Scaling Big Data and Computations on HPC Clusters

Mike Mikailov, Fu-Jyh Luo, Stuart Barkley, Dillip Emmanuel, Xian Wu, Trinity Cheng, Tsach Mackey, Kenny Cha, Nicholas Petrick
Center for Devices and Radiological Health (CDRH), Office of Science and Engineering Laboratories (OSEL), Division of Imaging, Diagnostics and Software Reliability (DIDSR)

Background

Exponential growth of data As part of FDA's data modernization [1], CDRH HPC clusters support a wide array of applications with exponentially growing data in: Bioinformatics analysis; Artificial intelligence/machine learning; Genomics; Next-generation sequence analysis and alignment; Modeling and simulation (M&S); Statistical analysis and more.

Data scaling **However**, even powerful HPC clusters could be overwhelmed if applications do not utilize adequate scaling techniques for taking advantage of the massively parallel HPC resources [2]. Our data partitioning and scaling techniques overcome the challenges for the application areas.

Exponential growth of computations Processing exponentially growing data along with large-scale modeling and simulation in turn lead to exponentially growing computations. Evaluating the immunogenicity claims of therapeutic protein products via MCMC simulation would require 28 years of computations [3] per product.

Computations scaling **However**, traditional software parallelization techniques are not adequate even on powerful HPC clusters [3]. Our computation partitioning and scaling techniques overcome the challenges for the application areas.

Materials & Methods

Bioinformatics For applications such as BLAST large query and reference datasets are partitioned into M and N subsets, which are combined into MxN unique pairs and processed in parallel using MxN tasks. A workflow is created to automate all the steps.

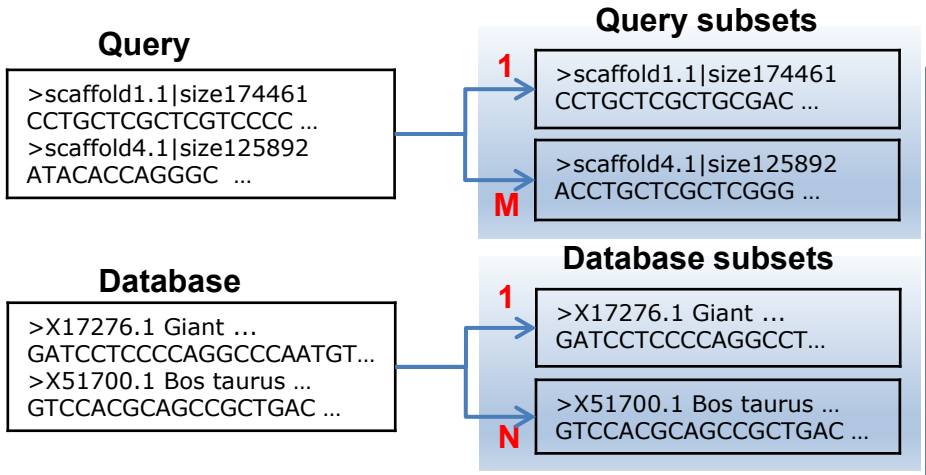

Partition and Scale

Quickly partition and **scale** your applications!

Join the discussions
9/22/21 12PM-1PM ET
Contact:
Mike.Mikailov@fda.hhs.gov

Digital Pathology Deep Learning AI Large-size gigapixel whole-slide images (WSI) are partitioned and grouped in hierarchical data format (HDF5) files, designed for concurrent access, and processed in parallel.

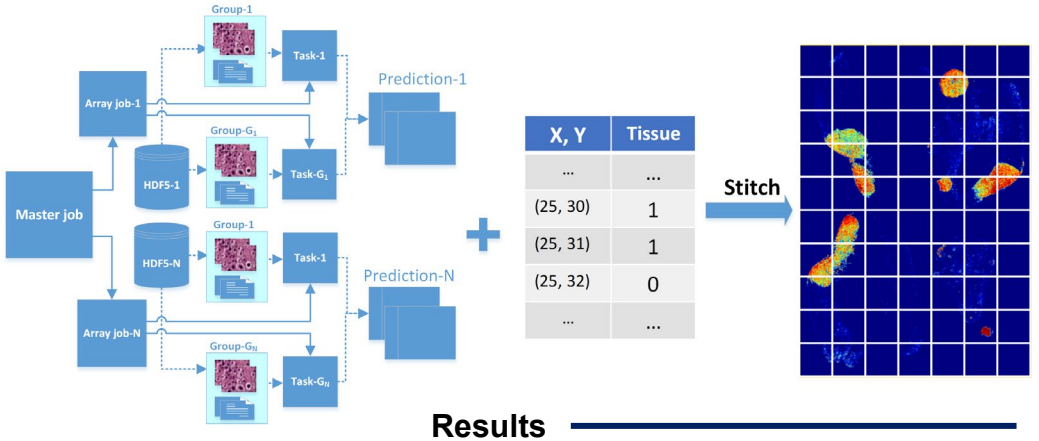

Modeling and Simulation A comprehensive approach for scaling M&S applications on HPC clusters has been developed and is called "computation segmentation.[3]" This enabled drops in run times from periods such as 28 years to days. The computation segmentation is based on the built-in array job facility.

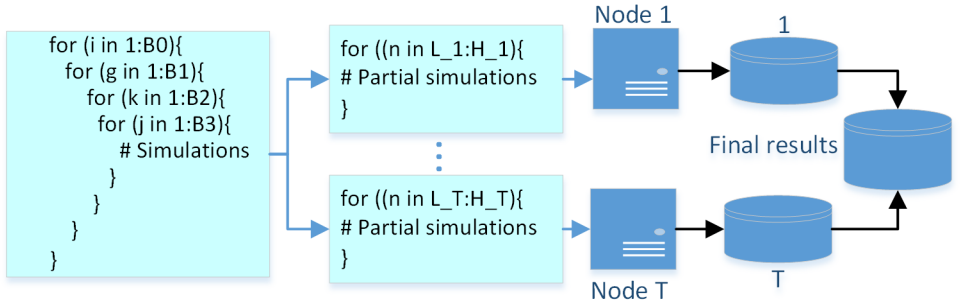

Results

Applying BLAST to one of the FDA Bioinformatics projects reduced run time from 27 days to a single day using 4,104 tasks (where M=152, N=27 subset size 100 MB), each task taking less than 7 minutes to complete.

CAMELYON datasets of 399 WSIs (>700 GB) were partitioned into 27,280 groups of ~230 MB in 399 HDF5 files for testing a DLNN. With 27,280 tasks this might take upwards of 18 years on a single CPU core, 30 days on a single GPU or less than 45 hours when implemented in parallel on the HPC.

Conclusions

The scaling techniques presented here are already in use by FDA scientists.

The techniques enable reduction of the data subset processed by each job task to a size that fits into the memory of the computing nodes where computations are performed. This avoids expensive I/O for swapping and produces excellent results, enabling substantial drops in run times.

The described methods use only open-source code with no additional hardware cost.

References

1. Data Modernization Action Plan, Data Modernization Action Plan (fda.gov)
2. Mikailov, M., Luo, F., Barkley, S. *et al.* Scaling bioinformatics applications on HPC. *BMC Bioinformatics* **18,** 501 (2017). https://doi.org/10.1186/s12859-017-1902-7
3. Mikailov, M., Qiu, J., Luo, F.-J., Whitney, S., & Petrick, N. (2020). Scaling modeling and simulation on high-performance computing clusters. *SIMULATION*, 96(2), 221–232. https://doi.org/10.1177/0037549719878249