TOWARDS I/O MONITORING AT SCALE

Designing a self-tuning I/O environment in HPC

ADMIRE

Adaptive Multi-tier Intelligent Data Manager for Exascale is a EuroHPC project aiming at addressing upcoming I/O challenges in Exascale systems. It gathers thirteen European partners to pioneer new ways of doing HPC I/Os.

Objective

The main objective of the ADMIRE project is the creation of an active I/O stack that dynamically adjusts computation and storage requirements through intelligent global coordination, elasticity of computation and I/O, and the scheduling of storage resources along all levels of the storage hierarchy, while offering quality-of-service (QoS), energy efficiency, and resilience for accessing extremely large data sets in very heterogeneous computing and storage environments.

Acknowledgement

This project has received funding from the European Union's Horizon 2020 under Grant Agreement number: 956748-ADMIRE-H2020-JTI-EuroHPC-2019-1.

ADMIRE
malleable data solutions for HPC

Applications

In the ADMIRE project, the dynamic allocation of resources to jobs is crucial. This is particularly true for applications that involve training Deep Learning (DL) models on large datasets. DL has unleashed advances in applications from various disciplines, such as physics or medicine, reaching unprecedented performances compared to traditional Machine Learning.

Remote sensing

One such application involves training a DL model on a large remote sensing (RS) dataset called BigEarthNet, with the goal of land cover classification. Given the size of the dataset and the deep DL model involved, many resources are required for training and inference.

EuroHPC Joint Undertaking

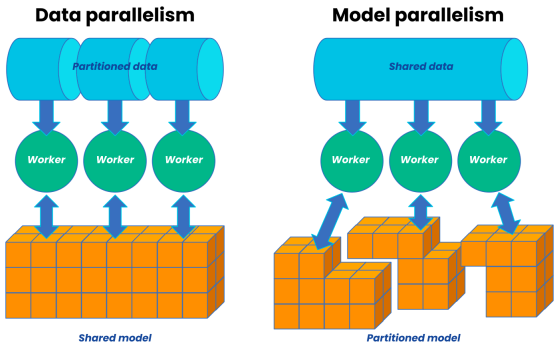

Fig I. RGB images acquired by the Sentinel-2 satellite mission.

Life sciences

Another application utilizing DL models is the Life-Sciences Application trained on sequences of mice cerebellum microscope images. The goal is to detect the early stages of polyglutamine diseases. Given that every mouse cerebellum consists of hundreds of sequences of high-resolution images, DL needs to be involved in processing such a vast amount of imaging data.

Distributed deep learning

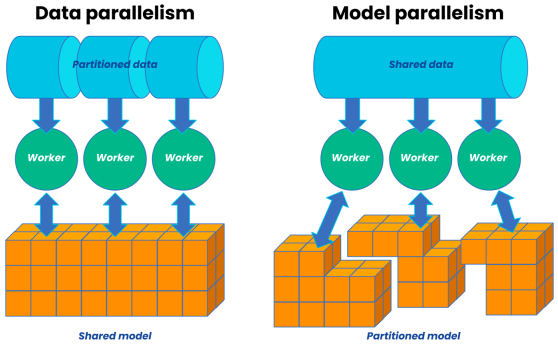

Fig II. Data and model parallelism.

This is where Distributed Deep Learning (DDL) comes into play, as it can significantly speed up the process and reduce the time needed for classification. However, running DL applications in a dynamic environment can be challenging. There are two prominent families of approaches to scale models on multiple workers in a parallel fashion: data parallelism and model parallelism. While the data is chunked in the former, and a subset is fed into a replica of the DL model on each device, the latter splits the DL model among workers.

Elastic Horovod

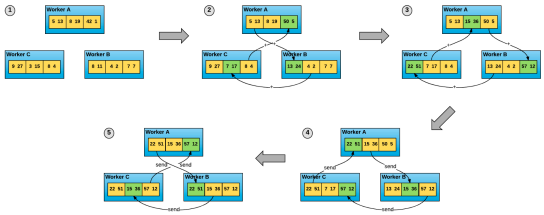

Fig III. Illustration of Ring-AllReduce algorithm.

Horovod is a framework for data parallelism that can work on top of existing DL libraries. It adopts an efficient distributed algorithm to exchange the local gradients among the workers, aggregates them into global gradients, and updates the weights of the DL model. The ADMIRE project plans to adopt Elastic Horovod to address this challenge. This software allows workers at runtime without requiring a restart or resuming from checkpoints saved to join or leave the Horovod job without interrupting the training process.

What's next...

By adopting Elastic training for Horovod, the ADMIRE project has found a means to provide the two DL applications with a strategy to respond to the dynamic requirements of the project. In this way, the applications can be continually scaled up or down depending on the needed resources, ensuring efficient resource utilization and optimal performance.

Join Us !

Feel feee to contact us for further information on the project !

admire_eurohpc@uc3m.es

Admire EURO-HPC Project

@admire_eurohpc

Captured Metrics

ADMIRE's monitoring infrastructure is capable of collecting a wide range of metrics today and we constantly keep improving the measurement layer to provide a detailed view of the system.

Message Passing Interface

We capture all MPI calls and their duration and respective sizes in an always-on fashion. It means in ADMIRE's configuration you can see MPI calls rate and throughput in real-time on the whole system.

I/O Monitoring

Thanks to a close collaboration with the I/O backend and dedicated instrumentation layers, ADMIRE can non-intrusively characterize I/O traffic. To do so, we instrumented both common interfaces (POSIX) and specialized (Ad-Hoc, MPI) I/O interfaces to track both the bandwidth and origin of the data traffic.

System Monitoring

System monitoring is done by the LIMITLESS daemons which precisely track node-level data with optimized data-cropping algorithms to only account for representative state changes. Thanks to this component all parameters relative to nodes' health and global resource usage are tracked in real-time thanks to the reduction Tree-Based Overlay Network (TBON).