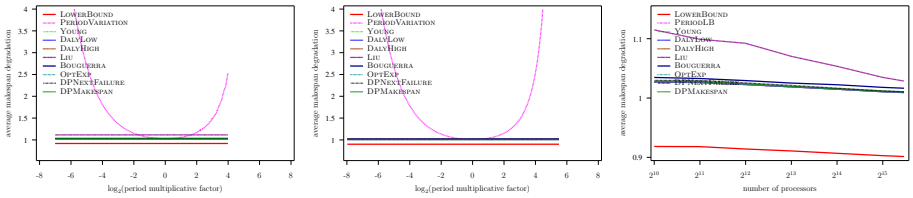

a) 1024 processors b) 45208 processors c)

Figure 48: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF = 125$ years), using Amdahl law with $\gamma = 10^{-4}$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

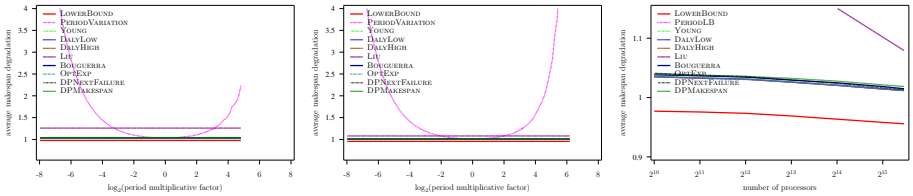

a) 1024 processors b) 45208 processors c)

Figure 49: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years), using Amdahl law with $\gamma = 10^{-4}$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

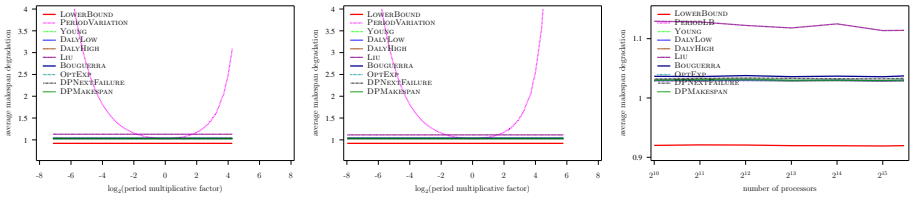

a) 1024 processors b) 45208 processors c)

Figure 50: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Amdahl law with $\gamma = 10^{-6}$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

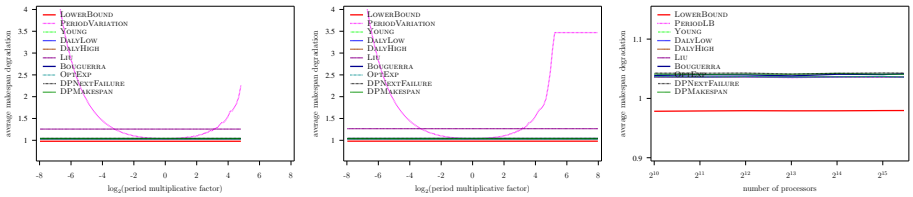

a) 1024 processors b) 45208 processors c)

Figure 51: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years), using Amdahl law with $\gamma = 10^{-6}$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

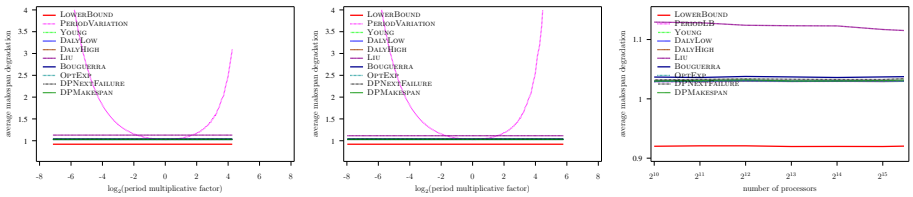

a) 1024 processors b) 45208 processors c)

Figure 52: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Numerical Kernel law with $\gamma = 0.1$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

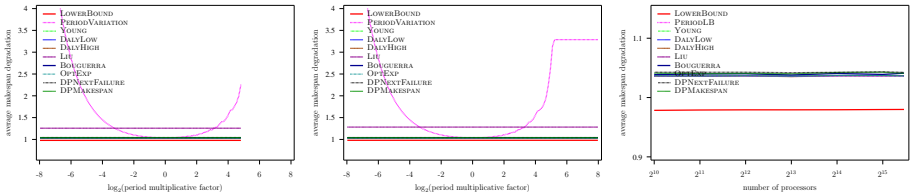

a) 1024 processors b) 45208 processors c)

Figure 53: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years), using Numerical Kernel law with $\gamma = 0.1$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

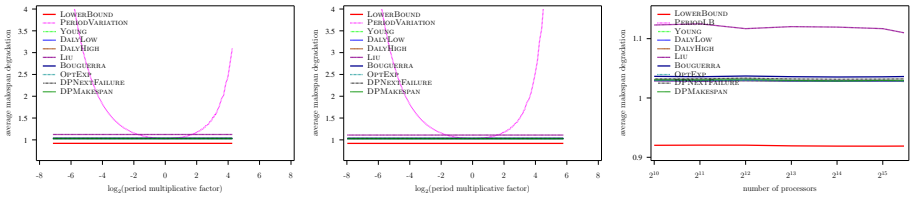

a) 1024 processors b) 45208 processors c)

Figure 54: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Numerical Kernel law with $\gamma = 1$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

