



Which Verification for Soft Error Detection?

Leonardo Bautista-Gomez, Anne Benoit, Aurélien Cavelan, Saurabh
K. Raina, Yves Robert, and Hongyang Sun

**RESEARCH
REPORT**

N° 8741

June 2015

Project-Team ROMA



Which Verification for Soft Error Detection?

Leonardo Bautista-Gomez^{*}, Anne Benoit^{†‡}, Aurélien Cavelan^{†‡},
Saurabh K. Raina[§], Yves Robert^{†‡¶}, and Hongyang Sun^{†‡}

Project-Team ROMA

Research Report n° 8741 — June 2015 — 17 pages

Abstract: Many methods are available to detect silent errors in high-performance computing (HPC) applications. Each comes with a given cost and recall (fraction of all errors that are actually detected). The main contribution of this paper is to show which detector(s) to use, and to characterize the optimal computational pattern for the application: how many detectors of each type to use, together with the length of the work segment that precedes each of them. We conduct a comprehensive complexity analysis of this optimization problem, showing NP-completeness and designing an FPTAS (Fully Polynomial-Time Approximation Scheme). On the practical side, we provide a greedy algorithm whose performance is shown to be close to the optimal for a realistic set of evaluation scenarios.

Key-words: fault tolerance, high performance computing, silent data corruption, partial verification, supercomputer, exascale.

^{*} Argonne National Laboratory, USA

[†] École Normale Supérieure de Lyon

[‡] INRIA, France

[§] Jaypee Institute of Information Technology, India

[¶] University of Tennessee Knoxville, USA

**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

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

Quelle vérification pour la détection d'erreurs silencieuses ?

Résumé : De nombreuses méthodes sont disponibles pour détecter les erreurs silencieuses dans les applications de Calcul Haute Performance (HPC). Chaque méthode a un coût et un rappel (fraction de toutes les erreurs qui sont effectivement détectées). La principale contribution de cet article est de montrer quel(s) détecteur(s) utiliser, et de caractériser le motif de calcul optimale pour une application: combien de détecteurs de chaque type utiliser, ainsi que la longueur du segment de travail qui les précède. Nous menons une analyse de complexité exhaustive de ce problème d'optimisation, montrant sa NP-complétude et la conception d'une FPTAS (Fully Polynomial-Time Approximation Scheme). Sur le plan pratique, nous fournissons un algorithme glouton dont la performance est montrée comme étant proche de l'optimal pour un ensemble réaliste de scénarios d'évaluation.

Mots-clés : tolérance aux pannes, Calcul Haute Performance, erreur silencieuse, corruption de donnée silencieuse, vérification partielle, checkpoint, supercalculateur, exascale.

1 Introduction

Failures in high-performance computing (HPC) systems have become a major issue as the number of components proliferate. Indeed, future exascale platforms are expected to be composed of hundreds of thousands of computing nodes [17]. Even if each individual node provides an optimistic mean time between failures (MTBF) of, say 100 years, the whole platform will experience a failure around every few hours on average, which is shorter than the execution time of most HPC applications. Thus, effective resilient protocols will be essential to achieve efficiency.

The de-facto general-purpose error recovery technique in HPC is checkpoint and rollback recovery [12, 20]. Such protocols employ checkpoints to periodically save the state of a parallel application, so that when an error strikes some process, the application can be restored to one of its former states. However, checkpoint/restart assumes instantaneous error detection, and therefore applies to fail-stop errors. Silent errors, a.k.a. silent data corruptions (SDC), constitute another source of failures in HPC, whose threat can no longer be ignored [30, 34, 28]. There are several causes of silent errors such as cosmic radiation, packaging pollution, among others. 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.

In order to avoid corrupted checkpoints, an effective approach consists in employing some verification mechanism and combining it with checkpointing [13, 31, 1]. The simplest protocol with this approach would be to execute a verification procedure before taking each checkpoint. If the verification succeeds, then one can safely store the checkpoint. Otherwise, it means that an error has struck since the last checkpoint, which was duly verified, and one can safely recover from that checkpoint to resume the execution of the application. Of course, more sophisticated protocols can be designed, by coupling multiple verifications with one checkpoint, or even interleaving multiple checkpoints and verifications [1, 6]. The optimal pattern (i.e., number of verifications per checkpoint) in these protocols would be determined by the cost of executing a verification.

In practice, not all verification mechanisms are 100% accurate and at the same time admit fast implementations. In fact, guaranteeing accurate and efficient detection of silent errors for scientific applications is one of the hardest challenges in extreme-scale computing [2]. Indeed, thorough error detection is usually very costly, and often involves expensive techniques, such as replication [21] or even triplication [27]. For many parallel applications, alternative techniques exist that are capable of detecting some but not all errors. We call these techniques *partial verifications*, while a *guaranteed verification* is capable of detecting all errors. One example is the lightweight SDC




Figure 1: A periodic pattern (highlighted in red) with two partial verifications and a verified checkpoint.

detector based on data dynamic monitoring [2], designed to recognize anomalies in HPC datasets based on physical laws and spatial interpolation. Similar fault filters have also been designed to detect silent errors based on time series predictions [8]. Although not completely accurate, these partial verification techniques nevertheless cover a substantial amount of silent data corruptions, and more importantly, they incur low overhead. These properties make them attractive candidates for designing more efficient resilient protocols.

Since checkpointing is often expensive in terms of both time and space required, to avoid saving corrupted data, we only keep *verified checkpoints* by placing a guaranteed verification right before each checkpoint. Such a combination ensures that the checkpoint contains valid data and can be written onto stable storage. The execution of the application is partitioned into *periodic patterns*, i.e. computational chunks that repeat over time, and that are delimited by verified checkpoints, possibly with a sequence of partial verifications in between. Figure 1 shows a periodic pattern with two partial verifications followed by a verified checkpoint.

The error detection accuracy of a partial verification can be characterized by its *recall* r , which is the ratio between the number of detected errors and the total number of errors that occurred during a computation. For example, a basic spatial based SDC detector has been shown to have a recall around 0.5 measured on synthetic benchmarks [2], which means that it is capable of detecting half of the errors. Note that a guaranteed verification can be considered as a special type of partial verification with a recall $r = 1$. Each partial verification also has an associated *cost* V , which is typically much smaller than the cost V^* of a guaranteed verification.

An application can use several types of detectors with different overheads. For instance, to detect silent errors in HPC datasets, one has the option of using either the detector based on time series prediction [8], or the detector using spatial multivariate interpolation [2]. The first one needs more data to make a prediction, hence comes at a higher cost. However, its recall is also better. In the example of Figure 1, the second verification may use a detector whose cost is lower than that of the first one ($V_2 < V_1$), but is expected to have a lower recall as well ($r_2 < r_1$).

In this paper, we assume to have several detector types, whose costs and recalls may differ. At the end of each segment inside the pattern, any detector can be used. The only constraint is to enforce a guaranteed verification after the last segment. Given the values of C (cost to checkpoint), V^* (cost of a guaranteed verification), the cost $V^{(j)}$

and recall $r^{(j)}$ of detector type $D^{(j)}$, the main question is which detector(s) to use? The objective is to find the optimal pattern that minimizes the expected execution time of the application. Intuitively, including more partial verifications in a pattern allows us to detect more errors, and earlier in the execution, thereby reducing the waste due to re-execution; but that comes at the price of additional overhead in an error-free execution. Therefore, an optimal strategy must seek a good tradeoff between error-induced waste and error-free overhead. The problem is intrinsically combinatorial, because there are many parameters to choose: the length of the pattern, the number of partial verifications, and the type and location of each partial verification within the pattern. Of course, the length of an optimal pattern will also depend on the platform MTBF μ .

Only very specific instances of the problem have received a solution yet. For instance, when there is a single segment in the pattern without intermediate verification, the only thing to determine is the size of the segment. In the classical protocol for fail-stop errors (where verification is not needed), the optimal checkpointing period is known to be $\sqrt{2\mu C}$ (where C is the checkpoint time), as given by Young [33] and Daly [16]. A similar result is known for silent errors when using only verified checkpoints [6, 5]: in that case, the optimal period is $\sqrt{\mu(C + V^*)}$. These formulas provide first-order approximations to the length of the optimal pattern in the corresponding scenario, and are valid only if resilience parameters satisfy $C, V^* \ll \mu$. To the best of your knowledge, the only analysis of patterns that include partial verifications is the recent work [11], which deals with patterns that may include one or several detector(s), but all of the same type. While most applications accept several detector types, there has been no attempt to determine which and how many of these detectors should be used. This paper is the first to investigate the use of different types of partial verifications, i.e., different detectors.

We prove that this optimization problem is NP-complete. We show that a detector is most useful when it offers verifications with a high *accuracy-to-cost ratio*, defined as $\phi^{(j)} = \frac{r^{(j)}}{2-r^{(j)}} / \frac{V^{(j)}}{V^*+C}$. We then propose a greedy algorithm and a fully polynomial-time approximation scheme (FPTAS) to solve the problem. Simulation results, based on a wide range of parameters from realistic detectors, corroborate the theoretical study by showing that the detector with the best accuracy-to-cost ratio should be favored. In some particular cases with close accuracy-to-cost ratios, an optimal pattern may use two different detectors. Finally, the greedy algorithm has been shown to work quite well in practice.

The rest of this paper is organized as follows. Section 2 surveys related work. Section 3 introduces the model, notations and assumptions. Section 4 presents key properties of optimal execution patterns. Section 5 provides a comprehensive complexity analysis. While the optimization problem is shown to be NP-complete, a simple greedy algorithm is presented, and a fully polynomial-time approximation

scheme is described. Simulation results are presented in Section 6. Finally, Section 7 provides concluding remarks and hints for future directions.

2 Related Work

Considerable efforts have been directed at detection techniques to reveal silent errors. Hardware mechanisms, such as ECC memory, can detect and even correct a fraction of errors. Unfortunately, future extreme scale systems are expected to observe an important increase in soft errors because of power constraints at increased system size. Most traditional resilient approaches maintain a single checkpoint. If the checkpoint file contains corrupted data, the application faces an irrecoverable failure and must restart from scratch. This is because error detection latency is ignored in traditional rollback and recovery schemes, which assume instantaneous error detection (therefore mainly targeting fail-stop errors) and are unable to accommodate SDC. This section describes some related work on detecting and handling silent errors.

2.1 Checkpoint versioning

One approach to deal with silent errors is by maintaining several checkpoints in memory [26]. This multiple-checkpoint approach, however, has three major drawbacks. First, it is very demanding in terms of stable storage: each checkpoint typically represents a copy of a large portion of the memory footprint of the application, which may well correspond to tens or even hundreds of terabytes. Second, the application cannot be recovered from fatal failures: suppose we keep k checkpoints in memory, and a silent error has struck before all of them. Then, all live checkpoints are corrupted, and one would have to re-execute the entire application from scratch. Third, even without memory constraints, we have to determine which checkpoint is the last valid one, which is needed to safely recover the application from. However, due to the detection latency, we do not know when the silent error has occurred, hence we cannot identify the last valid checkpoint.

