

Assessing General-Purpose Algorithms to Cope with Fail-Stop and Silent Errors

Anne Benoit, Aurélien Cavelan, Yves Robert, Hongyang Sun

► **To cite this version:**

Anne Benoit, Aurélien Cavelan, Yves Robert, Hongyang Sun. Assessing General-Purpose Algorithms to Cope with Fail-Stop and Silent Errors. ACM Transactions on Parallel Computing, Association for Computing Machinery, 2016, 3 (2), pp.1-36. 10.1145/2897189 . hal-01358146

HAL Id: hal-01358146

<https://hal.inria.fr/hal-01358146>

Submitted on 31 Aug 2016

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.

Assessing general-purpose algorithms to cope with fail-stop and silent errors

Anne Benoit, École Normale Supérieure de Lyon, CNRS & INRIA, France

Aurélien Cavelan, École Normale Supérieure de Lyon, CNRS & INRIA, France

Yves Robert, École Normale Supérieure de Lyon, CNRS & INRIA, France, and University of Tennessee Knoxville, USA

Hongyang Sun, École Normale Supérieure de Lyon, CNRS & INRIA, France

In this paper, we combine the traditional checkpointing and rollback recovery strategies with verification mechanisms to address both fail-stop and silent errors. The objective is to minimize either makespan or energy consumption. While DVFS is a popular approach for reducing the energy consumption, using lower speeds/voltages can increase the number of errors, thereby complicating the problem. We consider an application workflow whose dependence graph is a chain of tasks, and we study three execution scenarios: (i) a single speed is used during the whole execution; (ii) a second, possibly higher speed is used for any potential re-execution; (iii) different pairs of speeds can be used throughout the execution. For each scenario, we determine the optimal checkpointing and verification locations (and the optimal speeds for the third scenario) to minimize either objective. The different execution scenarios are then assessed and compared through an extensive set of experiments.

General Terms: Algorithms, Design, Experimentation

Additional Key Words and Phrases: HPC, resilience, checkpoint, verification, failures, fail-stop error, silent data corruption, silent error

1. INTRODUCTION

For HPC applications, scale is a major opportunity. Massive parallelism with 100,000+ nodes is the most viable path to achieving sustained petascale performance. Future platforms will enrol even more computing resources to enter the exascale era.

Unfortunately, scale is also a major threat. Resilience is the first challenge. Even if each node provides an individual MTBF (Mean Time Between Failures) of, say, one century, a machine with 100,000 such nodes will encounter a failure every 9 hours in average, which is smaller than the execution time of many HPC applications. Furthermore, a one-century MTBF per node is an optimistic figure, given that each node is composed of several hundreds of cores. Worse, several types of errors need to be considered when computing at scale. In addition to classical fail-stop errors (such as hardware failures), silent errors (a.k.a silent data corruptions) constitute another threat that cannot be ignored any longer [O’Gorman 1994; Ziegler et al. 1996a; Ziegler et al. 1996b; Ziegler et al. 1998; Moody et al. 2010].

Another challenge is energy consumption. The power requirement of current petascale platforms is that of a small town, hence measures must be taken to reduce the energy consumption of future platforms. A widely-used strategy is to use DVFS techniques: modern processors can run at different speeds, and lower speeds induce big savings in energy consumption. In a nutshell, this is because the dynamic power consumed when computing at speed s is proportional to s^3 , while execution time is proportional to $1/s$. As a result, computing energy (which is time times power) is proportional to s^2 . However, static power must be accounted for, and it is paid throughout the duration of the execution, which calls for a shorter execution (at higher speed). Overall there are tradeoffs to be found, but in most practical settings, using lower speeds reduces global energy consumption.

Author’s addresses: Anne Benoit, Aurélien Cavelan, Yves Robert and Hongyang Sun, École Normale Supérieure de Lyon, 46 allée d’Italie, 69364 Lyon Cedex 07, France.

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

To further complicate the picture, energy savings have an impact on resilience. Obviously, the longer the execution, the higher the expected number of errors, hence using a lower speed to save energy may well induce extra time and overhead to cope with more errors throughout execution. Even worse (again!), lower speeds are usually obtained via lower voltages, which themselves induce higher error rates and further increase the latter overhead.

In this paper, we introduce a model that addresses both challenges: resilience and energy-consumption. In addition, we address both fail-stop and silent errors, which, to the best of our knowledge, has only been achieved before through costly replication techniques [Ni et al. 2013]. While checkpoint/restart [Chandy and Lamport 1985; Elnozahy et al. 2002] is the de-facto recovery technique for dealing with fail-stop errors, there is no widely adopted general-purpose technique to cope with silent errors. The problem with silent errors is *detection latency*: contrarily to a fail-stop error whose detection is immediate, a silent error is identified only when the corrupted data is activated and/or leads to an unusual application behavior. However, checkpoint and rollback recovery assumes instantaneous error detection, and this raises a new difficulty: if the error stroke before the last checkpoint, and is detected after that checkpoint, then the checkpoint is corrupted and cannot be used to restore the application. To solve this problem, one may envision to keep several checkpoints in memory, and to restore the application from the last *valid* checkpoint, thereby rolling back to the last *correct* state of the application [Lu et al. 2013]. This multiple-checkpoint approach has three major drawbacks. First, it is very demanding in terms of stable storage: each checkpoint typically represents a copy of the entire memory footprint of the application, which may well correspond to several terabytes. The second drawback is the possibility of fatal failures. Indeed, if we keep k checkpoints in memory, the approach assumes that the error that is currently detected did not strike before all the checkpoints still kept in memory, which would be fatal: in that latter case, all live checkpoints are corrupted, and one would have to re-execute the entire application from scratch. The probability of a fatal failure is evaluated in [Aupy et al. 2013] for various error distribution laws and values of k . The third drawback of the approach is the most serious, and applies even without memory constraints, i.e., if we could store an infinite number of checkpoints in storage. The critical question is to determine which checkpoint is the last valid one. We need this information to safely recover from that point on. However, because of the detection latency, we do not know when the silent error has indeed occurred, hence we cannot identify the last valid checkpoint, unless some verification system is enforced.

We introduce such a verification system in this paper. This approach is agnostic of the nature of this verification mechanism (checksum, error correcting code, coherence tests, etc.). It is also fully general-purpose, although application-specific information, if available, can always be used to decrease the cost of verification: see the overview of related work in Section 2 for examples. In this context, the simplest protocol is to take only verified checkpoints (VC). This corresponds to performing a verification just before taking each checkpoint. If the verification succeeds, then one can safely store the checkpoint. If the verification fails, then a silent 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, if a fail-stop error strikes, we also safely recover from the last checkpoint, just as in the classical checkpoint and rollback recovery method. This VC-ONLY protocol basically amounts to replacing the cost C of a checkpoint by the cost $V + C$ of a verification followed by a checkpoint. However, because we deal with two sources of errors, one detected immediately and the other only when we reach the verification, the analysis of the optimal strategy is more involved. We extend both the classical bound by Young [Young 1974] or Daly [Daly

2006], and the dynamic programming algorithm of Toueg and Babaoglu [Toueg and Babaoglu 1984], to deal with these error sources.

While taking checkpoints without verifications seems a bad idea (because of the memory cost, and of the risk of saving corrupted data), taking a verification without checkpointing may be interesting. Indeed, if silent errors are frequent enough, it is worth verifying the data in between two (verified) checkpoints, so as to detect a possible silent error earlier in the execution, and thereby re-executing less work. We refer to this protocol as VC+V that allows for both verified checkpoints and isolated verifications.

One major objective of this paper is to study VC+V algorithms coupling verification and checkpointing, and to analytically determine the best balance of verifications between checkpoints so as to minimize either makespan (total execution time) or energy consumption. To achieve this ambitious goal, we restrict to a simplified, yet realistic, application framework. We consider application workflows that consist of a number of parallel tasks that execute on the platform, and that exchange data at the end of their executions. In other words, the task graph is a linear chain, and each task (except maybe the first one and the last one) reads data from its predecessor and produces data for its successor. This scenario corresponds to a high-performance computing application whose workflow is partitioned into a succession of (typically large) tightly-coupled computational kernels, each of them being identified as a task by the model. At the end of each task, we have the opportunity either to perform a verification of the task output or to perform a verification followed by a checkpoint.

In addition, we have to select a speed for each execution of each task. We envision three different execution scenarios. In the simple `SINGLE SPEED` scenario, a unique speed s is available throughout execution. In the intermediate `REEXEC SPEED` scenario, the same speed s is used for the first execution of each task, but another speed σ is available for re-execution after a fail-stop or silent error. Here, the first speed s can be considered as the regular speed, while the second speed σ corresponds to an adjusted speed to either speed up or to slow down the re-execution after an error strikes, depending on the optimization objective. Finally, in the advanced `MULTI SPEED` scenario, two different speeds s_i and σ_i can be used to execute the tasks in between two consecutive checkpoints (which we call a task segment). Each speed s_i or σ_i can be freely chosen among a set of K discrete speeds. Note that these speeds may well vary from one segment to another. For each execution scenario, we provide an optimal dynamic programming algorithm to determine the best locations of checkpoints and verifications (and for the `MULTI SPEED` scenario we also provide the corresponding optimal pair of speeds for each segment).

The main contributions of this paper are the following:

- We introduce a general-purpose model to deal with both fail-stop and silent errors, combining checkpoints with a verification mechanism.
- We consider three execution scenarios, first with a single speed, then in case of re-execution, and finally with several discrete speeds that can freely change after each checkpoint.
- For the three execution scenarios and for both makespan and energy objectives, we consider two approaches, one using verified checkpoints only and the other one using additional isolated verifications. We provide a dynamic programming algorithm that determines the best locations of checkpoints and verifications across a chain of tasks for each scenario/approach/objective combination.
- For both makespan and energy objectives, we derive the optimal checkpointing and verification periods for the divisible load application model, where checkpoints and verifications can be placed at any point in execution of the application.

— We conduct an extensive set of simulations to support the theory and to assess the performance of each algorithm.

The rest of the paper is organized as follows. Section 2 provides an overview of related work. Section 3 is devoted to formally defining the framework and all model parameters. The next three sections deal with the main algorithmic contributions: for all three execution scenarios, we design optimal algorithms for the VC-ONLY approach, and then for the VC+V approach, targeting either time or energy minimization. Then in Section 7, we report on a comprehensive set of experiments to assess the impact of each scenario and approach. Finally, we outline main conclusions and directions for future work in Section 8.

2. RELATED WORK

2.1. Fail-stop errors

The de-facto general-purpose error recovery technique in high performance computing is checkpoint and rollback recovery [Chandy and Lamport 1985; Elnozahy et al. 2002]. 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 back to one of its former states. There are several families of checkpointing protocols, but they share a common feature: each checkpoint forms a consistent recovery line, i.e., when an error is detected, one can rollback to the last checkpoint and resume execution, after a downtime and a recovery time.

Many models are available to understand the behavior of checkpoint/restart [Young 1974; Daly 2006; Ozaki et al. 2006; Bougeret et al. 2011]. For a divisible load application where checkpoints can be inserted at any point in execution for a nominal cost C , there exist well-known formulas due to Young [Young 1974] and Daly [Daly 2006] to determine the optimal checkpointing period. For an application composed of a linear chain of tasks, which is also the subject of this paper, the problem of finding the optimal checkpoint strategy, i.e., of determining which tasks to checkpoint, in order to minimize the expected execution time, has been solved by Toueg and Babaoglu [Toueg and Babaoglu 1984], using a dynamic programming algorithm.

One major contribution of this paper is to extend both the Young/Daly formulas [Young 1974; Daly 2006] and the result of Toueg and Babaoglu [Toueg and Babaoglu 1984] to deal with silent errors in addition to fail-stop errors, and with several discrete speeds instead of a single one.

2.2. Silent errors

Most traditional approaches maintain a single checkpoint. If the checkpoint file includes errors, 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 failures) and are unable to accommodate silent errors. We focus in this section on related work about silent errors. A comprehensive list of techniques and references is provided by Lu, Zheng and Chien [Lu et al. 2013].

Considerable efforts have been directed at error-checking to reveal silent errors. Error detection is usually very costly. Hardware mechanisms, such as ECC memory, can detect and even correct a fraction of errors, but in practice they are complemented with software techniques. The simplest technique is triple modular redundancy and voting [Lyons and Vanderkulk 1962], which induces a highly costly verification. For high-performance scientific applications, process replication (each process is equipped with a replica, and messages are quadruplicated) is proposed in the RedMPI library [Fiala et al. 2012]. Elliot et al. [Elliott et al. 2012] combine partial redundancy and checkpointing, and confirm the benefit of dual and triple redundancy. The drawback is that

twice the number of processing resources is required (for dual redundancy). As already mentioned, an approach based on checkpointing and replication is proposed in [Ni et al. 2013], in order to detect and enable fast recovery of applications from both silent errors and hard errors.

Application-specific information can be very useful to enable ad-hoc solutions, which dramatically decrease the cost of detection. Many techniques have been advocated. They include memory scrubbing [Hwang et al. 2012] and ABFT techniques [Huang and Abraham 1984; Bosilca et al. 2009; Shantharam et al. 2012], such as coding for the sparse-matrix vector multiplication kernel [Shantharam et al. 2012], and coupling a higher-order with a lower-order scheme for PDEs [Benson et al. 2014]. These methods can only detect an error but do not correct it. Self-stabilizing corrections after error detection in the conjugate gradient method are investigated by Sao and Vuduc [Sao and Vuduc 2013]. Heroux and Hoemmen [Heroux and Hoemmen 2011] design a fault-tolerant GMRES capable of converging despite silent errors. Bronevetsky and de Supinski [Bronevetsky and de Supinski 2008] provide a comparative study of detection costs for iterative methods.

A nice instantiation of the checkpoint and verification mechanism that we study in this paper is provided by Chen [Chen 2013], who deals with sparse iterative solvers. Consider a simple method such as the PCG, the Preconditioned Conjugate Gradient method: Chen’s approach performs a periodic verification every d iterations, and a periodic checkpoint every $d \times c$ iterations, which is a particular case of the VC+V approach with equi-distance verifications. For PCG, the verification amounts to checking the orthogonality of two vectors and to recomputing and checking the residual. The cost of the verification is small in front of the cost of an iteration, especially when the preconditioner requires much more flops than a sparse matrix-vector product.