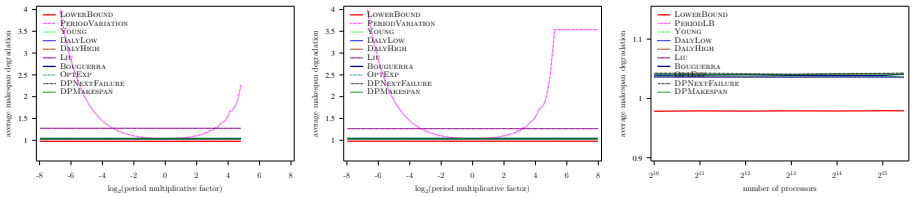

a) 1024 processors b) 45208 processors c)

Figure 55: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 500 years), using Numerical Kernel law with $\gamma = 1$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

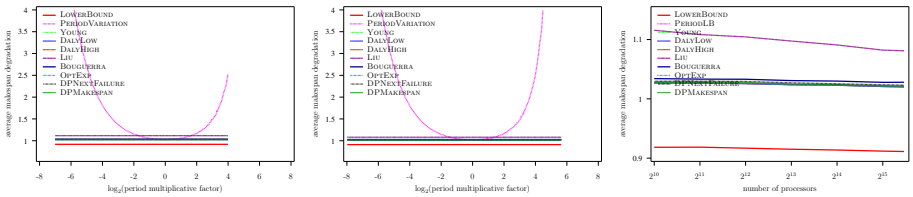

a) 1024 processors b) 45208 processors c)

Figure 56: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF$ = 125 years), using Numerical Kernel law with $\gamma = 10$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

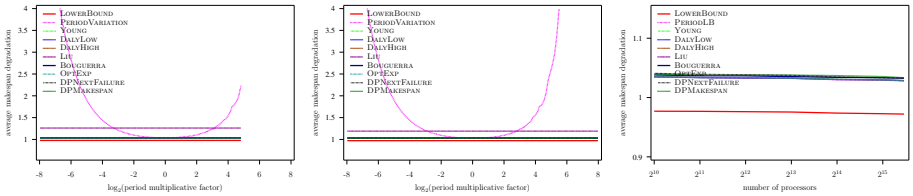

a) 1024 processors b) 45208 processors c)

Figure 57: Evaluation of the different heuristics on a Petascale platform with Exponential failures ($MTBF = 500$ years), using Numerical Kernel law with $\gamma = 10$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

B.3 Weibull law

B.3.1 Fixed checkpoint and recovery cost

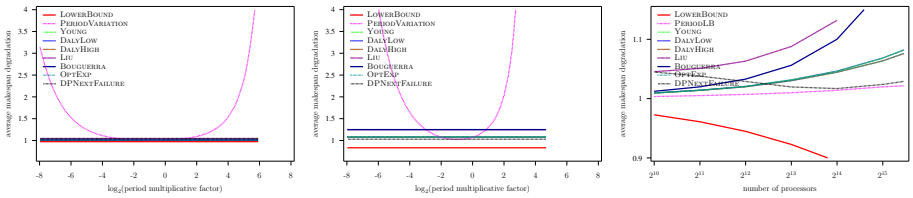

a) 1024 processors b) 45208 processors c)

Figure 58: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 125$ years), using embarrassingly parallel job, and constant overhead model

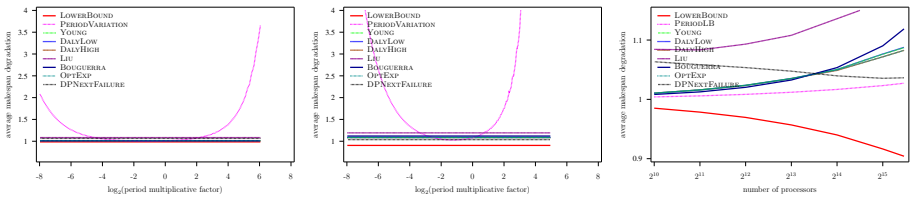

a) 1024 processors b) 45208 processors c)

Figure 59: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 500$ years), using embarrassingly parallel job, and constant overhead model

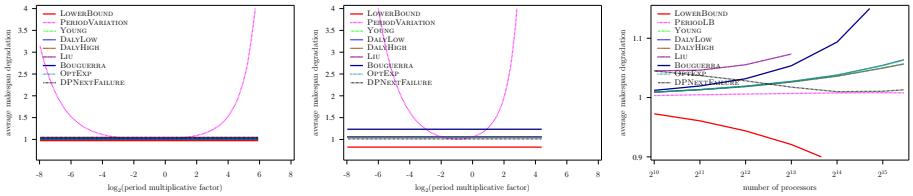

a) 1024 processors b) 45208 processors c)

Figure 60: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF$ = 125 years), using Amdahl law with $\gamma = 10^{-4}$, and constant overhead model

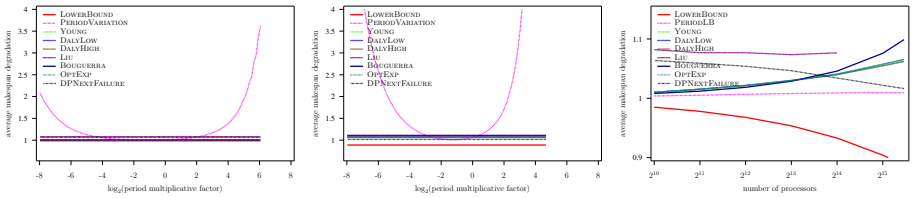

a) 1024 processors b) 45208 processors c)

Figure 61: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF$ = 500 years), using Amdahl law with $\gamma = 10^{-4}$, and constant overhead model

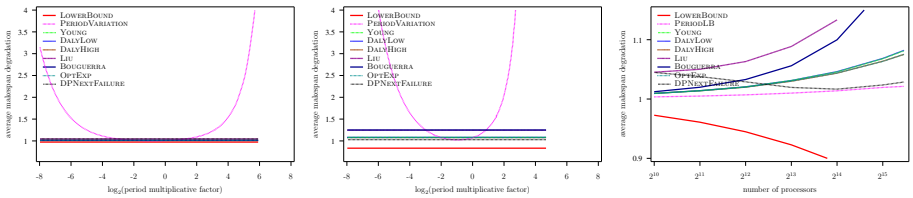

a) 1024 processors b) 45208 processors c)

