

A Software Framework for the Orchestration of Materials Science Workflows at Large Scale

Naweiluo Zhou, Giorgio Scorzelli, Jakob Luettgau, Rahul Reddy Kancharla, Joshua Kane, Robert Wheeler, Brendan Croom, Pania Newell, Valerio Pascucci, Michela Taufer

The University of Tennessee, Knoxville, TN, USA
The University of Utah, Salt Lake City, Utah, USA,
The Johns Hopkins University, Baltimore, Maryland, USA,
Idaho National Laboratory, Idaho Falls, USA,
MicroTesting Solutions LLC, Hilliard, Ohio, USA





Acknowledgment:

NSF: 2138811 (NSDF) and 2028923 (SOMOSPIE); IBM; XSEDE: TG-CIS210128; Chameleon: CHI-210923



Experimental Data Acquisition/Processing from Brookhaven National Laboratory



Reconstruction


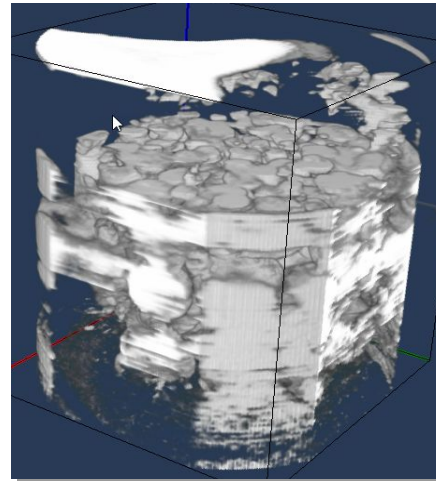




Segmentation



Motivation:

1. Support the work of materials scientists:
 - Design the materials science workflow
 - Parallelise computation
 - Scale the workflow for large-scale resources
 - Provide easy ways to execute the workflow
2. Provide the software framework to scale scientific workflows beyond materials science



Material Science Workflow

Challenges:

- Large storage:
Input size: 6GB x n, output size: 171GB x n (in TB)
- Intensive computation: several days




Fig. The six stages of the workflow represented in a linear way.

Material Science Workflow

Similarities of many scientific workflows:

- Intensive computation
- Large storage: big data
- Remote big-data retrieval: streaming services




Fig. The six stages of the workflow represented in a linear way.

Contribution

Provide a framework to support execution of scientific workflows beyond materials science :

- High-level workflow composition
- Job scheduling and resource management for heterogeneous resources:
GPU and CPU
- Scalable big-data storage
- Bring data close to computation



Material Science Workflow-- the framework

User interface



➤ Jupyter Notebook:
Data analysis

➤ Ansible:

