



HAL
open science

Checkpointing strategies for a fixed-length execution

Anne Benoit, Lucas Perotin, Yves Robert, Frédéric Vivien

► **To cite this version:**

Anne Benoit, Lucas Perotin, Yves Robert, Frédéric Vivien. Checkpointing strategies for a fixed-length execution. RR-9552, Inria. 2024, pp.26. hal-04668191

HAL Id: hal-04668191

<https://inria.hal.science/hal-04668191v1>

Submitted on 6 Aug 2024

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.



Distributed under a Creative Commons Attribution 4.0 International License

Inria

Checkpointing strategies for a fixed-length execution

Anne Benoit, Lucas Perotin, Yves Robert, Frédéric Vivien

**RESEARCH
REPORT**

N° 9552

July 2024

Project-Team ROMA

ISRN INRIA/RR--9552--FR+ENG

ISSN 0249-6399



Checkpointing strategies for a fixed-length execution

Anne Benoit*, Lucas Perotin[†], Yves Robert[‡], Frédéric Vivien[§]

Project-Team ROMA

Research Report n° 9552 — July 2024 — 26 pages

Abstract: This work considers checkpointing strategies for a parallel application executing on a large-scale platform whose nodes are subject to failures. The application executes for a fixed duration, namely the length of the reservation that it has been granted. We start with small examples that show the difficulty of the problem: it turns out that the optimal checkpointing strategy neither always uses periodic checkpoints nor always takes its last checkpoint exactly at the end of the reservation. Then, we introduce a dynamic heuristic that is periodic and decides for the checkpointing frequency based upon thresholds for the time left; we determine threshold times T_n such that it is best to plan for exactly n checkpoints if the time left (or initially the length of the reservation) is between T_n and T_{n+1} . Next, we use time discretization and design a (complicated) dynamic programming algorithm that computes the optimal solution, without any restriction on the checkpointing strategy. Finally, we report the results of an extensive simulation campaign that shows that the optimal solution is far more efficient than the Young/Daly periodic approach for short or mid-size reservations.

Key-words: Checkpointing, failures, fail-stop errors, fixed-length execution, fixed-size reservation.

* Univ Lyon, EnsL, UCBL, CNRS, Inria, LIP, F-69342, LYON Cedex 07, and Institut Universitaire de France, France

[†] Univ Lyon, EnsL, UCBL, CNRS, Inria, LIP, F-69342, LYON Cedex 07, France.

[‡] Univ Lyon, EnsL, UCBL, CNRS, Inria, LIP, F-69342, LYON Cedex 07, France and University of Tennessee, Knoxville, TN, USA

[§] Univ Lyon, EnsL, UCBL, CNRS, Inria, LIP, F-69342, LYON Cedex 07, France.

**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

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

Stratégies de checkpoint pour une exécution de longueur fixe

Résumé : Ce travail examine les stratégies de checkpoint pour une application parallèle s'exécutant sur une plateforme à grande échelle dont les noeuds sont sujets à des pannes. L'application s'exécute pendant une durée fixe, à savoir la durée de la réservation qui lui a été accordée. Nous commençons par de petits exemples qui montrent la difficulté du problème : il s'avère que la stratégie de checkpoint optimale n'utilise pas toujours des checkpoints périodiques et ne prend pas toujours son dernier checkpoint exactement à la fin de la réservation. Ensuite, nous introduisons une heuristique dynamique qui est périodique et décide de la fréquence des checkpoints en se basant sur des seuils pour le temps restant. Nous déterminons les seuils T_n de telle sorte qu'il est préférable de prévoir exactement n checkpoints si le temps restant (ou initialement la durée de la réservation) est compris entre T_n et T_{n+1} . Ensuite, nous utilisons la discrétisation temporelle et concevons un algorithme de programmation dynamique (compliqué) qui calcule la solution optimale, sans aucune restriction sur la stratégie de checkpoint. Enfin, nous présentons les résultats d'une vaste campagne de simulations qui montre que la solution optimale est beaucoup plus efficace que l'approche périodique de Young/Daly pour les réservations courtes ou moyennes.

Mots-clés : Checkpointing, défaillances, erreurs fatales, exécution de longueur fixe, réservation de taille fixe

Contents

1	Introduction	3
2	Related work	5
2.1	Checkpoint-restart	5
2.2	Fixed-length reservations	5
3	Framework	5
4	Difficulty of the problem	6
4.1	With a single checkpoint at the end of the reservation	6
4.2	With a single checkpoint in a short reservation	7
4.3	With two checkpoints	8
5	Threshold-based heuristic	9
6	Discretization with time quanta	12
7	Experiments	14
7.1	Simulation framework	14
7.2	Simulation results	15
8	Conclusion	17
A	All results	20

1 Introduction

This work considers checkpointing strategies for a parallel application executing on a large-scale platform whose nodes are subject to failures. The application executes for a fixed duration, namely the length T of the reservation that it has been granted. The large-scale platform may well experience several failures¹ per day [10, 8, 6, 11], hence, the application must be checkpointed during the reservation. After each failure, the application is interrupted and must be restarted. Without checkpointing, all the work executed for the application is lost. With checkpoint-restart [12, 9], the execution can resume from the last checkpoint, after some downtime (enroll a spare to replace the faulty processor) and a recovery (read the checkpoint).

The last checkpoint within the reservation plays a particular role, because all the work executed after that checkpoint and until the end of the reservation will be lost. The last checkpoint saves the state of the application and is usually taken at the very end of the reservation. For long-running High-Performance Computing (HPC) applications, it is common practice to make a series of reservations of the required resource: the execution is split into multiple smaller reservations, and checkpoint-restart is used to save intermediate steps of computation at the end of each reservation. For each actual reservation, the job needs to be checkpointed before the reservation time has elapsed, otherwise the progress of the execution during the reservation will be lost.

The optimal checkpointing strategy to mitigate the impact of failures has received considerable attention, but without considering fixed-size reservations. The problem is usually stated as follows: consider a parallel application executing on a platform whose processors are subject to

¹Failures are also called *fail-stop errors*.

failures. The application has a fixed amount of work W to execute. For simplicity, say W is expressed in seconds so that we can speak of work or time indifferently. The user has scheduled a long reservation T for this amount of work, typically over-estimating the duration of the reservation by a wide margin. In other words, we have $W \ll T$. The optimal checkpointing strategy is then to decide how frequently we should checkpoint, so that the expected total execution time $\mathbb{E}(W)$ is minimized. Note that the time left in the reservation after all the work W is completed is usually not billed to the user. We call this problem *fixed-work checkpointing*.

There is a well-known trade-off for *fixed-work checkpointing*: taking too many checkpoints leads to a high overhead, especially when there are few failures, while taking too few checkpoints leads to a large re-execution time after each failure. The optimal strategy to minimize the expected execution time is known when failure inter-arrival times, or IATs for short², obey an Exponential distribution on each processor. In that case, the optimal checkpointing strategy is periodic, the optimal period is known and can be expressed with a complicated formula that uses the Lambert function [9]. Fortunately, the optimal period can be approximated by the famous Young/Daly formula as $W_{YD} = \sqrt{2\mu C}$ [13, 7], where μ is the application Mean Time Between Failures (MTBF) and C is the checkpoint duration. The Young/Daly formula is widely used across a variety of applications and platforms (see [2] for a survey) and represents a major progress compared to naive strategies where each application would checkpoint, say, every hour, independently of the values of its MTBF μ and checkpoint duration C .

While the *fixed-work checkpointing* problem is well-understood, the problem to maximize the expected work that can be achieved during a fixed-length execution, namely the duration T of a reservation, has never been investigated. We call this problem *fixed-time checkpointing*. This new problem is the dual of the previous one: instead of minimizing the expected time to execute a fixed amount of work, we aim at maximizing the expected work achieved during a fixed amount of time.

This work provides the first study of *fixed-time checkpointing*, which we show to be much more difficult than *fixed-work checkpointing*. We start with examples that show that the optimal checkpointing strategy neither always uses periodic checkpoints nor always takes its last checkpoint exactly at the end of the reservation. Then, we introduce a dynamic heuristic strategy that is periodic and decides the checkpointing frequency based upon thresholds for the time left; we determine threshold times T_k such that it is best to plan for exactly k checkpoints if the time left (or initially the length of the reservation) is between T_k and T_{k+1} (letting $T_1 = 0$). Hence, the heuristic uses one (final) checkpoint for $T \leq T_2$, then two checkpoints (one intermediate and one final) for $T_2 \leq T \leq T_3$, and so on. Next, we use time discretization and design a (complicated) dynamic programming algorithm that computes the optimal solution, without any restriction on the checkpointing strategy. Finally, we report the results of an extensive simulation campaign that shows that the optimal solution can be far more efficient than the Young/Daly periodic approach for short or mid-size reservations.

The rest of the paper is organized as follows. Section 2 reviews related work. Section 3 details the framework. Section 4 provides examples that show the intrinsic difficulty of the problem. Section 5 is devoted to the dynamic heuristic based upon thresholds. Section 6 deals with the discretization of the problem into time quanta, and introduces the dynamic programming algorithm that computes the optimal solution without any restriction on the strategy. Section 7 reports simulation results for a comprehensive experimental campaign. Finally, Section 8 provides concluding remarks and directions for future work.

²IATs are the times elapsed between two consecutive failure events (or until the first failure at the start of the application).

2 Related work

We briefly survey related work in this section, first discussing checkpoint-restart techniques, then discussing fixed-length reservations.

2.1 Checkpoint-restart

Checkpoint-restart is one of the most widely used strategies to deal with failures. Several variants of this policy have been studied; see [9] for an overview. The literature focused on *fixed-work checkpointing*, for which a natural strategy is to checkpoint periodically; one must decide how often to checkpoint, i.e., derive the optimal checkpointing period. An optimal strategy is defined as a strategy that minimizes the expectation of the execution time of the application. For a preemptible application, where checkpointing can occur at any time, the classical formula due to Young [13] and Daly [7] states that the optimal checkpointing period is $W_{YD} = \sqrt{2\mu C}$, where μ is the application MTBF and C the checkpoint cost. This formula is a first-order approximation. For memoryless failures, Daly provides a second-order, more accurate, approximation in [7], while our previous work [5] provides the optimal value; both [7] and [5] use the Lambert function, whose Taylor expansion is key to assess the accuracy of the Young/Daly formula. The derivation in [5] is based on Equation (1) (see Section 4.1), a formula rediscovered ten years later, with a quite different proof based on a Markov model, in [11]. Finally, we point out that non-memoryless failures are more difficult to deal with for parallel applications, and we refer to the recent paper [4] for more details.

2.2 Fixed-length reservations

Long-running HPC applications usually make a series of reservations of the required resources: the execution is split into multiple smaller reservations, and checkpoint-restart is used to save intermediate steps of computation at the end of each reservation. There are multiple advantages to this approach, but the main one is that it lowers the wait time of the application, as the job scheduler can easily place a smaller reservation. On some platforms, a maximum reservation time is imposed on applications, forcing applications that run longer than this maximum time to split their reservation and rely on a form of checkpoint-restart. These scenarios occur in large scale High Performance Computing (HPC) platforms as well as on the Cloud. For each actual reservation, the job needs to be checkpointed before the reservation time has elapsed; otherwise the progress of the execution during the reservation will be lost.

To the best of our knowledge, the fixed-time checkpointing problem, namely, dealing with the impact of failures within a fixed-length reservation, has never been addressed in the literature. This is somewhat surprising, because it is a very natural problem that must be addressed for any large-scale HPC application that is granted one or several reservations by the batch scheduler. The closest related work is our paper [1], where we consider a failure-free scenario: an application executes for a fixed-length reservation and checkpoints at the end. The checkpoint duration is a stochastic random variable that obeys some well-known probability distribution law, and the question is to decide when to checkpoint so that the expected work is maximized. Adding failures to the picture was suggested by an anonymous reviewer of [1].

3 Framework

We consider a fixed-size reservation of length T . The parallel application executes and checkpoints to mitigate the impact of failures. There is always a final checkpoint, and the work after

that checkpoint is lost.

Failures Failures obey an Exponential distribution of parameter λ . The Mean-Time Between Failures (MTBF) is thus $\mu = \frac{1}{\lambda}$. If the parallel application enrolls p processors, then $\mu = \frac{\mu_{ind}}{p}$, where μ_{ind} is the individual MTBF on each processor. Given an execution of duration X , $\mathbb{P}_{succ}(X) = e^{-\lambda X}$ is the probability of success of that execution, and $\mathbb{P}_{fail}(X) = 1 - \mathbb{P}_{succ}(X) = 1 - e^{-\lambda X}$ is the probability of (at least) one failure during that execution. Failures can strike during checkpoint and recovery, but not during downtime (otherwise we would include it in the recovery). The checkpoint cost is C , the recovery cost R , and the downtime D .