Figure 62: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF$ = 125 years), using Amdahl law with $\gamma = 10^{-6}$, and constant overhead model

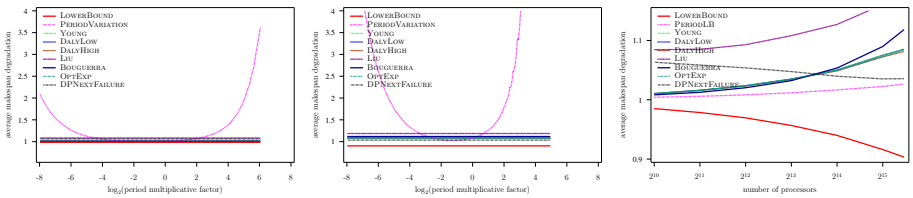

a) 1024 processors b) 45208 processors c)

Figure 63: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF$ = 500 years), using Amdahl law with $\gamma = 10^{-6}$, and constant overhead model

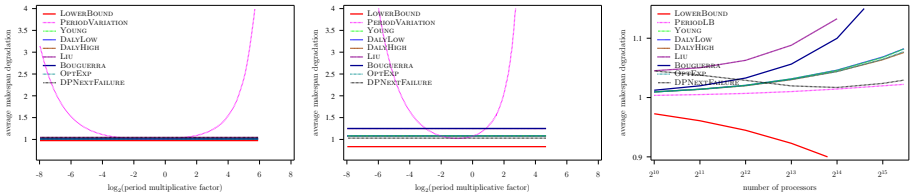

a) 1024 processors b) 45208 processors c)

Figure 64: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 125$ years), using Numerical Kernel law with $\gamma = 0.1$, and constant overhead model

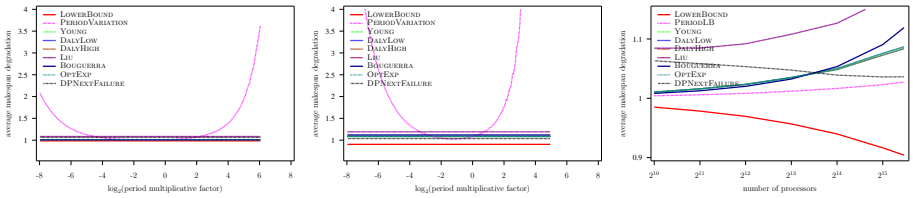

a) 1024 processors b) 45208 processors c)

Figure 65: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 500$ years), using Numerical Kernel law with $\gamma = 0.1$, and constant overhead model

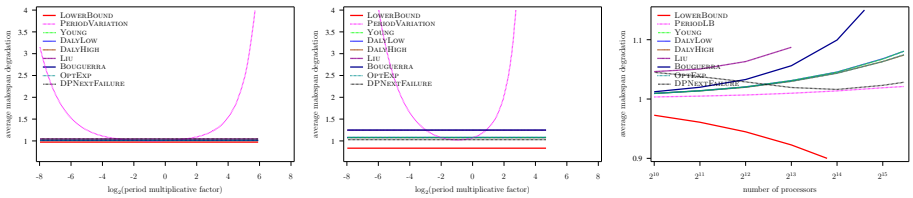

a) 1024 processors b) 45208 processors c)

Figure 66: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 125$ years), using Numerical Kernel law with $\gamma = 1$, and constant overhead model

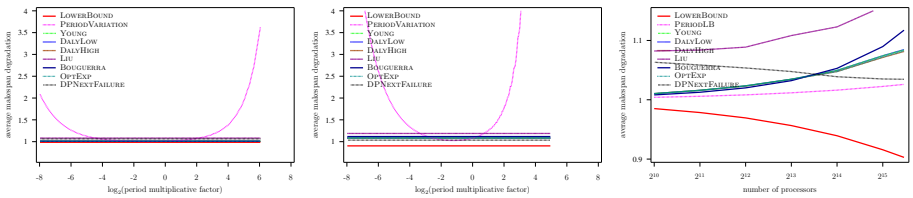

a) 1024 processors b) 45208 processors c)

Figure 67: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 500$ years), using Numerical Kernel law with $\gamma = 1$, and constant overhead model

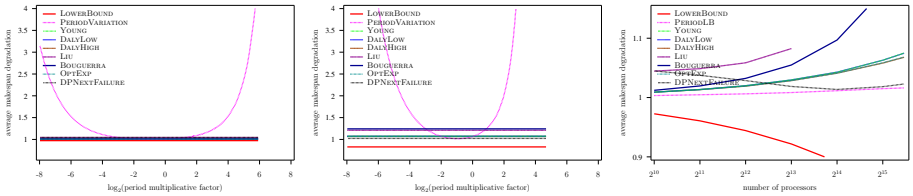

a) 1024 processors b) 45208 processors c)

Figure 68: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 125$ years), using Numerical Kernel law with $\gamma = 10$, and constant overhead model

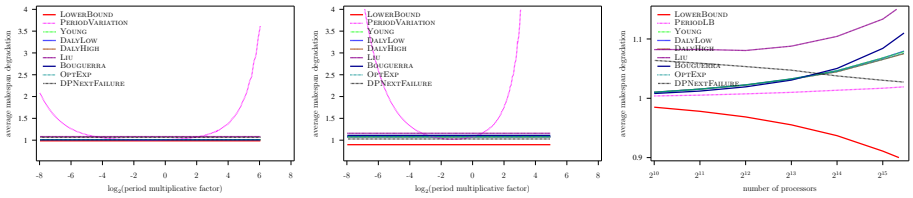

a) 1024 processors b) 45208 processors c)

