

Modeling the Allocation of Heterogeneous Storage Resources on HPC Systems

PhD Student:
Julien Monniot – julien.monniot@inria.fr
INRIA Rennes / Université Rennes 1

Advisors:
François Tessier – INRIA Rennes
Gabriel Antoniu – INRIA Rennes



Up to 3PB/year by 2027

SKA Telescope

1 PB/day by 2026

ECMWF Forecast & Digital Twin Earth

2 PB/day by 2035

ITER Reactor

Burst-buffers / staging area (SSD, NVMeoF, HDD, ...)

Node-local / Platform integrated (SSD, NVRAM, ...)

PFS / Archives (HDD, Tapes)

Hot

Cold

Examples:

Perlmutter 35PB all flash storage at NERSC (2021)

Aurora DAOS Storage with Intel Optane persistent memory at ALCF (2022)

Summit compute nodes with 1.6TB of NVRAM at OLCF (2018)

• Data deluge from new large-scale scientific workflows

• PFPlops TBps

• Deeper storage hierarchy

• New underlying technologies

• Hybrid platforms / workflows

Complexity & Underutilization of resources

Figure 1. Top500: I/O bandwidth to computing power ratio for the top 3 HPC platforms, last 10 years.

From HPC... + ... to Cloud

Optimal Resource Scheduling

MAIN AXES OF WORK

WHAT DO WE AIM FOR?

• Design of scheduling algorithms for storage allocations

• Representation of heterogeneous storage infrastructures

• Analysis of storage related metrics

Enable dynamic allocation of heterogeneous storage resources (real or simulated)

