Improving I/O Performance of HPC Applications Using Intra-Job Scheduling

Arnab K. Paul[†], Olaf Faaland[‡], Adam Moody[‡],
Elsa Gonsiorowski[‡], Kathryn Mohror[‡], Ali R. Butt[†]

[†]***Virginia Tech***, [‡]*Lawrence Livermore National Laboratory*

PDSW-DISCS 2019; collocated with SC'19, Denver, CO

Motivation: The Increasing Gap

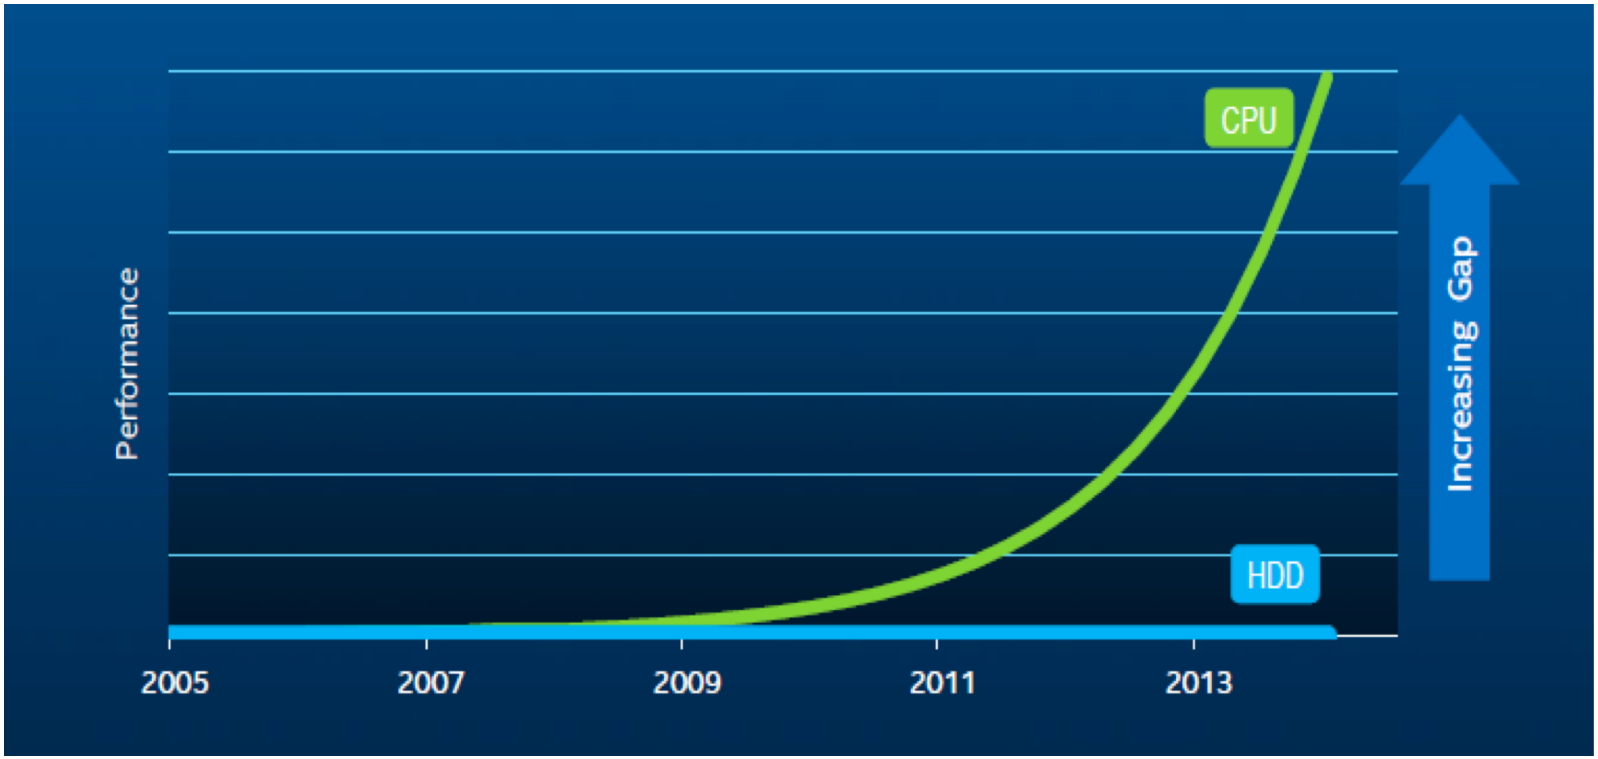

Processor Performance vs Disk Access Time

https://newsroom.intel.com/editorials/3d-xpoint-memory-storage/#gs.gqtcop

Motivation

I/O operations become a limiting factor in application efficiency.