Figure 69: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 500$ years), using Numerical Kernel law with $\gamma = 10$, and constant overhead model

B.3.2 Variable checkpoint and recovery cost

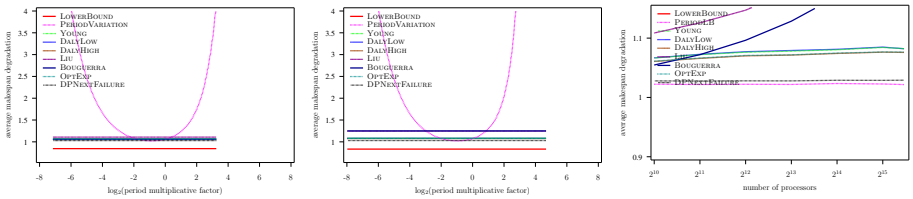

a) 1024 processors b) 45208 processors c)

Figure 70: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 125$ years), using embarrassingly parallel job, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

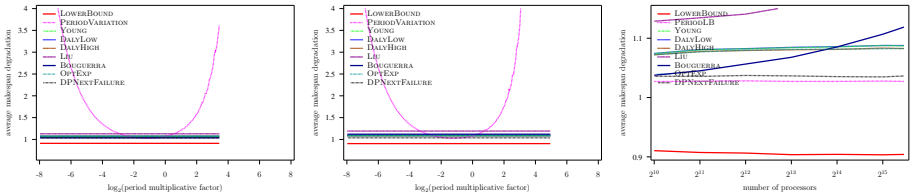

a) 1024 processors b) 45208 processors c)

Figure 71: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 500$ years), using embarrassingly parallel job, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

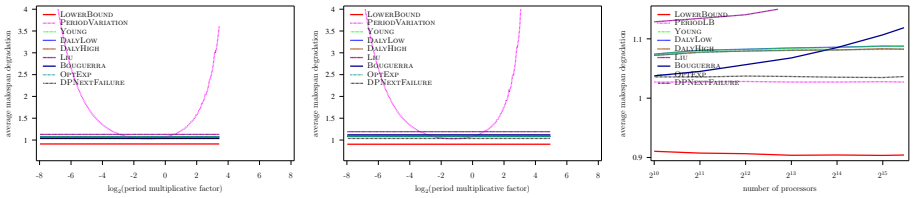

a) 1024 processors b) 45208 processors c)

Figure 72: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 500$ years), using embarrassingly parallel job, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

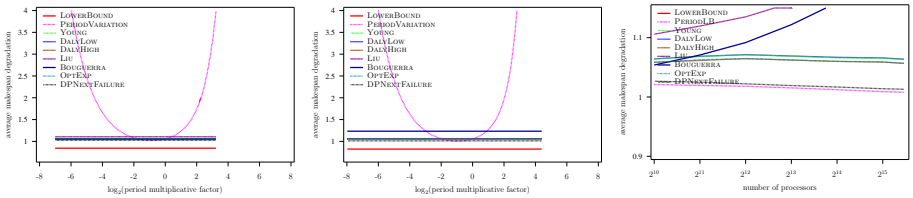

a) 1024 processors b) 45208 processors c)

Figure 73: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 125$ years), using Amdahl law with $\gamma = 10^{-4}$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

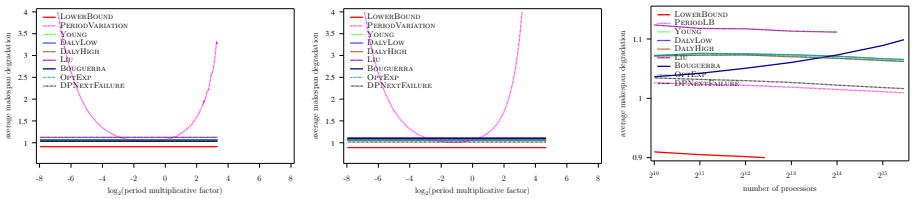

a) 1024 processors b) 45208 processors c)

Figure 74: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 500$ years), using Amdahl law with $\gamma = 10^{-4}$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

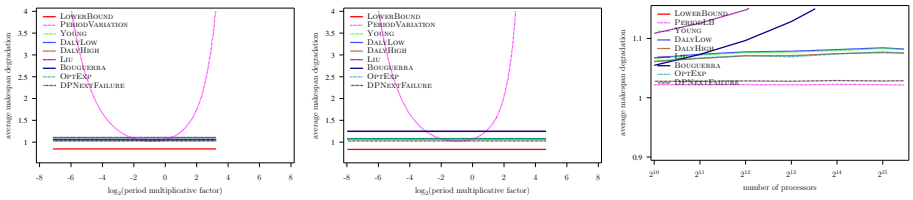

a) 1024 processors b) 45208 processors c)

Figure 75: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 125$ years), using Amdahl law with $\gamma = 10^{-6}$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

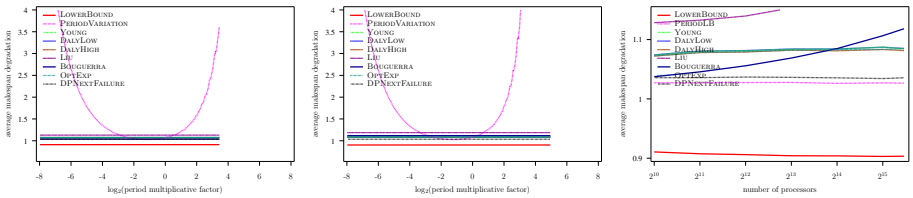

a) 1024 processors b) 45208 processors c)

