Enabling Scalability in the Cloud for Scientific Workflows: An Earth Science Use Case

Paula Olaya*, Camila Roa*, Jakob Luettgau*, Ricardo Llamas↟, Rodrigo Vargas↟,
Sophia Wen‡, I-Hsin Chung‡, Seetharami Seelam‡, Yoonho Park‡, Jay Lofstead†, and Michela Taufer*
*University of Tennessee, Knoxville,↟University of Delaware, ‡IBM, †Sandia National Laboratory

Data Scalability in Scientific Workflows

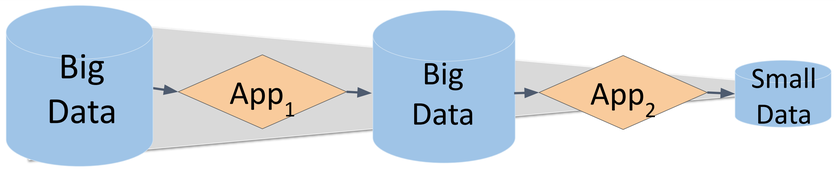

Scientific workflows (AI/ML) hide the complexity of large intermediate data and their overall execution can be affected by the I/O bandwidth of the underlying infrastructure

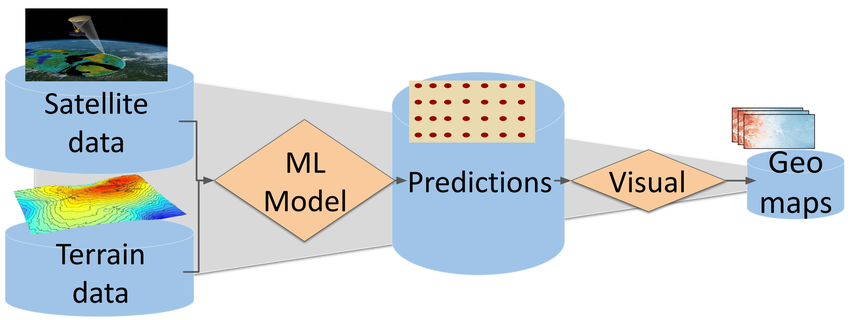

Exemplary ML-based earth science workflow (SOMOSPIE) that uses ML models to predict soil moisture data from low-resolution satellite data to high-resolution values

Data scalability is crucial in obtaining better models and predictions for scientific workflows (AI/ML)

Resolution Scalability

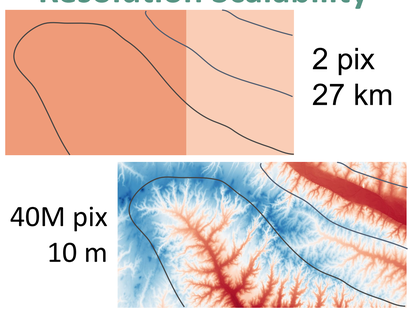

Data can scale as we move from low to high **resolutions** in a given region (scale-up)

Region Scalability

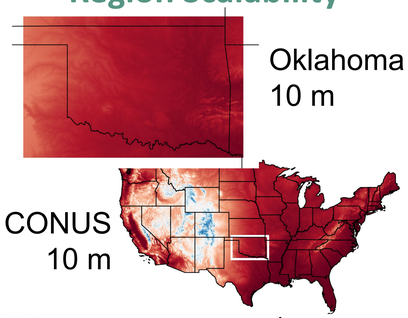

Data grows as we cover a larger **region** with the same resolution (scale-out)

Scientists need an infrastructure that **efficiently writes and reads large intermediate data** and **automatically scales** their scientific workflows' execution

Cloud-based HPC Infrastructure

We leverage cloud technology to integrate scientific workflows in cloud-based HPC services (LSF and Kubernetes) using Cloud Object Storage, enabling better I/O and data scalability

1. We **tune the cloud infrastructure's I/O parameters** at a single instance level to obtain better throughput performance
2. We **map our cloud infrastructure** to the parallel data nature of our scientific workflow granting I/O scalability

Cloud Object Storage (COS)
Data stored as objects in buckets

S3FS
Map COS into POSIX namespaces

LSF
HPC as a service on the cloud

Kubernetes
Open source
Cloud-native service

We tune the advanced **S3FS's I/O parameters (parallel count, multisize part, and caching)**

