

# Using group replication for resilience on exascale systems

Marin Bougeret, Henri Casanova, Yves Robert, Frédéric Vivien, Dounia Zaidouni

# ▶ To cite this version:

Marin Bougeret, Henri Casanova, Yves Robert, Frédéric Vivien, Dounia Zaidouni. Using group replication for resilience on exascale systems. [Research Report] RR-7876, 2012. hal-00668016v1

# HAL Id: hal-00668016 https://inria.hal.science/hal-00668016v1

Submitted on 8 Feb 2012 (v1), last revised 28 Jun 2013 (v2)

**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.



# Using group replication for resilience on exascale systems

Marin Bougeret, Henri Casanova, Yves Robert, Frédéric Vivien, Dounia Zaidouni

RESEARCH REPORT N° 7876 February 2012 Project-Team ROMA



# Using group replication for resilience on exascale systems

Marin Bougeret\*, Henri Casanova<sup>†</sup>, Yves Robert<sup>‡§¶</sup>, Frédéric Vivien<sup>||‡</sup>, Dounia Zaidouni<sup>||‡</sup>

Project-Team ROMA

Research Report $\,$ n° 7876 — February 2012 — 41 pages

- \* LIRMM Montpellier, France
- $^\dagger$ University of Hawai'i at Mānoa, Honolulu, USA
- <sup>‡</sup> LIP, Ecole Normale Supérieure de Lyon, France
- <sup>§</sup> University of Tennessee Knoxville, USA
- ¶ Institut Universitaire de France.
- || INRIA, Lyon, France

#### RESEARCH CENTRE GRENOBLE – RHÔNE-ALPES

Inovallée 655 avenue de l'Europe Montbonnot 38334 Saint Ismier Cedex **Abstract:** High performance computing applications must be resilient to faults, which are common occurrences especially in post-petascale settings. The traditional fault-tolerance solution is checkpoint-recovery, by which the application saves its state to secondary storage throughout execution and recovers from the latest saved state in case of a failure. An off studied research question is that of the optimal checkpointing strategy: when should state be saved? Unfortunately, even using an optimal checkpointing strategy, the checkpointing frequency must increase as platform scale increases, leading to higher checkpointing overhead. This overhead precludes high parallel efficiency for large-scale platforms, thus mandating other more scalable fault-tolerance mechanisms. One such mechanism is replication, which can be used in addition to checkpoint-recovery. Using replication, multiple processors perform the same computation so that a processor failure does not necessarily imply application failure. While at first glance replication may seem wasteful, it may be significantly more efficient than using solely checkpoint-recovery at large scale. In this work we investigate a simple approach where entire application instances are replicated. We provide a theoretical study of checkpoint-recovery with replication in terms of expected application execution time, under an exponential distribution of failures. We design dynamic-programming based algorithms to define checkpointing dates that work under any failure distribution. We also conduct simulation experiments assuming that failures follow Exponential or Weibull distributions, the latter being more representative of real-world systems, and using failure logs from production clusters. Our results show that replication is useful in a variety of realistic application and checkpointing cost scenarios for future exascale platforms.

Key-words: Fault-tolerance, replication, checkpointing, parallel job, Weibull, exascale

# La réplication pour l'amélioration de la résilience des applications sur systèmes exascales

Résumé : Les applications de calcul à haute-performance doivent être résiliantes aux pannes, car les pannes ne seront pas des évènements rares sur les plates-formes post-petascales. La tolérance aux pannes est traditionnellement réalisée par un mécanisme d'enregistrement et redémarrage, au moyen duquel l'application sauve son état sur un système de stockage secondaire et, en cas de panne, redémarre à partir du dernier état sauvegardé. Une question souvent étudiée est celle de la stratégie de sauvegarde optimale: quand l'état doit-il être sauvé ? Malheureusement, même quand on utilise une stratégie de sauvegarde optimale, la fréquence de sauvegarde doit augmenter avec la taille de la plateforme, augmentant mécaniquement le coût des sauvegardes. Ce coût interdit d'obtenir une très bonne efficacité sur des plates-formes à très large échelle, et requiert d'utiliser d'autres mécanismes de tolérance aux pannes, qui passent mieux à l'échelle. Un mécanisme potentiel est la réplication, qui peut être utilisée conjointement avec une solution de sauvegarde et redémarrage. Avec la réplication, plusieurs processeurs exécutent le même calcul de sorte que la panne de l'un d'entre eux n'implique pas nécessairement une panne pour l'application. Alors qu'à première vue une telle approche gaspille des ressources, la réplication peut être significativement plus efficace que la seule mise en œuvre de techniques de sauvegarde et redémarrage sur des plates-formes à très grande échelle. Dans la présente étude nous considérons une approche simple où une application toute entière est répliquée. Nous fournissons une étude théorique d'un schéma d'exécution avec réplication lorsque la distribution des pannes suit une loi exponentielle. Nous proposons des algorithmes de détermination des dates de sauvegarde quand la distribution des pannes suit une loi quelconque. Nous menons aussi une étude expérimentale, au moyen de simulations, basée sur une distribution de pannes suivant une loi exponentielle, de Weibull (ce qui est plus représentatif des systèmes réels), ou tirée de logs de clusters utilisés en production. Nos résultats montrent que la réplication est bénéfique pour un ensemble de modèles d'applications et de coût de sauvegardes réalistes, dans le cadre des futures plates-formes exascales.

**Mots-clés :** Tolérance aux pannes, réplication, checkpoint, tâche parallèle, Weibull, exascale

# 1 Introduction

As plans are made for deploying post-petascale high performance computing (HPC) systems [10, 22], solutions need to be developed to ensure resilience to failures that occur because not all faults can be automatically detected and corrected in hardware. For instance, the 224,162-core Jaguar platform is reported to experience on the order of 1 failure per day [20, 2], and its scale is modest compared to platforms in the plans for the next decade. For applications that enroll large numbers of, or perhaps all, processors a failure is the common case rather than the exception. One can recover from a failure by resuming execution from a previously saved fault-free execution state, or *checkpoint*. Checkpoints are saved to resilient storage throughout execution (usually periodically). More frequent checkpoints lead to less loss when a failure occurs but to higher overhead during fault-free execution. A *checkpointing strategy* specifies when checkpoints should be taken.

A large literature is devoted to developing efficient checkpointing strategies, i.e., ones that minimize expected job execution time, including both theoretical and practical efforts. The former typically rely on assumptions regarding the probability distributions of times to failure of the processors (e.g., Exponential, Weibull), while the latter rely on simulations driven by failure datasets obtained on real-world platforms. In a previous paper [4], we have made several contributions in this context, including optimal solutions for Exponential failures and dynamic programming solutions in the general case.

A major issue with checkpoint-recovery is scalability: the necessary checkpoint frequency for tolerating failures in large-scale platforms is so large that processors spend more time saving state than computing. It is thus expected that future platforms will lead to unacceptably low parallel efficiency if only checkpoint-recovery is used, no matter how good the checkpointing strategy. Consequently, additional mechanisms must be used. In this work we focus on *replication*: several processors perform the same computation synchronously, so that a fault on one of these processors does not lead to an application failure. Replication is an age-old fault-tolerant technique, but it has gained traction in the HPC context only relatively recently. While replication wastes compute resources in fault-free executions, it can alleviate the poor scalability of checkpoint-recovery.

Consider a parallel application that is *moldable*, meaning that it can be executed on an arbitrary number of processors, which each processor running one application process. In our *group replication* approach, multiple application instances are executed. One could, for instance, execute 2 distinct *n*-process application instances on a 2n-processor platform. Each instance runs at a smaller scale, meaning that it has better parallel efficiency than a single 2n-process instance due to a lower checkpointing frequency. Furthermore, once an instance saves a checkpoint, the other instance can use this checkpoint immediately. Given the above, our contributions in this work are:

- A theoretical analysis of the optimal number of processors to use for a checkpoint-recovery execution of a parallel application, for various parallel workload models for Exponential failure distributions;
- An effective approach for group replication, with a theoretical analysis bounding expected execution time for Exponential failure distribution, and several dynamic programming solutions working for general failure

distributions;

• Extensive simulations showing that group replication can indeed lower application running times, and that some of our proposed strategies deliver good performance.

The rest of this paper is organized as follows. Section 2 discusses related work. Section 3 defines the theoretical framework and states key assumptions. Section 4 discusses the optimal number of processors for a checkpoint-recovery execution of a parallel application under an Exponential distribution of failures. Section 5 provides several approaches for group replication, and the theoretical analysis of one of them. Section 6 describes the experimental methodology and Section 7 presents simulation results. Finally, Section 8 concludes with a summary of our findings and with future perspectives.

# 2 Related work

Checkpointing policies have been widely studied in the literature. In [9], Daly studies periodic checkpointing policies for Exponentially distributed failures, generalizing the well-known bound obtained by Young [28]. Daly extended his work in [15] to study the impact of sub-optimal checkpointing periods. In [25], the authors develop an "optimal" checkpointing policy, based on the popular assumption that optimal checkpointing must be periodic. In [6], Bouguerra et al. prove that the optimal checkpointing policy is periodic when checkpointing and recovery overheads are constant, for either Exponential or Weibull failures. But their results rely on the unstated assumption that all processors are rejuvenated after each failure and after each checkpoint. In our recent work [4], we have shown that this assumption is unreasonable for Weibull failures. We have developed optimal solutions for Exponential failures and dynamic programming solutions for any failure distribution, demonstrating performance improvements over checkpointing approaches proposed in the literature in the case of Weibull and log-based failures. The Weibull distribution is recognized as a reasonable approximation of failures in real-world systems [14, 24]. The work in this paper relates to checkpointing policies in the sense that we study a replication mechanism that is used as an addition to checkpointing. Part of our results build on the algorithms and results in [4].

In spite of all the above advances, several studies have questioned the feasibility of pure checkpoint-recovery for large-scale systems (see [12] for a discussion of this issue and for references to such studies). In this work, we study the use of replication as a mechanism complementary to checkpoint-recovery. Replication has long been used as a fault-tolerance mechanism in distributed systems [13] and more recently in the context of volunteer computing [17]. The idea to use replication together with checkpoint-recovery has been studied in the context of grid computing [27].One concern about replication in HPC is the induced resource waste. However, given the scalability limitations of pure checkpoint-recovery, replication has recently received more attention in the HPC literature [23, 29, 11].

In this work we study "group replication," by which multiple application instances are executed on different groups of processors. An orthogonal approach, "process replication," was recently studied by Ferreira et al. [12] in the context of MPI applications. To achieve fault-tolerance, each MPI process is replicated in a way that is transparent to the application developer. While this approach can lead to good fault-tolerance, one of its drawbacks is that it increases the number and the volume of communications. Let  $V_{tot}$  be the total volume of inter-processor communications for a traditional execution. With process replication using g replicas per replica-groups, each original communication now involves g sources and g destinations, hence the total communication volume becomes  $V_{tot} \times g^2$ . Instead, with group replication using g groups, each original communication takes place g times, hence the total communication volume increases only to  $V_{tot} \times g$ . Another drawback of process replication is that it requires the use of a customized MPI library (such as the prototype developed by the authors in [12]). By contrast, group replication is completely agnostic to the parallel runtime system and thus does not even require MPI. Nevertheless, even for MPI applications, group replication provides an out of the box faulttolerance solution that can be used until process replication possibly becomes a mainstream feature in MPI implementations.

# 3 Framework

We consider the execution of a tightly-coupled parallel application, or job, on a platform composed of p processors. 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 applies regardless of the granularity of the platform. We assume that system-level checkpoint-recovery is enabled.

The job must complete  $\mathcal{W}$  units of (divisible) work, which can be split arbitrarily into separate *chunks*. The job can execute on any number  $q \leq p$  processors. Letting  $\mathcal{W}(q)$  be the time required for a failure-free execution on q processors, we use three models:

- Perfectly parallel jobs: W(q) = W/q.
- Generic parallel jobs:  $W(q) = (1 \gamma)W/q + \gamma W$ . As in Amdahl's law [1],  $\gamma < 1$  is the fraction of the work that is inherently sequential.
- Numerical kernels:  $\mathcal{W}(q) = \mathcal{W}/q + \gamma \mathcal{W}^{2/3}/\sqrt{q}$ . 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],  $q = r^2$  and each processor receives 2r blocks of size  $N^2/r^2$  during the execution. Here  $\gamma$  is the communication-to-computation ratio of the platform.

Each participating processor is subject to *failures*. A failure causes a *down*time period of the failing processor, of duration D. When a processor fails, the whole execution is stopped, and all processors must recover from the previous checkpoint. We let C(q) denote the time needed to perform a checkpoint, and R(q) the time to perform a recovery. The downtime accounts for software rejuvenation (i.e., rebooting [16, 8]) 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 recovery period corresponds to the time needed to restore the last checkpoint. Assuming that the application's memory footprint is V bytes, with each processor holding V/q bytes, we consider two scenarios:

• Proportional overhead:  $C(q) = R(q) = \alpha V/q = C/q$  with  $\alpha$  some constant, for cases where the bandwidth of the network card/link at each processor is the I/O bottleneck.

• Constant overhead:  $C(q) = R(q) = \alpha V = C$  with  $\alpha$  some constant, for cases where the bandwidth to/from the resilient storage system is the I/O bottleneck.

We assume coordinated checkpointing [26], meaning that no message logging/replay is needed when recovering from failures. 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 assume that the parallel job is tightly coupled, meaning that all q processors operate synchronously throughout the job execution. These processors execute the same amount of work W(q) in parallel, chunk by chunk. The total time (on one processor) to execute a chunk of size  $\omega$ , and then checkpointing it, is  $\omega + C(q)$ . Finally, we assume that failure arrivals at all processors are independent and identically distributed (*iid*).

# 4 Optimal number of processors for execution

Let  $\mathbb{E}(q)$  be the expectation of the execution time, or makespan, when using q processors, and  $q_{opt}$  the value of q that minimizes  $\mathbb{E}(q)$ . Is it true that the optimal solution is to use all processors, i.e.,  $q_{opt} = p$ ? If not, what can we say about the value of  $q_{opt}$ ? This question was partially and empirically addressed in [25], via experiments for 4 MPI applications for up to 35 processors. Our approach here is radically different since we target large-scale platforms and seek theoretical results in the form of optimal solutions. The main objective of this section is to show analytically that, for Exponential failures,  $\mathbb{E}(q)$  may reach its minimum for some finite value of q (implying that  $q_{opt}$  is not necessarily equal to p).