2.2 Process replication

There are few methods that can guarantee a perfect detection recall. Process replication is one of them. The simplest technique is triple modular redundancy and voting [27]. Elliot et al. [19] propose combining partial redundancy and checkpointing, and confirm the benefit of dual and triple redundancy. Fiala et al. [21] apply process replication (each process is equipped with a replica, and messages are quadruplicated) in the RedMPI library for high-performance scientific applications. Ni et al. [29] use checkpointing and replication to detect and enable fast recovery of applications from both silent errors and hard errors. However, full process replication is too expensive to be used in extreme scale HPC systems and is usually avoided for this reason.

2.3 Application-specific techniques

Application-specific information can be very useful to enable ad-hoc solutions, which dramatically decrease the cost of detection. Algorithm-based fault tolerance (ABFT) [24, 9, 32] 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 inside the linear algebra kernels and it has to 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 [31] investigate self-stabilizing corrections after error detection in the conjugate gradient method. Heroux and Hoemmen [23] design a fault-tolerant GMRES capable of converging despite silent errors, and Bronevetsky and de Supinski [10] provide a comparative study of detection costs for iterative methods.

2.4 Analytics-based corruption detection

Recently, several SDC 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 [2, 3, 4]. Such techniques have the benefit of offering 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). 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 SDCs in the temperature data of the Orbital Thermal Imaging Spectrometer (OTIS) [14]. Most of the research work done in this domain focuses on how to increase the recall while keeping low overhead but there is no theoretical study trying to find the optimal protocol applications should use when multiple verification techniques are offered by the runtime.

2.5 Optimal strategies with guaranteed verifications

Theoretically, various protocols that couple verification and checkpointing have been studied. Aupy et al. [1] propose and analyze two simple patterns: one with k checkpoints and one verification, and the other with k verifications and one checkpoint. The latter pattern, which needs to maintain only one checkpoint, is also analyzed in [5] to accommodate both fail-stop and silent errors. Benoit et al. [6] extend the analysis of [1] by including p checkpoints and q verifications that are interleaved to form arbitrary patterns. All of these results assume the use of guaranteed verifications only.

As already mentioned, the only analysis of patterns that include partial verifications is the recent work of [11]. However, [11] restricts to a single type of partial verification. In this paper, we provide the first theoretical analysis that includes partial verifications of different types.

3 Model

We consider divisible-load applications, where checkpoints and verifications can be inserted anywhere in the execution of the application. The occurrence of silent errors follows a Poisson process with arrival rate $\lambda = 1/\mu$, where μ denotes the MTBF of the platform.

We enforce resilience through the use of a *pattern* that repeats periodically throughout the execution, as discussed in Section 1. When an error is detected inside a pattern, either by a partial verification or by the guaranteed verification, we roll back to the beginning of the pattern and recover from the last checkpoint (taken at the end of the execution of the previous pattern, or initial data for the first pattern). Since the last verification of the pattern is guaranteed, we need to maintain only one checkpoint at any time, and it is always valid. The objective is to find a pattern that minimizes the expected execution time of the application.

Let C denote the cost of checkpointing, R the cost of recovery and V^* the cost of guaranteed verification. Furthermore, there are k types of detectors available, and the detector type $D^{(j)}$, where $1 \leq j \leq k$, is characterized by its cost $V^{(j)}$ and recall $r^{(j)}$. For convenience, we also define $g^{(j)} = 1 - r^{(j)}$ (proportion of undetected errors) and let D^* be the guaranteed detector with cost V^* and recall $r^* = 1$.

A pattern $\text{PATTERN}(W, n, \alpha, \mathbf{D})$ is defined by its total length W , the number n of segments in the pattern, a vector $\alpha = [\alpha_1, \alpha_2, \dots, \alpha_n]^T$ containing the proportion of segment sizes, and a vector $\mathbf{D} = [D_1, D_2, \dots, D_{n-1}, D_n]^T$ containing the detectors used at the end of each segment. We also define the vector of segment sizes $\mathbf{w} = [w_1, w_2, \dots, w_n]^T$. Formally, for each segment i , where $1 \leq i \leq n$, w_i is the size of the segment, $\alpha_i = \frac{w_i}{W}$ is the proportion of the segment size in the whole pattern, and D_i is the detector used at the end of the segment. We have $\sum_{i=1}^n \alpha_i = 1$, and $\sum_{i=1}^n w_i = W$. If $i < n$, D_i has cost V_i and recall r_i (we have $D_i = D^{(j)}$ for some type j , $1 \leq j \leq k$), and $D_n = D^*$ of cost V^* and recall $r^* = 1$. Note that the same detector type $D^{(j)}$ may well be used at the end of several segments. For notational convenience, we let $g_i = 1 - r_i$ be the probability that the i -th detector of the pattern fails to detect an error (for $1 \leq i < n$), and $g_{[i,j]} = \prod_{k=i}^{j-1} g_k$ be the probability that the error remains undetected by the detectors D_i to D_{j-1} (for $1 \leq i < j < n$). In the example of Figure 1, we have $W = w_1 + w_2 + w_3$ and $n = 3$. The first partial verification has cost V_1 with recall r_1 , and the second one has cost V_2 with recall r_2 .

Let W_{base} denote the base time of an application without any overhead due to resilience techniques (without loss of

generality, we assume unit-speed execution). Suppose the execution is divided into periodic patterns, defined by $\text{PATTERN}(W, n, \alpha, \mathbf{D})$.

Let $\mathbb{E}(W)$ be the expected execution time of the pattern. Then, the expected makespan W_{final} of the application when taking silent errors into account can be bounded as follows:

$$\left\lfloor \frac{W_{\text{base}}}{W} \right\rfloor \times \mathbb{E}(W) \leq W_{\text{final}} \leq \left\lceil \frac{W_{\text{base}}}{W} \right\rceil \times \mathbb{E}(W).$$

This is because the execution involves $\lfloor \frac{W_{\text{base}}}{W} \rfloor$ full patterns, and terminates by a (possibly) incomplete one. For large jobs, we approximate the execution time as

$$W_{\text{final}} \approx \frac{\mathbb{E}(W)}{W} \times W_{\text{base}}.$$

Let $H(W) = \frac{\mathbb{E}(W)}{W} - 1$ denote the relative *overhead* of the pattern. We obtain $W_{\text{final}} \approx W_{\text{base}} + H(W) \times W_{\text{base}}$. Thus, minimizing the expected makespan is equivalent to minimizing the pattern overhead $H(W)$.

We assume that errors only strike during computations, while verifications and I/O transfers (checkpointing and recovery) are protected and are thus error-free.

4 Properties of Optimal Patterns

In this section, we show how to compute the expected execution time of any given pattern (Section 4.1), and we derive a closed-form formula which is exact up to second-order terms. Based on this key result, given a pattern, we are able to characterize its optimal length (Section 4.2), and to compute the optimal positions of the partial verifications (Section 4.3).

4.1 Expected execution time of a pattern

Consider a pattern $\text{PATTERN}(W, n, \alpha, \mathbf{D})$. The following proposition computes the expected execution time of this pattern:

Proposition 1. *The expected execution time to execute a pattern $\text{PATTERN}(W, n, \alpha, \mathbf{D})$ is*

$$\mathbb{E}(W) = W + \sum_{i=1}^n V_i + C + \lambda W (R + W \alpha^T A \alpha + \mathbf{d}^T \alpha) + o(\lambda), \quad (1)$$

where A is an $n \times n$ symmetric matrix defined by $A_{ij} = \frac{1}{2} (1 + g_{[i,j]})$ for $i \leq j$, and \mathbf{d} is an $n \times 1$ vector defined by $\mathbf{d}_i = \sum_{j=i}^n g_{[i,j]} V_j$ for $1 \leq i \leq n$.

Proof. Let $q_i = 1 - e^{-\lambda w_i}$ denote the probability that one error occurs in the execution of segment i . We can express the expected execution time of the pattern by the following

recursive formula:

$$\begin{aligned} \mathbb{E}(W) &= \left(\prod_{k=1}^n (1 - q_k) \right) C + \left(1 - \prod_{k=1}^n (1 - q_k) \right) (R + \mathbb{E}(W)) \\ &+ \sum_{i=1}^n \left(\sum_{j=1}^{i-1} \left(\prod_{k=1}^{j-1} (1 - q_k) \right) q_j g_{[j,i]} + \prod_{k=1}^{i-1} (1 - q_k) \right) (w_i + V_i). \end{aligned} \quad (2)$$

The first line shows that checkpointing will only be taken if no error has occurred in all the segments, which happens with probability $\prod_{k=1}^n (1 - q_k)$, and in all the other cases, the application needs to recover from the last checkpoint and then re-computes the entire pattern. The second line shows the expected cost involved in the execution of each segment of the pattern and the associated verification. To better understand it, consider the third segment of size w_3 and the verification D_3 right after it, which will be executed only when the following events happen (with the probability of each event in brackets):

- There is a fault in the first segment (q_1), which is missed by the first verification ($1 - r_1 = g_1$) and again missed by the second verification ($1 - r_2 = g_2$).
- There is no fault in the first segment ($1 - q_1$), and there is a fault in the second segment (q_2), which is missed by the second verification ($1 - r_2 = g_2$).
- There is no fault in the first segment ($1 - q_1$) and no fault in the second segment ($1 - q_2$).

Thus, the expected cost involved in the execution of this segment is given by

$$\begin{aligned} &\left(q_1 g_1 g_2 + (1 - q_1) q_2 g_2 + (1 - q_1)(1 - q_2) \right) (w_3 + V_3) \\ &= \left(q_1 g_{[1,3]} + (1 - q_1) q_2 g_{[2,3]} + (1 - q_1)(1 - q_2) \right) (w_3 + V_3). \end{aligned}$$

We can generalize this reasoning to express the expected cost to execute the i -th segment of the pattern, which leads to Equation (2).