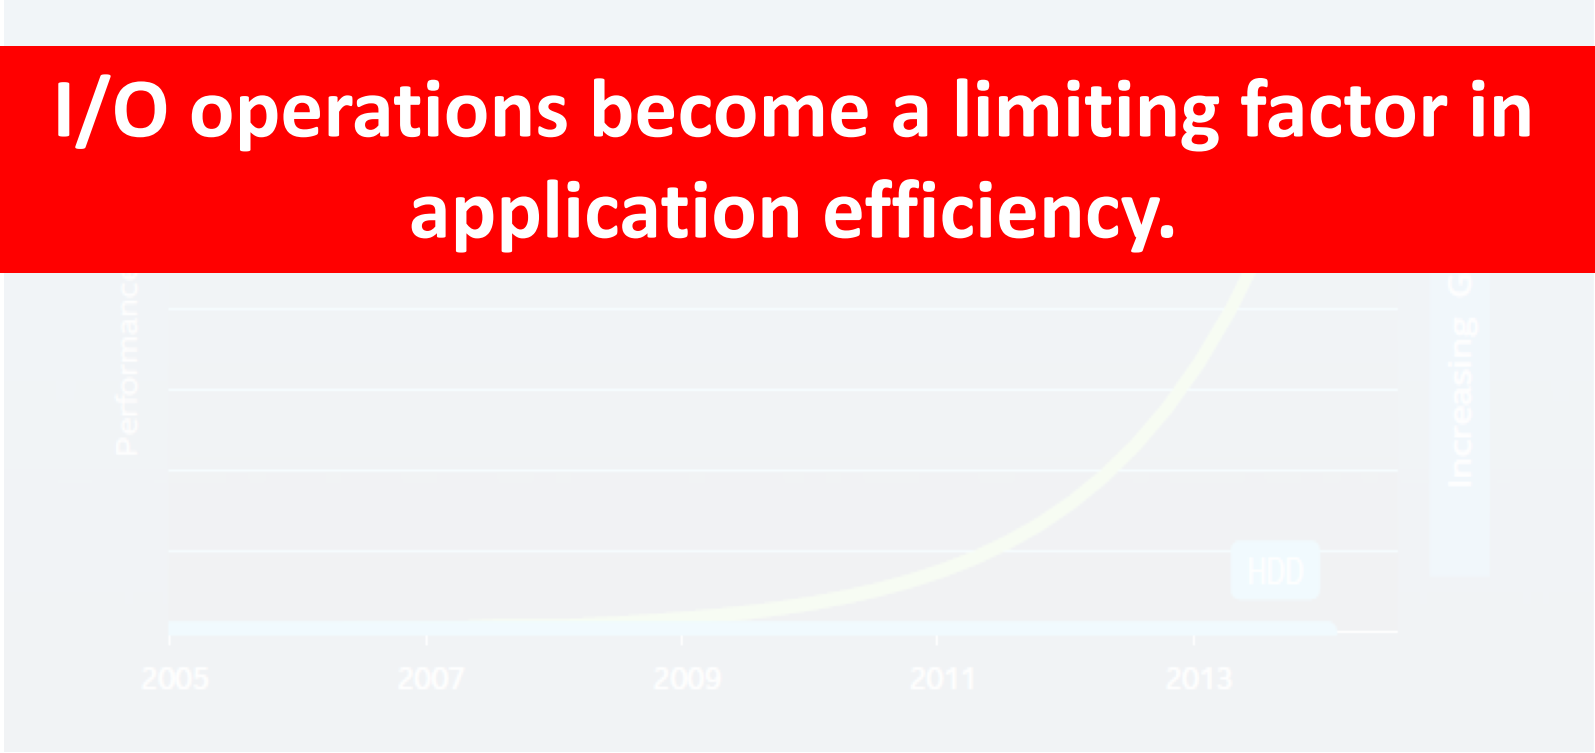

Processor Performance vs Disk Access Time

https://newsroom.intel.com/editorials/3d-xpoint-memory-storage/#gs.gqtcop

Motivation

I/O operations become a limiting factor in application efficiency.

