

The research of High-Performance Computing infrastructure for Artificial Intelligence and Big Data

Since the difference of software stack between traditional HPC applications, big data applications and artificial intelligence applications, infrastructure for each usually adopted completely different method and system to build and manage.

To meet the rapidly growing demand for computing and storage resources raised by big data and artificial intelligence application, and take fully advantage of infrastructure's computing and storage ability, we build our infrastructure based on traditional HPC technology and infrastructure, and make improvements on different layers of the infrastructure to meet the requirements of big data and artificial intelligence application, realizing a high-performance computing infrastructure with high efficiency and user-friendly.

For storage layer, besides the traditional high-performance parallel file system, we specifically equipped high-volume SSD storage for each computing node, making the infrastructure compatible with distributed storage file system like HDFS, and meet the data-locality requirements. Based on that, the infrastructure can handle big data applications built on top of Hadoop, Spark and other related framework. To manage the data of different I/O pattern efficiently, we build multiple data management interface to upload and write data, satisfying different application mode.

To adapt to the computing pattern of artificial application, like deep learning application, each computing node equipped with 8 Nvidia Tesla P100 GPUs, and inter-connected by 56Gbps InfiniBand network. We use high efficient and powerful scheduler system respectively based on LSF and Apache Yarn. For both system, leveraging the container and other resource isolation and technologies, the infrastructure implement CPU affinity, GPU affinity and other features to improve computing and communication efficiency. Multiple scheduling policy and queue management policy are supported by the infrastructure, so application with different resource requirement characteristics can be managed reasonably and fully utilize the computing resource. Based on the container technology, users can build their own software stack image, manage their software and framework by a more easy-to-use way than the traditional HPC user environment.

Specifically, we build interface between common deep learning framework and the scheduler system, like tensorflow, caffe, pytorch, mxnet. For distributed deep-learning applications, user need not to explicitly allocate resource detailly in their code, the scheduler will handle most of the work. Furthermore, we pay attention to applying MPI and NCCL technology in common deep-learning frameworks on the infrastructure, aiming to make a full use of the InfiniBand network and improving the communication performance.

To make the infrastructure more easy to use by artificial intelligence developers, different kinds of assistance service are being built, including visualization tools and training monitoring tools for the common deep learning frameworks, that uses can view their model structure and training process from web.

The infrastructure has been adopted by users from different fields. An atmospheric Science application using deep-learning networks to predict the weather, the test results shows that the infrastructure make a significant performance improvement than their own deep-learning clusters. Other applications include a research of applying reinforcement learning to poker games, a research of building algorithmic trader based on reinforcement learning leveraging large scale of financial trading data, and application of face recognition.

In the future, the infrastructure will be connected with other commercial computing service platform through grid technology, like Kingsoft's Cloud Service. Furthermore, the infrastructure will connected to the China Nation Grid, to offer service for more academic and enterprise users.

Primary author: Dr WANG, Jue (Computer Network Information Center, Chinese Academy of Sciences)

Co-authors: Ms LI, Chen (Computer Network Information Center, Chinese Academy of Sciences); Dr LIU, Fang (Computer Network Information Center, Chinese Academy of Sciences); Mr SUN, Kun (Computer Network Information Center, Chinese Academy of Sciences); Mr HU, Tengting (Computer Network Information Center, Chinese Academy of Sciences); Prof. WANG, Yangang (Computer Network Information Center, Chinese Academy of Sciences)

of Sciences); Mr SUN, Yongze (Computer Network Information Center, Chinese Academy of Sciences); Prof. LU, Zhonghua (Computer Network Information Center, Chinese Academy of Sciences)

Presenters: Dr WANG, Jue (Computer Network Information Center, Chinese Academy of Sciences); Prof. WANG, Yangang (Computer Network Information Center, Chinese Academy of Sciences); Prof. LU, Zhonghua (Computer Network Information Center, Chinese Academy of Sciences)

Track Classification: Infrastructure Clouds and Virtualisation