Now, substituting $q_i = 1 - e^{-\lambda w_i}$ into the recursive formula and solving for $\mathbb{E}(W)$, we obtain the expected execution time as

$$\begin{aligned} \mathbb{E}(W) &= C + (e^{\lambda W} - 1)R \\ &+ \sum_{i=1}^n \left(\sum_{j=1}^{i-1} (e^{\lambda W_{j,n}} - e^{\lambda W_{j+1,n}}) g_{[j,i]} + e^{\lambda W_{i,n}} \right) (w_i + V_i), \end{aligned}$$

where $W_{i,j} = \sum_{k=i}^j w_k$. Approximating $e^{\lambda x} = 1 + \lambda x + o(\lambda)$ to the first-order term, we can further simplify the expected

execution time as

$$\begin{aligned} \mathbb{E}(W) &= C + \lambda WR + o(\lambda) \\ &+ \sum_{i=1}^n \left(\sum_{j=1}^{i-1} \lambda w_j g_{[j,i]} + 1 + \lambda \sum_{j=i}^n w_j \right) (w_i + V_i) \\ &= W + \lambda WR + \sum_{i=1}^n V_i + C + o(\lambda) \\ &+ \lambda \sum_{i=1}^n \left(\sum_{j=1}^{i-1} w_j g_{[j,i]} + \sum_{j=i}^n w_j \right) (w_i + V_i). \end{aligned}$$

Letting $F = \sum_{i=1}^n \left(\sum_{j=1}^{i-1} w_j g_{[j,i]} + \sum_{j=i}^n w_j \right) (w_i + V_i)$, we have the following matrix form:

$$F = \mathbf{w}^T M \mathbf{w} + \mathbf{d}^T \mathbf{w},$$

where M is the following $n \times n$ matrix

$$M = \begin{bmatrix} 1 & g_{[1,2]} & g_{[1,3]} & \cdots & g_{[1,n]} \\ 1 & 1 & g_{[2,3]} & \cdots & g_{[2,n]} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \cdots & 1 \end{bmatrix}.$$

For instance, when $n = 4$ we have

$$M = \begin{bmatrix} 1 & g_1 & g_{1g_2} & g_{1g_2g_3} \\ 1 & 1 & g_2 & g_{2g_3} \\ 1 & 1 & 1 & g_3 \\ 1 & 1 & 1 & 1 \end{bmatrix}.$$

Replacing M by $A = \frac{M+M^T}{2}$ gives the same value for F , and we obtain the following symmetric matrix:

$$A = \frac{1}{2} \begin{bmatrix} 2 & 1 + g_{[1,2]} & \cdots & 1 + g_{[1,n]} \\ 1 + g_{[1,2]} & 2 & \cdots & 1 + g_{[2,n]} \\ \vdots & \vdots & \ddots & \vdots \\ 1 + g_{[1,n]} & 1 + g_{[2,n]} & \cdots & 2 \end{bmatrix}.$$

Now we use $\mathbf{w} = W\boldsymbol{\alpha}$, and obtain Equation (1). This completes the proof of the proposition. \square

4.2 Optimal length of a pattern to minimize the overhead

In this section, we compute the overhead $H(W)$ of a pattern $\text{PATTERN}(W, n, \boldsymbol{\alpha}, \mathbf{D})$ and show how to compute its length W to minimize this overhead. We introduce two key values:

Definition 1. Consider a pattern $\text{PATTERN}(W, n, \boldsymbol{\alpha}, \mathbf{D})$. The fault-free overhead o_{ff} of the pattern is

$$o_{\text{ff}} = \sum_{i=1}^n V_i + C, \quad (3)$$

and the fraction of re-executed work in case of faults is

$$f_{\text{re}} = \boldsymbol{\alpha}^T A \boldsymbol{\alpha}. \quad (4)$$

Proposition 2. The execution overhead of a pattern $\text{PATTERN}(W, n, \boldsymbol{\alpha}, \mathbf{D})$ is minimized when its length is

$$W^* = \sqrt{\frac{o_{\text{ff}}}{\lambda f_{\text{re}}}}. \quad (5)$$

In that case, we have

$$H(W^*) = 2\sqrt{\lambda o_{\text{ff}} f_{\text{re}}} + o(\sqrt{\lambda}). \quad (6)$$

Proof. From Equation (1), we directly derive

$$H(W) = \frac{\mathbb{E}(W)}{W} - 1 = \lambda R + \frac{o_{\text{ff}}}{W} + \lambda f_{\text{re}} W + \lambda \mathbf{d}^T \boldsymbol{\alpha} + o(\lambda).$$

The optimal pattern length that minimizes the execution overhead can now be computed from the above equation, and it is given by Equation (5).

Substituting W^* back into $H(W)$, the execution overhead is given by

$$\begin{aligned} H(W^*) &= 2\sqrt{\lambda o_{\text{ff}} f_{\text{re}}} + \lambda(R + \mathbf{d}^T \boldsymbol{\alpha}) + o(\lambda) \\ &= 2\sqrt{\lambda o_{\text{ff}} f_{\text{re}}} + o(\sqrt{\lambda}). \end{aligned}$$

\square

When the platform MTBF $\mu = 1/\lambda$ is large in front of the resilience parameters, Equation (6) shows that the expected execution overhead of the optimal size pattern is dominated by $2\sqrt{\lambda o_{\text{ff}} f_{\text{re}}}$, and the last term becomes negligible. The problem is then reduced to the minimization of the product $o_{\text{ff}} f_{\text{re}}$. Intuitively, this calls for a tradeoff between fault-free overhead and fault-induced re-execution, as a smaller value of o_{ff} with less overhead tends to induce a larger re-execution fraction f_{re} , and vice versa.

4.3 Optimal positions of verifications to minimize f_{re}

To fully characterize an optimal pattern, we have to determine its number of segments, and the type and position of each partial verification. In this section, we consider a pattern whose number of segments is given, as well as the type of all partial verifications, so that the value of o_{ff} (see Equation (3)) is given. We show how to determine the optimal length of each segment (or equivalently, the optimal position of each verification), so as to minimize the value of f_{re} (see Equation (4)). The result is technically involved and the reader may want to skip the lengthy proof.

Theorem 1. Consider a pattern $\text{PATTERN}(W, n, \boldsymbol{\alpha}, \mathbf{D})$ where W , n , and \mathbf{D} are given. The fraction of re-executed work f_{re} is minimized when $\boldsymbol{\alpha} = \boldsymbol{\alpha}^*$, where

$$\alpha_k^* = \frac{1}{U_n} \times \frac{1 - g_{k-1}g_k}{(1 + g_{k-1})(1 + g_k)} \quad \text{for } 1 \leq k \leq n, \quad (7)$$

where $g_0 = g_n = 0$ and

$$U_n = 1 + \sum_{i=1}^{n-1} \frac{1 - g_i}{1 + g_i}. \quad (8)$$

In that case, the value of f_{re} is

$$f_{re}^* = \frac{1}{2} \left(1 + \frac{1}{U_n} \right). \quad (9)$$

Proof. The goal is to minimize $f_{re} = \alpha^T A \alpha$ (Equation (4)) subject to the constraint $\sum_{k=1}^n \alpha_k = 1$, which we rewrite as $\mathbf{c}^T \alpha = 1$ with $\mathbf{c} = (1 \ 1 \ \dots \ 1)^T$. Hence, we have a quadratic minimization problem under a linear constraint. When A is symmetric positive definite (SPD), which we show later in the proof, there is a unique solution

$$f_{re}^{\text{opt}} = \frac{1}{\mathbf{c}^T A^{-1} \mathbf{c}}, \quad (10)$$

obtained for

$$\alpha^{\text{opt}} = \frac{A^{-1} \mathbf{c}}{\mathbf{c}^T A^{-1} \mathbf{c}}. \quad (11)$$

This result is shown as follows. Let a *valid* vector α be a vector such that $\mathbf{c}^T \alpha = 1$. We have $\mathbf{c}^T \alpha^{\text{opt}} = f_{re}^{\text{opt}} (\mathbf{c}^T A^{-1} \mathbf{c}) = 1$, hence α^{opt} is indeed a valid vector. Then, because A is SPD, we have $X = (\alpha - \alpha^{\text{opt}})^T A (\alpha - \alpha^{\text{opt}}) \geq 0$ for any valid vector α , and $X = 0$ if and only if $\alpha = \alpha^{\text{opt}}$. Developing X , we get

$$X = \alpha^T A \alpha - 2\alpha^T A \alpha^{\text{opt}} + (\alpha^{\text{opt}})^T A \alpha^{\text{opt}}.$$

We have $\alpha^T A \alpha^{\text{opt}} = f_{re}^{\text{opt}} \alpha^T \mathbf{c} = f_{re}^{\text{opt}}$ because $\mathbf{c}^T \alpha = 1$. Similarly, we get $(\alpha^{\text{opt}})^T A \alpha^{\text{opt}} = f_{re}^{\text{opt}}$. Hence, we derive that $X = \alpha^T A \alpha - f_{re}^{\text{opt}} \geq 0$, with equality if and only if $\alpha = \alpha^{\text{opt}}$. Hence the optimal value of f_{re} is achieved for α^{opt} , and is equal to f_{re}^{opt} .

In the following, we prove that A is symmetric positive definite (SPD), and that $\alpha^{\text{opt}} = \alpha^*$ and $f_{re}^{\text{opt}} = f_{re}^*$. To avoid ambiguity, we use superscripts like $A^{(n)}$ whenever needed to identify the problem size n (the number of work segments).

From Proposition 1, we can rewrite $A^{(n)}$ as:

$$A^{(n)} = \frac{1}{2} \left(J^{(n)} + B^{(n)} \right),$$

where $J^{(n)}$ is the $n \times n$ matrix whose entries are all 1, and $B^{(n)}$ is the $n \times n$ matrix defined by $B_{ij}^{(n)} = g_{[i,j]}$ for $i \leq j$.

We start by proving two properties of α^* .

Lemma 1. α^* is a valid vector, i.e., $\sum_{k=1}^n \alpha_k^* = 1$.

Proof. The proof is by induction on n . First, for $n = 1$ we do have $\sum_{k=1}^1 \alpha_k^{*(1)} = 1$. For $n = 2$, we have $\sum_{k=1}^2 \alpha_k^{*(2)} = \frac{1+g_1}{2} \left(\frac{1}{1+g_1} + \frac{1}{1+g_1} \right) = 1$, which is also correct. Assume that this result holds up to $n - 1$. We can express $\alpha^{*(n)}$ as:

$$\alpha^{*(n)} = \frac{U_{n-1}}{U_n} \begin{pmatrix} \alpha_1^{*(n-1)} \\ \alpha_2^{*(n-1)} \\ \vdots \\ \alpha_{n-2}^{*(n-1)} \\ \alpha_{n-1}^{*(n-1)} \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ -g_{n-1} \alpha_n^{*(n)} \\ \alpha_n^{*(n)} \end{pmatrix}. \quad (12)$$

Therefore we have:

$$\begin{aligned} \sum_{k=1}^n \alpha_k^{*(n)} &= \sum_{k=1}^{n-2} \alpha_k^{*(n)} + \alpha_{n-1}^{*(n)} + \alpha_n^{*(n)} \\ &= \frac{U_{n-1}}{U_n} \sum_{k=1}^{n-2} \alpha_k^{*(n-1)} + \frac{U_{n-1}}{U_n} \alpha_{n-1}^{*(n-1)} - g_{n-1} \alpha_n^{*(n)} + \alpha_n^{*(n)} \\ &= \frac{U_{n-1}}{U_n} \sum_{k=1}^{n-1} \alpha_k^{*(n-1)} + \alpha_n^{*(n)} (1 - g_{n-1}) \\ &= \frac{U_{n-1}}{U_n} \sum_{k=1}^{n-1} \alpha_k^{*(n-1)} + \frac{1}{U_n} \times \frac{1 - g_{n-1}}{1 + g_{n-1}}. \end{aligned}$$

Now, using the inductive hypothesis that $\sum_{k=1}^{n-1} \alpha_k^{*(n-1)} = 1$, we get:

$$\begin{aligned} \sum_{k=1}^n \alpha_k^{*(n)} &= \frac{1}{U_n} \left(U_{n-1} + \frac{1 - g_{n-1}}{1 + g_{n-1}} \right) \\ &= \frac{1}{U_n} \times U_n \\ &= 1, \end{aligned}$$

which concludes the proof. \square

Lemma 2. $A \alpha^* = f_{re}^* \mathbf{c}$.