As already mentioned, our work is agnostic of the underlying error-detection technique and takes the cost of verification as an input parameter to the model.

2.3. Energy model and error rate

Modern processors are equipped with *dynamic voltage and frequency scaling* (DVFS) capability. The total power consumption is the sum of the static/idle power and the dynamic power, which is proportional to the cube of the processing speed s [Yao et al. 1995; Bansal et al. 2007], i.e., $P(s) = P_{idle} + \beta \cdot s^3$, where $\beta > 0$. A widely used reliability model assumes that radiation-induced transient faults (soft errors) follow a Poisson process with an average arrival rate λ . The impact of DVFS on the error rate is, however, not completely clear.

On the one hand, lowering the voltage/frequency is believed to have an adverse effect on the system reliability [Dixit and Wood 2011; Zhu et al. 2004]. In particular, many papers (e.g., [Zhu et al. 2004; Zhao et al. 2008; Aupy et al. 2012; Das et al. 2014]) have assumed the following exponential error rate model:

$$\lambda(s) = \lambda_0 \cdot 10^{\frac{d(s_{\max} - s)}{s_{\max} - s_{\min}}}, \quad (1)$$

where λ_0 denotes the average error rate at the maximum speed s_{\max} , $d > 0$ is a constant indicating the sensitivity of error rate to voltage/frequency scaling, and s_{\min} is the minimum speed. This model suggests that the error rate increases exponentially with decreased processing speed, which is a result of decreasing the voltage/frequency and hence lowering the circuit’s critical charge (i.e., the minimum charge required to cause an error in the circuit).

On the other hand, the failure rates of computing nodes have also been observed to increase with temperature [Patterson 2008; Feng 2003; Hsu and Feng 2005; Sa-rood et al. 2013], which generally increases together with the processing speed (voltage/frequency). As a rule of thumb, Arrhenius’ equation when applied to microelec-

Table I. List of main notations.

Tasks	
$\{T_1, T_2, \dots, T_n\}$	Set of n tasks
w_i	Computational cost of task T_i
Speeds	
s	Regular speed
σ	Re-execution speed
$\{s_1, s_2, \dots, s_K\}$	Set of K discrete computing speeds (DVFS)
Time	
$T_{i,j}(s)$	Time needed to execute tasks T_i to T_j at speed s
$V_i(s)$	Time needed to verify task T_i at speed s
C_i	Time needed to checkpoint task T_i
R_i	Time needed to recover from task T_i
Resilience	
$\lambda^F(s)$	Fail-stop error rate for a given speed s
$\lambda^S(s)$	Silent error rate for a given speed s
$p_{i,j}^F(s)$	Probability that a fail-stop error strikes between tasks T_i and T_j
$p_{i,j}^S(s)$	Probability that a silent error strikes between tasks T_i and T_j
Energy	
P_{idle}	Static/idle power dissipated when the platform is switched on
$P_{cpu}(s)$	Dynamic power spent by operating the CPU at speed s
P_{io}	Dynamic power spent by I/O transfers (checkpoints and recoveries)
$E_{i,j}(s)$	Energy needed to execute tasks T_i to T_j at speed s
$E_i^V(s)$	Energy needed to verify task T_i at speed s
E_i^C	Energy needed to checkpoint task T_i
E_i^R	Energy needed to recover from task T_i

tronic devices suggests that the error rate doubles for every 10°C increase in the temperature [Feng 2003]. In general, the mean time between failure (MTBF) of a processor, which is the reciprocal of failure rate, can be expressed as [Sarood et al. 2013]:

$$MTBF = \frac{1}{\lambda} = A \cdot e^{-b \cdot T},$$

where A and b are thermal constants, and T denotes the temperature of the processor. Under the reasonable assumption that higher operating voltage/frequency leads to higher temperature, this model suggests that the error rate increases with increased processing speed.

Clearly, the two models above draw contradictory conclusions on the impact of DVFS on error rates. In practice, the impact of the first model may be more evident, as the temperature dependency in some systems has been observed to be linear (or even not exist) instead of being exponential [El-Sayed et al. 2012]. Generally speaking, the processing speed should have a composite effect on the average error rate by taking both voltage level and temperature into account. In the experimental section of this paper (Section 7), we adopt a tradeoff model and modify Equation (1) to include the impact of temperature.

3. FRAMEWORK

In this section we introduce all model parameters. For reference, main notations are summarized in Table I. We start with a description of the application workflows. Then we present parameters related to energy consumption. Next we detail the resilient model to deal with fail-stop and silent errors. We conclude by presenting the various execution scenarios.

3.1. Application workflows

We consider application workflows whose task graph is a linear chain $T_1 \rightarrow T_2 \cdots \rightarrow T_n$. Here n is the number of tasks, and each task T_i is weighted by its computational cost w_i . We target a platform with p identical processors. Each task is a parallel task that is executed on the whole platform. A fundamental characteristic of the application model is that it allows to view the platform as a single (albeit very powerful) *macro-processor*, thereby providing a tractable abstraction of the problem.

3.2. Energy consumption

When computing (including verification), we use DVFS to change the speed of the processors, and assume a set $S = \{s_1, s_2, \dots, s_K\}$ of K discrete computing speeds. During checkpointing and recovery, we assume a dedicated (constant) power consumption. Altogether, the total power consumption of the macro-processor is p times the power consumption of each individual resource. It is decomposed into three different components:

- P_{idle} , the static power dissipated when the platform is on (even idle);
- $P_{cpu}(s)$, the dynamic power spent by operating the CPU at speed s ;
- P_{io} , the dynamic power spent by I/O transfers (checkpoints and recoveries).

Assume w.l.o.g. that there is no overlap between CPU operations and I/O transfers. Then the total energy consumed during the execution of the application can be expressed as

$$Energy = P_{idle}(T_{cpu} + T_{io}) + \sum_{i=1}^K P_{cpu}(s_i)T_{cpu}(s_i) + P_{io}T_{io}$$

where $T_{cpu}(s_i)$ is the time spent on computing at speed s_i , $T_{cpu} = \sum_{i=1}^K T_{cpu}(s_i)$ is the total time spent on computing, and T_{io} is the total time spent on I/O transfers.

The time to compute tasks T_i to T_j at speed s is $T_{i,j}(s) = \frac{1}{s} \sum_{k=i}^j w_k$ and the corresponding energy is $E_{i,j}(s) = T_{i,j}(s)(P_{idle} + P_{cpu}(s))$.

3.3. Resilience

We assume that errors only strike during computations, and not during I/O transfers (checkpoints and recoveries) nor verifications. We consider two types of errors: *fail-stop* and *silent*.

To cope with fail-stop errors, we use checkpointing, and to cope with silent errors, an additional verification mechanism is used. The time to checkpoint (the output of) task T_i is C_i , the time to recover from (the checkpoint of) task T_i is R_i , and the time to verify (the output of) task T_i at speed s is $V_i(s)$. We assume that both fail-stop errors and silent errors follow an exponential distribution with average rates $\lambda^F(s)$ and $\lambda^S(s)$, respectively, where s denotes the current computing speed. Given an error rate λ , let $p(\lambda, L) = 1 - e^{-\lambda L}$ denote the probability that an error strikes during an execution of length L . For convenience, we define $p_{i,j}^F(s) = p(\lambda^F(s), T_{i,j}(s))$ to be the probability that a fail-stop error strikes when executing from T_i to T_j , and define $p_{i,j}^S(s) = p(\lambda^S(s), T_{i,j}(s))$ similarly for silent errors.

Resilience also has a cost in terms of energy consumption. Specifically, the energy to checkpoint task T_i is $E_i^C = C_i(P_{idle} + P_{io})$, to recover from task T_i is $E_i^R = R_i(P_{idle} + P_{io})$, and to verify task T_i at speed s is $E_i^V(s) = V_i(s)(P_{idle} + P_{cpu}(s))$.

3.4. Execution scenarios

We consider three different execution scenarios:

SINGLE SPEED A single speed s is used during the whole execution ($K = 1$).

REEXECSPEED There are two speeds, s for the first execution of each task, and σ for any potential re-execution ($K = 2$).

MULTISPEED We are given K discrete speeds, where K is arbitrary. The workflow chain is cut into subchains called segments and delimited by checkpoints. For each of these segments, we can freely choose the speed of the first execution, and the (possibly different) speed for any ulterior execution, among the K speeds. Note that these speeds may well vary from one segment to another.

3.5. Optimization problems

For each execution scenario, we deal with four problems:

TIME-VC-ONLY Minimize total execution time (or makespan) using the VC-ONLY approach.

TIME-VC+V Minimize total execution time (or makespan) using the VC+V approach.

ENERGY-VC-ONLY Minimize total energy consumption using the VC-ONLY approach.

ENERGY-VC+V Minimize total energy consumption using the VC+V approach.

For the **SINGLESPEED** and **REEXECSPEED** scenarios, we have to decide the optimal locations of the checkpoints (VC-ONLY) and of the verifications (VC+V). For the **MULTISPEED** scenario, we further have to select a pair of speeds (first execution and re-execution) for each segment.

4. SINGLESPEED SCENARIO

In this scenario, we are given a single processing speed s . We investigate the VC-ONLY and VC+V approaches. For each approach, we present an optimal polynomial-time dynamic programming algorithm.

4.1. VC-ONLY: Using verified checkpoints only

In this approach, we only place verified checkpoints. We aimed at finding the best positions for checkpoints in order to minimize the total execution time (TIME-VC-ONLY) or the total energy consumption (ENERGY-VC-ONLY).

4.1.1. TIME-VC-ONLY: MINIMIZING MAKESPAN

THEOREM 4.1. *For the SINGLESPEED scenario, the TIME-VC-ONLY problem can be solved by a dynamic programming algorithm in $O(n^2)$ time.*

PROOF. We define $Time_C(j, s)$ to be the optimal expected time to successfully execute tasks T_1, \dots, T_j , where T_j has a verified checkpoint, and there are possibly other verified checkpoints from T_1 to T_{j-1} . Note that we always verify and checkpoint the last task T_n to save the final result, so the goal is to find $Time_C(n, s)$.

To compute $Time_C(j, s)$, we formulate the following dynamic program by trying all possible locations for the last checkpoint before T_j (see Figure 1):

$$Time_C(j, s) = \min_{0 \leq i < j} \{Time_C(i, s) + T_C(i+1, j, s)\} + C_j,$$

where $T_C(i, j, s)$ is the expected time to successfully execute the tasks T_i to T_j , provided that T_{i-1} and T_j are both verified and checkpointed, and while no other task in between is verified nor checkpointed. Note that we also account for the checkpointing cost C_j for task T_j , which is not included in the definition of T_C . To initialize the dynamic program, we define $Time_C(0, s) = 0$.




Fig. 1. Illustration of the dynamic programming formulation for $Time_C(j, s)$.

In the following, we show how to compute $T_C(i, j, s)$ for each (i, j) pair with $i \leq j$. We start by considering only *silent errors* and use the notation $T_C^S(i, j, s)$ for that purpose. Silent errors can occur at any time during the computation but we only detect them after all tasks have been executed. Thus, we always have to pay $T_{i,j}(s) + V_j(s)$, the time to execute from task T_i to T_j at speed s and to verify T_j . If the verification fails, which happens with probability $p_{i,j}^S(s)$, a silent error has occurred and we have to recover from T_{i-1} and start anew. For convenience, we assume that there is a virtual task T_0 that is always verified and checkpointed, with a recovery cost $R_0 = 0$. Mathematically, we can express $T_C^S(i, j, s)$ as

$$T_C^S(i, j, s) = T_{i,j}(s) + V_j(s) + p_{i,j}^S(s) (R_{i-1} + T_C^S(i, j, s)) ,$$

which simplifies to

$$T_C^S(i, j, s) = e^{\lambda^S(s)T_{i,j}(s)} (T_{i,j}(s) + V_j(s)) + (e^{\lambda^S(s)T_{i,j}(s)} - 1)R_{i-1} .$$

Things are different when accounting for *fail-stop errors*, because the application will stop immediately when a fail-stop error occurs, even in the middle of the computation. Let $T_{lost,i,j}(s)$ denote the expected time lost during the execution from T_i to T_j if a fail-stop error strikes, and it can be expressed as

$$T_{lost,i,j}(s) = \int_0^\infty x \mathbb{P}(X = x | X < T_{i,j}(s)) dx = \frac{1}{\mathbb{P}(X < T_{i,j}(s))} \int_0^{T_{i,j}(s)} x \lambda^F(s) e^{-\lambda^F(s)x} dx ,$$

where $\mathbb{P}(X = x)$ denotes the probability that a fail-stop error strikes at time x . By definition, we have $\mathbb{P}(X < T_{i,j}(s)) = 1 - e^{-\lambda^F(s)T_{i,j}(s)}$. Integrating by parts, we can get

$$T_{lost,i,j}(s) = \frac{1}{\lambda^F(s)} - \frac{T_{i,j}(s)}{e^{\lambda^F(s)T_{i,j}(s)} - 1} . \quad (2)$$

Therefore, the expected execution time $T_C^F(i, j, s)$ when considering only fail-stop errors is given by

$$T_C^F(i, j, s) = p_{i,j}^F(s) (T_{lost,i,j}(s) + R_{i-1} + T_C^F(i, j, s)) + (1 - p_{i,j}^F(s)) T_{i,j}(s) ,$$

which simplifies to

$$T_C^F(i, j, s) = (e^{\lambda^F(s)T_{i,j}(s)} - 1) \left(\frac{1}{\lambda^F(s)} + R_{i-1} \right) .$$

We now account for both *fail-stop and silent errors*, and use the notation $T_C^{SF}(i, j, s)$ for that purpose. To this end, we consider fail-stop errors first. If the application stops, then we do not need to perform a verification since we must do a recovery anyway. If no

fail-stop error strikes during the execution, we proceed with the verification and check for silent errors. Therefore,

$$T_C^{SF}(i, j, s) = p_{i,j}^F(s) (T_{lost_{i,j}}(s) + R_{i-1} + T_C^{SF}(i, j, s)) \\ + (1 - p_{i,j}^F(s)) (T_{i,j}(s) + V_j(s) + p_{i,j}^S(s) (R_{i-1} + T_C^{SF}(i, j, s))) .$$

When plugging $p_{i,j}^F(s)$, $p_{i,j}^S(s)$ and $T_{lost_{i,j}}(s)$ into the above equation, we get

$$T_C^{SF}(i, j, s) = e^{\lambda^S(s)T_{i,j}(s)} \left(\frac{e^{\lambda^F(s)T_{i,j}(s)} - 1}{\lambda^F(s)} + V_j(s) \right) + \left(e^{(\lambda^F(s) + \lambda^S(s))T_{i,j}(s)} - 1 \right) R_{i-1} . \quad (3)$$

By setting $T_C(i, j, s) = T_C^{SF}(i, j, s)$, we can now compute $Time_C(j, s)$ for all $j = 1, \dots, n$. For the complexity, the computation of $T_C^{SF}(i, j, s)$ for all (i, j) pairs with $i \leq j$ takes $O(n^2)$ time. The computation of the dynamic programming table for $Time_C(j, s)$ also takes $O(n^2)$ time, as $Time_C(j, s)$ depends on at most j other entries in the same table, which are already computed. Therefore, the overall complexity is $O(n^2)$, and this concludes the proof. \square

Theorem 4.1 extends the result of Toueg and Babaoglu [Toueg and Babaoglu 1984] to a linear chain of tasks subject to both fail-stop and silent errors. For the sake of comparing with the case of a divisible load application, we extend Young/Daly's formula [Young 1974; Daly 2006] as follows.

