A hierarchical file system for efficient pipeline processing on HPC

Université Concordia University

Valerie Hayot-Sasson and Tristan Glatard

4th BigBrain Workshop
June 26, 2020

Introduction

- High-resolution neuroimaging data becoming increasingly popular
- Process data on HPC using portals (e.g. CBRAIN)
- Data accessed from costly shared parallel FS (e.g. Lustre)
- Big Data file systems reduce I/O costs through:
  - Data locality
  - In-memory computing

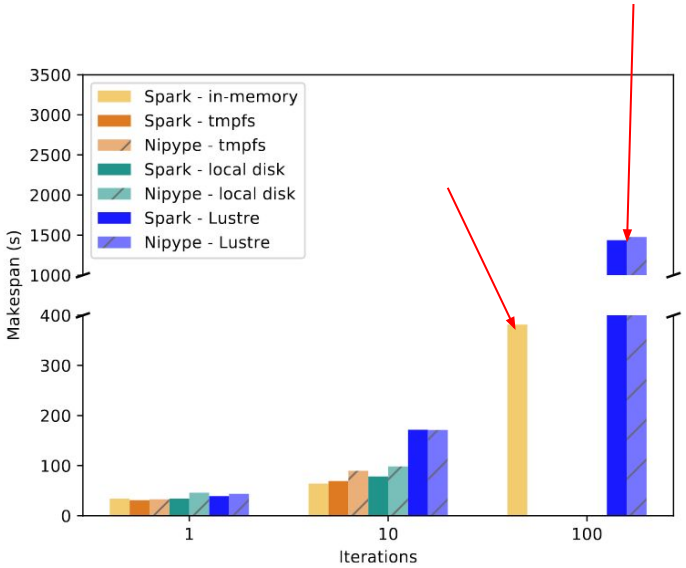

[Hayot-Sasson et al. 2019]

Objective: Incorporate Big Data optimizations in neuroimaging pipelines without instrumentation

Methodology: The Sea filesystem

- Lightweight library intercepting libc calls through LD_PRELOAD
- Access to filesystem enabled via mountpoint
- User-specified file system hierarchy
- Can be in-memory or with flushing to shared fs

