Massively Parallel Graph Analytics

Manycore graph processing, distributed graph layout, and supercomputing for graph analytics

George M. Slota[1,2,3] Kamesh Madduri[2]
Sivasankaran Rajamanickam[1]

[1]Sandia National Laboratories, [2]Penn State University, [3]Blue Waters Fellow
gslota@psu.edu, madduri@cse.psu.edu, srajama@sandia.gov

Blue Waters Symposium 11 May 2015

Research Motivation and Goals

- Graph analysis is key for the study of biological, chemical, social, and other networks
- Real-world graphs are big, irregular, complex
  - Graph analytics is one of DARPA's 23 toughest mathematical challenges
  - Web graph: 3.5B sites, 129B hyperlinks
  - Brain graph: 100B neurons, 1,000T synaptic connections
  - **Goal**: How can we analyze these massive graphs on supercomputers?
- Modern computational systems like Blue Waters are also big and complex
  - Multiple levels of parallelism, memory hierarchy, hardware configurations, GPUs and coprocessors
  - **Goal**: How can we generically optimize graph algorithms for varying computational hardware?

Methods and Approaches

- **Observation**: most graph algorithms follow a tri-nested loop structure
  - Optimize for this general algorithmic structure
  - Transform structure for more parallelism
- **Observation**: varying in-memory distributed graph layout affects total execution time
  - Partition graph to minimize per-task computation and communication
  - Order vertices within partition for optimal cache performance
- **Observation**: previous approaches for massive graph analytics have only considered external memory solutions
  - Use proper distributed layout to efficiently store graph in distributed memory supercomputer
  - Use algorithmic and layout optimizations to concurrently minimize intra-node execution times and inter-node communication times

Results - Improving Computation and Communication

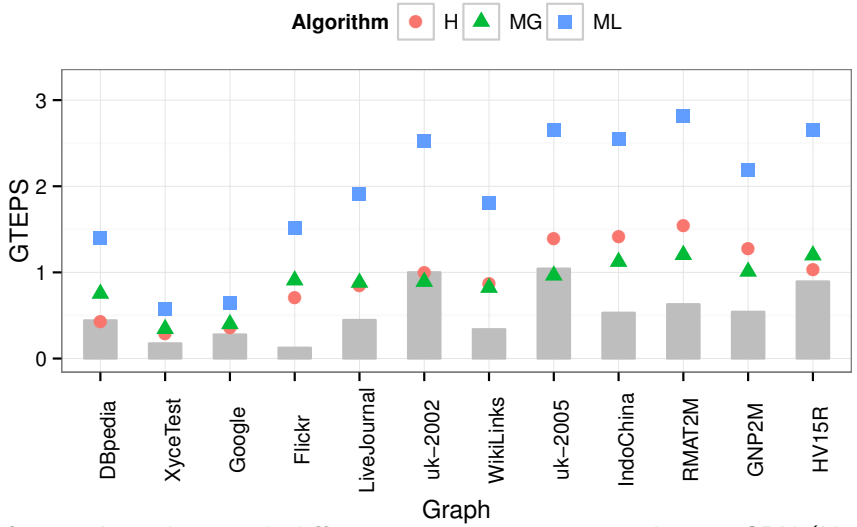

Computational performance rate of a graph analytic with different optimization approaches on GPU (H: hierarchical, MG: global approach, ML: Local approach, Grey bar: baseline)