PROPOSITION 4.2. *For a divisible load application subject to both fail-stop and silent errors, a first-order approximation of the optimal checkpointing period to minimize the expected execution time is*

$$T_{VCO}^{time}(s) = \sqrt{\frac{2(V(s) + C)}{\lambda^F(s) + 2\lambda^S(s)}} ,$$

where $V(s)$ is the verification cost at speed s , $\lambda^F(s)$ is the rate of fail-stop errors at speed s , $\lambda^S(s)$ is the rate of silent errors at speed s , and C is the checkpointing cost.

PROOF. In the presence of fail-stop errors only, Young/Daly's formula [Young 1974; Daly 2006] gives the optimal checkpointing period $\sqrt{\frac{2C}{\lambda^F(s)}}$ if we process the application with speed s . With both fail-stop and silent errors, let $T(s)$ denote the period. According to Equation (3), the expected execution time for a segment between two consecutive checkpoints (accounting for the cost of the checkpoint itself) is given by

$$Time(T(s)) = e^{\lambda^S(s)T(s)} \left(\frac{e^{\lambda^F(s)T(s)} - 1}{\lambda^F(s)} + V(s) \right) + \left(e^{(\lambda^F(s) + \lambda^S(s))T(s)} - 1 \right) R + C , \quad (4)$$

and we are interested in the value of $T(s)$ that minimizes the overhead $\frac{Time(T(s))}{T(s)}$ compared to the error-free and checkpoint-free execution time $T(s)$. Using Taylor expansion to approximate $e^{\lambda T} \approx 1 + \lambda T + \frac{\lambda^2 T^2}{2}$ up to the second term, we get the following first-order approximation for the overhead:

$$\frac{Time(T(s))}{T(s)} = \left(\frac{\lambda^F(s)}{2} + \lambda^S(s) \right) T(s) + \frac{V(s) + C}{T(s)} + 1 + \lambda^S(s)V(s) + (\lambda^F(s) + \lambda^S(s))R .$$

Differentiating the above expression with respect to $T(s)$, we find that $T_{VCO}^{time}(s) = \sqrt{\frac{2(V(s) + C)}{\lambda^F(s) + 2\lambda^S(s)}}$ minimizes the overhead, which extends Young/Daly's formula to include

both fail-stop and silent errors. We stress that, as in Young/Daly's formula, this result is a first-order approximation, which is valid only if all resilience parameters C , R and $V(s)$ are small in front of both MTBF values, namely $1/\lambda^F(s)$ for fail-stop errors and $1/\lambda^S(s)$ for silent errors. \square

In Section 7, we compare the makespan derived from this formula with the performance of the VC-ONLY algorithm. With a chain of tasks, we have less flexibility for checkpointing than with a divisible load application, and it is interesting to numerically evaluate the difference due to the application framework.

4.1.2. ENERGY-VC-ONLY: MINIMIZING ENERGY

PROPOSITION 4.3. *For the SINGLESPEED scenario, the ENERGY-VC-ONLY problem can be solved by a dynamic programming algorithm in $O(n^2)$ time.*

PROOF. The proof is similar to that of the TIME-VC-ONLY problem presented in Section 4.1.1. Here, we replace $Time_C(j, s)$ with $Energy_C(j, s)$, which is the optimal expected energy to successfully execute the tasks T_1 to T_j . Instead of $T_C^{SF}(i, j, s)$, we use $E_C^{SF}(i, j, s)$ to denote the total expected energy to successfully execute all the tasks from T_i to T_j without any checkpoint and verification in between, while T_{i-1} and T_j are both verified and checkpointed. The goal is to find $Energy_C(n, s)$ and the corresponding dynamic program is formulated as:

$$Energy_C(j, s) = \min_{0 \leq i < j} \{Energy_C(i, s) + E_C^{SF}(i+1, j, s)\} + E_j^C.$$

Previously we used $T_{lost_{i,j}}(s)$ to denote the expected time lost when a fail-stop error occurs. Now, we use $E_{lost_{i,j}}(s) = T_{lost_{i,j}}(s)(P_{cpu}(s) + P_{idle})$ to denote the expected energy lost when a fail-stop error stroke during the execution from task T_i to task T_j . Therefore, we can express $E_C^{SF}(i, j, s)$ as follows:

$$\begin{aligned} E_C^{SF}(i, j, s) &= p_{i,j}^F(s) \left(E_{lost_{i,j}}(s) + E_{i-1}^R + E_C^{SF}(i, j, s) \right) \\ &\quad + (1 - p_{i,j}^F(s)) \left(E_{i,j}(s) + E_j^Y(s) + p_{i,j}^S(s) \left(E_{i-1}^R + E_C^{SF}(i, j, s) \right) \right), \end{aligned}$$

which simplifies to

$$\begin{aligned} E_C^{SF}(i, j, s) &= (P_{idle} + P_{cpu}(s)) e^{\lambda^S(s)T_{i,j}(s)} \left(\frac{e^{\lambda^F(s)T_{i,j}(s)} - 1}{\lambda^F(s)} + V_j(s) \right) \\ &\quad + \left(e^{(\lambda^F(s) + \lambda^S(s))T_{i,j}(s)} - 1 \right) E_{i-1}^R. \end{aligned} \quad (5)$$

Clearly, the time complexity is the same as that of the makespan minimization algorithm (in Theorem 4.1). \square

Proposition 4.3 extends the dynamic programming algorithm of Toueg and Babaoglu to account for energy consumption. Now, for consistency, we present the optimal checkpointing period that minimizes the total energy consumption in the case of a divisible load application.

PROPOSITION 4.4. *For a divisible load application subject to both fail-stop and silent errors, a first-order approximation of the optimal checkpointing period to minimize the expected energy consumption is*

$$T_{VCO}^{energy}(s) = \sqrt{\frac{2(V(s) + C_e(s))}{\lambda^F(s) + 2\lambda^S(s)}}, \quad (6)$$

where $V(s)$ is the verification cost at speed s , $\lambda^F(s)$ is the rate of fail-stop errors at speed s , $\lambda^S(s)$ is the rate of silent errors at speed s . Here, $C_e(s)$ is defined as $C_e(s) = \frac{C(P_{idle}+P_{io})}{P_{idle}+P_{cpu}(s)}$, where C is the checkpointing cost, P_{idle} is the static idle power, $P_{cpu}(s)$ is the dynamic CPU power at speed s , and P_{io} is the dynamic I/O power.

PROOF. The proof is similar to that of Proposition 4.2. In the presence of both fail-stop and silent errors, let $T(s)$ denote the checkpointing period and let $E(s) = T(s)(P_{idle} + P_{cpu}(s))$ denote the error-free and checkpoint-free energy consumption in the period. According to Equation (5), the expected energy consumption for a segment between two consecutive checkpoints (accounting for the cost of the checkpoint itself) is given by:

$$\begin{aligned} Energy(T(s)) &= (P_{idle} + P_{cpu}(s)) e^{\lambda^S(s)T(s)} \left(\frac{e^{\lambda^F(s)T(s)} - 1}{\lambda^F(s)} + V(s) \right) \\ &\quad + \left(e^{(\lambda^F(s)+\lambda^S(s))T(s)} - 1 \right) (P_{idle} + P_{io})R + (P_{idle} + P_{io})C, \end{aligned}$$

which can be rewritten as

$$Energy(T(s)) = (P_{idle} + P_{cpu}(s)) F(T(s)),$$

where $F(T(s))$ is similar to $Time(T(s))$ in Equation (4), but with C replaced by $C_e(s) = \frac{C(P_{idle}+P_{io})}{P_{idle}+P_{cpu}(s)}$ and R replaced by $R_e(s) = \frac{R(P_{idle}+P_{io})}{P_{idle}+P_{cpu}(s)}$ as follows:

$$F(T(s)) = e^{\lambda^S(s)T(s)} \left(\frac{e^{\lambda^F(s)T(s)} - 1}{\lambda^F(s)} + V(s) \right) + \left(e^{(\lambda^F(s)+\lambda^S(s))T(s)} - 1 \right) R_e(s) + C_e(s).$$

Again, we are interested in the value of $T(s)$ that minimizes the overhead

$$\frac{Energy(T(s))}{E(s)} = \frac{F(T(s))}{T(s)}.$$

Using the same approach as in Proposition 4.2, we can get the following first-order approximation to the overhead:

$$\frac{F(T(s))}{T(s)} = \left(\frac{\lambda^F(s)}{2} + \lambda^S(s) \right) T(s) + \frac{V(s) + C_e(s)}{T(s)} + 1 + \lambda^S(s)V(s) + (\lambda^F(s) + \lambda^S(s))R_e(s).$$

Differentiating the above expression with respect to $T(s)$, we find that $T_{VCO}^{energy}(s) = \sqrt{\frac{2(V(s)+C_e(s))}{\lambda^F(s)+2\lambda^S(s)}}$ minimizes the overhead. This result is analogous to what we have in Proposition 4.2 and extends Young/Daly's formula to cover energy consumption in the presence of both fail-stop and silent errors. \square

4.2. VC+V: Using verified checkpoints and single verifications

In this approach, we can place additional verifications between two checkpoints, which allows to detect silent errors before reaching the next checkpoint, and hence to avoid wasted execution by performing early recoveries. We aim at finding the best positions for checkpoints and verifications in order to minimize the total execution time (TIME-VC+V) or the total energy consumption (ENERGY-VC+V). For both objectives, adding extra verifications between two checkpoints adds an extra step in the algorithm, which results in a higher complexity.

4.2.1. TIME-VC+V: MINIMIZING MAKESPAN

THEOREM 4.5. *For the SINGLESPEED scenario, the TIME-VC+V problem can be solved by a dynamic programming algorithm in $O(n^3)$ time.*

PROOF. In the TIME-VC-ONLY problem, we were only allowed to place verified checkpoints. Here, we can add single verifications that are not associated with a checkpoint. The main idea is to replace T_C in the dynamic program of Theorem 4.1 by another expression $Time_V(i, j, s)$, which denotes the optimal expected time to successfully execute from task T_i to task T_j (and to verify it), provided that T_{i-1} has a verified checkpoint and only single verifications are allowed within tasks T_i, \dots, T_{j-1} . Furthermore, we use $Time_{VC}(j, s)$ to denote the optimal expected time to successfully execute the first j tasks, where T_j has a verified checkpoint, and there are possibly other verified checkpoints and single verifications before T_j . The goal is to find $Time_{VC}(n, s)$. The dynamic program to compute $Time_{VC}(j, s)$ can be formulated as (see Figure 2):

$$Time_{VC}(j, s) = \min_{0 \leq i < j} \{Time_{VC}(i, s) + Time_V(i+1, j, s)\} + C_j .$$




Fig. 2. Illustration of the dynamic programming formulation for $Time_{VC}(j, s)$.

In particular, we try all possible locations for the last checkpoint before T_j , and for each location T_i , we compute the optimal expected time $Time_V(i+1, j, s)$ to execute tasks T_{i+1} to T_j with only single verifications in between. We also account for the checkpointing time C_j , which is not included in the definition of $Time_V$. By initializing the dynamic program with $Time_{VC}(0, s) = 0$, we can then compute the optimal solution as in the TIME-VC-ONLY problem.

It remains to compute $Time_V(i, j, s)$ for each (i, j) pair with $i \leq j$. To this end, we formulate another dynamic program by trying all possible locations for the last single verification before T_j (see Figure 3):

$$Time_V(i, j, s) = \min_{i-1 \leq l < j} \{Time_V(i, l, s) + T_V(l+1, j, i-1, s)\} ,$$

where $T_V(i, j, l_c, s)$ is the expected time to successfully execute all the tasks from T_i to T_j (and to verify T_j), knowing that if an error strikes, we can recover from T_{l_c} , the last task before T_i to have a verified checkpoint.




Fig. 3. Illustration of the dynamic programming formulation for $Time_V(i, j, s)$.