Assume that failure inter-arrival times follow an Exponential distribution with parameter  $\lambda$ . In our recent work [4], we have shown that the optimal strategy to minimize the expected makespan  $\mathbb{E}(q)$  is to split  $\mathcal{W}$  into  $K^* = \max(1, \lfloor K_0(q) \rfloor)$  or  $K^* = \lceil K_0(q) \rceil$  same-size chunks, whichever leads to the smaller value, where  $K_0(q) = \frac{q\lambda \mathcal{W}(q)}{1+\mathbb{L}(-e^{-q\lambda C(q)-1})}$  is the optimal (non integer) number of chunks.  $\mathbb{L}$  denotes the Lambert function, defined as  $\mathbb{L}(z)e^{\mathbb{L}(z)} = z$ . This result shows that the optimal strategy is periodic and that the optimal expectation of the makespan is:

$$\mathbb{E}^*(q) = K^*(q) \left(\frac{1}{q\lambda} + \mathbb{E}(X_D(q))\right) e^{q\lambda R(q)} \left(e^{\frac{q\lambda W(q)}{K^*(q)} + q\lambda C(q)} - 1\right)$$
(1)

where  $\mathbb{E}(X_D(q))$  denotes the expectation of the downtime. It turns out that, although we can compute the optimal number of chunks (and thus the chunk size), we cannot compute  $\mathbb{E}^*(q)$  analytically because  $\mathbb{E}(X_D(q))$  is difficult to compute. This is because a processor can fail while another one is down, thus prolonging the downtime. With a single processor (q = 1),  $X_D(q)$  has constant value D, but with several processors there could be cascading downtimes. It turns out that we can compute the following lower and upper bounds for  $\mathbb{E}(X_D(q))$ :

**Proposition 1.** Let  $X_D(q)$  denote the downtime of a group of q processors. Then

$$D \le \mathbb{E}(X_D(q)) \le \frac{e^{(q-1)\lambda D} - 1}{(q-1)\lambda}$$
(2)

RR n° 7876

*Proof.* In [4], we have shown that the optimal expectation of the makespan is computed as:

$$\mathbb{E}^*(q) = K^*(q) \left(\frac{1}{q\lambda} + \mathbb{E}(T_{rec}(q))\right) \left(e^{\frac{q\lambda\mathcal{W}(q)}{K^*(q)} + q\lambda C(q)} - 1\right)$$
(3)

where  $\mathbb{E}(T_{rec}(q))$  denotes the expectation of the recovery time, i.e., the time spent recovering from failure during the computation of a chunk. All chunks have the same recovery time because they all have the same size and because of the memoryless property of the Exponential distribution. It turns out that although we can compute the optimal number of chunks (and thus the chunk size), we cannot compute  $\mathbb{E}^*(q)$  analytically because  $\mathbb{E}(T_{rec}(q))$  is difficult to compute. We write the following recursion:

$$T_{rec}(q) = \begin{cases} X_D(q) + R(q) & \text{if no processor fails} \\ & \text{during } R(q) \text{ units of time,} \\ & X_D(q) + T_{lost}(R(q)) + T_{rec}(q) & \text{otherwise.} \end{cases}$$
(4)

 $X_D(q)$  is the downtime of a group of q processors, that is the time between the first failure of one of the processors and the first time at which all of them are available (accounting for the fact a processor can fail while another one is down, thus prolonging the downtime).  $T_{lost}(R(q))$  is the amount of time spent computing by these processors before a first failure, knowing that the next failure occurs within the next R(q) units of time. In other terms, it is the compute time that is wasted because checkpoint recovery was not completed. The time until the next failure of a group of q processors is the minimum of q *iid* Exponentially distributed variables, and is thus Exponential with parameter  $q\lambda$ . We can compute  $\mathbb{E}(T_{lost}(R(q))) = \frac{1}{q\lambda} - \frac{R(q)}{e^{q\lambda R(q)} - 1}$  (see [4] for details). Plugging this value into Equation 4 leads to:

$$\mathbb{E}(T_{rec}(q)) = e^{-q\lambda R(q)} (\mathbb{E}(X_D(q)) + R(q)) + (1 - e^{-q\lambda R(q)}) \left( \mathbb{E}(X_D(q)) + \frac{1}{q\lambda} - \frac{R(q)}{e^{q\lambda R(q)} - 1} + \mathbb{E}(T_{rec}(q)) \right)$$
(5)

Equation 5 reads as follows: after the downtime  $X_D(q)$ , either the recovery succeeds for everybody, or there is a failure during the recovery and another attempt must be made. Both events are weighted by their respective probabilities. Simplifying the above expression we get:

$$\mathbb{E}(T_{rec}(q)) = \mathbb{E}(X_D(q))e^{q\lambda R(q)} + \frac{1}{q\lambda}(e^{q\lambda R(q)} - 1)$$
(6)

Plugging back this expression in Equation 3, we obtain the value given in Equation 1.

Now we establish the desired bounds on  $\mathbb{E}(X_D(q))$  We always have  $X_D(q) \ge X_D(1) \ge D$ , hence the lower bound. For the upper bound, consider a date at which one of the q processors, say processor  $i_0$ , just had a failure and initiates its downtime period for D time units. Some other processors might be in the middle of their downtime period: for each processor  $i, 1 \le i \le q$ , let  $t_i$  denote the remaining duration of the downtime of processor i. We have  $0 \le t_i \le D$  for

 $1 \leq i \leq q, t_{i_0} = D$ , and  $t_i = 0$  means that processor *i* is up and running. Let  $X_D^{t_1,...,t_q}(q)$  be the *remaining* downtime of a group of *q* processors, knowing that processor *i*,  $1 \leq i \leq q$ , will still be down for a duration of  $t_i$ , and that a failure just happened (i.e., there exists  $i_0$  such that  $t_{i_0} = D$ ). Given the values of the  $t_i$ 's, we have the following equation for the random variable  $X_D^{t_1,...,t_q}(q)$ :

$$X_D^{t_1,..,t_q}(q) = \begin{cases} D \\ \text{if none of the processors of the group} \\ \text{fails during the next } D \text{ units of time} \\ T_{lost}^{t_1,..,t_q}(D) + X_D^{t_1',..,t_q'}(q) \\ \text{otherwise.} \end{cases}$$

In the second case of the equation, consider the next D time-units. Processor i can only fail in the last  $D - t_i$  of these time-units. Here the values of the  $t'_i$ 's depend on the  $t_i$ 's and on  $T^{t_1,..,t_q}_{lost}(D)$ . Indeed, except for the last processor to fail, say  $i_1$ , for which  $t'_{i_1} = D$ , we have  $t'_i = \max\{t'_i - T^{t_1,..,t_q}_{lost}(D), 0\}$ . More importantly we always have  $T^{t_1,..,t_q}_{lost}(D) \leq T^{D,0,..,0}_{lost}(D)$  and  $X^{t_1,..,t_q}_D(q) \leq X^{D,0,..,0}_D(q)$  because the probability for a processor to fail during D time units is always larger than that to fail during  $D - t_i$  time-units. Thus,  $\mathbb{E}(X^{t_1,..,t_q}_D(q)) \leq \mathbb{E}(X^{D,0,..,0}_D(q))$ . Following the same line of reasoning, we derive an upper-bound for  $X^{D,0,..,0}_D(q)$ :

$$X_D^{D,0,..,0}(q) \leq \begin{cases} D \\ \text{if none of the } q-1 \text{ running processors of the group} \\ \text{fails during the downtime } D \\ T_{lost}^{D,0,..,0}(D) + X_D^{D,0,..,0}(q) \\ \text{otherwise.} \end{cases}$$

Weighting both cases by their probability and taking expectations, we obtain

$$\begin{split} & \mathbb{E}\left(X_{D}^{D,0,..,0}(q)\right) \\ & \leq e^{-(q-1)\lambda D}D + (1 - e^{-(q-1)\lambda D})\left(E\left(T_{lost}^{D,0,..,0}(D)\right) + E\left(X_{D}^{D,0,..,0}(q)\right)\right) \\ & \text{hence } \mathbb{E}\left(X_{D}^{D,0,..,0}(q)\right) \leq D + (e^{(q-1)\lambda D} - 1)E\left(T_{lost}^{D,0,..,0}(D)\right), \text{ with } \end{split}$$

$$E\left(T^{D,0,..,0}_{lost}(D)\right) = \frac{1}{(q-1)\lambda} - \frac{D}{e^{(q-1)\lambda D} - 1}.$$

We derive

$$\mathbb{E}\left(X_D^{t_1,\dots,t_q}(q)\right) \le \mathbb{E}\left(X_D^{D,\dots,0}(q)\right) \le \frac{e^{(q-1)\lambda D} - 1}{(q-1)\lambda}$$

which concludes the proof. As a sanity check, we observe that the upper bound is at least D, using the identity  $e^x \ge 1 + x$  for  $x \ge 0$ .

While in a failure-free environment  $\mathbb{E}^*(q)$  would always decrease as q increases, using the above lower bound on  $\mathbb{E}(X_D(q))$  we obtain the following results:

RR n° 7876

**Theorem 1.** When the failure distribution follows an Exponential law,  $\mathbb{E}^*(q)$  reaches its minimum for some finite value of q in the following scenarios: all job types (perfectly parallel, generic and numerical) with constant overhead, and generic or numerical jobs with proportional overhead.

Note that the only open scenario is with perfectly parallel jobs and proportional overhead. In this case the lower bound on  $\mathbb{E}^*(q)$  decreases to some constant value while the upper bound goes to  $+\infty$  as q increases.

*Proof.* We show that  $\lim_{q\to+\infty} \mathbb{E}^*(q) = +\infty$  for the relevant scenarios. We first plug the lower-bound of Equation 2 into Equation 6 and obtain:

$$\mathbb{E}(T_{rec}(q)) \ge De^{q\lambda R(q)} + \frac{1}{q\lambda} \left( e^{q\lambda R(q)} - 1 \right).$$

From Equation 1 we then derive the lower-bound:

$$\mathbb{E}^*(q) \ge K_0(q) \left(\frac{1}{q\lambda} + D\right) e^{q\lambda R(q)} \left(e^{\frac{q\lambda W(q)}{K_0(q)} + q\lambda C(q)} - 1\right)$$

using the fact that, by definition, the expression in the right hand-side of Equation 1 is minimized by  $K_0$ , where  $K_0(q) = \frac{q\lambda \mathcal{W}(q)}{1 + \mathbb{L}(-e^{-q\lambda C(q)-1})}$ .

With constant overhead. Let us consider the case of perfectly parallel jobs  $(\mathcal{W}(q) = \mathcal{W}/q)$  with constant checkpointing overhead (C(q) = R(q) = C). We get the lower bound:

$$\mathbb{E}^*(q) \ge K_0(q) \left(\frac{1}{q\lambda} + D\right) e^{q\lambda C} \left(e^{\frac{\lambda W}{K_0(q)} + q\lambda C} - 1\right)$$

where  $K_0(q) = \frac{\lambda W}{1 + \mathbb{L}(-e^{-q\lambda C-1})}$ . When q tends to  $+\infty$ ,  $K_0(q)$  goes to  $\lambda W$ , while  $(\frac{1}{q\lambda} + D)e^{q\lambda C} \left(e^{\frac{\lambda W}{K_0(q)} + q\lambda C} - 1\right)$  goes to  $+\infty$ . Consequently,  $\mathbb{E}^*(q)$  is bounded below by a quantity that goes to  $+\infty$ , which concludes the proof. This result also implies that  $\mathbb{E}^*(q)$  reaches a minimum for a finite q value for other job types (generic, numerical) with constant overhead, just because the execution time is larger in those cases than with perfectly parallel jobs.

Generic parallel job with proportional overhead. Here we assume that  $\mathcal{W}(q) = \mathcal{W}/q + \gamma \mathcal{W}$ , and use proportional overhead:  $C(q) = R(q) = \frac{C}{q}$ . We get the lower bound:

$$\mathbb{E}^*(q) \ge K_0(q) \left(\frac{1}{q\lambda} + D\right) e^{\lambda C} \left(e^{\frac{\lambda W + q\lambda\gamma W}{K_0(q)} + \lambda C} - 1\right)$$

where  $K_0(q) = \frac{\lambda W + q \lambda \gamma W}{1 + \mathbb{L}(-e^{-\lambda C-1})}$ . As before, we show that  $\lim_{q \to +\infty} \mathbb{E}_{min}^*(q) = +\infty$  to get the result. When q tends to  $+\infty$ ,  $K_0(q)$  tends to  $+\infty$ , while  $(\frac{1}{q\lambda} + D)e^{\lambda C}\left(e^{\frac{\lambda W + q \lambda \gamma W}{K_0(q)} + \lambda C} - 1\right)$  tends to some positive constant. This concludes the proof. Note that this proof also serves for generic parallel jobs with constant overhead, simply because the execution time is larger in that case than with proportional overhead.

RR n° 7876

Numerical kernels with proportional overhead. Here we assume that  $\mathcal{W}(q) = \mathcal{W}/q + \gamma \mathcal{W}^{2/3}/\sqrt{q}$ , and use proportional overhead:  $C(q) = R(q) = \frac{C}{q}$ . We get the lower bound:

$$\mathbb{E}^*(q) \ge K_0(q) \left(\frac{1}{q\lambda} + D\right) e^{\lambda C} \left(e^{\frac{\lambda W + \lambda \gamma W^{2/3} \sqrt{q}}{K_0(q)} + \lambda C} - 1\right)$$

where  $K_0(q) = \frac{\lambda W + \lambda \gamma W^{2/3} \sqrt{q}}{1 + \mathbb{L}(-e^{-\lambda C} - 1)}$ . As before, we show that  $\lim_{q \to +\infty} \mathbb{E}_{min}^*(q) = +\infty$  to get the result. When q tends to  $+\infty$ ,  $K_0(q)$  tends to  $+\infty$ , while  $(\frac{1}{q\lambda} + D)e^{\lambda C} \left(e^{\frac{\lambda W + \lambda \gamma W^{2/3} \sqrt{q}}{K_0(q)} + \lambda C} - 1\right)$  tends to some positive constant. This concludes the proof.

## 5 Group replication

Since using all processors to run a single application instance may not make sense in light of Theorem 1, the group replication approach consists in executing multiple application instances on different processor groups, where the number of processors in a group is closer to  $q_{opt}$ . All groups compute the same chunk simultaneously, and do so until one of them succeeds, potentially after several failed trials. Then all other groups stop executing that chunk and recover from the checkpoint stored by the successful group. All groups then attempt to compute the next chunk. Group replication can be implemented easily with no modification to the application, provided that the recovery implementation allows a group to recover immediately from a checkpoint produced by another group. In this section we formalize group replication as an execution protocol called ASAP (As Soon As Possible), and analyze its performance for Exponential failures. We then introduce dynamic programming solutions that work with general failure distributions.

### 5.1 The ASAP execution protocol

We consider g groups, where each group has q processors, with  $g \times q \leq p$ . A group is available for execution if and only if all its q processors are available. In case of a failure, the downtime of a group is a random variable  $X_D(q) \geq D$ , whose expectation is bounded in Proposition 1. If a group encounters a first processor failure at time t, the group is *down* between t and  $t + X_D(q)$ .

The ASAP algorithm proceeds in k macro-steps. During macro-step  $j, 1 \leq j \leq k$ , each group independently attempts to execute the *j*-th chunk of size  $\omega_j$  and to checkpoint, restarting as soon as possible in case of a failure. As soon as one of the groups succeeds, say at time  $t_j^{end}$ , all the other groups are immediately stopped, macro-step j is over, and macro-step (j+1) starts (if j < k). Note that the value of k, the total number of chunks, as well as the chunk sizes, the  $\omega_j$ 's, are inputs to the algorithm (we always have  $\sum_{j=1}^{k} \omega_j = \mathcal{W}(q)$ ). We provide an analytical evaluation of ASAP for Exponential failure laws, and discuss how to choose these values, in Section 5.2.

Two important aspects must be mentioned. First, before being able to start macro-step (j + 1), a group that has been stopped must execute a recovery, in order to restart from the checkpoint of a successful group. Second, this recovery



Figure 1: Execution of chunks  $\omega_1$  and  $\omega_2$  (macro-steps 1 and 2) using the ASAP protocol. At time  $t_1^{end}$ , Group 1 is not ready, and Group 2 is the only one that does not need to recover.

may start later than time  $t_j^{end}$ , in the case where the group is down at time  $t_j^{end}$ . An example is shown in Figure 1, in which group 1 cannot start the recovery at time  $t_j^{end}$ . The only groups that do not need to recover at the beginning of the next step are the groups that were successful for the previous step, except during the first step at which all groups can start computing right away.

We now provide an analytical evaluation of ASAP for Exponential failure laws, and show how to compute the number of macro-steps k and the values of the chunk sizes  $\omega_j$ .

#### 5.2 Exponential failures

Let use assume that the failure rate of each processor obeys an Exponential law of parameter  $\lambda$ . For the sake of the theoretical analysis, we introduce a slightly modified version of the ASAP protocol in which all groups, including the successful ones, execute a recovery at the beginning of all macro-steps, including the first one. This new version of ASAP is described in Algorithm 1. It is completely symmetric, which renders its analysis easier: for macro-step jto be successful, one of the groups must be up and running for a duration of  $R(q) + \omega_i + C(q)$ .

Al	gorithm 1: ASAP $(\omega_1, \ldots, \omega_k)$
1 f	$\mathbf{pr} \ j = 1 \ to \ k \ \mathbf{do}$
2	for each group do in parallel
3	repeat
4	Finish current downtime (if any)
5	Try to perform a recovery, then a chunk of size $\omega_j$ , and finally to
	checkpoint
6	if execution successful then
7	Signal other groups to immediately stop their attempts
8	until one of the groups has a successful attempt

Consider the *j*-th macro step, number the attempts of all groups by their start time, and let  $N_j$  be the index of the earliest started attempt that successfully computes chunk  $\omega_j$ . In Figure 2, we have j = 2, the successful chunk of size  $R + \omega_2 + C$  is the fourth attempt, so  $N_2 = 4$ . To represent each attempt, we sample random variables  $X_i^j$  and  $Y_i^j$ ,  $1 \le i \le N_j$ , that correspond respectively



Figure 2: Zoom on macro-step 2 of the execution depicted in Figure 1, using the (X, Y) notation of Algorithm 2. Recall that  $Job_i$  has size  $X_i^2 + Y_i^2$  for  $1 \le i \le 3$ , and  $Job_4$  has size  $R(q) + \omega_2 + C(q)$ .

to the  $i^{th}$  tentative execution of the chunk and to the  $i^{th}$  downtime that follows it (if  $i \neq N_j$ ). Note that  $X_i^j < R + \omega_j + C$  for  $i < N_j$ , and  $X_{N_j}^j \ge R + \omega_j + C$ . All the  $X_i^j$ 's follow the same distribution  $D_X$ , namely an Exponential law of parameter  $q\lambda$ . And all the  $Y_i^j$ 's follow the same distribution  $D_{X_D}(q)$ , that of the the random variable  $X_D(q)$  corresponding to the downtime of a group of qprocessors.

