



HAL
open science

Assuming failure independence: are we right to be wrong?

Guillaume Aupy, Yves Robert, Frédéric Vivien

► **To cite this version:**

Guillaume Aupy, Yves Robert, Frédéric Vivien. Assuming failure independence: are we right to be wrong?. [Research Report] RR-9078, Inria. 2017. hal-01556292

HAL Id: hal-01556292

<https://inria.hal.science/hal-01556292v1>

Submitted on 4 Jul 2017

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.



Assuming failure independence: are we right to be wrong?

Guillaume Aupy, Yves Robert, Frédéric Vivien

**RESEARCH
REPORT**

N° 9078

July 2017

Project-Team ROMA



Assuming failure independence: are we right to be wrong?

Guillaume Aupy*, Yves Robert^{†‡}, Frédéric Vivien[†]

Project-Team ROMA

Research Report n° 9078 — July 2017 — 34 pages

Abstract: This report revisits the *failure temporal independence* hypothesis which is omnipresent in the analysis of resilience methods for HPC. We explain why a previous approach is incorrect, and we propose a new method to detect failure cascades, i.e., series of non-independent consecutive failures. We use this new method to assess whether public archive failure logs contain failure cascades. Then we design and compare several cascade-aware checkpointing algorithms to quantify the maximum gain that could be obtained, and we report extensive simulation results with archive and synthetic failure logs. Altogether, not only are there but a few logs that contain cascades, but we show that the gain that can be achieved from this knowledge is not significant. The conclusion is that we can wrongly, but safely, assume failure independence!

Key-words: checkpoint, failure independence, resilience.

* Inria & LaBRI, University of Bordeaux, France

† LIP, École Normale Supérieure de Lyon, CNRS & Inria, France

‡ Univ. Tenn. Knoxville, USA

**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

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

Au sujet de l'indépendance des fautes: avons-nous raison d'avoir tort?

Résumé : Ce rapport étudie l'indépendance temporelle des pannes qui surviennent dans les plates-formes de calcul scientifique à grande échelle. L'hypothèse d'indépendance est omniprésente dans les études. Un travail récent détecte des cascades, i.e., des séries de pannes consécutives et non-indépendantes, mais nous montrons que l'approche est incorrecte. Nous proposons une nouvelle approche, que nous mettons en oeuvre pour détecter des cascades dans les traces d'archive publiques. Certaines de ces traces contiennent bien des cascades. Puis nous concevons et comparons plusieurs algorithmes de checkpoint qui tiennent compte de la présence de cascades, dont plusieurs oracles, et nous évaluons leur performance par rapport à l'algorithme périodique classique, à la fois sur des traces d'archive et sur des traces synthétiques qui contiennent des cascades "artificielles". Le gain potentiel s'avère négligeable, et la conclusion est que nous pouvons supposer l'indépendance temporelle des pannes, à tort certes mais sans perte de performance avérée.

Mots-clés : checkpoint, indépendance temporelle des fautes, résilience.

1 Introduction

This work revisits the *failure*¹ *temporal independence* hypothesis in HPC (High Performance Computing) systems. Assuming failure temporal independence is mandatory to analyze resilience protocols. To give a single example: the well-known Young/Daly formula for the optimal checkpointing period [24, 7] is valid only if failure inter-arrival times, or IATs, are IID (Independent and Identically Distributed) random variables. We aim at providing a quantitative answer to the following question: to what extent are failures temporally independent? We base our analysis on publicly available failure logs from LANL [15, 14] and Tsubame [23]. We show that a previously proposed approach based on degraded intervals [3] leads to incorrect results, and we propose a new algorithm to detect failure cascades, based on the study of pairs of consecutive IATs. This new algorithm is used for the largest six public logs at our disposal, and we detect cascades in one log for sure, and possibly in a second one. The first conclusion is that it is wrong to assume failure independence everywhere!

The next question is to assess which gain can be achieved when equipped with the knowledge that failure cascades are present in some logs (and in real life, on some large-scale execution platforms). We design and compare several cascade-aware checkpointing algorithms. Four algorithms are simple periodic algorithms, with a constant checkpointing period obeying the Young/Daly formula $\sqrt{2C\mu}$; they differ by the value chosen for the MTBF μ . The choices for μ are: (i) the MTBF of the entire log; (ii) the MTBF of the log expunged of failures present in degraded intervals [3]; (iii) the MTBF of the log expunged of failures whose IATs belong to the first quantile (our new algorithm); and (iv) a brute-force algorithm that searches all possible periods (used for reference). The remaining seven algorithms are more sophisticated in that they use two different regimens: a *normal regimen* for failure-free segments, with a large checkpointing period; and a *degraded regimen* which is entered after a failure, with a smaller checkpointing period to cope with potential cascades. We compare different versions based upon the work in [3] and upon our new cascade detection algorithm. Finally, we use brute-force methods and oracle-based solutions to fully quantify the maximum gain that could be achieved with omniscient knowledge of forthcoming failures. Altogether, the overall conclusion is that there is not much to gain from the knowledge of failure cascades: we can just ignore them without any significant overhead. The second and final conclusion is that we can wrongly, but safely, assume failure independence!

The main contributions of this work are:

- The correct evaluation of the method in [3] to detect failure cascades (Section 3.1);
- A novel method to detect failure cascades, based on the quantile distribution of consecutive IAT pairs (Section 3.2);
- The design and comparison of several cascade-aware checkpointing algorithms to assess the potential gain of cascade detection (Sections 4.1 and 4.2);
- Extensive evaluation via archive and synthetic logs of all algorithms (Section 5).

In addition to the above sections, Section 2 reviews relevant related work, and Section 6 provides concluding remarks.

¹Throughout the report, we use the work *failure* and *fault* interchangeably, to denote an unrecoverable interruption of resource execution, a.k.a a fail-stop error.

2 Related Work

Reliability is key to future extreme-scale HPC systems. The de-facto standard approach, namely the coordinated checkpointing protocol [6] has been recently extended in several directions, such as hierarchical checkpointing [5], in-memory checkpointing [25, 18], or multi-level checkpointing [17, 4], to quote just a few. To develop better solutions and reliability techniques, it is important to understand and characterize failure behavior. Indeed, failure characteristics can be used to inform failure predictors [8, 2], or to improve fault-tolerance techniques [11, 10, 22].

In their seminal work, Young [24] and Daly [7] assume that failure IATs are IID and follow an Exponential probability distribution. This assumption was key to derive their famous formula for the optimal checkpoint interval. Other distributions have been considered, such as Weibull distributions [9, 22, 11]. Formulas for the optimal checkpointing period have been obtained by Gelenbe and Hernández [9], but still under the temporal independence assumption. In stochastic terms, each checkpoint is assumed to be a renewal point, which (unrealistically) amounts to rebooting the entire platform after each failure. Similarly, Tiwari et al. [22] and Heien et al [11] confirmed the observation that the Young/Daly formula is a very good approximation for Weibull distributions. In particular, Tiwari et al. [22] analyze a failure log, show that it matches well a Weibull distribution through Kolmogorov-Smirnov and other statistical tests, and report some gains when using the fact that much more than half of IATs are smaller than the expectation of the Weibull law (due to its infant mortality property).

It is important to understand the impact of the renewal condition on the works that deal with non-Exponential distribution. Consider a platform with several processors, each subject to failures whose IATs follow a Weibull (or, in fact, any other, except Exponential) probability distribution. Without the renewal condition, the distribution of failure IATs over the whole platform is no longer independent nor identically distributed, which complicates everything, including the mere definition of the platform MTBF (see [12] for details). This explains why all these previous works consider that failures are temporally independent.

However, it has been observed many times that when a failure occurs, it may trigger other failures that will strike different system components [11, 22, 3]. As an example, a failing cooling system may cause a series of successive crashes of different nodes. Also, an outstanding error in the file system will likely be followed by several others [20, 13]. Recently Bautista-Gomez et al. [3] studied nine systems, and they report periods of high failure density in all of them. They call these periods *cascade failures*. This observation has led them to revisit the temporal failure independence assumption, and to design bi-periodic checkpointing algorithms that use different periods in normal (failure-free) and degraded (with failure cascades) modes. See Sections 3.1 and 4.2 for a full description of their approach [3].

Finally, this report focuses on temporal properties of failures. Spatial properties of failures [10] are out of scope for this study.

3 Algorithms to Detect Failure Cascades

We informally define a failure cascade as a series of consecutive failures that strike closer in time than one would normally expect. We explain how to refine such an imprecise definition by describing two approaches below. The first approach is introduced in [3] and is based

upon *degraded intervals*. The second approach is a major contribution of this report and uses *quantiles of consecutive IAT pairs*.

3.1 Degraded Intervals

In their recent work, Bautista-Gomez et al. [3] provide Algorithm 1 to detect *degraded intervals*, i.e., intervals containing failure cascades. Consider a failure log of total duration L seconds, with n failures striking at time-steps t_i , where $0 \leq t_1 \leq \dots \leq t_n \leq L$. By definition, the MTBF (Mean Time Between Failures) of the platform log is $\mu = \frac{L}{n}$. Intuitively, we expect failures to strike every μ seconds in average. Algorithm 1 simply divides the log into n intervals of length L/n and checks for intervals that contain two or more failures. Such intervals are called *degraded* while intervals with zero or one failure are called *normal*. Algorithm 1 returns the set \mathcal{S}_c of degraded intervals. The percentage of degraded intervals is $P_{deg} = \frac{|\mathcal{S}_c|}{n}$.

Algorithm 1 Identifying the set \mathcal{S}_c of degraded intervals.

```

1: procedure DETCASCADES( $\{t_1, \dots, t_n\}, L, X$ )
2:   for  $i = 0$  to  $n - 1$  do
3:     if  $|\{j | t_j \in [i\frac{L}{n}, (i+1)\frac{L}{n}]\}| \geq 2$  then
4:        $\mathcal{S}_c = \mathcal{S}_c \cup \{j | t_j \in [i\frac{L}{n}, (i+1)\frac{L}{n}]\}$ 
5:     end if
6:   end for
7:   return  $\mathcal{S}_c$ 
8: end procedure

```

Table 1 summarizes the percentage of degraded interval P_{deg} found by Algorithm 1 for the LANL and Tsubame failure logs. Altogether, 20% to 30% of system time is spent in degraded regimen. The authors of [3] conclude that Algorithm 1 has identified failure cascades. This conclusion is incorrect:

- Assume first that failure IATs are IID and follow an Exponential probability distribution $\text{EXP}[\lambda]$ with parameter λ (we have $\mathbb{P}(X \leq t) = 1 - e^{-\lambda t}$). For a failure log of size L and with n failures, we directly have $\lambda = \frac{1}{\mu} = \frac{n}{L}$. Theorem 1 shows that in this case, the expected percentage of degraded intervals is $1 - \frac{2}{e} \approx 26.4$.
- Assume now that failure IATs are IID and follow a Weibull probability distribution $\text{WEIBULL}[k, \lambda]$ with parameter shape parameter k and scale parameter λ : we have $\mathbb{P}(X \leq t) = 1 - e^{-(\frac{t}{\lambda})^k}$, and the MTBF is $\mu = \lambda \Gamma(1 + \frac{1}{k})$. For a failure log of size L and with n failures, we have $\mu = \frac{L}{n}$, hence we let $\lambda = \frac{L}{n \Gamma(1 + \frac{1}{k})}$. Table 2 shows the value of P_{deg} for values of the shape parameter k ranging from 0.5 to 1. The value for $k = 1$ is the same as for the Exponential distribution, because $\text{WEIBULL}[1, \lambda]$ reduce to $\text{EXP}[\frac{1}{\lambda}]$. For all the values of k in Table 2, the percentage of degraded intervals lies between 26% and 27.5%. Note that these values are independent of λ and are obtained experimentally, using MonteCarlo simulations. Note also that typical values of k used to model failures in the literature are ranging from $k = 0.5$ to $k = 0.9$ [16, 19, 22].

