The 1st Parallel and Distributed Computing for Machine Learning and Inference Problems (*ParLearning* 2012)

In Conjunction with IPDPS 2012, Shanghai, China

Data-driven computing needs no introduction today. The case for using data for strategic advantages is exemplified by web search engines, online translation tools and many more examples. The past decade has seen 1) the emergence of multicore architectures and accelerators as GPGPUs, 2) widespread adoption of distributed computing via the map-reduce/hadoop eco-system and 3) democratization of the infrastructure for processing massive datasets ranging into petabytes by cloud computing. The complexity of the technological stack has grown to an extent where it is imperative to provide frameworks to abstract away the system architecture and orchestration of components for massive-scale processing. However, the growth in volume and heterogeneity in data seems to outpace the growth in computing power. A "collect everything" culture stimulated by cheap storage and ubiquitous sensing capabilities contribute to increasing the noise-to-signal ratio in all collected data. Thus, as soon as the data hits the processing infrastructure, determining the value of information, finding its rightful place in a knowledge representation and determining subsequent actions are of paramount importance. To use this data deluge to our advantage, a convergence between the field of Parallel and Distributed Computing and the interdisciplinary science of Artificial Intelligence seems critical. From application domains of national importance as cyber-security, health-care or smart-grid to providing real-time situational awareness via natural interface based smartphones, the fundamental AI tasks of Learning and Inference need to be enabled for large-scale computing across this broad spectrum of application domains.

Many of the prominent algorithms for learning and inference are notorious for their complexity. Adopting parallel and distributed computing appears as an obvious path forward, but the mileage varies depending on how amenable the algorithms are to parallel processing and secondly, the availability of rapid prototyping capabilities with low cost of entry. The first issue represents a wider gap as we continue to think in a sequential paradigm. The second issue is increasingly recognized at the level of programming models, and building robust libraries for various machine-learning and inferencing tasks will be a natural progression. As an example, scalable versions of many prominent graph algorithms written for distributed shared memory architectures or clusters look distinctly different from the textbook versions that generations of programmers have grown with. This reformulation is difficult to accomplish for an interdisciplinary field like Artificial Intelligence for the sheer breadth of the knowledge spectrum involved. The primary motivation of the proposed workshop is to invite leading minds from AI and Parallel & Distributed Computing communities for identifying research areas that require most convergence and assess their impact on the broader technical landscape.

978-0-7695-4676-6/12 $26.00 © 2012 IEEE
DOI 10.1109/IPDPSW.2012.352

IEEE computer society

Organization

General Co-chairs:

- Sutanay Choudhury, Pacific Northwest National Laboratory, USA
- George Chin, Pacific Northwest National Laboratory, USA
- Yinglong Xia, IBM T.J. Watson Research Center, USA

Local Chair:

- Yihua Huang, Nanjing University, China

Program Co-chairs:

- John Feo, Pacific Northwest National Laboratory, USA
- Chandrika Kamath, Lawrence Livermore National Laboratory, USA
- Anshul Gupta, IBM T.J. Watson Research Center, USA

Invited Speakers (keynote speech):

- Haixun Wang, Microsoft Research, China

  Title: *Probase + Trinity: a distributed graph knowledge database that knows our mental world*
- Hari Cadambi, NEC Laboratories America, USA

  Title: *Accelerating Machine Learning on Heterogeneous Clusters: Lessons Learnt and Future Directions*

Committee

- Arindam Banerjee, University of Minnesota, USA
- Arindam Pal, Indian Institute of Technology, Delhi, India
- Edmond Chow, Georgia Institute of Technology, USA
- Eric Goodman, Sandia National Laboratories, USA
- Feiping Nie, University of Texas, Arlington, USA
- Haixun Wang, Microsoft Research, China
- Hyeran Jeon, University of Southern California, USA
- Jatin Chhugani, Intel Corp., USA
- Jun Wang, IBM Research, USA
- Lawrence B. Holder, Washington State University, USA
- Lexing Xie, Australian National University, AU
- Oreste Villa, Pacific Northwest National Laboratory, USA
- Shuping Liu, eBay Inc., USA
- Tina Eliassi-Rad, Rutgers University, USA
- Weizhu Chen, Microsoft Research, China
- Yangqiu Song, Microsoft Research, China
- Yi Wang, Tencent Holdings Lt., China
- Yihua Huang, Nanjing University, China