Checkpointing strategy As stated above, the objective is to maximize the expected amount of work $\mathbb{E}(T)$ that can be achieved within the whole reservation of duration T . A checkpointing strategy is recursively defined as follows:

- Initially, the time left is $t_{left} = T$. The strategy decides how many checkpoints will be taken, and at which instants, if there is no failure during the whole execution. With k checkpoints taken, let $t_{end}(i) \leq T$ be the completion time of checkpoint number i , with $1 \leq i \leq k$.
- If there is no failure up to time $t_{end}(k)$ (when the last checkpoint completes), then the work achieved by the strategy will be equal to $W = t_{end}(k) - kC$.
- On the contrary, if a (first) failure strikes at time $t \leq T$, let $\ell \leq k$ be the number of the last checkpoint that completed before the failure. The work done after time $t_{end}(\ell)$ and up to time t is lost, and the work achieved by the strategy up to the failure will be equal to $W = t_{end}(\ell) - \ell C$.
- Now after the failure at time t , there is a downtime and a recovery; hence, the time left is $t_{left} = (T - t) - D - R$. if $t_{left} \geq C$, we call recursively the strategy and will add the work done then to W .

In other words, for any value $t_{left} \leq T$, a checkpointing strategy must decide how many checkpoints will be taken, and at which instants, if no failure strikes during the execution of length t_{left} . We point out that the strategy will not be the same for different values of t_{left} : the distances between consecutive checkpoints are recomputed after each failure: this is what renders the problem so difficult.

4 Difficulty of the problem

This section provides several examples and case-studies that demonstrate the difficulty of the problem.

4.1 With a single checkpoint at the end of the reservation

The first example deals with a simple checkpointing strategy: given any value of time left t_{left} , the strategy is to always work up to the end and take a unique checkpoint at the end, at time $t_{left} - C$. We let $\mathbb{E}^{end}(T, 1)$ denote the expected amount of work achieved with this strategy. We cannot provide a closed-form expression for $\mathbb{E}^{end}(T, 1)$, but we provide a recursive formula. To do so, we need an auxiliary quantity $\mathbb{E}_R^{end}(T, 1)$, where the only difference with $\mathbb{E}^{end}(T, 1)$ is that we start the execution with a recovery.

To compute $\mathbb{E}_R^{end}(T, 1)$, we distinguish two cases:

- Either there is no failure, which happens with probability $\mathbb{P}_{succ}(T)$, and we work for $T - R - C$ seconds, accounting for the initial recovery and the final checkpoint
- Or there is a failure before time T , and we have a recursion depending upon the time t at which the first failure strikes.

Overall, $\mathbb{E}_R^{end}(T, 1)$ can then be expressed as:

$$\begin{aligned} \mathbb{E}_R^{end}(T, 1) &= e^{-\lambda T}(T - R - C) \\ &+ \int_{t=0}^{T-D-R-C} \lambda t \frac{e^{-\lambda t}}{1-e^{-\lambda T}} \mathbb{E}_R^{end}(T - t - D, 1) dt \end{aligned}$$

In the integral, we use the conditional density probability $\lambda t \frac{e^{-\lambda t}}{1-e^{-\lambda T}}$ that the first failure strikes at time t , given that we know that a failure will strike before time T . We cap the range of t down to $T - D - R - C$ because no more work can be executed if there remains less than $D + R + C$ seconds after the failure.

Now the formula for $\mathbb{E}^{end}(T, 1)$ is quite similar:

$$\begin{aligned} \mathbb{E}^{end}(T, 1) &= e^{-\lambda T}(T - C) \\ &+ \int_{t=0}^{T-D-R-C} \lambda t \frac{e^{-\lambda t}}{1-e^{-\lambda T}} \mathbb{E}_R^{end}(T - t - D, 1) dt \end{aligned}$$

and simply accounts for the absence of the recovery at the beginning of the execution.

Unfortunately, the formulae for $\mathbb{E}^{end}(T, 1)$ and $\mathbb{E}_R^{end}(T, 1)$ do not lead to a closed-form expression. This is in sharp contrast with the corresponding instance of the dual fixed-work problem, i.e., the computation of the expected time $\mathbb{E}(W)$ to execute a given amount of work W followed by a checkpoint C . Indeed, the recursive formula for $\mathbb{E}(W)$ writes:

$$\begin{aligned} \mathbb{E}(W) &= e^{-\lambda(W+C)}(W + C) \\ &+ (1 - e^{-\lambda(W+C)})(\mathbb{E}(T_{lost}(W + C)) + \mathbb{E}(T_{rec}) + \mathbb{E}(W)), \end{aligned} \quad (1)$$

where $\mathbb{E}(T_{lost}(W + C))$ is the expected time lost before the first failure at time t . Similarly, $\mathbb{E}(T_{rec})$ is the expected time for the recovery (recall that failures may strike during recovery). Skipping details (see [5]), we derive the closed-form formula

$$\mathbb{E}(W) = \left(\frac{1}{\lambda} + D \right) \frac{1}{\lambda} e^{\lambda R} (e^{\lambda(W+C)} - 1).$$

The key to solving Equation (1) is that the term $\mathbb{E}(W)$ is on both sides of the recursive equation: after a failure, we resume in the same state, i.e., with the same amount W of work to execute. On the contrary, when computing $\mathbb{E}_R^{end}(T, 1)$, we are left with fewer time left after a failure at time t and call recursively for $\mathbb{E}_R^{end}(T - t - R)$. This is the key why computing $\mathbb{E}_R^{end}(T, 1)$ (or $\mathbb{E}^{end}(T, 1)$) is extremely difficult.

4.2 With a single checkpoint in a short reservation

The second example shows that the final checkpoint should not always be taken at the very end of the reservation. We consider a short reservation T , say $T = 6$, with $D = 0$ and $C = R = 4$: there is time for only one (final) checkpoint, because $2C > T$. Also, no extra work can be achieved after a failure, whenever it strikes, because $R + C > T$.

We compare the strategy STRAT₁ where we checkpoint at the end, from $t = 2$ to $t = 6$, with the strategy STRAT₂ where we checkpoint earlier, from $t = 1$ to $t = 5$. The expected gain GAIN of STRAT₁ over STRAT₂ can be either nonnegative or nonpositive, and depends upon whether a failure strikes and when; each case is weighted by its probability to occur:

- If there is no failure, the gain is $\mathbb{P}_{succ}(6) \times 1$: two seconds of work have been saved for STRAT₁, only one second for STRAT₂;
- If the first failure strikes after $t = 5$, the gain is $\mathbb{P}_{succ}(5)\mathbb{P}_{fail}(1) \times (-1)$: it is negative since no work is saved for STRAT₁, versus one second of work saved for STRAT₂;
- If the first failure strikes before $t = 5$, the gain is 0: no work has been saved for either strategy, and there is no time left after downtime and recovery after the failure to take a checkpoint.

Altogether, the gain is:

$$\begin{aligned} \text{GAIN} &= e^{-6\lambda} - e^{-5\lambda}(1 - e^{-\lambda}) = 2e^{-6\lambda} - e^{-5\lambda} \\ &= e^{\ln(2)-6\lambda} - e^{-5\lambda}, \end{aligned}$$

which is negative when $\lambda > \ln(2)$. Hence, if λ is large enough, it is better to checkpoint before the end of the reservation!

This example is quite simple, because no extra work can be achieved after a failure, so there is no recursive call, regardless of the checkpoint strategy. We deal with a more complicated example below.

4.3 With two checkpoints

This third example, and the heuristic approach described in Section 5, both involve checkpoint strategies with recursive calls after each failure, as long as there remains sufficient time left. To compare the performance of two different checkpoint strategies STRAT₁ and STRAT₂ for a reservation of size T , we compare the expected work that they achieve until the first failure, or until the end of the reservation if there is no failure. In other words, we compare the expected work until the first recursive call. This is because we can always assume that both strategies will proceed identically after the recursive call, i.e., for shorter reservations of size $t_{left} < T$. Indeed, if, say, STRAT₁ achieves more expected work than STRAT₂ until the first failure, then STRAT₁ can be used to improve STRAT₂ as follows: consider strategy STRAT₃, a modified version of STRAT₂, which starts the execution using the very same checkpoints as STRAT₁, and is identical to STRAT₂ after the first failure (if any); then STRAT₃ has better total performance than the initial STRAT₂. This explains why our comparison metric is the gain achieved by one strategy compared to the other until the first failure, if any.

Now, we move to the example and show that we should not always use same-size segments when provisioning two checkpoints. Consider the following problem: given a reservation $T \geq 2C$, we provision two checkpoints, the first scheduled to end at time $\alpha(T)T$, and the second one scheduled to end at time T . Here, $\alpha(T)$ is some function of T , and we ask whether it is optimal to have $\alpha(T) = \frac{1}{2}$, i.e., two segments of same size. We have two bounds, a lower bound $C \leq \alpha(T)T$ (so that the first segment is long enough to take a checkpoint) and an upper bound $\alpha(T)T \leq T - C$ (so that the second segment is long enough to take a checkpoint). Let STRAT₂($\alpha(T)$) be the strategy with two checkpoints, the first ending at time $\alpha(T)T$, and the second ending at time T . We aim at computing the best value of $\alpha(T)$. As stated above, we compare the gain (positive or negative) GAIN($\alpha(T)$) of STRAT₂($\alpha(T)$) over STRAT₁, the strategy where we execute only a single checkpoint at the end of the reservation. Recall that the gain GAIN($\alpha(T)$) is the difference of the expected work achieved until the first failure, if any.

The gain GAIN($\alpha(T)$) of STRAT₂($\alpha(T)$) over STRAT₁ is

- $\mathbb{P}_{succ}(T) \times (-C)$ if there is no failure, because there is one more checkpoint in STRAT₂($\alpha(T)$);

- 0 if the first failure strikes before time $\alpha(T)T$: no work saved for either strategy;
- $\mathbb{P}_{succ}(\alpha(T)T)\mathbb{P}_{fail}((1-\alpha(T))T) \times (\alpha(T)T - C)$ if the first failure strikes after time $\alpha(T)T$: no work saved for STRAT₁ versus $\alpha(T)T - C$ seconds of work for STRAT₂($\alpha(T)$).

Altogether, the gain is

$$\begin{aligned} \text{GAIN}(\alpha(T)) &= e^{-\lambda T}(-C) \\ &\quad + e^{-\lambda\alpha(T)T}(1 - e^{-\lambda(1-\alpha(T))T})(\alpha(T)T - C) \\ &= e^{-\lambda\alpha(T)T}(\alpha(T)T - C) - e^{-\lambda T}\alpha(T)T \end{aligned}$$

Given T , we consider the function $f(\alpha) = e^{-\lambda\alpha T}(\alpha T - C) - e^{-\lambda T}\alpha T$ and differentiate:

$$\begin{aligned} f'(\alpha) &= -\lambda T e^{-\lambda\alpha T}(\alpha T - C) + e^{-\lambda\alpha T}T - e^{-\lambda T}T \\ &= T e^{-\lambda\alpha T}(1 - \lambda(\alpha T - C) - e^{-\lambda(1-\alpha)T}) \end{aligned}$$

Letting $g(\alpha) = 1 - \lambda(\alpha T - C) - e^{-\lambda(1-\alpha)T}$ and differentiating again:

$$g'(\alpha) = -\lambda T(1 + e^{-\lambda(1-\alpha)T}) < 0$$

Hence, g is decreasing, and f' is decreasing too (product of two decreasing functions). We have $f'(0) = T(\lambda C + 1 - e^{-\lambda T}) > 0$ and $f'(1) = -\lambda T e^{-\lambda T}(T - C) < 0$. Hence, f' has a unique zero $\alpha_{opt}(T)$, which is the unique solution of the equation

$$1 = \alpha T - C - e^{-\lambda(1-\alpha)T}. \quad (2)$$

We see that $\alpha_{opt}(T)$ is not equal to $\frac{1}{2}$ in general, so that having two same-size segments is not optimal. The only case it is optimal is when T and λ satisfy to $1 = \frac{T}{2} - C - e^{-\lambda\frac{T}{2}}$. This is in sharp contrast with the dual fixed-work checkpointing problem, where having same-size segments is always optimal [9]. But we point out that when $\lambda \rightarrow 0$ and with $T = \Theta(\lambda^{-\frac{1}{2}})$, then the first order expansion of $\alpha_{opt}(T)$ does give $\alpha_{opt}(T) \rightarrow \frac{1}{2}$, which is comforting.