Proof. We have $A = \frac{1}{2} (J + B)$ and from Lemma 1 $J \alpha^* = (\sum_{k=1}^n \alpha_k^*) \mathbf{c} = \mathbf{c}$. The result will follow if we show

$$B \alpha^* = \frac{\mathbf{c}}{U_n}, \quad (13)$$

for all $n \geq 1$. Equivalently, letting $\gamma = U_n \alpha^*$, we prove by induction on n that $B^{(n)} \gamma^{(n)} = \mathbf{c}^{(n)}$. First, for $n = 1$ we have $B^{(1)} \gamma^{(1)} = 1$, and for $n = 2$ we get:

$$B^{(2)} \gamma^{(2)} = \begin{pmatrix} 1 & g_1 \\ g_1 & 1 \end{pmatrix} \begin{pmatrix} \frac{1}{1+g_1} \\ \frac{1}{1+g_1} \end{pmatrix} = \begin{pmatrix} \frac{1}{1+g_1} + \frac{g_1}{1+g_1} \\ \frac{g_1}{1+g_1} + \frac{1}{1+g_1} \end{pmatrix} = \mathbf{c}^{(2)}.$$

Now, suppose the result holds up to $n - 1$. We can write:

$$\begin{aligned} B^{(n)} \gamma^{(n)} &= \begin{pmatrix} B^{(n-1)} & \mathbf{x}^{(n-1)} \\ (\mathbf{x}^{(n-1)})^T & 1 \end{pmatrix} \begin{pmatrix} \bar{\gamma}^{(n-1)} \\ \gamma_n^{(n)} \end{pmatrix} \\ &= \begin{pmatrix} B^{(n-1)} \bar{\gamma}^{(n-1)} + \mathbf{x}^{(n-1)} \gamma_n^{(n)} \\ (\mathbf{x}^{(n-1)})^T \bar{\gamma}^{(n-1)} + \gamma_n^{(n)} \end{pmatrix}, \quad (14) \end{aligned}$$

where $\bar{\gamma}^{(n-1)}$ is the $(n - 1) \times 1$ truncated vector containing the first $n - 1$ elements of $\gamma^{(n)}$ (for a problem of size n), and $\mathbf{x}^{(n-1)}$ is an $(n - 1) \times 1$ vector defined as $\mathbf{x}^{(n-1)} = (g_{[1,n-1]} \ g_{[2,n-1]} \ \dots \ g_{n-1})^T$. For instance, for $n = 4$ we have $\mathbf{x}^{(3)} = (g_1 g_2 g_3 \ g_2 g_3 \ g_3)^T$. Then the goal is to show $B^{(n-1)} \bar{\gamma}^{(n-1)} + \mathbf{x}^{(n-1)} \gamma_n^{(n)} = \mathbf{c}^{(n-1)}$ and $(\mathbf{x}^{(n-1)})^T \bar{\gamma}^{(n-1)} + \gamma_n^{(n)} = 1$. From Equation (12), we can

derive:

$$\begin{aligned}
& B^{(n-1)} \bar{\gamma}^{(n-1)} \\
&= B^{(n-1)} \left(\gamma^{(n-1)} + \begin{pmatrix} \mathbf{0}^{(n-2)} \\ -g_{n-1} \gamma_n^{(n)} \end{pmatrix} \right) \\
&= B^{(n-1)} \gamma^{(n-1)} + \begin{pmatrix} B^{(n-2)} & \mathbf{x}^{(n-2)} \\ (\mathbf{x}^{(n-2)})^T & 1 \end{pmatrix} \begin{pmatrix} \mathbf{0}^{(n-2)} \\ -g_{n-1} \gamma_n^{(n)} \end{pmatrix} \\
&= B^{(n-1)} \gamma^{(n-1)} + \begin{pmatrix} -\mathbf{x}^{(n-2)} g_{n-1} \gamma_n^{(n)} \\ -g_{n-1} \gamma_n^{(n)} \end{pmatrix} \\
&= \mathbf{c}^{(n-1)} - \mathbf{x}^{(n-1)} \gamma_n^{(n)}.
\end{aligned}$$

The last line applies the inductive hypothesis $B^{(n-1)} \gamma^{(n-1)} = \mathbf{c}^{(n-1)}$ as well as the property that $\begin{pmatrix} \mathbf{x}^{(n-2)} \\ 1 \end{pmatrix} g_{n-1} = \mathbf{x}^{(n-1)}$. Putting this result back into Equation (14), we derive that

$$\begin{aligned}
& B^{(n-1)} \bar{\gamma}^{(n-1)} + \mathbf{x}^{(n-1)} \gamma_n^{(n)} \\
&= \mathbf{c}^{(n-1)} - \mathbf{x}^{(n-1)} \gamma_n^{(n)} + \mathbf{x}^{(n-1)} \gamma_n^{(n)} = \mathbf{c}^{(n-1)}.
\end{aligned}$$

Using the property $(\mathbf{x}^{(n-1)})^T = \begin{pmatrix} (\mathbf{x}^{(n-2)})^T & 1 \end{pmatrix} g_{n-1}$, we can write:

$$\begin{aligned}
& (\mathbf{x}^{(n-1)})^T \bar{\gamma}^{(n-1)} + \gamma_n^{(n)} \\
&= \begin{pmatrix} (\mathbf{x}^{(n-2)})^T & 1 \end{pmatrix} g_{n-1} \left(\gamma^{(n-1)} + \begin{pmatrix} \mathbf{0}^{(n-2)} \\ -g_{n-1} \gamma_n^{(n)} \end{pmatrix} \right) + \gamma_n^{(n)} \\
&= \begin{pmatrix} (\mathbf{x}^{(n-2)})^T & 1 \end{pmatrix} \gamma^{(n-1)} g_{n-1} - g_{n-1}^2 \gamma_n^{(n)} + \gamma_n^{(n)}.
\end{aligned}$$

Notice that $\begin{pmatrix} (\mathbf{x}^{(n-2)})^T & 1 \end{pmatrix} \gamma^{(n-1)}$ is actually the last row of the product $B^{(n-1)} \gamma^{(n-1)}$, which, by induction, is 1. Therefore we get:

$$\begin{aligned}
& (\mathbf{x}^{(n-1)})^T \bar{\gamma}^{(n-1)} + \gamma_n^{(n)} \\
&= g_{n-1} + \gamma_n^{(n)} (1 - g_{n-1}^2) \\
&= g_{n-1} + \frac{(1 + g_{n-1})(1 - g_{n-1})}{1 + g_{n-1}} \\
&= g_{n-1} + 1 - g_{n-1} \\
&= 1.
\end{aligned}$$

This concludes the proof. \square

We now prove that A is SPD. This requires several intermediate steps.

Lemma 3. B is nonsingular and $\alpha^* = \frac{1}{U_n} B^{-1} \mathbf{c}$.

Proof. To prove that B is nonsingular, we prove by induction on n that $B^{(n)} \mathbf{y}^{(n)} = \mathbf{0}^{(n)}$ has only one solution $\mathbf{y}^{(n)} = \mathbf{0}^{(n)}$. First, for $n = 1$ we have $y_1^{(1)} = 0$, which is correct. Then, for $n = 2$ we have the following equation:

$$\begin{pmatrix} 1 & g_1 \\ g_1 & 1 \end{pmatrix} \begin{pmatrix} y_1^{(2)} \\ y_2^{(2)} \end{pmatrix} = \mathbf{0}^{(2)},$$

from which we derive $y_1^{(2)}(1 - g_1^2) = 0$ and $y_2^{(2)}(1 - g_1^2) = 0$, hence $\mathbf{y}^{(2)} = \mathbf{0}^{(2)}$, which is also correct. Now, assume that the result holds up to $n - 1$. We want to solve the general equation:

$$\begin{pmatrix} B^{(n-1)} & \mathbf{x}^{(n-1)} \\ (\mathbf{x}^{(n-1)})^T & 1 \end{pmatrix} \begin{pmatrix} \bar{\mathbf{y}}^{(n-1)} \\ y_n^{(n)} \end{pmatrix} = \mathbf{0}^{(n)}, \quad (15)$$

which is equivalent to:

$$\begin{pmatrix} B^{(n-2)} & \mathbf{x}^{(n-2)} & \mathbf{x}^{(n-2)} g_{n-1} \\ (\mathbf{x}^{(n-2)})^T & 1 & g_{n-1} \\ (\mathbf{x}^{(n-2)})^T g_{n-1} & g_{n-1} & 1 \end{pmatrix} \begin{pmatrix} \bar{\mathbf{y}}^{(n-2)} \\ y_{n-1}^{(n)} \\ y_n^{(n)} \end{pmatrix} = \mathbf{0}^{(n)}, \quad (16)$$

where $\bar{\mathbf{y}}^{(n-1)}$ and $\bar{\mathbf{y}}^{(n-2)}$ are the truncated vectors containing respectively the first $n - 1$ and $n - 2$ elements of $\mathbf{y}^{(n)}$ (for a problem of size n). First, let us expand Equation (16) and consider only the last two equations of the system:

$$\begin{aligned}
& (\mathbf{x}^{(n-2)})^T \bar{\mathbf{y}}^{(n-2)} + y_{n-1}^{(n)} + g_{n-1} y_n^{(n)} = 0 \\
& g_{n-1} (\mathbf{x}^{(n-2)})^T \bar{\mathbf{y}}^{(n-2)} + g_{n-1} y_{n-1}^{(n)} + y_n^{(n)} = 0.
\end{aligned}$$

We can derive that $y_n^{(n)}(1 - g_{n-1}^2) = 0$, hence $y_n^{(n)} = 0$. Then, plugging $y_n^{(n)} = 0$ back into Equation (15), we derive that:

$$B^{(n-1)} \bar{\mathbf{y}}^{(n-1)} = \mathbf{0}^{(n-1)}.$$

Using the induction hypothesis for $B^{(n-1)} \bar{\mathbf{y}}^{(n-1)} = \mathbf{0}^{(n-1)}$, we have $\bar{\mathbf{y}}^{(n-1)} = \mathbf{0}^{(n-1)}$ and thus $\mathbf{y}^{(n)} = \mathbf{0}^{(n)}$, which implies that $B^{(n)}$ is nonsingular. Hence, from Equation (13), we can get:

$$\alpha^* = \frac{1}{U_n} B^{-1} \mathbf{c},$$

which concludes the proof. \square

Lemma 4. A is nonsingular.