First, we show how to compute $T_V(i, j, l_c, s)$. When accounting for only *silent errors* (with notation T_V^S), we always execute from task T_i to task T_j and then verify T_j . In case of failure, we recover from T_{l_c} and redo the entire computation from T_{l_c+1} to T_j , which contains a single verification after T_{i-1} and possibly other single verifications between T_{l_c+1} and T_{i-2} . Hence, we have

$$T_V^S(i, j, l_c, s) = T_{i,j}(s) + V_j(s) + p_{i,j}^S(s) (R_{l_c} + Time_V(l_c + 1, i - 1, s) + T_V^S(i, j, l_c, s)) ,$$

which simplifies to

$$T_V^S(i, j, l_c, s) = e^{\lambda^S(s)T_{i,j}(s)} (T_{i,j}(s) + V_j(s)) + (e^{\lambda^S(s)T_{i,j}(s)} - 1) (R_{l_c} + Time_V(l_c + 1, i - 1, s)) .$$

When there are only *fail-stop errors*, we do not need to perform any single verification, and hence the problem becomes simply the TIME-VC-ONLY problem. When accounting for both *silent and fail-stop errors* (with notation T_V^{SF}), we apply the same method as in the proof of Theorem 4.1. Specifically, if a fail-stop error strikes between two verifications, we directly perform a recovery; otherwise we check for silent errors:

$$\begin{aligned} T_V^{SF}(i, j, l_c, s) &= p_{i,j}^F(s) (T_{lost_{i,j}}(s) + R_{l_c} + Time_V(l_c + 1, i - 1, s) + T_V^{SF}(i, j, l_c, s)) \\ &+ (1 - p_{i,j}^F(s)) (T_{i,j}(s) + V_j(s) + p_{i,j}^S(s) (R_{l_c} + Time_V(l_c + 1, i - 1, s) + T_V^{SF}(i, j, l_c, s))) . \end{aligned}$$

When plugging $p_{i,j}^F$, $p_{i,j}^S$ and $T_{lost_{i,j}}$ into the above equation, we get

$$\begin{aligned} T_V^{SF}(i, j, l_c, s) &= e^{\lambda^S(s)T_{i,j}(s)} \left(\frac{e^{\lambda^F(s)T_{i,j}(s)} - 1}{\lambda^F(s)} + V_j(s) \right) \\ &+ (e^{(\lambda^F(s) + \lambda^S(s))T_{i,j}(s)} - 1) (R_{l_c} + Time_V(l_c + 1, i - 1, s)) . \end{aligned}$$

Note that $T_V(i, j, l_c, s)$ depends on the value of $Time_V(l_c + 1, i - 1, s)$, except when $l_c + 1 = i$, in which case we initialize $Time_V(i, i - 1, s) = 0$. Hence, in the dynamic program, $Time_V(i, j, s)$ can be expressed as a function of $Time_V(i, l, s)$ for all $l = i - 1, \dots, j - 1$.

Finally, the complexity is dominated by the computation of the second dynamic programming table for $Time_V(i, j, s)$, which contains $O(n^2)$ entries and each entry depends on at most n other entries that are already computed. Hence, the overall complexity of the algorithm is $O(n^3)$, and this concludes the proof. \square

Now, for the sake of comparing with a divisible load application, we assume that the verification pattern between two consecutive checkpoints is also periodic. The following result extends Proposition 4.2 by showing the optimal verification period.

PROPOSITION 4.6. *For a divisible load application subject to both fail-stop and silent errors, a first-order approximation of the optimal verification period to minimize*

the expected execution time is

$$T_{VCV}^{time}(s) = \sqrt{\frac{2(V(s) + C/\bar{k}^*)}{\bar{k}^* \lambda^F(s) + (\bar{k}^* + 1)\lambda^S(s)}}.$$

Here, \bar{k}^* denotes the optimal number of verifications between two consecutive checkpoints and it is equal to either $\max(1, \lfloor k^* \rfloor)$ or $\lceil k^* \rceil$, where $k^* \in (0, \sqrt{\frac{c}{a}})$ is the unique solution to the following equation:

$$2d\sqrt{ak^* + b + \frac{c}{k^*}} = \frac{c}{(k^*)^2} - a, \quad (7)$$

with $a = 2V(s)(\lambda^F(s) + \lambda^S(s))$, $b = 2C(\lambda^F(s) + \lambda^S(s)) + 2V(s)\lambda^S(s)$, $c = 2C\lambda^S(s)$ and $d = \frac{V(s)(\lambda^S(s) + \lambda^F(s))}{2}$.

PROOF. Let $T(s)$ denote the verification period and k the number of verifications between two consecutive checkpoints. For convenience, we refer to a segment between two consecutive verifications a V-segment and a segment between two consecutive checkpoints a C-segment. Let $p^F(s) = 1 - e^{-\lambda^F(s)T(s)}$ denote the probability that a fail-stop error strikes when executing a V-segment, and let $p^S(s) = 1 - e^{-\lambda^S(s)T(s)}$ denote the corresponding probability for silent errors. The probability that neither type of error occurs during the execution of a V-segment is then given by $q(s) = (1 - p^F(s))(1 - p^S(s)) = e^{-(\lambda^F(s) + \lambda^S(s))T(s)}$. The expected execution time for a C-segment can be expressed recursively by enumerating the failure possibilities for all V-segments inside it as follows:

$$\begin{aligned} Time(T(s), k) &= \sum_{i=1}^k q(s)^{i-1} \left(p^F(s) \left((i-1) \cdot (T(s) + V(s)) + T_{lost}(s) + R + Time(T(s), k) \right) \right. \\ &\quad \left. + (1 - p^F(s)) p^S(s) \left(i \cdot (T(s) + V(s)) + R + Time(T(s), k) \right) \right) \\ &\quad + q(s)^k \left(k \cdot (T(s) + V(s)) + C \right), \end{aligned} \quad (8)$$

where $T_{lost}(s)$ is the expected time lost during the execution of a V-segment, knowing that a fail-stop error has struck. According to Equation (2), we have $T_{lost}(s) = \frac{1}{\lambda^F(s)} - \frac{T(s)}{e^{\lambda^F(s)T(s)} - 1}$.

Let $h = \sum_{i=1}^k i \cdot q(s)^{i-1}$, and $q(s)h = \sum_{i=1}^k i \cdot q(s)^i$, so we have $(1 - q(s))h = \sum_{i=1}^k q(s)^{i-1} - kq(s)^k = \frac{1 - q(s)^k}{1 - q(s)} - kq(s)^k$, hence $h = \frac{1 - q(s)^k}{(1 - q(s))^2} - \frac{kq(s)^k}{1 - q(s)}$. Substituting h

and $T_{lost}(s)$ into Equation (8) and simplifying it, we get

$$\begin{aligned}
Time(T(s), k) &= \frac{1 - q(s)^k}{1 - q(s)} (p^F(s) + (1 - p^F(s))p^S(s)) (R + Time(T(s), k)) \\
&\quad + \frac{(1 - q(s)^k)p^F(s)}{1 - q(s)} \left(\frac{1}{\lambda^F(s)} - \frac{T(s)}{e^{\lambda^F(s)T(s)-1}} \right) + q(s)^k C \\
&\quad + (T(s) + V(s)) \left(\frac{p^F(s)(q(s) - q(s)^k) + (1 - p^F(s))p^S(s)(1 - q(s)^k)}{(1 - q(s))^2} \right. \\
&\quad \quad \left. - \frac{(k-1)p^F(s)q(s)^k + k(1 - p^F(s))p^S(s)q(s)^k}{1 - q(s)} \right) \\
&= (1 - q(s)^k)(R + Time(T(s), k)) + \frac{(1 - p^F(s))(1 - q(s)^k)}{1 - q(s)} (T(s) + V(s)) \\
&\quad + \frac{(1 - q(s)^k)p^F(s)}{1 - q(s)} \left(\frac{1}{\lambda^F(s)} - \frac{T(s)}{e^{\lambda^F(s)T(s)-1}} \right) + q(s)^k C,
\end{aligned}$$

which leads to

$$\begin{aligned}
Time(T(s), k) &= \frac{q(s)^{-k} - 1}{1 - q(s)} \left((1 - p^F(s))(T(s) + V(s)) \right. \\
&\quad \left. + p^F(s) \left(\frac{1}{\lambda^F(s)} - \frac{T(s)}{e^{\lambda^F(s)T(s)-1}} \right) \right) + (q(s)^{-k} - 1)R + C. \quad (9)
\end{aligned}$$

Applying Taylor expansion to Equation (9) by approximating $e^{\lambda T} \approx 1 + \lambda T + \frac{\lambda^2 T^2}{2}$ and

$$\begin{aligned}
\frac{e^{k\lambda T} - 1}{1 - e^{-\lambda T}} &\approx \frac{k\lambda T (1 + \frac{k\lambda T}{2})}{\lambda T (1 - \frac{\lambda T}{2})} \\
&\approx k \left(1 + \frac{(k+1)\lambda T}{2} \right),
\end{aligned}$$

we can approximate, up to the first-order term, the overhead of executing a C-segment as

$$\begin{aligned}
\frac{Time(T(s), k)}{kT(s)} &= \frac{k\lambda^F(s) + (k+1)\lambda^S(s)}{2} T(s) + \frac{V(s) + C/k}{T(s)} \\
&\quad + 1 + \frac{(k+1)\lambda^S(s) + (k-1)\lambda^F(s)}{2} V(s) + (\lambda^F(s) + \lambda^S(s))R. \quad (10)
\end{aligned}$$

Differentiating Equation (10) with respect to $T(s)$, we find that $T_{V_{CV}}^{time}(s) = \sqrt{\frac{2(V(s)+C/k)}{k\lambda^F(s)+(k+1)\lambda^S(s)}}$ minimizes the overhead.

Now, substituting $T_{V_{CV}}^{time}(s)$ back into Equation (10), we get

$$\frac{Time(T(s), k)}{kT(s)} = \sqrt{ak + b + \frac{c}{k}} + dk + e, \quad (11)$$

where $a = 2V(s)(\lambda^F(s) + \lambda^S(s))$, $b = 2C(\lambda^F(s) + \lambda^S(s)) + 2V(s)\lambda^S(s)$, $c = 2C\lambda^S(s)$, $d = \frac{V(s)(\lambda^S(s)+\lambda^F(s))}{2}$ and $e = 1 + (\lambda^F(s) + \lambda^S(s))R + \frac{V(s)(\lambda^S(s)-\lambda^F(s))}{2}$. Differentiating Equation (11) with respect to k , we find that the optimal k^* satisfies Equation (7), which has a unique solution in $(0, \sqrt{\frac{c}{a}})$ that can be solved numerically. As the number

of verifications in a C-segment must be an integer, the optimal strategy uses either $\max(1, \lfloor k^* \rfloor)$ or $\lceil k^* \rceil$, whichever leads to a smaller value for the overhead $\frac{Time(T(s), k)}{kT(s)}$. \square

Let us consider an example. Suppose $\lambda^F(s) = 0.001$, $\lambda^S(s) = 0.002$, $C = R = 20$ and $V(s) = 1$. The optimal value of k^* should satisfy

$$0.003\sqrt{0.006k^* + 0.124} + \frac{0.08}{k^*} = \frac{0.08}{(k^*)^2} - 0.006,$$

whose solution is $k^* \approx 3.3265 < 3.6515 \approx \sqrt{\frac{c}{a}}$. This leads to the optimal number of verifications $\bar{k}^* = 3$ and $T^*(s) \approx 37.3355$ with an overhead $\frac{Time(T(s), k)}{kT(s)} \approx 1.4757$. Figure 4 plots the overheads for different values of k and $T(s)$ around the optimal pair of solution. Finally, we observe in the example that

$$k^* \approx \sqrt{\frac{c}{a}} = \sqrt{\frac{\lambda^S(s)}{\lambda^F(s) + \lambda^S(s)} \times \frac{C}{V(s)}}.$$

This is a general result, which always holds when $\lambda^F(s)$ and $\lambda^S(s)$ are very small: letting $\lambda = \lambda^F(s) + \lambda^S(s)$, the left-hand side of Equation (7) has order $O(\lambda^{1.5})$ while the right-hand side has order $O(\lambda)$. Hence, the value of k^* that zeroes out the right-hand side, namely $\sqrt{\frac{c}{a}}$, is a good approximation when λ is small.




Fig. 4. The overheads for different values of $(k, T(s))$ around the optimal pair $(3, 37.3355)$ for the setting with $\lambda^F(s) = 0.001$, $\lambda^S(s) = 0.002$, $C = R = 20$ and $V(s) = 1$.

4.2.2. ENERGY-VC+V: MINIMIZING ENERGY

PROPOSITION 4.7. *For the SINGLESPEED scenario, the ENERGY-VC+V problem can be solved by a dynamic programming algorithm in $O(n^3)$ time.*

PROOF. The proof is similar to that of the TIME-VC+V problem presented in Section 4.2.1. Here, we replace $Time_{VC}(j, s)$ with $Energy_{VC}(j, s)$, which is the optimal expected energy to successfully execute the tasks T_1 to T_j . Moreover, instead of $Time_V(i, j, s)$, we use $Energy_V(i, j, s)$ and replace $T_V^{SF}(i, j, l_c, s)$ with $E_V^{SF}(i, j, l_c, s)$. The goal is to find $Energy_{VC}(n, s)$, and the two dynamic programs can be correspondingly formulated as:

$$Energy_{VC}(j) = \min_{0 \leq i < j} \{Energy_{VC}(i, s) + Energy_V(i+1, j, s)\} + E_j^C,$$

$$Energy_V(i, j, s) = \min_{i-1 \leq l < j} \{Energy_V(i, l, s) + E_V^{SF}(l+1, j, i-1, s)\} .$$

In particular, $E_V^{SF}(i, j, l_c, s)$ has similar construction as $T_V^{SF}(i, j, l_c, s)$, and it denotes the expected total energy consumption to successfully execute the tasks from T_i to T_j , knowing that the last checkpoint is at T_{l_c} and there might be additional single verifications between T_{l_c+1} and T_{i-2} . When accounting for both fail-stop and silent errors, $E_V^{SF}(i, j, l_c, s)$ can be expressed as:

$$E_V^{SF}(i, j, l_c, s) = p_{i,j}^F(s) \left(E_{lost_{i,j}}(s) + E_{l_c}^R + Energy_V(l_c+1, i-1, s) + E_V^{SF}(i, j, l_c, s) \right) \\ + (1 - p_{i,j}^F(s)) \left(E_{i,j}(s) + E_j^Y(s) + p_{i,j}^S(s) \left(E_{l_c}^R + Energy_V(l_c+1, i-1, s) + E_V^{SF}(i, j, l_c, s) \right) \right) ,$$

which simplifies to

$$E_V^{SF}(i, j, l_c, s) = (P_{idle} + P_{cpu}(s)) e^{\lambda^S(s)T_{i,j}(s)} \left(\frac{e^{\lambda^F(s)T_{i,j}(s)} - 1}{\lambda^F(s)} + V_j(s) \right) \\ + \left(e^{(\lambda^F(s) + \lambda^S(s))T_{i,j}(s)} - 1 \right) \left(E_{l_c}^R + Energy_V(l_c+1, i-1, s) \right) .$$

Clearly, the complexity is the same as that of the makespan minimization algorithm (in Theorem 4.5). \square

PROPOSITION 4.8. *For a divisible load application subject to both fail-stop and silent errors, a first-order approximation of the optimal verification period to minimize the expected energy consumption is*

$$T_{VCV}^{energy}(s) = \sqrt{\frac{2(V(s) + C_e(s)/\bar{k}^*)}{\bar{k}^* \lambda^F(s) + (\bar{k}^* + 1)\lambda^S(s)}} ,$$

where $C_e(s) = \frac{C(P_{idle} + P_{io})}{P_{idle} + P_{cpu}(s)}$. Here, \bar{k}^* denotes the optimal number of verifications between two consecutive checkpoints and it is equal to either $\max(1, \lfloor k^* \rfloor)$ or $\lceil k^* \rceil$, where $k^* \in (0, \sqrt{\frac{c}{a}})$ is the unique solution to the following equation:

$$2d \sqrt{ak^* + b + \frac{c}{k^*}} = \frac{c}{(k^*)^2} - a ,$$

with $a = 2V(s)(\lambda^F(s) + \lambda^S(s))$, $b = 2C_e(s)(\lambda^F(s) + \lambda^S(s)) + 2V(s)\lambda^S(s)$, $c = 2C_e(s)\lambda^S(s)$ and $d = \frac{V(s)(\lambda^S(s) + \lambda^F(s))}{2}$.