5 Threshold-based heuristic

We introduce a dynamic heuristic that relies upon thresholds. This heuristic always uses same-size segments, each finishing with a checkpoint, and the last checkpoint is scheduled to complete exactly at the end of the reservation (or what remains of it). The major difficulty is to decide how many checkpoints to provision in case of a failure-free execution. The key of the heuristic is to provide a sequence of thresholds $(T_n)_{n \geq 1}$ so that we provision exactly n checkpoints when the time left (initially the full length T of the reservation) satisfies

$$T_n \leq t_{left} \leq T_{n+1}. \quad (3)$$

Intuitively, introducing such thresholds makes sense: the longer the reservation, the more checkpoints one should provision. Also, we could expect that the longer the reservation, the more work should the heuristic achieve in expectation. Unfortunately, this is not true: revisiting the example of Section 4.2, the heuristic achieves more expected work when $T = 5$ than when $T = 6$ whenever $\lambda > \ln(2)$. And the example of Section 4.3 shows that relying on same-length segments is seldom optimal. Still, the heuristic is very natural, and we explain below how to determine the threshold sequence $(T_n)_{n \geq 1}$.

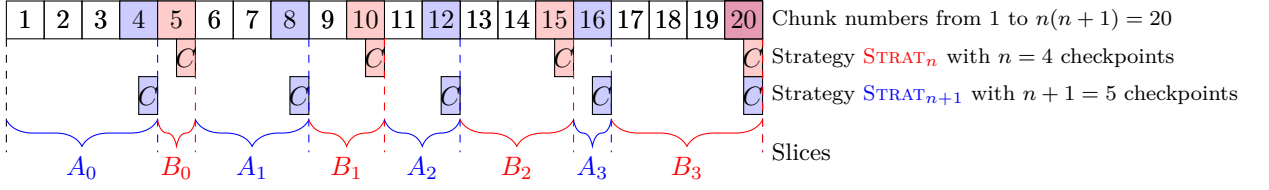


Figure 1: Comparing strategies with n and $n + 1$ checkpoints for $n = 4$.

We construct the sequence inductively and start with $T_1 = 0$: for short reservations, we always need a checkpoint at the end. Now, given the value of the first n thresholds T_i , for $1 \leq i \leq n$, consider a reservation of length

$$T \geq \max(T_n, (n + 1)C). \quad (4)$$

We determine at which length T_{n+1} it does become better to use $n + 1$ checkpoints instead of n . To do so, we compare the gain (either positive, zero or negative) when using $(n + 1)$ checkpoints (strategy STRAT_{n+1}) instead of using n checkpoints (strategy STRAT_n). Recall that the gain compares the expected work until the first failure. Hence, for a reservation T obeying Equation (4), we want to compare $\mathbb{E}(T, n + 1)$, the expected work with STRAT_{n+1} , with $\mathbb{E}(T, n)$, the expected work with STRAT_n , until the first failure. We compute the difference $\text{GAIN}(T, n + 1) = \mathbb{E}(T, n + 1) - \mathbb{E}(T, n)$ via a case analysis, depending upon when the first failure (if any) strikes in the execution.

In Figure 1, we give an example for a reservation of length T , using either $n = 4$ or $n + 1 = 5$ checkpoints. In the general case, because n and $n + 1$ are relatively prime, there are $n(n + 1)$ chunks where the first failure may strike. For the example, there are 20 chunks. Altogether, there are several cases to analyze, depending upon which chunk is struck by the first failure, plus the last case without any failure. This last case is the easiest to analyze: with one more checkpoint for STRAT_{n+1} than for STRAT_n , the gain is negative, and its expected value is $\mathbb{P}_{\text{succ}}(T) \times (-C)$.

Now, we discuss the gain for the case when the first failure strikes chunk number s , for $1 \leq s \leq n(n + 1)$. Let $U = \frac{T}{n(n+1)}$ be the length of a chunk. We start with a lemma:

Lemma 1. *The probability that the first failure strikes chunk number $s \geq 2$ is $\mathbb{P}_{\text{succ}}((s - 1)U) \times \mathbb{P}_{\text{fail}}(U)$. The probability that the first failure strikes a slice of t consecutive chunks composed of chunks number s to $s + t - 1$ with $s \geq 1$ and $t \geq 2$ is $\mathbb{P}_{\text{succ}}((s - 1)U) \times \mathbb{P}_{\text{fail}}(tU)$.*

Proof. This is a well-known result, due to the memoryless property of the Exponential probability distribution. It is based upon the fact that the probability that the first failure strikes a slice of k consecutive chunks is the sum of the probabilities that the first failure strikes any of these chunks. \square

Chunks can be partitioned into $2n$ slices A_m and B_m for $0 \leq m \leq n - 1$:

- Slices appear in the order

$$A_0, B_0, A_1, B_1, \dots, A_{n-1}, B_{n-1}.$$

- Slice A_m is composed of chunks number s with $m(n + 1) < s \leq (m + 1)n$: chunks after checkpoint number m (or the beginning of the execution if $m = 0$) of strategy STRAT_{n+1} , and up to checkpoint number $m + 1$ of strategy STRAT_n (see Figure 1). In the figure, slice A_0 includes the first $n = 4$ chunks, and slice A_1 includes $n - 1$ chunks after the first

checkpoint in strategy STRAT_{n+1} , i.e., chunks 6 to 8. Slice A_m is composed of $n - m$ chunks.

- Slice B_m is composed of chunks number s with $(m + 1)n < s \leq (m + 1)(n + 1)$: chunks after checkpoint number $m + 1$ of strategy STRAT_n and up to checkpoint number $m + 1$ of strategy STRAT_{n+1} . In Figure 1, slice B_0 includes chunk 5, and slice B_1 includes chunks 9 and 10. Slice B_m is composed of $m + 1$ chunks.

If the first failure strikes within slice A_0 , both strategies lose everything and there is no difference. If the first failure strikes within slice A_m for $m \geq 1$, then strategy STRAT_n has saved more chunks than strategy STRAT_{n+1} , namely all the chunks of slice B_{m-1} . Hence, the gain is negative, and its expected value is

$$-\mathbb{P}_{succ}(m(n + 1)U)\mathbb{P}_{fail}((n - m)U) \times (mU).$$

The probability is computed using the fact that slice A_m starts at chunk $m(n + 1) + 1$, and has length $n - m$ chunks. The extra number of chunks for STRAT_n is given by the number m of chunks of slice B_{m-1} . Indeed, since both strategies made the same number of checkpoints before the failure, the gain is exactly the difference between the completion time of the last checkpoint for both strategies, which corresponds to the number of chunks of slice B_{m-1} .

If the first failure strikes within slice B_m , then strategy STRAT_n has saved fewer chunks than strategy STRAT_{n+1} , namely all the chunks of slice A_m , but had to do one less checkpoint. Hence, the expected value of the gain is $\mathbb{P}_{succ}((m + 1)nU)\mathbb{P}_{fail}((m + 1)U) \times ((n - m)U - C)$. The probability is computed using the fact that slice B_m starts at chunk $m(n + 1) + 1$, and has length $m + 1$ chunks. The extra number of chunks for STRAT_{n+1} is given by the number $n - m$ of chunks of slice A_m , but this time we need to deduce the additional checkpoint taken by STRAT_{n+1} at the end of A_m . Altogether, we have derived that

$$\begin{aligned} \text{GAIN}(T, n + 1) &= \mathbb{E}(T, n + 1) - \mathbb{E}(T, n) \\ &= -\mathbb{P}_{succ}(T) \times C \\ &\quad (\text{loss if no failure}) \\ &\quad - \sum_{m=1}^{n-1} \mathbb{P}_{succ}(m(n + 1)U)\mathbb{P}_{fail}((n - m)U) \times (mU) \\ &\quad (\text{gain if first failure strikes slice } A_m) \\ &\quad + \sum_{m=0}^{n-1} \mathbb{P}_{succ}((m + 1)nU)\mathbb{P}_{fail}((m + 1)U) \times ((n - m)U - C) \\ &\quad (\text{loss if first failure strikes slice } B_m). \end{aligned}$$

We are unable to prove that there exists a single solution $T_{n+1} \geq (n + 1)C$ to the equation $\text{GAIN}(T_{n+1}, n + 1) = 0$, but we conjecture that this is always the case (as one can intuitively expect, there is no reason to use more checkpoints in a shorter reservation than in a longer one). In the experiments reported in Section 7, we have solved the equation numerically and we constructed a unique sequence of thresholds $(T_n)_{n \geq 1}$.

We proceed to a first-order approximation of T_{n+1} when $\lambda \rightarrow 0$. As for the Young/Daly approximation [9], the failure-free overhead and the failure-induced overhead must be of the same order of magnitude for this approximation. The failure-free overhead is the relative cost of checkpoints $\Theta(\frac{C}{T})$. The failure-induced overhead is the re-execution cost and is of the order $\Theta(\lambda T)$: indeed, we re-execute a fraction of the reservation every time a failure occurs, which happens a fraction $\frac{1}{\lambda}$ (the MTBF) of the time. This gives that $T = \Theta(\frac{1}{\sqrt{\lambda}})$, hence, $T \rightarrow \infty$ when $\lambda \rightarrow 0$, but also $\lambda T = \Theta(\sqrt{\lambda}) \rightarrow 0$ when $\lambda \rightarrow 0$, and we can safely use the first order approximation $\mathbb{P}_{succ}(T) = e^{-\lambda T} = 1 - \lambda T + o(\frac{1}{\lambda})$. We rewrite the gain $\text{GAIN}(T, n + 1)$ as

$$\begin{aligned} \text{GAIN}(T, n+1) &= -C(1 - n(n+1)\lambda U) \\ &\quad - \sum_{m=1}^{n-1} (1 - m(n+1)\lambda U)(n-m)\lambda U m U \\ &\quad + \sum_{m=0}^{n-1} (1 - (m+1)n\lambda U)(m+1)\lambda U((n-m)U - C) + O(\lambda) \end{aligned}$$

We have $\lambda^2 U^2 = O(\lambda)$ while $\lambda^2 U = O(\sqrt{\lambda})$ because T (and U) are $\Theta(\frac{1}{\sqrt{\lambda}})$. Keeping only first-order terms, we derive that

$$\begin{aligned} \text{GAIN}(T, n+1) &= -C + \frac{n(n+1)}{2}\lambda U^2 + O(\sqrt{\lambda}) \\ &= -C + \frac{\lambda T^2}{2n(n+1)} + O(\sqrt{\lambda}), \end{aligned}$$

which leads to the first-order approximation

$$T_{n+1} \approx \sqrt{\frac{2n(n+1)C}{\lambda}}. \quad (5)$$

It is interesting to compare with the Young/Daly formula:

- When $n = 1$; introducing the MTBF $\mu = \frac{1}{\lambda}$, we have $T_2 \approx \sqrt{4C\mu}$ while the optimal Young/Daly period is $W_{YD} = \sqrt{2C\mu}$: in limited time, better keep a unique checkpoint up to T_2 , whose value is $\sqrt{2}$ larger than W_{YD} .
- More generally, the length of n Young/Daly segments is $f(n) = n\sqrt{2\mu C}$; hence, a first approximation for T_{n+1} , the threshold limit between n or $n+1$ segments, is the geometric mean of $f(n)$ and $f(n+1)$, which gives

$$\sqrt{f(n)f(n+1)} = \sqrt{n(n+1)2\mu C} \approx T_{n+1}.$$

6 Discretization with time quanta

In this section, we show that time discretization enables us to derive the optimal solution, without any restriction on the solution: we can indeed have segments with different lengths and we can checkpoint a few moments before the end of the reservation if it is better (in expectation) than working up to the very last moment. The main idea is to discretize time into small quanta, and to use dynamic programming to compute the best checkpointing strategy. Obviously, the smaller the quanta, the more accurate the results, but the more costly the dynamic programming algorithm.

Technically, we introduce a time quantum u , and we discretize time into quanta. This means that all quantities (reservation length, segment sizes, checkpoint and recovery times, downtime) are integer multiples of u , and that failures strike at the end of a quantum. More precisely, if a variable y is defined in work units, $y^* = y/u$ is the corresponding number of quanta, which we always suppose integer. The reservation length becomes the total number of quanta $T^* = T/u$, the checkpoint time is $C^* = C/u$ quanta, the recovery time is $R^* = R/u$ quanta and the downtime $D^* = D/u$ quanta. The probability that a failure happens during (at the end of) quantum i , $1 \leq i \leq T^*$, is $p_i^* = \int_{(i-1)u}^{iu} \lambda e^{-\lambda x} dx = e^{-\lambda(i-1)u} - e^{-\lambda iu}$. We denote the probability that a failure happens during the next j quanta of execution as $\mathbb{P}_{fail}^*(j) = 1 - e^{-\lambda j u}$. We define $\mathbb{P}_{succ}^*(j) = e^{-\lambda j u}$ similarly. And we retrieve that the probability that the first failure strikes during quantum i is $\mathbb{P}_{succ}^*(i-1)\mathbb{P}_{fail}^*(1) = p_i^*$. This discretization restricts the search for an optimal execution to a finite set of possible executions. The trade-off is that a smaller value of u

leads to a more accurate solution, but also to a higher number of states in the algorithm, hence, to a higher computing time.

