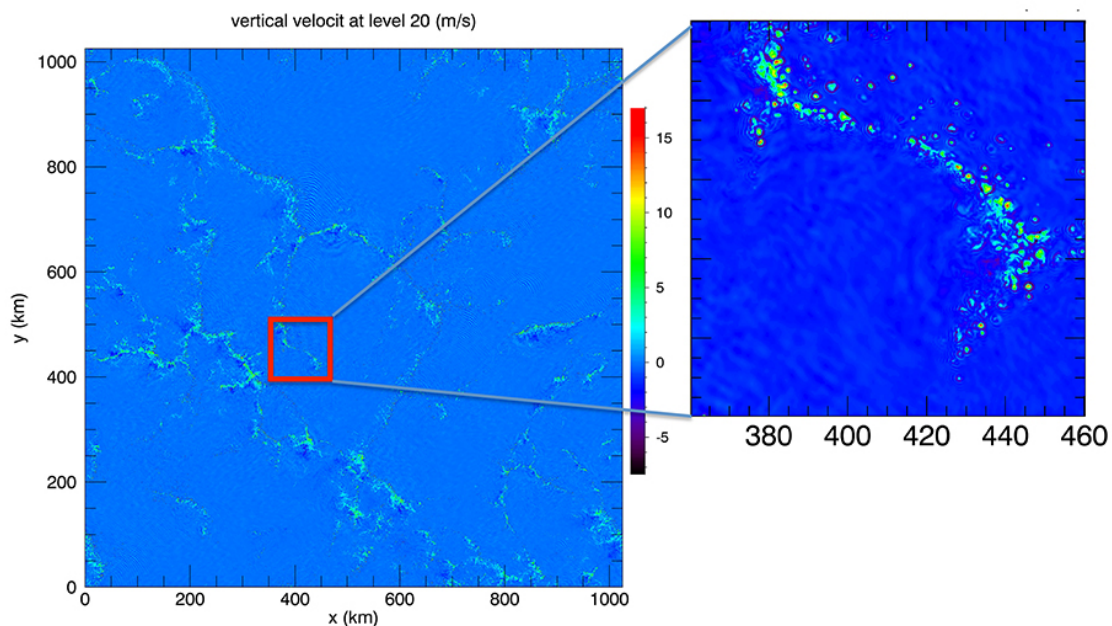


Reaching for the stormy cloud with Chameleon

May 10 2017, by Nsf-Funded Chameleon Cloud Testbed Speeds
Development Of Porthadoop Reader For Nasa Cloud Library



The Chameleon cloud testbed has sped development of PortHadoop-R, a portable Hadoop reader for parallel file systems which can integrate data transfer with data analysis. The NASA cloud library uses PortHadoop-R as part of its MapReduce analytics. Goddard Cumulus Ensemble simulation shown here utilizing PortHadoop-R to enable real-time dynamic visualization cloud vertical velocities. Credit: Wei-Kuo Tao, NASA.

Some scientists dream about big data. The dream bridges two divided realms. One realm holds lofty peaks of number-crunching scientific computation. Endless waves of big data analysis line the other realm. A deep chasm separates the two. Discoveries await those who cross these estranged lands.

Unfortunately, data cannot move seamlessly between Hadoop (HDFS) and parallel file systems (PFS). Scientists who want to take advantage of the [big data analytics](#) available on Hadoop must copy data from parallel file systems. That can slow workflows to a crawl, especially those with terabytes of data.

Computer Scientists working in Xian-He Sun's group are bridging the file system gap with a cross-platform Hadoop reader called PortHadoop, short for portable Hadoop. "PortHadoop, the system we developed, moves the data directly from the parallel file system to Hadoop's memory instead of copying from disk to disk," said Xian-He Sun, Distinguished Professor of Computer Science at the Illinois Institute of Technology. Sun's PortHadoop research was funded by the National Science Foundation and the NASA Advanced Information Systems Technology Program (AIST).

The concept of 'virtual blocks' helps bridge the two systems by mapping data from parallel file systems directly into Hadoop memory, creating a virtual HDFS environment. These 'virtual blocks' reside in the centralized namespace in HDFS NameNode. The HDFS MapReduce application cannot see the 'virtual blocks'; a map task triggers the MPI file read procedure and fetches the data from the remote PFS before its Mapper function processes its data. In other words, a dexterous slight-of-hand from PortHadoop tricks the HDFS to skip the costly I/O operations and data replications it usually expects.

Sun said he sees PortHadoop as the consequence of the strong desire for

scientists to merge [high performance computing](#) with cloud computing, which companies such as Facebook and Amazon use to 'divide and conquer' data-intensive MapReduce tasks among its sea of servers. "Traditional scientific computing is merging with big data analytics," Sun said. "It creates a bigger class of scientific computing that is badly needed to solve today's problems."