PROOF. The proof is similar to that of Proposition 4.6. Let $T(s)$ denote the verification period and k the number of verifications between two consecutive checkpoints. We can rewrite Equation (8) to compute the expected energy consumption instead of the expected execution time for a C-segment by replacing C with $C_e(s) = \frac{C(P_{idle} + P_{io})}{P_{idle} + P_{cpu}(s)}$ and replacing R with $R_e(s) = \frac{R(P_{idle} + P_{io})}{P_{idle} + P_{cpu}(s)}$. Then accounting for P_{idle} and $P_{cpu}(s)$, we can get the following expression:

$$Energy(T(s), k) = (P_{idle} + P_{cpu}(s)) F(T(s), k) , \quad (12)$$

where $F(T(s), k)$ can be expressed as

$$F(T(s), k) = \sum_{i=1}^k q(s)^{i-1} \left(p^F(s) \left((i-1) \cdot (T(s) + V(s)) + T_{lost}(s) + R_e(s) + F(T(s), k) \right) \right. \\ \left. + (1 - p^F(s)) p^S(s) \left(i \cdot (T(s) + V(s)) + R_e(s) + F(T(s), k) \right) \right) \\ + q(s)^k \left(k \cdot (T(s) + V(s)) + C_e(s) \right).$$

Let $E(s) = (P_{idle} + P_{cpu}(s))T(s)$ denotes the error-free, checkpoint-free and verification-free energy consumption during a verification period $T(s)$. Again, we are interested in the value of $T(s)$ that minimizes the overhead

$$\frac{Energy(T(s), k)}{kE(s)} = \frac{F(T(s), k)}{kT(s)}.$$

Using the same approach as in Proposition 4.6, we can get the following first-order approximation to the overhead:

$$\frac{F(T(s), k)}{kT(s)} = \frac{k\lambda^F(s) + (k+1)\lambda^S(s)}{2} T(s) + \frac{V(s) + C_e(s)/k}{T(s)} \\ + 1 + \frac{(k+1)\lambda^S(s) + (k-1)\lambda^F(s)}{2} V(s) + (\lambda^F(s) + \lambda^S(s))R_e(s). \quad (13)$$

Note that Equation (13) is very similar to Equation (10). Hence, by differentiating Equation (13) with respect to $T(s)$ and k , we can get the optimal verification period $T_{VCV}^{energy}(s)$ and the optimal number k^* as claimed in the proposition. The results are analogous to what we have in Proposition 4.6. \square

5. REEXECSPEED SCENARIO

In the REEXECSPEED scenario, we are given two CPU speeds s and σ , where s is the regular speed and σ is the adjusted speed. The regular speed s is used for the first execution of the tasks, while σ is used to speed up or to slow down any potential re-execution in order to improve makespan or energy. We always account for both *silent and fail-stop errors*.

In this scenario, and for both VC-ONLY and VC+V approaches, we need to derive two independent expressions to compute the expected execution time (or energy consumption). The first expression is for the first execution of the tasks with the first speed until the first error is encountered. Once the first error strikes, we recover from the last checkpoint and start re-executing the tasks with the second speed until we reach the next checkpoint. This latter expression for either the time or energy is essentially the same as that in the SINGLESPEED scenario, but with speed σ .

5.1. VC-ONLY: Using verified checkpoints only

In this approach, we aim at finding the best positions of verified checkpoints in order to minimize the makespan (TIME-VC-ONLY) or the total energy consumption (ENERGY-VC-ONLY).

5.1.1. TIME-VC-ONLY: MINIMIZING MAKESPAN

THEOREM 5.1. *For the REEXECSPEED scenario, the TIME-VC-ONLY problem can be solved by a dynamic programming algorithm in $O(n^2)$ time.*

PROOF. The proof extends that of Theorem 4.1 in Section 4.1.1. To account for the two speeds, we replace $Time_C(j, s)$ with $Time_{C_{re}}(j, s, \sigma)$, which denotes the optimal

expected time to successfully execute the tasks T_1 to T_j , where T_j has a verified checkpoint. Similarly, we replace $T_C(i, j, s)$ with $T_{C_{re}}(i, j, s, \sigma)$ as the expected time to successfully execute the tasks T_i to T_j , where both T_{i-1} and T_j are verified and checkpointed. The goal is to find $Time_{C_{re}}(n, s, \sigma)$, and the dynamic program is formulated as:

$$Time_{C_{re}}(j, s, \sigma) = \min_{0 \leq i < j} \{Time_{C_{re}}(i, s, \sigma) + T_{C_{re}}(i+1, j, s, \sigma)\} + C_j.$$

Note that the checkpointing cost after T_j is included in $Time_{C_{re}}(j, s, \sigma)$ but not in $T_{C_{re}}(i, j, s, \sigma)$. We initialize the dynamic program with $Time_{C_{re}}(0, s, \sigma) = 0$.

To compute $T_{C_{re}}(i, j, s, \sigma)$ for each (i, j) pair with $i \leq j$, we need to distinguish the first execution (before the first error) and all the potential re-executions (after at least one error). Let $T_{C_{first}}^{SF}(i, j, s)$ denote the expected time to execute the tasks T_i to T_j for the very first time, before the first error is encountered, and let $T_C^{SF}(i, j, \sigma)$ denote the expected time to successfully execute the tasks T_i to T_j in the re-executions but using speed σ . While $T_C^{SF}(i, j, \sigma)$ can be computed in the same way as in Section 4.1.1, $T_{C_{first}}^{SF}(i, j, s)$ is computed by considering two possible scenarios: either a fail-stop error has occurred during the execution from T_i to T_j , and we lose $T_{lost_{i,j}}(s)$ time as given in Equation (2), or there was no fail-stop error and we check whether a silent error has occurred. Note that in both cases we do not account for the re-executions, as they are handled by T_C^{SF} separately (with the second speed). Therefore,

$$T_{C_{first}}^{SF}(i, j, s) = (1 - p_{i,j}^F(s))(T_{i,j}(s) + V_j(s)) + p_{i,j}^F(s) \left(\frac{1}{\lambda^F(s)} - \frac{T_{i,j}(s)}{e^{\lambda^F(s)T_{i,j}(s)} - 1} \right).$$

Let $p_{i,j}^{Err}(s)$ denote the probability that at least one error is detected in the first execution of the tasks from T_i to T_j at speed s . Note that we account for both *silent and fail-stop errors*, and we can only detect silent errors if no fail-stop error has occurred, thus $p_{i,j}^{Err}(s) = p_{i,j}^F(s) + (1 - p_{i,j}^F(s))p_{i,j}^S(s)$. If no error strikes during the first execution, i.e., $p_{i,j}^{Err}(s) = 0$, then the time to execute from T_i to T_j is exactly $T_{C_{first}}^{SF}(i, j, s)$, which means that all the tasks have been executed successfully with the first speed. If at least one error occurs with probability $p_{i,j}^{Err}(s) > 0$, then $T_{C_{first}}^{SF}(i, j, s)$ is the time lost trying to execute the tasks with the first speed. In this case, we need to recover from the last checkpoint and use $T_C^{SF}(i, j, \sigma)$ to re-execute all the tasks from T_i to T_j with the second speed until we pass the next checkpoint. Therefore,

$$T_{C_{re}}(i, j, s, \sigma) = T_{C_{first}}^{SF}(i, j, s) + p_{i,j}^{Err}(s) (R_{i-1} + T_C^{SF}(i, j, \sigma)). \quad (14)$$

Despite the two steps needed to compute $T_{C_{re}}(i, j, s, \sigma)$, the complexity remains the same as in the SINGLESPEED scenario (Theorem 4.1). This concludes the proof. \square

5.1.2. ENERGY-VC-ONLY: MINIMIZING ENERGY

PROPOSITION 5.2. *For the REEXEC SPEED scenario, the ENERGY-VC-ONLY problem can be solved by a dynamic programming algorithm in $O(n^2)$ time.*

PROOF. The proof is similar to that of the TIME-VC-ONLY problem presented in Section 5.1.1. We replace $Time_{C_{re}}(j, s, \sigma)$ with $Energy_{C_{re}}(j, s, \sigma)$ while accounting for both speeds. The goal is to find $Energy_{C_{re}}(n, s, \sigma)$, and the dynamic program is formulated as:

$$Energy_{C_{re}}(j, s, \sigma) = \min_{0 \leq i < j} \{Energy_{C_{re}}(i, s, \sigma) + E_{C_{re}}(i+1, j, s, \sigma)\} + E_j^C.$$

Similarly to $T_{C_{re}}(i, j, s, \sigma)$, we introduce $E_{C_{re}}(i, j, s, \sigma)$, which is also composed of two parts: one for the first execution with the first speed s and the other for all potential

re-executions with the second speed σ . Again, the first part only includes the energy spent on executing the tasks until the first error is encountered and it does not account for any recovery cost. As a result, the associated energy cost is only contributed by the CPU and is directly proportional to the time $T_{C_{first}}^{SF}(i, j, s)$ spent on executing the tasks in the first execution. We can express $E_{C_{re}}(i, j, s, \sigma)$ as:

$$E_{C_{re}}(i, j, s, \sigma) = (P_{idle} + P_{cpu}(s)) T_{C_{first}}^{SF}(i, j, s) + p_{i,j}^{Err}(s) \left(E_{i-1}^R + E_C^{SF}(i, j, \sigma) \right), \quad (15)$$

where $E_C^{SF}(i, j, \sigma)$ can be computed in the same way as in Section 4.1.2. Clearly, the complexity is the same as that of the makespan minimization algorithm (in Theorem 5.1). \square

5.2. VC+V: Using verified checkpoints and single verifications

In the VC+V approach, we can place single verifications between two verified checkpoints to achieve early detection of silent errors. As two speeds are used in this scenario, we will place two sets of single verifications. The first set is used during the first execution of the tasks until the first error is encountered, in which case we recover from the last checkpoint and start re-executing the tasks using the second set of verifications with the second speed. The problem is to find the best positions for the verified checkpoints as well as the best positions for the two sets of single verifications in order to minimize the total execution time (TIME-VC+V) or the total energy consumption (ENERGY-VC+V).

Because two sets of single verifications need to be placed, we formulate two independent dynamic programs to determine their respective optimal positions. The overall complexity, however, remains the same as that of the SINGLESPEED scenario.

5.2.1. TIME-VC+V: MINIMIZING MAKESPAN

THEOREM 5.3. *For the REEXECSPEED scenario, the TIME-VC+V problem can be solved by a dynamic programming algorithm in $O(n^3)$ time.*

PROOF. The VC+V approach follows the same reasoning as the VC-ONLY approach (see Section 5.1.1). Here, we replace $Time_{C_{re}}(j, s, \sigma)$ with $Time_{VC_{re}}(j, s, \sigma)$, and replace $T_{C_{re}}(i, j, s, \sigma)$ with $T_{VC_{re}}(i, j, s, \sigma)$. Note that both expressions follow the new re-execution model and account for both sets of single verifications. The goal is to find $Time_{VC_{re}}(n, s, \sigma)$, and the dynamic program is formulated as:

$$Time_{VC_{re}}(j, s, \sigma) = \min_{0 \leq i < j} \{Time_{VC_{re}}(i, s, \sigma) + T_{VC_{re}}(i+1, j, s, \sigma)\} + C_j.$$

Note that the checkpointing cost after T_j is included in $Time_{VC_{re}}(j, s, \sigma)$ but not in $T_{VC_{re}}(i, j, s, \sigma)$. We initialize the dynamic program with $Time_{VC_{re}}(0, s, \sigma) = 0$.