To express the dynamic programming algorithm, we need to distinguish whether the current segment starts with a recovery or not. Let $\mathbb{E}(n, k, \delta)$ be the optimal work that can be achieved (in expectation) during n quanta and when taking exactly k checkpoints initially, with $\delta = 0$ if there is no initial recovery with the work, and $\delta = 1$ otherwise. Let $k_{max} = \lfloor \frac{T^*}{C^*} \rfloor$ be the maximum number of checkpoints that can be taken within T^* quanta, i.e., if we checkpoint all the time. The objective is to use a dynamic programming algorithm to determine

$$\mathbb{E}_{opt}(T^*) = \max_{1 \leq k \leq k_{max}} \mathbb{E}(T^*, k, 0). \quad (6)$$

In Equation (6), we have $k \geq 1$ because we need at least one checkpoint. We start with $k = 1$, i.e., having exactly one checkpoint. We derive the recursion formula, where i is the index of the quantum when the (unique) checkpoint completes:

$$\begin{aligned} \mathbb{E}(n, 1, \delta) &= \max_{\delta R^* + C^* + 1 \leq i \leq n} \mathbb{P}_{succ}^*(i)(i - C^* - \delta R^*) \\ &\quad + \sum_{f=1}^{i - D^* - R^* - C^*} p_f^* \mathbb{E}(n - f - D^*, 1, 1). \end{aligned} \quad (7)$$

The formula reads as follows: we look for every quantum i where to complete the checkpoint, meaning that if no failure strikes, the actual execution would have last $i - C^*$ quanta, minus again R^* quanta if there was an initial recovery; this gives the lower bound for i , because we aim at doing at least one quantum of work. Now, if the execution does succeed, which happens with probability $\mathbb{P}_{succ}^*(i)$, the remaining quanta will be lost, since we cannot take a second checkpoint by hypothesis. However, if the first failure strikes during quantum f (which happens with probability p_f^*), we have to restart from scratch. There will remain $n - f - D^*$ quanta for execution, starting with a recovery, and still aiming at taking one checkpoint. The upper bound for f comes from the fact that we will need a downtime, a recovery and a checkpoint before saving any work.

To understand how the algorithm is initialized, we point out that we have

$$\mathbb{E}(n, 1, 0) = \max_{C^* + 1 \leq i \leq n} \mathbb{P}_{succ}^*(i)(i - C^*)$$

for $n \leq 1 + C^* + D^* + R^*$. Indeed, without an initial recovery, we try all possible quanta i where to end the checkpoint. If there is a failure, we will need a downtime, a recovery and a checkpoint before saving any more work; in the best case the failure strikes during the first quantum, hence, the upper bound for n . The case with an initial recovery is similar, we have $\mathbb{E}(n, 1, 1) = \max_{R^* + C^* + 1 \leq i \leq n} \mathbb{P}_{succ}^*(i)(i - C^* - R^*)$ for $n \leq 1 + C^* + D^* + R^*$.

In fact, we can use the simpler initialization

$$\begin{aligned} \mathbb{E}(n, 1, 0) &= 0 && \text{for } n \leq C^* \text{ (where possibly } n \leq 0) \\ \mathbb{E}(n, 1, 1) &= 0 && \text{for } n \leq R^* + C^* \text{ (where possibly } n \leq 0) \end{aligned}$$

and recursively compute the values $\mathbb{E}(n, 1, 0)$ and $\mathbb{E}(n, 1, 1)$ for larger values of n using Equation (7).

We now detail the recursion for arbitrary values of the number $k \geq 2$ of checkpoints:

$$\begin{aligned} \mathbb{E}(n, k, \delta) &= \\ &\max_{\delta R^* + C^* + 1 \leq i \leq n - (k-1)C^*} \left(\mathbb{P}_{succ}^*(i)(i - C^* - \delta R^* + \mathbb{E}(n - i, k - 1, 0)) \right. \\ &\quad \left. + \sum_{f=1}^i p_f^* \max_{1 \leq m \leq k} \mathbb{E}(n - f - D^*, m, 1) \right). \end{aligned} \quad (8)$$

The main differences with the case $k = 1$ are the following:

Algorithm 1: Dynamic programming algorithm

```

1 for  $k = 1$  to  $k_{max}$  do
2   for  $n = 1$  to  $T^*$  do
3     Compute  $\mathbb{E}(n, k, 0)$  and  $\mathbb{E}(n, k, 1)$ 
4     using Equation (7) if  $k = 1$ ,
5     and Equation (8) if  $k \geq 2$ 

```

- After completing a checkpoint at quantum i , if no failure has struck by then, there remains the possibility to save more work, because we are still planning for $k - 1$ more checkpoints, hence, the recursive call to $\mathbb{E}(n - 1, k - 1, 0)$.
- On the contrary, if a failure has struck at quantum f , there is no reason to continue with k checkpoints, we can try using fewer checkpoints instead, hence the recursive call to the maximum value $\mathbb{E}(n - f - D^*, m, 1)$ achieved with m checkpoints, where $1 \leq m \leq k$.

The initialization now writes exactly as for $k = 1$:

$$\begin{aligned} \mathbb{E}(n, k, 0) &= 0 && \text{for } n \leq kC^* \text{ (where possibly } n \leq 0) \\ \mathbb{E}(n, k, 1) &= 0 && \text{for } n \leq R^* + kC^* \text{ (where possibly } n \leq 0) \end{aligned}$$

Algorithm 1 informally presents the dynamic programming algorithm. Its complexity is $O((T^*)^2(k_{max})^2)$ because there are $k_{max}T^*$ iterations, each calling up to $k_{max}T^*$ already computed values.

7 Experiments

In this section, we assess through simulations the performance of four checkpointing strategies:

- DYNAMICPROGRAMMING is the quantum-based dynamic programming algorithm (Algorithm 1).
- FIRSTORDER uses the first-order approximation of the thresholds T_n defined by Equation (5).
- NUMERICALOPTIMUM uses a numerical approximation of the thresholds T_n , with a numerical search for the zero of $\text{GAIN}(T, n)$.
- YOUNGDALY is the baseline reference, which takes checkpoints periodically following the Young/Daly formula W_{YD} . If the remaining reservation length after the last checkpoint or last recovery after a failure is shorter than W_{YD} , then a checkpoint is taken right at the end of the reservation.

7.1 Simulation framework

To assess the performance of the different checkpointing policies, we considered the following set of parameters:

- $C \in \{10, 20, 40, 80, 160\}$, and $R = C$;
- $D \in \{0, 5\}$;

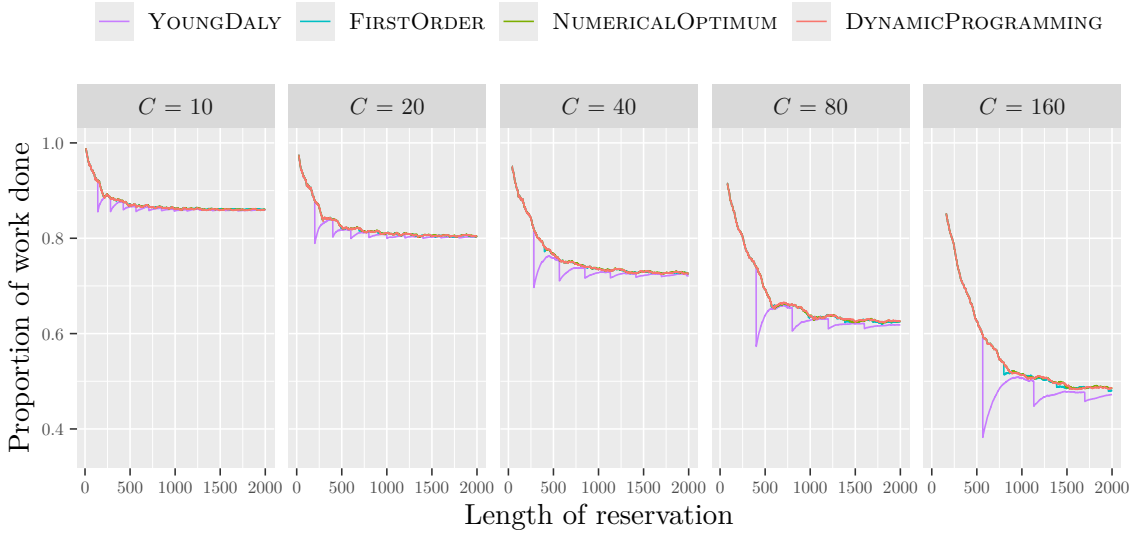


Figure 2: Proportion of work completed by the four checkpointing strategies when $\lambda = 0.001$ and $D = 0$.

- $\lambda \in \{0.01, 0.001, 0.0001\}$;
- $T \in [C, 2000]$.

All these values are expressed in an arbitrary time-unit (seconds, minutes or hours), thereby allowing for variation in the granularity of the scenarios.

For each combination of parameter values, we randomly generate 1000 failure traces. On any given instance, the work achieved by any heuristic is upper bounded by $T - C$. In every figure, we report the *proportion of work* achieved by each heuristic, i.e., the amount of work achieved by this heuristic divided by the $T - C$ upper bound. Hence, the proportion of work takes values between 0 and 1, and the higher the better.

We first report performance for DYNAMICPROGRAMMING when the quantum is set at 1. Then we study the influence of the quantum value.

7.2 Simulation results

We focus here on a subset of the results, while the whole results can be found in Appendix A. In Figure 2, we compare the performance of the different strategies when $\lambda = 0.001$. In Figure 3, we illustrate an extreme case where $C \in \{80, 160\}$ and $\lambda = 0.0001$. For both figures, we use $D = 0$, since the value of D has no significant impact on the performance of all strategies (see Appendix A for details).

Below is a summary of the main observations from these figures, which hold true for the entire simulation campaign (see [3]):

- NUMERICALOPTIMUM delivers better, or equivalent, performance than FIRSTORDER. This is natural because NUMERICALOPTIMUM uses a finer approximation for the thresholds T_n . In most cases, the difference is negligible. It can nevertheless be easily spotted on Figure 3 when $C = 160$.

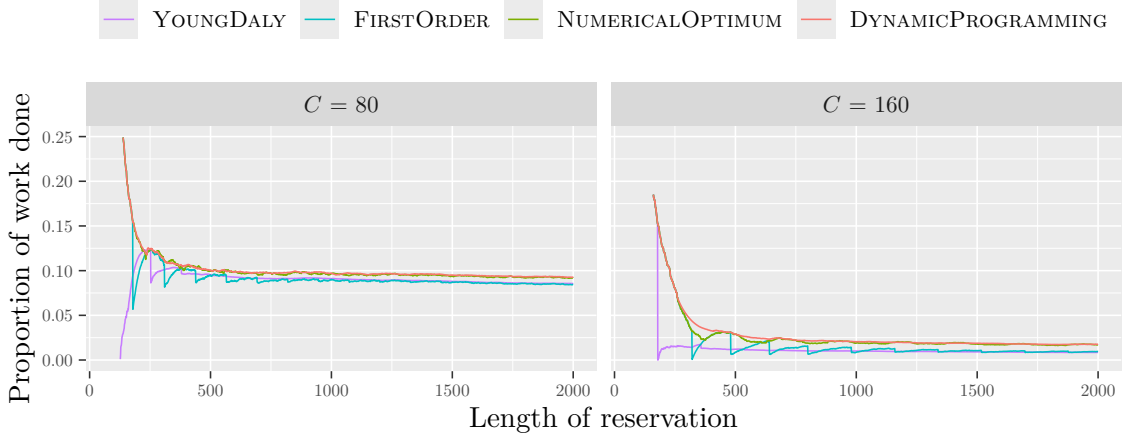


Figure 3: Proportion of work completed by the four checkpointing strategies when $\lambda = 0.01$, $D = 0$, and $C \in \{80, 160\}$.