Figure 76: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 500$ years), using Amdahl law with $\gamma = 10^{-6}$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

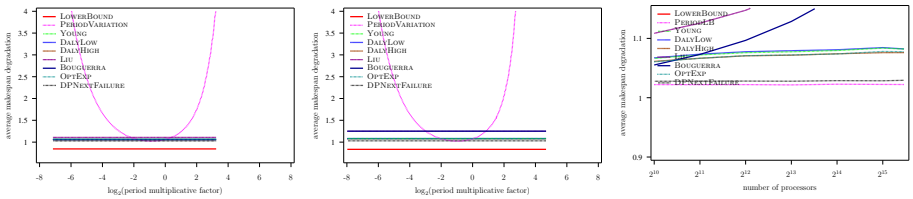

a) 1024 processors b) 45208 processors c)

Figure 77: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 125$ years), using Numerical Kernel law with $\gamma = 0.1$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

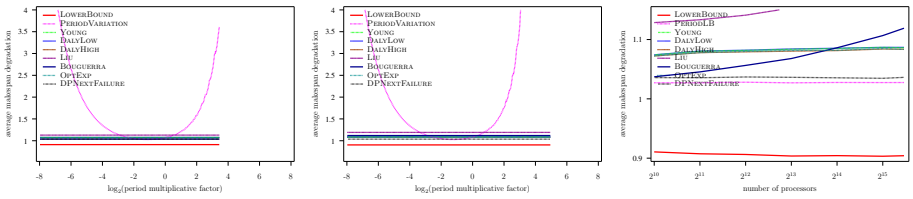

a) 1024 processors b) 45208 processors c)

Figure 78: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 500$ years), using Numerical Kernel law with $\gamma = 0.1$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

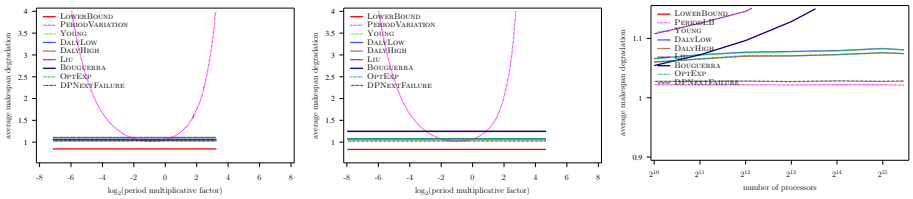

a) 1024 processors b) 45208 processors c)

Figure 79: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 125$ years), using Numerical Kernel law with $\gamma = 1$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

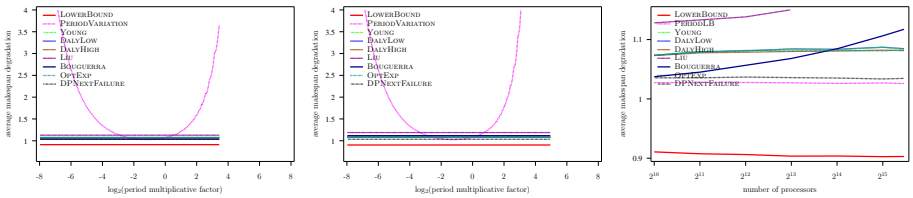

a) 1024 processors b) 45208 processors c)

Figure 80: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 500$ years), using Numerical Kernel law with $\gamma = 1$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

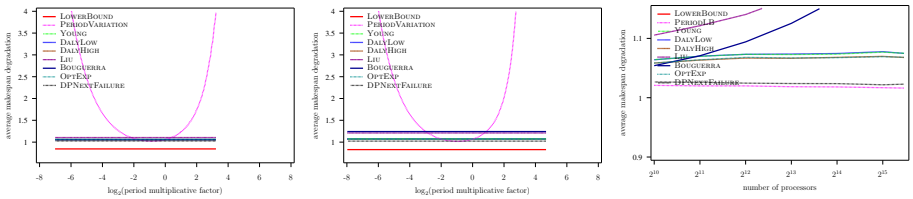

a) 1024 processors b) 45208 processors c)

Figure 81: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 125$ years), using Numerical Kernel law with $\gamma = 10$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

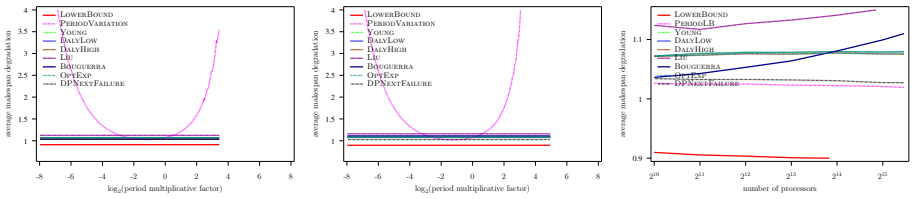

a) 1024 processors b) 45208 processors c)

Figure 82: Evaluation of the different heuristics on a Petascale platform with Weibull failures ($MTBF = 500$ years), using Numerical Kernel law with $\gamma = 10$, and variable overhead model ($C(p) = 600\frac{45208}{p}$)

C Appendix : Results for Exascale platforms

C.1 Exponential law, rejuvenating all processors

C.1.1 Fixed checkpoint and recovery cost

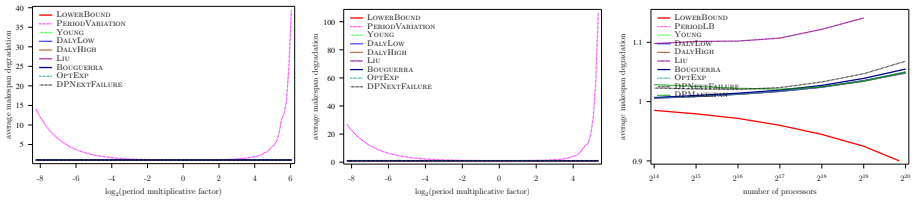

a) 262144 processors b) 1048576 processors c)

