Automated I/O Parameter Tuning of Scientific Applications with Parametrizable Workload Replays

Azat Nurgaliev and Marcus Paradies
German Aerospace Center, Institute of Data Science
USENIX FAST'21 WiP Reports, 24 February 2021

Azat.Nurgaliev@dlr.de

DLR

Problem Description

- I/O is a major bottleneck in scientific data-intensive applications
- SSDs oftentimes used “as-is” without proper tuning of the application to fully leverage performance
- SSD has different internal parameters (memory types, internal architecture, controllers)
- Tuning the application (e.g., parallelism level and block size) requires access to source code (oftentimes not available because of containerization)

How to identify optimal I/O application parameter configurations for a given storage device?

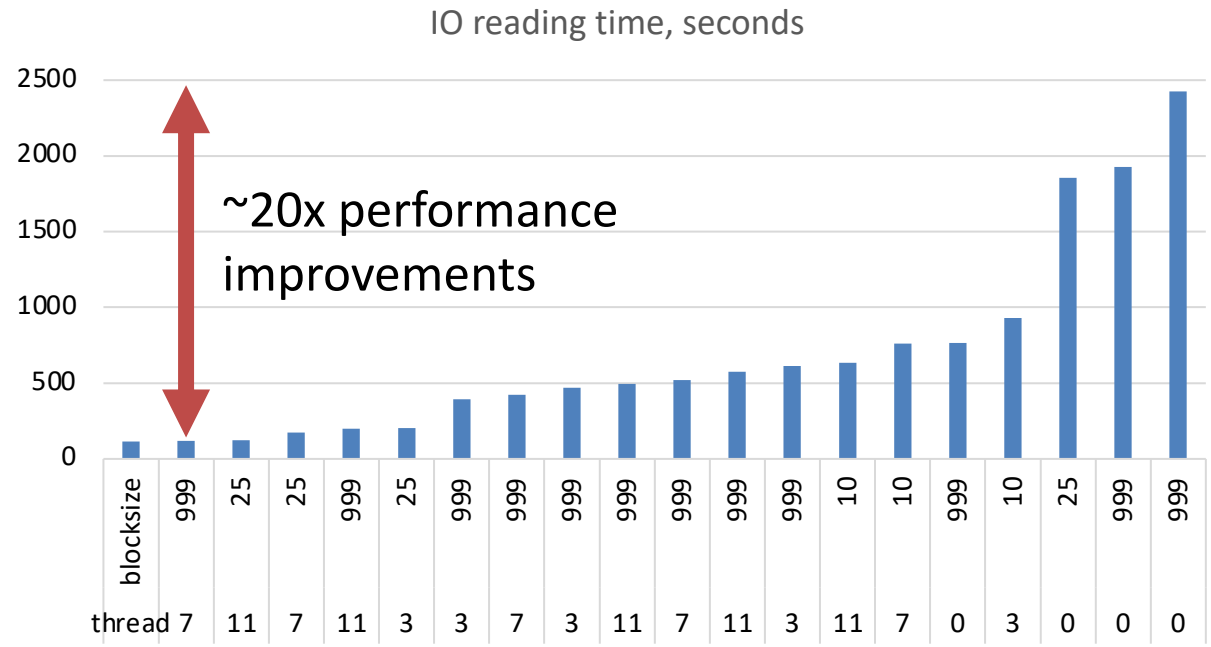

Real application using GDAL with different I/O parameters

Existing I/O Analysis and Tracing Tools

Workload Replay tools (e.g. btreplay, hfrplayer)	Pros	High precision mimicking
	Cons	No flexibility on parameter changing
		Destructive for file system
Microbenchmarking tools (e.g. fio, dd)	Pros	High flexibility on parameter changing
	Cons	Workload far too distinct from the real workload

Approach & Future work

Open questions:

- How to run application correctly? (sample run)
- How to accurately represent the computational part of the application?