- DYNAMICPROGRAMMING delivers better, or equivalent, performance than NUMERICALOPTIMUM. If we temporarily forget about the issue of the quantum size (which will be discussed later), DYNAMICPROGRAMMING is supposed to achieve better results than NUMERICALOPTIMUM because it searches for the best solution in a larger set of potential solutions. Indeed, DYNAMICPROGRAMMING is not mandated to complete its last checkpoint exactly at the end of the reservation. Once again, in most cases, the difference of performance with NUMERICALOPTIMUM is negligible. It can nevertheless be seen on Figure 3 when $C = 160$. Note that, even on this figure, when $C = 80$ and when $C = 160$, there are reservation lengths for which NUMERICALOPTIMUM achieves a better performance than DYNAMICPROGRAMMING. This can be explained by the limitation of the quantization and/or by the variability of random instances.
- The higher the failure rate and the shorter the reservation, the more significant the differences between the strategies. On the contrary, when the reservation length grows towards infinity or when the failure rate is small, then the problem more closely resembles that with an infinite reservation, and the approximations made by YOUNGDALY and FIRSTORDER are fully justified. When the length of the reservation grows, the performance of all strategies converge toward the asymptotic performance of YOUNGDALY.
- YOUNGDALY achieves significantly lower performance when only a handful of checkpoints are required. The most significant loss of performance happens when the reservation length becomes slightly larger than W_{YD} : then YOUNGDALY uses a second checkpoint, while the other strategies still use a single checkpoint. The second most significant loss of performance happens when YOUNGDALY starts performing a second checkpoint, e.g., when $1.2W_{YD} \leq T \leq 1.6W_{YD}$. Its first checkpoint completes at time W_{YD} and the second one at the end of the reservation length, so they are not equally spaced. In such circumstances, the three other strategies achieve significantly better performance by using two segments of equal size.

Figure 4 presents the impact of the choice of the quantum size on DYNAMICPROGRAMMING when $C = 20$. The choice of the quantum size only has a significant impact on the performance

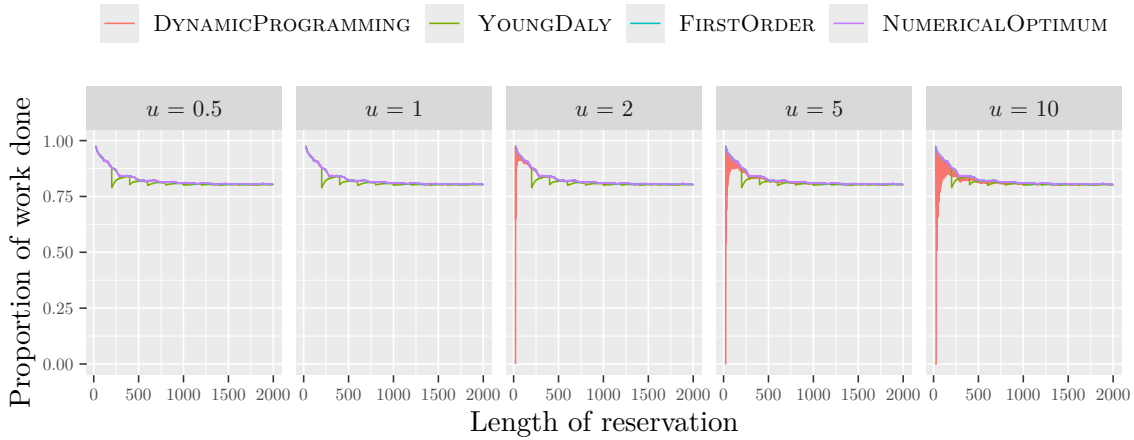


Figure 4: Impact of quantum size on the performance of DYNAMICPROGRAMMING when $\lambda = 0.001$, $D = 0$, and $C = 20$.

of DYNAMICPROGRAMMING when the reservation length is quite small, being of the order of a handful of quanta. This can be seen on Figure 5, which presents a subset of the data of Figure 4 corresponding to $T \leq 100$. When the quantum does not exceed 2, the differences become quickly negligible. We do not observe any difference when the quantum is equal to 1, and choosing a smaller quantum (say, 0.5) does not lead to better performance. This justifies our choice of quantum for Figures 2 and 3.

Overall, when the reservation length is of the order of a few Young/Daly periods (say, less than half a dozen such periods), using the Young/Daly period can lead to significant loss of performance. One should rather use evenly distributed checkpoints whose number is derived using a numerical approximation of the thresholds T_n (the NUMERICALOPTIMUM policy). The DYNAMICPROGRAMMING policy is more general and more expensive to run but does not lead to significantly better performance. However, it fully validates using the NUMERICALOPTIMUM policy as a lightweight and efficient checkpointing strategy.

8 Conclusion

This work has focused on the *fixed-time checkpointing* problem, where the objective is to maximize the expected work achieved during a fixed-length reservation. The *fixed-time checkpointing* problem is the dual of the classical *fixed-work checkpointing* problem, where the objective is to minimize the expected time to execute a fixed amount of work. To the best of our knowledge, this is the first work to study the *fixed-time checkpointing* problem, despite its importance for scientific applications that execute on large-scale, hence, failure-prone platforms; typically, these applications split their execution into a series of fixed-length reservations, a policy that is safer and more friendly to the batch scheduler than using a single (necessarily over-provisioned) reservation.

Our first contribution is to show that the *fixed-time checkpointing* problem is much more difficult than the classical and well-studied dual problem. The optimal solution seems out of reach, and we have provided several examples to explain why. Our second contribution is to provide two variants of a dynamic threshold-based strategy that outperform the standard Young/Daly

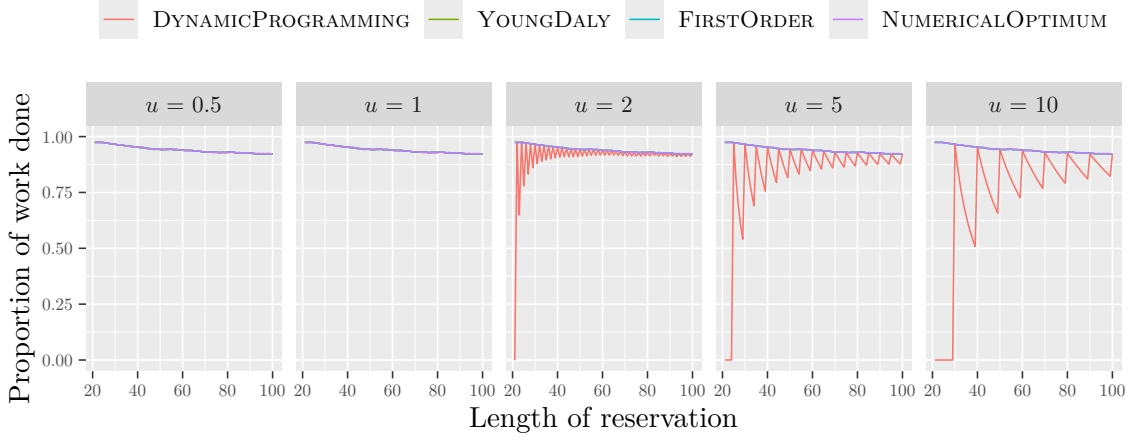


Figure 5: Impact of quantum size on the performance of `DYNAMICPROGRAMMING` when $\lambda = 0.001$, $D = 0$, and $C = 20$: detail from Figure 4.

approach for short-size or medium-size reservations, as demonstrated by a comprehensive simulation campaign. Finally, our third contribution is to use time discretization, and to design a dynamic programming algorithm that computes the optimal solution, without any restriction on the checkpointing strategy. The dynamic programming algorithm is optimal for small quanta, and achieves only slightly better performance than the dynamic threshold strategy, which assesses the quality of the latter strategy.

Altogether, these contributions enable us to conclude with two good news. The first one is that the Young/Daly approach performs well for long reservations lasting at least, say, a dozen Young/Daly periods $W_{YD} = \sqrt{2\mu C}$. The second one is that the dynamic threshold strategy is very easy to use and performs remarkably well for shorter reservations, as demonstrated by its performance on par with the optimal (discrete) solution.

Future work will be devoted to extending the results in a stochastic framework, where both checkpoint durations and application progress rate (think of iterations in sparse linear algebra problems) are no longer fully deterministic.

Acknowledgment

We would like to thank the anonymous reviewer of our FTXS'2023 paper [1], whose question motivated this work on the fixed-time checkpointing problem.

References

- [1] Q. Barbut, A. Benoit, T. Herault, Y. Robert, and F. Vivien. When to checkpoint at the end of a fixed-length reservation? In *FTXS'2023, the 13th Workshop on Fault-Tolerance for HPC at Extreme Scale, in conjunction with SC'2023*. ACM Press, 2023.
- [2] L. Bautista-Gomez, A. Benoit, S. Di, T. Herault, Y. Robert, and H. Sun. A survey on checkpointing strategies: Should we always checkpoint à la Young/Daly? *Future Generation Computer Systems*, 161:315–328, 2024.
- [3] A. Benoit, L. Perotin, Y. Robert, and F. Vivien. Checkpointing strategies for a fixed-length execution. Research report RR-9552, Inria, July 2024. Available at hal.inria.fr.
- [4] A. Benoit, L. Perotin, Y. Robert, and F. Vivien. Checkpointing strategies to tolerate non-memoryless failures on HPC platforms. *ACM Trans. Parallel Computing*, 11(1), 2024.
- [5] M. Bougeret, H. Casanova, M. Rabie, Y. Robert, and F. Vivien. Checkpointing strategies for parallel jobs. In *Proceedings of SC'11*, 2011.
- [6] F. Cappello, A. Geist, W. Gropp, S. Kale, B. Kramer, and M. Snir. Toward exascale resilience: 2014 update. *Supercomputing frontiers and innovations*, 1(1), 2014.
- [7] J. T. Daly. A higher order estimate of the optimum checkpoint interval for restart dumps. *Future Generation Comp. Syst.*, 22(3):303–312, 2006.
- [8] K. Ferreira, J. Stearley, J. H. I. Laros, R. Oldfield, K. Pedretti, R. Brightwell, R. Riesen, P. G. Bridges, and D. Arnold. Evaluating the Viability of Process Replication Reliability for Exascale Systems. In *SC'11*. ACM, 2011.
- [9] T. Herault and Y. Robert, editors. *Fault-Tolerance Techniques for High-Performance Computing*, Computer Communications and Networks. Springer Verlag, 2015.
- [10] B. Schroeder and G. A. Gibson. Understanding Failures in Petascale Computers. *Journal of Physics: Conference Series*, 78(1), 2007.
- [11] P. Sigdel, X. Yuan, and N. Tzeng. Realizing best checkpointing control in computing systems. *IEEE TPDS*, 32(2):315–329, 2021.
- [12] Wikipedia. Application checkpointing, 2023. https://en.wikipedia.org/wiki/Application_checkpointing.
- [13] J. W. Young. A first order approximation to the optimum checkpoint interval. *Comm. of the ACM*, 17(9):530–531, 1974.