Altogether, the correct conclusion from Theorem 1 and the results in Table 1 is that public logs exhibit the same number of degraded intervals as pure $\text{EXP}[\lambda]$ and $\text{WEIBULL}[k, \lambda]$ renewal processes. Hence we cannot conclude anything!

Id	Log		Approach in [3]		
	Number of faults	MTBF in hours	MTBF in hours	Degraded intervals: P_{deg}	Faults in cascades
LANL 2	5351	14.1	36.4	25.3%	71.1%
LANL 3	294	59.3	142.3	26.3%	69.4%
LANL 4	298	56.2	126.7	24.9%	66.8%
LANL 5	304	54.7	119.6	26.4%	66.4%
LANL 6	63	279.9	604.1	33.9%	69.8%
LANL 7	126	311.1	943.1	21.6%	74.6%
LANL 8	448	84.5	226.6	26.2%	72.5%
LANL 9	278	58.8	216.0	23.1%	79.1%
LANL 10	234	68.7	218.3	23.6%	76.1%
LANL 11	265	60.8	230.5	23.9%	80.0%
LANL 12	254	64.2	192.7	25.3%	75.2%
LANL 13	193	83.4	380.7	24.0%	83.4%
LANL 14	120	103.7	410.6	20.0%	80.0%
LANL 15	53	124.1	292.0	23.1%	67.9%
LANL 16	2262	21.9	56.2	25.2%	70.9%
LANL 17	125	216.9	526.1	21.8%	68.0%
LANL 18	3900	7.5	17.9	26.0%	68.9%
LANL 19	3222	7.9	17.1	26.4%	66.0%
LANL 20	2389	13.7	41.5	21.3%	74.1%
LANL 21	105	24.2	79.9	26.9%	78.1%
LANL 22	235	272.1	696.6	27.8%	71.9%
LANL 23	448	147.3	348.8	23.7%	67.9%
LANL 24	150	412.7	1040.6	22.1%	69.3%
Tsubame	884	14.8	36.5	23.9%	69.2%

Table 1: Percentage of degraded intervals in failure logs. The expected value of this percentage for an Exponential distribution is 26.4%. **Red** entries are within 5% of this value. **Pink** entries are within 10% of this value.

Shape parameter k	Degraded intervals: P_{deg}	Faults in cascades
0.50	26.0%	84.7%
0.52	26.3%	83.7%
0.54	26.6%	82.7%
0.56	26.8%	81.6%
0.58	27.0%	80.6%
0.60	27.2%	79.7%
0.62	27.3%	78.7%
0.64	27.4%	77.7%
0.66	27.4%	76.8%
0.68	27.5%	75.9%
0.70	27.5%	75.0%
0.72	27.5%	74.1%
0.74	27.5%	73.2%
0.76	27.5%	72.4%
0.78	27.4%	71.5%
0.80	27.4%	70.7%
0.82	27.3%	69.9%
0.84	27.3%	69.1%
0.86	27.2%	68.3%
0.88	27.1%	67.5%
0.90	27.0%	66.8%
0.92	26.9%	66.0%
0.94	26.8%	65.3%
0.96	26.7%	64.6%
0.98	26.6%	63.9%
1.00	26.4%	63.2%

Table 2: Expected values of P_{deg} and of the number of faults in cascades for IID IATs following a Weibull distribution $WEIBULL[k, \lambda]$.

Theorem 1. For a log duration L , and IID failure IATs with Exponential distribution $\text{EXP}[\lambda]$, the percentage of degraded intervals P_{deg} converges to

$$\lim_{L \rightarrow \infty} P_{deg} = 100\left(1 - \frac{2}{e}\right) \approx 26.4\%$$

Proof. Assume Exponentially-distributed failure IATs with a MTBF $\mu = \frac{1}{\lambda}$. The expected number of failures is $\frac{L}{\mu}$. Let $A(T)$ be the event ‘there are at least two failures in an interval of length T ’:

$$\mathbb{P}(A(T)) = 1 - e^{-\lambda T} - \lambda T e^{-\lambda T}$$

We want to compute the expectation of $B(L)$, the ratio of intervals with at least two failures for a log size of L . Of course we have $P_{deg} = 100B(L)$ in terms of percentage. We define X the random variable that represents the number of failures during a log of size L . Let $B(L|X = i)$ be the random variable: the ratio of the number of intervals with at most two failures in a log of size L when there were i failures. Finally, let Y_0, \dots, Y_{i-1} denote the random variables such that Y_j is equal to 1 if there are at least two failures in the interval $[jL/i, (j+1)L/i]$ and 0 otherwise. Then $\mathbb{E}[B(L|X = i)] = \frac{\sum_{j=1}^i \mathbb{E}[Y_j]}{i}$ if $i > 1$ and 0 otherwise. To compute $\mathbb{E}[Y_j]$, we write:

$$\begin{aligned} \mathbb{E}[Y_j] &= \sum_{k=2}^i \mathbb{P}(\text{there are } k \text{ faults during the } j^{\text{th}} \text{ interval}) \\ &= 1 - \mathbb{P}(\text{there are 0 fault during the } j^{\text{th}} \text{ interval}) \\ &\quad - \mathbb{P}(\text{there are 1 fault during the } j^{\text{th}} \text{ interval}). \end{aligned}$$

The event ‘‘there are k faults during the j^{th} interval’’ is equivalent to saying ‘‘there are k faults during L/i units of time and $i - k$ faults during $(i - 1)L/i$ units of time, given that there are i faults during L units of time’’. Let $C(k, T)$ denote the event ‘‘there are k failures in an interval of length T ’’. By definition $\mathbb{P}(C(k, T)) = \frac{(\lambda T)^k}{k!} e^{-\lambda T}$. Finally,

$$\begin{aligned} &\mathbb{P}(\text{there are } k \text{ faults during the } j^{\text{th}} \text{ interval}) \\ &= \mathbb{P}(C(k, L/i) \wedge C(i - k, (i - 1)L/i) | C(i, L)) \\ &= \frac{\mathbb{P}(C(k, L/i)) \mathbb{P}(C(i - k, (i - 1)L/i))}{\mathbb{P}(C(i, L))} \\ &= \frac{(\lambda L/i)^k e^{-\lambda L/i} (\lambda (i - 1)L/i)^{i-k} e^{-\lambda (i-1)L/i}}{k!(i - k)! (\lambda L)^i e^{-\lambda L}} \\ &= \binom{i}{k} \frac{(i - 1)^{i-k}}{i^i}, \end{aligned}$$

and we have: $\mathbb{E}[Y_j] = 1 - (1 - \frac{1}{i})^i - (1 - \frac{1}{i})^{i-1}$.

We can now return to computing the expectation of $B(L)$, the ratio of intervals with at

least two failures for a log size of L , as follows:

$$\begin{aligned}\mathbb{E}[B(L)] &= \sum_{i=0}^{+\infty} \mathbb{P}(X = i) \mathbb{E}[B(L|X = i)] \\ &= \sum_{i=2}^{+\infty} \frac{(\lambda L)^i}{i!} e^{-\lambda L} \left(1 - \left(1 - \frac{1}{i}\right)^i - \left(1 - \frac{1}{i}\right)^{i-1} \right) \\ &= \sum_{i=2}^{+\infty} \frac{(\lambda L)^i}{i!} e^{-\lambda L} \\ &\quad - \sum_{i=2}^{+\infty} \frac{(\lambda L)^i}{i!} e^{-\lambda L} \left(\left(1 - \frac{1}{i}\right)^i + \left(1 - \frac{1}{i}\right)^{i-1} \right)\end{aligned}$$

We show that when L tends to infinity, $\mathbb{E}[B(L)]$ converges to $1 - 2/e$. Consider $\varepsilon > 0$. Both functions $f : x \mapsto \left(1 - \frac{1}{x}\right)^x$ and $g : x \mapsto \left(1 - \frac{1}{x}\right)^{x-1}$ converge to $1/e$ when $x \rightarrow +\infty$, hence there exists i_0 such that for all $i \geq i_0$,

$$\left| \left(1 - \frac{1}{i}\right)^i + \left(1 - \frac{1}{i}\right)^{i-1} - 2/e \right| < \frac{\varepsilon}{4}.$$

Recall that $\sum_{i=i_0}^{+\infty} \frac{(\lambda L)^i}{i!} e^{-\lambda L} \leq \sum_{i=1}^{+\infty} \frac{(\lambda L)^i}{i!} e^{-\lambda L} = 1$. Thus, there exists L_1 such that for all $L \geq L_1$,

$$\left| e^{-\lambda L} \sum_{i=2}^{i_0} \frac{(\lambda L)^i}{i!} \left(\left(1 - \frac{1}{i}\right)^i + \left(1 - \frac{1}{i}\right)^{i-1} - 2/e \right) \right| < \frac{\varepsilon}{4}.$$

For the same reason, there exists L_2 such that for all $L \geq L_2$,

$$\left| (\lambda L) e^{-\lambda L} \right| < \frac{\varepsilon}{4}.$$

Altogether, for all $L \geq L_0 = \max(L_1, L_2)$,

$$\begin{aligned}
& |\mathbb{E}[B(L)] - (1 - 2/e)| \\
&= |\mathbb{E}[B(L)] - (1 - 2/e) \sum_{i=1}^{+\infty} \frac{(\lambda L)^i}{i!} e^{-\lambda L}| \\
&= \left| -\lambda L e^{-\lambda L} \right. \\
&\quad \left. - e^{-\lambda L} \sum_{i=2}^{+\infty} \frac{(\lambda L)^i}{i!} \left(\left(1 - \frac{1}{i}\right)^i + \left(1 - \frac{1}{i}\right)^{i-1} - \frac{2}{e} \right) \right. \\
&\quad \left. + \frac{2}{e} \lambda L e^{-\lambda L} \right| \\
&\leq \left| \lambda L e^{-\lambda L} \right| \\
&\quad + \left| e^{-\lambda L} \sum_{i=2}^{i_0} \frac{(\lambda L)^i}{i!} \left(\left(1 - \frac{1}{i}\right)^i + \left(1 - \frac{1}{i}\right)^{i-1} - \frac{2}{e} \right) \right| \\
&\quad + \left| e^{-\lambda L} \sum_{i=i_0}^{+\infty} \frac{(\lambda L)^i}{i!} \left(\left(1 - \frac{1}{i}\right)^i + \left(1 - \frac{1}{i}\right)^{i-1} - \frac{2}{e} \right) \right| \\
&\quad + \left| \frac{2}{e} \lambda L e^{-\lambda L} \right| \\
&\leq \varepsilon
\end{aligned}$$

We have derived that for all $\varepsilon > 0$, there exists L_0 such that for all $L > L_0$, $|\mathbb{E}[B(L)] - (1 - 2/e)| \leq \varepsilon$. This concludes the proof. \square

3.2 Quantile distribution of consecutive IAT pairs

In this section, we introduce a new method to detect failure cascades. In a nutshell, the method analyzes the distribution of pairs of two consecutive IATs. Intuitively, consider a failure log and its IATs. If small values are scattered across the log, we do not conclude anything. On the contrary, if a small value follows another small value, we may have found the beginning of a cascade. Our approach checks the frequency of having two consecutive small values, and compares this frequency with the expected value when IATs are independent. Our approach proceeds as follows. Consider a failure log with $n = N + 2$ failures, i.e., with $N + 1$ IATs z_i , $1 \leq i \leq N + 1$. Note that the first failure in the log does not correspond to an IAT since its predecessor is not included in the log. Note also that $z_i = t_{i+1} - t_i$ where t_i is the time where the i -th failure strikes, as defined in Algorithm 1. Finally, note that there are N pairs (z_i, z_{i+1}) of consecutive IATs, $1 \leq i \leq N$.