To compute $T_{VC_{re}}(i, j, s, \sigma)$, where both T_{i-1} and T_j are verified and checkpointed, we again distinguish two parts: the optimal expected time $Time_{V_{first}}(i, j, s)$ to execute the tasks T_i to T_j in the first execution using the first set of single verifications with speed s , and the optimal expected time $Time_V(i, j, \sigma)$ to successfully execute the tasks T_i to T_j in all subsequent re-executions using the second set of single verifications with speed σ . Similarly to the proof of the TIME-VC-ONLY problem in Section 5.1.1, $T_{VC_{re}}(i, j, s, \sigma)$ always includes the cost $Time_{V_{first}}(i, j, s)$ regardless of whether an error strikes during the first execution. Let $p_{i,j}^{Err}(s)$ denote the probability that at least one error is detected in the first execution, and it is again given by $p_{i,j}^{Err}(s) = p_{i,j}^F(s) + (1 - p_{i,j}^F(s))p_{i,j}^S(s)$. If an error indeed strikes during the first execution, then we need to recover from the last checkpoint and use $Time_V(i, j, \sigma)$ to re-execute all the tasks from T_i to T_j with the

second speed until we pass the next checkpoint. Therefore,

$$T_{VC_{re}}(i, j, s, \sigma) = Time_{V_{first}}(i, j, s) + p_{i,j}^{Err}(s) (R_{i-1} + Time_V(i, j, \sigma)) . \quad (16)$$

Here, $Time_V(i, j, \sigma)$ follows the same dynamic programming formulation as in Section 4.2.1 but using speed σ . $Time_{V_{first}}(i, j, s)$, on the other hand, denotes the optimal expected time to execute the tasks T_i to T_j at speed s until the first error strikes. Hence, it should not include the recovery cost nor the re-executions. The following describes a dynamic programming formulation to compute $Time_{V_{first}}(i, j, s)$:

$$Time_{V_{first}}(i, j, s) = \min_{i-1 \leq l < j} \left\{ Time_{V_{first}}(i, l, s) + (1 - p_{i,l}^{Err}(s)) (T_{V_{first}}^{SF}(l+1, j, s)) \right\} ,$$

where $p_{i,l}^{Err}(s) = p_{i,l}^F(s) + (1 - p_{i,l}^F(s))p_{i,l}^S(s)$ is the probability that at least one error is detected when executing the tasks T_i to T_l , and $T_{V_{first}}^{SF}(i, j, s)$ denotes the expected time to execute the tasks T_i to T_j with both T_{i-1} and T_j verified. In this formulation, we include the second term only when no error has happened during the first term, otherwise we have to recover and re-execute the tasks with the second speed, which is handled by $Time_V(i, j, \sigma)$.

To compute $T_{V_{first}}^{SF}(i, j, s)$, we consider two possible scenarios: either a fail-stop error has occurred and we lose $T_{lost_{i,j}}(s)$ time, or there was no fail-stop error and we have to check whether a silent error has occurred. Recall that we do not account for the re-executions. Therefore,

$$T_{V_{first}}^{SF}(i, j, s) = (1 - p_{i,j}^F(s)) (T_{i,j}(s) + V_j(s)) + p_{i,j}^F(s) \left(\frac{1}{\lambda^F(s)} - \frac{T_{i,j}(s)}{e^{\lambda^F(s)T_{i,j}(s)} - 1} \right) .$$

Finally, we initialize this dynamic program with $Time_{V_{first}}(i, i-1, s) = 0$ for all $i = 1, \dots, n$.

The complexity is dominated by the computation of $Time_V(i, j, s)$ and $Time_{V_{first}}(i, j, s)$, both of which take $O(n^3)$ time. Therefore, the overall complexity remains the same as in the SINGLESPEED scenario (Theorem 4.5). \square

5.2.2. ENERGY-VC+V: MINIMIZING ENERGY

PROPOSITION 5.4. *For the REEXECSPPEED scenario, the ENERGY-VC+V problem can be solved by a dynamic programming algorithm in $O(n^3)$ time.*

PROOF. The proof is similar to that of the TIME-VC+V problem in Section 5.2.1. We replace $Time_{VC_{re}}(j, s, \sigma)$ with $Energy_{VC_{re}}(j, s, \sigma)$, and replace $T_{VC_{re}}(i, j, s, \sigma)$ with $E_{VC_{re}}(i, j, s, \sigma)$. Note that both expressions account for the two sets of single verifications with the corresponding speed for each set, as it was done for the makespan problem. The goal is to find $Energy_{VC_{re}}(n, s, \sigma)$, and the dynamic program is formulated as:

$$Energy_{VC_{re}}(j, s, \sigma) = \min_{0 \leq i < j} \{ Energy_{VC_{re}}(i, s, \sigma) + E_{VC_{re}}(i+1, j, s, \sigma) \} + E_j^C .$$

As with the minimization of the makespan, we compute $E_{VC_{re}}(i, j, s, \sigma)$ by breaking it into two parts: the optimal expected energy $Energy_{V_{first}}(i, j, s)$ to execute the tasks in the first execution and the optimal expected energy $Energy_V(i, j, \sigma)$ to successfully execute the tasks in the re-executions. Hence, analogous to Equation (16), $E_{VC_{re}}(i, j, s, \sigma)$ can be expressed as

$$E_{VC_{re}}(i, j, s, \sigma) = Energy_{V_{first}}(i, j, s) + p_{i,j}^{Err}(s) \left(E_{i-1}^R + Energy_V(i, j, \sigma) \right) . \quad (17)$$

Since the first execution does not include for any recovery cost, the optimal energy $Energy_{V_{first}}(i, j, s)$ spent during this time comes only from CPU and it is directly proportional to the optimal time $Time_{V_{first}}(i, j, s)$. Hence, we can express:

$$Energy_{V_{first}}(i, j, s) = (P_{idle} + P_{cpu}(s))Time_{V_{first}}(i, j, s).$$

For the re-executions, the optimal energy $Energy_V(i, j, \sigma)$ can be obtained by following the same dynamic programming formulation as in Section 4.2.2 but using speed σ . Clearly, the complexity is the same as that of the makespan minimization algorithm (in Theorem 5.3). \square

6. MULTISPEED SCENARIO

In this section, we investigate the most flexible scenario, MULTISPEED, which is built upon the REEXEC SPEED scenario, to get even more control over the execution time and the energy consumption, but at the cost of a higher complexity. Instead of having two fixed speeds, we are given a set $S = \{s_1, s_2, \dots, s_K\}$ of K discrete speeds. We call a sequence of tasks between two checkpoints a *segment* of the chain, and we allow each segment to use one speed for the first execution, and a second speed for all potential re-executions, where both speeds belong to S . The two speeds can well be different for different segments.

6.1. VC-ONLY: Using verified checkpoints only

In this approach, we aim at finding the best positions for the checkpoints, as well as the best speed pair for each segment, in order to minimize the expected makespan (TIME-VC-ONLY) or the total energy consumption (ENERGY-VC-ONLY).

6.1.1. TIME-VC-ONLY: MINIMIZING MAKESPAN

THEOREM 6.1. *For the MULTISPEED scenario, the TIME-VC-ONLY problem can be solved by a dynamic programming algorithm in $O(n^2 K^2)$ time.*

PROOF. The proof is built upon that of Theorem 5.1 in Section 5.1.1. Here, we use $Time_{C_{mul}}(j)$ to denote the optimal expected time to successfully execute tasks T_1 to T_j , where T_j has a verified checkpoint and there are possibly other verified checkpoints in between. Also, we use $T_{C_{mul}}(i, j)$ to denote the optimal expected time to successfully execute the tasks T_i to T_j , where both T_{i-1} and T_j are verified and checkpointed. In both expressions, the two execution speeds for each segment can be arbitrarily chosen from the discrete set S . The goal is to find $Time_{C_{mul}}(n)$, and the dynamic program can be formulated as:

$$Time_{C_{mul}}(j) = \min_{0 \leq i < j} \{Time_{C_{mul}}(i) + T_{C_{mul}}(i+1, j)\} + C_j,$$

which is initialized with $Time_{C_{mul}}(0) = 0$. Recall that $T_{C_{re}}(i, j, s, \sigma)$ from the REEXEC SPEED scenario (see Equation (14)) already accounts for two speeds that are fixed. We can use it to compute $T_{C_{mul}}(i, j)$ by trying all possible speed pairs:

$$T_{C_{mul}}(i, j) = \min_{s, \sigma \in S} T_{C_{re}}(i, j, s, \sigma).$$

The complexity is now dominated by the computation of $T_{C_{mul}}(i, j)$ for all (i, j) pairs with $i \leq j$, and it takes $O(n^2 K^2)$ time. After $T_{C_{mul}}(i, j)$ is computed, the dynamic programming table can then be filled in $O(n^2)$ time. \square

6.1.2. ENERGY-VC-ONLY: MINIMIZING ENERGY

PROPOSITION 6.2. *For the MULTISPEED scenario, the ENERGY-VC-ONLY problem can be solved by a dynamic programming algorithm in $O(n^2 K^2)$ time.*

PROOF. The proof is similar to that of the TIME-VC-ONLY problem in Section 6.1.1. We replace $Time_{C_{mul}}(j)$ with $Energy_{C_{mul}}(j)$ and replace $T_{C_{mul}}(i, j)$ with $E_{C_{mul}}(i, j)$. The goal is to find $Energy_{C_{mul}}(n)$ and the dynamic program is formulated as:

$$Energy_{C_{mul}}(j) = \min_{0 \leq i < j} \{Energy_{C_{mul}}(i) + E_{C_{mul}}(i+1, j)\} + E_j^C.$$

Similarly, $E_{C_{mul}}(i, j)$ can be computed from $E_{C_{re}}(i, j, s, \sigma)$ (see Equation (15)) by trying all possible speed pairs :

$$E_{C_{mul}}(i, j) = \min_{s, \sigma \in S} E_{C_{re}}(i, j, s, \sigma).$$

Clearly, the complexity is the same as that of the makespan minimization algorithm (in Theorem 6.1). \square

6.2. VC+V: Using verified checkpoints and single verifications

In this approach, we aim at finding the best positions for the checkpoints and verifications, as well as the best speed pair for each segment, in order to minimize the expected makespan (TIME-VC-ONLY) or the total energy consumption (ENERGY-VC-ONLY).

6.2.1. TIME-VC+V: MINIMIZING MAKESPAN

THEOREM 6.3. *For the MULTISPEED scenario, the TIME-VC+V problem can be solved by a dynamic programming algorithm in $O(n^3 K^2)$ time.*

PROOF. The proof is similar to that of the TIME-VC-ONLY problem in Theorem 6.1. Here, we replace $Time_{C_{mul}}(j)$ with $Time_{VC_{mul}}(j)$ and replace $T_{C_{mul}}(i, j)$ with $T_{VC_{mul}}(i, j)$. Again, the two expressions denote the optimal execution times with the best speed pair chosen from S for each segment. The goal is to find $Time_{VC_{mul}}(n)$, and the dynamic program is formulated as:

$$Time_{VC_{mul}}(j) = \min_{0 \leq i < j} \{Time_{VC_{mul}}(i) + T_{VC_{mul}}(i+1, j)\} + C_j,$$

which is initialized with $Time_{VC_{mul}}(0) = 0$. We can compute $T_{VC_{mul}}(i, j)$ from $T_{VC_{re}}(i, j, s, \sigma)$ (see Equation (16)) by trying all possible speed pairs:

$$T_{VC_{mul}}(i, j) = \min_{s, \sigma \in S} T_{VC_{re}}(i, j, s, \sigma).$$

The complexity is still dominated by the computation of $T_{VC_{mul}}(i, j)$, which amounts to computing $T_{VC_{re}}(i, j, s, \sigma)$ for all (i, j) pairs and for K^2 possible pairs of speeds (see Theorem 5.3). Therefore, the overall complexity is $O(n^3 K^2)$. \square

6.2.2. ENERGY-VC+V: MINIMIZING ENERGY

PROPOSITION 6.4. *For the MULTISPEED scenario, the ENERGY-VC+V problem can be solved by a dynamic programming algorithm in $O(n^3 K^2)$ time.*

PROOF. The proof follows that of the TIME-VC+V problem in Section 6.2.1. We replace $Time_{VC_{mul}}(j)$ with $Energy_{VC_{mul}}(j)$ and replace $T_{VC_{mul}}(i, j)$ with $E_{VC_{mul}}(i, j)$. The goal is to find $Energy_{VC_{mul}}(n)$, and the dynamic program is formulated as:

$$Energy_{VC_{mul}}(j) = \min_{0 \leq i < j} \{Energy_{VC_{mul}}(i) + E_{VC_{mul}}(i+1, j)\} + E_j^C,$$

where $E_{VC_{mul}}(i, j)$ is computed from $E_{VC_{re}}(i, j, s, \sigma)$ (see Equation (17)) by trying all possible speed pairs:

$$E_{VC_{mul}}(i, j) = \min_{s, \sigma \in S} E_{VC_{re}}(i, j, s, \sigma).$$

Clearly, the complexity is the same as that of the makespan minimization algorithm (in Theorem 6.3). \square

7. EXPERIMENTS

We conduct simulations to evaluate the performance of the dynamic programming algorithms under different execution scenarios and parameter settings. We instantiate the model parameters with realistic values taken from the literature, and we point out that the code for all algorithms and simulations is publicly available at <http://graal.ens-lyon.fr/~yrobert/failstop-silent>, so that interested readers can build relevant scenarios of their choice.

7.1. Simulation settings

We generate linear chains with different number n of tasks while keeping the total computational cost at $W = 5 \times 10^4$ seconds ≈ 14 hours. The total amount of computation is distributed among the tasks in three different patterns: (1) *Uniform*, all tasks share the same cost $\frac{W}{n}$, as in matrix multiplication or in some iterative stencil kernels; (2) *Decrease*, task T_i has cost $\alpha \cdot (n+1-i)^2$, where $\alpha \approx \frac{3W}{n^3}$. This quadratically decreasing function resembles some dense matrix solvers, e.g., by using LU or QR factorization. (3) *HighLow*, a set of identical tasks with large cost is followed by tasks with small cost. This distribution is created to distinguish the performance of different execution scenarios. In the default setting, we assume that 10% of the tasks are large and they contain 60% of the total computational cost. We will also vary these parameters to evaluate their impacts on the performance.

We adopt the set of speeds from the Intel Xscale processor. Following [Rizvandi et al. 2012], the normalized speeds are $\{0.15, 0.4, 0.6, 0.8, 1\}$ and the fitted power function is given by $P(s) = 1550s^3 + 60$. From the discussion in Section 2.3, we assume the following model for the average error rate of fail-stop errors:

$$\lambda^F(s) = \lambda_{\text{ref}}^F \cdot 10^{\frac{d \cdot |s_{\text{ref}} - s|}{s_{\text{max}} - s_{\text{min}}}}, \quad (18)$$

where $s_{\text{ref}} \in [s_{\text{min}}, s_{\text{max}}]$ denotes the reference speed with the lowest error rate λ_{ref}^F among all possible speeds in the range. The above equation allows us to account for higher fail-stop error rates when the CPU speed is either too low or too high. In the simulations, the reference speed is set to be $s_{\text{ref}} = 0.6$ with an error rate of $\lambda_{\text{ref}}^F = 10^{-5}$ for fail-stop errors, and the sensitivity parameter is set to be $d = 3$. These parameters represent realistic settings reported in the literature [Assayad et al. 2013; Aupy et al. 2012; Zhao et al. 2008], and they correspond to $0.83 \sim 129$ errors over the entire chain of computation depending on the processing speed chosen.