A All results

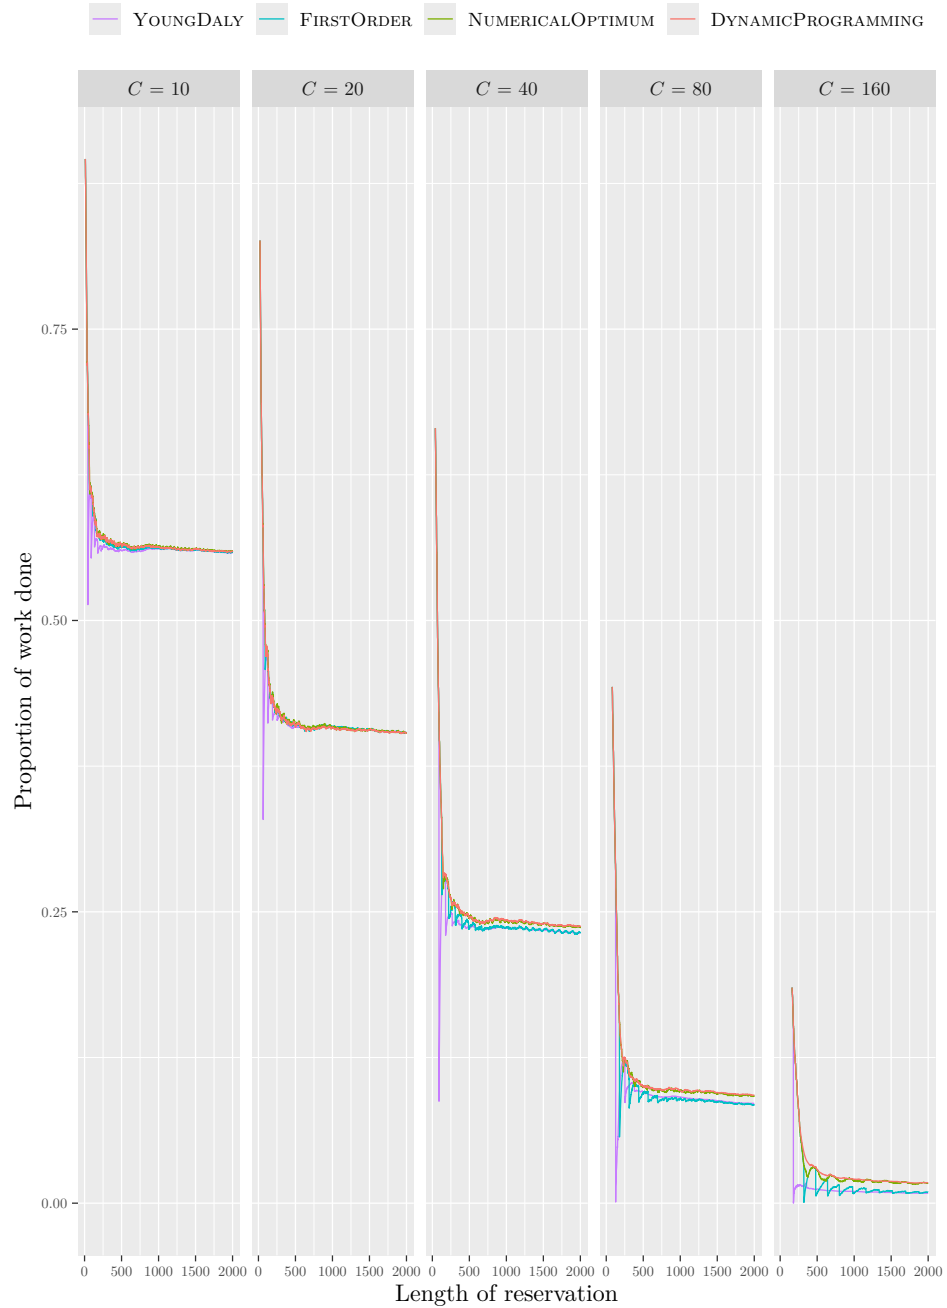
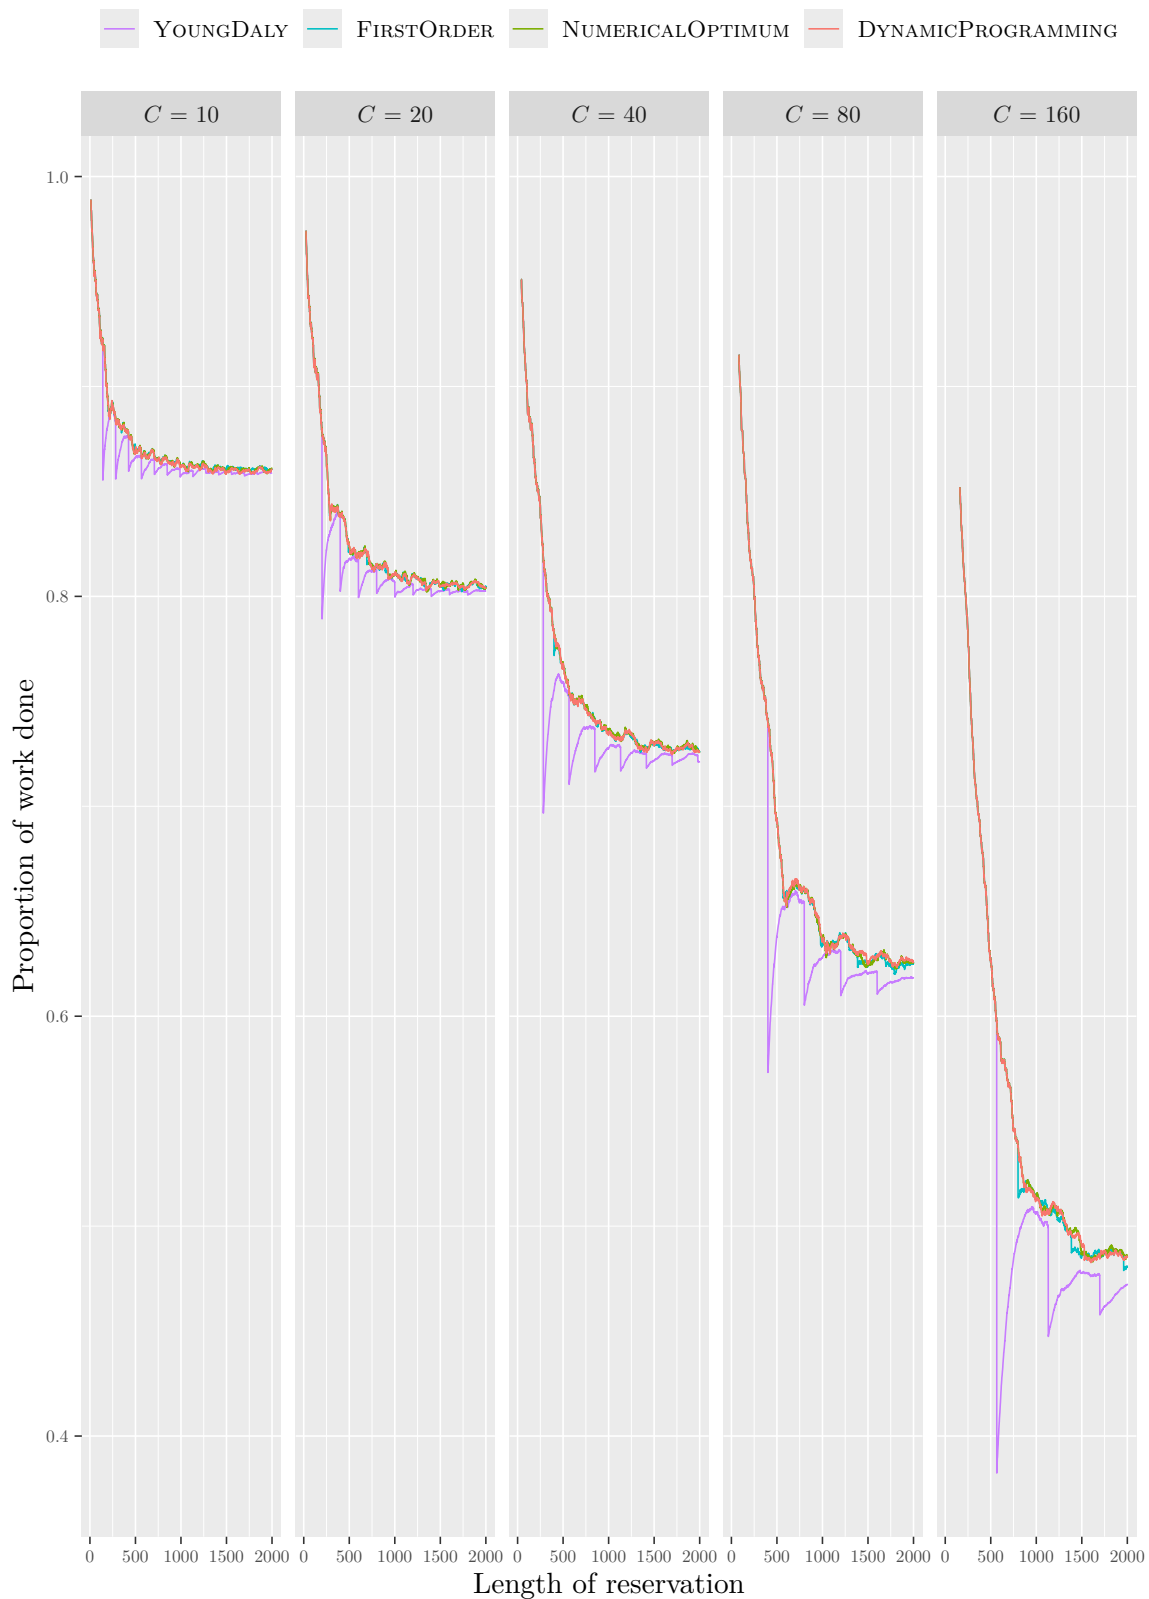
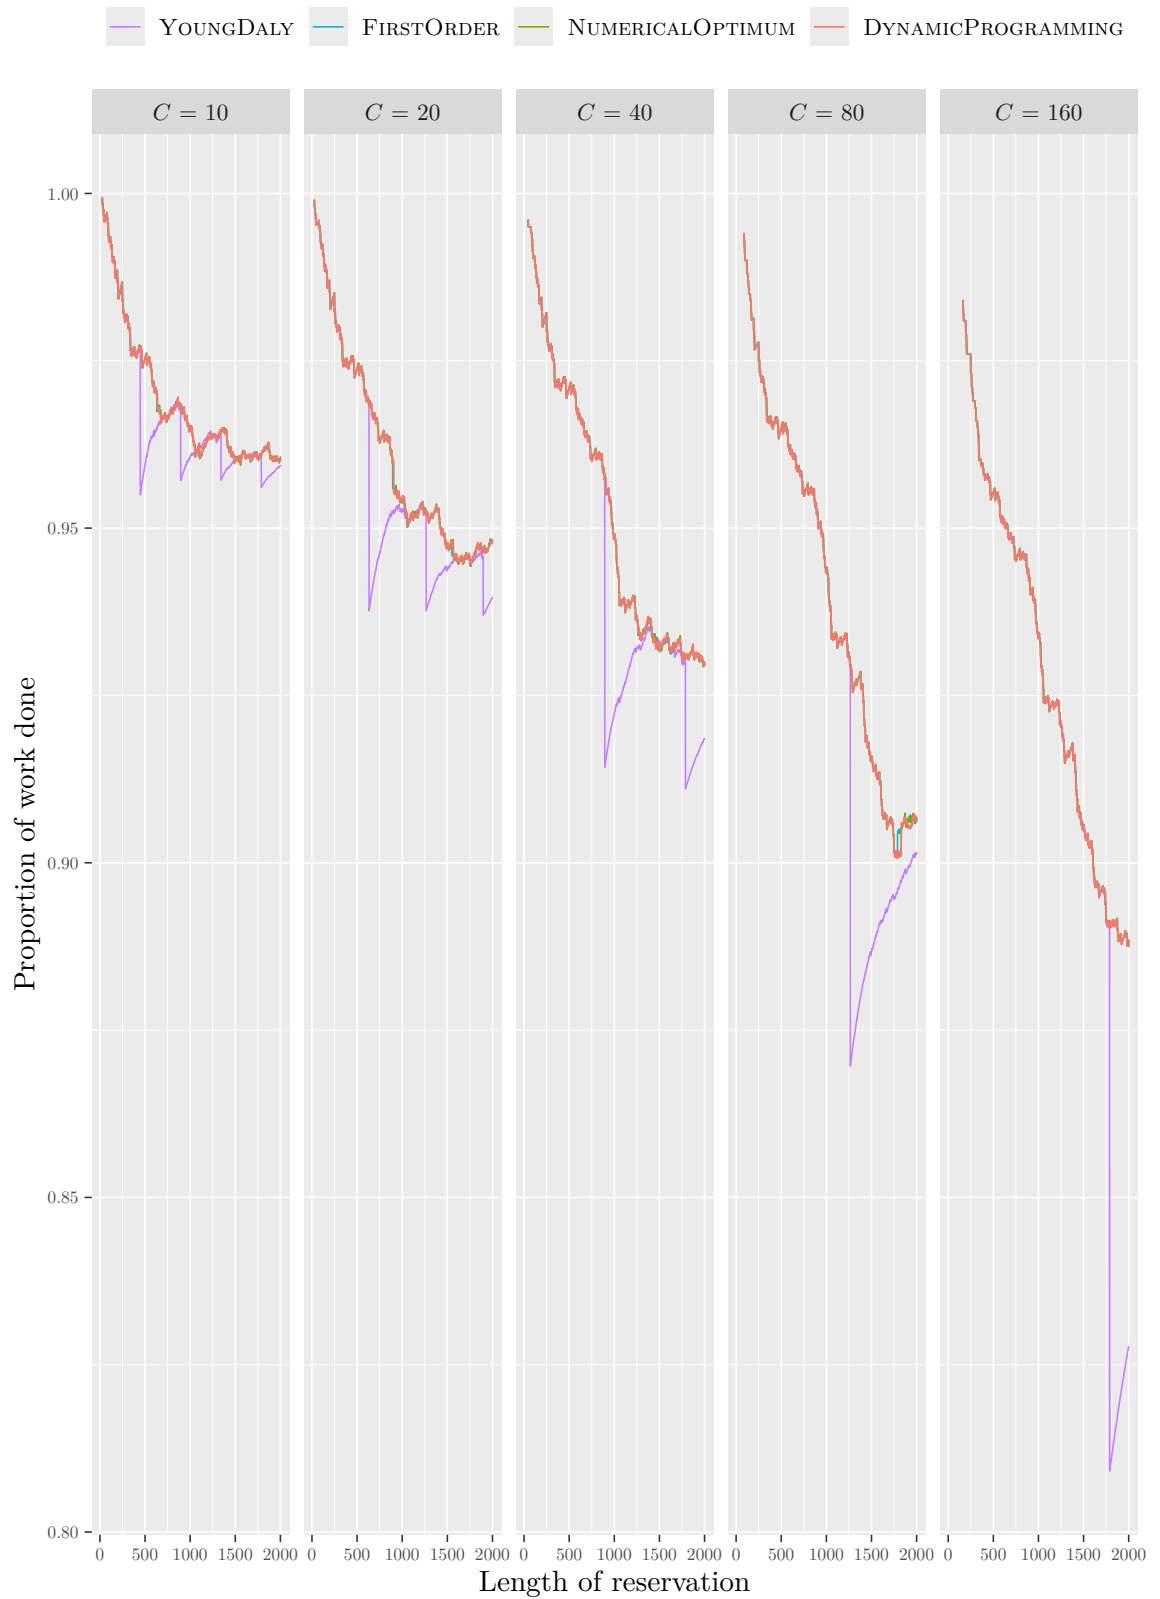


Figure 6: Proportion of work completed by the different checkpointing strategies when $\lambda = 0.01$ and $D = 0$.



