Identifying the right replication level to detect and correct silent errors at scale

Anne Benoit, Aurélien Cavelan, Franck Cappello, Padma Raghavan, Yves Robert, Hongyang Sun

RESEARCH REPORT

N° 9047

March 2017

Project-Team ROMA

ISSN 0249-6399 ISRN INRIA/RR--9047--FR+ENG

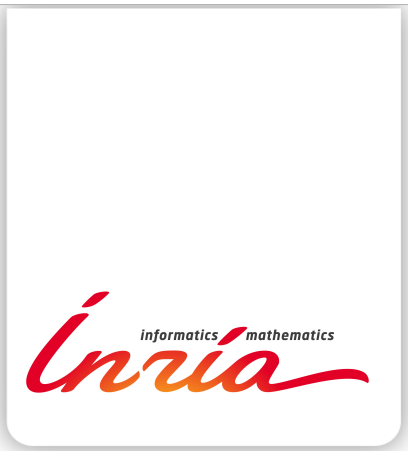

Identifying the right replication level to detect and correct silent errors at scale

Anne Benoit*, Aurélien Cavelan*, Franck Cappello†, Padma Raghavan‡, Yves Robert*§, Hongyang Sun‡

Project-Team ROMA

Research Report n° 9047 — March 2017 — 26 pages

* Ecole Normale Superieure de Lyon and Inria, France
† Argonne National Laboratory, USA
‡ Vanderbilt University, USA
§ Univ. Tenn. Knoxville, USA

RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Abstract: This paper provides a model and an analytical study of replication as a technique to detect and correct silent errors. Although other detection techniques exist for HPC applications, based on algorithms (ABFT), invariant preservation or data analytics, replication remains the most transparent and least intrusive technique. We explore the right level (duplication, triplication or more) of replication needed to efficiently detect and correct silent errors. Replication is combined with checkpointing and comes with two flavors: *process replication* and *group replication*. Process replication applies to message-passing applications with communicating processes. Each process is replicated, and the platform is composed of process pairs, or triplets. Group replication applies to black-box applications, whose parallel execution is replicated several times. The platform is partitioned into two halves (or three thirds). In both scenarios, results are compared before each checkpoint, which is taken only when both results (duplication) or two out of three results (triplication) coincide. If not, one or more silent errors have been detected, and the application rolls back to the last checkpoint. We provide a detailed analytical study of both scenarios, with formulas to decide, for each scenario, the optimal parameters as a function of the error rate, checkpoint cost, and platform size. We also report a set of extensive simulation results that corroborates the analytical model.

Key-words: resilience, replication, silent errors, silent data corruptions, SDC, detection, correction, duplication, triplication, voting, optimal, number of processors.

Quel est le bon niveau de réplication pour détecter et corriger les erreurs silencieuses?

Résumé : Ce rapport propose un modèle et une étude analytique de la réplication en tant que technique pour détecter et corriger les erreurs silencieuses. Bien que d'autres techniques existent pour les applications HPC, basées sur des algorithmes (ABFT), préservation d'invariant, ou analyse de données, la réplication reste la technique la plus transparente et la moins intrusive. Nous explorons le bon niveau (duplication, triplication ou plus) de réplication nécessaire pour détecter et corriger les erreurs silencieuses de manière efficace. La réplication est combinée avec des checkpoints et se présente sous deux formes : *réplication de processus* et *réplication de groupes*. La réplication de processus s'applique aux applications à passage de messages avec des processus communicants. Chaque processus est répliqué, et la plate-forme est composée de paires, ou triplets de processus. La réplication de groupe s'applique à des applications type boîte noire, dont l'exécution parallèle est répliquée plusieurs fois. La plate-forme est alors partitionnée en deux moitiés (ou trois tiers). Dans les deux scénarios, les résultats sont comparés avant chaque checkpoint, qui est effectué seulement lorsque les deux résultats (duplication) ou deux sur trois (triplication) coïncident. Sinon, une ou plusieurs erreurs silencieuses ont été détectées, et l'application redémarre depuis le dernier checkpoint. Nous proposons une étude analytique détaillée des deux scénarios ainsi que les paramètres optimaux fonction du taux d'erreur, du coût du checkpoint, et de la taille de la plate-forme. Nous donnons également les résultats d'un ensemble de simulations qui viennent corroborer le modèle analytique.

Mots-clés : résilience, réplication, erreurs silencieuses, duplication, triplication, détection, correction, nombre de processeurs, optimal.

1 Introduction

Triple Modular Redundancy, or TMR [33], is the standard fault-tolerance approach for critical systems, such as embedded or aeronautical devices [1]. With TMR, computations are executed three times, and a majority voting is conducted to select the correct result out of the three available ones. Indeed, if two or more results agree, they are declared correct, because the probability of two or more errors leading to the same wrong result is assumed so low that it can be ignored. While triplication seems very expensive in terms or resources, anybody sitting in a plane would heartily agree that it is worth the price.

On the contrary, duplication, let alone triplication, has a bad reputation in the High Performance Computing (HPC) community. Who would be ready to waste half or two-thirds of precious computing resources? However, despite its high cost, several authors have been advocating the use of duplication in HPC in the recent years [40, 48, 24, 26]. In a nutshell, this is because platform sizes have become so large that fail-stop errors are likely to strike at a high rate during application execution. More precisely, the MTBF (Mean Time Between Failures) μ_P of the platform decreases linearly with the number of processors P, since $\mu_P = \frac{\mu_{\text{ind}}}{P}$, where μ_{ind} is the MTBF of each individual component (see Proposition 1.2 in [30]). Take $\mu_{\text{ind}} = 10$ years as an example. If $P = 10^5$ then $\mu_P \approx 50$ minutes and if $P = 10^6$ then $\mu_P \approx 5$ minutes: from the point of view of fault-tolerance, scale is the enemy. Given any value of μ_{ind}, there is a threshold value for the number of processors above which platform throughput will decrease [23, 37, 41, 26]: the platform MTBF becomes so small that the applications experience too many failures, hence too many recoveries and re-execution delays, to progress efficiently. All this explains why duplication has been considered for HPC applications despite its cost. The authors in [26] propose *process replication* by which each process in a parallel MPI (Message Passing Interface) application is duplicated on multiple physical processors while maintaining synchronous execution of the replicas. This approach is effective because the MTBF of a set of two replicas (which is the average delay for failures to strike *both* processors in the replica set) is much larger than the MTBF of a single processor.

Process replication may not always be a feasible option. Process replication features must be provided by the application. Some prototype MPI implementations [26, 27] are convincing proofs of concept and do provide such capabilities. However, many other programming frameworks (not only MPI-like frameworks, but also concurrent objects, distributed components, workflows, algorithmic skeletons) do not provide an equivalent to transparent process replication for the purpose of fault-tolerance, and enhancing them with transparent replication may be non-trivial. When transparent replication is not (yet) provided by the runtime system, one solution could be to implement it explicitly within the application, but this is a labor-intensive process especially for legacy applications. Another approach introduced in [14] is *group replication*, a technique that can be used whenever process replication is not available. Group replication is agnostic to the parallel programming model, and thus views the application as an unmodified black box. The only requirement is that the application be startable from a saved checkpoint file. Group replication consists in executing multiple application instances concurrently. For example, 2 distinct P-process application instances could be executed on a $2P$-processor platform. At first glance, it may seem paradoxical that better performance can be achieved by using group duplication. After all, in the above example, 50% of the platform is "wasted" to perform redundant computation. The key point here is that each application instance runs at a smaller scale. As a result each instance can use lower checkpointing frequency, and can thus have better parallel efficiency in the presence of faults, when compared

to a single application instance running at full scale. In some cases, the application makespan can then be comparable to, or even shorter than that obtained when running a single application instance. In the end, the cost of wasting processor power for redundant computation can be offset by the benefit of reduced checkpointing frequency. Furthermore, in group replication, once an instance saves a checkpoint, the other instance can use this checkpoint immediately to "jump ahead" in its execution. Hence, group replication is more efficient than the mere independent execution of several instances: each time one instance successfully completes a given "chunk of work", all the other instances immediately benefit from this success. To implement group replication the runtime system needs to perform the typical operations needed for system-assisted checkpoint/restart: determining checkpointing frequencies for each application instance, causing checkpoints to be saved, detecting application failures, and restarting an application instance from a saved checkpoint after a failure. The only additional feature is that the system must be able to stop an instance and cause it to resume execution from a checkpoint file produced by another instance as soon as it is produced.

Process or group replication has been mainly proposed in HPC to cope with fail-stop errors. However, another challenge is represented by silent errors, or silent data corruption, whose threat can no longer be ignored [36, 49, 34]. There are several causes of silent errors, such as cosmic radiation, packaging pollution, among others. Silent errors can strike the cache and memory (bit flips) as well as CPU operations; in the latter case they resemble floating-point errors due to improper rounding, but have a dramatically larger impact because any bit of the result, not only low-order mantissa bits, can be corrupted. In contrast to a fail-stop error whose detection is immediate, a silent error is identified only when the corrupted data leads to an unusual application behavior. Such detection latency raises a new challenge: if the error struck before the last checkpoint, and is detected after that checkpoint, then the checkpoint is corrupted and cannot be used for rollback. To distinguish from fail-stop failures, we use MTBE instead of MTBF to characterize the rate of silent errors.

