I/O interference benchmarking

Méline Trochon

F. Boito, F. Tessier, B. Goglin, J. Acquaviva

Per3S – May 23rd, 2025

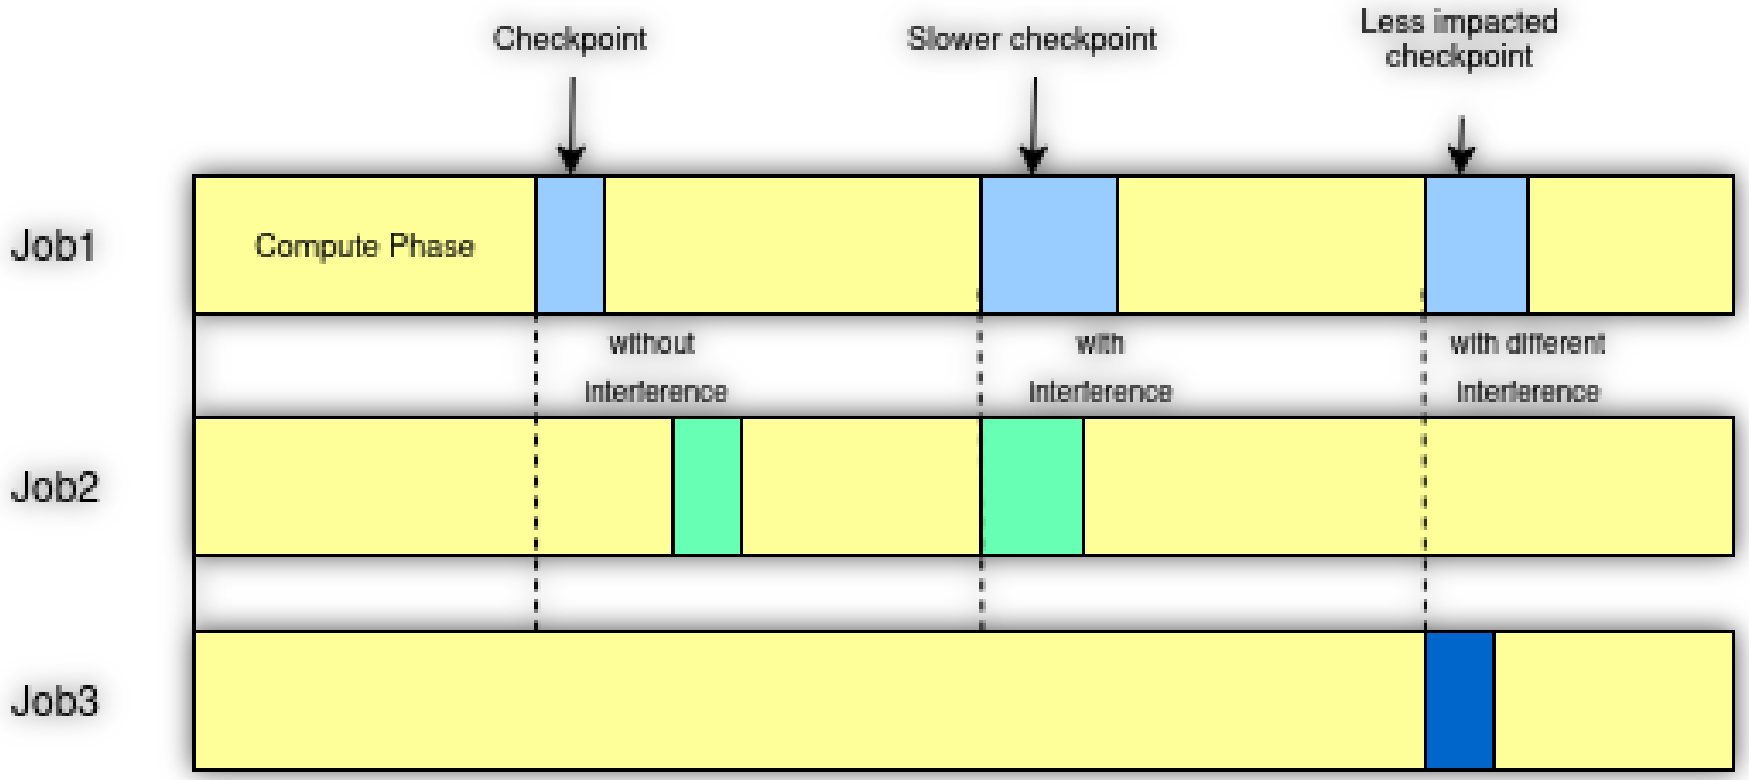

- Tracing concurrent I/O-intensive workloads
- Study the impact of different parameters

We identified that small I/O accesses harm overall I/O performance

I/O Interference Benchmarking

Méline Trochon - F. Boito, F. Tessier, B. Goglin, J. Acquaviva
Inria - DataDirect Networks

1. Context: Checkpointing and I/O contention

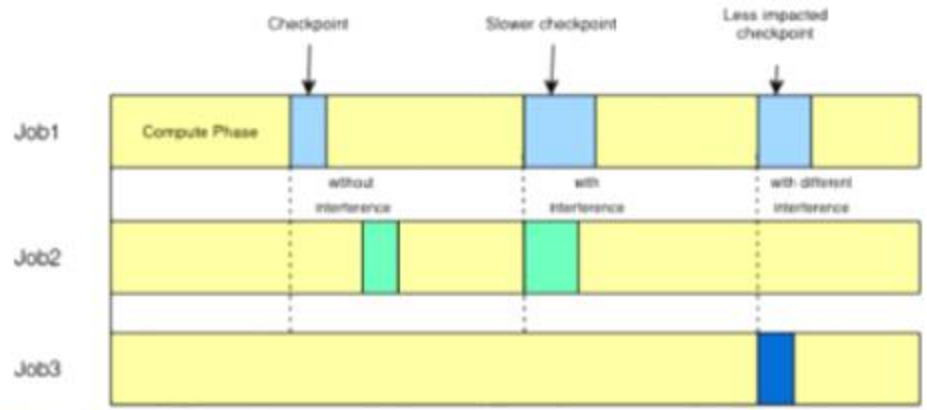

Figure: Job execution pattern, with the compute phase in yellow, the I/O phase in blue. Each line represents a different job execution.

On HPC systems, jobs are executed concurrently. Concurrent I/O activity leads to **increased duration and high variability** in I/O phases. This work aims at **characterizing the interference** generated by different application behaviors.

2. I/O characterisation - Methodology

We executed two simultaneous I/O-intensive workloads. The first, called **checkpoint**, is a contiguous write with large transfer sizes to mimic a one-file-per-process C/R mechanism. The second, **interference**, simulates typical parallel I/O patterns or stresses known PFS weaknesses like lock contention.

Table: Checkpoint and interference parameters.
Partitioning corresponds to the number of nodes and processes per nodes.
File strat. is either file per process (FPP) or single shared file (SSF).

	Partitioning	File strat.	Block size	Transfer size
ckpt	8 x 8	FPP	6 GiB	8 MiB
interf.	4-16 x 4-16	FPP/SSF	3-12 GiB	4KiB - 8 MiB

Checkpoint parameters were tuned to **saturate the system's bandwidth**. For each benchmark, the interference parameters are equal to the checkpoint parameters, unless specified in the plot. Benchmarks are run on a local cluster at Inria Bordeaux (around 40 nodes), using **BeeGFS** with an interconnection of 100Gbit/s.