The 7th International Workshop on Parallel and Distributed Computing for Large Scale Machine Learning and Big Data Analytics (ParLearning 2018)

May 21, 2018

In conjunction with
The 32nd IEEE International Parallel & Distributed Processing Symposium (IPDPS 2018)
May 21-25, 2018
Vancouver, British Columbia CANADA

The ever increasing size of real-world datasets provides opportunities to find insights in new and emerging types of data and content, to make businesses more agile, and to answer questions that were previously considered beyond our reach – the promise of Big Data. However, scaling up Machine Learning, Data Mining, and reasoning algorithms from Artificial Intelligence for massive datasets remains a major technical challenge. Many algorithms for machine learning and statistical inference are notorious for their computational complexity with respect to the size of the data. Adopting parallel and distributed computing to scale such algorithms appears as an obvious path forward. But the reward varies depending on how amenable the algorithms are to parallel processing and the availability of rapid prototyping capabilities with low cost of entry. To use this data deluge to our advantage, a convergence between the fields of Parallel and Distributed Computing and Machine Learning is necessary.

The primary motivation of the ParLearning workshops is to bring together researchers from Artificial Intelligence, Analytics, Machine Learning, Data Mining, and Parallel & Distributed Computing communities to share current work and innovative ideas, identify areas that require the most convergence, and assess the impact on the broader technical landscape. In 2018, the papers address a variety of issues relevant to scaling up machine learning and data mining algorithms on parallel and distributed computing architectures. These include:

- Methods for mitigating stragglers (outlier nodes that run very slowly) in distributed gradient descent.
- Flexible implementation of widely-used convolutional neural networks for many-core architectures such as Epiphany.
- Out-of-core matrix multiplication using a cluster of GPUs.
- Hierarchical matrix approximations for scalable and efficient kernel ridge regression.

ParLearning 2018 will include four keynote lectures. The first talk will be delivered by Abhinav Vishnu from AMD. The title of his talk is “Scaling Deep Learning Algorithms on Extreme Scale Architectures”. The second talk will be given by Azalia Mirhoseini from Google Brain, the title of her talk is “Model Parallelism optimization with deep reinforcement learning”. The third talk will be given by Thomas Parnell from IBM Research and the title of his talk will be “Introduction to Snap Machine Learning”. The fourth and final talk will be given by T.B.A., the title of their talk will be “T.B.A”.

ParLearning 2018 represents the cutting edge in the multidisciplinary area of Parallel and Distributed Computing and Machine Learning. Succeeding workshops are expected to highlight innovations in this growing area.

Organizing Committee:

- General co-chairs: Henri Bal (Vrije Universiteit, The Netherlands) and Arindam Pal (TCS Research and Innovation, India)
- Technical Program co-chairs: Azalia Mirhoseini (Google Brain, USA) and Thomas Parnell (IBM Research – Zurich, Switzerland)
- Publicity chair: Yanik Ngoko (Université Paris XIII, France)
- Steering Committee: Sutanay Choudhury (Pacific Northwest National Laboratory, USA), Anand Panangadan (California State University, Fullerton, USA), and Yinglong Xia (Huawei Research America, USA)

978-1-5386-5555-9/18/$31.00 ©2018 IEEE
DOI 10.1109/IPDPSW.2018.00133

Technical Program Committee:

- Vito Giovanni Castellana, Pacific Northwest National Laboratory, USA
- Tanmoy Chakraborty, IIIT Delhi, India
- Sutanay Choudhury, Pacific Northwest National Laboratory, USA
- Erich Elsen, Google Brain, USA
- Dinesh Garg, IIT Gandhinagar and IBM Research, India
- Kripabandhu Ghosh, IIT Kanpur, India
- Saptarshi Ghosh, IIT Kharagpur, India
- Kazuaki Ishizaki, IBM Research - Tokyo, Japan
- Debnath Mukherjee, TCS Research, India
- Francesco Parisi, University of Calabria, Italy
- Saurabh Paul, PayPal, USA
- Jianting Zhang, City College of New York, USA