To address the problem of silent errors, many application-specific detectors, or verification mechanisms, have been proposed (see Section 2 for a survey). It is not clear, however, whether a special-purpose detector can be designed for each scientific application. In addition, application-specific verification mechanisms only protect from some types of error sources, and fail to provide accurate and efficient detection of all silent errors. In fact, providing such detectors for scientific applications has been identified as one of the hardest challenges[1] towards extreme-scale computing [12, 13].

Altogether, silent errors call for revisiting replication in the framework of scientific application executing on large-scale HPC platforms. Because replication is now applied at the process level, scale becomes an even harder-to-fight enemy. Processor count ranges to about 10^5 on the K-computer and TaihuLight systems. The number of processors could increase further to 10^6 (hence 10^6 or more processes) on Exascale systems, with billions of threads [20]. In addition, the probability of several errors striking during an execution can get significant, depending upon whether or not circuit manufacturers increase significantly the protection of the logic, latch/flip-flops and static arrays in the processor. In a recent paper [43], the authors consider that with significant more protection (more hardware, more power consumption), the FIT[2] rate for undetected errors on a processor circuit could be maintained to around 20.

[1] More generally, trustworthy computing, which aims at guaranteeing the correctness of the results of a long-lasting computation on a large-scale supercomputer, has received considerable attention recently [11].

[2] The Failures in Time (FIT) rate of a device is the number of failures that can be expected in one billion (10^9) device-hours of operation.

But without additional protection compared to the current situation, the FIT rate for undetected errors could be as high as 5,000 (or 1 error every 200,000 hours). Combining 10 million of devices with this FIT rate would result in a silent error in the system every 72 seconds. This work aims at providing a quantitative assessment of the potential of duplication and triplication to mitigate such a threat. Specifically, the main contributions of this work are:

- an analytical model to study the performance of all replication scenarios against silent errors, namely, duplication, triplication, or more for process and group replications;
- closed-form formulas that give the optimal checkpointing period and optimal process number as a function of the error rate, checkpoint cost, and platform size;
- a set of simulation results that corroborate the analytical model.

The rest of the paper is organized as follows. Section 2 surveys the related work. We introduce the performance model in Section 3, and derive the general expected execution time in Section 4. The analysis for process replication is presented in Section 5, followed by the analysis for group replication in Section 6. Section 7 is devoted to the simulation results. Finally, we provide concluding remarks and directions for future work in Section 8.

2 Related work

We survey related work in this section. We start with replication for HPC applications in Section 2.1 and cover application-specific detectors in Section 2.2.

2.1 Replication for fail-stop errors

Checkpointing policies have been widely studied. We refer to [30] for a survey of various protocols and the derivation of the Young's and Daly's formula [46, 18] for the optimal checkpointing periods. Recent advances include multi-level approaches, or the use of SSD or NVRAM as secondary storage [13]. Combining replication with checkpointing has been proposed in [41, 48, 24] for HPC platforms, and in [32, 45] for grid computing.

The use of redundant MPI processes is analyzed in [25, 26, 15]. In particular, the work by Ferreira et al. [26] has studied the use of process replication for MPI applications, using 2 replicas per MPI process. They provide a theoretical analysis of parallel efficiency, an MPI implementation that supports transparent process replication (including failure detection, consistent message ordering among replicas, etc.), and a set of experimental and simulation results. Partial redundancy is studied in [22, 44] (in combination with coordinated checkpointing) to decrease the overhead of full replication. Adaptive redundancy is introduced in [28], where a subset of processes is dynamically selected for replication.

Thread-level replication has been investigated in [47, 17, 38]. This paper targets process-level replication, in order to be able to detect (and correct) silent errors striking in all communication-related operations.

Finally, Ni et al [35] introduce process duplication to cope with both fail-stop and silent errors. Their pioneering paper contains many interesting results but differs from this work as follows: (i) they limit themselves to perfectly parallel applications while we investigate speedup profiles that obey Amdahl's law; (ii) they do not investigate triplication; and (iii) they compute an upper bound on the optimal period and do not determine optimal processor counts.

2.2 Silent error detection and correction

Application-specific information enables ad-hoc solutions, which dramatically decrease the cost of error detection. Algorithm-based fault tolerance (ABFT) [31, 9, 42] is a well-known technique, which uses checksums to detect up to a certain number of errors in linear algebra kernels. Unfortunately, ABFT can only protect datasets in linear algebra kernels, and it must be implemented for each different kernel, which incurs a large amount of work for large HPC applications. Other techniques have also been advocated. Benson, Schmit and Schreiber [7] compare the result of a higher-order scheme with that of a lower-order one to detect errors in the numerical analysis of ODEs and PDEs. Sao and Vuduc [39] investigate self-stabilizing corrections after error detection in the conjugate gradient method. Bridges et al. [29] propose linear solvers to tolerant soft faults using selective reliability. Elliot et al. [21] design a fault-tolerant GMRES capable of converging despite silent errors. Bronevetsky and de Supinski [10] provide a comparative study of detection costs for iterative methods.

Recently, several silent error detectors based on data analytics have been proposed, showing promising results. These detectors use several interpolation techniques such as time series prediction [8] and spatial multivariate interpolation [3, 4, 5]. Such techniques offer large detection coverage for a negligible overhead. However, these detectors do not guarantee full coverage; they can detect only a certain percentage of corruptions (i.e., partial verification with an imperfect recall). Nonetheless, the accuracy-to-cost ratios of these detectors are high, which makes them interesting alternatives at large scale. Similar detectors have also been designed to detect silent errors in the temperature data of the Orbital Thermal Imaging Spectrometer (OTIS) [16].

Again, all the papers quoted in this section provide application-specific detectors, while our approach is agnostic of the application characteristics. The only information is whether we can use either process replication. If not, we see the application as a black box and can use only group replication.

3 Model

This section presents the analytical model for evaluating the performance of different replication scenarios. The model is classical, similar to those of the literature for replication [26], only with a different objective (quantifying replication for silent errors). Table 1 summarizes the main notations.

Recall that μ_{ind} denotes the MTBE of an individual processor or process[3] of the system, and let $\lambda = \frac{1}{\mu_{\text{ind}}}$ denote the silent error rate of the processor. The error rate for a collection of P processors is then given by $\lambda_P = \frac{1}{\mu_P} = \frac{P}{\mu_{\text{ind}}} = \lambda P$ [30]. Assuming that the error arrivals follow *Exponential* distribution, the probability that a computation hit by a silent error during time T on P processes is given by $\mathbb{P}(T, P) = 1 - e^{-\lambda P T}$.

Consider long-latsting HPC applications that execute for hours or even days on a large-scale platform. Resilience is enforced by the combined use of replication and periodic checkpointing. Before each checkpoint, the results of different replicas are compared. Only when both

[3]We assume that each process is executed by a dedicated processor, hence use "processor" and "process" interchangeably. We also use MTBE instead of MTBF to emphasize that we deal with (silent) errors, not failures.

Table 1: List of Notations.

Parameters	
T	Length (or period) of a pattern
P	Number of processes allocated to an application
n	Number of (process or group) replicas
$S(P)$	Speedup function of an application
$H(P) = \frac{1}{S(P)}$	Error-free execution overhead
$\mathbb{E}_n(T,P)$	Expected execution time of a pattern
$\mathbb{H}_n(T,P)$	Expected execution overhead of a pattern
$\mathbb{S}_n(T,P)$	Expected speedup function of a pattern
$\lambda = \frac{1}{\mu_{\text{ind}}}$	Silent error rate of an individual process
$\mathbb{P}_n(T,P)$	Silent error probability of a pattern
C	Checkpointing cost
R	Recovery cost
V	Verification cost (comparison of replicas)

results (for duplication) or two out of three results (for triplication) coincide[4], in which case a *consensus* is said to be reached, the checkpoint is taken. Otherwise, silent errors are assumed to have been detected, and they cannot be corrected through consensus. The application then rolls back to the last checkpoint. There are two different types of replications:

- *Process replication*: Each process of the application is replicated, and the results of different processes are independently compared. A rollback is needed when at least one process has failed to reach a consensus;
- *Group replication*: The entire application (as a black box) is replicated, and the results of all replicas (as a whole) are compared. A rollback is needed when these group replicas fail to reach a consensus.

The computational chunk between two checkpoints is called a *periodic pattern*. For a replication scenario with n replicas, the objective is to minimize the expected total execution time (or makespan) of an application by finding the optimal pattern parameters:

- T: length (or period) of the pattern;
- P: number of processes allocated to the application.

Indeed, for long-lasting applications, it suffices to focus on just one pattern, since the pattern repeats itself over time. To see this, let W_{total} denote the total amount of work of the application and suppose the application has a speedup function $S(P)$ when executed on P processors. In this paper, we focus on a speedup function that obeys Amdahl's law[5]:

$$S(P) = \frac{1}{\alpha + \frac{1-\alpha}{P}} , \tag{1}$$

where $\alpha \in [0,1]$ denotes the sequential fraction of the application that cannot be parallelized. For convenience, we also define $H(P) = \frac{1}{S(P)}$ to be the execution overhead. For a pattern

[4]For $n > 3$ replicas, the results of k replicas should coincide, where $2 \leq k < n$ is a design parameter set by the system to control the level of reliability. $k = \lfloor \frac{n}{2} \rfloor + 1$ is a widely-used choice (majority voting).