The main idea here is to view the  $N_j$  execution attempts as jobs, where the size of job *i* is  $X_i^j + Y_i^j$ , and to distribute them across the *g* groups using the classical online *list scheduling* algorithm for independent jobs [21, section 5.6]. This formulation (see Proposition 2) allows us to provide an upper bound for the starting time of job  $N_j$ , and hence for the length of macro-step *j*, using a well-known scheduling argument (see Proposition 3). We then derive an upper bound for the expected execution time of ASAP (see Theorem 2).

Algorithm 2: Step j of ASAP  $(\omega_1, \ldots, \omega_k)$ 

 $i \quad i \leftarrow 1$ /\* i represents the number of attempts for the job \*/ /\*  $\mathcal L$  represents the list of attempts for the job \*/ 2  $\mathcal{L} \leftarrow \emptyset$ **3** Sample  $X_i^j$  and  $Y_i^j$  using  $D_X$  and  $D_{X_D(q)}$  respectively while  $X_i^j < R(q) + \omega_j + C(q)$  do 4 Add  $Job_i$ , with processing time  $X_i^j + Y_i^j$ , to  $\mathcal{L}$ 5 6  $i \leftarrow i + 1$ Sample  $X_i^j$  and  $Y_i^j$  using  $D_X$  and  $D_{X_D(q)}$  respectively 7 **8**  $N_j \leftarrow i$ 9 Add  $Job_{N_j}$ , with processing time  $R(q) + \omega_j + C(q)$ , to  $\mathcal{L}$ /\* the first successful job has size  $R(q) + \omega_j + C(q)$ , not  $X_{N_j}^j + Y_{N_j}^j$ \*/ 10 From time  $t_{j-1}^{end}$  on, execute a List Scheduling algorithm to distribute jobs of  $\mathcal{L}$ 

to the different groups (recall that some groups may not be ready at time  $t_{j-1}^{end}$ )

**Proposition 2.** The *j*-th macro-step of the ASAP protocol can be simulated using Algorithm 2: the last job scheduled by Algorithm 2 ends exactly at time  $t_i^{end}$ .

*Proof.* The List Scheduling algorithm distributes the next job to the first available group. Because of the memoryless property of Exponential laws, it is



Figure 3: Notations used in Proposition 3.

equivalent (i) to generate the attempts *a priori* and greedily schedule them, or (ii) to generate them independently within each group.  $\Box$ 

**Proposition 3.** Let  $T_{truestart}^{(R(q)+\omega_j+C(q))}$  be the time elapsed between  $t_{j-1}^{end}$  and the beginning of  $Job_{N_j}$  (see Figure 3). We have  $\mathbb{E}\left(T_{truestart}^{(R(q)+\omega_j+C(q))}\right) \leq \mathbb{E}(Y) + \frac{\mathbb{E}(N_j)\mathbb{E}(X)-\mathbb{E}(X_j^{N_j})+(\mathbb{E}(N_j)-1)\mathbb{E}(Y)}{g}$  where X and Y are random variables corresponding to an attempt (sampled using  $D_X$  and  $D_{X_D(q)}$  respectively). Moreover, we have  $\mathbb{E}(N_j) = e^{\lambda q(R(q)+\omega_j+C(q))}$  and  $\mathbb{E}(X_j^{N_j}) = \frac{1}{q\lambda} + R(q) + \omega_j + C(q)$ .

Proof. For group  $x, 1 \leq x \leq g$ , let  $\tilde{Y}_x$  denote the time elapsed before it is ready for macro-step j. For example in Figure 2, we have  $\tilde{Y}_1 > 0$  (group 1 is down at time  $t_{j-1}^{end}$ ), while  $\tilde{Y}_2 = \tilde{Y}_3 = 0$  (groups 2 and 3 are ready to compute at time  $t_{j-1}^{end}$ ). Proposition 2 has shown that executing macro-step j can be simulated by executing a List Schedule on a job list  $\mathcal{L}$  (see Algorithm 2). We now consider g "jobs"  $Job_x, x = 1, \ldots, g$ , so that  $Job_x$  has duration  $\tilde{Y}_x$ . We now consider the augmented job list  $\mathcal{L}' = \mathcal{L} \cup \bigcup_{x=1}^g Job_x$ . Note that  $\mathcal{L}'$  may contain more jobs than macro-step j: the jobs that start after the successful job  $Job_{N_j}$  are discarded from the list  $\mathcal{L}'$ . However, both schedules have the same makespan, and jobs common to both systems have the same start and completion dates. Thus, we have  $T_{\text{truestart}}^{(R(q)+\omega_j+C(q))} \leq \frac{\sum_{x=1}^g (\tilde{Y_x}) + \sum_{i=1}^{N_j-1} (X_i^j + Y_i^j)}{g}$ : this key inequality is due to the property of list scheduling: the group which is assigned the last job is the least loaded when this assignment is decided, hence its load does not exceed the average load (which is the total load divided by the number of groups). Given that  $\mathbb{E}(\tilde{Y_x}) \leq \mathbb{E}(Y)$ , we derive

$$\mathbb{E}\left(T_{\text{truestart}}^{(R(q)+\omega_j+C(q))}\right) \le \mathbb{E}(Y) + \frac{\mathbb{E}\left(\sum_{i=1}^{N_j-1} X_i^j\right) + \mathbb{E}\left(\sum_{i=1}^{N_j-1} (Y_i^j)\right)}{g}$$

But  $N_j$  is the stopping criterion of the  $(X_i^j)$  sequence, hence using Wald's theorem we have  $\mathbb{E}(\sum_{i=1}^{N_j} X_i^j) = \mathbb{E}(N_j)\mathbb{E}(X)$  which leads to  $\mathbb{E}(\sum_{i=1}^{N_j-1} X_i^j) = \mathbb{E}(N_j)\mathbb{E}(X) - \mathbb{E}(X_j^{N_j})$ . Moreover, as  $N_j$  and  $Y_i^j$  are independent variables,

we have  $\mathbb{E}(\sum_{i=1}^{N_j-1} Y_i^j) = (\mathbb{E}(N_j) - 1)\mathbb{E}(Y)$ , and we get the desired bound for  $\mathbb{E}(T_{\text{truestart}}^{(R(q)+\omega_j+C(q))})$ .

Finally, as the expected number of attempts when repeating independently until success an event of probability  $\alpha$  is  $\frac{1}{\alpha}$  (geometric law), we get  $\mathbb{E}(N_j) = e^{\lambda q(R(q)+\omega_j+C(q))}$ . The value of  $\mathbb{E}(X_j^{N_j})$  can be directly computed from the definition, recalling that  $X_j^{N_j} \geq R(q) + \omega_j + C(q)$  and each  $X_j^i$  follows an Exponential distribution of parameter  $q\lambda$ .

**Theorem 2.** The expected execution time of ASAP has the following upper bound:

$$\begin{split} \frac{g-1}{g}\mathcal{W}(q) &+ \frac{1}{g}\left(\frac{1}{q\lambda} + \mathbb{E}(Y)\right)e^{\lambda q(R(q) + C(q))}k^*e^{\lambda q\frac{\mathcal{W}(q)}{k^*}} \\ &+ k^*\left(\frac{g-1}{g}\left(\mathbb{E}(Y) + R(q) + C(q)\right) - \frac{1}{g}\frac{1}{q\lambda}\right) \end{split}$$

which is obtained when using  $k^* = \max(1, \lfloor k_0 \rfloor)$  or  $k^* = \lceil k_0 \rceil$  same-size chunks, whichever leads to the smaller value, where:

$$k_0 = \frac{\lambda q \mathcal{W}(q)}{1 + \mathbb{L}\left(\left(g - 1 + \frac{(g - 1)q\lambda(R(q) + C(q)) - g}{1 + q\lambda\mathbb{E}(Y)}\right)e^{-(1 + \lambda q(R(q) + C(q)))}\right)}$$

*Proof.* From Proposition 3, the expected execution time of ASAP has upper bound  $T_{ASAP} = \sum_{j=1}^{k} \alpha_j$ , where

$$\alpha_j = \mathbb{E}(Y) + \frac{\mathbb{E}(N_j)\mathbb{E}(X) - \mathbb{E}(X_j^{N_j}) + (\mathbb{E}(N_j) - 1)\mathbb{E}(Y)}{g} + (R(q) + \omega_j + C(q)).$$

Our objective now is to find the inputs to the ASAP algorithm, namely the number k of macro-steps together with the chunk sizes  $(\omega_1, \ldots, \omega_k)$ , that minimize this  $T_{ASAP}$  bound.

We first have to prove that any optimal (in expectation) policy uses only a finite number of chunks. Let  $\alpha$  be the expectation of the ASAP makespan using a unique chunk of size  $\mathcal{W}(q)$ . According to Proposition 3,

$$\alpha = \mathbb{E}(T_{\text{truestart}}^{(R(q)+\mathcal{W}(q)+C(q))}) + C(q) + \mathcal{W}(q) + R(q))$$

and is finite. Thus, if an optimal policy uses  $k^*$  chunks, we must have  $k^*C(q) \leq \alpha$ , and thus  $k^*$  is bounded.

In the proof of Theorem 1 in [4], we have shown that any deterministic strategy uses the same sequence of chunk sizes, whatever the failure scenario, thanks to the memoryless property of the Exponential distribution. We cannot prove such a result in the current context. For instance, the number of groups performing a downtime at time  $t_1^{end}$  depends on the scenario. There is thus no reason a priori for the size of the second chunk to be independent of the scenario. To overcome this difficulty, we restrict our analysis to strategies that use the same sequence of chunk sizes whatever the failure scenario. We optimize  $T_{ASAP}$  in that context, at the possible cost of finding a larger upper bound.

We thus suppose that we have a fixed number of chunks, k, and a sequence of chunk sizes  $(\omega_1, \ldots, \omega_k)$ , and we look for the values of  $(\omega_1, \ldots, \omega_k)$  that minimize

 $T_{ASAP} = \sum_{j=1}^{k} \alpha_j$ . Let us first compute one of the  $\alpha_j$  term. Replacing  $\mathbb{E}(N_j)$  and  $\mathbb{E}(X_j^{N_j})$  by the values given in Proposition 3, and  $\mathbb{E}(X)$  by  $\frac{1}{q\lambda}$ , we get

$$\begin{split} \alpha_j &= \frac{g-1}{g} \omega_j + \frac{1}{g} e^{\lambda q(R(q) + \omega_j + C(q))} \left( \frac{1}{q\lambda} + \mathbb{E}(Y) \right) \\ &+ \frac{g-1}{g} \left( \mathbb{E}(Y) + R(q) + C(q) \right) - \frac{1}{g} \frac{1}{q\lambda} \end{split}$$

$$T_{ASAP} = \frac{g-1}{g} \mathcal{W} + \frac{1}{g} \left( \frac{1}{q\lambda} + \mathbb{E}(Y) \right) e^{\lambda q (R(q) + C(q))} \sum_{j=1}^{k} e^{\lambda q \omega_j} + k \left( \frac{g-1}{g} \left( \mathbb{E}(Y) + R(q) + C(q) \right) - \frac{1}{g} \frac{1}{q\lambda} \right)$$

By convexity, the expression  $\sum_{j=1}^{k} e^{\lambda q \omega_j}$  is minimal when all  $\omega_j$ 's are equal (to  $\mathcal{W}(q)/k$ ). Hence all the chunks should be equal for  $T_{ASAP}$  to be minimal. We obtain:

$$T_{ASAP} = \frac{g-1}{g} \mathcal{W} + \frac{1}{g} \left( \frac{1}{q\lambda} + \mathbb{E}(Y) \right) e^{\lambda q (R(q) + C(q))} k e^{\lambda q \frac{\mathcal{W}(q)}{k}} + k \left( \frac{g-1}{g} \left( \mathbb{E}(Y) + R(q) + C(q) \right) - \frac{1}{g} \frac{1}{q\lambda} \right).$$

Let  $f(x) = \tau_1 x e^{\lambda q \frac{W(q)}{x}} + \tau_2 x$ , where

$$\tau_1 = \frac{1}{g} \left( \frac{1}{q\lambda} + \mathbb{E}(Y) \right) e^{\lambda q(R(q) + C(q))} \quad \text{and}$$
$$\tau_2 = \left( \frac{g-1}{g} \left( \mathbb{E}(Y) + R(q) + C(q) \right) - \frac{1}{g} \frac{1}{q\lambda} \right).$$

A simple analysis using differentiation shows that f has a unique minimum, and solving f'(x) = 0 leads to  $\tau_1 e^{\lambda q \frac{W(q)}{k}} \left(1 - \frac{\lambda q W(q)}{k}\right) + \tau_2 = 0$ , and thus to  $k = \frac{\lambda q W(q)}{1 + \mathbb{L}\left(\frac{\tau_2}{\tau_1 \cdot e}\right)} = k^*$ , which concludes the proof.  $\Box$ 

Using the upper bound of  $\mathbb{E}(Y) = \mathbb{E}(X_D(q))$  in Proposition 1, we can compute numerically the number of chunks and the expectation of the upper bound given by Theorem 2.

### 5.3 Group replication heuristics for general failure distributions

The results of the previous section are limited to Exponential failures. We now address turn to the general case. In Section 5.3.1 we recall the dynamic program designed in [4] to define checkpointing dates in a context *without* replication. We then discuss how to use this dynamic program in the context of ASAP (Section 5.3.2). Finally, in Section 5.3.3 we propose heuristics that correspond to more general execution protocols than ASAP.

Algorithm 3: DPNEXTFAILURE  $(W, \tau_1, ..., \tau_p)$ 1 if W = 0 then return 02 best  $\leftarrow 0$ ; chunksize  $\leftarrow 0$ 3 for  $\omega =$  quantum to W step quantum do4 (expected\_work, 1st\_chunk)  $\leftarrow$  DP-<br/>NEXTFAILURE( $W - \omega, \tau_1 + \omega + C(p), ..., \tau_p + \omega + C(p)$ )5 cur\_exp\_work  $\leftarrow$ <br/> $P_{suc}(\tau_1 + \omega + C(p), ..., \tau_p + \omega + C(p) | \tau_1, ..., \tau_p) \times (\omega + expected_work)$ 6 if cur\_exp\_work > best then best  $\leftarrow$  cur\_exp\_work; chunksize  $\leftarrow \omega$ <br/> $\tau$  return (best, chunksize)

#### 5.3.1 Solution without replication

According to [4], the most efficient algorithm to define checkpointing dates, for *general* failure distributions and when no replication is used, is the dynamic programming approach called DPNEXTFAILURE, shown in Algorithm 3. This algorithm works on a failure by failure basis, maximizing the expectation of the amount of work completed (and checkpointed) before the next failure occurs. This algorithm only provides an approximation of the optimal solution as it relies on a time discretization. (For the sake of simplicity, we present all dynamic programs as recursive algorithms.)

#### 5.3.2 Implementing the ASAP execution protocol

A key question when implementing ASAP is that of the chunk size. A naive approach would be to use DPNEXTFAILURE. Each group would call DPNEXT-FAILURE to compute what would be the optimal chunk size for itself, as if there were no other groups. Then we have to *merge* these individual chunk sizes to obtain a common chunk size. This heuristic, DPNEXTFAILUREASAP, is shown in Algorithm 4 (the ALIVE function returns, for a list of q processors, the amount of time each has been up and running since its last downtime). For the MERGE operator, one could be pessimistic and take the minimum of the chunk sizes, or be optimistic and take the maximum, or attempt a trade-off by taking the average. Two important limitations of this heuristic are: 1) the MERGE operator that has no theoretical justification; and 2) the use of chunk sizes defined using an obsolete failure history. The latter limitation shows after a group is victim of a failure: the failure history has changed significantly but the chunk size is not recomputed (this has no consequences with an Exponential distribution as the Exponential is memoryless).