For silent errors, we assume that its error rate is related to that of the fail-stop errors as $\lambda^S(s) = \eta \cdot \lambda^F(s)$, where $\eta > 0$ is the relative parameter. To achieve realistic scenarios, we try to vary η to assess the impact of both error sources on the performance. However, we point out that our approach is completely independent of the evolution of the error rates as a function of the speed. In a practical setting, we are given a set of discrete speeds and two error rates for each speed, one for fail-stop errors and one for silent errors. This is enough to instantiate our model.

In addition, we define cr to be the ratio between the checkpointing/recovery cost and the computational cost for the tasks, and define vr to be the ratio between the verification cost and the computational cost. By default, we execute the tasks using the reference speed s_{ref} , and we set $\eta = 1$, $cr = 1$ and $vr = 0.01$. This setting corresponds to the case where fail-stop and silent errors have similar probabilities to strike the system. Moreover, the tasks have costly checkpoints (same order of magnitude as the computational costs) and lightweight verifications (average cost 1% of computational costs); examples of such tasks are data-oriented kernels processing large files




Fig. 5. Impact of n and cost distribution on the performance of the TIME-VC-ONLY and TIME-VC+V algorithms. In (b), the thick bars represent the number of verified checkpoints and the yellow thin bars represent the total number of verifications. In (c), the number of tasks is fixed at $n = 100$. The long vertical bars mark the positions of the verified checkpoints within the task chain, whereas the short vertical bars mark the positions of the additional verifications.

and checksumming for verification. We will also vary these parameters to study their impacts.

7.2. Results

7.2.1. SINGLESPEED scenario for makespan. The first set of experiments is devoted to the evaluation of the time-optimal algorithms in the SINGLESPEED scenario.

Impact of number of tasks and cost distribution. Figure 5(a) shows the expected makespan (normalized by the error-free execution time at the reference speed, i.e., $\frac{W}{s_{\text{ref}}}$) with different number n of tasks and different cost distributions. The results show that having more tasks reduces the expected makespan (for a fixed total computation), since it enables the algorithms to place more checkpoints and verifications, as can be seen in Figure 5(b). The distribution that renders a larger variation in the costs of the tasks create more difficulty in the placement of checkpoints/verifications, thus resulting in worse makespan. Figure 5(c) shows, for $n = 100$ tasks, that the checkpoints and verifications are placed evenly within the task chain for the *Uniform* distribution, while for *Decrease* and *HighLow* distributions, they are placed more densely at the beginning, where large tasks need to be checkpointed and/or verified for better resilience. The figure also compares the performance of the TIME-VC-ONLY algorithm with that of TIME-VC+V. The latter algorithm, being more flexible, naturally leads to improved makespan under all cost distributions. Because of the additionally placed verifications, it also reduces the number of checkpoints in the optimal solution.

In the rest of this section, we will mainly focus on the evaluation of the TIME-VC+V algorithm.

Impact of error mode and relative ratio. Figure 6(a) compares the performance of the TIME-VC+V algorithm under different error modes, namely, fail-stop (F) only, silent (S) only, and fail-stop plus silent with different values of η . As silent errors are harder to detect and hence to deal with, the S-only case leads to larger makespan than the F-only case. In the presence of both types of errors, the makespan becomes worse with larger η , i.e., with increased rate for silent errors, despite the algorithm's effort to place more checkpoints and more verifications as shown in Figure 6(b). Similar results (not shown) are also observed for the other cost distributions.

In the subsequent experiments, we will concentrate on $n = 100$ tasks in the presence of both fail-stop and silent errors with $\eta = 1$.




Fig. 6. Impact of η on the performance. F denotes fail-stop error only and S denotes silent error only.

Impact of checkpointing and verification ratios. Figure 7(a) presents the impact of checkpointing/recovery ratio (cr) and verification ratio (vr) on the performance of the TIME-VC+V algorithm under different CPU speeds. For a given speed, small cr (or vr) enables the algorithm to place more checkpoints (or verifications), which leads to better makespan. Moreover, the performance degrades significantly as the CPU speed is set below the reference speed s_{ref} for the error rate increases exponentially. A higher CPU speed also increases the error rate, but it improves the makespan, at least for small cr , by executing the tasks faster with more checkpoints. Finally, if checkpointing cost is on par with verification cost (e.g., $cr = 0.1$), reducing the verification cost can additionally increase the number of checkpoints (e.g., at $s = 0.6$), since each checkpoint also has a verification cost associated with it. For high checkpointing cost, however, reducing the verification cost no longer influences the algorithm's checkpointing decisions.




Fig. 7. Impact of cr and vr on the performance with different CPU speeds. Speed $s = 0.15$ yields extremely large makespan, which is omitted in the figure.

Comparison with divisible load application. Figure 8(a) compares the makespan of the TIME-VC+V algorithm with that of the periodic checkpointing and verification algorithm for a divisible load application. Figure 8(b) also shows the number of checkpoints and verifications placed for the two cases. Note that, in the divisible load application, the total computational cost, the checkpointing cost and the verification cost are set to be the same as the corresponding costs of a discrete task under the *Uniform* distribution. We see that the makespan for uniform tasks is almost identical to that of the divisible load under all CPU speeds, while the performance degrades significantly for the other two distributions with larger variations in the costs of the tasks. Similar results (not shown) are also observed when comparing the TIME-VC-ONLY algorithm with the periodic checkpointing algorithm for divisible load. Moreover, because divisible load does not impose restrictions in the checkpointing and verification positions,




Fig. 8. Performance comparison of the TIME-VC+V algorithm and the periodic checkpointing and verification algorithm for divisible load application.

there tends to be more verifications (or checkpoints in the case of periodic checkpointing algorithm) than for discrete tasks, especially when the CPU speed is further away from the reference speed and hence the error rate is high.

In view of these results, we could imagine the following greedy algorithm as an alternative to the TIME-VC-ONLY and TIME-VC+V algorithms for a linear chain of tasks with *Uniform* cost: position the next checkpoint or verification as soon as the time spent on computing since the last checkpoint or verification exceeds the optimal periods given by Propositions 4.2 and 4.6. The results here suggest that this linear-time algorithm (with complexity $O(n)$) would give a good approximation to the optimal solution (returned by the TIME-VC-ONLY algorithm with complexity $O(n^2)$ or the TIME-VC+V algorithm with complexity $O(n^3)$).

Performance with independent checkpointing cost. We now consider the case where the checkpointing costs are independent of the tasks' computational costs. To assess the impact, we generate different patterns by varying the checkpointing costs linearly within the task chain while keeping the sum a constant. Specifically, we use the following function to generate the checkpointing cost $C_i = C \left(1 + \frac{\delta}{n} \left(i - \frac{n}{2}\right)\right)$ for each task T_i , where $1 \leq i \leq n$ and $C = \frac{W}{n}$ denotes the average checkpointing cost. Here, δ represents the gradient of the linear function: $\delta = 0$ means all tasks have the same checkpointing cost, $\delta > 0$ increases the checkpointing cost as more tasks are processed, and $\delta < 0$ decreases the checkpointing cost.

Figure 9 shows the performance of the TIME-VC+V algorithm for different values of δ under the reference speed $s_{\text{ref}} = 0.6$. When δ increases, the makespan for tasks with *Uniform* computational cost is barely affected, while significant improvements are observed for the other two distributions. The reason is because the algorithm for the *Decrease* and the *HighLow* distributions places more checkpoints at the beginning of the task chain, where the large tasks are located. Therefore, reducing the checkpointing costs for these tasks decreases the execution overhead and hence the overall makespan.

7.2.2. SINGLESPEED scenario for energy. This set of experiments evaluates the energy-optimal algorithms in the SINGLESPEED scenario. The default power parameters are set to be $P_{\text{idle}} = 60$ and $P_{\text{cpu}}(s) = 1550s^3$ according to [Rizvandi et al. 2012]. The dynamic power consumption P_{io} due to I/O is equal to the dynamic power of the CPU at the lowest discrete speed 0.15. We will vary these parameters to study their impacts.

Impact of number of tasks and cost distribution. Figure 10(a) shows the expected energy (normalized by the error-free energy at the reference speed, i.e., $\frac{W(P_{\text{idle}} + P_{\text{cpu}}(s_{\text{ref}}))}{s_{\text{ref}}}$). As with the time-optimal algorithms (see Figure 5), a larger number of tasks improves




Fig. 9. Impact of δ on the performance when the checkpointing cost for task T_i is given by $C_i = C \left(1 + \frac{\delta}{n} \left(i - \frac{n}{2}\right)\right)$, where $C = \frac{W}{n}$.




Fig. 10. Impact of n and cost distribution on the performance of the ENERGY-VC-ONLY and ENERGY-VC+V algorithms. In (c), the number of tasks is fixed at $n = 100$.

the performance, while a larger variation in the tasks' costs worsens the performance. Unlike the makespan case, the performance difference between ENERGY-VC-ONLY and ENERGY-VC+V is less evident. The reason is because the checkpointing cost is much smaller in terms of energy consumption, so more checkpoints are placed (compare Figures 10(b) and 5(b)), and this reduces the number of additional verifications required and hence their benefits.

Comparison with divisible load application. Figure 11 compares the performance of the ENERGY-VC+V algorithm with that of the periodic checkpointing and verification algorithm for a divisible load application when executing $n = 100$ tasks. Similarly to the makespan case (see Figure 8), the energy consumed for tasks with *Uniform* distribution is very close to that for the divisible load, which admits more verifications and/or checkpoints due to the flexible task model. In addition, the optimal energy is achieved by setting the CPU speed below the reference, i.e., at $s = 0.4$. This is in contrast to the TIME-VC+V algorithm, which achieves the optimal makespan at $s = 0.8$, a higher speed than the reference.

Energy-makespan tradeoff. We now study the energy-makespan tradeoff exhibited by the ENERGY-VC+V and TIME-VC+V algorithms. Figure 12 compares their relative performance for makespan and energy when executing $n = 100$ tasks. At speed $s = 0.15$, the power consumed by the CPU is identical to that by I/O. This yields the same number of checkpoints placed by the two algorithms, which in turn leads to the same performance for makespan and energy. As the speed s increases, the I/O power becomes relatively cheaper, so ENERGY-VC+V tends to place more checkpoints to improve the energy at the expense of makespan. Since the makespan of TIME-VC+V




Fig. 11. Performance comparison of the ENERGY-VC+V algorithm and the periodic checkpointing and verification algorithm for divisible load application.




Fig. 12. Relative performance of the ENERGY-VC+V and TIME-VC+V algorithms with different CPU speeds.

degrades at speed $s = 1$ under the default setting (see Figure 8(a)), this diminishes its makespan advantage at the highest discrete speed. Figure 12 indicates that the TIME-VC+V algorithm running at speed $s = 0.8$ offers a good energy-makespan tradeoff: compared to the ENERGY-VC+V algorithm, it provides more than 25% improvement in makespan with about 10% degradation in energy under the default settings.

Impact of P_{idle} and P_{io} . Figures 13(a) and 13(b) show the relative performance of TIME-VC+V and ENERGY-VC+V by varying P_{idle} and P_{io} separately according to the dynamic power function $1550s^3$, while keeping the other one at the smallest CPU power, i.e., $1550 \cdot 0.15^3$. The CPU speed is fixed at $s = 0.6$. Figure 13(c) further shows the number of checkpoints placed by the ENERGY-VC+V algorithm at different P_{idle} and P_{io} values. (The TIME-VC+V algorithm is not affected by these two parameters and always places 11 checkpoints in this experiment.) First, setting the smallest value for both parameters creates a big gap between the CPU and I/O power consumptions. This leads to a large number of checkpoints placed by the ENERGY-VC+V algorithm. Increasing P_{idle} closes this gap and hence reduces the number of checkpoints, which leads to the performance convergence of the two algorithms. While increasing P_{io} has the same effect, a larger value than $P_{cpu} = 1550 \cdot 0.6^3$ further reduces the number of checkpoints below 11, since checkpointing is now less energy-efficient. This again gives the ENERGY-VC+V algorithm advantage in terms of energy.

7.2.3. REEXEC SPEED and MULTISPEED scenarios. This set of experiments evaluates the REEXEC SPEED and MULTISPEED scenarios for makespan and energy. To distinguish them from the SINGLESPEED scenario, we consider the *HighLow* distribution, which yields a larger variance among the computational costs of the tasks. In the simulation, we again focus on the VC+V algorithms for $n = 100$ tasks, and vary the *cost ratio* γ ,




Fig. 13. Impact of P_{idle} and P_{io} on the relative performance of the ENERGY-VC+V and TIME-VC+V algorithms at the reference speed $s = 0.6$. The number of checkpoints placed by ENERGY-VC+V with different P_{io} , P_{idle} values ($= 1550s^3$) is shown in (c), while TIME-VC+V always places 11 checkpoints in this experiment.

which is the percentage of the large tasks' computational cost in the total computational cost.

Comparing different execution scenarios. Figure 14(a) compares the makespan of the TIME-VC+V algorithms under the three execution scenarios. For the SINGLESPEED and REEXECSPPEED scenarios, only $s = 0.6$ and $s = 0.8$ are drawn, since the other speeds lead to much larger makespan. For small cost ratios, no task has a very large computational cost, so the distribution is close to *Uniform*. In this case, the faster speed $s = 0.8$ offers the best performance despite its higher error rate, as we have seen in Figure 8(a). When the cost ratio increases, tasks with larger cost start to emerge. At the high error rate of $s = 0.8$, these tasks will experience much more errors and re-executions, so their execution time will dominate the overall makespan. Therefore, for large cost ratios, $s = 0.6$ becomes the best speed due to its smaller error rate, which was also observed in Figure 8(a) under the *HighLow* distribution.