We start by sampling the failure log and computing quantiles. In Figure 1 we plot the cumulative distribution of IATs for failure log LANL2. In Figure 1, we use $Q = 10$ quantiles. The 10% smallest values constitute the first quantile, or *limit quantile* Q_{limit} , and are between 0 and 2,220 seconds. The next 10% smallest values (second quantile) are between 2,221 and 6,000 seconds, and so on. By definition, the probability that an IAT belongs to a given quantile is $\frac{1}{Q} = 0.1$. Now, if we assume (temporal) failure independence, the probability that

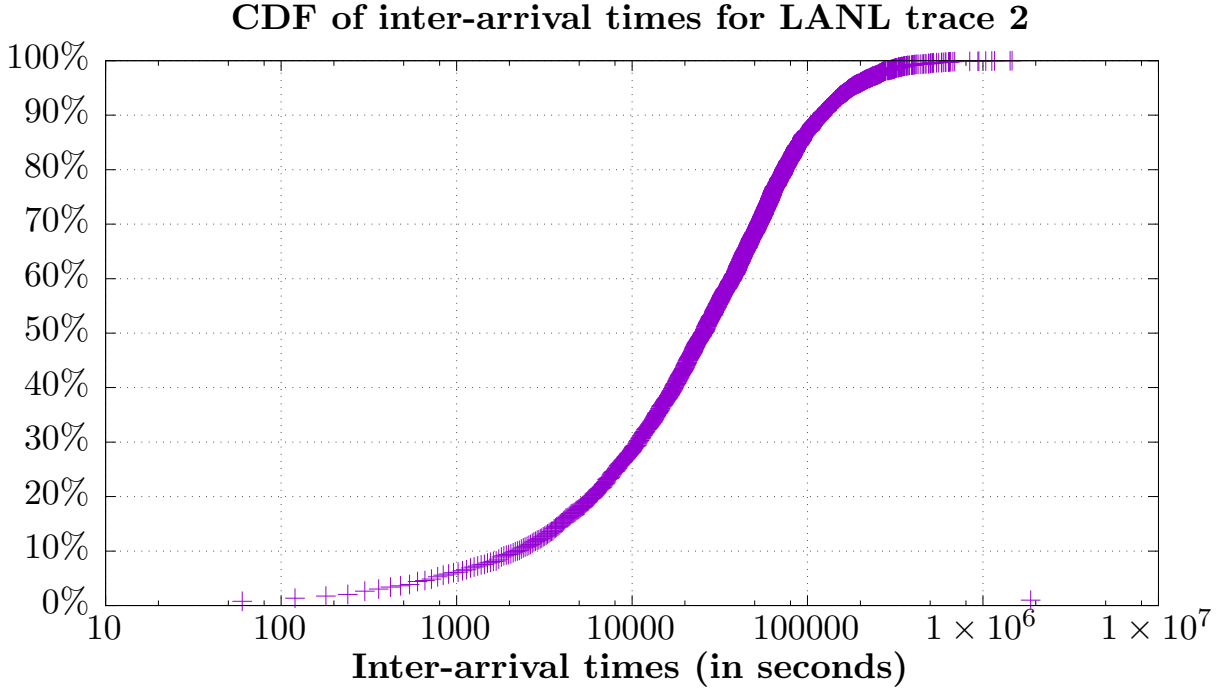


Figure 1: Cumulative plot of IATs for failure log LANL 2.

both components of a pair (z_i, z_{i+1}) of consecutive IATs belongs to the same given quantile is $\frac{1}{Q^2}$, and the expected number of such pairs is $\frac{N}{Q^2}$. We need N to be reasonably large so that this expectation is accurate. Out of the 24 logs in Table 1, we keep only the five LANL logs with $N \geq 1000$ (namely LANL 2, 16, 18, 19, 20) and the Tsubame log: see Table 3. For these 6 logs, in Figures 2 to 7, we plot the ratio of the actual number of pairs in each quantile over the expected value. This ratio is called the lag plot density [1, 21]. We expect a log without cascades to exhibit a flat surface up to a few statistical artefacts. The figures lead to the conclusions reported in Table 3: only log LANL2 contains cascades for sure, because the ratio for the first quantile is greater than four times its expected value; maybe LANL20 does too (ratio between 2 and 3); the other logs do not include any cascade.

Note that another phenomenon can be observed in LANL2 and, to some extent, in LANL20: there is an *anti-cascade* behavior, where long IATs are followed by other long IATs more often than expected. We are not able to give any explanation to this phenomenon.

Altogether, there are indeed some cascades, albeit not very frequent, in some failure logs. Hence we were wrong to assume failure independence everywhere. The next question is whether the knowledge that cascades are present may help reduce the overhead due to the standard checkpoint/recovery approach. The rest of the report is devoted to answering this question.

We conclude this section with three more lag plots to help the reader understand and assess temporal correlation in failure logs. First, Figures 8 and 9 show the lag plots for an $\text{EXP}[\lambda]$ and $\text{WEIBULL}[0.7, \lambda]$ IAT distributions: as expected, there is no cascade in these renewal processes. Then, Figure 10 shows the lag plot for LANL2 after randomly shuffling the log IATs; compare with Figure 2: there is no more cascade!

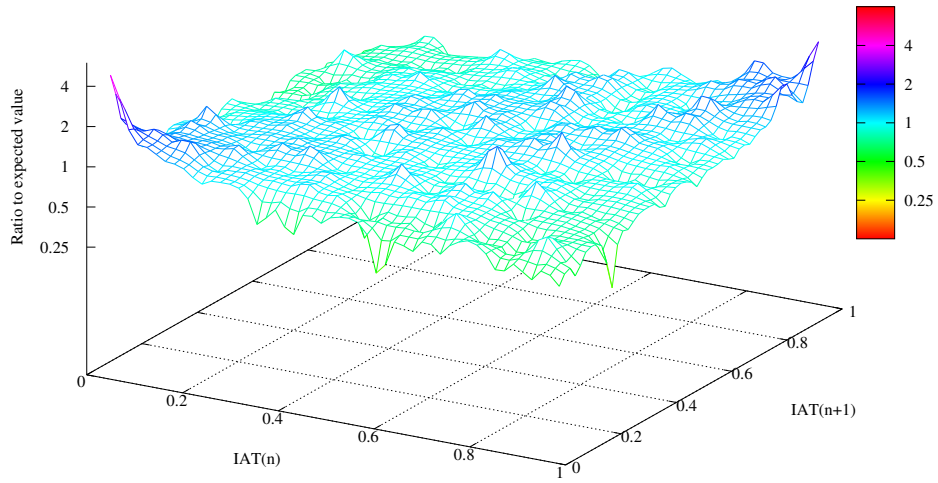


Figure 2: Lag plot for failure log LANL2.




Figure 3: Lag plot for failure log LANL16.

Log	Number of Faults	Cascades
LANL 2	5351	Yes
LANL 16	2262	No
LANL 18	3900	No
LANL 19	3222	No
LANL 20	2389	Maybe
Tsubame	884	No

Table 3: Presence of cascades in large logs.




Figure 4: Lag plot for failure log LANL18.




Figure 5: Lag plot for failure log LANL19.




Figure 6: Lag plot for failure log LANL 20.




Figure 7: Lag plot for failure log Tsubame.




Figure 8: Lag plot for an $\text{EXP}[\lambda]$ IAT log.




Figure 9: Lag plot for a $\text{WEIBULL}[0.7, \lambda]$ IAT log.




Figure 10: Lag plot for failure log LANL2 after randomized shuffling.

4 Cascade-aware checkpointing

In this section we provide a quantitative assessment of many algorithms that can be used to improve the classical periodic checkpointing algorithm, whose period is given by the Young/Daly formula [24, 7]. We use both the previous public logs and synthetic logs to generate simulation results.

When playing an algorithm against a log, it is common practice to have the algorithm learn data from, say, the first half of the log, and play against the second half of the log. More precisely, we would start the algorithm from a randomly generated instant in the second half of the log to avoid bias. This instant must not too close to the end of the log to ensure that there remain enough failures to strike, e.g., any instant in the thirist quarter of the log would be fine. In the following, we do not adopt this strategy. Instead, we let the algorithms learn form the entire log, and replay them from a randomly chosen instant. This gives a somewhat unfair advantage to the algorithms, but our goal is to provide an upper bound of the maximum gains that can be achieved.

4.1 Periodic algorithms

What can we learn from a log? The Young/Daly checkpointing algorithm needs to know the checkpoint time C and the MTBF μ_{log} , and then uses the optimal checkpointing period $T = \sqrt{2\mu_{log}C}$. As already mentioned, one can directly use $\mu_{log} = \frac{L}{N}$ for a log of length L with N failures. However, inspecting the log for cascades can lead to refined values of the MTBF:

- Remember that [3] defines two categories of intervals in the log, *normal* and *degraded*, and computes μ_{normal_int} and $\mu_{degraded_int}$, the MTBF for each set of intervals. Typically, μ_{normal_int} will be significantly larger than μ_{log} , and using this value instead of μ_{log} will decrease the failure-free overhead incurred by checkpointing too frequently outside the cascades.
- Following the approach in Section 3.2, we can divide the log into Q quantiles and separate IATs that belong to the first quantile Q_{limit} from the other ones. IATs that belong

to the first quantile Q_{limit} constitute the cascades and have mean value $\mu_{cascade}(Q)$. The other IATs have mean value $\mu_{non-cascade}(Q)$. Just as μ_{normal_int} , $\mu_{non-cascade}(Q)$ will be larger than μ_{log} , with the same desired impact. We use different values of $|Q_{limit}|$ for the simulations: $|Q_{limit}| = 10\%$, and $|Q_{limit}| = 5\%$ in the core of this report, and, in addition, $|Q_{limit}| \in \{2\%, 1\%, 0.5\%, 0.2\%, 0.1\%\}$ in Appendix A.