[5]The model is generally applicable to other speedup functions as well.

of length T and run by P processes, the amount of work done in a pattern is therefore $W_{\text{pattern}} = T{\cdot}S(P)$, and the total number of patterns in the application can be approximated as $m = \frac{W_{\text{total}}}{W_{\text{pattern}}} = \frac{W_{\text{total}}}{T \cdot S(P)} = \frac{W_{\text{total}}}{T} H(P)$. Now, let $\mathbb{E}_n(T,P)$ denote the expected execution time of the pattern with n replicas in either replication scenario. Define $\mathbb{H}_n(T,P) = \frac{\mathbb{E}_n(T,P)}{T} H(P)$ to be the expected execution overhead of the pattern, and $\mathbb{S}_n(T,P) = \frac{1}{\mathbb{H}_n(T,P)}$ the expected speedup. The expected makespan of the application can then be written as $\mathbb{E}_{\text{total}} \approx \mathbb{E}_n(T,P)m = \mathbb{E}_n(T,P)\frac{W_{\text{total}}}{T} H(P) = \mathbb{H}_n(T,P) \cdot W_{\text{total}} = \frac{W_{\text{total}}}{\mathbb{S}_n(T,P)}$. This shows that the optimal expected makespan can be achieved by minimizing the expected execution overhead of a pattern, or equivalently, maximizing the expected speedup.

Now, we describe a model for the costs of checkpoint, recovery and consensus verification. First, the checkpoint cost clearly depends on the protocol and storage type. Note that only the result of one replica needs to be checkpointed, so the cost does not increase with the number of replicas. To save the application's memory footprint M to the storage system using P processes, we envision the following two scenarios:

- $C = \frac{M}{\tau_{io}}$: In this case, checkpoints are being written to the remote storage system, whose bandwidth is the I/O bottleneck. Here, τ_{io} is the remote I/O bandwidth.
- $C = \frac{M}{\tau_{net} P}$: This case corresponds to in-memory checkpoints, where each process stores $\frac{M}{P}$ data locally (e.g., on SSDs). Here, τ_{net} is the process network bandwidth.

The recovery cost is assumed to be the same as the checkpointing cost, i.e., $R = C$, as it involves the same I/O operations. This is a common assumption [34], although practical recovery cost can be somewhat smaller than the checkpoint cost [19]. Finally, verifying consensus is performed by communicating and comparing $\frac{M}{P}$ data stored on each process, which can be executed concurrently by all process pairs (or triplets). Hence, the verification cost satisfies $V = O(\frac{M}{P})$. Overall, we use the following general expression to account for the combined cost of verification and checkpoint/recovery:

$$V + C = c + \frac{d}{P}\,, \tag{2}$$

where c and d are constants that depend on the application memory footprint, checkpointing protocol, network or I/O bandwidth, etc. Equation (2) is convenient in terms of analysis as we will see in the subsequent sections. Here, $c = 0$ corresponds to the second checkpointing scenario discussed above.

4 Expected execution time

In this section, we compute the expected execution time of a periodic pattern, which will be used in the next two sections to derive the optimal pattern parameters.

Theorem 1. *The expected time to execute a periodic pattern of length T using P processes and n replicas can be expressed as*

$$\mathbb{E}_n(T,P) = T + V + C + \frac{\mathbb{P}_n(T,P)}{1 - \mathbb{P}_n(T,P)} (T + V + R)\,, \tag{3}$$

where $\mathbb{P}_n(T,P)$ denotes the probability that the execution fails due to silent errors striking during the pattern and we have to roll back to the last checkpoint.

Proof. Since replicas are synchronized, we can generally express the expected execution time as follows:

$$\mathbb{E}_n(T,P) = T + V + \mathbb{P}_n(T,P)\big(R + \mathbb{E}_n(T,P)\big) + \big(1 - \mathbb{P}_n(T,P)\big)C \ . \tag{4}$$

First, the pattern of length T is executed followed by the verification (through comparison and/or voting), which incurs cost V. With probability $\mathbb{P}_n(T,P)$, the pattern fails due to silent errors. In this case, we need to re-execute the pattern after performing a recovery from the last checkpoint with cost R. Otherwise, with probability $1-\mathbb{P}_n(T,P)$, the execution succeeds and the checkpoint with cost C is taken at the end of the pattern. Now, solving for $\mathbb{E}_n(T,P)$ from Equation (4), we can obtain the expected execution time of the pattern as shown in Equation (3). □

Remarks. Theorem 1 is applicable to both process replication and group replications. The only difference lies in the computation of failure probability $\mathbb{P}_n(T,P)$, which depends not only on the replication scenario but also on the number of replicas n.

5 Process replication

In this section, we consider process replication. We first derive the optimal computing patterns when each process of the application is duplicated (Section 5.1) and triplicated (Section 5.2), respectively. Finally, we generalize the results to an arbitrary but constant number of replications per process under a general process replication framework (Section 5.3).

5.1 Process duplication

We start with process duplication, that is, each process has two replicas. The following lemma shows the failure probability of a given computing pattern in this case.

Lemma 1. *Using process duplication, the failure probability of a computing pattern of length T and with P processes is given by*

$$\mathbb{P}_2^{\text{prc}}(T,P) = 1 - e^{-2\lambda TP} \ . \tag{5}$$

Proof. With duplication, errors cannot be corrected (no consensus), hence a process fails if either one of its replicas fails or both replicas fail. In other words, there is an error if the results of both replicas do not coincide (we neglect the quite unlikely scenario with one error in each replica leading to the same wrong result). Let $\mathbb{P}_1^{\text{prc}}(T,1) = 1 - e^{-\lambda T}$ denote the probability of a single process failure. Therefore, we can write the failure probability of any duplicated process as follows:

$$\begin{aligned}
\mathbb{P}_2^{\text{prc}}(T,1) &= \binom{2}{1}\left(1 - \mathbb{P}_1^{\text{prc}}(T,1)\right)\mathbb{P}_1^{\text{prc}}(T,1) + \mathbb{P}_1^{\text{prc}}(T,1)^2 \\
&= 2e^{-\lambda T}\left(1 - e^{-\lambda T}\right) + \left(1 - e^{-\lambda T}\right)^2 \\
&= 1 - e^{-2\lambda T} \ .
\end{aligned}$$

Now, because we have P independent processes, the probability that the application gets interrupted by silent errors is the probability that at least one process fails because of silent errors, which can be expressed as:

$$\begin{aligned}\mathbb{P}_2^{\text{prc}}(T,P) &= 1 - \mathbb{P}(\text{“No process fails”}) \\ &= 1 - \left(1 - \mathbb{P}_2^{\text{prc}}(T,1)\right)^P \\ &= 1 - e^{-2\lambda PT} \ .\end{aligned}$$

□

Using the failure probability in Lemma 1, we derive the optimal computing pattern for process duplication as shown in the following theorem. Recall that the application speedup follows Amdahl's law as shown in Equation (1) and the cost of verification and checkpoint is modeled by Equation (2).

Theorem 2. *A first-order approximation to the optimal number of processes for an application with 2 replicas per process is given by*

$$P_{\text{opt}} = \min\left\{\frac{Q}{2}, \left(\frac{1}{2}\left(\frac{1-\alpha}{\alpha}\right)^2 \frac{1}{c\lambda}\right)^{\frac{1}{3}}\right\} \ , \tag{6}$$

where Q denotes the total number of available processes in the system. The associated optimal checkpointing period and the expected speedup function of the application are

$$T_{\text{opt}}(P_{\text{opt}}) = \left(\frac{V+C}{2\lambda P_{\text{opt}}}\right)^{\frac{1}{2}} \ , \tag{7}$$

$$\mathbb{S}_2^{\text{prc}}(P_{\text{opt}}) = \frac{S(P_{\text{opt}})}{1 + 2\big(2\lambda(V+C)P_{\text{opt}}\big)^{\frac{1}{2}}} \ . \tag{8}$$

Proof. First, we can derive, from Theorem 1 and Lemma 1, the expected execution time of a pattern with length T and P duplicated processes as follows:

$$\begin{aligned}\mathbb{E}_2^{\text{prc}}(T,P) &= T + V + C + \left(e^{2\lambda PT} - 1\right)(T + V + R) \\ &= T + V + C + 2\lambda PT\,(T + V + R) + o(\lambda PT^2) \ .\end{aligned}$$

The second equation above is obtained by applying Taylor series to approximate $e^z = 1 + z + o(z)$ for $z < 1$, while assuming $\lambda PT = \Theta(\lambda^\epsilon)$, where $\epsilon > 0$.