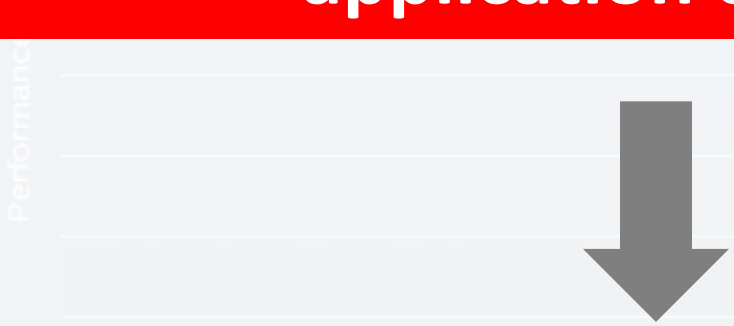

Improve I/O Performance of HPC Applications Using Intra-Job Scheduling

https://newsroom.intel.com/editorials/3d-xpoint-memory-storage/#gs.gqtcop

Lustre Parallel File System

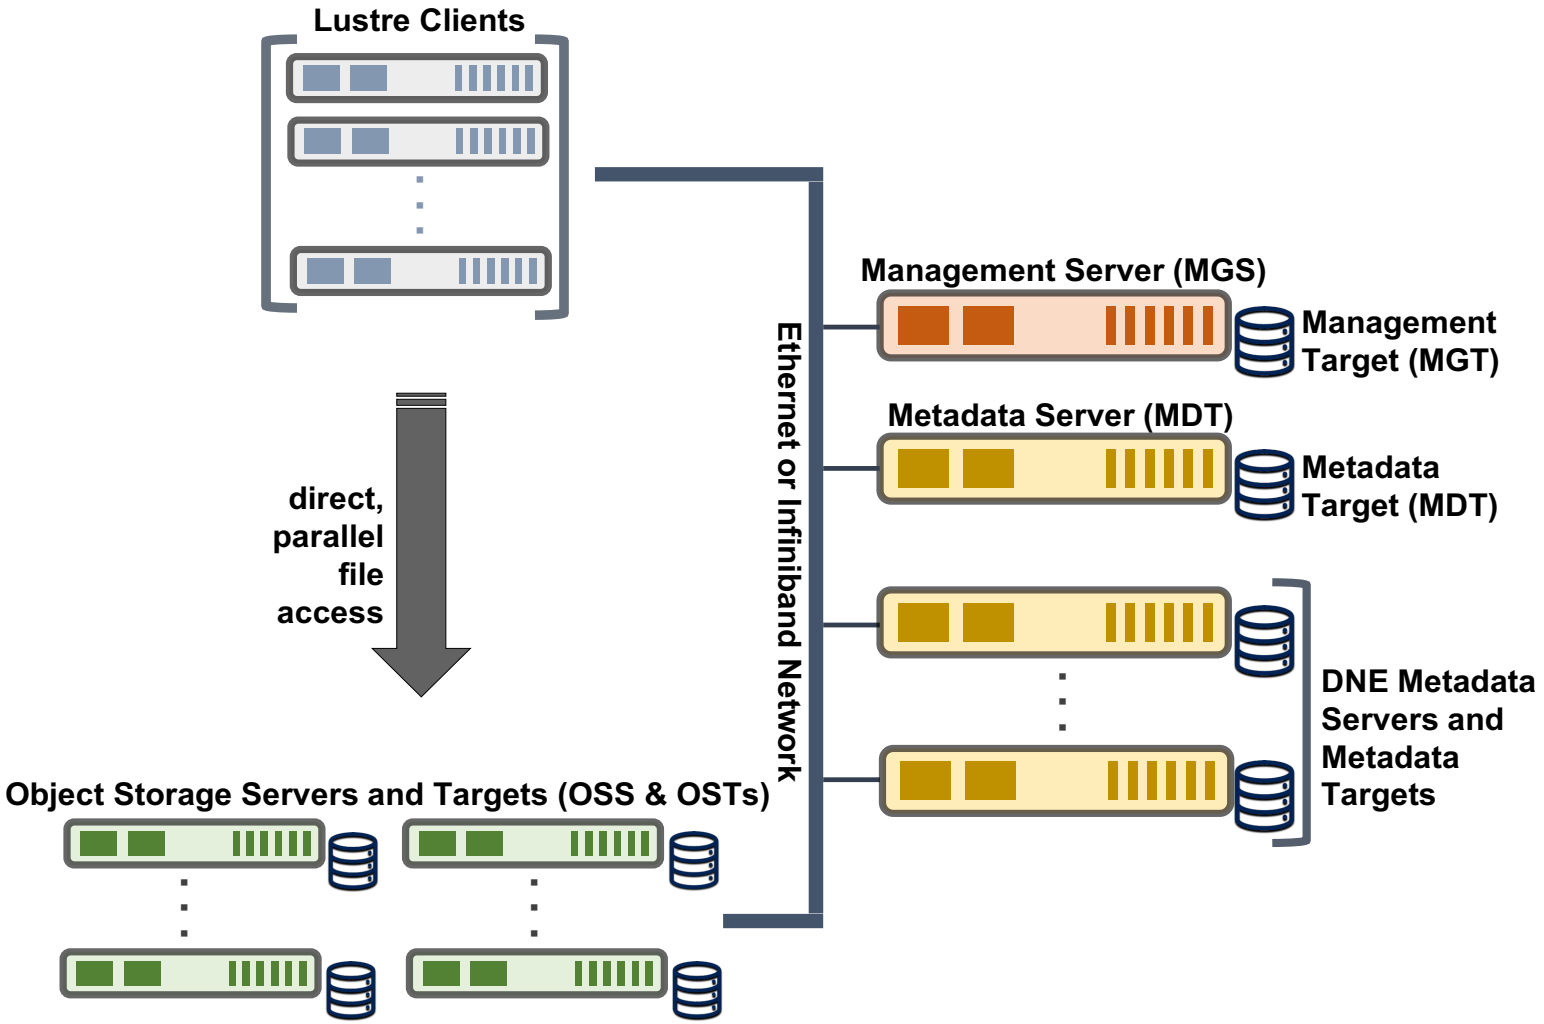

System Design

