

Perspectives on High-Performance Computing in a Big Data World

Geoffrey Fox
Indiana University
Bloomington, IN, USA
gcf@indiana.edu

ABSTRACT

High-Performance Computing (HPC) and Cyberinfrastructure have played a leadership role in computational science even since the start of the NSF computing centers program. Thirty years ago parallel computing was a centerpiece of computer science research. Naively Big Data surely requires HPC to be processed, and transformational Big Data technology such as Hadoop and Spark exploit parallelism to success. Nevertheless, the HPC community does not appear to be thriving as a leader in Data Science while parallel computing is no longer a centerpiece. Some reasons for this are the dominant presence of Industry in technology futures and the universal fascination with Artificial Intelligence and Machine Learning. Maybe the pendulum will swing back a bit, but I expect the "AI first" philosophy to dominate in the foreseeable future. Thus I describe a future where HPC thrives in collaboration with Industry and AI. In particular, I discuss the promise of MLforHPC (AI for systems) and HPCforML (systems for AI).

CCS Concepts/ACM Classifiers

• Computing methodologies~Artificial intelligence
• Computing methodologies~Massively parallel and high-performance simulations
• Computer systems organization~Parallel architectures

Author Keywords

HPC, Big Data, Computational Science, Data Science

BIOGRAPHY

Fox received a Ph.D. in Theoretical Physics from Cambridge University where he was Senior Wrangler. He is now a distinguished professor of Engineering, Computing, and Physics at Indiana University where he is the director of the Digital Science Center. He previously held positions at Caltech, Syracuse University, and Florida State University after being a postdoc at the Institute for Advanced Study at Princeton, Lawrence Berkeley Laboratory, and Peterhouse College Cambridge. He has supervised the Ph.D. of 72 students and published around 1300 papers (over 500 with at least ten citations) in physics and computing with an hindex of 77 and over 35000 citations. He is a Fellow of APS (Physics) and ACM (Computing) and works on the interdisciplinary interface between computing and applications. Current work is in

Biology, Pathology, Sensor Clouds and Ice-sheet Science, Image processing, Deep Learning, and Particle Physics. His architecture work is built around High-performance computing enhanced Software Defined Big Data Systems on Clouds and Clusters. The analytics focuses on scalable parallel machine learning. He is an expert on streaming data and robot-cloud interactions. He is involved in several projects to enhance the capabilities of Minority Serving Institutions. He has experience in online education and its use in MOOCs for areas like Data and Computational Science.



REFERENCES

1. Geoffrey Fox, James A. Glazier, JCS Kadupitiya, Vikram Jadhao, Minje Kim, Judy Qiu, James P. Sluka, Endre Somogyi, Madhav Marathe, Abhijin Adiga, Jiangzhuo Chen, Oliver Beckstein, and Shantenu Jha, "Learning Everywhere: Pervasive Machine Learning for Effective High-Performance Computation", HPBDC workshop at IPDPS, Rio de Janeiro, Brazil, Monday, May 20th, 2019

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

HPDC '19, June 22–29, 2019, Phoenix, AZ, USA.

© 2019 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-6670-0/19/06.

DOI: <https://doi.org/10.1145/3307681.3325410>