Now, we have a closed-form expression for $\mathbb{E}_2^{\text{prc}}(T,P)$. Substituting it into $\mathbb{H}_2^{\text{prc}}(T,P) = H(P)\frac{\mathbb{E}_2^{\text{prc}}(T,P)}{T}$, we can get the expected execution overhead as:

$$\mathbb{H}_2^{\text{prc}}(T,P) = H(P)\left(1 + \frac{V+C}{T} + 2\lambda PT + o(\lambda PT)\right) \ . \tag{9}$$

The optimal overhead can then be achieved by balancing (or equating) the two terms $\frac{V+C}{T}$ and $2\lambda PT$ above, which gives the following optimal checkpointing period as a function of the process count:

$$T_{\text{opt}}(P) = \left(\frac{V+C}{2\lambda P}\right)^{\frac{1}{2}} \ . \tag{10}$$

Now, substituting $T_{\text{opt}}(P)$ back into Equation (9), we get the execution overhead as a function of the process count as follows (lower-order terms ignored):

$$\mathbb{H}_2^{\text{prc}}(P) = H(P)\left(1 + 2\big(2\lambda(V+C)P\big)^{\frac{1}{2}}\right) . \tag{11}$$

Note that Equations (10) and (11) hold true regardless of the form of the function $H(P)$ or the cost $V + C$. Recall that we consider Amhdal's law $H(P) = \alpha + \frac{1-\alpha}{P}$ and a cost model $V + C = c + \frac{d}{P}$. In order to derive the optimal process count, we consider two cases:

Case (1). $c > 0$ and $\alpha > 0$ are both constants: we can expand Equation (11) to be

$$\mathbb{H}_2^{\text{prc}}(P) = \alpha + 2\alpha\big(2\lambda cP\big)^{\frac{1}{2}} + \frac{1-\alpha}{P} + o(\lambda^{\frac{1}{2}}) . \tag{12}$$

The optimal overhead can then be achieved by setting

$$\frac{\partial \mathbb{H}_2^{\text{prc}}(P)}{\partial P} = \alpha\left(\frac{2\lambda c}{P}\right)^{\frac{1}{2}} - \frac{1-\alpha}{P^2} = 0 ,$$

which leads to $P^* = \left(\frac{1}{2}\left(\frac{1-\alpha}{\alpha}\right)^2 \frac{1}{c\lambda}\right)^{\frac{1}{3}}$. Since the total number of processes in the system is Q and each application process is duplicated, the optimal process count is upper-bounded by $\frac{Q}{2}$ if $P^* > \frac{Q}{2}$, due to the convexity of $\mathbb{H}_2^{\text{prc}}(P)$ as shown in Equation (11). Hence, the optimal process count P_{opt} is given by Equation (6).

Case (2). $c = 0$ or $\alpha = 0$: In either case, we can see that Equation (11) becomes a decreasing function of P. Therefore, the optimal strategy is to utilize all the available Q processes, i.e., $P_{\text{opt}} = \frac{Q}{2}$, which again satisfies Equation (6), since $\left(\frac{1}{2}\left(\frac{1-\alpha}{\alpha}\right)^2 \frac{1}{c\lambda}\right)^{\frac{1}{3}} = \infty$.

In either case, the expected application speedup is then given by the reciprocal of the overhead as shown in Equation (11) with the optimal process count P_{opt}. □

Remarks. For fully parallelizable applications, i.e., $\alpha = 0$, the optimal pattern on a Q-process platform is characterized by

$$P_{\text{opt}} = \frac{Q}{2} , \quad T_{\text{opt}} = \begin{cases} \sqrt{\frac{c}{\lambda Q}} & \text{for } V + C = c \\ \frac{1}{Q}\sqrt{\frac{2d}{\lambda}} & \text{for } V + C = \frac{d}{P} \end{cases} ,$$

$$\mathbb{S}_2^{\text{prc}}(P_{\text{opt}}) = \begin{cases} \frac{Q}{2(1+2\sqrt{\lambda c Q})} & \text{for } V + C = c \\ \frac{Q}{2(1+2\sqrt{2\lambda d})} & \text{for } V + C = \frac{d}{P} \end{cases} .$$

5.2 Process triplication

Now, we consider process duplication, that is, each process has three replicas. This is the smallest number of replicas that allows an application to recover from silent errors through majority voting instead of rolling back to the last checkpoint.

Lemma 2. *Using process triplication, the failure probability of a computing pattern of length T and with P processes is given by*

$$\mathbb{P}_3^{\text{prc}}(T, P) = 1 - \left(3e^{-2\lambda T} - 2e^{-3\lambda T}\right)^P . \tag{13}$$

Proof. Using triplication, if only one replica fails, the silent error can be masked by the two successful replicas. Hence, in this case, a process fails if at least two of its replicas are hit by silent errors. Let $\mathbb{P}_1^{\text{prc}}(T,1) = 1 - e^{-\lambda T}$ denote the probability of a single process failure. Therefore, we can write the failure probability of any triplicated process as follows:

$$\begin{aligned}\mathbb{P}_3^{\text{prc}}(T,1) &= \binom{3}{2}\left(1-\mathbb{P}_1^{\text{prc}}(T,1)\right)\mathbb{P}_1^{\text{prc}}(T,1)^2 + \mathbb{P}_1^{\text{prc}}(T,1)^3 \\ &= 3e^{-\lambda T}\left(1-e^{-\lambda T}\right)^2 + \left(1-e^{-\lambda T}\right)^3 \\ &= 1 - 3e^{-2\lambda T} + 2e^{-3\lambda T}\ .\end{aligned}$$

For P independent processes, the application fails when at least one of its processes fails. Hence, we have:

$$\begin{aligned}\mathbb{P}_3^{\text{prc}}(T,P) &= 1 - \mathbb{P}(\text{“No process fails”}) \\ &= 1 - \left(1-\mathbb{P}_3^{\text{prc}}(T,1)\right)^P \\ &= 1 - \left(3e^{-2\lambda T} - 2e^{-3\lambda T}\right)^P\ .\end{aligned}$$

□

The following theorem derives the optimal computing pattern for process triplication.

Theorem 3. *A first-order approximation to the optimal number of processes for an application with* 3 *replicas per process is given by*

$$P_{\text{opt}} = \min\left\{\frac{Q}{3}, \left(\frac{4}{3}\left(\frac{1-\alpha}{\alpha}\right)^3\left(\frac{1}{c\lambda}\right)^2\right)^{\frac{1}{4}}\right\}\ , \tag{14}$$

where Q denotes the total number of available processes in the system. The associated optimal checkpointing period and the expected speedup function of the application are

$$T_{\text{opt}}(P_{\text{opt}}) = \left(\frac{V+C}{6\lambda^2 P_{\text{opt}}}\right)^{\frac{1}{3}}\ , \tag{15}$$

$$\mathbb{S}_3^{\text{prc}}(P_{\text{opt}}) = \frac{S(P_{\text{opt}})}{1 + 3\left(\frac{3}{4}\left(\lambda(V+C)\right)^2 P_{\text{opt}}\right)^{\frac{1}{3}}}\ . \tag{16}$$

Proof. From Theorem 1 and Lemma 2, and applying Taylor series, we can derive the expected

execution time of a pattern as follows:

$$\begin{aligned}
\mathbb{E}_3^{\text{prc}}(T,P) &= T+V+C+\frac{1-\left(3e^{-2\lambda T}+2e^{-3\lambda T}\right)^P}{\left(3e^{-2\lambda T}-2e^{-3\lambda T}\right)^P}\left(T+V+R\right)\\
&= T+V+C+\left(\left(\frac{e^{3\lambda T}}{3e^{\lambda T}-2}\right)^P-1\right)\left(T+V+R\right)\\
&\approx T+V+C+\left(\left(\frac{1+3\lambda T+\frac{(3\lambda T)^2}{2}}{1+3\lambda T+\frac{3(\lambda T)^2}{2}}\right)^P-1\right)\left(T+V+R\right)\\
&\approx T+V+C+\left(\left(1+3(\lambda T)^2\right)^P-1\right)\left(T+V+R\right)\\
&= T+V+C+\left(\sum_{j=0}^{P}\binom{P}{j}\left(3(\lambda T)^2\right)^j-1\right)\left(T+V+R\right)\\
&= T+V+C+3P(\lambda T)^2(T+V+R)+o(\lambda^2PT^3)\ .
\end{aligned}$$

The execution overhead can then be expressed as:

$$\mathbb{H}_3^{\text{prc}}(T,P)=H(P)\left(1+\frac{V+C}{T}+3P(\lambda T)^2+o(\lambda^2PT^2)\right)\ . \tag{17}$$

The optimal checkpointing period is then obtained by setting

$$\frac{\partial\mathbb{H}_3^{\text{prc}}(T,P)}{\partial T}=-\frac{V+C}{T^2}+6\lambda^2PT=0\ ,$$

which gives

$$T_{\text{opt}}(P)=\left(\frac{V+C}{6\lambda^2P}\right)^{\frac{1}{3}}\ .$$

Substituting $T_{\text{opt}}(P)$ back into Equation (17), we get the following execution overhead (with lower-order terms ignored):

$$\mathbb{H}_3^{\text{prc}}(P)=H(P)\left(1+3\left(\frac{3}{4}\left(\lambda(V+C)\right)^2P\right)^{\frac{1}{3}}\right)\ . \tag{18}$$

To derive the optimal process count, consider $V+C=c$ and $H(P)=\alpha+\frac{1-\alpha}{P}$ for $\alpha>0$. Then, Equation (11) can be expanded as

$$\mathbb{H}_3^{\text{prc}}(P)=\alpha+3\alpha\left(\frac{3}{4}\left(\lambda c\right)^2P\right)^{\frac{1}{3}}+\frac{1-\alpha}{P}+o(\lambda^{\frac{2}{3}})\ . \tag{19}$$

The optimal overhead is achieved by setting

$$\frac{\partial\mathbb{H}_3^{\text{prc}}(P)}{\partial P}=\alpha\left(\frac{3}{4}\left(\lambda c\right)^2\frac{1}{P^2}\right)^{\frac{1}{3}}-\frac{1-\alpha}{P^2}=0\ ,$$

which gives rise to $P^*=\left(\frac{4}{3}\left(\frac{1-\alpha}{\alpha}\right)^3\left(\frac{1}{c\lambda}\right)^2\right)^{\frac{1}{4}}$. Now, the optimal process count is upper-bounded by $\frac{Q}{3}$. Thus, P_{opt} is given by Equation (14), which again holds true when $c=0$ or $\alpha=0$, and the optimal expected speedup satisfies $\mathbb{S}_3^{\text{prc}}(P_{\text{opt}})=\frac{1}{\mathbb{H}_3^{\text{prc}}(P_{\text{opt}})}$, as shown in Equation (16). □