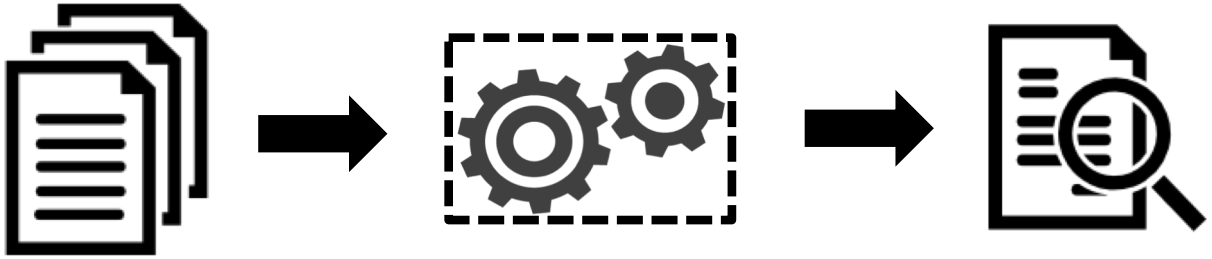

Job Statistics Dataset

Machine Learning Modeling

Validation

Models are stored

System Design

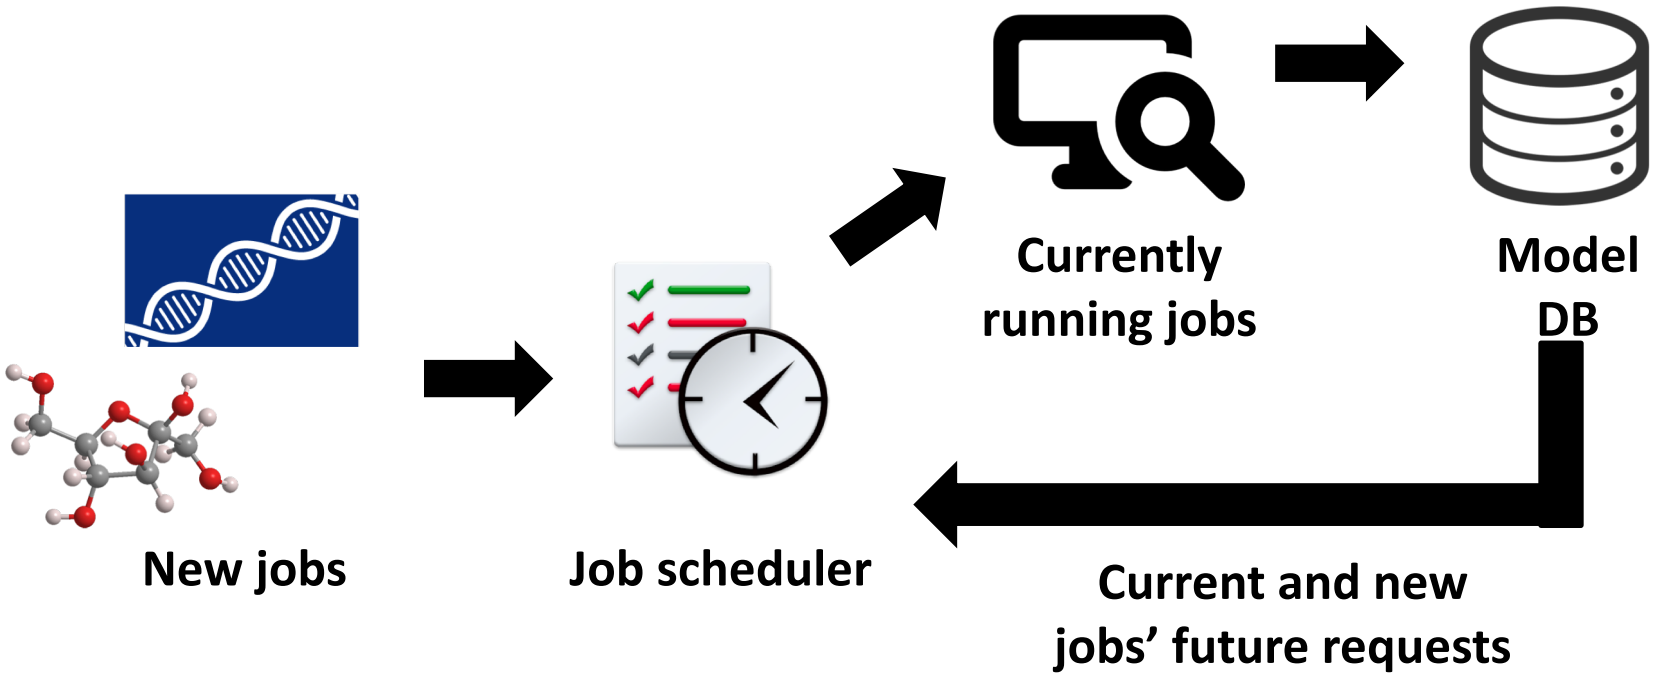

Preliminary Results

- Built a **Lustre Simulator** on NS3.

- Results from **time-series modeling** show an accuracy of 95% in predicting job write bursts.

Next Steps

- Modify the scheduler to reduce I/O contention.
- Measure the I/O performance of the jobs as well as the overall performance of the system.

Thank You! Q & A

akpaul@vt.edu

http://research.cs.vt.edu/dssl/