- Perform large-scale deployment
- Workflow definition
- Environment provision

Workflow management



➤ Dask:
Scheduling and resource management

Resource management



➤ Kubernetes:
● Streaming service orchestration
● Fault tolerance
● Auto-scaling

CPU

GPU

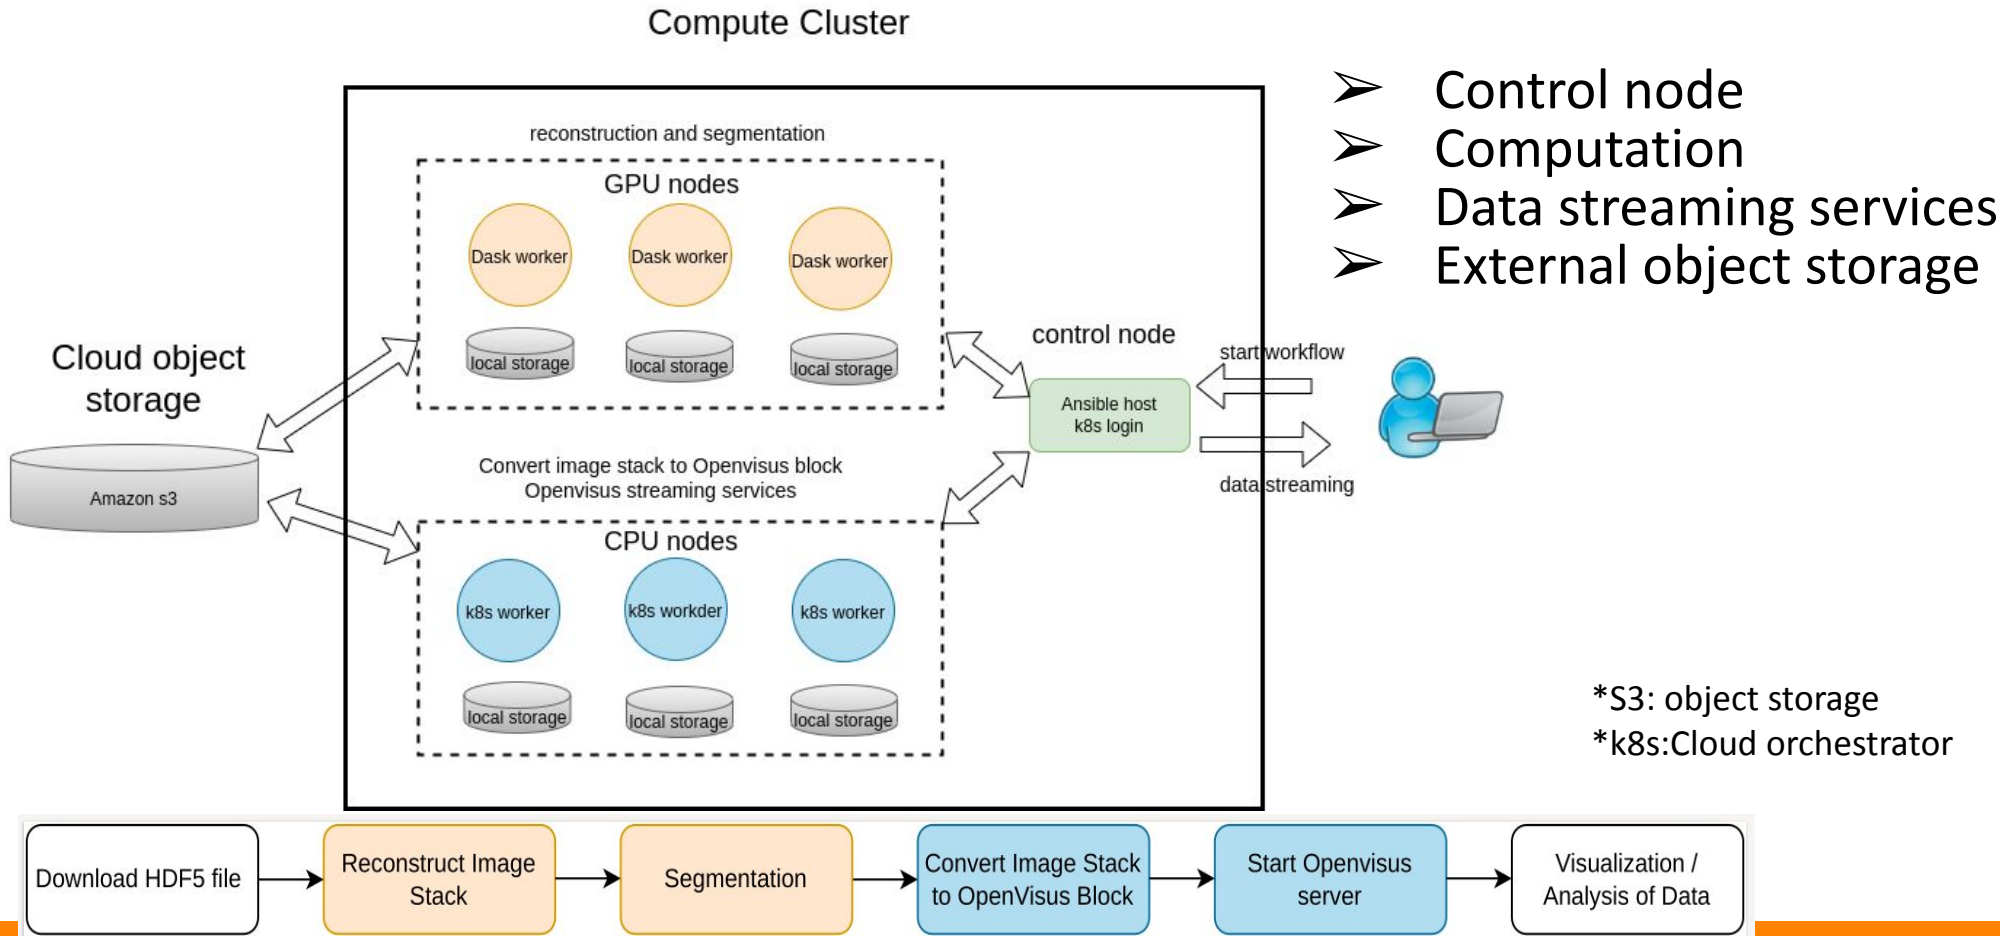
S3

PFS

*S3: object storage


*PFS: parallel filesystem

Material Science Workflow-- the architecture of the cluster



Preliminary Results

Recon and Seg stages with 16 hdf5 as inputs



GPU per node	memory capacity	GPU vendor	local disk capacity
1	23.4GB each	Quadro RTX 6000	210GB

Configurations of 16 GPU nodes

Conclusion Remarks and Future Work

Summary:

- A hierarchy framework to perform resource management & orchestration
- High-level workflow composition
- Orchestration of Materials science workflow
- Large-scale data retrieve via orchestrated streaming services
- Preliminary results

Future work

- Portability supports: containerisation of the workflow
- Automatic container building and deployment