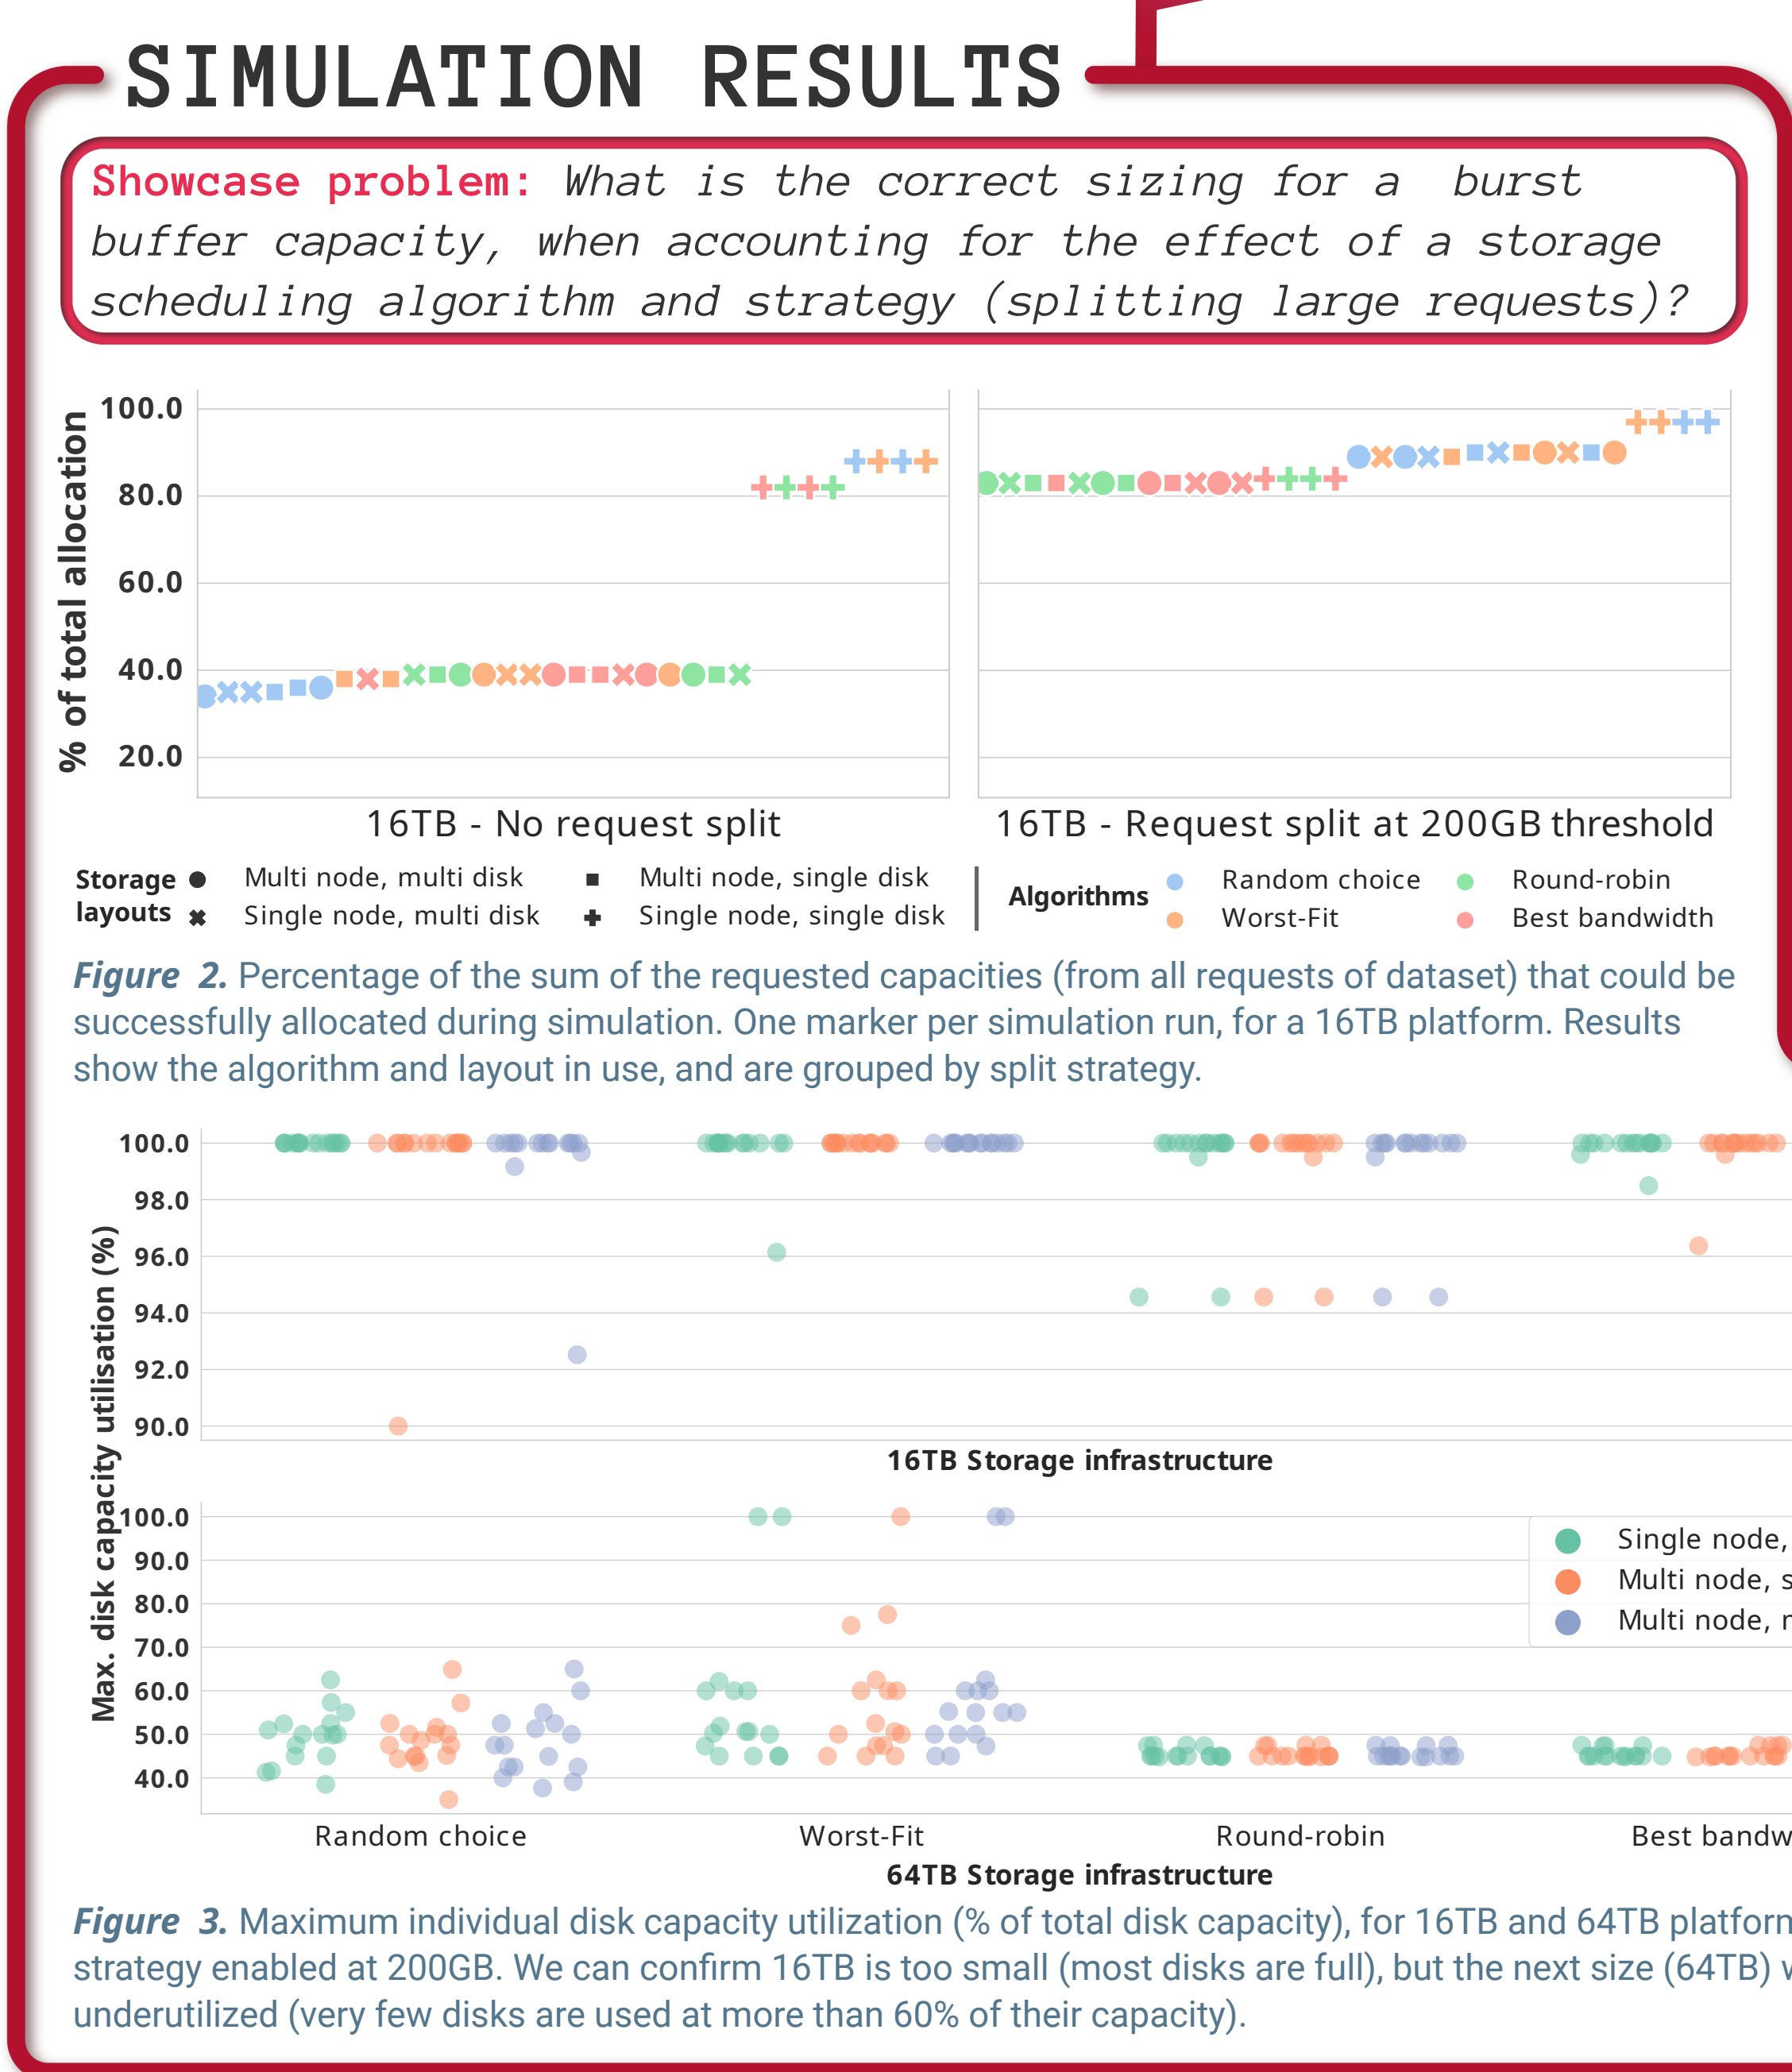
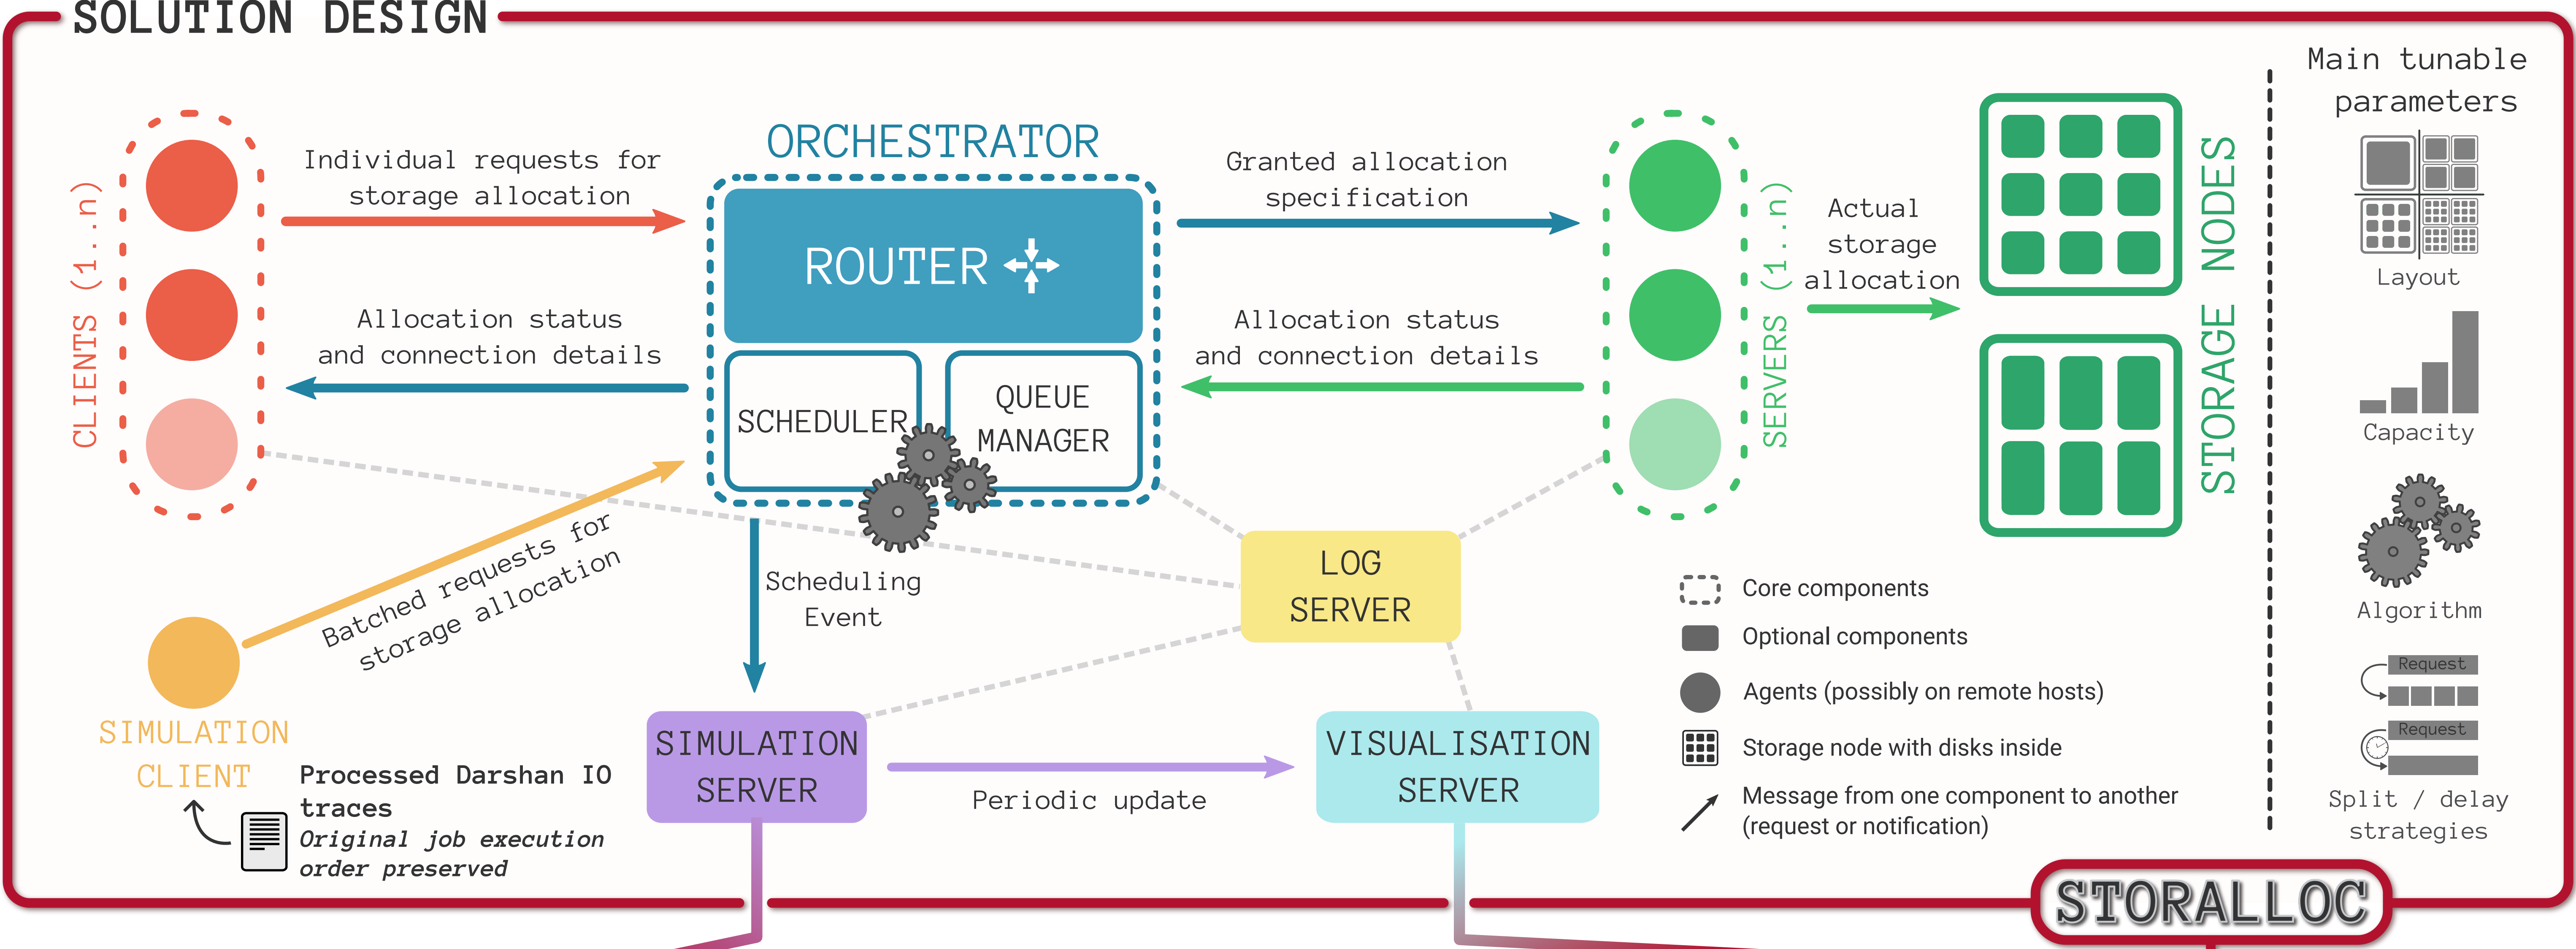
Dealing with resource heterogeneity

1

2

3

Fair and efficient use of resources



Storage-aware job scheduler simulator

Extensible with new algorithms

Independent components design

Common messaging interface

Abstraction of heterogeneous storage hardware

IMPLEMENTATION

Python3

SimPy (DES simulation)

ZMQ (Messaging)

bokeh (Real-time plotting)

seaborn (Plotting)

Darshan (IO traces)

GitHub

hephtaicie/storalloc