

Berkeley Lab, Intel, Cray harness power of deep learning to study the universe

September 5 2018



Example simulation of dark matter in the universe, used as input to the CosmoFlow network. CosmoFlow is the first large-scale science application to use the TensorFlow framework on a CPU-based high performance computing



platform with synchronous training. Credit: Lawrence Berkeley National Laboratory

A Big Data Center collaboration between computational scientists at Lawrence Berkeley National Laboratory's (Berkeley Lab) National Energy Research Scientific Computing Center (NERSC) and engineers at Intel and Cray has yielded another first in the quest to apply deep learning to data-intensive science: CosmoFlow, the first large-scale science application to use the TensorFlow framework on a CPU-based high performance computing platform with synchronous training. It is also the first to process three-dimensional (3-D) spatial data volumes at this scale, giving scientists an entirely new platform for gaining a deeper understanding of the universe.

Cosmological "big data" problems go beyond the simple volume of data stored on disk. Observations of the universe are necessarily finite, and the challenge that researchers face is how to extract the most information from the observations and simulations available. Compounding the issue is that cosmologists typically characterize the distribution of matter in the universe using statistical measures of the structure of matter in the form of two- or three-point functions or other reduced statistics. Methods such as deep learning that can capture all features in the distribution of matter would provide greater insight into the nature of dark energy. First to realize that deep learning could be applied to this problem were Siamak Ravanbakhsh and his colleagues, as referenced in proceedings of <u>The 33rd International Conference on</u> <u>Machine Learning</u>. However, computational bottlenecks when scaling up the network and dataset limited the scope of the problem that could be tackled.

Motivated to address these challenges, CosmoFlow was designed to be



highly scalable; to process large, 3-D cosmology datasets; and to improve deep learning training performance on modern HPC supercomputers such as the Intel processor-based Cray XC40 Cori supercomputer at NERSC. CosmoFlow is built on top of the popular TensorFlow machine learning framework and uses Python as the front end. The application leverages the Cray PE Machine Learning Plugin to achieve unprecedented scaling of the TensorFlow Deep Learning framework to more than 8,000 nodes. It also benefits from Cray's DataWarp I/O accelerator technology, which provides the I/O throughput required to reach this level of scalability.

In a technical paper to be presented at SC18 in November, the CosmoFlow team describes the application and initial experiments using dark matter N-body simulations produced using the MUSIC and pycola packages on the Cori supercomputer at NERSC. In a series of singlenode and multi-node scaling experiments, the team was able to demonstrate fully synchronous data-parallel training on 8,192 of Cori with 77% parallel efficiency and 3.5 Pflop/s sustained performance.

"Our goal was to demonstrate that TensorFlow can run at scale on multiple nodes efficiently," said Deborah Bard, a big data architect at NERSC and a co-author of the technical paper. "As far as we are aware, this is the largest ever deployment of TensorFlow on CPUs, and we think it is the largest attempt to run TensorFlow on the largest number of CPU nodes."

Early on, the CosmoFlow team laid out three primary goals for this project: science, single-node optimization and scaling. The science goal was to demonstrate that deep learning can be used on 3-D volumes to learn the physics of the universe. The team also wanted to ensure that TensorFlow ran efficiently and effectively on a single Intel Xeon Phi processor node with 3-D volumes, which are common in science but not so much in industry, where most deep learning applications deal with



2-D image data sets. And finally, ensure high efficiency and performance when scaled across 1000's of nodes on the Cori supercomputer system.

As Joe Curley, Sr. Director of the Code Modernization Organization in Intel's Data Center Group, noted, "The Big Data Center collaboration has produced amazing results in computer science through the combination of Intel technology and dedicated software optimization efforts. During the CosmoFlow project, we identified framework, kernel and communication optimization that led to more than 750x performance increase for a single node. Equally as impressive, the team solved problems that limited scaling of deep learning techniques to 128 to 256 nodes—to now allow the CosmoFlow application to scale efficiently to the 8,192 nodes of the Cori supercomputer at NERSC."

"We're excited by the results and the breakthroughs in artificial intelligence applications from this collaborative project with NERSC and Intel," said Per Nyberg, vice president of market development, artificial intelligence and cloud at Cray. "It is exciting to see the CosmoFlow team take advantage of unique Cray technology and leverage the power of the a Cray supercomputer to effectively scale deep learning models. It is a great example of what many of our customers are striving for in converging traditional modeling and simulation with new deep learning and analytics algorithms, all on a single, scalable platform."

Prabhat, Group Leader of Data & Analytics Services at NERSC, added, "From my perspective, CosmoFlow is an exemplar project for the Big Data Center collaboration. We've truly leveraged competencies from various institutions to solve a hard scientific problem and enhance our production stack, which can benefit the broader NERSC user community."



In addition to Bard and Prabhat, co-authors on the SC18 paper include Amrita Mathuriya, Lawrence Meadows, Lei Shao, Tuomas Karna, John Pennycook, Jason Sewall, Nalini Kumar and Victor Lee from Intel; Peter Mendygral, Diana Moise, Kristyn Maschhoff and Michael Ringenburg from Cray; Siyu He and Shirley Ho from the Flatiron Institute; and James Arnemann from UC Berkeley.

Provided by Lawrence Berkeley National Laboratory

Citation: Berkeley Lab, Intel, Cray harness power of deep learning to study the universe (2018, September 5) retrieved 27 April 2024 from <u>https://phys.org/news/2018-09-berkeley-lab-intel-cray-harness.html</u>

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.