#### 5.3.3 Other execution protocols

In order to circumvent both limitations of DPNEXTFAILUREASAP, we relax the constraint that all groups work with the same chunk size. Each group now works with its own chunk size that is recomputed each time the group fails, and each time one of the groups successfully complete its own chunk. This leads to the heuristic DPNEXTFAILURESYNCHRO in Algorithm 5. Each time a group successfully works for the duration of its chunk size and checkpoints its work, it signals its success to all groups which then interrupt their own work. This

Al	gorithm 4: DPNEXTFAILUREASAP $(W)$ .
1 W	while $W \neq 0$ do
2	for each group $x = 1g$ do in parallel
3	$(\tau_{(x-1)q+1}, \dots, \tau_{(x-1)q+q}) \leftarrow \text{ALIVE}((x-1)q+1, \dots, (x-1)q+g)$
4	$(exp\_work_x, \omega_x) \leftarrow DPNEXTFAILURE(W, \tau_{(x-1)q+1}, \dots, \tau_{(x-1)q+q})$
5	$\omega \leftarrow \operatorname{Merge}(\omega_1, \ldots, \omega_g)$
6	for each group do in parallel
7	repeat
8	Try to execute a chunk of size $\omega$ and then checkpoint
9	if successful then
10	Signal other groups to immediately stop their attempts
11	else if <i>failure</i> then Complete downtime and perform recovery
12	until One of the groups signals its success
13	$W \leftarrow W - \omega$
14	for each group do in parallel
15	if not successful on last chunk then Perform recovery from last
	successfully completed checkpoint

behavior is clearly sub-optimal if the successful completion is for a very small chunk and an interrupted group that was computing a large chunk was close to completion. The problems are that: 1) chunk sizes are defined as if each group was alone; and 2) the definition of the chunk sizes does not account for the fact that the first group to complete its checkpoint defines a mandatory checkpointing date for all groups.

**Algorithm 5:** DPNEXTFAILURESYNCHRO(W).

1 fe	or each group $x = 1g$ do in parallel
2	while $W \neq 0$ do
3	$(\tau_{(x-1)q+1}, \dots, \tau_{(x-1)q+q}) \leftarrow \text{ALIVE}((x-1)q+1, \dots, (x-1)q+g)$
4	$(exp\_work_x, \omega_x) \leftarrow \text{DPNExtFAILURE}(W, \tau_{(x-1)q+1}, \dots, \tau_{(x-1)q+q})$
5	Try to execute a chunk of size $\omega_x$ and then checkpoint
6	if successful then
7	Signal other groups to immediately stop their attempts
8	$W \leftarrow W - \omega_x$
9	if <i>failure</i> then Complete downtime
10	if failure or signal then
11	Perform recovery from last successfully completed checkpoint

To address these problems, rather than doing another attempt at reusing DPNEXTFAILURE, we design a brand new dynamic program. In [4], without replication, we found that a dynamic program that minimizes the expectation of the makespan would require an intractable exponential number of states to record which processors fail and when. Hence we aimed instead at maximizing the expectation of the work completed before the next failure. We now extend this approach to the context of replication. The first failure will only interrupt a single group. Therefore, the objective should be to maximize the expectation of the work completed before all groups have failed. This approach ignores that once a group has failed, it will eventually restart and resume computing. However, keeping track of such restarts would require recording which processors

have failed and when, thus once again leading to an exponential number of states.

To avoid having the first completed checkpoint force a checkpoint for all other groups we design a new dynamic program, DPNEXTCHECKPOINT (Algorithm 6). DPNEXTCHECKPOINT does not define chunk sizes, i.e., amount of work to be processed before a checkpoint is taken, but instead it defines checkpoint dates. The rationale is that one checkpointing date can correspond to different amounts of work for each group, depending on when the group has started to process its chunk, after either its last failure and recovery, or its last checkpoint, or its last recovery based on another group's checkpoint. The function WORKALREADYDONE (Line 3) returns, for each group, the time since it started processing its current chunk.

DPNEXTCHECKPOINT proceeds as follows. At the checkpointing date, the amount of work completed is the maximum of the amount of work done by the different groups that successfully complete the checkpoint. Therefore, we consider all the different cases (Line 8), that is, which group x, among the successful groups, has done the most work. We compute the probability of each case (Line 11). All groups that started to work earlier than group x have failed (i.e., at least one processor in each of them has failed) but not group x (i.e., none of its processors have failed). We compute the expectation of the amount of work completed in each case (Lines 12 and 13). We then sum the contributions of all the cases (Line 14) and record the checkpointing date leading to the largest expectation (Line 15). Note that the probability computed at Line 11 explicitly states which groups have successfully completed the checkpoint and which groups have not. We choose not to take this information into account when computing the expectation (recursive call at Line 13). This is to avoid keeping track of which group had failed, thereby lowering the complexity of the dynamic program. This explains why the conditions do not evolve in the conditional probability at Line 11.

Finally, Algorithm 7 shows the algorithm, called DPNEXTFAILUREGLOBAL, that uses DPNEXTCHECKPOINT. Each time a group is affected by an event (a failure, a successful checkpoint by itself or by another group), it computes the next checkpoint date and signals the result of its computation to the other groups (e.g., by broadcasting it to the g group leaders). Hence, a group may have computed the next checkpoint date to be t, and that date can be either un-modified, or postponed, or advanced by events occurring on other groups and by their re-computation of the best next checkpoint date. In practice, as these algorithms rely on a time discretization, at each time quantum a group can check whether the current time is a checkpoint date or not.

## 6 Simulation framework

In this section we detail our simulation methodology. We use both synthetic and real-world failure distributions. The source code and all simulation results are publicly available at: http://perso.ens-lyon.fr/frederic.vivien/Data/Resilience/SPAA2012/.

Algorithm 6: DPNEXTCHECKPOINT $(W, T, T_0, \tau_1, ..., \tau_{gq})$ 

1 if W = 0 then return 0 **2** best work  $\leftarrow 0$ ;  $next\_chkpt \leftarrow date$ 3  $(W_1, ..., W_g) \leftarrow WORKALREADYDONE(T)$  /\* Time since last recovery or checkpoint \*/ 4 Reorder groups in non-increasing availabilities ( $W_1$  is maximum) 5 for t = T to  $T + W - W_g$  step quantum /\* Loop on checkpointing date \*/ 6 do 7  $cur\_work \gets 0$ /\* Loop on the first group to successfully work 8 for x = 1 to g until t + C(q) \*/do 9  $\delta \leftarrow (t+C(q)) - T_0$  /\* Total time elapsed until the checkpoint 10 completion \*/  $\begin{aligned} proba \leftarrow \prod_{y=1}^{x-1} P_{fail}(\tau_{(y-1)q+1} + \delta, ..., \tau_{(y-1)q+p} + \delta \,|\, \tau_{(y-1)q+1}, ..., \tau_{(y-1)q+p}) \\ \times P_{suc}(\tau_{(x-1)q+1} + \delta, ..., \tau_{(x-1)q+p} + \delta \,|\, \tau_{(x-1)q+1}, ..., \tau_{(x-1)q+p}) \end{aligned}$ 11  $\omega \leftarrow \min\{W - W_x, t - T\}$  /\* Work done between T and t by group 12 x \*/  $(rec\_\omega, rec\_t) \leftarrow DP$ -13 NEXTCHECKPOINT $(W - W_x - \omega, T + \omega + C(q) + R(q), T_0, \tau_1, ..., \tau_p)$  $cur\_work \leftarrow cur\_work + proba \times (W_x + \omega + rec\_\omega)$  $\mathbf{14}$ if  $cur\_work > best\_work$  then 15 $best\_work \leftarrow cur\_work; \quad next\_chkpt \leftarrow t$ 16 return (best\_work, next\_chkpt)

**Algorithm 7:** DPNEXTFAILUREGLOBAL(W).

1 f	or each group $x = 1g$ do in parallel
2	while $W \neq 0$ do
3	$(\tau_1,, \tau_{gq}) \leftarrow \text{ALIVE}(1,, gq)$
4	$T_0 \leftarrow \text{TIME}()$ /* Current time */
5	$date \leftarrow \text{DPNExtCheckpoint}(W, T_0, T_0, \tau_1, \dots, \tau_{gq})$
6	Signal all processors that the next checkpoint date is now $date$
7	Try to work until <i>date</i> and then checkpoint
8	if successful work until date and checkpoint then
9	Let $y$ be the longest running group without failure among the
	successful groups
10	Let $\omega$ be the work performed by $y$ since its last recovery or
	checkpoint
11	$W \leftarrow W - \omega$
12	if group $x$ last recovery or checkpoint was strictly later than that of
	y then
13	Perform a recovery
14	if <i>failure</i> then Complete downtime
15	if failure or signal then Perform recovery from last successfully
	completed checkpoint

#### 6.1 Heuristics

Our simulator implements the following eight checkpointing policies:

- Two versions of the ASAP protocol: OPTEXP, that uses the optimal and periodic policy established in [4] for Exponential failure distributions and no replication; OPTEXPGROUP, that uses the periodic policy defined by Theorem 2 for Exponential distributions.
- The six dynamic programming approaches in Section 5.3: DPNEXT-FAILURE, the 3 variants of DPNEXTFAILUREASAP, DPNEXTFAILURESYN-CHRO, and DPNEXTFAILUREGLOBAL.

Our simulator also implements PERIODLB, which is a numerical search for the optimal period for ASAP. We evaluate each candidate period on 50 randomly generated scenarios. To build the candidate periods, the period computed for 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. The evaluation of PERIODLB on a configuration requires running 24,000 simulations (which would be prohibitive in practice), but we include it for reference. Based on the results in [4], we do not consider any additional checkpointing policy, such as those defined by Young [28] or Daly [9] for instance. We point out that OPTEXP and OPTEXPGROUP compute the checkpointing period based solely on the MTBF, implicitly assuming 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 are not exponentially distributed. To use OPTEXP with g groups we use the period from [4] computed with |p/g| processors.

#### 6.2 Platforms

We target two types of platforms, depending on the type of the failure distribution. For synthetic distributions, we consider platforms containing from 32,768 to 1,048,576 processors. For platforms with failures based on failure logs from production clusters, because of the limited scale of those clusters, we restrict the size of the platforms to a maximum of 131,072 processors, starting with 4,096 processors. For both platform types, we determine the job size W so that a job using the whole platform would use it for a significant amount of time in the absence of failures, namely  $\approx 3.5$  days on the largest platforms for synthetic failures (W = 10,000 years), and  $\approx 2.8$  days on those for log-based failures (W = 1,000 years). Otherwise, we use the same parameters as in [4]:  $C = R = 600 \ s, \ D = 60 \ s, \ \gamma = 10^{-6}$  for generic parallel jobs, and  $\gamma = 0.1$ for numerical kernels. Note that the checkpointing overheads come from the scenarios in [7] and are smaller than those used in [12].

#### 6.3 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 contains 45, 208 processors and is said to experience on the order of 1 failure per day [20, 2]. Assuming a 1-day platform MTBF gives us a processor MTBF equal to  $\frac{45,208}{365} \approx 125$  years. For the Exponential distribution, we then have  $\lambda$  as  $\lambda = \frac{1}{MTBF}$  and for Weibull, which requires two

parameters k and  $\lambda$ , we have  $\lambda = MTBF/\Gamma(1+1/k)$  and we fix k to 0.7 based on the results of [24]. (We have shown in [4] that the general trends were not influenced by the exact values used for k nor for the MTBF.)

Log-based failure distributions – We also consider failure distributions based on failure logs from production clusters. We used logs from the largest clusters among the preprocessed logs in the *Failure trace archive* [18], i.e., from clusters at the Los Alamos National Laboratory [24]. 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 only two clusters with more than 1,000 nodes, as we needed a sample history sufficiently large to simulate platforms with more than 10,000 nodes. The two chosen logs are for clusters 18 and 19 in the archive (referred to as 7 and 8 in [24]). For each log, we record the set S of availability intervals. A discrete failure distributionfor the simulation is then generated as follows: the conditional probability  $\mathbb{P}(X \ge t \mid X \ge \tau)$  that a node stays up for a duration t, knowing that it has been up for a duration  $\tau$ , is set to the ratio of the number of availability durations in S greater than or equal to t, over the number of availability durations in Sgreater than or equal to  $\tau$ .

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 (set to 2 years). The job start time is assumed to be 1 year for synthetic distribution platforms, and 0.25 year for log-based distribution platforms. We use a non-null start time to avoid side-effects related to the synchronous initialization of all 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.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. Finally, the two clusters used for computing our log-based failure distributions consist of 4-processor nodes. Hence, to simulate a 131,072-processor platform we generate 32,768 failure traces, one for each four-processor node.

# 7 Simulation results

In this section we discuss simulation results, but only show graphs for perfectly parallel applications under the constant overhead scenario. All results can be found in the appendix. All simulations with replication are for two groups (using three groups never leads to improvements in our experiments).

**Optimal number of processors for execution** – For Weibull and log-based failures with the constant overhead scenario, using all the available processors to run a single application instance leads to significantly larger makespans. This is seen in Figures 5-7 where the curves for OPTEXP, PERIODLB, and DPNEXT-FAILURE, the no-replication strategies, shoot upward when p reaches a large enough value. For instance, with traces based on the logs of LANL cluster 18, the increase in makespan is more than 37% when going from  $p = 2^{16}$  to  $p = 2^{17}$ . This behavior is in line with the result in Theorem 1. However, in our experiments



Figure 4: Exponential failures with MTBF=125 years.





Figure 5: Weibull failures with MTBF=125 years and k = 0.70.



Figure 6: Failures based on the failure log of LANL cluster 18.

Figure 7: Failures based on the failure log of LANL cluster 19.

with Exponential failures the lowest makespan when running a single application instance is obtained when using all the processors, regardless of the application model. This is seen plainly in Figure 4 at least up to  $2^{20}$  processors. This likely means that our experimental scenarios have not reached the scale at which the makespan would increase. Furthermore, running a single application instance is not always best in our results even for Exponential failures. For instance, for generic parallel jobs OPTEXP (with g = 1) begins being outperformed by strategies that use replication at  $2^{20}$  processors (see the full results in [5]). We conclude that on exascale platforms the optimal makespan may not be obtained by running a single application instance. This is similar to the result obtained in [12] in the context of process replication.

The ASAP protocol with Exponential – Considering only those results obtained for Exponential failures and using group replication, we find that OPT-EXPGROUP with g = 2 delivers performance very close to that of the numerical lower bound on any periodic policy under ASAP (PERIODLB with g = 2) and slightly better than the performance of OPTEXP with g = 2. These results corroborates our theoretical study. Therefore, determining chunk sizes according to Theorem 2 is more efficient than using two instances that each use the chunk size determined by OPTEXP.

**The dynamic programming approaches -** The performance of the three variants of DPNEXTFAILUREASAP (with MERGE being the minimum, the average, or the maximum) are indistinguishable on Exponential and Weibull failure distributions. For log-based failure distributions, their performance are still very close, with the average and the maximum seemingly achieving better performance (the performance differences, however, may not be significant in light of the standard deviations).

DPNEXTFAILURESYNCHRO achieves worse performance than DPNEXTFAIL-UREASAP for Exponential and Weibull failure distributions. For log-based failure distributions, in some rare settings, the opposite is true. Overall, this heuristic provides disappointing results. By contrast, depending on the scenario, DPNEXTFAILUREGLOBAL is equivalent to the other dynamic programming approaches or outperforms them significantly. Furthermore, its performance is relatively close to that of PERIODLB with g = 2.