Proof. To prove that A is nonsingular, we solve $A\mathbf{y} = \mathbf{0}$ and show that $\mathbf{y} = \mathbf{0}$. First, we can write:

$$J\mathbf{y} + B\mathbf{y} = \mathbf{0},$$

$$B\mathbf{y} = -J\mathbf{y} = -\left(\sum_{i=1}^n y_i \right) \mathbf{c}.$$

From Lemma 3, we know that B is nonsingular and $B^{-1} \mathbf{c} = U_n \alpha^*$. Therefore, we get:

$$\mathbf{y} = -U_n \left(\sum_{i=1}^n y_i \right) \alpha^*. \quad (17)$$

Summing the components of both sides of Equation (17), we obtain:

$$\left(\sum_{i=1}^n y_i \right) = -U_n \left(\sum_{i=1}^n y_i \right) \left(\sum_{i=1}^n \alpha_i^* \right).$$

Since $\sum_{i=1}^n \alpha_i^* = 1$ from Lemma 1, we have:

$$\left(\sum_{i=1}^n y_i \right) (1 + U_n) = 0,$$

$$\sum_{i=1}^n y_i = 0,$$

which implies $\mathbf{y} = \mathbf{0}$ from Equation (17), and this concludes the proof that A is nonsingular. \square

Lemma 5. *The last column of B^{-1} is given by:*

$$\mathbf{z} = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ -g_{n-1}z_n \\ z_n \end{pmatrix}, \text{ with } z_n = \frac{1}{1 - g_{n-1}^2}.$$

Proof. Because we do not need the whole inverse of B , we solve $B\mathbf{z} = \mathbf{d}$, where $\mathbf{d} = (0 \ 0 \ \dots \ 1)^T$, hence \mathbf{z} will be the last column of B^{-1} . We can write:

$$\begin{pmatrix} B^{(n-2)} & \mathbf{x}^{(n-2)} & \mathbf{x}^{(n-2)}g_{n-1} \\ (\mathbf{x}^{(n-2)})^T & 1 & g_{n-1} \\ (\mathbf{x}^{(n-2)})^T g_{n-1} & g_{n-1} & 1 \end{pmatrix} \begin{pmatrix} \bar{\mathbf{z}}^{(n-2)} \\ z_{n-1}^{(n)} \\ z_n^{(n)} \end{pmatrix} = \mathbf{d}^{(n)},$$

where $\bar{\mathbf{z}}^{(n-2)}$ is the truncated vector containing the first $n-2$ elements of $\mathbf{z}^{(n)}$. Expanding the product, we get the following system of equations:

$$B^{(n-2)}\bar{\mathbf{z}}^{(n-2)} + \mathbf{x}^{(n-2)}z_{n-1}^{(n)} + \mathbf{x}^{(n-2)}g_{n-1}z_n^{(n)} = \mathbf{0}^{(n-2)},$$

$$(\mathbf{x}^{(n-2)})^T \bar{\mathbf{z}}^{(n-2)} + z_{n-1}^{(n)} + g_{n-1}z_n^{(n)} = 0,$$

$$(\mathbf{x}^{(n-2)})^T g_{n-1}\bar{\mathbf{z}}^{(n-2)} + g_{n-1}z_{n-1}^{(n)} + z_n^{(n)} = 1.$$

Since $B^{(n)}$ is nonsingular, there is a unique solution. We can check that $\bar{\mathbf{z}}^{(n-2)} = \mathbf{0}^{(n-2)}$, $z_{n-1}^{(n)} = \frac{-g_{n-1}}{1-g_{n-1}^2}$ and $z_n^{(n)} = \frac{1}{1-g_{n-1}^2}$ is indeed a solution, which concludes the proof. \square

Remark. The matrix B is an extension of the famous KMS symmetric matrix K [18], where $K_{ij} = g^{j-i}$ for $i \leq j$ (recall that $B_{ij} = g_{[i,j]}$). The inverse of B turns out to be tridiagonal, just as that of K , and we get:

$$B_{ij}^{-1} = \begin{cases} -\frac{g_j}{1-g_j^2} & \text{if } i = j+1 \\ -\frac{g_i}{1-g_i^2} & \text{if } i = j-1 \\ \frac{1-g_{i-1}^2 g_i^2}{(1-g_{i-1}^2)(1-g_i^2)} & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}.$$

For instance, when $n = 4$ we have

$$B^{-1} = \begin{bmatrix} \frac{1}{1-g_1^2} & -\frac{g_1}{1-g_1^2} & 0 & 0 \\ -\frac{g_1}{1-g_1^2} & \frac{1-g_1^2 g_2^2}{(1-g_1^2)(1-g_2^2)} & -\frac{g_2}{1-g_2^2} & 0 \\ 0 & -\frac{g_2}{1-g_2^2} & \frac{1-g_2^2 g_3^2}{(1-g_2^2)(1-g_3^2)} & -\frac{g_3}{1-g_3^2} \\ 0 & 0 & -\frac{g_3}{1-g_3^2} & \frac{1}{1-g_3^2} \end{bmatrix}.$$

The proof of this result is very similar to the proof of Lemma 5.

Lemma 6. $A_{nn}^{-1} = 2 \frac{U_n(1+g_{n-1})+2g_{n-1}}{(U_n+1)(1-g_{n-1})(1+g_{n-1})^2}.$

Proof. As in the proof of Lemma 5, we compute the last column of A^{-1} , which we call β , by solving $A\beta = \mathbf{d}$. Because we already solved $B\mathbf{z} = \mathbf{d}$, we have:

$$A\beta = B\mathbf{z} = \mathbf{d},$$

$$\frac{1}{2}(J+B)\beta = B\mathbf{z},$$

$$J\beta = B(2\mathbf{z} - \beta).$$

Remember that J is the matrix whose entries are all 1. Hence, we have $J\beta = (\sum_{i=1}^n \beta_i) \mathbf{c}$. Also, from Lemma 3, we have $B\alpha^* = \frac{\mathbf{c}}{U_n}$. Therefore, we can derive:

$$2\mathbf{z} - \beta = \left(\sum_{i=1}^n \beta_i \right) U_n \alpha^*. \quad (18)$$

Summing the components of both sides of Equation (18), we get $2 \sum_{i=1}^n z_i - \sum_{i=1}^n \beta_i = (\sum_{i=1}^n \beta_i) U_n (\sum_{i=1}^n \alpha_i^*)$. Since $\sum_{i=1}^n \alpha_i^* = 1$ from Lemma 1, we get

$$\sum_{i=1}^n \beta_i = \frac{2}{U_n + 1} \sum_{i=1}^n z_i.$$

From Lemma 5, we can easily compute $\sum_{i=1}^n z_i = -g_{n-1}z_n + z_n = \frac{1}{1+g_{n-1}}$. Hence, we have

$$\sum_{i=1}^n \beta_i = \frac{2}{(U_n + 1)(1 + g_{n-1})}.$$

Finally, from Equation (18), we derive that

$$\begin{aligned} \beta_n &= 2z_n - \left(\sum_{i=1}^n \beta_i \right) U_n \alpha_n^* \\ &= \frac{2}{1 - g_{n-1}^2} - \frac{2}{(U_n + 1)(1 + g_{n-1})^2} \\ &= 2 \frac{U_n(1 + g_{n-1}) + 2g_{n-1}}{(U_n + 1)(1 - g_{n-1})(1 + g_{n-1})^2}, \end{aligned}$$

which concludes the proof. \square

Lemma 7. *A is symmetric positive definite (SPD).*

Proof. Note that by construction, A , J and B are all symmetric matrices. To show that A is positive definite, we show that all its principal minors are strictly positive. Recall that the principal minor of order k of $A^{(n)}$ is the determinant of the submatrix of size k that consists of the first k rows and columns of $A^{(n)}$. But this submatrix is exactly $A^{(k)}$, the matrix for the problem of size k , so the result will follow if we show that $\det(A^{(n)}) > 0$ for all $n \geq 1$. We prove by induction on n that

$$\det(A^{(n)}) = \frac{U_n + 1}{2^n} \prod_{k=1}^{n-1} (1 - g_k^2) > 0. \quad (19)$$

For $n = 1$, Equation (19) gives $\det(A^{(1)}) = 1$, which is correct. Suppose the result holds up to $n - 1$. Since $A^{(n)}$ is nonsingular, using the co-factor method, we get that

$$\left(A^{(n)}\right)_{nn}^{-1} = \frac{\det(A^{(n-1)})}{\det(A^{(n)})}.$$

Therefore, using the definition of U_n and the induction hypothesis for $\det(A^{(n-1)})$, we can get:

$$\begin{aligned} \det(A^{(n)}) &= \frac{\det(A^{(n-1)})}{\left(A^{(n)}\right)_{nn}^{-1}} \\ &= \frac{1}{\left(A^{(n)}\right)_{nn}^{-1}} \cdot \frac{U_{n-1} + 1}{2^{n-1}} \prod_{k=1}^{n-2} (1 - g_k^2) \\ &= \frac{1}{\left(A^{(n)}\right)_{nn}^{-1}} \cdot \frac{1}{2^{n-1}} \left(U_n - \frac{1 - g_{n-1}}{1 + g_{n-1}} + 1 \right) \prod_{k=1}^{n-2} (1 - g_k^2) \\ &= \frac{1}{\left(A^{(n)}\right)_{nn}^{-1}} \cdot \frac{1}{2^{n-1}} \cdot \frac{U_n(1 + g_{n-1}) + 2g_{n-1}}{1 + g_{n-1}} \prod_{k=1}^{n-2} (1 - g_k^2). \end{aligned} \quad (20)$$

Now, plugging $\left(A^{(n)}\right)_{nn}^{-1}$ from Lemma 6 into Equation (20), we get:

$$\begin{aligned} \det(A^{(n)}) &= \frac{1}{2^n} \cdot \frac{(U_n + 1)(1 - g_{n-1})(1 + g_{n-1})^2}{U_n(1 + g_{n-1}) + 2g_{n-1}} \\ &\quad \times \frac{U_n(1 + g_{n-1}) + 2g_{n-1}}{1 + g_{n-1}} \prod_{k=1}^{n-2} (1 - g_k^2) \\ &= \frac{U_n + 1}{2^n} (1 - g_{n-1}^2) \prod_{k=1}^{n-2} (1 - g_k^2) \\ &= \frac{U_n + 1}{2^n} \prod_{k=1}^{n-1} (1 - g_k^2), \end{aligned}$$

which shows that Equation (19) holds for $\det(A^{(n)})$ and completes the proof that $A^{(n)}$ is SPD. \square

We are almost done! There remains to show that $\alpha^{\text{opt}} = \alpha^*$ and $f_{\text{re}}^{\text{opt}} = f_{\text{re}}^*$. But Lemma 2 shows that $A\alpha^* = f_{\text{re}}^* \mathbf{c}$, hence $\alpha^* = f_{\text{re}}^* A^{-1} \mathbf{c}$ and $1 = \mathbf{c}^T \alpha^* = f_{\text{re}}^* (\mathbf{c}^T A^{-1} \mathbf{c})$, which leads to $f_{\text{re}}^{\text{opt}} = f_{\text{re}}^*$, and finally $\alpha^{\text{opt}} = \alpha^*$. This concludes the proof of Theorem 1. \square

When all the partial verifications in the pattern have the same type, i.e., $g_k = g$ for all $1 \leq k \leq n$, we retrieve the result of [11], obtaining $f_{\text{re}}^* = \frac{1}{2} \left(1 + \frac{1+g}{n(1-g)+2g} \right)$ with

$$\alpha_k^* = \begin{cases} \frac{1}{n(1-g)+2g} & \text{for } k = 1, n \\ \frac{1-g}{n(1-g)+2g} & \text{for } k = 2, \dots, n-1 \end{cases}.$$

Theorem 1 also shows that, for a given set of partial verifications in the pattern, the minimal value of f_{re} does not depend upon their ordering within the pattern:

Corollary 1. *For a given set of partial verifications within a pattern, the minimum fraction of re-executed work f_{re}^* is independent of their ordering.*

5 Complexity

This section builds upon previous results to provide a comprehensive complexity analysis. We introduce the *accuracy-to-cost ratio* of a detector and show that it is the key parameter to compute the optimal rational solution (Section 5.1). Then we establish the NP-completeness to determine the optimal integer solution (Section 5.2). On the positive side, we design a simple algorithm whose performance is guaranteed, and sketch the construction of an FPTAS for the problem (Section 5.3).