Figure 83: Evaluation of the different heuristics on a Exascale platform with Exponential failures ($MTBF = 1250$ years), using embarrassingly parallel job, and constant overhead model

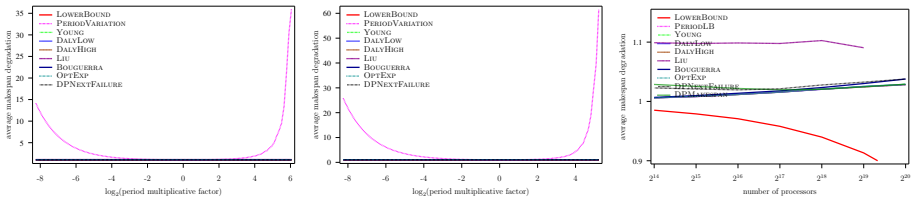

a) 262144 processors b) 1048576 processors c)

Figure 84: Evaluation of the different heuristics on a Exascale platform with Exponential failures ($MTBF = 1250$ years), using Amdahl law with $\gamma = 10^{-6}$, and constant overhead model

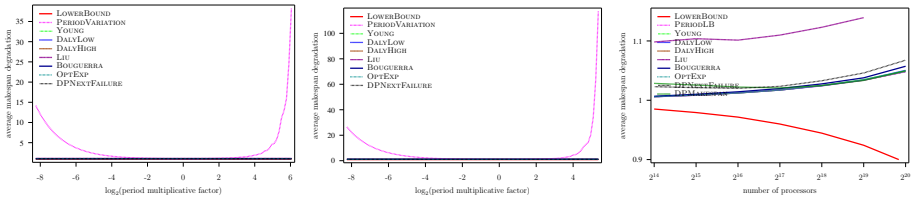

a) 262144 processors b) 1048576 processors c)

Figure 85: Evaluation of the different heuristics on a Exascale platform with Exponential failures ($MTBF = 1250$ years), using Numerical Kernel law with $\gamma = 0.1$, and constant overhead model

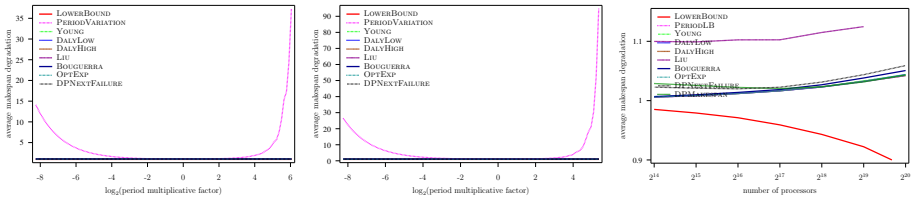

a) 262144 processors b) 1048576 processors c)

Figure 86: Evaluation of the different heuristics on a Exascale platform with Exponential failures ($MTBF$ = 1250 years), using Numerical Kernel law with $\gamma = 1$, and constant overhead model

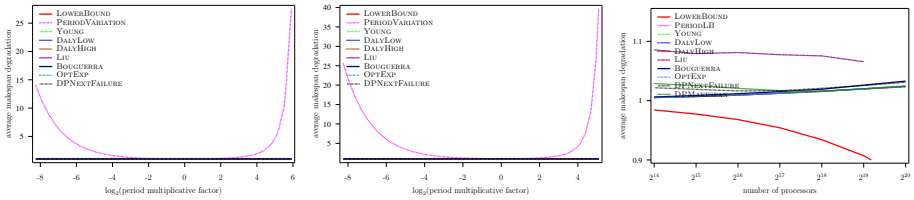

a) 262144 processors b) 1048576 processors c)

Figure 87: Evaluation of the different heuristics on a Exascale platform with Exponential failures ($MTBF$ = 1250 years), using Numerical Kernel law with $\gamma = 10$, and constant overhead model

C.2 Exponential law, rejuvenating only the faulty processor(s)

C.2.1 Fixed checkpoint and recovery cost

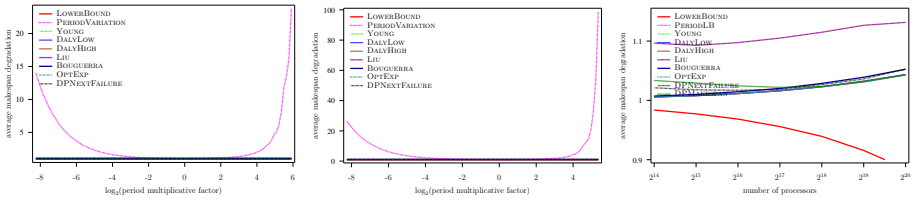

a) 262144 processors b) 1048576 processors c)

Figure 88: Evaluation of the different heuristics on a Exascale platform with Exponential failures ($MTBF$ = 1250 years), using embarrassingly parallel job, and constant overhead model

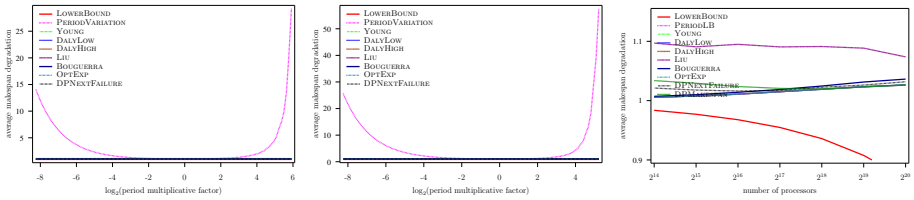

a) 262144 processors b) 1048576 processors c)