The very perplexing result here is that, regardless of the underlying failure distribution, DPNEXTFAILUREGLOBAL is outperformed by the periodic policy OPTEXP with g = 2. Recall that this policy erroneously assumes Exponential failures and computes the checkpointing period ignoring that replication is used! By contrast, DPNEXTFAILUREGLOBAL is based on strong theoretical foundations and was designed specifically to handle multiple groups efficiently. And yet, in non-Exponential failure cases, OPTEXP with q = 2 outperforms DPNEXTFAILUREGLOBAL (see Figures 5-7). In the case of log-based experiments (Figures 6 and 7), it may be that our datasets are not sufficiently large. We use failure data from the largest production clusters in the Failure trace archive [18], but these clusters are orders of magnitude smaller than our simulated platforms. Consequently, we are forced to "replay" the same failure data many times for many of our simulated nodes, which biases experiments. In the case of Weibull failures with k = 0.7 (Figure 5), it turns out that the failure distribution is sufficiently close to an Exponential, i.e., that k is sufficiently close to 1, so that OPTEXP does well. However, smaller values of k are reasonable as seen for instance in [19]  $(k \approx 0.5)$  and in [24] (0.33 < k < 0.49). In Table 1 we show results for smaller values of k for a platform with  $2^{20}$  processors. We see that, as k decreases, the performance of OPTEXP with q = 2degrades when compared to PERIODLB, reaching more than 100% degradation for k = 0.5. However, DPNEXTFAILUREGLOBAL still achieves close-to-optimal makespans in such cases, as can be seen on Figure 8 (for k = 0.5). Overall, DPNEXTFAILUREGLOBAL achieves good to very good performance all across the board. Finally, remark on Figure 8 that PERIODLB achieves better performance using a replication factor of 3 (g = 3) than with a replication factor of 2.

## 8 Conclusion

In this paper we have presented a study of replication techniques for large-scale platforms. These platforms are subject to failures, the frequency of which increase dramatically with platform scale. For a variety of job types (perfectly parallel, generic, or numerical) and checkpoint cost models (constant or proportional overhead), we have shown that using the largest possible number of processors does not always lead to the smallest execution time. This is because



Figure 8: Weibull failures with MTBF=125 years and k = 0.50.

using more resources implies facing more failures during execution, hence wasting more time tolerating them (via an increase in checkpointing frequency) and recovering from them. This waste results in a slow-down despite the additional hardware resources.

This observation led us to investigate replication as a technique to better use all the resources provided by the platform. While others have studied process replication [12], in this work we have focused on group replication, a more general and less intrusive approach. Group replication consists in partitioning the platform into several groups, which each executes an instance of the application concurrently. All groups coordinate to take advantage of the most advanced application state available. We have conducted an analytical study of group replication for Exponential failures when using the ASAP execution protocol, using an analogy with a list schedule to bound the number of attempts and the execution time of each group. This analogy makes it possible to compute an upper bound on the expected execution time, which can be computed numerically. We have also proposed several dynamic programming algorithms to minimize application makespan, whatever the failure distribution. We have compared all these approaches, along with no-replication approaches proposed in previous work, in simulation for both synthetic failure distributions and failure data from production clusters. Our main findings are 1) that replication can significantly lower the execution time of applications on very large scale platforms, for failure and checkpointing characteristics corresponding to today's platforms; 2) that our DPNEXTFAILUREGLOBAL dynamic programming approach is close to the optimal periodic solution (in fact close to PERIODLB, which uses a prohibitively expensive numerical search for the best period).

A clear direction for future work, which we are currently exploring and that could lead to results in the final version of the article, is the reasons for the unexpected good results for OPTEXP with g = 2. In terms of longer-term objectives, it would be interesting to generalize this work beyond the case of coordinated checkpointing. For instance, we plan to study hierarchical checkpointing schemes with message logging. Another interesting direction would be to study process replication and compare it to group replication, both theoretically and through simulations.

#### Acknowledgements

This work was supported in part by the ANR RESCUE project, and by the INRIA-Illinois Joint Laboratory for Petascale Computing.

# References

- 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] Marin Bougeret, Henri Casanova, Mikael Rabie, Yves Robert, and Frédéric Vivien. Checkpointing strategies for parallel jobs. In Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, SC '11, pages 33:1–33:11, New York, NY, USA, 2011. ACM.
- [5] Marin Bougeret, Henri Casanova, Yves Robert, Frédéric Vivien, and Dounia Zaidouni. Using group replication for resilience on exascale systems. Research Report RR-xxxx, INRIA, ENS Lyon, France, 2011. Available at graal.ens-lyon.fr/~yrobert/.
- [6] 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.
- [7] Franck Cappello, Henri Casanova, and Yves Robert. Checkpointing vs. migration for post-petascale supercomputers. In *ICPP'2010*. IEEE Computer Society Press, 2010.
- [8] 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.
- [9] J. T. Daly. A higher order estimate of the optimum checkpoint interval for restart dumps. *Future Generation Computer Systems*, 22(3):303–312, 2004.
- [10] 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.

- [11] C. Engelmann, H. H. Ong, and S. L. Scorr. The case for modular redundancy in large-scale highh performance computing systems. In Proc. of the 8th IASTED Infernational Conference on Parallel and Distributed Computing and Networks (PDCN), pages 189–194, 2009.
- [12] K. Ferreira, J. Stearley, J. H. III Laros, R. Oldfield, K. Pedretti, R. Brightwell, R. Riesen, P. G. Bridges, and D. Arnold. Evaluating the Viability of Process Replication Reliability for Exascale Systems. In *Pro*ceedings of the 2011 ACM/IEEE Conference on Supercomputing, 2011.
- [13] F. Gärtner. Fundamentals of fault-tolerant distributed computing in asynchronous environments. ACM Computing Surveys, 31(1), 1999.
- [14] T. Heath, R. P. Martin, and T. D. Nguyen. Improving cluster availability using workstation validation. SIGMETRICS Perf. Eval. Rev., 30(1):217– 227, 2002.
- [15] 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.
- [16] Nick Kolettis and N. Dudley Fulton. Software rejuvenation: Analysis, module and applications. In *FTCS '95*, page 381, Washington, DC, USA, 1995. IEEE CS.
- [17] D. Kondo, A. Chien, and H. Casanova. Scheduling Task Parallel Applications for Rapid Application Turnaround on Enterprise Desktop Grids. *Journal of Grid Computing*, 5(4):379–405, 2007.
- [18] 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.
- [19] 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.
- [20] 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.
- [21] Michael Pinedo. Scheduling: theory, algorithms, and systems (3rd edition). Springer, 2008.
- [22] 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.
- [23] B. Schroeder and G. Gibson. Understanding failures in petascale computers. Journal of Physics: Conference Series, 78(1), 2007.

k	PeriodLB	OptExp	Relative
	(g=2)	(g=2)	degradation
0.5	35.57	81.20	128.30%
0.6	21.50	30.10	39.99%
0.7	15.38	17.14	11.44%
0.8	12.27	12.45	1.46%
0.9	10.45	10.47	0.21%

Table 1: Makespans (in days) obtained by PERIODLB with g = 2 and OPTEXP with g = 2, and relative degradation of OPTEXP, as a function of k, the shape parameter of the Weibull distribution of failures, on platforms with 1,048,576 processors.

- [24] B. Schroeder and G. A. Gibson. A large-scale study of failures in highperformance computing systems. In Proc. of DSN, pages 249–258, 2006.
- [25] K. Venkatesh. Analysis of Dependencies of Checkpoint Cost and Checkpoint Interval of Fault Tolerant MPI Applications. *Analysis*, 2(08):2690– 2697, 2010.
- [26] L. Wang, P. Karthik, Z. Kalbarczyk, R.K. Iyer, L. Votta, C. Vick, and A. Wood. Modeling Coordinated Checkpointing for Large-Scale Supercomputers. In Proc. of the International Conference on Dependable Systems and Networks, pages 812–821, June 2005.
- [27] S. Yi, D. Kondo, B. Kim, G. Park, and Y. Cho. Using Replication and Checkpointing for Reliable Task Management in Computational Grids. In Proc. of the International Conference on High Performance Computing & Simulation, 2010.
- [28] John W. Young. A first order approximation to the optimum checkpoint interval. Communications of the ACM, 17(9):530–531, 1974.
- [29] Z. Zheng and Z. Lan. Reliability-aware scalability models for high performance computing. In Proc. of the IEEE Conference on Cluster Computing, 2009.

# A Detailed simulation results

	_		_		_	r—		
	DPNEXTFAILUREGLOBAL		$226.89 \pm 0.59$	$114.94 \pm 0.34$	$58.94 \pm 0.30$	$30.89 \pm 0.24$	$16.92 \pm 0.20$	$9.93 \pm 0.22$
	DPNEXTFAILURESYNCHRO		$254.27 \pm 1.32$	$133.61 \pm 0.88$	$71.65 \pm 0.50$	$39.05 \pm 0.35$	$21.79 \pm 0.31$	$12.42 \pm 0.22$
	SAP	MAX	$238.75 \pm 0.29$	$122.93 \pm 0.26$	$65.20 \pm 0.39$	$34.80 \pm 0.27$	$19.12 \pm 0.27$	$11.07 \pm 0.17$
g = 2	DPNEXTFAILUREAS	AVG	$238.75 \pm 0.56$	$122.99 \pm 0.31$	$65.14 \pm 0.26$	$34.80 \pm 0.24$	$19.10 \pm 0.16$	$11.07 \pm 0.19$
		MIN	$238.74 \pm 0.49$	$122.99 \pm 0.28$	$65.14 \pm 0.17$	$34.77 \pm 0.15$	$19.14 \pm 0.17$	$11.05 \pm 0.13$
	PERIODLB		$228.53 \pm 0.58$	$116.07 \pm 0.54$	$59.49 \pm 0.34$	$30.98 \pm 0.29$	$16.62 \pm 0.20$	$9.42 \pm 0.17$
	OPTEXPGROUP		$228.74 \pm 0.56$	$116.11 \pm 0.51$	$59.53 \pm 0.44$	$31.01 \pm 0.36$	$16.68 \pm 0.33$	$9.45 \pm 0.39$
	OPTEXP		$231.72 \pm 0.33$	$117.96 \pm 0.18$	$60.61 \pm 0.15$	$31.60 \pm 0.16$	$16.96 \pm 0.18$	$9.55 \pm 0.23$
	DPNEXTFAILURE		$124.12 \pm 0.77$	$65.38 \pm 0.57$	$35.27 \pm 0.47$	$19.73 \pm 0.32$	$11.79 \pm 0.35$	$7.87 \pm 0.31$
g = 1	PeriodLB		$124.14 \pm 0.86$	$65.21 \pm 0.60$	$35.12 \pm 0.56$	$19.63 \pm 0.35$	$11.72 \pm 0.32$	$7.82 \pm 0.31$
	OPTEXP		$124.14 \pm 0.86$	$65.21 \pm 0.60$	$35.16 \pm 0.53$	$19.71 \pm 0.37$	$11.74 \pm 0.33$	$7.82 \pm 0.31$
	d		32768	65536	131072	262144	524288	1048576

jobs.
parallel
Perfectly
(a)

	DPNEXTFAILUREGLOBAL		$230.65 \pm 0.60$	$118.68 \pm 0.34$	$62.82 \pm 0.33$	$34.97 \pm 0.28$	$21.33 \pm 0.28$	$15.07 \pm 0.27$
	DPNEXTFAILURESYNCHRO		$258.59 \pm 1.30$	$137.95 \pm 0.86$	$76.37 \pm 0.53$	$44.23 \pm 0.40$	$27.47 \pm 0.38$	$18.95 \pm 0.28$
	SAP	MAX	$242.70 \pm 0.34$	$126.96 \pm 0.22$	$69.40 \pm 0.24$	$39.35 \pm 0.16$	$24.13 \pm 0.25$	$16.86 \pm 0.30$
g = 2	NEXTFAILUREA	AVG	$242.59 \pm 0.46$	$126.98 \pm 0.28$	$69.42 \pm 0.23$	$39.42 \pm 0.39$	$24.13 \pm 0.29$	$16.77 \pm 0.15$
	DPI	MIN	$242.71 \pm 0.25$	$127.01 \pm 0.32$	$69.45 \pm 0.27$	$39.43 \pm 0.31$	$24.12 \pm 0.19$	$16.83 \pm 0.25$
	PERIODLB		$232.29 \pm 0.62$	$119.87 \pm 0.53$	$63.40 \pm 0.34$	$35.05 \pm 0.26$	$20.97 \pm 0.26$	$14.33 \pm 0.20$
	OPTEXPGROUP		$232.45 \pm 0.56$	$119.94 \pm 0.52$	$63.46 \pm 0.49$	$35.07 \pm 0.38$	$21.05 \pm 0.42$	$14.38 \pm 0.55$
	OPTEXP		$235.53 \pm 0.35$	$121.81 \pm 0.18$	$64.58 \pm 0.16$	$35.76 \pm 0.19$	$21.40 \pm 0.21$	$14.54 \pm 0.32$
	DPNEXTFAILURE		$128.22 \pm 0.76$	$69.64 \pm 0.61$	$39.90 \pm 0.46$	$24.90 \pm 0.39$	$17.90 \pm 0.39$	$16.02 \pm 0.46$
g = 1	PERIODLB		$128.22 \pm 0.86$	$69.48 \pm 0.60$	$39.76 \pm 0.60$	$24.75 \pm 0.41$	$17.83 \pm 0.35$	$15.92 \pm 0.48$
	OPTEXP		$128.22 \pm 0.86$	$69.48 \pm 0.60$	$39.78 \pm 0.49$	$24.81 \pm 0.41$	$17.83 \pm 0.37$	$15.95 \pm 0.48$
	đ		32768	65536	131072	262144	524288	1048576

jobs.
parallel
Generic
(q)

	DPNEXTFAILUREGLOBAL		$227.36 \pm 0.59$	$115.22 \pm 0.35$	$59.19 \pm 0.34$	$31.06 \pm 0.25$	$17.05 \pm 0.21$	$10.03 \pm 0.22$
	DPNEXTFAILURESYNCHRO		$254.78 \pm 1.28$	$134.02 \pm 0.84$	$71.98 \pm 0.52$	$39.25 \pm 0.36$	$21.98 \pm 0.32$	$12.55 \pm 0.21$
	AP	MAX	$239.20 \pm 0.56$	$123.29 \pm 0.29$	$65.35 \pm 0.17$	$35.02 \pm 0.27$	$19.31 \pm 0.25$	$11.19 \pm 0.20$
g = 2	NEXTFAILUREAS	AVG	$239.27 \pm 0.37$	$123.25 \pm 0.23$	$65.39 \pm 0.31$	$34.97 \pm 0.16$	$19.32 \pm 0.32$	$11.19 \pm 0.21$
	DPI	MIN	$239.16 \pm 0.47$	$123.34 \pm 0.42$	$65.42 \pm 0.25$	$35.00 \pm 0.25$	$19.28 \pm 0.17$	$11.16 \pm 0.13$
	PERIODLB		$228.95 \pm 0.60$	$116.39 \pm 0.51$	$59.75 \pm 0.45$	$31.15 \pm 0.31$	$16.73 \pm 0.19$	$9.52 \pm 0.17$
	OPTEXPGROUP		$229.18 \pm 0.57$	$116.43 \pm 0.49$	$59.80 \pm 0.49$	$31.15 \pm 0.36$	$16.81 \pm 0.32$	$9.56 \pm 0.40$
	OPTEXP		$232.14 \pm 0.34$	$118.26 \pm 0.17$	$60.86 \pm 0.17$	$31.78 \pm 0.16$	$17.09 \pm 0.18$	$9.65 \pm 0.23$
	DPNEXTFAILURE		$124.49 \pm 0.87$	$65.64 \pm 0.58$	$35.48 \pm 0.47$	$19.87 \pm 0.34$	$11.92 \pm 0.35$	$7.99 \pm 0.31$
g = 1	PERIODLB		$124.44 \pm 0.83$	$65.46 \pm 0.60$	$35.36 \pm 0.56$	$19.77 \pm 0.35$	$11.88 \pm 0.33$	$7.95 \pm 0.32$
	OPTEXP		$124.44 \pm 0.83$	$65.46 \pm 0.60$	$35.38 \pm 0.51$	$19.86 \pm 0.37$	$11.88 \pm 0.34$	$7.95 \pm 0.32$
	d		32768	65536	131072	262144	524288	1048576

(c) Numerical kernels.

Table 2: Evaluation of the different heuristics on a platform with **Exponential failures** (MTBF = 125 years) under the constant overhead model.