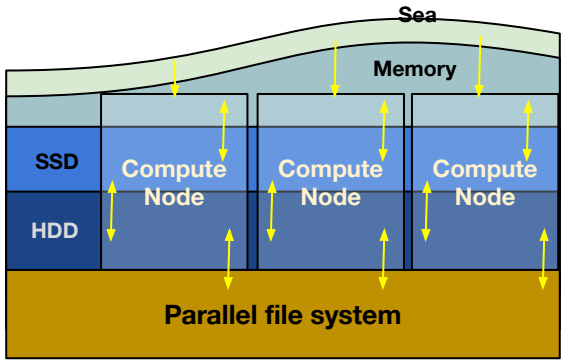

Sea overview

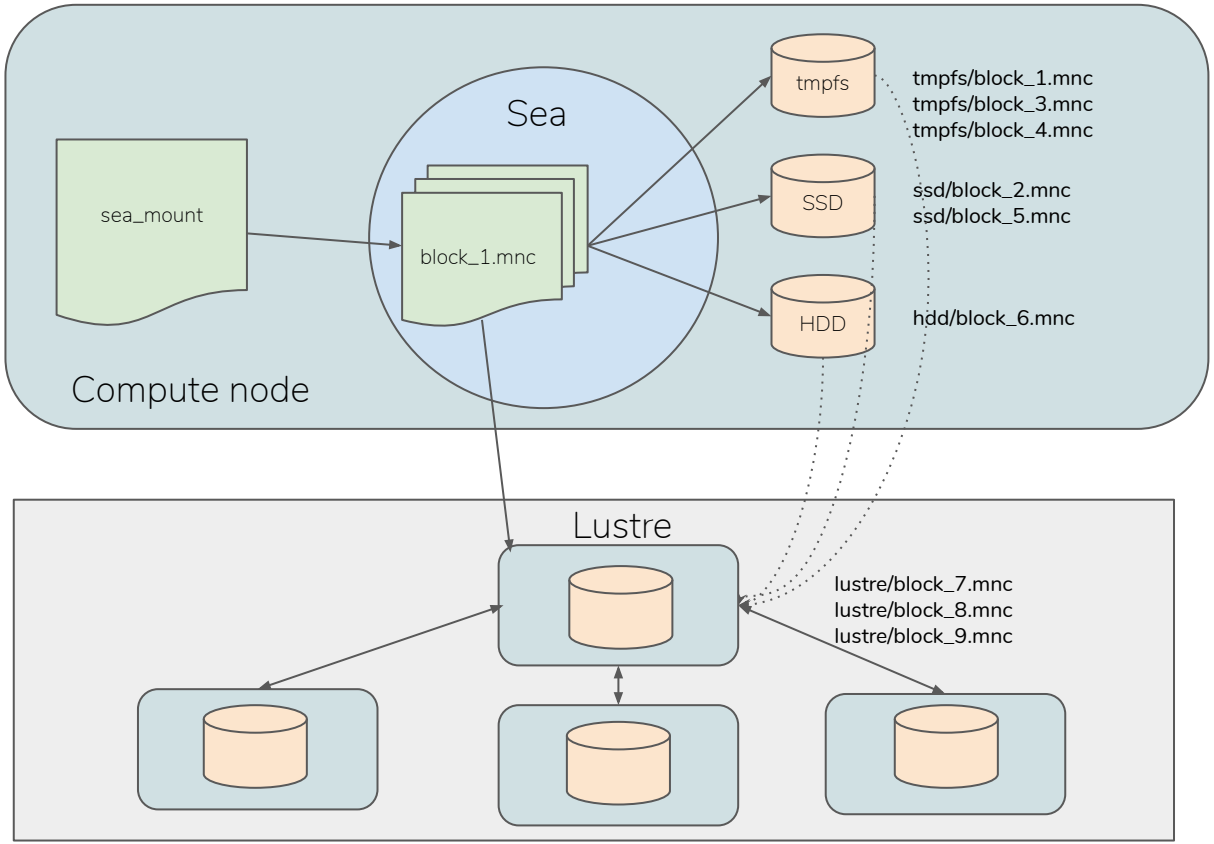


```
$ ls sea_mount/
block_1.mnc
block_2.mnc
block_3.mnc
block_4.mnc
block_5.mnc
block_6.mnc
block_7.mnc
block_8.mnc
block_9.mnc
```

Preliminary results

Experiment:

- Incremented the value of a 617MiB Big Brain block 20x
- First read always on Lustre

Infrastructure:
Shared lab cluster with 126GB tmpfs on each node

Sea and Memory (tmpfs) performance **comparable**

Sea speedup ~2x Lustre

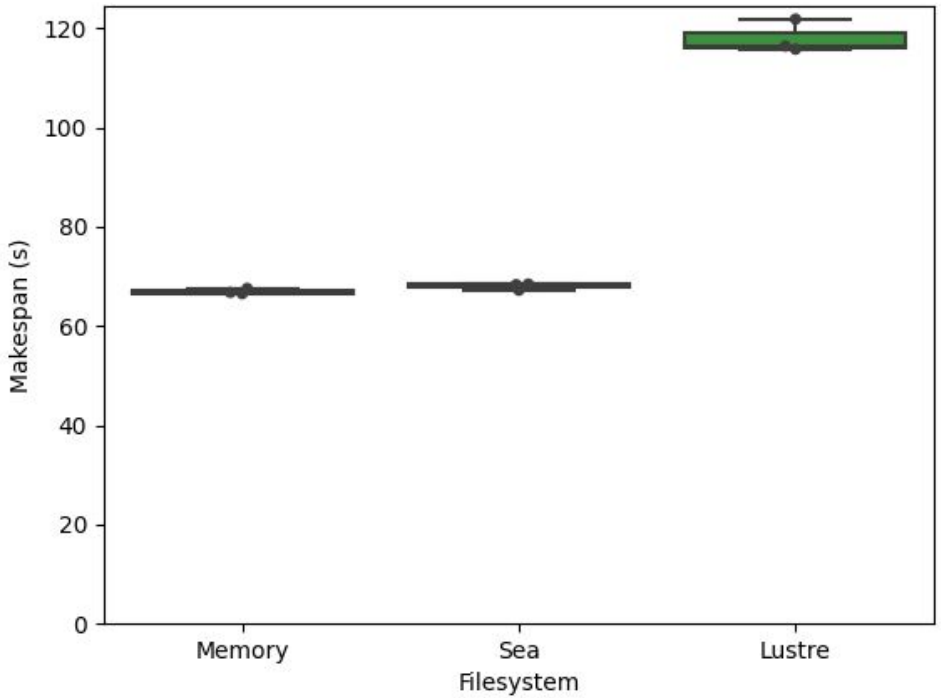

Thank you!

Follow the project: https://github.com/ValHayot/Sea

valeriehayot@gmail.com