

MapGraph - Graph processing at 30 billion edges per second on NVIDIA GPUs

Zhisong Fu
SYSTAP, LLC
fuzhisong@systap.com

Harish Kumar Dasari
University of Utah
hdasari@sci.utah.edu

Bradley Bebee
SYSTAP, LLC
beebbs@systap.com

Martin Berzins
University of Utah
mb@sci.utah.edu

Bryan Thompson
SYSTAP, LLC
bryan@systap.com
(Presenting)

MapGraph is a disruptive technology that delivers extreme performance for graph problems on many-core hardware. MapGraph can be run on a laptop, on EC2 HPC GPU compute nodes, and on large GPU compute clusters. With processing speeds of up to 3 billion edges per second on a single GPU, MapGraph changes what is possible with your data.

Many-core computing is the future. CPU architectures are not getting any faster. Continued performance gains must come from many-core technologies such as GPUs or the Intel Xeon Phi. GPUs are widely known for their role in games, high-performance computing, and high FLOPS/watt ratio. However, graph algorithms are data-intensive, not compute intensive, and have degree dependent parallelism. As a consequence, graph algorithms place an extreme burden on the memory bus and communications network.

SYSTAP and the Scientific Computing and Imaging Institute have developed a capability for extreme performance parallel graph algorithms on GPUs from laptops to large GPU clusters. MapGraph provides a scalable technology for data-intensive workloads that addresses the data-dependent parallelism, memory, and communication bottlenecks. I will review this research and present our roadmap for this technology.

Learn more at <http://mapgraph.io>.

This work was (partially) funded by the DARPA XDATA program under AFRL Contract #FA8750-13-C-0002. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. D14PC00029. The authors would like to thank Dr. White, NVIDIA, and the MVAPICH group at Ohio State University for their support of this work.