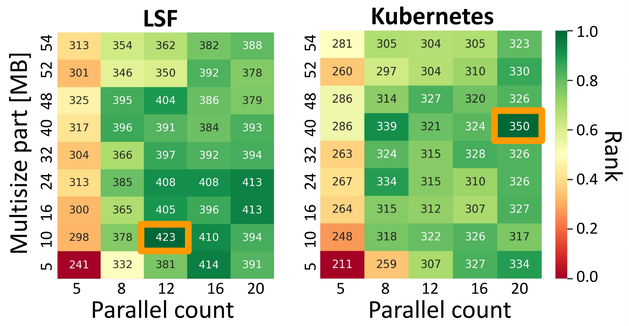

We mount the three COS buckets using the S3FS tuned parameters and **map each data partition (tiles) into an independent VM instance which reads and writes concurrently to the bucket**

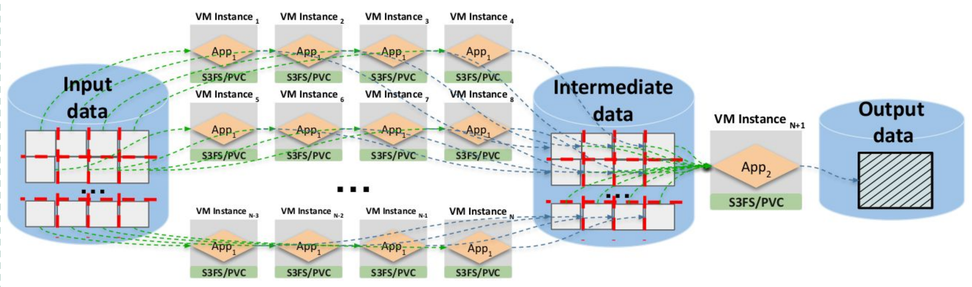

I/O Scalability on the Cloud

Weak scaling: We measure the I/O performance when **writing 685 MB** and **reading 1.2 GB** of data per VM instance as we increase from 8 to 225 concurrent VM instances reading and writing to two COS buckets

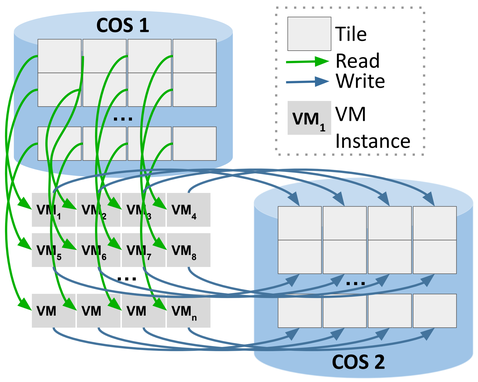

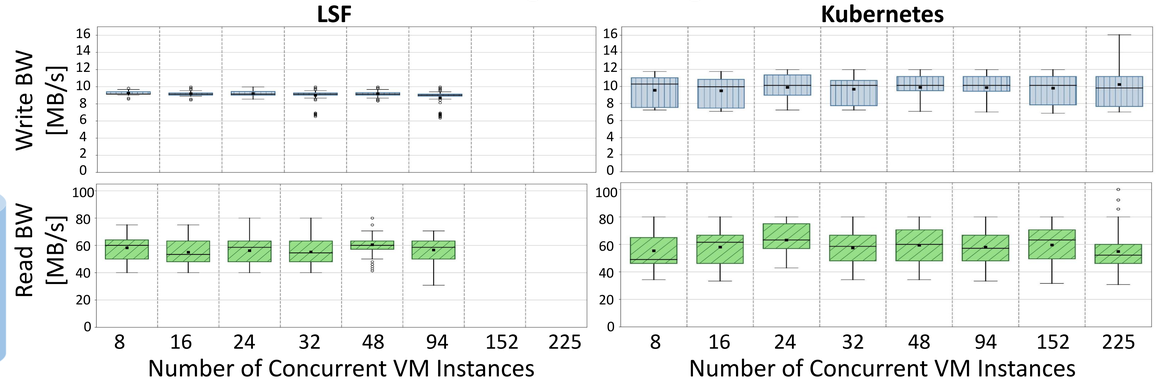

We observe **no I/O performance degradation in the object storage as we increase the number of VM instances** of writing and reading in parallel for LSF and Kubernetes (K8s)

We reach an accumulated write and read bandwidth
- LSF: **864.8 MB/s write BW** and **5.6 GB/s read BW** with 94 VMs
- K8s: **2.4 GB/s write BW** and **11.2 GB/s read BW** with 225 VMs

The authors acknowledge the support of IBM through a Shared University Research Award; Sandia National Laboratories; the National Science Foundation through the awards 1841758, 1941443, 2028923, 2103845, 2103836, and 2138811; and the XSEDE program through the NSF grant 1548562

Check more of our work