р	OptExp	PERIODLB	DPNEXTFAILURE	OPTEXP	OPTEXPGROUP	PERIODLB	DP	NEXTFAILUREA.	SAP	DPNEXTFAILURESYNCHRO	DPNEXTFAILUREGLOBAL
							MIN	AVG	MAX		
65536	$123.19 \pm 5.87$	$123.08 \pm 6.04$	$123.62 \pm 5.87$	$185.64 \pm 4.80$	$185.88 \pm 7.68$	$182.76 \pm 5.73$	$218.73 \pm 4.51$	$218.27 \pm 4.67$	$218.31 \pm 4.69$	$236.11 \pm 6.68$	$202.78 \pm 7.13$
131072	$61.89 \pm 2.91$	$61.79 \pm 2.49$	$61.96 \pm 2.89$	$92.79 \pm 2.20$	$92.55 \pm 3.60$	$91.22 \pm 3.45$	$109.71 \pm 2.14$	$109.45 \pm 2.11$	$109.48 \pm 2.14$	$118.43 \pm 2.55$	$101.39 \pm 3.43$
262144	$30.55 \pm 1.21$	$30.55 \pm 1.21$	$30.66 \pm 1.14$	$46.65 \pm 1.01$	$46.69 \pm 1.68$	$45.70 \pm 1.50$	$54.74 \pm 1.14$	$54.58 \pm 1.15$	$54.57\pm1.15$	$58.85 \pm 1.44$	$50.90 \pm 1.69$
524288	$15.46 \pm 0.58$	$15.45 \pm 0.47$	$15.51 \pm 0.58$	$23.14 \pm 0.60$	$23.22 \pm 0.89$	$22.89 \pm 0.71$	$27.33 \pm 0.60$	$27.26 \pm 0.60$	$27.26 \pm 0.60$	$29.43 \pm 0.55$	$25.27 \pm 0.64$
1048576	$7.82 \pm 0.31$	$7.82 \pm 0.31$	$7.87 \pm 0.31$	$9.55 \pm 0.23$	$9.45 \pm 0.39$	$9.42 \pm 0.17$	$11.05 \pm 0.13$	$11.07 \pm 0.19$	$11.07 \pm 0.17$	$12.42 \pm 0.22$	$9.93 \pm 0.22$

g = 2

jobs
parallel
Perfectly
(a)

g = 2	DPNEXTFAILUREGLOBAL		$209.88 \pm 6.78$	$107.85 \pm 4.11$	$57.51 \pm 2.03$	$31.90 \pm 0.83$	$15.07 \pm 0.27$
	DPNEXTFAILURESYNCHRO		$243.36 \pm 6.37$	$126.29 \pm 2.76$	$66.41 \pm 1.43$	$37.20 \pm 0.70$	$18.95 \pm 0.28$
	SAP	MAX	$225.78 \pm 4.89$	$116.68 \pm 2.20$	$61.76 \pm 1.35$	$34.47 \pm 0.65$	$16.86 \pm 0.30$
	DPNEXTFAILUREAS	AVG	$225.77 \pm 4.89$	$116.68 \pm 2.21$	$61.76 \pm 1.35$	$34.47 \pm 0.65$	$16.77 \pm 0.15$
		MIN	$225.85 \pm 4.83$	$116.71 \pm 2.26$	$61.83 \pm 1.29$	$34.52 \pm 0.69$	$16.83 \pm 0.25$
	PERIODLB		$188.77 \pm 5.90$	$97.18 \pm 3.25$	$51.71 \pm 1.53$	$28.85 \pm 0.74$	$14.33 \pm 0.20$
	OPTEXPGROUP		$191.80 \pm 7.74$	$98.82 \pm 3.90$	$52.64 \pm 1.82$	$29.27 \pm 0.87$	$14.38 \pm 0.55$
	OPTEXP		$192.02 \pm 4.74$	$99.09 \pm 2.18$	$52.79 \pm 1.16$	$29.27 \pm 0.60$	$14.54 \pm 0.32$
	DPNEXTFAILURE		$131.29 \pm 6.10$	$69.82 \pm 2.89$	$38.63 \pm 1.32$	$23.62 \pm 0.78$	$16.02 \pm 0.46$
g = 1	PERIODLB		$130.99 \pm 5.79$	$69.62 \pm 2.36$	$38.55 \pm 1.29$	$23.55 \pm 0.78$	$15.92 \pm 0.48$
	OPTEXP		$131.05 \pm 6.31$	$69.65 \pm 2.83$	$38.61 \pm 1.40$	$23.59 \pm 0.81$	$15.95 \pm 0.48$
	d		65536	131072	262144	524288	1048576

jobs.
parallel
Generic ]
(q)

	DPNEXTFAILUREGLOBAL		$203.63 \pm 7.03$	$102.19 \pm 3.63$	$51.24 \pm 1.82$	$25.45 \pm 0.68$	$10.03 \pm 0.22$
	DPNEXTFAILURESYNCHRO		$236.52 \pm 6.59$	$118.77 \pm 2.60$	$59.21 \pm 1.41$	$29.64 \pm 0.59$	$12.55 \pm 0.21$
	SAP	MAX	$219.10 \pm 4.63$	$109.95 \pm 2.18$	$54.90 \pm 1.20$	$27.48 \pm 0.61$	$11.19 \pm 0.20$
g = 2	NEXTFAILUREAS	AVG	$219.10 \pm 4.63$	$109.97 \pm 2.20$	$54.89 \pm 1.20$	$27.48 \pm 0.61$	$11.19 \pm 0.21$
	DPI	MIN	$219.24 \pm 4.56$	$110.01 \pm 2.22$	$55.05 \pm 1.20$	$27.51 \pm 0.60$	$11.16 \pm 0.13$
	PERIODLB		$183.65 \pm 5.74$	$91.67 \pm 3.32$	$45.91 \pm 1.51$	$23.08 \pm 0.70$	$9.52 \pm 0.17$
	OPTEXPGROUP		$186.73 \pm 7.69$	$93.03 \pm 3.72$	$46.89 \pm 1.71$	$23.44 \pm 0.92$	$9.56 \pm 0.40$
	OPTEXP		$186.39 \pm 4.73$	$93.32 \pm 2.17$	$46.91 \pm 1.04$	$23.37 \pm 0.58$	$9.65\pm0.23$
	DPNEXTFAILURE		$124.05 \pm 5.98$	$62.37 \pm 2.89$	$30.88 \pm 1.14$	$15.69 \pm 0.59$	$7.99 \pm 0.31$
g = 1	PeriodLB		$123.60 \pm 6.08$	$62.26 \pm 2.76$	$30.83 \pm 1.21$	$15.65 \pm 0.52$	$7.95 \pm 0.32$
	OPTEXP		$123.76 \pm 5.94$	$62.33 \pm 2.83$	$30.83 \pm 1.21$	$15.65 \pm 0.58$	$7.95\pm0.32$
	d		65536	131072	262144	524288	1048576

(c) Numerical kernels.

Table 3: Evaluation of the different heuristics on a platform with **Exponential failures** (MTBF = 125 years) under the proportional overhead model.

g = 1

	_						_	
g = 2	DPNEXTFAILUREGLOBAL		$232.66 \pm 0.75$	$121.11 \pm 0.63$	$65.51 \pm 0.56$	$37.85 \pm 0.52$	$24.37 \pm 0.50$	$17.98 \pm 0.76$
	DPNEXTFAILURESYNCHRO		$277.78 \pm 1.27$	$150.89 \pm 0.91$	$84.39 \pm 0.72$	$48.52 \pm 0.56$	$28.73 \pm 0.38$	$18.84 \pm 0.44$
	NEXTFAILUREASAP	MAX	$252.38 \pm 1.52$	$133.53 \pm 0.55$	$74.22 \pm 0.57$	$42.87 \pm 0.88$	$26.43 \pm 0.56$	$18.71 \pm 0.98$
		AVG	$252.17 \pm 1.06$	$133.52 \pm 0.52$	$74.41 \pm 0.94$	$42.76 \pm 0.34$	$26.43 \pm 0.58$	$18.52 \pm 0.43$
	DPI	MIN	$251.98 \pm 1.46$	$133.43 \pm 0.49$	$74.34 \pm 0.84$	$42.83 \pm 0.50$	$26.60 \pm 0.89$	$18.53 \pm 0.43$
	PERIODLB		$236.06 \pm 0.82$	$122.49 \pm 0.49$	$65.29 \pm 0.45$	$36.41 \pm 0.33$	$21.98 \pm 0.34$	$15.38 \pm 0.43$
	OPTEXPGROUP		$241.59 \pm 2.61$	$128.82 \pm 2.50$	$72.56 \pm 2.60$	$44.13 \pm 1.63$	$30.85 \pm 1.23$	$26.73 \pm 1.92$
	OPTEXP		$236.16 \pm 0.87$	$122.54 \pm 0.85$	$65.51 \pm 0.95$	$37.07 \pm 0.53$	$23.00 \pm 0.58$	$17.16 \pm 0.77$
g = 1	DPNEXTFAILURE		$137.51 \pm 1.22$	$76.37 \pm 0.81$	$45.21 \pm 0.66$	$29.52 \pm 0.59$	$22.72 \pm 0.70$	$23.96 \pm 1.04$
	PeriodLB		$137.19 \pm 1.14$	$76.17 \pm 0.79$	$44.92 \pm 0.64$	$29.16 \pm 0.56$	$22.49 \pm 0.62$	$23.67 \pm 1.01$
	OPTEXP		$142.66 \pm 1.91$	$80.44 \pm 1.45$	$48.93 \pm 1.25$	$33.15 \pm 1.25$	$27.43 \pm 1.45$	$31.83 \pm 1.93$
	d		32768	65536	131072	262144	524288	1048576

RR n° 7876

jobs.
parallel
Perfectly
۳

_	_	_	_	_	_			_
	DPNEXTFAILUREGLOBAL		$236.52 \pm 0.69$	$125.06 \pm 0.62$	$69.77 \pm 0.59$	$42.77 \pm 0.53$	$30.66 \pm 0.49$	$27.28 \pm 0.73$
	DPNEXTFAILURESYNCHRO		$282.32 \pm 1.29$	$155.87 \pm 1.02$	$89.92 \pm 0.77$	$54.77 \pm 0.59$	$36.17 \pm 0.45$	$28.70 \pm 0.51$
	SAP	MAX	$256.59 \pm 0.76$	$137.79 \pm 0.35$	$79.01 \pm 0.29$	$48.54 \pm 0.82$	$33.22 \pm 0.33$	$28.18 \pm 0.46$
g = 2	NEXTFAILUREAS	AVG	$256.44 \pm 0.96$	$137.74 \pm 0.36$	$79.06 \pm 0.40$	$48.38 \pm 0.40$	$33.22 \pm 0.31$	$28.26 \pm 1.02$
	DP	MIN	$255.91 \pm 0.80$	$137.70 \pm 0.37$	$79.25 \pm 1.06$	$48.43 \pm 0.57$	$33.16 \pm 0.32$	$28.47 \pm 1.46$
	PERIODLB		$239.98 \pm 0.84$	$126.52 \pm 0.52$	$69.54 \pm 0.49$	$41.11 \pm 0.40$	$27.70 \pm 0.34$	$23.41 \pm 0.52$
	OPTEXPGROUP		$245.59 \pm 2.66$	$132.99 \pm 2.50$	$77.40 \pm 2.75$	$50.12 \pm 1.71$	$38.49 \pm 1.37$	$39.94 \pm 2.13$
	OPTEXP		$240.03 \pm 0.87$	$126.54 \pm 0.89$	$69.78 \pm 1.00$	$41.84 \pm 0.61$	$29.02 \pm 0.50$	$25.99 \pm 0.98$
	DPNEXTFAILURE		$142.06 \pm 1.22$	$81.30 \pm 0.78$	$51.13 \pm 0.74$	$37.16 \pm 0.77$	$34.37 \pm 0.86$	$48.12 \pm 1.48$
g = 1	PeriodLB		$141.72 \pm 1.18$	$81.10 \pm 0.81$	$50.76 \pm 0.66$	$36.69 \pm 0.68$	$34.07 \pm 0.78$	$47.57 \pm 1.38$
	OPTEXP		$147.35 \pm 1.87$	$85.76 \pm 1.43$	$55.16 \pm 1.29$	$41.62 \pm 1.47$	$41.25 \pm 1.93$	$63.14 \pm 2.71$
	d		32768	65536	131072	262144	524288	1048576

jobs.
parallel
Generic 1
(q)

	DPNEXTFAILUREGLOBAL		$233.11 \pm 0.74$	$121.45 \pm 0.63$	$65.74 \pm 0.54$	$38.10 \pm 0.48$	$24.53 \pm 0.53$	$18.23 \pm 0.68$	
	DPNEXTFAILURESYNCHRO		$278.28 \pm 1.32$	$151.34 \pm 0.90$	$84.73 \pm 0.71$	$48.75 \pm 0.51$	$28.95 \pm 0.38$	$19.02 \pm 0.42$	
	SAP	MAX	$252.92 \pm 1.29$	$133.75 \pm 0.46$	$74.71 \pm 1.28$	$43.07 \pm 0.82$	$26.64 \pm 0.48$	$18.75 \pm 0.46$	
g = 2	NEXTFAILUREAS	AVG	$252.74 \pm 0.93$	$133.80 \pm 0.47$	$74.57 \pm 0.62$	$43.04 \pm 0.51$	$26.59 \pm 0.35$	$18.68 \pm 0.41$	
	DPI	MIN	$252.38 \pm 0.71$	$133.61 \pm 0.34$	$74.52 \pm 0.68$	$43.14 \pm 0.76$	$26.47 \pm 0.29$	$18.74 \pm 0.44$	
	PERIODLB		$236.51 \pm 0.85$	$122.83 \pm 0.50$	$65.53 \pm 0.45$	$36.59 \pm 0.36$	$22.13 \pm 0.32$	$15.56 \pm 0.41$	
	OPTEXPGROUP		$242.03 \pm 2.62$	$129.19 \pm 2.47$	$72.96 \pm 2.63$	$44.42 \pm 1.61$	$31.09 \pm 1.22$	$27.07 \pm 2.02$	
	OPTEXP		$236.61 \pm 0.88$	$122.85 \pm 0.84$	$65.77 \pm 0.94$	$37.27 \pm 0.54$	$23.18\pm0.58$	$17.34\pm0.78$	
	DPNEXTFAILURE		$137.93 \pm 1.27$	$76.67 \pm 0.78$	$45.39 \pm 0.64$	$29.70 \pm 0.63$	$22.99 \pm 0.75$	$24.37 \pm 1.09$	
g = 1	PERIODLB		$137.57 \pm 1.14$	$76.51 \pm 0.79$	$45.14 \pm 0.62$	$29.36 \pm 0.60$	$22.72 \pm 0.64$	$24.00 \pm 1.06$	
	OPTEXP		$143.09 \pm 1.86$	$80.83 \pm 1.45$	$49.15 \pm 1.29$	$33.37 \pm 1.27$	$27.73 \pm 1.45$	$32.34 \pm 1.96$	
	d		32768	65536	131072	262144	524288	1048576	

(c) Numerical kernels.

Table 4: Evaluation of the different heuristics on a platform with Weibull failures (MTBF = 125 years and k = 0.70) under the constant overhead model.

