PortHadoop-R: Support the Merging of HPC and Cloud

Xian-He Sun

Illinois Institute of Technology, Chicago IL, USA, sun@iit.edu

ABSTRACT

The merging of HPC and big data analytics is inevitable.

- High Performance Computing (HPC) is becoming data intensive.
- Big data applications are requiring more and more computing power.
- Two ecosystems are designed for different applications and with different design principles.

Two ecosystems will co-exist.

- By the CAP theory, neither cannot have all the merits of the other.
- The best we can have is a merged system which can provide the functionality and merits of both ecosystems.

We propose PortHadoop-R to support the merging at the file level.

- Allow reading data directly from PFS to the memory of Hadoop nodes.
- Integrate the data transfer with R data analysis and visualization.
- Optimized to utilize the merits of PFS and MapReduce to achieve concurrent data transfer and latency hiding.
- Tested on NASA climate modeling applications.
- Experimental results show 15x speedup.
- Significantly faster than MPI clusters on processing climate data.

BACKGROUND

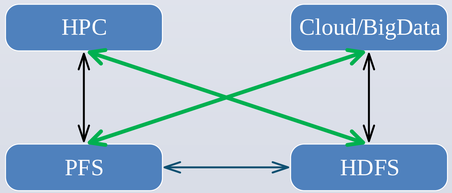

Trend of merging of HPC and Cloud

- Two camps of computing environment
  - MPI for HPC (e.g. large scale simulation)
  - MapReduce for cloud/BigData (e.g. visualization and data analysis)
- Native data access exists
  - HPC applications read from/write to PFS (POSIX I/O and MPI-IO)
  - Cloud applications read from/write to HDFS (HDFS API)
- Data dependency in-between
  - Cloud/BigData applications analyze data generated by simulations
  - Analysis results guide simulation runs
  - Use explicit **serial** copy
  - Redundant data store in HDFS

USE CASE

Typical scientific research process

- Iterative process
  - Simulation generates large data
  - Analyzed and visualized iteratively
  - Analysis results help steering simulations
- Simulation
  - MPI-based
  - HPC platform
  - Use PFS as data storage
- Visualization
  - Coarse visual analysis
  - Quick identification
- Data analysis
  - Statistical analysis
  - SQL query
  - Use HDFS as data storage

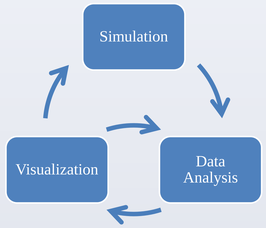

DESIGN

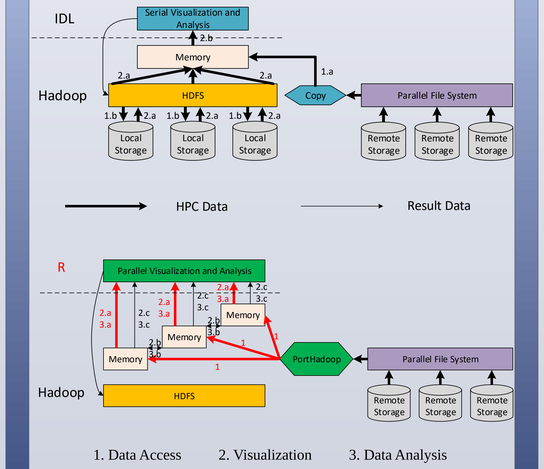

Figure 1. Architecture comparison between the conventional solution and PortHadoop-R

EVALUATIONS

- Testbed
  - Chameleon at TACC
  - Dual Intel Xeon E5-2670 V3 (2.3 GHz, 24 cores and 48 threads in total)
  - 128 GB memory, 10 Gbps Ethernet
  - 4, 8 and 16 Hadoop nodes, 8 OrangeFS nodes
- Software
  - CentOS 7.3
  - Java 1.7.0_79
  - PortHadoop-R (modified from Hadoop-2.5.0-CDH5.3.3)
- Workload

Name	Image Plotting	Animation	Analysis
Vis-only	Yes	Yes	No
Img-only	Yes	No	No
Anlys	Yes	Yes	Yes