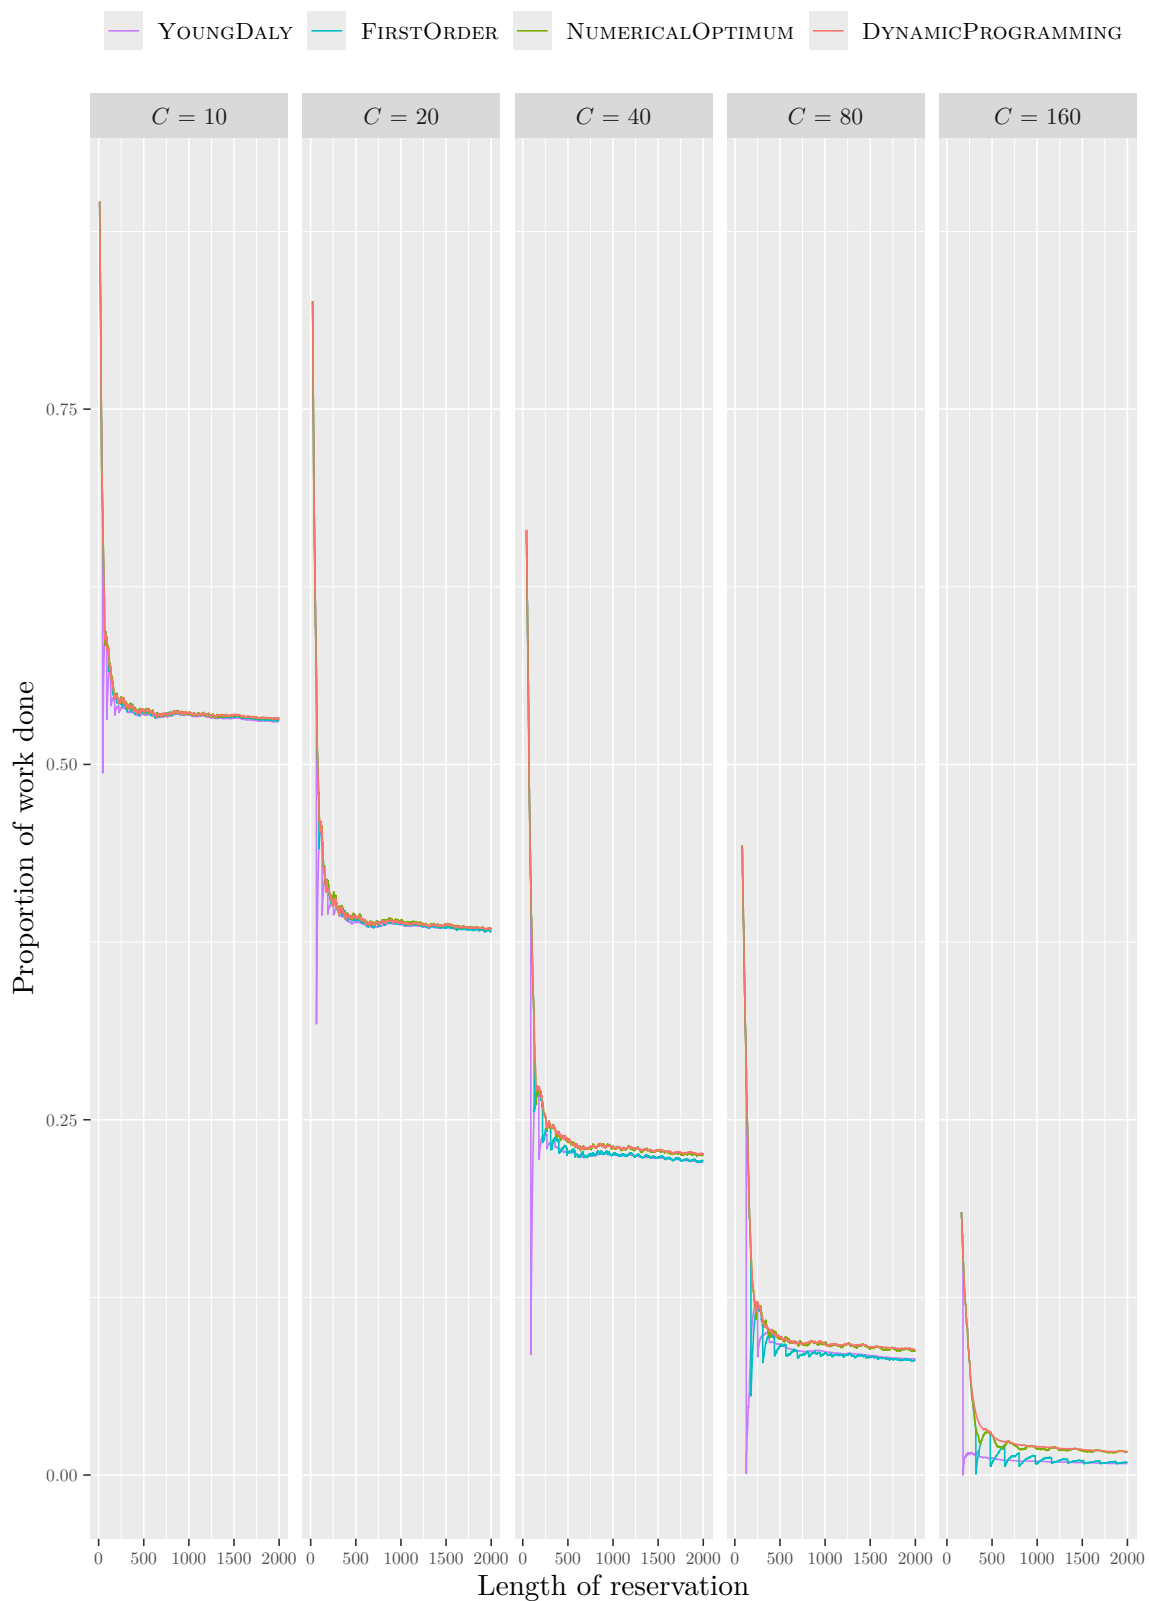
RR n° 9552

Figure 7: Proportion of work completed by the different checkpointing strategies when $\lambda = 0.001$ and $D = 0$.



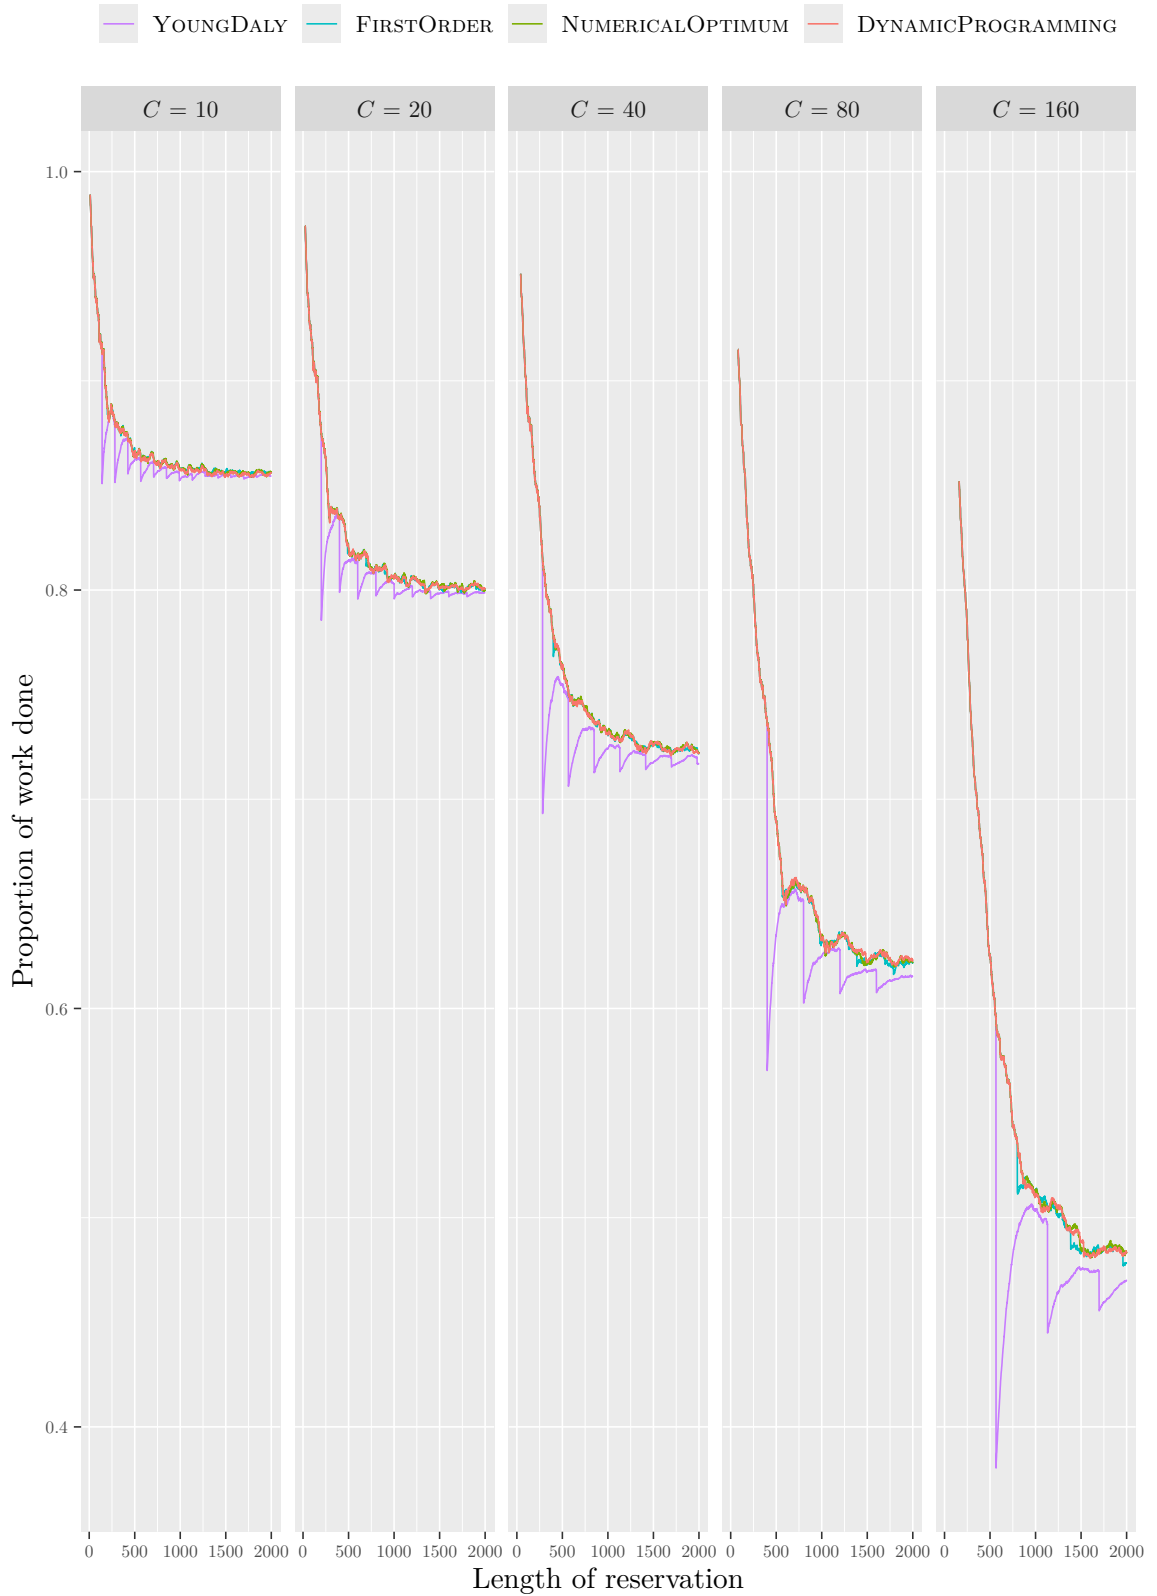
Inria

Figure 8: Proportion of work completed by the different checkpointing strategies when $\lambda = 0.0001$ and $D = 0$.