Remarks. For fully parallelizable applications, i.e., $\alpha = 0$, the optimal pattern on a Q-process platform is characterized by

$$P_{\text{opt}} = \frac{Q}{3}, \quad T_{\text{opt}} = \begin{cases} \sqrt[3]{\frac{c}{2\lambda^2 Q}} & \text{for } V + C = c \\ \sqrt[3]{\frac{3d}{2\lambda^2 Q^2}} & \text{for } V + C = \frac{d}{P} \end{cases},$$

$$\mathbb{S}_2^{\text{prc}}(P_{\text{opt}}) = \begin{cases} \frac{Q}{3\left(1+3\sqrt[3]{\left(\frac{\lambda c}{2}\right)^2 Q}\right)} & \text{for } V + C = c \\ \frac{Q}{3\left(1+3\sqrt[3]{\left(\frac{3\lambda c}{2}\right)^2 \frac{1}{Q}}\right)} & \text{for } V + C = \frac{d}{P} \end{cases}.$$

Compared with duplication, the ability to correct errors in triplication allows checkpoints to be taken less frequently (i.e., larger checkpointing period). In terms of the expected speedup, triplication suffers from a smaller error-free speedup ($\frac{Q}{3}$ vs $\frac{Q}{2}$) due to the use of fewer concurrent processes to perform useful work, but also has a smaller error-induced denominator, especially on platforms with a large number of processes Q. In Section 7, we will conduct simulations to evaluate this trade-off and compare the performance of duplication and triplication.

5.3 General process replication

In this section, we consider a general resilience framework and derive the optimal pattern using n replicas per process, where n is an arbitrary constant. Moreover, let k denote the number of "good" replicas (not hit by silent errors) that is required to reach a consensus through voting. Optimistically, assuming any two replicas that are hit by silent errors will produce different results, we can set $k = 2$, i.e., at least two replicas should agree on the result to avoid a rollback. Under a more pessimistic assumption, we will need a majority of the n replicas to agree on the result, so in this case we need $k = \lfloor \frac{n}{2} \rfloor + 1$. Our results are independent of the choice of k.

As for duplication and triplication, for a given (n, k) pair, we can compute the failure probability of a pattern with length T and P processes as follows:

$$\begin{aligned} \mathbb{P}_{n,k}^{\text{prc}}(T, P) &= 1 - \mathbb{P}(\text{"No process fails"}) \\ &= 1 - \left(1 - \mathbb{P}_{n,k}^{\text{prc}}(T, 1)\right)^P, \end{aligned} \tag{20}$$

where

$$\begin{aligned} \mathbb{P}_{n,k}^{\text{prc}}(T, 1) &= \sum_{j=0}^{k-1} \binom{n}{j} \left(1 - \mathbb{P}_1^{\text{prc}}(T, 1)\right)^j \mathbb{P}_1^{\text{prc}}(T, 1)^{n-j} \\ &= \sum_{j=0}^{k-1} \binom{n}{j} e^{-\lambda j T} \left(1 - e^{-\lambda T}\right)^{n-j} \end{aligned} \tag{21}$$

denotes the failure probability of a single process with n replicas due to less than k of them surviving silent errors.

The following theorem shows the general result for (n, k)-process replication.

Theorem 4. *On a system with a total number of Q available processors, a first-order approximation to the optimal number of processes for an application with n replicas per process (k of which must concur to avoid a rollback) is given by*

$$P_{\text{opt}} = \min\left\{\frac{Q}{n}, \left(\gamma_{n,k}\left(\frac{1-\alpha}{\alpha}\right)^{n-k+2}\left(\frac{1}{c\lambda}\right)^{n-k+1}\right)^{\frac{1}{n-k+3}}\right\} . \tag{22}$$

The associated optimal checkpointing period and the expected speedup function of the application are

$$T_{\text{opt}}(P_{\text{opt}}) = \left(\frac{V+C}{\beta_{n,k}\lambda^{n-k+1}P_{\text{opt}}}\right)^{\frac{1}{n-k+2}} , \tag{23}$$

$$\mathbb{S}^{\text{prc}}_{n,k}(P_{\text{opt}}) = \frac{S(P_{\text{opt}})}{1+(n-k+2)\left(\frac{((V+C)\lambda)^{n-k+1}P_{\text{opt}}}{\gamma_{n,k}}\right)^{\frac{1}{n-k+2}}} . \tag{24}$$

Here, $\beta_{n,k} = \binom{n}{k-1}(n-k+1)$ and $\gamma_{n,k} = \frac{(n-k+1)^{n-k+1}}{\binom{n}{k-1}}$.

Proof. As in the preceding two cases, we start by approximating the error probability. First, we can approximate the probability of single process failure as

$$\begin{aligned}\mathbb{P}^{\text{prc}}_{n,k}(T,1) &= \sum_{j=0}^{k-1}\binom{n}{j}(1-\lambda T)^j(\lambda T)^{n-j} \\ &\approx \binom{n}{k-1}(\lambda T)^{n-k+1} + o((\lambda T)^{n-k+1}) .\end{aligned}$$

We can now approximate

$$\begin{aligned}\frac{\mathbb{P}^{\text{prc}}_{n,k}(T,P)}{1-\mathbb{P}^{\text{prc}}_{n,k}(T,P)} &\approx \left(\frac{1}{1-\mathbb{P}^{\text{prc}}_{n,k}(T,1)}\right)^P - 1 \\ &\approx \left(1+\binom{n}{k-1}(\lambda T)^{n-k+1}\right)^P - 1 \\ &= \sum_{j=0}^{P}\binom{P}{j}\left(\binom{n}{k-1}(\lambda T)^{n-k+1}\right)^j - 1 \\ &= \binom{n}{k-1}P(\lambda T)^{n-k+1} + o(P(\lambda T)^{n-k+1}) .\end{aligned}$$

Thus, the expected execution time of a pattern can be expressed as

$$\begin{aligned}\mathbb{E}^{\text{grp}}_n(T,P)k = T+V+C+\binom{n}{k-1}P(\lambda T)^{n-k+1}(T+V+R) \\ + o(\lambda^{n-k+1}PT^{n-k+2}) .\end{aligned}$$

The derivation of the optimal pattern then follows exactly the same steps as in the proofs of Theorems 2 and 3, and the detailed derivation steps are omitted here. □

Remarks. Theorem 4 encompasses Theorem 2 and Theorem 3 as special cases. We point out that it even holds for the case without replication, i.e., when $n = k = 1$. In this case, Theorem 4 evaluates to

$$T_{\text{opt}}(P) = \sqrt{\frac{V+C}{\lambda P}} \ ,$$
$$\mathbb{S}_1^{\text{prc}}(P) = \frac{S(P)}{1 + 2\sqrt{(V+C)\lambda P}} \ ,$$

which is consistent with the results obtained in [6, 2], provided that a reliable silent error detector is available. However, as mentioned previously, such a detector is only known in some application-specific domains. For general-purpose computations, replication appears to be the only viable approach to detect/correct silent errors so far.

6 Group replication

In this section, we consider group replication. Recall that, unlike process replication where the results of each process from different replicas are independently compared, group replication compares the outputs of the different groups viewed as independent black-box applications. First, we make the following technical observation, which establishes the relationship between the two replication mechanisms from the resilience point of view.

Observation 1. *Running an application using group replication with n replicas, where each replica has P processes and each process has error rate λ, has the same failure probability as running it using process replication with one process, which has error rate λP and is replicated n times.*

The above observation allows us to compute the failure probability for group replication by deriving from the corresponding formulas under process replication while setting $P = 1$ and $\lambda = \lambda P$. The rest of this section shows the results for duplication, triplication, and a general group replication framework. Proofs are similar to those in process replication, and are hence omitted.

6.1 Group duplication

By applying Observation 1 on Lemma 1, we can get the failure probability for a given pattern under group duplication as follows.

Lemma 3. *Using group duplication, the failure probability of a computing pattern of length T and with P processes is given by*

$$\mathbb{P}_2^{\text{grp}}(T, P) = 1 - e^{-2\lambda TP} \ . \tag{25}$$

This leads us to the following theorem on the optimal pattern:

Theorem 5. *A first-order approximation to the optimal number of processes for an application with 2 replica groups is given by*

$$P_{\text{opt}} = \min\left\{\frac{Q}{2}, \left(\frac{1}{2}\left(\frac{1-\alpha}{\alpha}\right)^2 \frac{1}{c\lambda}\right)^{\frac{1}{3}}\right\} \ , \tag{26}$$

where Q denotes the total number of available processes in the system. The associated optimal checkpointing period and the expected speedup function of the application are

$$T_{\text{opt}}(P_{\text{opt}}) = \left(\frac{V+C}{2\lambda P_{\text{opt}}}\right)^{\frac{1}{2}} , \tag{27}$$

$$\mathbb{S}_2^{\text{grp}}(P_{\text{opt}}) = \frac{S(P_{\text{opt}})}{1+2\big(2\lambda(V+C)P_{\text{opt}}\big)^{\frac{1}{2}}} . \tag{28}$$

Remarks. The result is identical to that of process duplication. Indeed, in both cases, a single silent error that strikes any of the running processes will cause the whole application to fail.

6.2 Group triplication

Again, applying Observation 1 on Lemma 2, we can get the failure probability for a given pattern under group triplication.