	NEXTFAILUREGLOBAL		$222.85 \pm 6.33$	$118.00 \pm 4.89$	$60.17 \pm 2.20$	$17.98 \pm 0.76$	
	DPNEXTFAILURESYNCHRO DI		$227.19 \pm 6.11$	$118.59 \pm 4.07$	$60.66 \pm 1.93$	$18.84 \pm 0.44$	
	AP ]	MAX	$242.39 \pm 5.46$	$126.06 \pm 3.83$	$64.40 \pm 1.44$	$18.71 \pm 0.98$	
g = 2	JEXTFAILUREAS.	AVG	$242.39 \pm 5.46$	$126.01 \pm 3.85$	$64.41 \pm 1.44$	$18.52 \pm 0.43$	
	DPN	NIIN	$238.97 \pm 5.09$	$125.47 \pm 3.52$	$64.42 \pm 1.44$	$18.53 \pm 0.43$	
	PERIODLB		$203.08 \pm 8.24$	$106.70 \pm 3.79$	$54.61 \pm 1.82$	$15.38 \pm 0.43$	
	OPTEXPGROUP		$350.20 \pm 9.66$	$205.92 \pm 13.43$	$113.18 \pm 9.20$	$26.73 \pm 1.92$	
	OPTEXP		$237.38 \pm 8.86$	$129.28 \pm 6.33$	$66.97 \pm 3.79$	$17.16 \pm 0.77$	
g = 1	DPNEXTFAILURE		$167.82 \pm 7.82$	$87.41 \pm 3.84$	$45.36 \pm 2.03$	$23.96 \pm 1.04$	
	PERIODLB		$167.58 \pm 7.60$	$87.20 \pm 3.61$	$45.14 \pm 2.15$	$23.67 \pm 1.01$	
	OPTEXP		$217.06 \pm 12.68$	$115.85 \pm 7.53$	$59.86 \pm 4.36$	$31.83 \pm 1.93$	
	d		131072	262144	524288	1048576	

jobs.	
parallel	
Perfectly	

	DPNEXTFAILUREGLOBAL		$232.91 \pm 5.22$	$132.56 \pm 4.34$	$75.59 \pm 2.54$	$27.28 \pm 0.73$
	DPNEXTFAILURESYNCHRO		$239.30 \pm 4.57$	$133.53 \pm 3.84$	$76.21 \pm 2.09$	$28.70 \pm 0.51$
	SAP	MAX	$255.12 \pm 3.11$	$142.10 \pm 3.67$	$80.77 \pm 1.74$	$28.18 \pm 0.46$
g = 2	NEXTFAILUREA	AVG	$255.12 \pm 3.11$	$142.12 \pm 3.66$	$80.80 \pm 1.71$	$28.26 \pm 1.02$
	DPI	MIN	$252.06 \pm 3.99$	$141.29 \pm 3.54$	$80.41 \pm 1.79$	$28.47 \pm 1.46$
	PERIODLB		$211.96 \pm 5.08$	$120.47 \pm 4.12$	$68.58 \pm 2.13$	$23.41 \pm 0.52$
	OPTEXPGROUP		$351.98 \pm 10.85$	$228.57 \pm 12.62$	$140.11 \pm 10.47$	$39.94 \pm 2.13$
	OPTEXP		$247.34 \pm 5.93$	$144.93 \pm 6.64$	$83.63 \pm 4.15$	$25.99 \pm 0.98$
	DPNEXTFAILURE		$188.28 \pm 8.08$	$109.18 \pm 4.25$	$68.04 \pm 2.58$	$48.12 \pm 1.48$
g = 1	PERIODLB		$188.05 \pm 7.78$	$108.86 \pm 3.98$	$67.80 \pm 2.38$	$47.57 \pm 1.38$
	OPTEXP		$241.53 \pm 13.87$	$143.31 \pm 8.11$	$90.43 \pm 5.74$	$63.14 \pm 2.71$
	d		131072	262144	524288	1048576

	222	
Darale	TOTTO TOOL	-
Jududt)		
5	2	

	NEXTFAILUREGLOBAL		$223.56 \pm 6.34$	$118.53 \pm 4.84$	$60.70 \pm 2.22$	$18.23 \pm 0.68$
	DPNEXTFAILURESYNCHRO D		$228.14 \pm 5.83$	$119.17 \pm 4.04$	$61.13 \pm 1.88$	$19.02 \pm 0.42$
	SAP I	MAX	$243.41 \pm 5.41$	$126.80 \pm 3.77$	$64.86 \pm 1.49$	$18.75 \pm 0.46$
g = 2	VEXTFAILUREAS	AVG	$243.39 \pm 5.39$	$126.77 \pm 3.78$	$64.85 \pm 1.51$	$18.68 \pm 0.41$
	DPI	MIN	$239.94 \pm 4.67$	$125.80 \pm 3.67$	$64.67 \pm 1.53$	$18.74 \pm 0.44$
	PERIODLB		$204.15 \pm 7.91$	$107.25 \pm 3.84$	$54.99 \pm 1.88$	$15.56 \pm 0.41$
	OPTEXPGROUP		$349.39 \pm 10.05$	$206.03 \pm 13.37$	$113.66 \pm 9.18$	$27.07 \pm 2.02$
	OPTEXP		$237.60 \pm 9.32$	$129.85 \pm 6.32$	$67.27 \pm 3.76$	$17.34 \pm 0.78$
	DPNEXTFAILURE		$168.34 \pm 7.78$	$88.00 \pm 3.78$	$45.85 \pm 2.02$	$24.37 \pm 1.09$
g = 1	PERIODLB		$168.14 \pm 7.49$	$87.85 \pm 3.52$	$45.61 \pm 2.06$	$24.00 \pm 1.06$
	OPTEXP		$217.99 \pm 12.89$	$116.18 \pm 7.16$	$60.63 \pm 4.37$	$32.34 \pm 1.96$
	d		131072	262144	524288	1048576

(c) Numerical kernels.

Table 5: Evaluation of the different heuristics on a platform with Weibull failures (MTBF = 125 years and k = 0.70) under the proportional overhead model.

DPNEXTFAILURESYNCHRO DPNEXTFAILUREGLOBAL			$\begin{array}{ c c c c c c c c c c c c c c c c c c c$	$\begin{array}{ c c c c c c c c c c c c c c c c c c c$				4.02 $234.28 \pm 1.75$ $197.38 \pm 0.87$ 2.63 $131.25 \pm 1.58$ $106.44 \pm 0.79$ $1.67$ $72.44 \pm 0.67$ $60.08 \pm 0.65$ $1.32$ $43.67 \pm 0.56$ $50.71 \pm 0.51$ $1.32$ $28.57 \pm 0.56$ $25.62 \pm 0.51$ $1.18$ $22.28 \pm 0.58$ $25.62 \pm 0.54$
ILUREASAP	G MAX	$\pm 4.22$ 216.48 $\pm 4$	$\pm 3.65$ 117.79 $\pm 2$	$\pm 1.75$ 67.34 $\pm 1$	$\pm 1.54$ 41.55 $\pm 1$	$\pm 1.23$ 28.06 $\pm 1$	$\pm 1.08$ 22.74 $\pm 1$	
DPNEXTF	MIN AV	$21.99 \pm 3.39$ 218.18	$22.11 \pm 2.80$ 119.21	$71.23 \pm 2.52$ 68.32	$44.06 \pm 1.63$ $42.26$	$30.02 \pm 1.49$ 28.43	$23.67 \pm 0.91$ 22.99	allel jobs.
PERIODLB		$198.95 \pm 0.90$ 2	$106.88 \pm 0.69$ 1	$60.23 \pm 0.59$	$35.74 \pm 0.50$	$23.64 \pm 0.45$	$18.94 \pm 0.44$	erfectly par
OPTEXPGROUP		$198.94 \pm 1.41$	$107.04 \pm 1.49$	$60.27 \pm 1.43$	$35.91 \pm 0.59$	$23.82 \pm 0.45$	$19.09 \pm 0.53$	(a) F
OPTEXP		$203.13 \pm 0.58$	$109.17 \pm 0.74$	$61.41 \pm 0.87$	$36.92 \pm 0.39$	$24.53 \pm 0.36$	$19.92 \pm 0.45$	
DPNEXTFAILURE		$120.94 \pm 1.45$	$69.80 \pm 0.92$	$44.12 \pm 0.80$	$31.76 \pm 0.88$	$28.22 \pm 1.15$	$38.75 \pm 1.70$	
y = 1 PERIODLB		$124.07 \pm 1.42$	$72.71 \pm 0.94$	$45.97 \pm 0.94$	$32.53 \pm 0.84$	$28.12 \pm 1.00$	$37.42 \pm 1.59$	
OPTEXP		$124.15 \pm 1.33$	$72.71 \pm 0.94$	$46.38 \pm 0.82$	$32.88 \pm 0.78$	$29.12 \pm 0.96$	$40.30 \pm 1.70$	
d		4096	8192	16384	32768	65536	131072	

	<b>DPNEXTFAILUREGLOBAL</b>		$197.50 \pm 0.87$	$106.57 \pm 0.72$	$60.21 \pm 0.70$	$36.88 \pm 0.52$	$25.78 \pm 0.54$	$22.97 \pm 0.74$
	DPNEXTFAILURESYNCHRO ]		$234.33 \pm 1.78$	$131.26 \pm 1.44$	$72.69 \pm 0.77$	$43.91 \pm 0.58$	$28.79 \pm 0.58$	$22.59 \pm 0.59$
	SAP	MAX	$216.09 \pm 3.64$	$117.90 \pm 3.01$	$67.99 \pm 1.70$	$41.87 \pm 1.37$	$28.47 \pm 1.55$	$22.95 \pm 1.18$
g = 2	NEXTFAILUREAS	AVG	$218.33 \pm 3.89$	$119.00 \pm 3.09$	$68.71 \pm 1.77$	$42.74 \pm 1.63$	$28.72 \pm 1.37$	$23.29 \pm 1.11$
	DPI	MIN	$222.15 \pm 4.37$	$122.45 \pm 3.50$	$70.85 \pm 1.97$	$44.31 \pm 1.74$	$30.21 \pm 1.74$	$24.05 \pm 1.07$
	PERIODLB		$199.02 \pm 0.89$	$107.00 \pm 0.67$	$60.35 \pm 0.47$	$35.91 \pm 0.51$	$23.77 \pm 0.45$	$19.15 \pm 0.60$
	OPTEXPGROUP		$199.05 \pm 1.40$	$107.14 \pm 1.49$	$60.37 \pm 1.44$	$36.08 \pm 0.58$	$23.98 \pm 0.44$	$19.39 \pm 0.54$
	OptExp		$203.23 \pm 0.57$	$109.27 \pm 0.74$	$61.54 \pm 0.87$	$37.07 \pm 0.40$	$24.71 \pm 0.36$	$20.40 \pm 0.48$
	DPNEXTFAILURE		$121.15 \pm 1.63$	$69.95 \pm 0.92$	$44.33 \pm 0.82$	$31.97 \pm 0.88$	$28.67 \pm 1.18$	$40.27 \pm 1.85$
g = 1	PERIODLB		$124.22 \pm 1.41$	$72.88 \pm 0.93$	$46.17 \pm 0.93$	$32.76 \pm 0.85$	$28.56 \pm 1.02$	$38.82\pm1.78$
	OPTEXP		$124.27 \pm 1.32$	$72.88 \pm 0.93$	$46.57 \pm 0.80$	$33.13 \pm 0.79$	$29.94 \pm 1.01$	$41.80 \pm 1.71$
	р		4096	8192	16384	32768	65536	131072

jobs.
parallel
Generic ]
(q

	DPNEXTFAILUREGLOBAL		$197.66 \pm 0.82$	$106.73 \pm 0.89$	$60.22 \pm 0.66$	$36.84 \pm 0.50$	$25.73 \pm 0.51$	$22.85 \pm 0.77$
	DPNEXTFAILURESYNCHRO		$234.46 \pm 1.61$	$131.47 \pm 1.58$	$72.60 \pm 0.68$	$43.90 \pm 0.54$	$28.68 \pm 0.57$	$22.41 \pm 0.58$
	AP	MAX	$216.89 \pm 3.71$	$118.16 \pm 3.06$	$67.83 \pm 1.58$	$41.88 \pm 1.49$	$28.07 \pm 1.26$	$22.83 \pm 1.02$
g = 2	NEXTFAILUREAS	AVG	$218.15 \pm 3.97$	$119.05 \pm 3.12$	$68.65 \pm 1.92$	$42.34 \pm 1.56$	$28.56 \pm 1.27$	$23.18 \pm 1.17$
	DPI	MIN	$221.89 \pm 3.01$	$122.93 \pm 2.61$	$71.36 \pm 1.90$	$44.58 \pm 1.37$	$29.87 \pm 1.39$	$23.89 \pm 0.98$
	PERIODLB		$199.15 \pm 0.89$	$107.01 \pm 0.68$	$60.36 \pm 0.59$	$35.85 \pm 0.51$	$23.72 \pm 0.45$	$18.99 \pm 0.60$
	OPTEXPGROUP		$199.19 \pm 1.40$	$107.21 \pm 1.49$	$60.40 \pm 1.42$	$36.00 \pm 0.60$	$23.90 \pm 0.45$	$19.23 \pm 0.53$
	OPTEXP		$203.37 \pm 0.57$	$109.33 \pm 0.73$	$61.56 \pm 0.85$	$37.05 \pm 0.39$	$24.63 \pm 0.36$	$20.22 \pm 0.46$
	DPNEXTFAILURE		$121.01 \pm 1.48$	$70.03 \pm 0.94$	$44.28 \pm 0.83$	$31.86 \pm 0.88$	$28.40 \pm 1.17$	$39.14 \pm 1.73$
g = 1	PERIODLB		$124.23 \pm 1.40$	$72.90 \pm 0.92$	$46.19 \pm 0.94$	$32.64 \pm 0.84$	$28.24 \pm 0.99$	$37.86 \pm 1.66$
	OPTEXP		$124.28 \pm 1.32$	$72.90 \pm 0.92$	$46.58 \pm 0.80$	$32.99 \pm 0.80$	$29.63 \pm 0.98$	$40.73 \pm 1.69$
	d		4096	8192	16384	32768	65536	131072

(c) Numerical kernels.

Table 6: Evaluation of the different heuristics on a platform with failures based on the failure log of LANL cluster 18 and under under the constant overhead model.

		_		<b></b>		_	,,			r	,,		_						
DPNEXTFAILUREGLOBAL		$331.38 \pm 17.67$	$131.63 \pm 7.76$	$22.68 \pm 0.72$			DPNEXTFAILUREGLOBAL		$332.71 \pm 18.34$	$132.50 \pm 7.13$	$22.97 \pm 0.74$			DPNEXTFAILUREGLOBAL		$332.25 \pm 18.59$	$132.28 \pm 8.05$	$22.85 \pm 0.77$	
DPNEXTFAILURESYNCHRO		$310.68 \pm 15.33$	$123.99 \pm 5.89$	$22.28 \pm 0.58$			DPNEXTFAILURESYNCHRO		$312.25 \pm 15.23$	$124.70 \pm 5.72$	$22.59 \pm 0.59$			DPNEXTFAILURESYNCHRO		$311.75 \pm 15.37$	$124.56 \pm 5.92$	$22.41 \pm 0.58$	
AP	MAX	$352.10 \pm 25.20$	$127.85 \pm 5.90$	$22.74 \pm 1.18$			AP	MAX	$353.95 \pm 25.44$	$129.12 \pm 5.76$	$22.95 \pm 1.18$			AP	MAX	$352.72 \pm 25.39$	$128.63 \pm 5.90$	$22.83 \pm 1.02$	
NEXTFAILUREAS	AVG	$339.30 \pm 24.99$	$128.67 \pm 5.64$	$22.99 \pm 1.08$		g = 2	NEXTFAILUREAS	AVG	$342.01 \pm 25.19$	$129.59 \pm 5.42$	$23.29 \pm 1.11$		g = 2	NEXTFAILUREAS	AVG	$340.34 \pm 24.77$	$129.21 \pm 5.85$	$23.18 \pm 1.17$	
DF	NIIN	$337.50 \pm 20.77$	$134.09 \pm 5.64$	$23.67 \pm 0.91$	trallel jobs.		DP	MIN	$340.51 \pm 20.16$	$135.11 \pm 5.63$	$24.05 \pm 1.07$	rallel jobs.		DP	MIN	$337.58 \pm 19.12$	$134.37 \pm 5.35$	$23.89 \pm 0.98$	l kernels.
PERIODLB		$285.64 \pm 13.78$	$112.43 \pm 4.79$	$18.94 \pm 0.44$	erfectly pe		PeriodLB		$287.54 \pm 13.83$	$113.28 \pm 4.74$	$19.15 \pm 0.60$	Generic pa		PERIODLB		$287.23 \pm 13.58$	$112.79 \pm 4.85$	$18.99 \pm 0.60$	Numerica
OPTEXPGROUP		$286.19 \pm 13.42$	$114.67 \pm 4.69$	$19.09 \pm 0.53$	(a) F		OPTEXPGROUP		$287.96 \pm 13.35$	$115.60 \pm 4.65$	$19.39 \pm 0.54$	(p) (q)		OPTEXPGROUP		$287.83 \pm 13.45$	$115.18 \pm 4.72$	$19.23 \pm 0.53$	(c)
OPTEXP		$292.56 \pm 11.30$	$119.10 \pm 4.48$	$19.92 \pm 0.45$			OPTEXP		$293.69 \pm 11.52$	$120.10 \pm 4.45$	$20.40 \pm 0.48$			OPTEXP		$293.58 \pm 11.58$	$119.55 \pm 4.50$	$20.22 \pm 0.46$	
DPNEXTFAILURE		$211.07 \pm 11.97$	$84.09 \pm 4.40$	$38.75 \pm 1.70$			DPNEXTFAILURE		$212.77 \pm 12.42$	$85.64 \pm 4.55$	$40.27 \pm 1.85$			DPNEXTFAILURE		$211.32 \pm 12.86$	$84.70 \pm 4.42$	$39.14 \pm 1.73$	
PERIODLB		$218.36 \pm 13.79$	$84.15 \pm 4.39$	$37.42 \pm 1.59$		g = 1	PeriodLB		$219.93 \pm 13.89$	$86.16 \pm 4.88$	$38.82 \pm 1.78$		g = 1	PERIODLB		$219.04 \pm 13.36$	$85.12 \pm 4.69$	$37.86 \pm 1.66$	
OPTEXP		$226.77 \pm 12.39$	$89.48 \pm 4.12$	$40.30 \pm 1.70$			OPTEXP		$228.66 \pm 12.25$	$91.06 \pm 4.29$	$41.80 \pm 1.71$			OPTEXP		$226.94 \pm 12.48$	$89.88 \pm 4.15$	$40.73 \pm 1.69$	
d		32768	65536	131072			d		32768	65536	131072			d		32768	65536	131072	
	_	-	_				_	_	_					_	_	-		_	

del.	
l mo	
thead	
ovei	
onal	
orti	
prot	
	proportional overhead model.