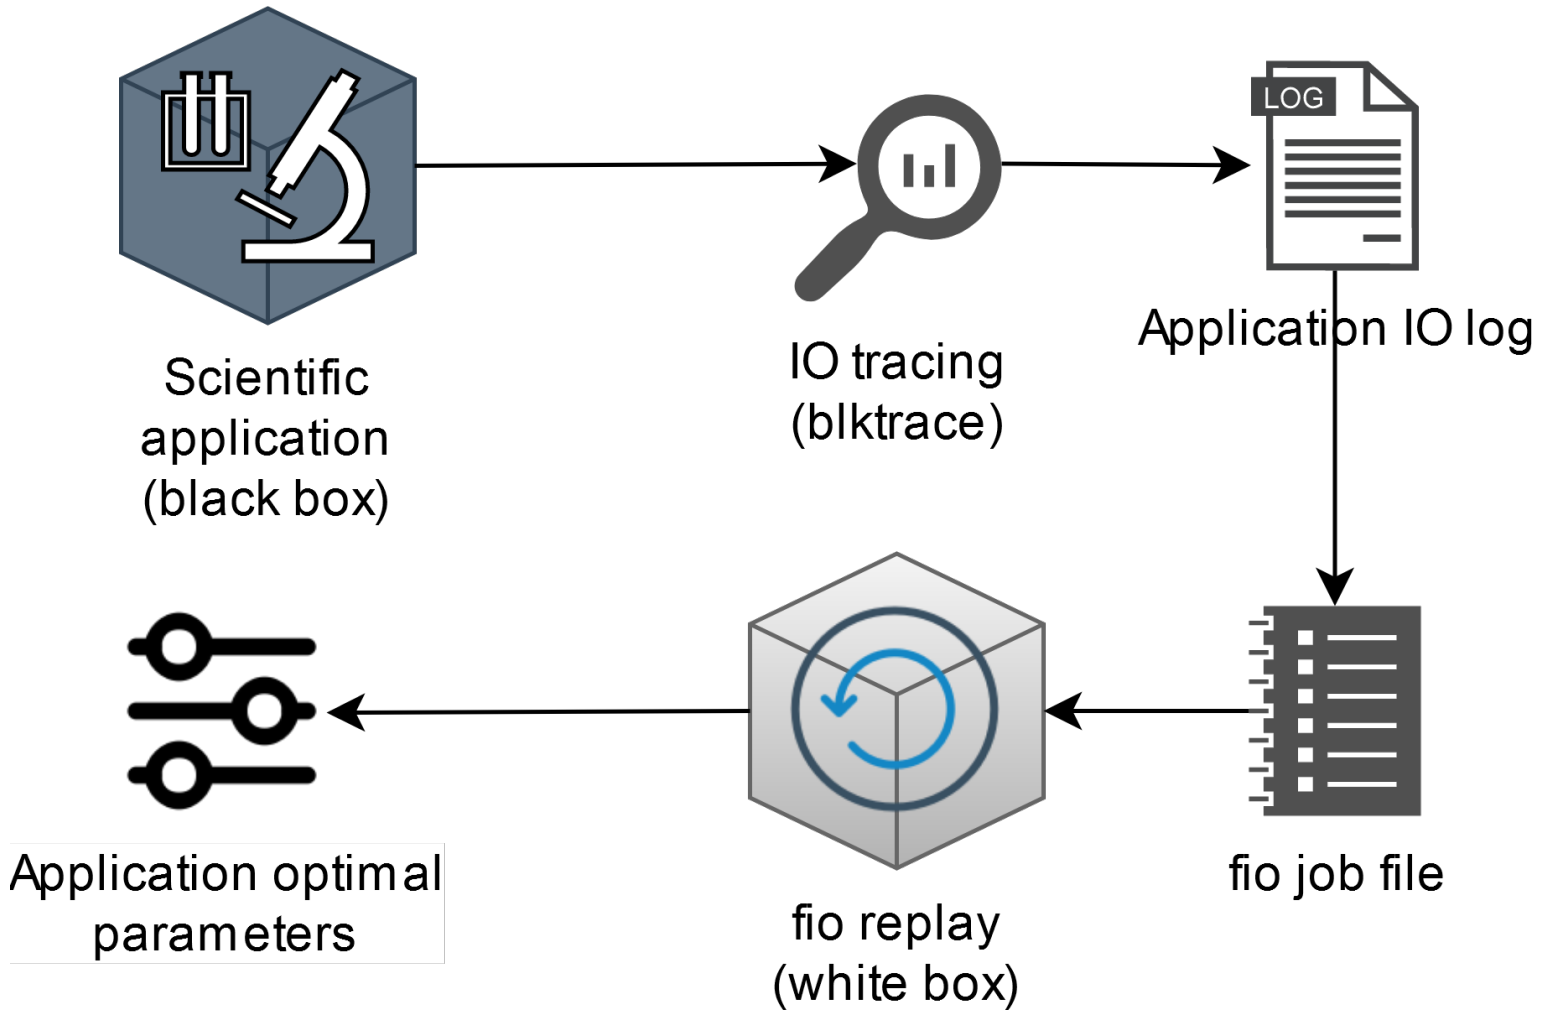

Automated I/O Parameter Tuning of Scientific Applications with Parametrizable Workload Replays

Azat Nurgaliev and Marcus Paradies
German Aerospace Center, Institute of Data Science
USENIX FAST'21 WiP Reports, 24 February 2021

Azat.Nurgaliev@dlr.de

DLR

Problem Description

- I/O is a major bottleneck in scientific data-intensive applications
- SSDs oftentimes used "as-is" without proper tuning of the application to fully leverage performance
- SSD has different internal parameters (memory types, internal architecture, controllers)
- Tuning the application (e.g., parallelism level and block size) requires access to source code (oftentimes not available because of containerization)

How to identify optimal I/O application parameter configurations for a given storage device?

Real application using GDAL with different I/O parameters

DLR

Existing I/O Analysis and Tracing Tools

Workload Replay tools (e.g. btreplay, hfrplayer)	Pros	High precision mimicking
	Cons	No flexibility on parameter changing
		Destructive for file system
Microbenchmarking tools (e.g. fio, dd)	Pros	High flexibility on parameter changing
	Cons	Workload far too distinct from the real workload

DLR

Approach & Future work

Open questions:

- How to run application correctly? (sample run)
- How to accurately represent the computational part of the application?

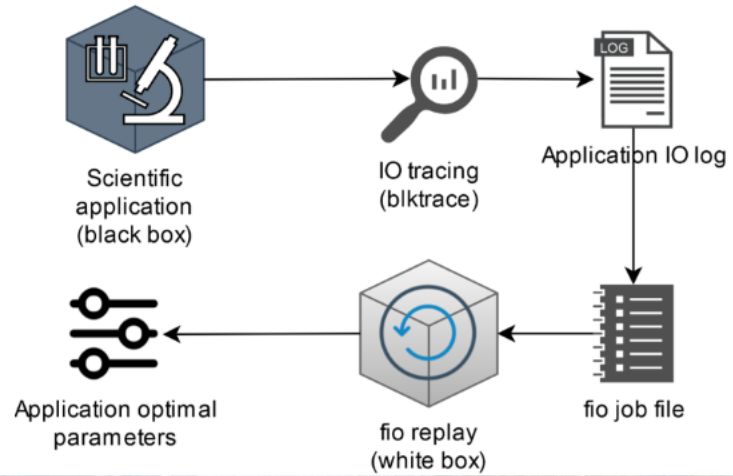

DLR