Figure 89: Evaluation of the different heuristics on a Exascale platform with Exponential failures ($MTBF = 1250$ years), using Amdahl law with $\gamma = 10^{-6}$, and constant overhead model

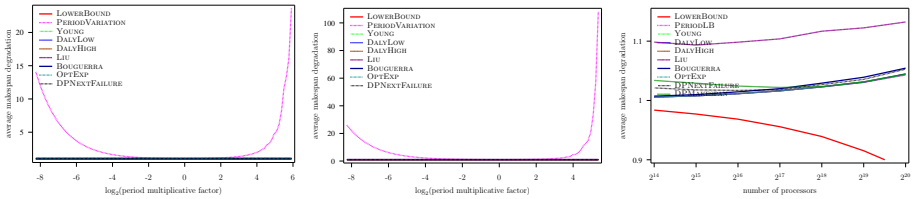

a) 262144 processors b) 1048576 processors c)

Figure 90: Evaluation of the different heuristics on a Exascale platform with Exponential failures ($MTBF = 1250$ years), using Numerical Kernel law with $\gamma = 0.1$, and constant overhead model

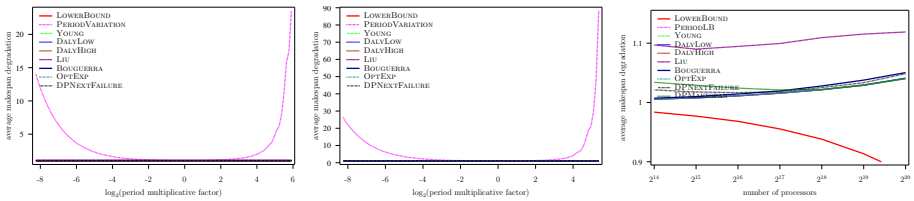

a) 262144 processors b) 1048576 processors c)

Figure 91: Evaluation of the different heuristics on a Exascale platform with Exponential failures ($MTBF = 1250$ years), using Numerical Kernel law with $\gamma = 1$, and constant overhead model

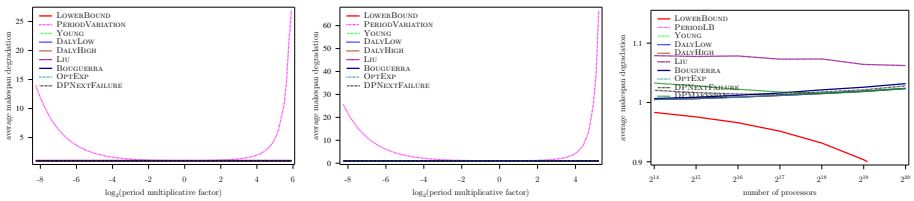

a) 262144 processors b) 1048576 processors c)

Figure 92: Evaluation of the different heuristics on a Exascale platform with Exponential failures ($MTBF = 1250$ years), using Numerical Kernel law with $\gamma = 10$, and constant overhead model

C.3 Weibull law

C.3.1 Fixed checkpoint and recovery cost

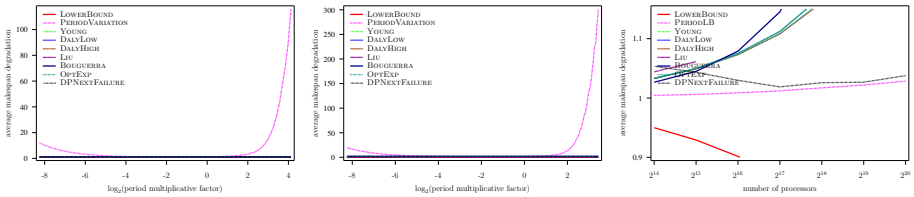

a) 262144 processors b) 1048576 processors c)

Figure 93: Evaluation of the different heuristics on a Exascale platform with Weibull failures ($MTBF = 1250$ years), using embarrassingly parallel job, and constant overhead model

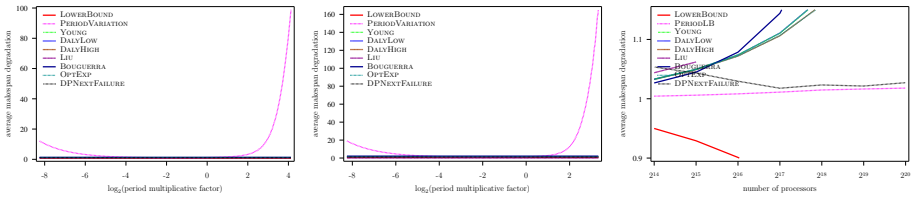

a) 262144 processors b) 1048576 processors c)

Figure 94: Evaluation of the different heuristics on a Exascale platform with Weibull failures ($MTBF = 1250$ years), using Amdahl law with $\gamma = 10^{-6}$, and constant overhead model

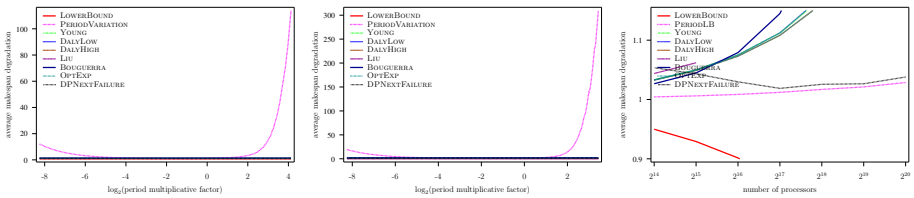

a) 262144 processors b) 1048576 processors c)

Figure 95: Evaluation of the different heuristics on a Exascale platform with Weibull failures ($MTBF = 1250$ years), using Numerical Kernel law with $\gamma = 0.1$, and constant overhead model