	5		ŝ	5	6	~	-1	5	(		5		5	5	2
	DPNEXTFAILUREGLOBAL		$196.71 \pm 0.85$	$106.39 \pm 0.8$	$59.81 \pm 0.50$	$36.84 \pm 0.5$	$25.21 \pm 0.5$	$21.45 \pm 0.7i$			DPNEXTFAILUREGLOBAL		$196.76 \pm 0.8^{\circ}$	$106.34 \pm 0.8$	50.80 + 0.55
	DPNEXTFAILURESYNCHRO		$233.54 \pm 1.80$	$130.38 \pm 1.57$	$72.15 \pm 0.63$	$43.54 \pm 0.54$	$28.22 \pm 0.56$	$21.48 \pm 0.72$			DPNEXTFAILURESYNCHRO		$233.77 \pm 1.86$	$130.28 \pm 1.52$	$79.97 \pm 0.63$
	SAP	MAX	$216.53 \pm 4.32$	$118.55 \pm 2.32$	$67.59 \pm 2.00$	$41.32 \pm 1.21$	$27.99 \pm 1.40$	$21.32 \pm 0.85$			SAP	MAX	$216.99 \pm 5.19$	$118.02 \pm 2.04$	$67.40 \pm 1.01$
g = 2	NEXTFAILUREA	AVG	$219.66 \pm 3.81$	$119.69 \pm 1.92$	$68.21 \pm 1.98$	$41.86 \pm 1.33$	$28.40 \pm 1.39$	$21.68 \pm 0.94$	÷	g = 2	NEXTFAILUREA	AVG	$219.79 \pm 3.74$	$120.09 \pm 1.92$	$68.41 \pm 2.00$
	DP	MIN	$224.81 \pm 5.58$	$122.92 \pm 2.72$	$71.57 \pm 2.56$	$44.26 \pm 2.14$	$29.78 \pm 1.08$	$22.95 \pm 1.24$	trallel jobs		DP	MIN	$224.16 \pm 4.98$	$123.09 \pm 2.93$	$71.60 \pm 9.19$
	PERIODLB		$198.97 \pm 0.92$	$106.98 \pm 0.57$	$59.88 \pm 0.46$	$35.96 \pm 0.42$	$23.01 \pm 0.48$	$17.92 \pm 0.57$	erfectly pe		PERIODLB		$199.07 \pm 0.91$	$107.09 \pm 0.58$	$60.09 \pm 0.47$
	<b>OPTEXPGROUP</b>		$198.87 \pm 1.39$	$107.05 \pm 1.36$	$60.08 \pm 1.37$	$36.00 \pm 0.42$	$23.06 \pm 0.50$	$18.24 \pm 0.52$	(a) P <sub>(</sub>		<b>OPTEXPGROUP</b>		$199.00 \pm 1.40$	$107.20 \pm 1.38$	$60.91 \pm 1.30$
	OPTEXP		$203.31 \pm 0.55$	$109.42 \pm 0.77$	$61.36 \pm 0.87$	$36.97 \pm 0.29$	$24.18 \pm 0.40$	$19.26 \pm 0.49$			OPTEXP		$203.42 \pm 0.55$	$109.52 \pm 0.76$	$61.47 \pm 0.88$
	DPNEXTFAILURE		$120.13 \pm 1.46$	$69.28 \pm 0.95$	$43.57 \pm 0.78$	$31.20 \pm 0.96$	$27.54 \pm 1.25$	$38.08 \pm 2.12$			DPNEXTFAILURE		$120.24 \pm 1.51$	$69.43 \pm 0.91$	$43.70 \pm 0.78$
g = 1	PERIODLB		$123.67 \pm 1.41$	$71.84 \pm 0.96$	$45.27 \pm 0.79$	$31.63 \pm 0.84$	$27.57 \pm 1.15$	$36.93 \pm 1.85$		g = 1	PERIODLB		$123.80 \pm 1.44$	$72.01 \pm 0.95$	$45.48 \pm 0.78$
	OPTEXP		$123.81 \pm 1.28$	$72.18 \pm 0.82$	$45.50 \pm 0.74$	$32.54 \pm 0.76$	$28.74 \pm 1.10$	$39.69 \pm 1.47$			OPTEXP		$123.93 \pm 1.30$	$72.33 \pm 0.81$	$45.73 \pm 0.68$
	d		4096	8192	16384	32768	65536	131072			d		4096	8192	16384
_	_		_	_	_	_	_	_		_			_		_

р	OPTEAP	LEKIODLE .	DFINEXTFAILURE	OPTEXP	OPTEXPGROUP	LEKIODLE	DF	NEXTFAILUREAU	SAP	DFINEXTFAILURES YNCHRO	DFINEXTFAILUREGLOBAL	
							MIN	AVG	MAX			
4096	$123.93 \pm 1.30$	$123.80 \pm 1.44$	$120.24 \pm 1.51$	$203.42 \pm 0.55$	$199.00 \pm 1.40$	$199.07 \pm 0.91$	$224.16 \pm 4.98$	$219.79 \pm 3.74$	$216.99 \pm 5.19$	$233.77 \pm 1.86$	$196.76 \pm 0.82$	
8192	$72.33 \pm 0.81$	$72.01 \pm 0.95$	$69.43 \pm 0.91$	$109.52 \pm 0.76$	$107.20 \pm 1.38$	$107.09 \pm 0.58$	$123.09 \pm 2.93$	$120.09 \pm 1.92$	$118.02 \pm 2.04$	$130.28 \pm 1.52$	$106.34 \pm 0.82$	
16384	$45.73 \pm 0.68$	$45.48 \pm 0.78$	$43.79 \pm 0.78$	$61.47 \pm 0.88$	$60.21 \pm 1.39$	$60.02 \pm 0.47$	$71.60 \pm 2.12$	$68.41 \pm 2.00$	$67.49 \pm 1.91$	$72.27 \pm 0.63$	$59.89 \pm 0.55$	
32768	$32.82 \pm 0.77$	$31.90 \pm 0.85$	$31.49 \pm 0.94$	$37.11 \pm 0.29$	$36.07 \pm 0.41$	$36.07 \pm 0.41$	$44.16 \pm 2.01$	$42.09 \pm 1.65$	$41.51 \pm 1.34$	$43.70 \pm 0.59$	$36.98 \pm 0.53$	
65536	$29.27 \pm 1.11$	$28.04 \pm 1.20$	$27.98 \pm 1.26$	$24.39 \pm 0.40$	$23.26 \pm 0.51$	$23.19 \pm 0.50$	$30.06 \pm 1.07$	$28.56 \pm 1.43$	$28.16 \pm 1.40$	$28.44 \pm 0.57$	$25.40 \pm 0.59$	
131072	$40.93 \pm 1.57$	$38.12 \pm 1.79$	$39.26 \pm 2.20$	$19.59 \pm 0.48$	$18.54 \pm 0.52$	$18.23 \pm 0.56$	$23.35 \pm 1.14$	$22.09 \pm 0.96$	$21.56 \pm 0.89$	$21.82 \pm 0.73$	$21.90 \pm 0.84$	

jobs.	,
oarallel	
eneric p	
9 (9	`

	OBAL		± 0.74	± 0.74	$\pm 0.55$	$\pm 0.50$	$\pm 0.59$	± 0.84	
	DPNEXTFAILUREGI		196.94 -	106.51 -	59.87 -	36.92 =	25.33 =	21.60 =	
	DPNEXTFAILURESYNCHRO		$234.13 \pm 1.92$	$130.39 \pm 1.53$	$72.28 \pm 0.65$	$43.63 \pm 0.54$	$28.36 \pm 0.57$	$21.60 \pm 0.73$	
	SAP	MAX	$217.95 \pm 4.53$	$118.64 \pm 2.19$	$67.73 \pm 2.00$	$41.54 \pm 1.31$	$28.14 \pm 1.41$	$21.37 \pm 0.84$	
g = 2	NEXTFAILUREA	AVG	$219.70 \pm 3.98$	$120.10 \pm 2.08$	$68.95 \pm 1.83$	$41.93 \pm 1.28$	$28.46 \pm 1.33$	$21.80\pm0.87$	
	ΔU	MIN	$224.18 \pm 4.70$	$123.37 \pm 2.12$	$71.17 \pm 2.71$	$44.01 \pm 1.59$	$29.94 \pm 0.99$	$23.21 \pm 1.11$	
	PeriodLB		$199.20 \pm 0.90$	$107.14 \pm 0.56$	$60.02 \pm 0.46$	$36.05 \pm 0.36$	$23.10 \pm 0.49$	$18.03 \pm 0.57$	
	OPTEXPGROUP		$199.13 \pm 1.39$	$107.24 \pm 1.33$	$60.21 \pm 1.36$	$36.07 \pm 0.41$	$23.17 \pm 0.50$	$18.34 \pm 0.51$	
	OPTEXP		$203.53 \pm 0.56$	$109.59 \pm 0.76$	$61.48 \pm 0.86$	$37.08 \pm 0.29$	$24.30 \pm 0.40$	$19.40 \pm 0.48$	
	DPNEXTFAILURE		$120.39 \pm 1.51$	$69.42 \pm 0.96$	$43.73 \pm 0.78$	$31.38 \pm 0.96$	$27.70 \pm 1.23$	$38.43 \pm 2.22$	
g = 1	PERIODLB		$123.90 \pm 1.44$	$72.00 \pm 0.95$	$45.42 \pm 0.79$	$31.78 \pm 0.85$	$27.77\pm1.18$	$37.25 \pm 1.82$	
	OPTEXP		$124.01 \pm 1.31$	$72.31 \pm 0.82$	$45.68 \pm 0.70$	$32.69 \pm 0.77$	$28.98 \pm 1.09$	$40.00 \pm 1.50$	
	d		4096	8192	16384	32768	65536	131072	

(c) Numerical kernels.

Table 8: Evaluation of the different heuristics on a platform with failures based on the failure log of LANL cluster 19 and under the constant overhead model.

TUTOTOT		$11 \pm 20.91$	$.70 \pm 6.52$	$.45 \pm 0.75$				reglobal		$95 \pm 20.76$	$3.05 \pm 6.62$	$.90 \pm 0.84$			reglobal		$87 \pm 20.83$	$0.37 \pm 6.45$	$.60 \pm 0.84$	
IOTIVITY IOT		329.	121	21				DPNEXTFAILUI		329.	125	21			DPNEXTFAILUI		329.	122	21	
CHICALL CHICALLY INTELL 17		$299.02 \pm 16.21$	$114.52 \pm 5.61$	$21.48 \pm 0.72$				DPNEXTFAILURESYNCHRO		$300.54 \pm 16.25$	$115.56 \pm 5.61$	$21.82 \pm 0.73$			DPNEXTFAILURESYNCHRO		$300.28 \pm 16.44$	$115.05 \pm 5.59$	$21.60 \pm 0.73$	
e.	MAX	$319.79 \pm 21.61$	$117.21 \pm 5.33$	$21.32 \pm 0.85$				AP	MAX	$321.32 \pm 21.78$	$118.34 \pm 5.34$	$21.56 \pm 0.89$			AP.	MAX	$322.73 \pm 23.52$	$117.84 \pm 5.47$	$21.37 \pm 0.84$	
	AVG	$311.87 \pm 16.09$	$119.30 \pm 5.83$	$21.68 \pm 0.94$			g = 2	NEXTFAILUREAS	AVG	$314.54 \pm 16.59$	$120.23 \pm 7.00$	$22.09 \pm 0.96$		g = 2	NEXTFAILUREAS	AVG	$313.46 \pm 16.92$	$120.07 \pm 6.93$	$21.80 \pm 0.87$	
14	MIN	$319.55 \pm 15.68$	$125.88 \pm 10.22$	$22.95 \pm 1.24$	rallel jobs.	5		DP	MIN	$320.80 \pm 16.98$	$125.85 \pm 9.64$	$23.35 \pm 1.14$	allel jobs.		DP	NIN	$322.35 \pm 16.87$	$125.56 \pm 9.92$	$23.21 \pm 1.11$	kernels.
THOOPIN T		$270.90 \pm 14.84$	$103.92 \pm 5.44$	$17.92 \pm 0.57$	erfectly pa	1		PERIODLB		$272.17 \pm 14.98$	$104.74 \pm 5.04$	$18.23 \pm 0.56$	deneric paı		PERIODLB		$272.25 \pm 15.13$	$104.24 \pm 5.41$	$18.03 \pm 0.57$	Numerical
IOCHIC IVET I IO		$270.90 \pm 14.84$	$105.29 \pm 4.59$	$18.24 \pm 0.52$	(a) P			OPTEXPGROUP		$272.17 \pm 14.98$	$108.24 \pm 4.90$	$18.54 \pm 0.52$	(b) (d)		OPTEXPGROUP		$272.25 \pm 15.13$	$107.79 \pm 4.85$	$18.34 \pm 0.51$	(c)
01 117		$282.16 \pm 15.71$	$109.64 \pm 4.65$	$19.26 \pm 0.49$				OPTEXP		$283.03 \pm 15.83$	$110.56 \pm 4.61$	$19.59 \pm 0.48$			OPTEXP		$283.17 \pm 15.92$	$110.15 \pm 4.60$	$19.40 \pm 0.48$	
		$191.61 \pm 12.91$	$77.98 \pm 4.38$	$38.08 \pm 2.12$				DPNEXTFAILURE		$192.73 \pm 13.45$	$79.50 \pm 4.66$	$39.26 \pm 2.20$			DPNEXTFAILURE		$192.61 \pm 12.83$	$78.57 \pm 4.51$	$38.43 \pm 2.22$	
angonin i		$194.84 \pm 9.98$	$77.91 \pm 4.04$	$36.93 \pm 1.85$			g = 1	PERIODLB		$196.63 \pm 10.59$	$79.36 \pm 4.30$	$38.12 \pm 1.79$		g = 1	PERIODLB		$196.02 \pm 10.15$	$78.42 \pm 4.07$	$37.25 \pm 1.82$	
WHI IO		$202.75 \pm 10.31$	$83.80 \pm 4.13$	$39.69 \pm 1.47$				OPTEXP		$205.09 \pm 10.62$	$85.36 \pm 4.23$	$40.93 \pm 1.57$			OPTEXP		$204.05 \pm 10.33$	$84.39 \pm 4.19$	$40.00 \pm 1.50$	
2		32768	65536	131072				d		32768	65536	131072			d		32768	65536	131072	

Table 9: Evaluation of the different heuristics on a platform with failures based on the failure log of LANL cluster 19 and under the proportional overhead model.



Figure 9: Evaluation of the different heuristics on a platform with **Exponential** failures (MTBF = 125 years).



Figure 10: Evaluation of the different heuristics on a platform with Weibull failures (MTBF = 125 years, and k = 0.70).



Figure 11: Evaluation of the different heuristics on a platform with failures based on the failure log of LANL cluster 18.



Figure 12: Evaluation of the different heuristics on a platform with failures based on the failure log of LANL cluster 19.



#### RESEARCH CENTRE GRENOBLE – RHÔNE-ALPES

Inovallée 655 avenue de l'Europe Montbonnot 38334 Saint Ismier Cedex Publisher Inria Domaine de Voluceau - Rocquencourt BP 105 - 78153 Le Chesnay Cedex inria.fr

ISSN 0249-6399