5.1 Accuracy-to-cost ratio and rational solution

Consider a pattern $\text{PATTERN}(W, n, \alpha, \mathbf{D})$. Let m_j denote the number of partial verifications using detector type $D^{(j)}$ in the pattern (the number of indices $i < n$ such that D_i is of type $D^{(j)}$), and define $\mathbf{m} = (m_1, m_2, \dots, m_k)$. Section 4.2 shows that minimizing the execution overhead of the pattern is equivalent to minimizing the product $o_{\text{ff}} f_{\text{re}}$. From Equations (3) and (9), we have $o_{\text{ff}} f_{\text{re}} = \frac{V^* + C}{2} f(\mathbf{m})$, where

$$f(\mathbf{m}) = \left(1 + \frac{1}{1 + \sum_{j=1}^k m_j a_j} \right) \left(1 + \sum_{j=1}^k m_j b_j \right). \quad (21)$$

In Equation (21), we use the *accuracy* $a_j = \frac{1-g^{(j)}}{1+g^{(j)}}$ and the *relative cost* $b_j = \frac{V^{(j)}}{V^* + C}$ of detector $D^{(j)}$. We further define $\phi^{(j)} = \frac{a_j}{b_j}$ to be the *accuracy-to-cost ratio* of $D^{(j)}$, and show that this ratio will play a key role in selecting the best detectors.

Altogether, minimizing the pattern overhead amounts to finding the solution $\mathbf{m} = (m_1, m_2, \dots, m_k)$ that minimizes $f(\mathbf{m})$, with $m_j \in \mathbb{N}_0$ for all $1 \leq j \leq k$. Indeed, once \mathbf{m} is given, Proposition 2 and Theorem 1 completely characterize the optimal pattern, giving its length W , the number of segments $n = \sum_{j=1}^k m_j + 1$, and the locations α of all partial detectors (whose ordering does not matter).

We first derive the optimal solution if we relax the integer constraint on \mathbf{m} . A rational solution in this case is denoted by $\bar{\mathbf{m}} = (\bar{m}_1, \bar{m}_2, \dots, \bar{m}_k)$, where $\bar{m}_j \geq 0$ for all $1 \leq j \leq k$. The optimal value of $f(\bar{\mathbf{m}})$ is a lower bound on the optimal integer solution.

Lemma 8. *Suppose there are k types of detectors sorted in non-increasing order of accuracy-to-cost ratio, i.e., $\phi^{(1)} \geq \phi^{(2)} \geq \dots \geq \phi^{(k)}$. Then,*

$$f^*(\bar{\mathbf{m}}) = \begin{cases} \left(\sqrt{\frac{1}{\phi^{(1)}}} + \sqrt{1 - \frac{1}{\phi^{(1)}}} \right)^2 & \text{if } \phi^{(1)} > 2 \\ 2 & \text{otherwise} \end{cases}.$$

Proof. First, we prove that the optimal rational solution is achieved when only the detector with largest accuracy-to-cost ratio $\phi^{(1)}$ is used. Specifically, given any rational solution $\bar{\mathbf{m}} = (\bar{m}_1, \bar{m}_2, \dots, \bar{m}_k)$, we show that

there exists a solution $\bar{\mathbf{m}}' = (\bar{m}'_1, 0, \dots, 0)$, which satisfies $f(\bar{\mathbf{m}}') \leq f(\bar{\mathbf{m}})$. We have:

$$\begin{aligned} f(\bar{\mathbf{m}}) &= \left(1 + \frac{1}{1 + \sum_{j=1}^k \bar{m}_j a_j}\right) \left(1 + \sum_{j=1}^k \bar{m}_j b_j\right) \\ &= \left(1 + \frac{1}{1 + a_1 \sum_{j=1}^k \frac{\bar{m}_j a_j}{a_1}}\right) \left(1 + b_1 \sum_{j=1}^k \frac{\bar{m}_j b_j}{b_1}\right). \end{aligned} \quad (22)$$

Let $\bar{m}'_1 = \sum_{j=1}^k \frac{\bar{m}_j a_j}{a_1}$ and $\bar{n}'_1 = \sum_{j=1}^k \frac{\bar{m}_j b_j}{b_1}$. Since $\frac{b_j}{a_j} \geq \frac{a_j}{a_1}$ for all $1 \leq j \leq k$, we get $\bar{n}'_1 = \sum_{j=1}^k \frac{\bar{m}_j b_j}{b_1} \geq \sum_{j=1}^k \frac{\bar{m}_j a_j}{a_1} = \bar{m}'_1$. Hence, Equation (22) can be written as

$$\begin{aligned} f(\bar{\mathbf{m}}) &= \left(1 + \frac{1}{1 + a_1 \bar{m}'_1}\right) (1 + b_1 \bar{n}'_1) \\ &= \left(1 + \frac{1}{1 + a_1 \bar{m}'_1}\right) \left(1 + \frac{b_1 \bar{n}'_1}{\bar{m}'_1} \cdot \bar{m}'_1\right) \\ &\geq \left(1 + \frac{1}{1 + a_1 \bar{m}'_1}\right) (1 + b_1 \bar{m}'_1) \\ &= f(\bar{\mathbf{m}}'). \end{aligned}$$

Now, define $f(\bar{m}) = \left(1 + \frac{1}{1 + a_1 \bar{m}}\right) (1 + b_1 \bar{m})$. The following derives the minimum value of $f(\bar{m})$. Differentiating $f(\bar{m})$ with respect to \bar{m} and solving $\partial f(\bar{m})/\partial \bar{m} = 0$, we get:

$$\bar{m}^* = -\frac{1}{a_1} + \sqrt{\frac{1}{a_1} \left(\frac{1}{b_1} - \frac{1}{a_1}\right)}, \quad (23)$$

which is positive (hence a potential solution) if $\phi^{(1)} = a_1/b_1 > 2$. Taking the second-order derivative of $f(\bar{m})$, we get:

$$\frac{\partial^2 f(\bar{m})}{\partial \bar{m}^2} = \frac{2a_1(a_1 - b_1)}{(a_1 \bar{m} + 1)^3},$$

which is positive (hence ensures that the solution is the unique minimum) for all $\bar{m} \in [0, \infty)$ if $\phi^{(1)} = a_1/b_1 > 1$.

Thus, when $\phi^{(1)} > 2$, the optimal solution is obtained by substituting \bar{m}^* into $f(\bar{m})$, and we get:

$$\begin{aligned} f(\bar{m}^*) &= \left(1 + \frac{1}{1 + a_1 \bar{m}^*}\right) (1 + b_1 \bar{m}^*) \\ &= \left(1 + \frac{1}{\sqrt{\phi^{(1)} - 1}}\right) \left(1 - \frac{1}{\phi^{(1)}} + \sqrt{\frac{1}{\phi^{(1)}} \left(1 - \frac{1}{\phi^{(1)}}\right)}\right) \\ &= \left(\sqrt{\frac{1}{\phi^{(1)}}} + \sqrt{1 - \frac{1}{\phi^{(1)}}}\right)^2. \end{aligned}$$

When $\phi^{(1)} \leq 2$, the minimum value of $f(\bar{m})$ is achieved at $\bar{m} = 0$, which gives $f(0) = 2$. \square

Lemma 8 shows that the optimal rational solution is achieved with only one detector, namely, the one with highest accuracy-to-cost ratio. The optimal integer solution, however, may use more than one detector. The following shows that finding the optimal integer solution is NP-complete.

5.2 NP-completeness

We show that finding the optimal solution \mathbf{m} is NP-complete, even when all detectors share the same accuracy-to-cost ratio. In particular, we consider the following decision problem.

Definition 2 (Multiple Partial Verifications (MPV)). *Given k detectors with the same accuracy-to-cost ratio ϕ ($a_j/b_j = \phi$ for $1 \leq j \leq k$) and a bound K , is there a solution \mathbf{m} that satisfies*

$$\left(1 + \frac{1}{1 + \sum_{j=1}^k m_j a_j}\right) \left(1 + \sum_{j=1}^k m_j b_j\right) \leq K? \quad (24)$$

Theorem 2. *The MPV problem is NP-complete.*

Proof. The MPV problem is obviously in NP. We prove the completeness by a reduction from the *Unbounded Subset Sum (USS)* problem, which is known to be NP-complete [22]. Given a multiset $S = \{s_1, s_2, \dots, s_k\}$ of k positive integers and a positive integer I , the USS problem asks if there exists a subset $S' \subseteq S$ whose sum is exactly I , i.e., $\sum_{j=1}^k m_j s_j = I$, where $m_j \in \mathbb{N}_0$. We can further assume that I/s_j is not an integer for $1 \leq j \leq k$, since otherwise we would have a trivial solution.

Given an instance of the USS problem, we construct an instance of the MPV problem with k detectors. First, choose any $\phi \in \left(2, (I/s_{\max} + 1)^2 + 1\right)$, where $s_{\max} = \max_{j=1}^k s_j$. Then, let $\frac{a}{b} = \phi$ and $-\frac{1}{a} + \sqrt{\frac{1}{a} \left(\frac{1}{b} - \frac{1}{a}\right)} = I$, so we can get $a = \frac{\sqrt{\phi-1}-1}{I}$ and $b = \frac{\sqrt{\phi-1}-1}{\phi I}$. For each $1 \leq j \leq k$, define $a_j = s_j a$ and $b_j = s_j b$. According to the range of ϕ , we have $a_j < 1$ and $b_j < 1$ for all $1 \leq j \leq k$. Finally, let $K = \left(\sqrt{\frac{1}{\phi}} + \sqrt{1 - \frac{1}{\phi}}\right)^2$.

If we use only one detector, say $D^{(j)}$, then Lemma 8 shows that Equation (24) is satisfied with the following unique solution:

$$\begin{aligned} m_j^* &= -\frac{1}{a_j} + \sqrt{\frac{1}{a_j} \left(\frac{1}{b_j} - \frac{1}{a_j}\right)} \\ &= \frac{1}{s_j} \left(-\frac{1}{a} + \sqrt{\frac{1}{a} \left(\frac{1}{b} - \frac{1}{a}\right)}\right) = \frac{I}{s_j}, \end{aligned}$$

which is not an integer by hypothesis, but achieves the lower bound $\left(\sqrt{\frac{1}{\phi}} + \sqrt{1 - \frac{1}{\phi}}\right)^2 = K$. Now, we show that, by using multiple detectors, an integer solution to the MPV instance exists if and only if there is an integer solution to the USS instance.