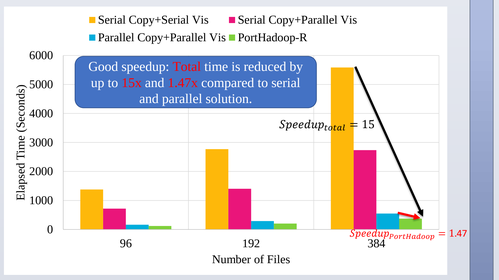

Performance comparison with other solutions for Vis-only

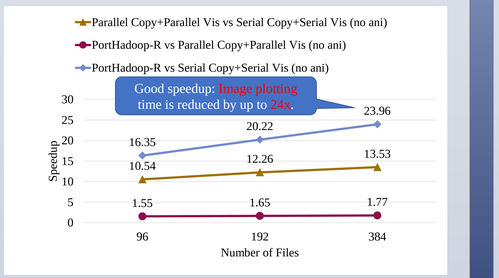

Speedup compared with other solutions for Img-only

EVALUATIONS

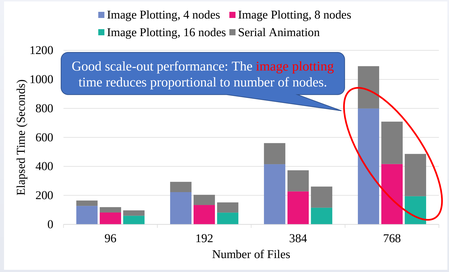

Scale-out performance for Vis-only

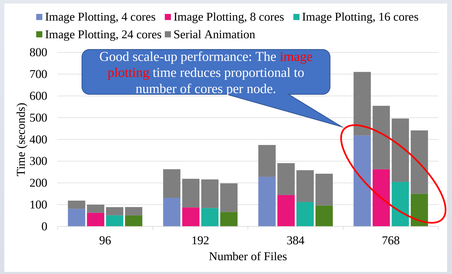

Scale-up performance for Vis-only

Data analysis performance for Anlys

CONTRIBUTIONS AND HIGHLIGHTS

- Effective and efficient analysis of HPC data.
- Read data from PFS **without copying**. Data are fetched into memory of Hadoop node dynamically on demand.
- Take advantage of **R** and **Hadoop** to visualize and analyze. R libraries are used for visualization and analysis capability. Hadoop is used to parallelize the process.
- Utilize **parallelism** of both PFS and Hadoop. Data are read out from PFS **concurrently** and visualization and analysis are carried out **in parallel** in Hadoop.
- Evaluation on NASA cloud-resolve model (WRF) simulation data. Data is generated from a 48-hour real simulation run. Up to 23GB data is analyzed.
- Great performance improvement. Up to **15x** speedup compared with the conventional solution.

CONCLUSIONS

- There is a need to merge the currently separated HPC ecosystem and cloud computing ecosystem.
- PortHadoop-R supports this merging in data access level.
- PortHadoop-R utilizes merits of both systems to speedup the performance.
- PortHadoop-R reaches a 15x speedup on one NASA climate modeling application in comparison with the currently used approach.

FUTURE WORK

- Support more frameworks built on top of Hadoop.
- Support more scientific data format.

ACKNOWLEDGEMENT

Grant Info: CSR: Small: Empower Data-Intensive Computing: the integrated data management approach," NSF, CNS-1526887, September 1, 2015 – August 31, 2018. I would like to thank Chameleon for providing testbed for testing this work.

REFERENCES

1. "The Weather Research&Forecasting Model Website." [Online]. Available: http://www.wrf-model.org/index.php.
2. S. Zhou, X. Yang, X. Li, T. Matsui, S. Liu, X. H. Sun, and W. Tao, "A Hadoop-based visualization and diagnosis framework for earth science data," Proc. - 2015 IEEE Int. Conf. Big Data, IEEE Big Data 2015, pp. 1972–1977, 2015.
3. X. Yang, N. Liu, B. Feng, X.-H. Sun, and S. Zhou, "PortHadoop: Support direct HPC data processing in Hadoop," in 2015 IEEE International Conference on Big Data (Big Data), 2015, pp. 223–232.
4. X. Yang, S. Liu, K. Feng, S. Zhou, and X.-H. Sun, "Visualization and Adaptive Subsetting of Earth Science Data in HDFS: A Novel Data Analysis Strategy with Hadoop and Spark," 2016 IEEE Int. Conf. Big Data Cloud Comput. (BDCloud), Soc. Comput. Netw. (SocialCom), Sustain. Comput. Commun., pp. 89–96, Oct. 2016.
5. "ImageMagick." [Online]. Available: https://www.imagemagick.org/script/index.php.
6. "RHadoop." [Online]. Available: https://github.com/RevolutionAnalytics/RHadoop.
7. "Chameleon." [Online]. Available: https://www.chameleoncloud.org/.