Altogether, we evaluate 4 periodic checkpointing algorithms, which use the period $T = \sqrt{2\mu C}$, where μ is chosen from the following:

- Algorithm Π_{Daly} uses $\mu = \mu_{log}$
- Algorithm $\Pi_{Intervals}$ uses $\mu = \mu_{normal_int}$
- Algorithm $\Pi_{Quantiles}$ uses $\mu = \mu_{non-cascade}(Q)$
- Algorithm Π_{Best_period} uses a brute-force search and returns the best period

For each algorithm, we report the WASTE, defined as the fraction of time where the platform does not perform useful work. Experimental values of WASTE are averaged from many Monte Carlo simulations. As a side note, a first-order approximation is given by the formula

$$\text{WASTE} = \frac{C}{T} + \frac{1}{\mu} \left(R + \frac{T}{2} \right) \quad (1)$$

and is obtained as follows [12]: the first term $\frac{C}{T}$ is the overhead in a failure free execution, since we lose C seconds to checkpoint, every period of T seconds. The second term is the overhead due to failures, which strike every μ seconds in expectation; for each failure, we lose R seconds for recovery (letting $R = C$ everywhere) and re-execute half the period on average. Equation (1) explains that checkpointing algorithms aim at finding the best trade-off between checkpointing too often (large failure-free overhead) and not often enough (large overhead after a failure); the waste is minimized when $T = \sqrt{2\mu C}$. All this helps understand how WASTE depends on the value chosen for the MTBF μ .

4.2 Bi-periodic algorithms

We compare the following seven bi-periodic algorithms. Each bi-periodic algorithm uses two different regimens, namely the *normal* and *degraded* regimens, to borrow the terminology of [3]. In the *normal* regimen, which is the regimen by default, the algorithm runs in the absence of failures, hence uses a larger checkpointing period. In the *degraded* regimen, the algorithm uses a shorter checkpointing period, to cope with potential cascades.

The seven algorithms differ by several parameters:

- the MTBF value μ used for each regimen, which dictates the corresponding checkpointing period. In Table 4, we report the MTBF for the normal regimen μ_{normal} and for the MTBF for the degraded regimen $\mu_{degraded}$. Again, the checkpointing period for the normal regimen is $T = \sqrt{2\mu_{normal}C}$ and that for the degraded regimen is $T = \sqrt{2\mu_{degraded}C}$.
- the criteria used to enter and exit the cascade regimen. Most algorithms enter the cascade regimen as soon as a failure strikes, but *lazy* variants enter the cascade regimen only after a second failure has struck, and provided that the IAT belongs to the first quantile Q_{limit} . All algorithms exit the degraded regimen when enough time has elapsed since the last failure. Following [3], we set this timeout to $2\mu_{degraded}$.

All these parameters are listed in Table 4. For reference, here follows a detailed description of each bi-periodic algorithm. First, we have 3 algorithms based on refined values of the MTBF:

Algorithm	μ_{normal}	$\mu_{degraded}$	Enter criterion
BI-II _{Intervals}	μ_{normal_int}	$\mu_{degraded_int}$	First failure
BI-II _{Quantiles}	$\mu_{cascade}(Q)$	$\mu_{non-cascade}(Q)$	First failure
BI-II-BEST	Best value for μ_{normal}	Best value for $\mu_{degraded}$	First failure
BI-II _{Quantiles} LAZY	$\mu_{cascade}(Q)$	$\mu_{non-cascade}(Q)$	Second failure in first quantile
BI-II _{Quantiles} LAZY-BEST	Best value for μ_{normal}	Best value for $\mu_{degraded}$	Second failure in first quantile
BI-II _{Quantiles} ORACLE	$\mu_{cascade}(Q)$	Omniscient	Omniscient
BI-II-ORACLE-BEST	Best value for μ_{normal}	Omniscient	Omniscient

Table 4: Bi-periodic algorithms

- BI-II_{Intervals} uses $\mu_{normal} = \mu_{normal_int}$ and $\mu_{degraded} = \mu_{degraded_int}$ for the checkpointing periods, as proposed by [3]. It enters the degraded mode as soon as a failure strikes, and exits it with timeout $2\mu_{degraded}$.
- BI-II_{Quantiles} works similarly, but with different MTBF values. $\mu_{normal} = \mu_{non-cascade}(Q)$ and $\mu_{degraded} = \mu_{cascade}(Q)$. These values are computed from the quantiles of the log IATs.
- BI-II_{Quantiles}LAZY is a variant of BI-II_{Quantiles} where the degraded mode is entered only after two successive failures, provided that the second failure strikes shortly (IAT in first quantile Q_{limit}) after the first one.

The next 2 algorithms use brute-force search:

- BI-II-BEST uses a brute-force method for everything: it computes the waste for all values of μ_{normal} , $\mu_{degraded}$ and timeout value, and retains the best triple. BI-II-BEST is agnostic of the cascade detection algorithm; its only limitation is that it enters the degraded mode after the first failure.
- BI-II_{Quantiles}LAZY-BEST is the lazy variant of BI-II-BEST. However, to decide whether to enter the degraded mode, just as BI-II_{Quantiles}LAZY, it needs to know the quantiles of the distribution to check whether the IAT of the second failure belongs to the first quantile Q_{limit} .

Finally, the last two algorithms are included for reference. They use oracles that know everything in cascades, including the future! Specifically:

- BI-II_{Quantiles}ORACLE uses $\mu_{normal} = \mu_{non-cascade}(Q)$ in normal mode, just as BI-II_{Quantiles}. However, as soon as a failure strikes, it knows exactly whether there will be a cascade, and when the next failures in that cascade will strike. It can thus checkpoint as late as possible, completing the checkpoint right before the failure. And it also knows in advance when the next failure is far away (IAT not in the first quantile Q_{limit}), so that it can immediately switch back to normal mode.
- BI-II-ORACLE-BEST is the variant of BI-II_{Quantiles}ORACLE that tests all possible values of μ_{normal} in normal mode, not just $\mu_{non-cascade}(Q)$. It behaves exactly the same after a failure. A comparison with BI-II_{Quantiles}ORACLE will help assess whether using $\mu_{normal} = \mu_{non-cascade}(Q)$ is a good decision or not.

5 Simulation results

5.1 Simulation setup

In addition to the six large public logs of Table 3, we generate synthetic logs. We first generate failures according to an Exponential distribution of MTBF $\mu_1 = 3,600$ seconds. Such an MTBF of 1 hour is much smaller than archive logs MTBFs in Table 3, which are of the order of 10 hours. This is because we want the waste to be higher for synthetic logs than for archive logs, so that there is a potential significant gain with cascade detection. Next we perturb the log by randomly generating cascades: after each failure, we generate a cascade with frequency (probability that this failure is the start of a cascade) $f = 1\%$, $f = 5\%$ or $f = 10\%$. The length of a cascade, defined as the number of additional failures generated (thus not counting the original failure), is a random value between 3 and 5 (we write $\ell = 3 - 5$), or between 3 and 10 (we write $\ell = 3 - 10$). Finally, the failures in the cascades follow another Exponential distribution of MTBF $\mu_2 = \frac{\mu_1}{\rho}$, where the ratio ρ is chosen in $\{10, 100, 1000\}$. Altogether, a synthetic log is tagged 'Synth. $\rho|f|\ell$ '. For instance, 'Synth. $\rho = 100|1\%|3 - 5$ ' means that a cascade is generated every 100 failures in average ($f = 1\%$), with MTBF 36 seconds, one hundredth of the original distribution ($\rho = 100$), and a number of additional failures, not including the first one, uniformly drawn between 3 and 5. The lag plot for this log is shown in Figure 11.




Figure 11: Lag plot for Synth. $\rho = 100|1\%|3 - 5$.

For each archive and synthetic log, we average results over 100 executions. We draw a starting point uniformly between the beginning of the log and $200\mu_{log}$ seconds before its end, and we run the simulation during $100\mu_{log}$ seconds, not counting checkpoint and re-execution overhead. Finally, we use a wide range of checkpoint values, namely $C = 300$, $C = 30$ and $C = 3$ seconds, in order to cover the widest range of scenarios.

5.2 Waste values

We report results for $|Q_{limit}| = 10\%$ and 5% . Values for $|Q_{limit}| \in \{2\%, 1\%, 0.5\%, 0.2\%, 0.1\%\}$ are available in Appendix A.

Log	Id	Number of faults	MTBF in hours	Approach: Intervals			Approach: Quantiles					
				MTBF in hours	Degraded intervals	Faults in cascades	MTBF in hours	Faults in cascades	Common faults	Number of cascades	MTBF	Avg. length
LANL 2	5351	14.14	36.45	25.3%	71.1%	15.71	16.9%	16.9%	370	0.250	2.4	11
LANL 16	2262	21.85	56.19	25.2%	70.9%	24.30	18.5%	18.5%	191	0.303	2.2	7
LANL 18	3900	7.52	17.88	26.0%	68.9%	8.36	18.7%	18.6%	337	0.240	2.2	7
LANL 19	3222	7.87	17.06	26.4%	66.0%	8.77	19.0%	18.8%	282	0.320	2.2	5
LANL 20	2389	13.66	41.46	21.3%	74.1%	15.19	17.7%	17.7%	182	0.193	2.3	6
Tsubame	884	14.78	36.50	23.9%	69.2%	16.47	18.6%	18.6%	73	0.271	2.2	5
Synth. $\rho = 10 - 1.00\% - 3-5$	3136	0.98	1.99	25.1%	63.1%	1.09	18.0%	17.5%	250	0.036	2.3	7
Synth. $\rho = 10 - 1.00\% - 3-10$	3182	0.95	1.98	25.0%	64.1%	1.06	18.0%	17.9%	254	0.033	2.3	8
Synth. $\rho = 10 - 5.00\% - 3-5$	3591	0.85	1.87	22.4%	64.9%	0.94	17.7%	17.7%	277	0.020	2.3	6
Synth. $\rho = 10 - 5.00\% - 3-10$	4075	0.75	1.77	20.5%	66.2%	0.84	18.1%	18.1%	330	0.014	2.2	7
Synth. $\rho = 10 - 10.00\% - 3-5$	4133	0.72	1.81	22.2%	69.0%	0.80	18.2%	18.1%	338	0.014	2.2	5
Synth. $\rho = 10 - 10.00\% - 3-10$	5067	0.59	1.65	19.5%	71.2%	0.66	18.4%	18.4%	426	0.009	2.2	4
Synth. $\rho = 100 - 1.00\% - 3-5$	3136	0.98	1.99	24.9%	62.9%	1.09	16.3%	16.0%	197	0.020	2.6	7
Synth. $\rho = 100 - 1.00\% - 3-10$	3182	0.95	1.98	24.6%	63.9%	1.06	15.0%	14.9%	157	0.014	3.0	12
Synth. $\rho = 100 - 5.00\% - 3-5$	3591	0.85	1.89	21.6%	64.8%	0.94	16.0%	15.9%	214	0.003	2.7	6
Synth. $\rho = 100 - 5.00\% - 3-10$	4075	0.75	1.79	18.3%	65.7%	0.84	17.1%	17.1%	288	0.002	2.4	6
Synth. $\rho = 100 - 10.00\% - 3-5$	4133	0.72	1.84	20.2%	68.7%	0.80	17.5%	17.5%	308	0.002	2.3	5
Synth. $\rho = 100 - 10.00\% - 3-10$	5067	0.59	1.69	15.8%	70.5%	0.66	18.0%	18.0%	407	0.001	2.2	5
Synth. $\rho = 1000 - 1.00\% - 3-5$	3136	0.98	1.99	24.9%	63.0%	1.09	16.4%	16.1%	199	0.017	2.6	7
Synth. $\rho = 1000 - 1.00\% - 3-10$	3182	0.95	1.98	24.6%	63.9%	1.06	15.0%	14.9%	157	0.009	3.0	12
Synth. $\rho = 1000 - 5.00\% - 3-5$	3591	0.85	1.89	21.3%	64.8%	0.94	15.5%	15.5%	196	0.000	2.8	6
Synth. $\rho = 1000 - 5.00\% - 3-10$	4075	0.75	1.79	18.1%	65.6%	0.84	16.9%	16.9%	280	0.000	2.5	6
Synth. $\rho = 1000 - 10.00\% - 3-5$	4133	0.72	1.83	19.9%	68.6%	0.80	17.3%	17.3%	302	0.000	2.4	5
Synth. $\rho = 1000 - 10.00\% - 3-10$	5067	0.59	1.70	15.3%	70.4%	0.66	18.0%	18.0%	403	0.000	2.3	5

Table 5: Statistics when $|Q_{limit}| = 10\%$.