(\Rightarrow) Suppose there is an integer solution $\mathbf{m} = (m_1, m_2, \dots, m_k)$ such that $\sum_{j=1}^k m_j s_j = I$. Then, by employing m_j partial verifications of detector type $D^{(j)}$ for $1 \leq j \leq k$, we get

$$\begin{aligned}
& \left(1 + \frac{1}{1 + \sum_{j=1}^k m_j a_j}\right) \left(1 + \sum_{j=1}^k m_j b_j\right) \\
&= \left(1 + \frac{1}{1 + a \sum_{j=1}^k m_j s_j}\right) \left(1 + b \sum_{j=1}^k m_j s_j\right) \\
&= \left(1 + \frac{1}{1 + aI}\right) (1 + bI) = \left(\sqrt{\frac{1}{\phi}} + \sqrt{1 - \frac{1}{\phi}}\right)^2 = K.
\end{aligned}$$

(\Leftarrow) Suppose there is an integer solution $\mathbf{m} = (m_1, m_2, \dots, m_k)$ to the MPV problem such that

$$\left(1 + \frac{1}{1 + \sum_{j=1}^k m_j a_j}\right) \left(1 + \sum_{j=1}^k m_j b_j\right) = K.$$

This implies

$$\begin{aligned}
& \left(1 + \frac{1}{1 + a \sum_{j=1}^k m_j s_j}\right) \left(1 + b \sum_{j=1}^k m_j s_j\right) \\
&= 1 + 2\sqrt{\frac{1}{\phi} \left(1 - \frac{1}{\phi}\right)}.
\end{aligned}$$

Let $T = \sum_{j=1}^k m_j s_j$. Solving T from the equation above, we get the following unique solution:

$$T = -\frac{1}{a} + \sqrt{\frac{1}{a} \left(\frac{1}{b} - \frac{1}{a}\right)} = I.$$

This completes the proof of the theorem. \square

5.3 Greedy algorithm and FPTAS

To cope with the NP-completeness of minimizing $o_{\text{ff}} f_{\text{re}}$, there is a simple and intuitive algorithm. This greedy algorithm uses only the detector with highest accuracy-to-cost ratio $\phi^{(1)}$. We compute the optimal rational number of partial verifications \bar{m}^* and round it up if it is not an integer. In Section 6, we show that this algorithm performs quite well in practice.

Interestingly, we can guarantee the performance of this simple algorithm. From Lemma 8, we can assume $\phi^{(1)} = a_1/b_1 > 2$. Since $a_1 < 1$, we can get $b_1 < 1/2$. If the optimal fractional solution \bar{m}^* given in Equation (23) happens to be an integer, then we get the optimal solution. Otherwise, rounding it to $\lceil \bar{m}^* \rceil$ increases the objective function $f(\mathbf{m})$ shown in Equation (21) by at most a factor of $\delta = 1 + b_1 < 3/2$. According to Equation (6), this gives a $\sqrt{3/2}$ -approximation algorithm for minimizing the expected execution overhead (and hence the makespan).

In the following, we show that it is possible to have a fully polynomial-time approximation scheme (FPTAS), which ensures, for any $\epsilon > 0$, that the solution is within $1 + \epsilon$ times the optimal, and that the running time of the algorithm

is polynomial in the input size and $1/\epsilon$. To develop the FPTAS, we perform the following transformations to the problem.

First, we convert all parameters in Equation (21) to integers. Since $a_j = \frac{1-g^{(j)}}{1+g^{(j)}} = \frac{r^{(j)}}{2-r^{(j)}} \leq 1$ and $r^{(j)}$ is rational, we can write $a_j = \frac{p_j}{q_j}$, where p_j and q_j are positive integers with $p_j \leq q_j$. Thus, minimizing $f(\mathbf{m})$ is equivalent to minimizing the following function:

$$F(\mathbf{m}) = \left(1 + \frac{L}{L + \sum_{j=1}^k m_j L^{(j)}}\right) \left(C + V^* + \sum_{j=1}^k m_j V^{(j)}\right),$$

where L denotes the least common multiple (LCM) of q_1, q_2, \dots, q_k , and $L^{(j)} = \frac{p_j}{q_j} L \leq L$. Clearly, L and all the $L^{(j)}$'s can be represented by a polynomial function of the original input size.

Next, we compute an upper bound on the number of partial verifications. Observe that $F(\mathbf{0}) = 2(C + V^*)$ and $F(\mathbf{m}) \geq C + V^* + \sum_{j=1}^k m_j V^{(j)}$. This implies that the optimal solution must satisfy $m_j \leq \frac{C+V^*}{V^{(j)}}$ for all $1 \leq j \leq k$. Therefore, it follows that $\sum_{j=1}^k m_j V^{(j)} \leq k(C + V^*)$. The bound on m_j allows us to transform the unbounded problem to the 0-1 problem by providing $\lceil \log m_j \rceil$ additional copies of each item type j with doubling $V^{(j)}$ and $L^{(j)}$ values. This is a standard technique also used in transforming the bounded and unbounded knapsack problems [25]. The total number of items becomes $K = \sum_{j=1}^k (1 + \lceil \log m_j \rceil) = O(k \log(C + V^*))$, which stays polynomial in the input size.

Define $\mathbf{x} = (x_1, x_2, \dots, x_K)$, and let L_j and V_j be the ratio and cost for item j . We can now formulate the optimization problem as follows:

$$\begin{aligned}
& \text{minimize } F(\mathbf{x}) = \left(1 + \frac{L}{L + \sum_{j=1}^K x_j L_j}\right) \left(C + V^* + \sum_{j=1}^K x_j V_j\right) \\
& \text{subject to } \sum_{j=1}^K x_j V_j \leq k(C + V^*) \\
& \quad x_j \in \{0, 1\} \quad \forall j = 1, 2, \dots, K
\end{aligned}$$

and the size of all parameters is a polynomial function of the input size of the original problem. To find an FPTAS for the problem above, we adopt the technique used in [15] for designing an FPTAS for the *Maximum Density Knapsack (MDK)* problem described below.

Maximum Density Knapsack (MDK): Given a set $S = \{s_1, s_2, \dots, s_K\}$ of K items, where each item $s_j \in S$ has a positive integer profit p_j and a positive integer weight w_j , a total capacity W , and an initial weight w_0 , the MDK problem is formulated as:

$$\begin{aligned}
& \text{maximize } \frac{\sum_{j=1}^K x_j p_j}{w_0 + \sum_{j=1}^K x_j w_j} \\
& \text{subject to } \sum_{j=1}^K x_j w_j \leq W \\
& \quad x_j \in \{0, 1\} \quad \forall j = 1, 2, \dots, K
\end{aligned}$$

Cohen and Katzir [15] give an FPTAS for the MDK problem by using the existing FPTAS for the knapsack problem [25]. In particular, their algorithm relies on the property that, for every profit P , a minimum weight solution \mathbf{x} is found such that $P(\mathbf{x}) = \sum_{j=1}^K x_j p_j \geq \lfloor \frac{P}{1+\epsilon'} \rfloor$, for any $\epsilon' > 0$. This immediately gives rise to an FPTAS for MDK.

We can apply the same technique to construct an FPTAS for minimizing $F(\mathbf{x})$. Let \mathbf{x}_{opt} denote the optimal solution. By considering V_j as weight and L_j as profit, we can run the FPTAS for knapsack and return in polynomial time a solution \mathbf{x} that satisfies $P(\mathbf{x}) \geq \lfloor \frac{P(\mathbf{x}_{opt})}{1+\epsilon'} \rfloor$ and $W(\mathbf{x}) \leq W(\mathbf{x}_{opt})$. By setting carefully the value of ϵ' as a function of ϵ , the solution yields $F(\mathbf{x}) \leq (1 + \epsilon)F(\mathbf{x}_{opt})$. The detail is similar to the one presented in [15] and is omitted here.

6 Simulations

Simulations are conducted to corroborate the theoretical study and to evaluate the performance improvement that partial verifications can provide. Unlike with a single detector type, the number of different patterns to consider, while searching for the optimal one, grows exponentially with multiple-type verifications (e.g., there are 64 different patterns with three partial verifications of four types). Since Corollary 1 shows that ordering of partial verifications does not matter, only unique patterns are considered to find the optimal one. The optimal algorithms for single-type as well as multiple-type detectors are implemented in Maple, and simulations are conducted with realistic detector values. The expected overhead of the algorithms employing partial verifications is compared with the overhead of the baseline algorithm, which uses only a single guaranteed verification. The reduction in overhead illustrates the usefulness of partial verifications. Section 6.1 describes the experimental setup, including the scenarios and the range of parameter values. Results are presented in Section 6.2 by exploring two problem instances.

6.1 Simulation framework

This section provides information about the parameters used for instantiating the performance model. We have chosen realistic parameters that depict a typical expected exascale platform. The target platform consists of 10^5 nodes whose individual MTBF is 100 years, which amounts to a platform MTBF of $\mu = 31,536$ seconds (i.e., about 8.7 hours). The global size of the memory for an exascale machine is expected to be between 32 PB and 64 PB; divided by the number of nodes (10^5), the memory size per node goes from 320 GB to 640 GB. Most HPC applications try to populate 90% of the node memory but only between 10% – 50% of the node memory is checkpointed. That makes checkpoints between 30 GB to 300 GB. At exascale, most checkpoints will be done in local non-volatile memory

(NVM) which is known to be slower than DRAM. We assume checkpoint throughput between 0.5 GB/s and 1 GB/s.

Concerning the SDC detectors, we assume a first detector $D^{(1)}$ with a throughput of about 200 MB/s/process and a recall of 0.5 [2, 4]. The second one $D^{(2)}$ has a throughput of about 20 MB/s/process and has a recall of 0.95 [8]. If we assume 512 processes per node at exascale, then the node throughput of the detectors becomes 100 GB/s for $D^{(1)}$ and 10 GB/s for $D^{(2)}$. Finally, we assume a third detector $D^{(3)}$ which is an optimized version that combines features of the first two detectors, achieving a recall of 0.8 and a throughput of 50 GB/s. Concerning the perfect detector D^* , we assume a throughput of 0.5 GB/s based on the fact that usually application-specific detectors are based on physical properties such as mass conservation or energy conservation, which requires global communications and it is more expensive than purely local checks.

In the first instance, we evaluate our model with the three detectors $D^{(1)}$, $D^{(2)}$ and $D^{(3)}$ (with respective costs and recall values $V^{(1)} = 3$, $V^{(2)} = 30$, $V^{(3)} = 6$ and $r^{(1)} = 0.5$, $r^{(2)} = 0.95$, $r^{(3)} = 0.8$) by computing expected overhead for each detector as a function of number of partial verifications (m from 0 to up to 35). Also, the checkpointing cost and the perfect detector cost (with recall $r^* = 1$) are fixed at $C = V^* = 600$ seconds. We then assess the performance of a detector and its effect on execution overhead.