Lemma 4. *Using group triplication, the failure probability of a computing pattern of length T and with P processes is given by*

$$\mathbb{P}_3^{\text{grp}}(T,P) = 1 - \left(3e^{-2\lambda TP} - 2e^{-3\lambda TP}\right) . \tag{29}$$

The following theorem shows the optimal pattern.

Theorem 6. *A first-order approximation to the optimal number of processes for an application with 3 replica groups is given by*

$$P_{\text{opt}} = \min\left\{\frac{Q}{3}, \left(\frac{1}{6}\left(\frac{1-\alpha}{\alpha}\right)^3\left(\frac{1}{c\lambda}\right)^2\right)^{\frac{1}{5}}\right\} , \tag{30}$$

where Q denotes the total number of available processes in the system. The associated optimal checkpointing period and the expected execution overhead are

$$T_{\text{opt}}(P_{\text{opt}}) = \left(\frac{V+C}{6(\lambda P_{\text{opt}})^2}\right)^{\frac{1}{3}} , \tag{31}$$

$$\mathbb{S}_3^{\text{grp}}(P_{\text{opt}}) = \frac{S(P_{\text{opt}})}{1+3\left(\frac{3}{4}\big(\lambda(V+C)P_{\text{opt}}\big)^2\right)^{\frac{1}{3}}} . \tag{32}$$

Remarks. Compared to the result of process triplication (Theorem 3) and under the same condition (e.g., $\alpha = 0$ so both scenarios allocate the same number of $P_{\text{opt}} = \frac{Q}{3}$ processes to each replica), group triplication needs to place checkpoints more frequently yet enjoys a smaller execution speedup. This provides a theoretical explanation to the common understanding that group replication in general cannot recover from some error combinations that its process counterpart is capable of, making the latter a superior replication mechanism provided that it can be feasibly implemented.

6.3 General group replication

Finally, we consider a general group replication framework and derive the optimal pattern using a constant number of n replica groups, out of which k of them must agree to avoid a rollback. Again, the results work for any choice of k.

Now, applying Observation 1 on Equations (20) and (21), we can compute the failure probability of a pattern with length T and P processes under a (n,k) group replication model:

$$\mathbb{P}^{\text{grp}}_{n,k}(T,P) = \sum_{j=0}^{k-1} \binom{n}{j} \left(e^{-\lambda PT}\right)^j \left(1-e^{-\lambda PT}\right)^{n-j} . \tag{33}$$

The following theorem shows the general result for this case.

Theorem 7. *On a system with a total number of Q available processors, a first-order approximation to the optimal number of processes for an application with n replica groups (k of which must concur to avoid a rollback) is given by*

$$P_{\text{opt}} = \min\left\{\frac{Q}{n}, \left(\frac{1}{\beta_{n,k}}\left(\frac{1-\alpha}{\alpha}\right)^{n-k+2}\left(\frac{1}{c\lambda}\right)^{n-k+1}\right)^{\frac{1}{2n-2k+3}}\right\} . \tag{34}$$

The associated optimal checkpointing period and the expected speedup function of the application are

$$T_{\text{opt}}(P_{\text{opt}}) = \left(\frac{C+V}{\beta_{n,k}(\lambda P_{\text{opt}})^{n-k+1}}\right)^{\frac{1}{n-k+2}} , \tag{35}$$

$$\mathbb{S}^{\text{grp}}_{n,k}(P_{\text{opt}}) = \frac{S(P_{\text{opt}})}{1+(n-k+2)\left(\frac{1}{\gamma_{n,k}}\left((V+C)\lambda P_{\text{opt}}\right)^{n-k+1}\right)^{\frac{1}{n-k+2}}} . \tag{36}$$

Here, $\beta_{n,k} = \binom{n}{k-1}(n-k+1)$ and $\gamma_{n,k} = \frac{(n-k+1)^{n-k+1}}{\binom{n}{k-1}}$.

7 Simulations

We conduct a set of simulations whose goal is twofold: (i) validate the accuracy of the theoretical study; and (ii) evaluate the efficiency of both process and group replication under different scenarios at extreme scale. The simulator is publicly available at `http://perso.ens-lyon.fr/aurelien.cavelan/replication.zip` so that interested readers can instantiate their preferred scenarios and repeat the same simulations for reproducibility purpose.

7.1 Simulation setup

The simulator has been designed to simulate each process individually, and each process has its own error trace. A simulation works as follows: we feed the simulator with the model parameters μ_{ind}, Q, C, V, R, and α, and we compute the associated optimal number of processes P_{opt} and the optimal checkpointing period $T_{\text{opt}}(P_{\text{opt}})$ using the corresponding model equations. For each run, the simulator outputs the efficiency, defined as $\frac{\mathbb{S}(P_{\text{opt}})}{Q}$, as well as the average number of errors and the average number of recoveries per million CPU hours of work.

Then, for each of the following scenarios, we compare the simulated efficiency to the theoretical value, obtained using the model equations for $\mathbb{S}(P_{\text{opt}})$. As suggested by Observation 1, process and group replications with $n = 2$ lead to identical results, so we have merged them together.

In the following, we set the cost of recovery to be the same as the checkpoint cost (as discussed in Section 3), and we set the cost $V + C$ according the values of c and d as in Equation (2). We consider different Mean Time Between Errors (MTBE), ranging from 10^6 seconds ($\approx$ 11 days) down to 10^2 seconds ($<$ 2 minutes) for $Q = 10^6$ processes, matching the numbers in [43].

7.2 Impacts of $MTBE$ and checkpoint cost

Figure 1 presents the impact of the $MTBE$ on the efficiency of both duplication and triplication for three different checkpoint costs, but using the same value $\alpha = 10^{-6}$ for the sequential fraction of the application (see next section for the impact of varying α). The first row of plots is obtained with a cost of 30 minutes (i.e. $c = 1,800, d = 0$), the second row with a cost of 60 seconds (i.e. $c = 60, d = 0$), and the last row with $c = 0, d = 10^7$, which correspond to a checkpoint cost of 20 seconds for duplication with $\frac{Q}{2}$ processes and 30 seconds for triplication with $\frac{Q}{3}$ processes. In addition to the efficiency, we provide the average number of errors and recoveries per million hours of work, the optimal checkpointing period $T_{\text{opt}}(P_{\text{opt}})$ and the optimal number of processes P_{opt}.

Efficiency. First, we observe in the first column that the difference between the theoretical efficiency and the simulated efficiency remains small ($< 5\%$ absolute difference), which shows the accuracy of the first-order approximation. Then, with very few errors ($MTBE = 10^6$), we observe that duplication is always better than triplication. This is as expected, since the maximum efficiency for duplication is 0.5 (assuming $\alpha = 0$ and no error), while the maximum efficiency for triplication is 0.33. However, as the $MTBE$ decreases, triplication becomes more attractive and eventually outperforms duplication. With a checkpoint cost of 30 minutes (first row), the $MTBE$ required is around 28 hours for process triplication to win and 20 hours for group triplication to win. With smaller checkpoint costs, such as 60 seconds (second row) and 30 seconds (third row), checkpoints can be more frequent and the $MTBE$ required for triplication to win is pushed down to a couple of hours and a couple of minutes, respectively.

Number of errors and recoveries. The second column presents the number of errors and the corresponding number of recoveries per million hours of work. The number of errors is always higher than the number of recoveries, because multiple errors can occur during a period (before the checkpoint, which is the point of detection), causing a single recovery. At $MTBE = 10^2$, almost half of the errors that occurred with duplication were actually hidden behind another error. Even more errors were hidden with group triplication, since one more error (in a different replica) is required to cause a recovery. Finally, (almost) all errors were hidden with process replication, which is able to handle many errors, as long as they strike in different processes.

Optimal checkpointing period. The third column shows the optimal length of the pattern. In order to cope with the increasing number of errors and recoveries, the length of the optimal period becomes smaller. Note that the length of the period for group triplication is comparable to that for duplication, around one day when $MTBE = 10^6$ down to a couple of minutes when $MTBE = 10^2$. However, the length of the pattern for process triplication is always higher by

several orders of magnitude, from more than 10 days when $MTBE = 10^6$ down to a couple of hours when $MTBE = 10^2$.

Optimal number of processes. With $\alpha = 10^{-6}$, the application has ample parallelism, so the optimal number of processes to use is always $\frac{Q}{2} = 5 \cdot 10^5$ for duplication and $\frac{Q}{3} \approx 3.3 \cdot 10^5$ for triplication, except when $MTBE = 10^2$ and $c = 1,800$, where the optimal number of processes for duplication is $\approx 3 \cdot 10^5$ and the optimal number of processes for group triplication is $\approx 2 \cdot 10^5$.

