

Accelerate Parallelism in Large Scale AI

Ji Liu, PhD

Computer Science & Electrical and Computer Engineering
University of Rochester

Abstract: Parallelism is the key strategy to improve the efficiency for solving large scale machine learning tasks, by involving multiple workers to work in parallel. In this talk, three frameworks are introduced to accelerate the parallel computation in minimizing ML objectives. The first framework is to study the asynchronous parallelism to avoid the synchronization overhead in the traditional synchronization framework. Our work solves an open theoretical problem by showing the convergence and speedup properties of the asynchronous stochastic gradient method used in modern machine learning tools such as Tensorflow, CNTK, and MXnet. The second framework is to design the decentralized computational topology to avoid the communication traffic at the central node in the traditional centralized computational topology. This work is selected to be an oral paper at NIPS2017. The third framework is to study how to design algorithms using low precision data to reduce the communication and computation complexity while ensure the correctness. All three frameworks are validated in both theory and practice in our work.

Bio: *Ji Liu* is currently an assistant professor in Computer Science, Electrical Computer Engineering, and Goergen Institute for Data Science at University of Rochester (UR). He received his Ph.D., Masters, and B.S. degrees from University of Wisconsin-Madison, Arizona State University, and University of Science and Technology of China respectively. His research interests cover a broad scope of machine learning, optimization, and their applications in other areas such as healthcare, bioinformatics, computer vision, and many other data analytics related areas. His recent research focus is on reinforcement learning, structural model estimation, asynchronous parallel optimization, sparse learning (compressed sensing) theory and algorithm, healthcare, etc. He founded the machine learning and optimization group at UR and published more than 40 papers in top conferences and journals. He won the award of Best Paper honorable mention at SIGKDD 2010, the award of Facebook Best Student Paper award at UAI 2015, the IBM faculty award 2017, and the MIT TR35 award China.

Tuesday, April 3, 2018

1:30 - 2:20 p.m.

MEB 235