Log statistics are provided in Tables 5 and 6. The column "Common faults" reports the percentage of failures that are detected as belonging to cascades, by both the interval-based and the quantile-based approaches.

For the waste, we report the improvement or degradation with respect to the reference periodic checkpointing algorithm Π_{Daly} . The color code in Tables 7 to 12 is the following:

- **Green:** Improvement by at least 10%
- **Lime:** Improvement between 5 and 10%
- **Black:** Improvement between 0 and 5%
- **Pink:** Degradation between 0 and 5%
- **Red:** Degradation larger than 10%

5.3 Discussion

Overall, the main take-away is that cascade-aware algorithms achieve negligible gains, except for a few scenarios where the waste is already very low with the standard Π_{Daly} approach. This is true even when considering the best scenarios with (i) very short checkpointing time; (ii) high frequency of cascade failures: and (iii) knowing exactly when the next cascade failures are going to strike. In fact, small gains could be achieved only if cascade failures strike both *frequently*, say more than 10% of the time after an initial failure, and *with a relatively large MTBF*, say, not less than 10% of the log MTBF. We further discuss these statements in the following.

5.3.1 Cascade detection algorithms

Tables 7 to 12 show that both cascade-aware algorithms based on intervals and on quantiles, are not more efficient than the simple Π_{Daly} approach. While the quantile-based approach ($\Pi_{Quantiles}$, $\text{Bi-}\Pi_{Quantiles}$, $\text{Bi-}\Pi_{Quantiles}\text{LAZY}$) seems slightly better than the interval-based approach ($\Pi_{Intervals}$, $\text{Bi-}\Pi_{Intervals}$), the gain (or loss) is still within an error margin (most are between -1% and +1%). Furthermore, as one can see from Tables 7 to 12, the quantile-based approach seems to perform better than the interval-based approach: this is because it detects

Table 12: Waste (and gain) with $C=3s$ and $|Q_{limit}| = 5\%$.

Log	Periodic algorithms				Bi-periodic algorithms					Omniscient oracle	
	Π_{Daly}	$\Pi_{Intervals}$	$\Pi_{Quantiles}$	$\Pi_{BestPeriod}$	Bi-II _{Intervals}	Bi-II _{Quantiles}	Bi-II _{BEST}	Bi-II _{Quantiles} _{LAZY}	Bi-II _{Quantiles} _{LAZY} _{BEST}	Bi-II _{Quantiles} _{ORACLE}	Bi-II _{ORACLE} _{BEST}
LANL 2	0.012	0.013 (-11.94%)	0.011 (0.21%)	0.011 (0.83%)	0.013 (-10.41%)	0.012 (-5.79%)	0.011 (1.35%)	0.011 (0.22%)	0.011 (1.29%)	0.011 (2.12%)	0.011 (2.41%)
LANL 16	0.009	0.011 (-14.17%)	0.009 (-1.11%)	0.009 (0.54%)	0.010 (-11.29%)	0.010 (-5.53%)	0.009 (1.40%)	0.009 (0.07%)	0.009 (1.33%)	0.009 (1.98%)	0.009 (2.97%)
LANL 18	0.015	0.017 (-10.27%)	0.015 (0.42%)	0.015 (0.06%)	0.016 (-9.07%)	0.016 (-4.90%)	0.015 (2.74%)	0.015 (-1.74%)	0.015 (1.79%)	0.015 (2.67%)	0.015 (3.54%)
LANL 19	0.015	0.016 (-8.52%)	0.015 (0.28%)	0.015 (1.05%)	0.016 (-8.57%)	0.016 (-10.12%)	0.015 (1.95%)	0.015 (-1.25%)	0.015 (0.95%)	0.014 (2.99%)	0.014 (3.32%)
LANL 20	0.011	0.012 (-13.09%)	0.011 (0.99%)	0.011 (1.97%)	0.013 (-19.57%)	0.011 (-1.22%)	0.010 (3.00%)	0.011 (0.50%)	0.011 (2.07%)	0.010 (3.05%)	0.010 (3.05%)
Tsubame	0.011	0.012 (-7.71%)	0.011 (-0.55%)	0.011 (4.63%)	0.013 (-19.97%)	0.012 (-8.89%)	0.011 (3.16%)	0.011 (-1.41%)	0.011 (3.36%)	0.011 (1.98%)	0.011 (6.33%)
Synth. $\rho = 10 - 1.00\% - 3-5$	0.043	0.045 (-4.65%)	0.043 (0.03%)	0.043 (0.97%)	0.048 (-11.18%)	0.047 (-8.47%)	0.043 (0.90%)	0.043 (-0.83%)	0.043 (0.96%)	0.043 (0.85%)	0.042 (2.72%)
Synth. $\rho = 10 - 1.00\% - 3-10$	0.044	0.046 (-5.89%)	0.044 (0.57%)	0.044 (0.45%)	0.049 (-11.10%)	0.047 (-6.36%)	0.043 (1.10%)	0.044 (-1.10%)	0.043 (1.29%)	0.043 (2.07%)	0.042 (3.46%)
Synth. $\rho = 10 - 5.00\% - 3-5$	0.045	0.049 (-8.99%)	0.045 (0.22%)	0.045 (0.16%)	0.053 (-15.57%)	0.049 (-4.93%)	0.045 (1.79%)	0.046 (-0.92%)	0.045 (1.01%)	0.042 (7.77%)	0.042 (8.11%)
Synth. $\rho = 10 - 5.00\% - 3-10$	0.048	0.052 (-8.30%)	0.048 (-0.11%)	0.048 (0.37%)	0.058 (-20.33%)	0.050 (-5.20%)	0.047 (1.20%)	0.048 (-0.68%)	0.047 (1.18%)	0.042 (12.02%)	0.042 (12.39%)
Synth. $\rho = 10 - 10.00\% - 3-5$	0.050	0.055 (-9.99%)	0.050 (-0.11%)	0.050 (0.05%)	0.059 (-18.80%)	0.053 (-5.20%)	0.049 (1.83%)	0.050 (-0.63%)	0.049 (1.45%)	0.042 (15.22%)	0.042 (15.22%)
Synth. $\rho = 10 - 10.00\% - 3-10$	0.054	0.061 (-11.63%)	0.055 (-0.60%)	0.054 (0.00%)	0.068 (-24.08%)	0.057 (-1.02%)	0.053 (1.90%)	0.055 (-0.92%)	0.054 (0.79%)	0.043 (20.70%)	0.043 (21.94%)
Synth. $\rho = 100 - 1.00\% - 3-5$	0.043	0.045 (-4.65%)	0.043 (0.07%)	0.043 (0.00%)	0.048 (-11.46%)	0.045 (-3.20%)	0.042 (0.42%)	0.043 (-0.11%)	0.042 (1.01%)	0.043 (0.01%)	0.042 (1.93%)
Synth. $\rho = 100 - 1.00\% - 3-10$	0.043	0.045 (-4.65%)	0.043 (0.00%)	0.043 (0.61%)	0.049 (-13.07%)	0.045 (-5.17%)	0.043 (0.68%)	0.043 (-0.29%)	0.043 (1.17%)	0.043 (0.52%)	0.042 (1.94%)
Synth. $\rho = 100 - 5.00\% - 3-5$	0.044	0.045 (-4.05%)	0.044 (0.28%)	0.043 (0.51%)	0.052 (-17.56%)	0.044 (-1.31%)	0.044 (0.28%)	0.043 (0.33%)	0.043 (1.20%)	0.042 (3.83%)	0.042 (4.68%)
Synth. $\rho = 100 - 5.00\% - 3-10$	0.045	0.047 (-4.22%)	0.045 (0.55%)	0.044 (2.06%)	0.058 (-28.89%)	0.046 (-1.97%)	0.045 (1.53%)	0.045 (0.66%)	0.044 (3.24%)	0.042 (7.34%)	0.042 (7.34%)
Synth. $\rho = 100 - 10.00\% - 3-5$	0.047	0.049 (-4.29%)	0.047 (-0.49%)	0.046 (1.28%)	0.059 (-26.69%)	0.047 (-1.42%)	0.046 (1.88%)	0.047 (0.21%)	0.045 (3.48%)	0.043 (8.97%)	0.042 (9.25%)
Synth. $\rho = 100 - 10.00\% - 3-10$	0.050	0.050 (0.00%)	0.050 (-0.19%)	0.048 (3.13%)	0.070 (-41.47%)	0.050 (-0.19%)	0.047 (4.71%)	0.050 (-0.19%)	0.047 (5.47%)	0.043 (13.32%)	0.043 (14.11%)
Synth. $\rho = 1000 - 1.00\% - 3-5$	0.042	0.044 (-4.12%)	0.042 (-0.18%)	0.042 (0.90%)	0.047 (-11.95%)	0.043 (-1.72%)	0.042 (0.99%)	0.042 (-0.10%)	0.042 (0.53%)	0.043 (-0.55%)	0.042 (1.18%)
Synth. $\rho = 1000 - 1.00\% - 3-10$	0.043	0.045 (-5.02%)	0.043 (0.39%)	0.042 (0.60%)	0.048 (-13.59%)	0.043 (-1.07%)	0.042 (0.78%)	0.043 (-0.29%)	0.042 (0.79%)	0.043 (-0.09%)	0.042 (0.79%)
Synth. $\rho = 1000 - 5.00\% - 3-5$	0.042	0.044 (-4.76%)	0.042 (0.39%)	0.042 (0.63%)	0.051 (-21.47%)	0.042 (0.39%)	0.042 (0.39%)	0.042 (0.39%)	0.042 (0.45%)	0.042 (1.47%)	0.041 (1.81%)
Synth. $\rho = 1000 - 5.00\% - 3-10$	0.043	0.044 (-3.95%)	0.042 (0.46%)	0.042 (2.19%)	0.057 (-32.79%)	0.042 (0.46%)	0.042 (1.75%)	0.042 (0.46%)	0.042 (1.75%)	0.042 (1.98%)	0.042 (1.98%)
Synth. $\rho = 1000 - 10.00\% - 3-5$	0.044	0.045 (-4.05%)	0.044 (0.23%)	0.043 (2.12%)	0.057 (-30.84%)	0.044 (0.23%)	0.043 (2.37%)	0.044 (0.23%)	0.043 (2.37%)	0.042 (3.24%)	0.042 (3.24%)
Synth. $\rho = 1000 - 10.00\% - 3-10$	0.045	0.044 (0.96%)	0.044 (0.53%)	0.042 (4.95%)	0.068 (-54.43%)	0.044 (0.53%)	0.043 (4.28%)	0.044 (0.53%)	0.043 (4.28%)	0.043 (3.89%)	0.042 (4.57%)
Avg.	0.037	0.039 (-6.51%)	0.037 (0.03%)	0.037 (1.17%)	0.045 (-18.81%)	0.044 (-3.80%)	0.036 (1.73%)	0.037 (-0.31%)	0.036 (1.80%)	0.035 (4.84%)	0.034 (5.74%)
Weighted Avg.	0.039	0.041 (-4.36%)	0.039 (0.04%)	0.038 (1.16%)	0.048 (-21.43%)	0.040 (-1.40%)	0.038 (1.80%)	0.039 (-0.10%)	0.038 (1.88%)	0.036 (5.59%)	0.036 (6.37%)