PortHadoop was extended to PortHadoop-R to seamlessly link cross-platform data transfer with [data analysis](#) and virtualization. Sun and colleagues developed PortHadoop-R specifically with the needs of NASA's high-resolution cloud and regional scale modeling applications in mind. High performance computing has served NASA well for their simulations, which crunch data through various climate models. Sun said the data generated from models combined with observational data are unmanageably huge and have to be analyzed and also visualized to more fully understand chaotic phenomena like hurricanes and hail storms in a timely fashion.

PortHadoop faced a major problem in preparation to work with NASA applications. NASA's production environment doesn't allow any testing and development on its live data.

PortHadoop developers overcame the problem with the Chameleon cloud testbed system, funded by the National Science Foundation (NSF). Chameleon is a large-scale, reconfigurable environment for cloud computing research co-located at the Texas Advanced Computing Center of the University of Texas at Austin and also at the the Computation Institute of the University of Chicago. Chameleon allows researchers bare-metal access, i.e., allows them to fully reconfigure the environment on its nodes including support for operations such as customizing the operating system kernel and console access.

What's more, the Chameleon system of ~15,000 cores with Infiniband

interconnect and 5 petabytes of storage adeptly blends in a variety of heterogeneous architectures, such as low-power processors, graphical processing units, and field-programmable gate arrays.

"Chameleon helped us in different ways," Sun said. "First, it made it possible for us to create two different environments," each on a separate computer cluster of the bare metal system to mimic the NASA environment. "We are really happy that we were able to use Chameleon. The system helped us a great deal in our development," Sun added.

Sun and colleagues installed some nodes with the traditional MPI, and on the other cluster they installed MapReduce. "Then we ran programs on these two different clusters and did the data integration, cross-platform data access, data analysis and visualization, all on Chameleon," Sun added. "Chameleon provides all functionality and scale of this computing facility, as well as the option of creating different programming environments on hardware resources with both HPC and data-intensive characteristics for our integration research."

Sun and colleagues put PortHadoop-R to the test by using it in the NASA Cloud library project for empowering data management, diagnostics, and visualization of Cloud-Resolving Models (CRMs) of climate modeling. The NASA Cloud library has big data from satellites and human observations, with more than 70,000 datasets downloaded since April 2010 by 155 distinct users. The data are used for real-time forecasts of hurricane and other natural disasters; and for long-term climate prediction.

"The ultimate goal is to generate a core cloud library that is dynamic and interactive with the user," said Wei-Kuo Tao, a senior research meteorologist at the NASA Goddard Space Flight Center. He leads the Goddard Mesoscale Dynamic and Modeling Group and is the principal investigator of the NASA AIST program. Tao and colleagues combine

large-scale CRM simulation data at real time for data analysis and visualization. "The idea of the dynamic visualization and analysis with the Hadoop reader is that you don't have to copy the data," Tao said. "You can produce the visualization at the same time."

Xian-He Sun spoke of the work that bridged the two types of storage systems. "We tested our PortHadoop-R strategy on Chameleon, and later confirmed these tests on NASA machines in practice. The result is fantastic, and beyond our expectation. We expected a 2-fold speedup. The result is a 15x speedup. The reason is that PortHadoop-R not only reduced one round disk copy but also utilized the concurrency of parallel file systems and Hadoop systems in a level which for general users would be difficult to achieve. In other words, PortHadoop-R has integrated the MPI and Hadoop systems," Sun said.

"Chameleon was really helpful to provide the flexible environment so we can install, or simulate different environments. We have an HPC environment. We have a cloud environment. And we have them together and test them together. Chameleon in this sense provides us the ability to scale our computing resource and the privilege to control, optimize, and install our custom designed software environment to develop our software systems," Sun said.

Building on the success with PortHadoop-R on real applications at NASA, Sun added that "the next step is to make PortHadoop-R more user-friendly. And, we would like to expand PortHadoop-R to support different application interfaces, so different users can use it easily."

"In the long run, we would like to extend the merge at the OS level and at the user level. So, there are still a lot of things we need to do to support the seamless integration of the high-performance computing and [big data](#) analytics," Sun said.

More information: — [DOI: 10.1109/BigData.2015.7363759](https://doi.org/10.1109/BigData.2015.7363759)
— [DOI: 10.1109/BigData.2016.7840949](https://doi.org/10.1109/BigData.2016.7840949),
— [DOI: 10.1109/BDCLOUD-SOCIALCOM-SUSTAINCOM.2016.24](https://doi.org/10.1109/BDCLOUD-SOCIALCOM-SUSTAINCOM.2016.24)

Provided by Texas Advanced Computing Center

Citation: Reaching for the stormy cloud with Chameleon (2017, May 10) retrieved 13 March 2024 from <https://phys.org/news/2017-05-stormy-cloud-chameleon.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.