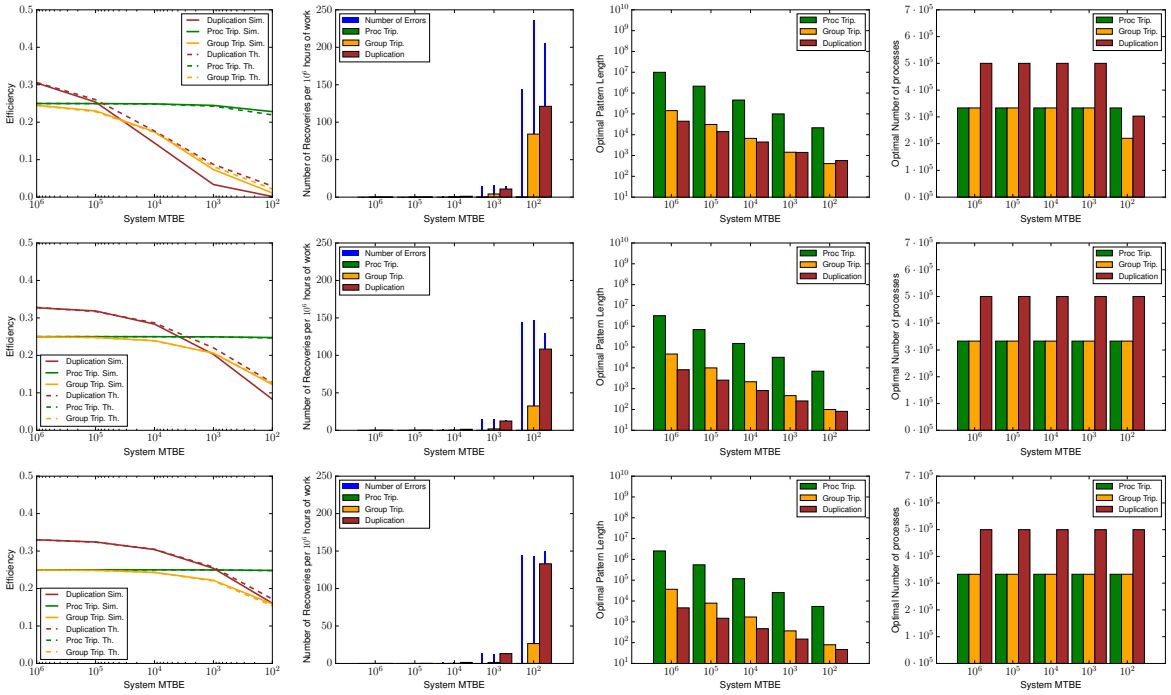

Figure 1: Impact of System MTBE on the efficiency with $c = 1,800, d = 0$ (top), $c = 60, d = 0$ (middle), $c = 0, d = 10^7$ (bottom) and $\alpha = 10^{-6}$.

7.3 Impact of sequential fraction (Amdahl)

Figure 2 presents two additional simulation results for $\alpha = 10^{-7}$ and $\alpha = 10^{-5}$. With a small fraction of sequential work (left plots), the efficiency is improved ($\approx 85\%$ of the maximum efficiency for duplication and $\approx 95\%$ for triplication at $MTBE = 10^6$), and both duplication and triplication use all processes available. On the contrary, with a higher sequential fraction of work (right plots), the efficiency drops ($< 20\%$ of the maximum efficiency for duplication and $< 30\%$ for triplication at $MTBE = 10^6$), and using more processes does not improve the efficiency and only contributes to increasing the number of errors. Therefore, these results suggest that even when using replication or triplication, there comes a point where it is no longer beneficial to use all available processes. In this example, when $MTBE = 10^2$, duplication and group triplication would use fewer than $2 \cdot 10^5$ processes (one fifth of the available resources). Process triplication, on the other hand, still utilizes all the resources and outperforms the other two schemes in terms of the efficiency across the whole range of system MTBE.

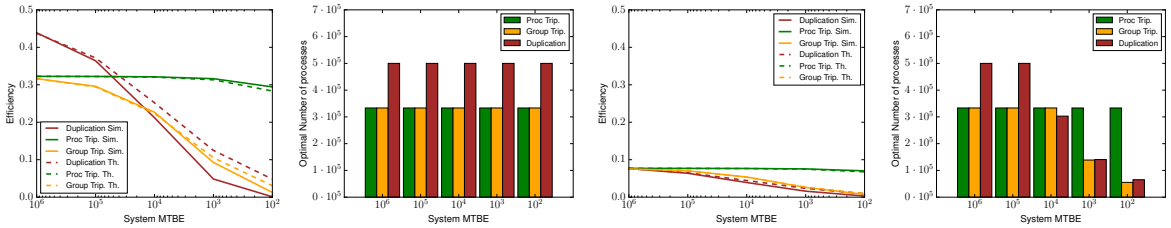

Figure 2: Impact of sequential fraction (in Amdahl's Law) on efficiency and optimal number of processes with $\alpha = 10^{-7}$ (left) and $\alpha = 10^{-5}$ (right).

7.4 Impact of number of processes

Figure 3 shows the impact of the number of processes on the simulated efficiency of different replication scenarios. In addition, we also show (as big dots) the theoretical efficiency obtained with the optimal number of processes from Theorems 2, 3 and 6. By varying the number of processes, we find that the simulated optimum (that yields the best efficiency) matches our theoretical optimal number of processes closely. We can also see that process triplication scales very well with increasing number of processes. As opposed to group triplication, which has to recover from a checkpoint if just two errors strike in two different replicas, process triplication benefits from the additional process: from a resilience point of view, each process acts as a buffer to handle one more error; in other words, the probability that two errors strike the two replicas of the same process decreases, thereby improving the efficiency.

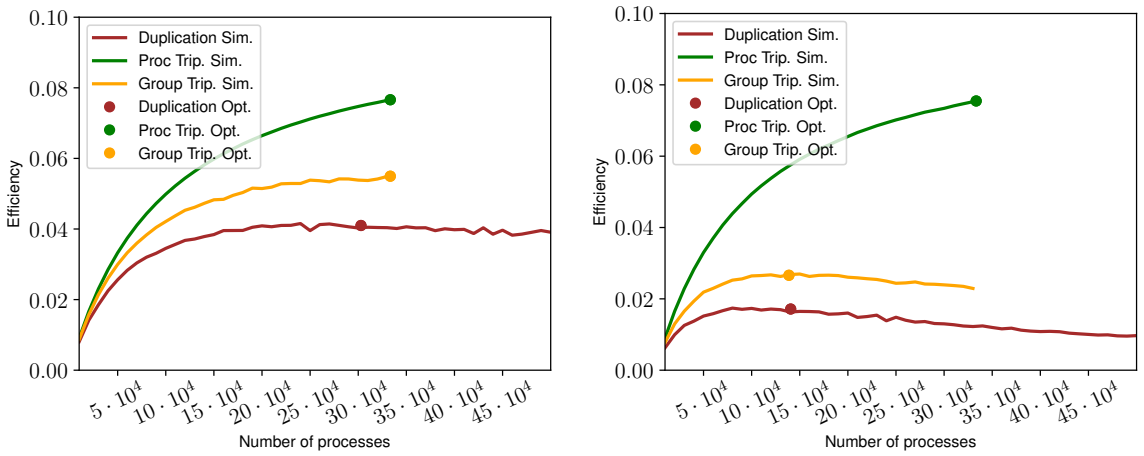

Figure 3: Impact of the number of processes on the efficiency with $MTBE = 10^4$ (left), $MTBE = 10^3$ (right), $Q = 10^6$, $c = 1n800, d = 0$, and $\alpha = 10^{-5}$.

7.5 Summary

Results suggest that duplication is more efficient than triplication for high $MTBE$ (e.g. 10^5 seconds for $C = 30$ minutes). If process triplication is available, then it is always more efficient for smaller $MTBE$: its efficiency remains stable despite the increasing number of failures. If

process triplication is not available, we show that group triplication is slightly more efficient than duplication for small $MTBE$, but the gain is small. Furthermore, the impact of the sequential fraction of work α (in Amdahl's Law) is twofold: it limits the efficiency (e.g. 15% of the maximum with $\alpha = 10^{-5}$ for for both duplication and triplication), and it is a major factor in limiting the optimal number of processes (e.g. one tenth of the platform with $\alpha = 10^{-5}$ and $Q = 10^6$ at $MTBE = 10^2$).

8 Conclusion

Silent-errors represent a major threat to the HPC community. In the absence of application-specific detectors, replication is the only solution. Unfortunately, it comes with high cost: by definition, the efficiency is upper-bounded by 0.5 for duplication, and by 0.333 for triplication. Are these upper bounds likely to be achieved? If yes, it means that duplication should always be preferred to triplication. If not, it means that in some scenarios, the striking of errors is so frequent that duplication, and in particular group duplication, is not the right choice.

The major contribution of this paper is to provide an in-depth analysis of process and group duplication, and of process and group triplication. Given a level n of replication, and a set of application/platform parameters (speedup profile, total number or processes, process MTBE, checkpoint time, etc), we derive closed-form formulas for the optimal period size and optimal resource usage, and for the overall efficiency of the approach. This allows to choose the best value of n. A set of simulations demonstrate the accuracy of the model and analysis. Our computer-algebra sheets and simulator code are made publicly available, so that one can instantiate their preferred scenario. Altogether, this paper has laid the foundations for a better understanding of the impact of silent errors on HPC computing at scale.

Future work will be devoted to combining replication and checkpointing to mitigate both fail-stop failures and silent errors. Partial replication is another topic to explore, if the application comes as a workflow whose tasks are atomic components: one could assign different replication levels (duplication, triplication or more) to the different tasks, depending upon their criticality in terms of longest paths, number of successors, etc.

References

[1] A. Avizienis, J. Laprie, B. Randell, and C. E. Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Trans. Dependable Sec. Comput.*, 1(1):11–33, 2004.

[2] L. Bautista-Gomez, A. Benoit, A. Cavelan, S. K. Raina, Y. Robert, and H. Sun. Which verification for soft error detection? In *HiPC*. IEEE, 2015.

[3] L. Bautista Gomez and F. Cappello. Detecting silent data corruption through data dynamic monitoring for scientific applications. In *PPoPP*. ACM, 2014.