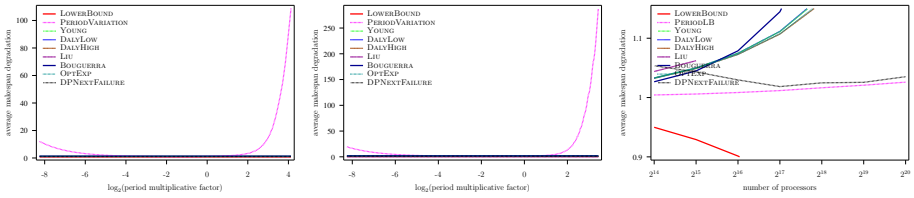

a) 262144 processors b) 1048576 processors c)

Figure 96: Evaluation of the different heuristics on a Exascale platform with Weibull failures ($MTBF = 1250$ years), using Numerical Kernel law with $\gamma = 1$, and constant overhead model

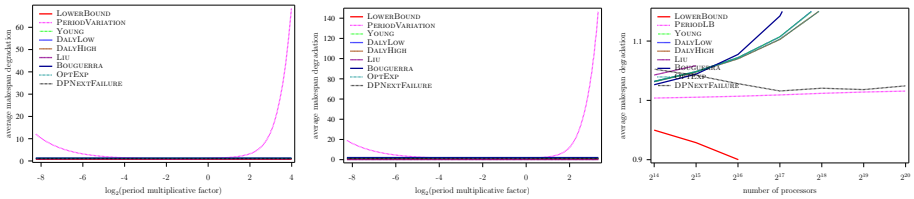

a) 262144 processors b) 1048576 processors c)

Figure 97: Evaluation of the different heuristics on a Exascale platform with Weibull failures ($MTBF = 1250$ years), using Numerical Kernel law with $\gamma = 10$, and constant overhead model

D Appendix: results for fixed heuristic according to application model variation

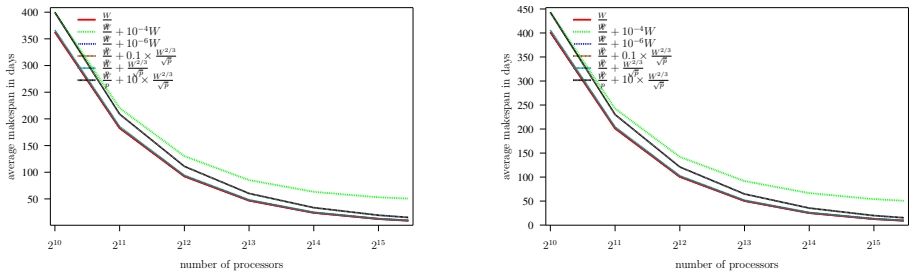

a) b)

Figure 98: Evolution of the makespan for OptExp as a function of the platform size for the different application profiles and a constant (a)) or platform-dependent (b)) checkpoint cost (exponential law, $MTBF = 125$ years).

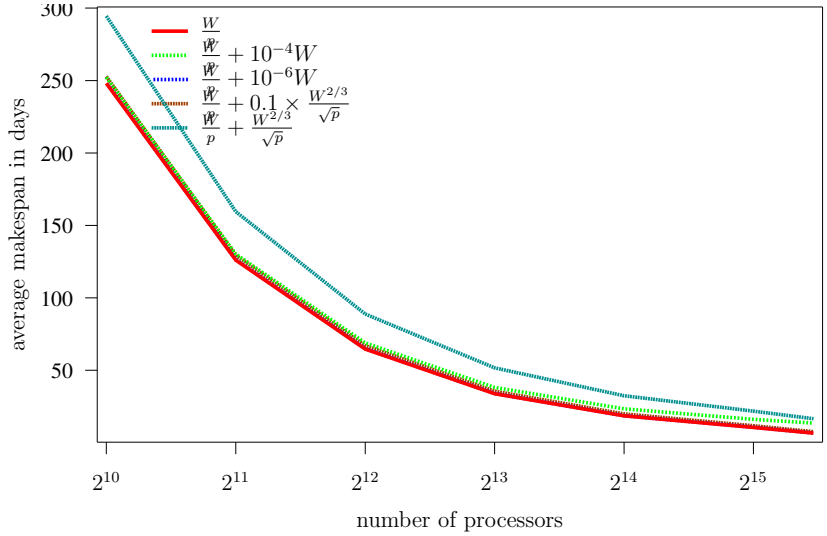

a)

Figure 99: Evolution of the makespan for DPNextFailure as a function of the platform size for the different application profiles and a constant checkpoint cost (Weibull law, $MTBF = 1250$ years).

E Appendix : Results for log-based failures

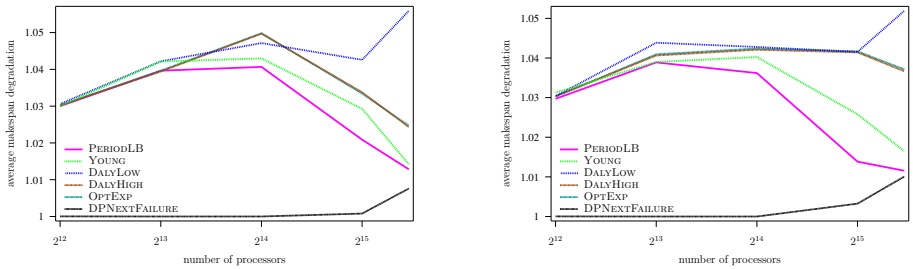

a) b)

Figure 100: Evaluation of the different heuristics on a Petascale platform with failures based on the failure log of LANL cluster 18 (a)) and 19 (b)).

Centre de recherche INRIA Grenoble – Rhône-Alpes
655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399