

# Towards Optimal Multi-Level Checkpointing

Anne Benoit<sup>1</sup>, Aurélien Cavelan<sup>1</sup>,  
Valentin Le Fèvre<sup>1</sup>, Yves Robert<sup>1,2</sup>, **Hongyang Sun<sup>1</sup>**

1. ENS Lyon & INRIA, France
2. University of Tennessee Knoxville, USA

[hongyang.sun@ens-lyon.fr](mailto:hongyang.sun@ens-lyon.fr)

JLESC Workshop, Lyon  
June 28, 2016

# Single-Level Checkpointing

Minimize expected execution overhead  $H(W) = \frac{\mathbb{E}(W)}{W} - 1$




Figure: periodic computing pattern

**What is the optimal checkpointing interval?**

# Single-Level Checkpointing

Minimize expected execution overhead  $H(W) = \frac{\mathbb{E}(W)}{W} - 1$




Figure: periodic computing pattern

## What is the optimal checkpointing interval?

- Exact solution:

$$H(W) = \frac{e^{\lambda R} \left( \frac{1}{\lambda} + D \right) e^{\lambda(W+C)}}{W} - 1, \text{ use Lambert function}$$

- First-order approximation [Young/Daly]:

$$W_{\text{opt}} = \sqrt{\frac{2C}{\lambda}}$$

$$H_{\text{opt}} = \sqrt{2\lambda C} + \Theta(\lambda)$$

# Single-Level Checkpointing

Minimize expected execution overhead  $H(W) = \frac{\mathbb{E}(W)}{W} - 1$




Figure: periodic computing pattern

## What is the optimal checkpointing interval?

- Exact solution:

$$H(W) = \frac{e^{\lambda R} \left( \frac{1}{\lambda} + D \right) e^{\lambda(W+C)}}{W} - 1, \text{ use Lambert function}$$

- First-order approximation [Young/Daly]:

$$W_{\text{opt}} = \sqrt{\frac{2C}{\lambda}}$$
$$H_{\text{opt}} = \sqrt{2\lambda C} + \Theta(\lambda)$$

## Scalability problem for large-scale platforms

# Multi-Level Checkpointing

E.g., Scalable Checkpoint/Restart (SCR) library, Fault Tolerance Interface (FTI)

- Local memory/SSD, Partner copy/XOR, RS-coding, PFS


# Multi-Level Checkpointing

E.g., Scalable Checkpoint/Restart (SCR) library, Fault Tolerance Interface (FTI)


- Local memory/SSD, Partner copy/XOR, RS-coding, PFS

Two approaches:

- Independent checkpointing:



- Synchronized checkpointing:



Easy because pattern repeats (memoryless property)




Figure: with  $N$  level-1 checkpoints

- Exact solution: very complicated (which error type occurs first?), equal-length chunks, see [1]

---

[1] S. Di, Y. Robert, F. Vivien, F. Cappello. Toward an optimal online checkpoint solution under a two-level HPC checkpoint model, *IEEE TPDS*, 2016.

Easy because pattern repeats (memoryless property)




Figure: with  $N$  level-1 checkpoints

- Exact solution: very complicated (which error type occurs first?), equal-length chunks, see [1]
- First-order approximation:

$$H_{\text{opt}} = \sqrt{2\lambda_1 C_1} + \sqrt{2\lambda_2 C_2} + \Theta(\lambda)$$

(obtained for some optimal pattern)


---

[1] S. Di, Y. Robert, F. Vivien, F. Cappello. Toward an optimal online checkpoint solution under a two-level HPC checkpoint model, *IEEE TPDS*, 2016.