fewer cascades, hence its MTBF for the normal regimen is very similar to the one used by Π_{Daly} .

To better understand why cascade-aware algorithms achieve little gain, we come back to Equation(1). There are two sources of waste, one due to checkpointing overhead $\frac{C}{T}$, and one due to failures. The intuition is the following: when there are cascade failures, the work wasted remains low: in average, it is approximately the MTBF of the degraded regimen. Additional (more frequent) checkpointing can reduce this waste, but can be an overkill too when there is no actual cascade.

Finally, recall that the MTBF μ_{log} of archive logs is approximately 10 hours, while it is 1 hour for synthetic logs. For the latter logs, when $\rho = 100$ means that $\mu_{degraded} = 36$ seconds, so there is little hope to gain anything except for $C = 3$ seconds. In that case, we do achieve some gain, up to 20%, but the waste was already low, around 5%, with the standard approach Π_{Daly} : overall, the absolute diminution of the waste reduces to 1%.

5.3.2 Assessing potential gain

Brute-force algorithms that search for the optimal MTBF in normal and degraded regimens allow us to quantify the potential gain that could be achieved with better cascade-aware algorithms. First, we observe that Bi-II-BEST, that supersedes both Bi-II_{Intervals} and Bi-II_{Quantiles}, is not significantly better than Π_{Daly} . Second, we make a similar observation for Bi-II_{Quantiles}_{LAZY}_{BEST}, that supersedes Bi-II_{Quantiles}_{LAZY} without performing significantly better. Third, Bi-II-ORACLE-BEST and Bi-II_{Quantiles}_{ORACLE} perform quite similarly, which is reassuring for the choice of $\mu_{non-cascade}(Q)$ as the MTBF in normal regimen. Overall, the results show that a significant gain is possible only for the latter two algorithms equipped with an omniscient oracle: when entering a cascade, Bi-II-ORACLE-BEST and Bi-II_{Quantiles}_{ORACLE} checkpoint right on time before the failures. Even so, one can see that the gains are very limited.

The only case where the gain becomes larger is when (i) the degraded MTBF is not too small in front of the log MTBF, i.e., when $\rho = 10$, and (ii) the proportion of failures that turn out to be actual cascade failures is large, i.e., when $f = 10\%$. In that case, with $|Q_{limit}| = 5\%$, the maximal potential gain is 8% with $C = 300$, 19% with $C = 30$ and 22% with $C = 3$ (the latter two values leading to small absolute gains, because the waste with Π_{Daly} is small in these cases). However, this scenario may not be realistic, because (i) cascade failures, as their

name merely indicates, are expected to strike within a few instants, hence $\mu_{degraded}$ will likely be less than one minute; and (ii) cascades are expected to be (more or less) rare events, so experiencing large values of f will probably never happen in practice.

6 Conclusion

In this report, we have revisited failure temporal independence. Recent work [3] has proposed a cascade-detection method, and we have shown that their approach was inconclusive. Then we have introduced a new approach based on pairs of consecutive IATs, and we have been able to put in evidence the presence of cascade failures. A few publicly available failure logs do contain cascades.

In a second step, we have discussed the usefulness of cascade-aware checkpointing algorithms. For this, we have used both public and synthetic logs. We used the latter to explicitly create "artificial" cascades. We have shown that current cascade-aware bi-periodic checkpointing algorithms are not really more efficient than the standard periodic checkpointing approach that considers failures to be independent. Finally, by using a brute-force search over all possible bi-periodic algorithms and considering omniscient oracles that know exactly when cascade failures will strike, we have shown that only insignificant gain should be expected from designing future cascade-aware checkpointing algorithms. The conclusion is that we can wrongly, but safely, assume failure independence!

Acknowledgements

We thank Leonardo Bautista-Gomez, Paolo Gonçalves and Arnaud Legrand for useful discussions.

References

- [1] A. Anderson and D. Semmelroth. *Statistics for Big Data For Dummies*. For Dummies, 2015.
- [2] G. Aupy, Y. Robert, F. Vivien, and D. Zaidouni. Checkpointing algorithms and fault prediction. *Journal of Parallel and Distributed Computing*, 74(2):2048–2064, 2014.
- [3] L. Bautista-Gomez, A. Gaineru, S. Perarnau, D. Tiwari, S. Gupta, C. Engelmann, F. Cappello, and M. Snir. Reducing waste in extreme scale systems through introspective analysis. In *IPDPS*, pages 212–221. IEEE, 2016.
- [4] L. Bautista-Gomez, S. Tsuboi, D. Komatitsch, F. Cappello, N. Maruyama, and S. Matsuoka. FTI: High performance fault tolerance interface for hybrid systems. In *Proc. SC'11*, 2011.
- [5] G. Bosilca, A. Bouteiller, E. Brunet, F. Cappello, J. Dongarra, A. Guermouche, T. Herault, Y. Robert, F. Vivien, and D. Zaidouni. Unified model for assessing checkpointing protocols at extreme-scale. *Concurrency and Computation: Practice and Experience*, 26(17):2772–2791, 2014.

-
- [6] K. M. Chandy and L. Lamport. Distributed snapshots: Determining global states of distributed systems. *ACM Transactions on Computer Systems*, 3(1):63–75, 1985.
 - [7] J. T. Daly. A higher order estimate of the optimum checkpoint interval for restart dumps. *FGCS*, 22(3):303–312, 2006.
 - [8] A. Gainaru, F. Cappello, M. Snir, and W. Kramer. Fault prediction under the microscope: A closer look into hpc systems. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, page 77. IEEE Computer Society Press, 2012.
 - [9] E. Gelenbe and M. Hernández. Optimum checkpoints with age dependent failures. *Acta Informatica*, 27(6):519–531, 1990.
 - [10] S. Gupta, D. Tiwari, C. Jantzi, J. Rogers, and D. Maxwell. Understanding and exploiting spatial properties of system failures on extreme-scale hpc systems. In *Dependable Systems and Networks (DSN), 2015 45th Annual IEEE/IFIP International Conference on*, pages 37–44. IEEE, 2015.
 - [11] E. Heien, D. Kondo, A. Gainaru, D. LaPine, B. Kramer, and F. Cappello. Modeling and tolerating heterogeneous failures in large parallel systems. In *High Performance Computing, Networking, Storage and Analysis (SC), 2011 International Conference for*, pages 1–11. IEEE, 2011.
 - [12] T. Héroult and Y. Robert, editors. *Fault-Tolerance Techniques for High-Performance Computing*, Computer Communications and Networks. Springer Verlag, 2015.
 - [13] S. Y. Ko, I. Hoque, B. Cho, and I. Gupta. Making cloud intermediate data fault-tolerant. In *Proc. 1st ACM Symposium on Cloud Computing*, SoCC '10. ACM, 2010.
 - [14] D. Kondo, B. Javadi, A. Iosup, and D. Epema. The failure trace archive: Enabling comparative analysis of failures in diverse distributed systems. *Cluster Computing and the Grid, IEEE International Symposium on*, pages 398–407, 2010.
 - [15] LANL. Computer failure data repository. <https://www.usenix.org/cfdr-data>, 2006.
 - [16] Y. Liu, R. Nassar, C. Leangsuksun, N. Naksinehaboon, M. Paun, and S. Scott. An optimal checkpoint/restart model for a large scale high performance computing system. In *IPDPS'08*. IEEE, 2008.
 - [17] A. Moody, G. Bronevetsky, K. Mohror, and B. R. d. Supinski. Design, modeling, and evaluation of a scalable multi-level checkpointing system. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC'10)*, pages 1–11, 2010.
 - [18] X. Ni, E. Meneses, and L. V. Kalé. Hiding checkpoint overhead in HPC applications with a semi-blocking algorithm. In *Cluster Computing (CLUSTER), 2012 IEEE International Conference on*, pages 364–372. IEEE Computer Society, 2012.
 - [19] B. Schroeder and G. A. Gibson. A large-scale study of failures in high-performance computing systems. In *Proc. of DSN*, pages 249–258, 2006.

-
- [20] K. Schroiff, P. Gemsjaeger, and C. Bolik. Cascading failover of a data management application for shared disk file systems in loosely coupled node clusters, 2006. US Patent 6,990,606.
 - [21] Y. A. Shardt. *Statistics for chemical and process engineers : a modern approach*. Springer, 2015.
 - [22] D. Tiwari, S. Gupta, and S. S. Vazhkudai. Lazy checkpointing: Exploiting temporal locality in failures to mitigate checkpointing overheads on extreme-scale systems. In *44th Int. Conf. Dependable Systems and Networks*, pages 25–36. IEEE, 2014.
 - [23] Tsubame. Failure history. <http://mon.g.gsic.titech.ac.jp/trouble-list/index.htm>, 2017.
 - [24] J. W. Young. A first order approximation to the optimum checkpoint interval. *Comm. of the ACM*, 17(9):530–531, 1974.
 - [25] G. Zheng, L. Shi, and L. V. Kale. FTC-Charm++: an in-memory checkpoint-based fault tolerant runtime for Charm++ and MPI. In *Cluster Computing, 2004 IEEE International Conference on*, pages 93–103. IEEE Computer Society, 2004.