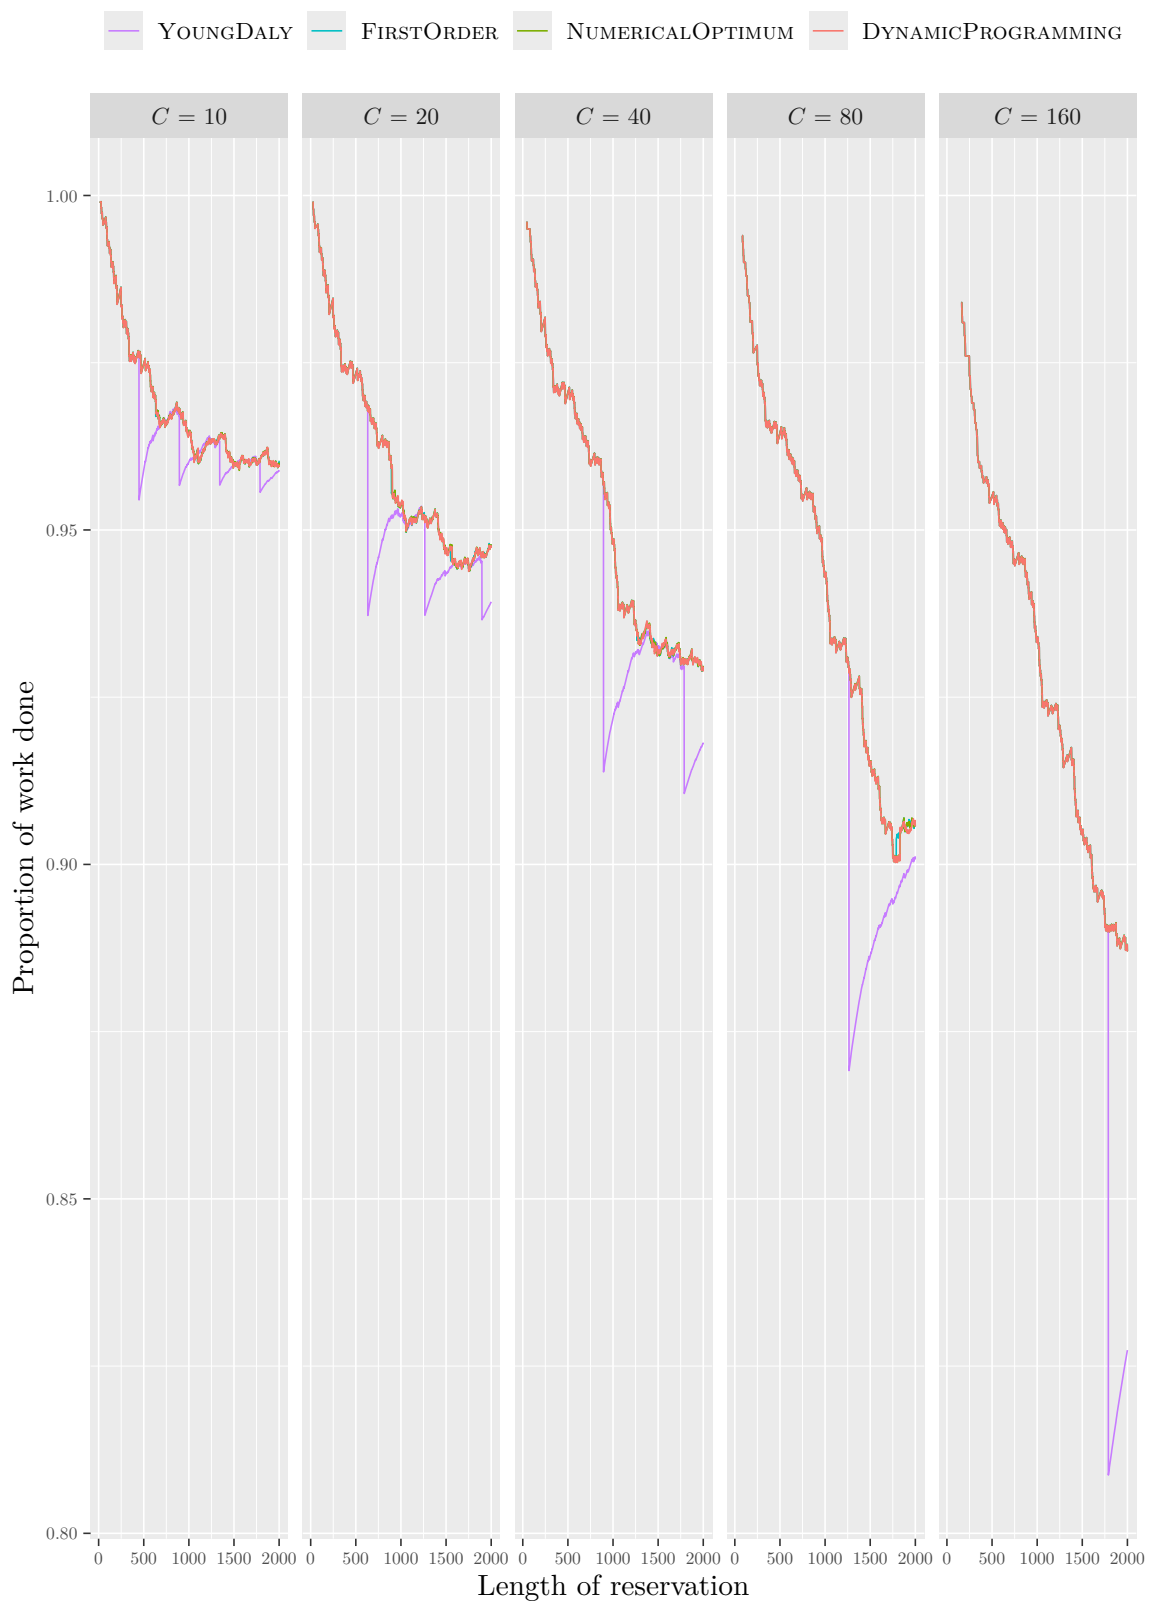
RR n° 9552

Figure 9: Proportion of work completed by the different checkpointing strategies when $\lambda = 0.01$ and $D = 5$.



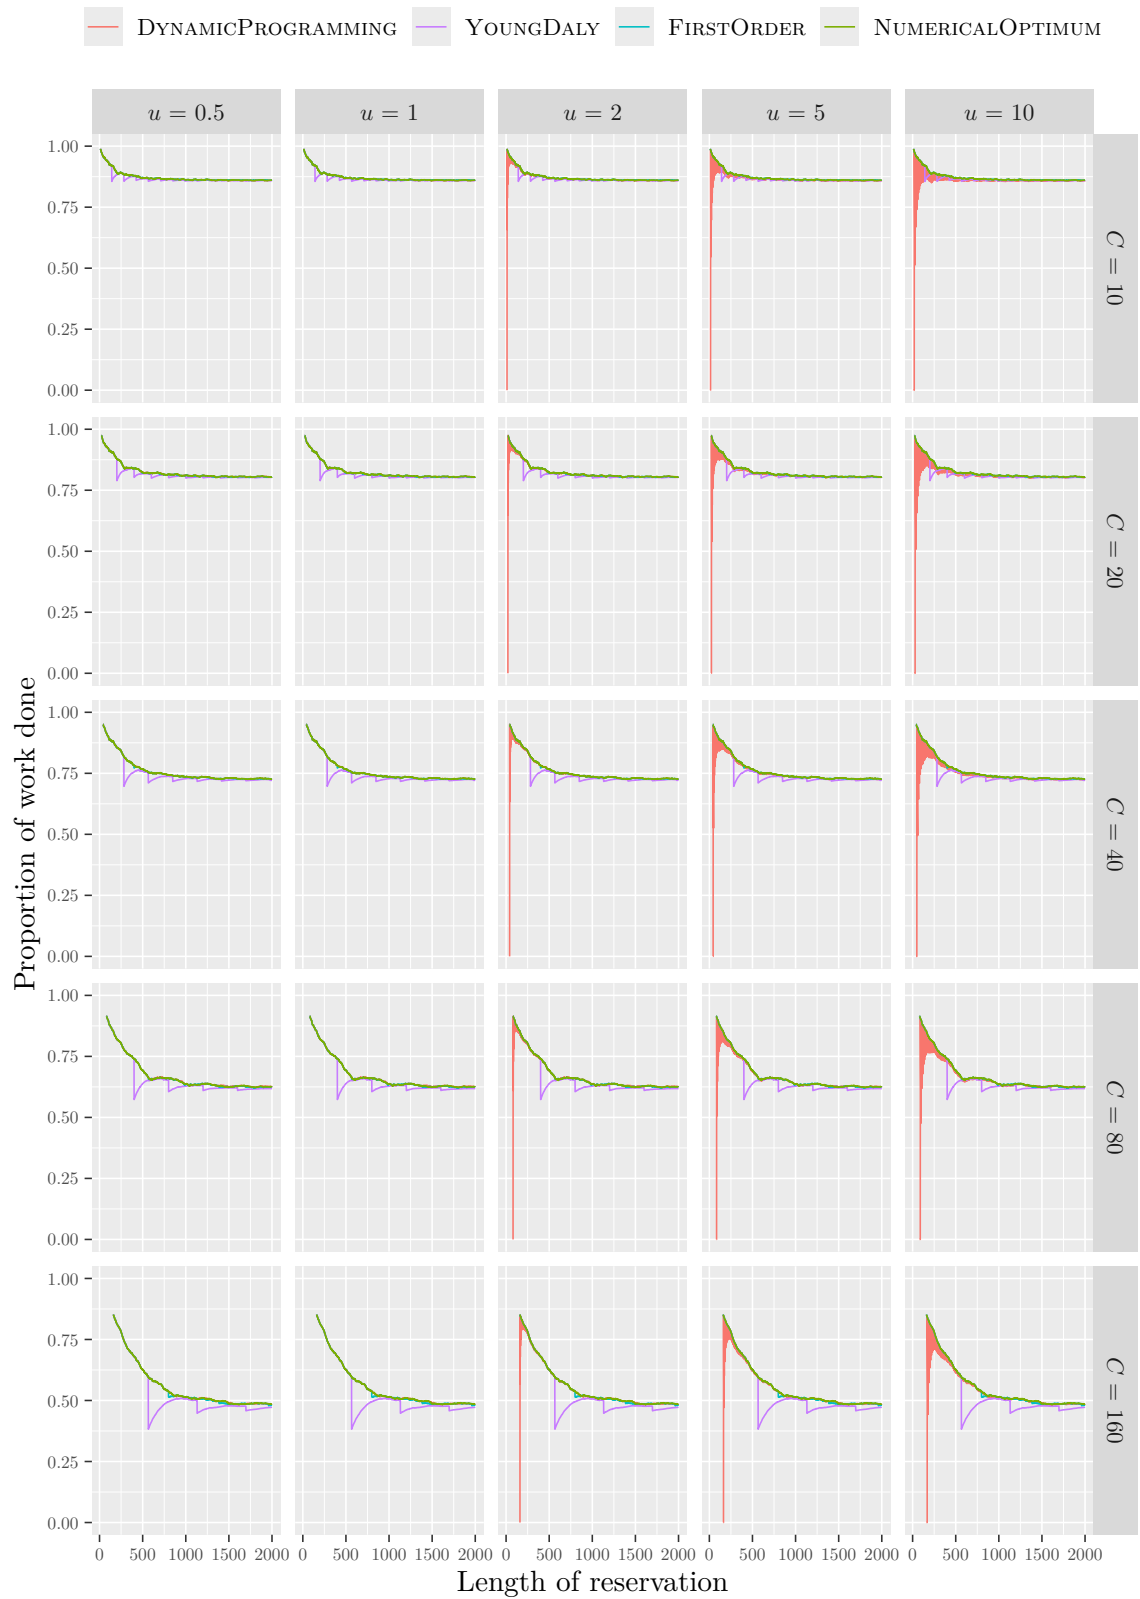
Inria

Figure 10: Proportion of work completed by the different checkpointing strategies when $\lambda = 0.001$ and $D = 5$.



RR n° 9552

Figure 11: Proportion of work completed by the different checkpointing strategies when $\lambda = 0.0001$ and $D = 5$.



Inria

Figure 12: Proportion of work completed by the different checkpointing strategies when $\lambda = 0.0001$ and $D = 0$.

Inria

**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

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

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399