[4] L. Bautista Gomez and F. Cappello. Detecting and correcting data corruption in stencil applications through multivariate interpolation. In *FTS*. IEEE, 2015.

[5] L. Bautista Gomez and F. Cappello. Exploiting Spatial Smoothness in HPC Applications to Detect Silent Data Corruption. In *HPCC*. IEEE, 2015.

[6] A. Benoit, A. Cavelan, Y. Robert, and H. Sun. Assessing general-purpose algorithms to cope with fail-stop and silent errors. In *PMBS*. ACM, 2014.

[7] A. R. Benson, S. Schmit, and R. Schreiber. Silent error detection in numerical time-stepping schemes. *Int. J. High Performance Computing Applications*, 2014.

[8] E. Berrocal, L. Bautista-Gomez, S. Di, Z. Lan, and F. Cappello. Lightweight silent data corruption detection based on runtime data analysis for HPC applications. In *HPDC*. ACM, 2015.

[9] G. Bosilca, R. Delmas, J. Dongarra, and J. Langou. Algorithm-based fault tolerance applied to high performance computing. *J. Parallel Distrib. Comput.*, 69(4):410–416, 2009.

[10] G. Bronevetsky and B. de Supinski. Soft error vulnerability of iterative linear algebra methods. In *ICS*. ACM, 2008.

[11] F. Cappello, E. M. Constantinescu, P. D. Hovland, T. Peterka, C. Phillips, M. Snir, and S. M. Wil. Improving the trust in results of numerical simulations and scientific data analytics. White paper MCS-TM-352, ANL, 2015.

[12] F. Cappello, A. Geist, B. Gropp, L. Kale, B. Kramer, and M. Snir. Toward Exascale Resilience. *Int. J. High Performance Computing Applications*, 23(4):374–388, 2009.

[13] F. Cappello, A. Geist, W. Gropp, S. Kale, B. Kramer, and M. Snir. Toward exascale resilience: 2014 update. *Supercomputing frontiers and innovations*, 1(1), 2014.

[14] H. Casanova, M. Bougeret, Y. Robert, F. Vivien, and D. Zaidouni. Using group replication for resilience on exascale systems. *Int. Journal of High Performance Computing Applications*, 28(2):210–224, 2014.

[15] H. Casanova, Y. Robert, F. Vivien, and D. Zaidouni. On the impact of process replication on executions of large-scale parallel applications with coordinated checkpointing. *Future Generation Comp. Syst.*, 51:7–19, 2015.

[16] E. Ciocca, I. Koren, Z. Koren, C. M. Krishna, and D. S. Katz. Application-level fault tolerance in the orbital thermal imaging spectrometer. In *PRDC*. IEEE, 2004.

[17] S. P. Crago, D. I. Kang, M. Kang, R. Kost, K. Singh, J. Suh, and J. P. Walters. Programming models and development software for a space-based many-core processor. In *4th Int. Conf. onon Space Mission Challenges for Information Technology*, pages 95–102. IEEE, 2011.

[18] J. T. Daly. A higher order estimate of the optimum checkpoint interval for restart dumps. *Future Generation Comp. Syst.*, 22(3):303–312, 2006.

[19] S. Di, M. S. Bouguerra, L. Bautista-Gomez, and F. Cappello. Optimization of multi-level checkpoint model for large scale HPC applications. In *IPDPS*. IEEE, 2014.

[20] J. Dongarra and et al. The international exascale software project roadmap. *Int. J. High Perform. Comput. Appl.*, 25(1):3–60, 2011.

[21] J. Elliott, M. Hoemmen, and F. Mueller. Evaluating the impact of SDC on the GMRES iterative solver. In *IPDPS*. IEEE, 2014.

[22] J. Elliott, K. Kharbas, D. Fiala, F. Mueller, K. Ferreira, and C. Engelmann. Combining partial redundancy and checkpointing for HPC. In *ICDCS*. IEEE, 2012.

[23] E. Elnozahy and J. Plank. Checkpointing for Peta-Scale Systems: A Look into the Future of Practical Rollback-Recovery. *IEEE Transactions on Dependable and Secure Computing*, 1(2):97—108, 2004.

[24] C. Engelmann, H. H. Ong, and S. L. Scorr. The case for modular redundancy in large-scale highh performance computing systems. In *PDCN*. IASTED, 2009.

[25] C. Engelmann and B. Swen. Redundant execution of HPC applications with MR-MPI. In *PDCN*. IASTED, 2011.

[26] K. Ferreira, J. Stearley, J. H. I. 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 *SC'11*. ACM, 2011.

[27] D. Fiala, F. Mueller, C. Engelmann, R. Riesen, K. Ferreira, and R. Brightwell. Detection and correction of silent data corruption for large-scale high-performance computing. In *SC*. ACM, 2012.

[28] C. George and S. S. Vadhiyar. Adft: An adaptive framework for fault tolerance on large scale systems using application malleability. *Procedia Computer Science*, 9:166 – 175, 2012.

[29] M. Heroux and M. Hoemmen. Fault-tolerant iterative methods via selective reliability. Research report SAND2011-3915 C, Sandia Nat. Lab., 2011.

[30] T. Hérault and Y. Robert, editors. *Fault-Tolerance Techniques for High-Performance Computing*, Computer Communications and Networks. Springer Verlag, 2015.

[31] K.-H. Huang and J. A. Abraham. Algorithm-based fault tolerance for matrix operations. *IEEE Trans. Comput.*, 33(6):518–528, 1984.

[32] T. Leblanc, R. Anand, E. Gabriel, and J. Subhlok. Volpexmpi: An mpi library for execution of parallel applications on volatile nodes. In *16th European PVM/MPI Users' Group Meeting*, pages 124–133. Springer-Verlag, 2009.

[33] R. E. Lyons and W. Vanderkulk. The use of triple-modular redundancy to improve computer reliability. *IBM J. Res. Dev.*, 6(2):200–209, 1962.

[34] A. Moody, G. Bronevetsky, K. Mohror, and B. R. d. Supinski. Design, modeling, and evaluation of a scalable multi-level checkpointing system. In *SC*. ACM, 2010.

[35] X. Ni, E. Meneses, N. Jain, and L. V. Kalé. ACR: Automatic Checkpoint/Restart for Soft and Hard Error Protection. In *SC*. ACM, 2013.

[36] T. O'Gorman. The effect of cosmic rays on the soft error rate of a DRAM at ground level. *IEEE Trans. Electron Devices*, 41(4):553–557, 1994.

[37] R. A. Oldfield, S. Arunagiri, P. J. Teller, S. Seelam, M. R. Varela, R. Riesen, and P. C. Roth. Modeling the Impact of Checkpoints on Next-Generation Systems. In *24th IEEE Conf. Mass Storage Systems and Technologies*. IEEE, 2007.

[38] M. W. Rashid and M. C. Huang. Supporting highly-decoupled thread-level redundancy for parallel programs. In *14th Int. Conf. on High-Performance Computer Architecture (HPCA)*, pages 393–404. IEEE, 2008.

[39] P. Sao and R. Vuduc. Self-stabilizing iterative solvers. In *ScalA '13*, 2013.

[40] B. Schroeder and G. Gibson. Understanding failures in petascale computers. *Journal of Physics: Conference Series*, 78(1), 2007.

[41] B. Schroeder and G. A. Gibson. Understanding Failures in Petascale Computers. *Journal of Physics: Conference Series*, 78(1), 2007.

[42] M. Shantharam, S. Srinivasmurthy, and P. Raghavan. Fault tolerant preconditioned conjugate gradient for sparse linear system solution. In *ICS*. ACM, 2012.

[43] M. Snir and et al. Addressing failures in exascale computing. *Int. J. High Perform. Comput. Appl.*, 28(2):129–173, 2014.

[44] J. Stearley, K. B. Ferreira, D. J. Robinson, J. Laros, K. T. Pedretti, D. Arnold, P. G. Bridges, and R. Riesen. Does partial replication pay off? In *FTXS*. IEEE, 2012.

[45] S. Yi, D. Kondo, B. Kim, G. Park, and Y. Cho. Using Replication and Checkpointing for Reliable Task Management in Computational Grids. In *SC*. ACM, 2010.

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

[47] J. Yu, D. Jian, Z. Wu, and H. Liu. Thread-level redundancy fault tolerant cmp based on relaxed input replication. In *ICCIT*. IEEE, 2011.

[48] Z. Zheng and Z. Lan. Reliability-aware scalability models for high performance computing. In *Cluster Computing*. IEEE, 2009.

[49] J. F. Ziegler, H. W. Curtis, H. P. Muhlfeld, C. J. Montrose, and B. Chin. IBM experiments in soft fails in computer electronics. *IBM J. Res. Dev.*, 40(1):3–18, 1996.

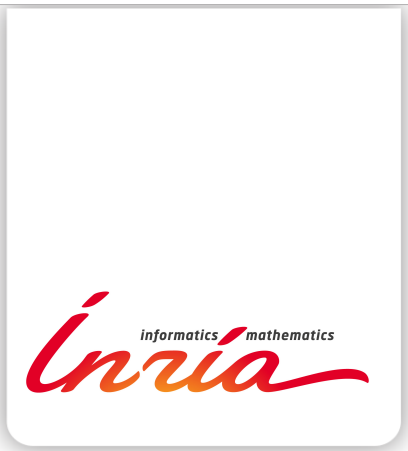

RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES

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