Log	Id	Number of faults	MTBF in hours	Approach: Intervals			Approach: Quantiles						
				MTBF in hours	Degraded intervals	Faults in cascades	MTBF in hours	Faults in cascades	Common faults	Number of cascades	in Cascades		
											MTBF	Avg. length	Max length
LANL 2	5351	14.14	36.45	25.3%	71.1%	14.43	3.6%	3.6%	82	0.035	2.3	8	
LANL 16	2262	21.85	56.19	25.2%	70.9%	22.43	4.8%	4.8%	50	0.046	2.2	5	
LANL 18	3900	7.52	17.88	26.0%	68.9%	7.68	4.3%	4.3%	81	0.050	2.0	3	
LANL 19	3222	7.87	17.06	26.4%	66.0%	8.05	4.3%	4.3%	69	0.074	2.0	2	
LANL 20	2389	13.66	41.46	21.3%	74.1%	13.95	4.0%	4.0%	45	0.040	2.1	4	
Tsubame	884	14.78	36.50	23.9%	69.2%	15.14	4.5%	4.5%	19	0.029	2.1	4	
Synth. $\rho = 10 - 1.00\% - 3-5$	3136	0.98	1.99	25.1%	63.1%	1.00	4.0%	4.0%	62	0.007	2.0	3	
Synth. $\rho = 10 - 1.00\% - 3-10$	3182	0.95	1.98	25.0%	64.1%	0.97	4.0%	4.0%	64	0.006	2.0	2	
Synth. $\rho = 10 - 5.00\% - 3-5$	3591	0.85	1.87	22.4%	64.9%	0.86	3.9%	3.9%	68	0.003	2.1	3	
Synth. $\rho = 10 - 5.00\% - 3-10$	4075	0.75	1.77	20.5%	66.2%	0.77	3.9%	3.9%	77	0.002	2.1	3	
Synth. $\rho = 10 - 10.00\% - 3-5$	4133	0.72	1.81	22.2%	69.0%	0.73	3.9%	3.9%	78	0.003	2.1	3	
Synth. $\rho = 10 - 10.00\% - 3-10$	5067	0.59	1.65	19.5%	71.2%	0.60	3.9%	3.9%	96	0.002	2.1	3	
Synth. $\rho = 100 - 1.00\% - 3-5$	3136	0.98	1.99	24.9%	62.9%	1.00	3.7%	3.6%	52	0.002	2.2	4	
Synth. $\rho = 100 - 1.00\% - 3-10$	3182	0.95	1.98	24.6%	63.9%	0.97	3.8%	3.8%	56	0.002	2.1	5	
Synth. $\rho = 100 - 5.00\% - 3-5$	3591	0.85	1.89	21.6%	64.8%	0.86	3.8%	3.8%	66	0.001	2.1	4	
Synth. $\rho = 100 - 5.00\% - 3-10$	4075	0.75	1.79	18.3%	65.7%	0.77	4.0%	4.0%	81	0.000	2.0	3	
Synth. $\rho = 100 - 10.00\% - 3-5$	4133	0.72	1.84	20.2%	68.7%	0.73	3.9%	3.9%	77	0.000	2.1	3	
Synth. $\rho = 100 - 10.00\% - 3-10$	5067	0.59	1.69	15.8%	70.5%	0.60	3.9%	3.9%	98	0.000	2.0	3	
Synth. $\rho = 1000 - 1.00\% - 3-5$	3136	0.98	1.99	24.9%	63.0%	1.00	3.3%	3.3%	41	0.000	2.5	6	
Synth. $\rho = 1000 - 1.00\% - 3-10$	3182	0.95	1.98	24.6%	63.9%	0.97	3.5%	3.5%	47	0.000	2.4	6	
Synth. $\rho = 1000 - 5.00\% - 3-5$	3591	0.85	1.89	21.3%	64.8%	0.86	3.8%	3.8%	66	0.000	2.1	4	
Synth. $\rho = 1000 - 5.00\% - 3-10$	4075	0.75	1.79	18.1%	65.6%	0.77	4.0%	4.0%	81	0.000	2.0	3	
Synth. $\rho = 1000 - 10.00\% - 3-5$	4133	0.72	1.83	19.9%	68.6%	0.73	3.8%	3.8%	75	0.000	2.1	4	
Synth. $\rho = 1000 - 10.00\% - 3-10$	5067	0.59	1.70	15.3%	70.4%	0.60	3.9%	3.9%	98	0.000	2.0	3	

Table 13: Statistics when $|Q_{limit}| = 2\%$.

A Additional results

In this section, we report results for $|Q_{limit}| \in \{2\%, 1\%, 0.5\%, 0.2\%, 0.1\%\}$.

A.1 Log statistics

We report log statistics in Tables 13 to 17.

A.2 Waste values

In Tables 18 to 32, we report the improvement or degradation of the waste with respect to the reference periodic checkpointing algorithm Π_{Daly} . The color code is the same as in Section 5.2.

Log	Id	Number of faults	MTBF in hours	Approach: Intervals			Approach: Quantiles						
				MTBF in hours	Degraded intervals	Faults in cascades	MTBF in hours	Faults in cascades	Common faults	Number of cascades	MTBF	in Cascades	
												Avg. length	Max length
LANL 2	5351	14.14	36.45	25.3%	71.1%	14.33	2.5%	2.5%	61	0.024	2.2	5	
LANL 16	2262	21.85	56.19	25.2%	70.9%	22.15	2.5%	2.5%	27	0.023	2.1	3	
LANL 18	3900	7.52	17.88	26.0%	68.9%	7.61	2.4%	2.4%	46	0.024	2.0	2	
LANL 19	3222	7.87	17.06	26.4%	66.0%	7.98	2.8%	2.8%	45	0.052	2.0	2	
LANL 20	2389	13.66	41.46	21.3%	74.1%	13.82	2.3%	2.3%	27	0.024	2.0	3	
Tsubame	884	14.78	36.50	23.9%	69.2%	15.00	2.9%	2.9%	13	0.017	2.0	2	
Synth. $\rho = 10 - 1.00\% - 3-5$	3136	0.98	1.99	25.1%	63.1%	0.99	2.0%	2.0%	32	0.003	2.0	2	
Synth. $\rho = 10 - 1.00\% - 3-10$	3182	0.95	1.98	25.0%	64.1%	0.96	2.0%	2.0%	32	0.003	2.0	2	
Synth. $\rho = 10 - 5.00\% - 3-5$	3591	0.85	1.87	22.4%	64.9%	0.86	2.0%	2.0%	36	0.002	2.0	2	
Synth. $\rho = 10 - 5.00\% - 3-10$	4075	0.75	1.77	20.5%	66.2%	0.76	1.9%	1.9%	38	0.001	2.1	3	
Synth. $\rho = 10 - 10.00\% - 3-5$	4133	0.72	1.81	22.2%	69.0%	0.73	2.0%	2.0%	41	0.001	2.0	3	
Synth. $\rho = 10 - 10.00\% - 3-10$	5067	0.59	1.65	19.5%	71.2%	0.60	2.0%	2.0%	49	0.001	2.0	3	
Synth. $\rho = 100 - 1.00\% - 3-5$	3136	0.98	1.99	24.9%	62.9%	0.99	1.9%	1.9%	28	0.001	2.1	3	
Synth. $\rho = 100 - 1.00\% - 3-10$	3182	0.95	1.98	24.6%	63.9%	0.96	1.9%	1.9%	29	0.001	2.1	3	
Synth. $\rho = 100 - 5.00\% - 3-5$	3591	0.85	1.89	21.6%	64.8%	0.86	2.0%	2.0%	36	0.000	2.0	2	
Synth. $\rho = 100 - 5.00\% - 3-10$	4075	0.75	1.79	18.3%	65.7%	0.76	2.0%	2.0%	40	0.000	2.0	3	
Synth. $\rho = 100 - 10.00\% - 3-5$	4133	0.72	1.84	20.2%	68.7%	0.73	2.0%	2.0%	41	0.000	2.0	3	
Synth. $\rho = 100 - 10.00\% - 3-10$	5067	0.59	1.69	15.8%	70.5%	0.60	2.0%	2.0%	49	0.000	2.0	3	
Synth. $\rho = 1000 - 1.00\% - 3-5$	3136	0.98	1.99	24.9%	63.0%	0.99	1.9%	1.9%	28	0.000	2.1	3	
Synth. $\rho = 1000 - 1.00\% - 3-10$	3182	0.95	1.98	24.6%	63.9%	0.96	1.9%	1.9%	29	0.000	2.1	3	
Synth. $\rho = 1000 - 5.00\% - 3-5$	3591	0.85	1.89	21.3%	64.8%	0.86	2.0%	2.0%	36	0.000	2.0	2	
Synth. $\rho = 1000 - 5.00\% - 3-10$	4075	0.75	1.79	18.1%	65.6%	0.76	2.0%	2.0%	40	0.000	2.0	3	
Synth. $\rho = 1000 - 10.00\% - 3-5$	4133	0.72	1.83	19.9%	68.6%	0.73	2.0%	2.0%	41	0.000	2.0	3	
Synth. $\rho = 1000 - 10.00\% - 3-10$	5067	0.59	1.70	15.3%	70.4%	0.60	2.0%	2.0%	49	0.000	2.0	3	

Table 14: Statistics when $|Q_{limit}| = 1\%$.

Log	Id	Number of faults	MTBF in hours	Approach: Intervals			Approach: Quantiles						
				MTBF in hours	Degraded intervals	Faults in cascades	MTBF in hours	Faults in cascades	Common faults	Number of cascades	MTBF	in Cascades	
												Avg. length	Max length
LANL 2	5351	14.14	36.45	25.3%	71.1%	14.24	1.4%	1.4%	35	0.017	2.1	5	
LANL 16	2262	21.85	56.19	25.2%	70.9%	22.03	1.5%	1.5%	17	0.017	2.1	3	
LANL 18	3900	7.52	17.88	26.0%	68.9%	7.57	1.3%	1.3%	25	0.017	2.0	2	
LANL 19	3222	7.87	17.06	26.4%	66.0%	7.91	1.1%	1.1%	17	0.019	2.0	2	
LANL 20	2389	13.66	41.46	21.3%	74.1%	13.74	1.3%	1.3%	15	0.017	2.0	2	
Tsubame	884	14.78	36.50	23.9%	69.2%	15.00	2.9%	2.9%	13	0.017	2.0	2	
Synth. $\rho = 10 - 1.00\% - 3-5$	3136	0.98	1.99	25.1%	63.1%	0.99	1.0%	1.0%	16	0.002	2.0	2	
Synth. $\rho = 10 - 1.00\% - 3-10$	3182	0.95	1.98	25.0%	64.1%	0.95	1.0%	1.0%	16	0.002	2.0	2	
Synth. $\rho = 10 - 5.00\% - 3-5$	3591	0.85	1.87	22.4%	64.9%	0.85	1.0%	1.0%	18	0.001	2.0	2	
Synth. $\rho = 10 - 5.00\% - 3-10$	4075	0.75	1.77	20.5%	66.2%	0.76	1.0%	1.0%	21	0.000	2.0	2	
Synth. $\rho = 10 - 10.00\% - 3-5$	4133	0.72	1.81	22.2%	69.0%	0.72	1.0%	1.0%	21	0.001	2.0	2	
Synth. $\rho = 10 - 10.00\% - 3-10$	5067	0.59	1.65	19.5%	71.2%	0.60	1.0%	1.0%	26	0.000	2.0	2	
Synth. $\rho = 100 - 1.00\% - 3-5$	3136	0.98	1.99	24.9%	62.9%	0.99	1.0%	1.0%	15	0.001	2.1	3	
Synth. $\rho = 100 - 1.00\% - 3-10$	3182	0.95	1.98	24.6%	63.9%	0.95	1.0%	1.0%	16	0.000	2.0	2	
Synth. $\rho = 100 - 5.00\% - 3-5$	3591	0.85	1.89	21.6%	64.8%	0.85	1.0%	1.0%	18	0.000	2.0	2	
Synth. $\rho = 100 - 5.00\% - 3-10$	4075	0.75	1.79	18.3%	65.7%	0.76	1.0%	1.0%	21	0.000	2.0	2	
Synth. $\rho = 100 - 10.00\% - 3-5$	4133	0.72	1.84	20.2%	68.7%	0.72	1.0%	1.0%	21	0.000	2.0	2	
Synth. $\rho = 100 - 10.00\% - 3-10$	5067	0.59	1.69	15.8%	70.5%	0.60	1.0%	1.0%	25	0.000	2.0	3	
Synth. $\rho = 1000 - 1.00\% - 3-5$	3136	0.98	1.99	24.9%	63.0%	0.99	1.0%	1.0%	15	0.000	2.1	3	
Synth. $\rho = 1000 - 1.00\% - 3-10$	3182	0.95	1.98	24.6%	63.9%	0.95	1.0%	1.0%	16	0.000	2.0	2	
Synth. $\rho = 1000 - 5.00\% - 3-5$	3591	0.85	1.89	21.3%	64.8%	0.85	1.0%	1.0%	18	0.000	2.0	2	
Synth. $\rho = 1000 - 5.00\% - 3-10$	4075	0.75	1.79	18.1%	65.6%	0.76	1.0%	1.0%	21	0.000	2.0	2	
Synth. $\rho = 1000 - 10.00\% - 3-5$	4133	0.72	1.83	19.9%	68.6%	0.72	1.0%	1.0%	21	0.000	2.0	2	
Synth. $\rho = 1000 - 10.00\% - 3-10$	5067	0.59	1.70	15.3%	70.4%	0.60	1.0%	1.0%	25	0.000	2.0	3	