The second instance targets applications with various datasets that expose a change in detection recall [4, 3]. Therefore, it assumes a range for the recall of each detector rather than a single value: $r^{(1)} = [0.5, 0.9]$, $r^{(2)} = [0.75, 0.95]$, and $r^{(3)} = [0.8, 0.99]$. Given a dataset, we obtain a value of recall for each detector, within the interval. This is because different datasets might expose different levels of entropy and therefore the detectors might expose different prediction accuracy, hence different recall. We note that although the recall might be different for these datasets, the work done, hence the detection cost, is the same. Then, we compare the greedy algorithm to the optimal solution for different datasets, keeping the same values for C , V^* and μ .

6.2 Results and analysis

Based on the above framework, we report the simulation results highlighting the behavior of the overhead, and the improvements achieved by employing partial verifications.

6.2.1 Performance of detectors and impact of m

Figure 2 shows the performance of three different detectors $D^{(1)}$, $D^{(2)}$ and $D^{(3)}$, and assesses the impact of the number m of partial verifications on the execution overhead. The plot also shows the overhead of the baseline algorithm, which is approximately 39%, represented by $m = 0$. The expected overhead is reduced for all three detectors by em-




Figure 2: Expected overhead against partial verifications when a single type of detector is employed for three different detectors (with their respective recall values) and $C = 600$, $V^* = 600$, $V^{(1)} = 3$, $V^{(2)} = 30$, and $V^{(3)} = 6$, and $\mu = 100y/10^5$.

ploying partial verifications. For each detector, the optimal overhead is attained for a particular value of m after which it starts rising again due to the fact that forcing too many verifications will eventually lead the error-free overhead to rise. The improvement in overhead over the baseline algorithm is 9% for detectors $D^{(1)}$ and $D^{(3)}$ (optimal overhead for both is $\approx 30\%$), and 7% for detector $D^{(2)}$ (optimal overhead is $\approx 32\%$). Concerning performance of detectors, we can see that $D^{(1)}$ and $D^{(3)}$ are slightly better than $D^{(2)}$, due to their high accuracy-to-cost ratio. However, for $m \leq 2$, $D^{(2)}$ is better due to its highest recall value among the three detectors, while its performance degrades as more $D^{(2)}$ detectors are employed due to its high cost.

6.2.2 Performance of the greedy algorithm

In this second instance, the recall depends on the dataset and its value is within the detector's recall interval. Even

Table 1: Performance comparison of greedy algorithm and optimal solution. In all scenarios, $C = V^* = 600$, $V^{(1)} = 3$, $V^{(3)} = 6$.

	m	overhead H	diff. from opt.
Scenario 1: $r^{(1)} = 0.51$, $r^{(3)} = 0.82$, $\phi^{(1)} \approx 137$, $\phi^{(3)} \approx 139$			
Optimal solution	(1, 15)	29.828%	0%
Greedy with $D^{(3)}$	(0, 16)	29.829%	0.001%
Scenario 2: $r^{(1)} = 0.58$, $r^{(3)} = 0.9$, $\phi^{(1)} \approx 163$, $\phi^{(3)} \approx 164$			
Optimal solution	(1, 14)	29.659%	0%
Greedy with $D^{(3)}$	(0, 15)	29.661%	0.002%
Scenario 3: $r^{(1)} = 0.64$, $r^{(3)} = 0.97$, $\phi^{(1)} \approx 188$, $\phi^{(3)} \approx 188$			
Optimal solution	(1, 13)	29.523%	0%
Greedy with $D^{(1)}$	(27, 0)	29.524%	0.001%
Greedy with $D^{(3)}$	(0, 14)	29.525%	0.002%

with the intervals, because of their high cost, $D^{(2)}$ detectors always have a lower accuracy-to-cost ratio when compared to $D^{(1)}$ and $D^{(3)}$, which share similar ratios. We show that the greedy algorithm presented in Section 5.3 performs extremely well in practice, even though the optimal pattern may employ both $D^{(1)}$ and $D^{(3)}$ in the solution.

Indeed, we identify three scenarios where a combination of $D^{(1)}$ and $D^{(3)}$ constitutes the optimal pattern, as presented in Table 1. In all three scenarios, the greedy algorithm, which uses only the detector with the highest accuracy-to-cost ratio, performs within 0.002% of the optimal solution.

7 Conclusion and Future Work

In this paper, we provided a comprehensive analysis of a computing pattern that employs different types of partial verifications for detecting silent errors in HPC applications. We showed that the optimization problem is NP-complete in general, and proposed a greedy algorithm as well as an FPTAS for choosing the number of detectors to be used, as well as their types and locations in the pattern. Simulations based on realistic detector settings show that the greedy algorithm works well in practice, and confirm the usefulness of these detectors in tackling SDC on exascale systems.

In future work, we will investigate the impact of false positives/alarms on the performance of partial verification patterns. False positives are characterized by the *precision* value, which is the number of true errors over the total number of errors raised by a detector. In many partial detectors, a tradeoff exists between the recall and precision by adjusting the detection parameters. Analyzing a pattern in the presence of both false positives and false negatives will refine the assessment of the usefulness of partial detectors.

References

- [1] G. Aupy, A. Benoit, T. Hérault, Y. Robert, F. Vivien, and D. Zaidouni. On the combination of silent error detection and checkpointing. In *Proceedings of the 19th IEEE Pacific Rim International Symposium on Dependable Computing (PRDC)*, pages 11–20, 2013.
- [2] L. Bautista Gomez and F. Cappello. Detecting silent data corruption through data dynamic monitoring for scientific applications. In *Proceedings of the 19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP '14*, pages 381–382, New York, NY, USA, 2014. ACM.
- [3] L. Bautista Gomez and F. Cappello. Detecting and correcting data corruption in stencil applications through multivariate interpolation. In *Proceedings of the 1st International Workshop on Fault Tolerant Systems, FTS'15*. IEEE, 2015 (To appear).

- [4] L. Bautista Gomez and F. Cappello. Exploiting Spatial Smoothness in HPC Applications to Detect Silent Data Corruption. In *Proceedings of the 17th IEEE International Conference on High Performance Computing and Communications, HPCC'15*. IEEE, 2015 (To appear).
- [5] A. Benoit, A. Cavelan, Y. Robert, and H. Sun. Assessing general-purpose algorithms to cope with fail-stop and silent errors. In *Proceedings of the 5th International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, pages 215–236, 2014.
- [6] A. Benoit, Y. Robert, and S. K. Raina. Efficient checkpoint/verification patterns. *ICL Research report RR-1403. To appear in IJHPCA (2015)*, 2014.
- [7] A. R. Benson, S. Schmit, and R. Schreiber. Silent error detection in numerical time-stepping schemes. *Int. J. High Performance Computing Applications*, DOI: 10.1177/1094342014532297, 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 *Proceedings of The ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC)*, HPDC '15, 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 *Proceedings of the International Conference on Supercomputing (ICS)*, pages 155–164, 2008.
- [11] A. Cavelan, S. K. Raina, Y. Robert, and H. Sun. Assessing the impact of partial verifications against silent data corruptions. In *Proceedings of the 44th Annual International Conference on Parallel Processing (ICPP)*, 2015.
- [12] K. M. Chandy and L. Lamport. Distributed snapshots: Determining global states of distributed systems. *ACM Transactions on Computer Systems*, 3(1):63–75, 1985.
- [13] Z. Chen. Online-ABFT: An online algorithm based fault tolerance scheme for soft error detection in iterative methods. In *Proceedings of the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*, pages 167–176, 2013.
- [14] E. Ciocca, I. Koren, Z. Koren, C. M. Krishna, and D. S. Katz. Application-level fault tolerance in the orbital thermal imaging spectrometer. In *Proceedings of the 10th IEEE Pacific Rim International Symposium on Dependable Computing (PRDC'04)*, pages 43–48, 2004.
- [15] R. Cohen and L. Katzir. The generalized maximum coverage problem. *Inf. Process. Lett.*, 108(1):15–22, 2008.
- [16] J. T. Daly. A higher order estimate of the optimum checkpoint interval for restart dumps. *Future Generation Comp. Syst.*, 22(3):303–312, 2006.
- [17] J. Dongarra, P. Beckman, P. Aerts, F. Cappello, T. Lippert, S. Matsuoka, P. Messina, T. Moore, R. Stevens, A. Trefethen, and M. Valero. The international exascale software project: a call to cooperative action by the global high-performance community. *IJHPCA*, 23(4):309–322, 2009.
- [18] M. Dow. Explicit inverses of Toeplitz and associated matrices. *ANZIAM J.*, 44(E):E185–E215, Jan. 2003.
- [19] J. Elliott, K. Kharbas, D. Fiala, F. Mueller, K. Ferreira, and C. Engelmann. Combining partial redundancy and checkpointing for HPC. In *Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS)*, pages 615–626, 2012.
- [20] E. N. M. Elnozahy, L. Alvisi, Y.-M. Wang, and D. B. Johnson. A survey of rollback-recovery protocols in message-passing systems. *ACM Computing Survey*, 34:375–408, 2002.
- [21] 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 *Proc. SC'12*, page 78, 2012.
- [22] M. R. Garey and D. S. Johnson. *Computers and Intractability, a Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.
- [23] M. Heroux and M. Hoemmen. Fault-tolerant iterative methods via selective reliability. Research report SAND2011-3915 C, Sandia National Laboratories, 2011.
- [24] K.-H. Huang and J. A. Abraham. Algorithm-based fault tolerance for matrix operations. *IEEE Trans. Comput.*, 33(6):518–528, 1984.
- [25] H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack Problems*. Springer, 2004.
- [26] G. Lu, Z. Zheng, and A. A. Chien. When is multi-version checkpointing needed? In *Proc. 3rd Workshop on Fault-tolerance for HPC at extreme scale (FTXS)*, pages 49–56, 2013.
- [27] 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.

- [28] A. Moody, G. Bronevetsky, K. Mohror, and B. R. d. Supinski. Design, Modeling, and Evaluation of a Scalable Multi-level Checkpointing System. In *Proc. of the ACM/IEEE SC Conf.*, pages 1–11, 2010.
- [29] X. Ni, E. Meneses, N. Jain, and L. V. Kalé. ACR: Automatic Checkpoint/Restart for Soft and Hard Error Protection. In *Proc. SC'13*. ACM, 2013.
- [30] 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.
- [31] P. Sao and R. Vuduc. Self-stabilizing iterative solvers. In *Proceedings of the Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems (ScalA)*, 2013.
- [32] M. Shantharam, S. Srinivasmurthy, and P. Raghavan. Fault tolerant preconditioned conjugate gradient for sparse linear system solution. In *Proceedings of the ACM International Conference on Supercomputing (ICS)*, pages 69–78, 2012.
- [33] J. W. Young. A first order approximation to the optimum checkpoint interval. *Comm. of the ACM*, 17(9):530–531, 1974.
- [34] 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