# Three Levels


Difficult because sub-patterns may differ



- Exact solution: unknown

# Three Levels

Difficult because sub-patterns may differ




- Exact solution: unknown
- First-order approximation:

$$H_{\text{opt}} = \sqrt{2\lambda_1 C_1} + \sqrt{2\lambda_2 C_2} + \sqrt{2\lambda_3 C_3} + \Theta(\lambda)$$

(obtained for some optimal pattern)

Difficult because sub-patterns may differ



- Exact solution: unknown
- First-order approximation:

$$H_{\text{opt}} = \sqrt{2\lambda_1 C_1} + \sqrt{2\lambda_2 C_2} + \sqrt{2\lambda_3 C_3} + \Theta(\lambda)$$

(obtained for some optimal pattern)

- Choose optimal set of levels:

Level	Overhead
1, 2, 3	$\sqrt{2C_1\lambda_1} + \sqrt{2C_2\lambda_2} + \sqrt{2C_3\lambda_3}$
1, 3	$\sqrt{2C_1\lambda_1} + \sqrt{2C_3(\lambda_2 + \lambda_3)}$
2, 3	$\sqrt{2C_2(\lambda_1 + \lambda_2)} + \sqrt{2C_3\lambda_3}$
3	$\sqrt{2C_3(\lambda_1 + \lambda_2 + \lambda_3)}$

## Theorem

The optimal  $k$ -level pattern, under the first-order approximation, has equal-length chunks at all levels:

$$\text{Optimal pattern length: } W^{\text{opt}} = \sqrt{\frac{\sum_{l=1}^k N_l^{\text{opt}} C_l}{\frac{1}{2} \sum_{l=1}^k \frac{\lambda_l}{N_l^{\text{opt}}}}}$$

$$\text{Optimal \#chkpts at level } l: N_l^{\text{opt}} = \sqrt{\frac{\lambda_l}{C_l} \cdot \frac{C_k}{\lambda_k}}, \quad \forall l = 1, \dots, k$$

$$\text{Optimal pattern overhead: } H_{\text{opt}} = \sum_{l=1}^k \sqrt{2\lambda_l C_l} + \Theta(\lambda)$$

## Theorem

The optimal  $k$ -level pattern, under the first-order approximation, has equal-length chunks at all levels:

$$\text{Optimal pattern length: } W^{\text{opt}} = \sqrt{\frac{\sum_{l=1}^k N_l^{\text{opt}} C_l}{\frac{1}{2} \sum_{l=1}^k \frac{\lambda_l}{N_l^{\text{opt}}}}}$$

$$\text{Optimal \#chkpts at level } l: N_l^{\text{opt}} = \sqrt{\frac{\lambda_l}{C_l} \cdot \frac{C_k}{\lambda_k}}, \quad \forall l = 1, \dots, k$$


$$\text{Optimal pattern overhead: } H_{\text{opt}} = \sum_{l=1}^k \sqrt{2\lambda_l C_l} + \Theta(\lambda)$$

- Dynamic programming algorithm to choose set of levels


- Rounding for integer solution:  $n_l^{\text{opt}} = \frac{N_l^{\text{opt}}}{N_{l+1}^{\text{opt}}} = \sqrt{\frac{\lambda_l}{\lambda_{l+1}} \cdot \frac{C_{l+1}}{C_l}}$

# Simulations

Set	Source	Level	1	2	3	4
(A)	Moody et al. [1]	C (s)	0.5	4.5	1051	-
		MTBF (s)	5.00e6	5.56e5	2.50e6	-
(B)	Balaprakash et al. [2]	C (s)	10	20	20	100
		MTBF (s)	3.60e4	7.20e4	1.44e5	7.20e5



(A)



(B)

[1] A. Moody, G. Bronevetsky, K. Mohror, and B. R. de Supinski. Design, modeling, and evaluation of a scalable multi-level checkpointing system. *Supercomputing*, 2010.

[2] P. Balaprakash, L. A. Bautista-Gomez, M.-S. Bouguerra, S. M. Wild, F. Cappello, and P. D. Hovland. Analysis of the tradeoffs between energy and run time for multilevel checkpointing. *PMBS*, 2014.

Explicit formulas for (almost) optimal multi-level checkpointing

$$H_{\text{opt}} = \sum_{\ell=1}^k \sqrt{2\lambda_{\ell} C_{\ell}} + \Theta(\lambda)$$

Limitations:

- First-order approximation (accurate for 10,000s of nodes with MTBF in hours; beyond?)
- Independent errors (correlated failure?)