Table 15: Statistics when $|Q_{limit}| = 0.5\%$.

Log	Id	Number of faults	MTBF in hours	Approach: Intervals			Approach: Quantiles						
				MTBF in hours	Degraded intervals	Faults in cascades	MTBF in hours	Faults in cascades	Common faults	Number of cascades	MTBF	in Cascades	
												Avg. length	Max length
LANL 2	5351	14.14	36.45	25.3%	71.1%	14.24	1.4%	1.4%	35	0.017	2.1	5	
LANL 16	2262	21.85	56.19	25.2%	70.9%	22.03	1.5%	1.5%	17	0.017	2.1	3	
LANL 18	3900	7.52	17.88	26.0%	68.9%	7.57	1.3%	1.3%	25	0.017	2.0	2	
LANL 19	3222	7.87	17.06	26.4%	66.0%	7.91	0.9%	0.9%	15	0.017	2.0	2	
LANL 20	2389	13.66	41.46	21.3%	74.1%	13.74	1.3%	1.3%	15	0.017	2.0	2	
Tsubame	884	14.78	36.50	23.9%	69.2%	15.00	2.9%	2.9%	13	0.017	2.0	2	
Synth. $\rho = 10 - 1.00\% - 3-5$	3136	0.98	1.99	25.1%	63.1%	0.98	0.4%	0.4%	7	0.001	2.0	2	
Synth. $\rho = 10 - 1.00\% - 3-10$	3182	0.95	1.98	25.0%	64.1%	0.95	0.4%	0.4%	7	0.001	2.0	2	
Synth. $\rho = 10 - 5.00\% - 3-5$	3591	0.85	1.87	22.4%	64.9%	0.85	0.4%	0.4%	8	0.000	2.0	2	
Synth. $\rho = 10 - 5.00\% - 3-10$	4075	0.75	1.77	20.5%	66.2%	0.75	0.4%	0.4%	9	0.000	2.0	2	
Synth. $\rho = 10 - 10.00\% - 3-5$	4133	0.72	1.81	22.2%	69.0%	0.72	0.4%	0.4%	9	0.000	2.0	2	
Synth. $\rho = 10 - 10.00\% - 3-10$	5067	0.59	1.65	19.5%	71.2%	0.59	0.4%	0.4%	11	0.000	2.0	2	
Synth. $\rho = 100 - 1.00\% - 3-5$	3136	0.98	1.99	24.9%	62.9%	0.98	0.4%	0.4%	7	0.000	2.0	2	
Synth. $\rho = 100 - 1.00\% - 3-10$	3182	0.95	1.98	24.6%	63.9%	0.95	0.4%	0.4%	7	0.000	2.0	2	
Synth. $\rho = 100 - 5.00\% - 3-5$	3591	0.85	1.89	21.6%	64.8%	0.85	0.4%	0.4%	8	0.000	2.0	2	
Synth. $\rho = 100 - 5.00\% - 3-10$	4075	0.75	1.79	18.3%	65.7%	0.75	0.4%	0.4%	9	0.000	2.0	2	
Synth. $\rho = 100 - 10.00\% - 3-5$	4133	0.72	1.84	20.2%	68.7%	0.72	0.4%	0.4%	9	0.000	2.0	2	
Synth. $\rho = 100 - 10.00\% - 3-10$	5067	0.59	1.69	15.8%	70.5%	0.59	0.4%	0.4%	11	0.000	2.0	2	
Synth. $\rho = 1000 - 1.00\% - 3-5$	3136	0.98	1.99	24.9%	63.0%	0.98	0.4%	0.4%	7	0.000	2.0	2	
Synth. $\rho = 1000 - 1.00\% - 3-10$	3182	0.95	1.98	24.6%	63.9%	0.95	0.4%	0.4%	7	0.000	2.0	2	
Synth. $\rho = 1000 - 5.00\% - 3-5$	3591	0.85	1.89	21.3%	64.8%	0.85	0.4%	0.4%	8	0.000	2.0	2	
Synth. $\rho = 1000 - 5.00\% - 3-10$	4075	0.75	1.79	18.1%	65.6%	0.75	0.4%	0.4%	9	0.000	2.0	2	
Synth. $\rho = 1000 - 10.00\% - 3-5$	4133	0.72	1.83	19.9%	68.6%	0.72	0.4%	0.4%	9	0.000	2.0	2	
Synth. $\rho = 1000 - 10.00\% - 3-10$	5067	0.59	1.70	15.3%	70.4%	0.59	0.4%	0.4%	11	0.000	2.0	2	

Table 16: Statistics when $|Q_{limit}| = 0.2\%$.

Log	Id	Number of faults	MTBF in hours	Approach: Intervals			Approach: Quantiles						
				MTBF in hours	Degraded intervals	Faults in cascades	MTBF in hours	Faults in cascades	Common faults	Number of cascades	MTBF	in Cascades	
												Avg. length	Max length
LANL 2	5351	14.14	36.45	25.3%	71.1%	14.24	1.4%	1.4%	35	0.017	2.1	5	
LANL 16	2262	21.85	56.19	25.2%	70.9%	22.03	1.5%	1.5%	17	0.017	2.1	3	
LANL 18	3900	7.52	17.88	26.0%	68.9%	7.57	1.3%	1.3%	25	0.017	2.0	2	
LANL 19	3222	7.87	17.06	26.4%	66.0%	7.91	0.9%	0.9%	15	0.017	2.0	2	
LANL 20	2389	13.66	41.46	21.3%	74.1%	13.74	1.3%	1.3%	15	0.017	2.0	2	
Tsubame	884	14.78	36.50	23.9%	69.2%	15.00	2.9%	2.9%	13	0.017	2.0	2	
Synth. $\rho = 10 - 1.00\% - 3-5$	3136	0.98	1.99	25.1%	63.1%	0.98	0.3%	0.3%	4	0.001	2.0	2	
Synth. $\rho = 10 - 1.00\% - 3-10$	3182	0.95	1.98	25.0%	64.1%	0.95	0.3%	0.3%	4	0.001	2.0	2	
Synth. $\rho = 10 - 5.00\% - 3-5$	3591	0.85	1.87	22.4%	64.9%	0.85	0.2%	0.2%	4	0.000	2.0	2	
Synth. $\rho = 10 - 5.00\% - 3-10$	4075	0.75	1.77	20.5%	66.2%	0.75	0.2%	0.2%	5	0.000	2.0	2	
Synth. $\rho = 10 - 10.00\% - 3-5$	4133	0.72	1.81	22.2%	69.0%	0.72	0.2%	0.2%	5	0.000	2.0	2	
Synth. $\rho = 10 - 10.00\% - 3-10$	5067	0.59	1.65	19.5%	71.2%	0.59	0.2%	0.2%	6	0.000	2.0	2	
Synth. $\rho = 100 - 1.00\% - 3-5$	3136	0.98	1.99	24.9%	62.9%	0.98	0.3%	0.3%	4	0.000	2.0	2	
Synth. $\rho = 100 - 1.00\% - 3-10$	3182	0.95	1.98	24.6%	63.9%	0.95	0.3%	0.3%	4	0.000	2.0	2	
Synth. $\rho = 100 - 5.00\% - 3-5$	3591	0.85	1.89	21.6%	64.8%	0.85	0.2%	0.2%	4	0.000	2.0	2	
Synth. $\rho = 100 - 5.00\% - 3-10$	4075	0.75	1.79	18.3%	65.7%	0.75	0.2%	0.2%	5	0.000	2.0	2	
Synth. $\rho = 100 - 10.00\% - 3-5$	4133	0.72	1.84	20.2%	68.7%	0.72	0.2%	0.2%	5	0.000	2.0	2	
Synth. $\rho = 100 - 10.00\% - 3-10$	5067	0.59	1.69	15.8%	70.5%	0.59	0.2%	0.2%	6	0.000	2.0	2	
Synth. $\rho = 1000 - 1.00\% - 3-5$	3136	0.98	1.99	24.9%	63.0%	0.98	0.3%	0.3%	4	0.000	2.0	2	
Synth. $\rho = 1000 - 1.00\% - 3-10$	3182	0.95	1.98	24.6%	63.9%	0.95	0.3%	0.3%	4	0.000	2.0	2	
Synth. $\rho = 1000 - 5.00\% - 3-5$	3591	0.85	1.89	21.3%	64.8%	0.85	0.2%	0.2%	4	0.000	2.0	2	
Synth. $\rho = 1000 - 5.00\% - 3-10$	4075	0.75	1.79	18.1%	65.6%	0.75	0.2%	0.2%	5	0.000	2.0	2	
Synth. $\rho = 1000 - 10.00\% - 3-5$	4133	0.72	1.83	19.9%	68.6%	0.72	0.2%	0.2%	5	0.000	2.0	2	
Synth. $\rho = 1000 - 10.00\% - 3-10$	5067	0.59	1.70	15.3%	70.4%	0.59	0.2%	0.2%	6	0.000	2.0	2	

Table 17: Statistics when $|Q_{limit}| = 0.1\%$.



**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