In the REEXECSPPEED scenario, we observe that the best re-execution speed σ , regardless of the initial speed s , is also determined by the computational costs of the tasks, which are in turn decided by the cost ratio. Figure 14(b) shows, for cost ratio $\gamma = 0.6$, that setting $\sigma = 0.6$ improves the execution of the big tasks but degrades the performance of the small tasks. On the other hand, setting $\sigma = 0.8$ helps the small tasks but hurts the big tasks. These experiments suggest that, despite the ability to select a more appropriate speed for the reexecutions, the REEXECSPPEED scenario presents only limited benefits compared to the best performance achievable in the SINGLESPEED scenario. The MULTISPEED scenario, with its flexibility to choose different speeds depending on the costs of the tasks, offers clear performance gains. The advantage is especially evident at medium cost ratio, where a good mix of large and small tasks coexist — a situation that is very hard to cope with by using fixed speed(s).

Similar results can also be observed for the ENERGY-VC+V algorithms under the three scenarios, which are shown in Figure 15. Note that, in terms of energy consumption, speed $s = 0.4$ is more suitable for small tasks due to its better power efficiency, while speed $s = 0.6$ is more suitable for big tasks due to its lower error rate. Again, the most flexible MULTISPEED scenario is able to choose between the two speeds depending on the costs of the tasks, and hence it offers the best overall performance.

Energy-makespan tradeoff under MULTISPEED scenario. Figure 16 shows the relative performance of the ENERGY-VC+V and TIME-VC+V algorithms under the MULTISPEED scenario. As small cost ratios favor speed $s = 0.4$ for energy and $s = 0.8$ for makespan, it distinguishes the two algorithms in terms of their respective optimiza-




Fig. 14. Performance comparison of the TIME-VC+V algorithms in MULTISPEED, REEXECSPEED and SINGLESPEED scenarios for $n = 100$ tasks under *HighLow* cost distribution. In (b), the cost ratio is fixed at $\gamma = 0.6$, and the yellow part at the bottom of each bar represents the expected execution time for the large tasks.




Fig. 15. Performance comparison of the ENERGY-VC+V algorithms in MULTISPEED, REEXECSPEED and SINGLESPEED scenarios for $n = 100$ tasks with *HighLow* cost distribution. In (b), the cost ratio is fixed at $\gamma = 0.6$, and the yellow part at the bottom of each bar represents the expected energy for the large tasks.




Fig. 16. Impact of cost ratio on the relative performance of the ENERGY-VC+V and TIME-VC+V algorithms under the MULTISPEED scenario.

tion objectives, by up to 100% in makespan and even more in energy consumption. Increasing the cost ratio creates more computationally demanding tasks, which need to be executed at speed 0.6 for both makespan and energy efficiency as it incurs fewer errors. This closes the performance gap of the two algorithms as well as the number of checkpoints placed by them. For both algorithms, the number of checkpoints reduces with the cost ratio, because the total computational cost in the small tasks shrinks, thus fewer checkpoints are needed among them.

7.2.4. Summary. We have evaluated and compared various algorithms under different execution scenarios and parameter settings. In general, the algorithms under the most

flexible VC+V and MULTISPEED scenario provide the best overall performance, which in practice translates to shorter makespan or lower energy consumption.

For tasks with similar computational cost as in the *Uniform* distribution, the SINGLESPEED algorithm, or the greedy approximation in the context of divisible load application, could provide comparable solutions with lower computational complexity. The REEXEC SPEED algorithms provide marginal benefit compared to SINGLESPEED, but clear performance gains are observed from the MULTISPEED algorithms, especially for tasks with very different costs. The results also show that the optimal solutions are often achieved by processing around the reference speed that yields the least number of failures.

Finally, for the complexity of computing the optimal solutions, we point out that application workflows rarely exceed a few tens of tasks. In such practical contexts, even the most advanced algorithms have a very fast execution time, of a few seconds. To give a number, VC+V in the MULTISPEED scenario requires less than one second for 100 tasks, using a 3.7Ghz single-core. Hence all our algorithms can be applied to determine the optimal checkpointing and verification locations with almost negligible cost.

8. CONCLUSION

In this paper, we have presented a general-purpose solution that combines checkpointing and verification mechanisms to cope with both fail-stop errors and silent data corruptions. By using dynamic programming, we have devised polynomial-time algorithms that decide the optimal checkpointing and verification positions for a linear chain of tasks. The algorithms can be applied to several execution scenarios to minimize the expected execution time (makespan) or energy consumption. Moreover, we have extended the classical bound of Young/Daly for a divisible load application to handle both fail-stop and silent errors, and to include both checkpoints and verifications. The results are supported by a set of extensive simulations, which demonstrate the quality and tradeoff of our optimal algorithms under a wide range of parameter settings. One future direction is to extend our study from linear chains to other application workflows, such as tree graphs, fork-join graphs, series-parallel graphs, or even general directed acyclic graphs (DAGs).

References

- Ismail Assayad, Alain Girault, and Hamoudi Kalla. 2013. Tradeoff exploration between reliability, power consumption, and execution time for embedded systems. *International Journal on Software Tools for Technology Transfer* 15, 3 (2013), 229–245.
- Guillaume Aupy, Anne Benoit, Thomas Hérault, Yves Robert, Frédéric Vivien, and Dounia Zaidouni. 2013. On the combination of silent error detection and checkpointing. In *Proceedings of the 19th IEEE Pacific Rim International Symposium on Dependable Computing (PRDC)*. 11–20.
- Guillaume Aupy, Anne Benoit, and Yves Robert. 2012. Energy-aware scheduling under reliability and makespan constraints. In *Proceedings of the International Conference on High Performance Computing (HiPC)*. 1–10.
- Nikhil Bansal, Tracy Kimbrel, and Kirk Pruhs. 2007. Speed Scaling to Manage Energy and Temperature. *J. ACM* 54, 1 (2007), 3:1–3:39.
- Austin R. Benson, Sven Schmit, and Robert Schreiber. 2014. Silent error detection in numerical time-stepping schemes. *The International Journal of High Performance Computing Applications* DOI: 10.1177/1094342014532297 (2014).
- George Bosilca, Rémi Delmas, Jack Dongarra, and Julien Langou. 2009. Algorithm-based fault tolerance applied to high performance computing. *J. Parallel Distrib. Comput.* 69, 4 (2009), 410–416.
- Marin Bougeret, Henri Casanova, Mikael Rabie, Yves Robert, and Frédéric Vivien. 2011. Checkpointing strategies for parallel jobs. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. 1–11.

- Greg Bronevetsky and Bronis de Supinski. 2008. Soft error vulnerability of iterative linear algebra methods. In *Proceedings of the International Conference on Supercomputing (ICS)*. 155–164.
- K. M. Chandy and L. Lamport. 1985. Distributed Snapshots: Determining Global States of Distributed Systems. *ACM Transactions on Computer Systems* 3, 1 (1985), 63–75.
- Zizhong Chen. 2013. 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)*. 167–176.
- J. T. Daly. 2006. A higher order estimate of the optimum checkpoint interval for restart dumps. *Future Generation Comp. Syst.* 22, 3 (2006), 303–312.
- A. Das, A. Kumar, B. Veeravalli, C. Bolchini, and A. Miele. 2014. Combined DVFS and Mapping Exploration for Lifetime and Soft-error Susceptibility Improvement in MPSoCs. In *Proceedings of the Conference on Design, Automation & Test in Europe (DATE)*. 61:1–61:6.
- Anand Dixit and Alan Wood. 2011. The impact of new technology on soft error rates. In *IEEE International on Reliability Physics Symposium (IRPS)*. 5B.4.1–5B.4.7.
- Nosayba El-Sayed, Ioan A. Stefanovici, George Amvrosiadis, Andy A. Hwang, and Bianca Schroeder. 2012. Temperature Management in Data Centers: Why Some (Might) Like It Hot. *SIGMETRICS Perform. Eval. Rev.* 40, 1 (2012), 163–174.
- James Elliott, Kishor Kharbas, David Fiala, Frank Mueller, Kurt Ferreira, and Christian Engelmann. 2012. Combining partial redundancy and checkpointing for HPC. In *Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS)*. 615–626.
- E. N. Mootaz Elnozahy, Lorenzo Alvisi, Yi-Min Wang, and David B. Johnson. 2002. A survey of rollback-recovery protocols in message-passing systems. *ACM Computing Survey* 34 (2002), 375–408. Issue 3.
- Wu-Chun Feng. 2003. Making a Case for Efficient Supercomputing. *Queue* 1, 7 (Oct. 2003), 54–64.
- David Fiala, Frank Mueller, Christian Engelmann, Rolf Riesen, Kurt Ferreira, and Ron Brightwell. 2012. Detection and correction of silent data corruption for large-scale high-performance computing. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis (SC)*. 78.
- M.A. Heroux and M Hoemmen. 2011. *Fault-tolerant iterative methods via selective reliability*. Research report SAND2011-3915 C. Sandia National Laboratories.
- C.-H Hsu and W.-C Feng. 2005. A Power-Aware Run-Time System for High-Performance Computing. In *Proceedings of the ACM/IEEE Supercomputing Conference (SC)*. 1–9.
- Kuang-Hua Huang and J. A. Abraham. 1984. Algorithm-Based Fault Tolerance for Matrix Operations. *IEEE Trans. Comput.* 33, 6 (1984), 518–528.
- Andy A. Hwang, Ioan A. Stefanovici, and Bianca Schroeder. 2012. Cosmic rays don't strike twice: understanding the nature of DRAM errors and the implications for system design. *SIGARCH Comput. Archit. News* 40, 1 (2012), 111–122.
- Guoming Lu, Ziming Zheng, and Andrew A. Chien. 2013. When is Multi-version Checkpointing Needed?. In *Proceedings of the 3rd Workshop on Fault-tolerance for HPC at extreme scale (FTXS)*. 49–56.
- R. E. Lyons and W. Vanderkulk. 1962. The use of triple-modular redundancy to improve computer reliability. *IBM J. Res. Dev.* 6, 2 (1962), 200–209.
- Adam Moody, Greg Bronevetsky, Kathryn Mohror, and Bronis R. de Supinski. 2010. Design, Modeling, and Evaluation of a Scalable Multi-level Checkpointing System. In *Proc. of the ACM/IEEE SC Conf.* 1–11.
- Xiang Ni, Esteban Meneses, Nikhil Jain, and Laxmikant V. Kalé. 2013. ACR: Automatic Checkpoint/Restart for Soft and Hard Error Protection. In *Proc. Int. Conf. High Performance Computing, Networking, Storage and Analysis (SC '13)*. ACM.
- T.J. O’Gorman. 1994. The effect of cosmic rays on the soft error rate of a DRAM at ground level. *IEEE Trans. Electron Devices* 41, 4 (1994), 553–557.
- T. Ozaki, T. Dohi, H. Okamura, and N. Kaio. 2006. Distribution-free checkpoint placement algorithms based on min-max principle. *IEEE Transactions on Dependable and Secure Computing* 3, 2 (2006), 130–140.
- M.K Patterson. 2008. The effect of data center temperature on energy efficiency. In *Proceedings of 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*. 1167–1174.
- Nikzad Babaii Rizvandi, Albert Y. Zomaya, Young Choon Lee, Ali Javadzadeh Bolori, and Javid Taheri. 2012. Multiple Frequency Selection in DVFS-Enabled Processors to Minimize Energy Consumption. In *Energy-Efficient Distributed Computing Systems*, A. Y. Zomaya and Y. C. Lee (Eds.). John Wiley & Sons, Inc., Hoboken, NJ, USA.
- Piyush Sao and Richard Vuduc. 2013. Self-stabilizing Iterative Solvers. In *Proceedings of the Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems (SCALA)*.

- Osman Sarood, Esteban Meneses, and Laxmikant V. Kale. 2013. A ‘Cool’ Way of Improving the Reliability of HPC Machines. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis (SC)*. 58:1–58:12.
- Manu Shantharam, Sowmyalatha Srinivasmurthy, and Padma Raghavan. 2012. Fault Tolerant Preconditioned Conjugate Gradient for Sparse Linear System Solution. In *Proceedings of the ACM International Conference on Supercomputing (ICS)*. 69–78.
- Sam Toueg and Özalp Babaoglu. 1984. On the Optimum Checkpoint Selection Problem. *SIAM J. Comput.* 13, 3 (1984), 630–649.
- F. Yao, A. Demers, and S. Shenker. 1995. A Scheduling Model for Reduced CPU Energy. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science (FOCS)*. 374.
- John W. Young. 1974. A first order approximation to the optimum checkpoint interval. *Comm. of the ACM* 17, 9 (1974), 530–531.
- Baoxian Zhao, Hakan Aydin, and Dakai Zhu. 2008. Reliability-Aware Dynamic Voltage Scaling for Energy-Constrained Real-Time Embedded Systems. In *Proceedings of the IEEE International Conference on Computer Design (ICCD)*. 633–639.
- Dakai Zhu, R. Melhem, and D. Mosse. 2004. The Effects of Energy Management on Reliability in Real-time Embedded Systems. In *Proceedings of the IEEE / ACM International Conference on Computer-Aided Design (ICCAD)*. 35–40.
- J.F. Ziegler, H.P. Muhlfield, C.J. Montrose, H.W. Curtis, T.J. O’Gorman, and J.M. Ross. 1996b. Accelerated testing for cosmic soft-error rate. *IBM J. Res. Dev.* 40, 1 (1996), 51–72.
- J.F. Ziegler, M.E. Nelson, J.D. Shell, R.J. Peterson, C.J. Gelderloos, H.P. Muhlfield, and C.J. Montrose. 1998. Cosmic ray soft error rates of 16-Mb DRAM memory chips. *IEEE Journal of Solid-State Circuits* 33, 2 (1998), 246–252.
- J. F. Ziegler, H. W. Curtis, H. P. Muhlfield, C. J. Montrose, and B. Chin. 1996a. IBM Experiments in Soft Fails in Computer Electronics. *IBM J. Res. Dev.* 40, 1 (1996), 3–18.