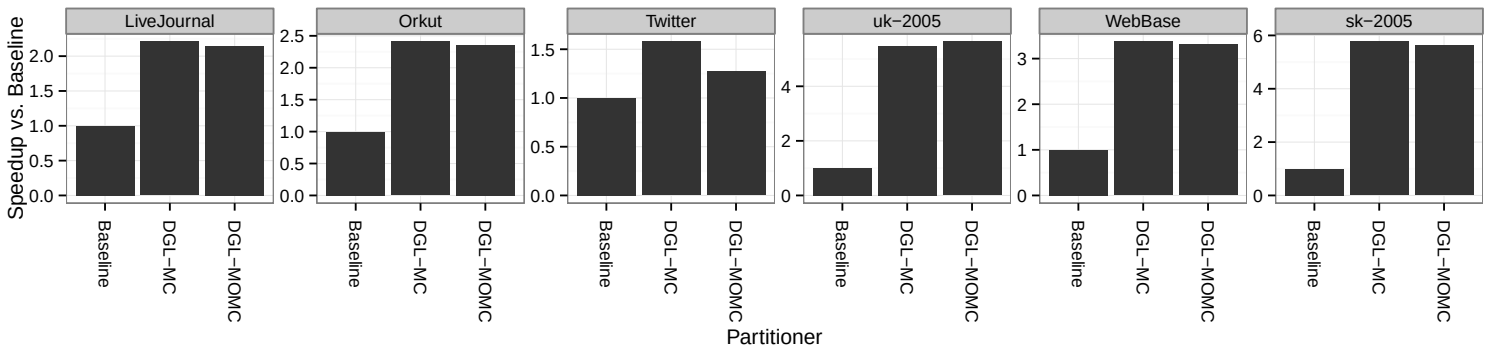

Communication speedups for a complex analytic relative to a random baseline with different distributed layout approaches (DGL-MC: multi-constraint, DGL-MOMC: multi-object)

Results - Analyzing the Internet

- Using performance optimization approaches, we can find communities and most important pages by centrality measures in minutes using Blue Waters

Largest Communities Discovered (numbers in millions)

Pages	Internal Links	External Links	Representative Page
112	2126	32	YouTube
18	548	277	Tumblr
9	516	84	Creative Commons
8	186	85	WordPress
7	57	83	Amazon
6	41	21	Flickr

Individual Page Centrality Rankings

In Degree	PageRank	Harmonic
YouTube	YouTube	WordPress
WordPress	YouTube/t/..	Twitter
YouTube/t/..	YouTube/testtube	Twitter/privacy
YouTube/..	YouTube/..	Twitter/About
YouTube/..	Tumblr	Twitter/account
YouTube/t/..	Google/..	Twitter/about

Publications Based on Fellowship Work

- **Distributed Graph Layout for Scalable Small-world Network Analysis**
  - George M. Slota, Kamesh Madduri, Sivasankaran Rajamanickam
  - In submission
- **Supercomputing for Web Graph Analytics**
  - George M. Slota, Sivasankaran Rajamanickam, Kamesh Madduri
  - Under Review
- **High-performance Graph Analytics on Manycore Processors**
  - George M. Slota, Sivasankaran Rajamanickam, Kamesh Madduri
  - To appear in the Proceedings of the 29th IEEE International Parallel and Distributed Processing Symposium (IPDPS15)

Summary of Accomplishments

- Optimizations for manycore parallelism result in up to a **3.25× performance improvement** for graph analytics executing on GPU
- Modifications to in-memory storage of graph structure results in up to a **1.48× performance improvement** for distributed analytics running with MPI+OpenMP on Blue Waters
- **First-ever analysis** of largest to-date web crawl (129B hyperlinks) on a distributed memory system
- Running on 256 nodes of Blue Waters, we are able to run several complex graph analytics on the web crawl **in only minutes of execution time**
- **These approaches will allow further scaling to analyze even larger graphs**, such as our brain's neural network (1K trillion connections)

Future Work

- **Implement more graph analytic algorithms**
  - Subgraph counting
  - Other community detection approaches
  - etc.
- **Further improve scaling and performance**
  - Explore parameter space of optimizations
  - Vary layout objectives and constraints per-algorithm
- **Acquire and analyze larger and more complex networks on Blue Waters**
- **Planned future presentations of fellowship work**:
  - Presentation of manycore-based optimizations strategies at IPDPS15
  - Poster presentation of overall layout approach at IPDPS15
  - Presentation and poster presentation of web graph analytics at SC15 (tentative)

Acknowledgments

This research is part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (awards OCI-0725070, ACI-1238993, and ACI-1444747) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications. This work is also supported by NSF grants ACI-1253881, CCF-1439057, and the DOE Office of Science through the FASTMath SciDAC Institute. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.