Extreme Big Data Processing in Large-Scale Graph Analytics and Billion-Scale Social Simulation

Toyotaro Suzumura
IBM Research
Tokyo, Japan
suzumurat@gmail.com

ABSTRACT

This paper introduces some of the example applications handling extremely big data with supercomputers such as large-scale network analysis, X10-based large-scale graph analytics library, Graph500 benchmark, and billion-scale social simulation.

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: C.5.1 [**Computer System Implementation**]*;* D.2.11 [**Software Architectures**]; E.1 [**Data Structures**]

Keywords: Distributed computing, graph algorithms, social simulation, X10

1. Large-Scale Graph Analysis

Recently, social network services such as Twitter, Facebook, MySpace, LinkedIn have been remarkably growing. Haewoon performed the analysis of the Twitter network on June 2009 and showed the degree of separation in Twitter network. However, the number of users on 2009 is about 41.7 million, the graph scale is not very large compared to continuously growing current graph. Our study in [1] shows the transition of the number of users on Twitter from June 2006 to September 2012. The number of users on September 2012 attains 469.9 million and the number of relationships attains 28.7 billion. This data collection is obtained by our series of crawling for 3 months conducted in late 2012. Therefore, it is considered that with increasing users, the graph characteristics has changed greatly and we analyzed for the current large graph. The motivation of our work [1] is to understand how such characteristics is changed and evolving from the results in 2009. We used HyperANF API as an analysis tool to compute approximate degree of separation and diameter in a sampling based fashion. To obtain the precise result as much as possible, the tool requires more memory. Our computation was performed on a 64-core machine with 512 GB memory that is one of the TSUBAME 2.0 super computer located Tokyo Institute of Technology.

In order to obtain the degree of separation and diameter with HyperANF from the crawled Twitter network, a list of preprocessing are required. In this preprocessing and the resulting output consumes around 20 TB approximately. Currently Hadoop and Web Graph APIs are used for the preprocessing.

1. Prepare user id lists containing serial number from zero to renumber
2. Create edge lists from follower-friend data and renumber with the user id lists.
3. Convert the renumbered edge lists to adjacency lists formatted Ascii Graph.
4. Finally, Convert the adjacency lists to compressed data with BV compression API .

Note that the adjacency lists and the compressed data size are 263GB and 73GB, respectively. We conduct a Twitter network analysis in terms growth by region, scale-free, reciprocity, degree of separation and diameter using Twitter user data with 469.9 million users and 28.7 billion relationships. We report that the value of degree of separation is 4.59 in current Twitter network through our experiments. Some of the hints such as storage requirements, tools, workflow towards exascale computing in this domain can be obtained by this kind of analysis.

2. ScaleGraph: X10-Based Large-Scale Graph Analytic Library

Graphs will be a prominent computational workload in Exascale era as demonstrated in our work [8][9][13][14]. Large graph analysis is a dilemma faced by programmers in various domains such as scientific applications, biology, national security, business analytics described in previous section. We have developed ScaleGraph [2] which is an open-source X10 library available from [3] for massive graph analytics targeting large scale graph analysis scenarios. The differentiating features of ScaleGraph is as follows:

(1) High Productive HPC Graph Analysis : ScaleGraph library has been developed to reduce complexity and increase programmer productivity involved in use of HPC systems for large graph analysis. We provide an object oriented interface for users of ScaleGraph while preserving scalability in large scale distributed environments.

(2) Comprehensive PGAS Library for Large Graph Analysis : Our library is designed from ground up with aiming complex network analysis community.

(3) Scalability Analysis in Distributed Environment : We demonstrated scalability of our library in distributed environments such as TSUBAE 2.0. We are currently expanding a set of graph analytics and also doing the same web mining analysis as the work in previous section.

3. Graph500 Benchmark

As an alternative to Linpack, Graph500 [4] was recently developed. We conducted a thorough study of the algorithms of

ICPE'14, March 22–26, 2014, Dublin, Ireland.
ACM 978-1-4503-2733-6/14/03.
http://dx.doi.org/10.1145/2568088.2576096

the reference implementations and their performance in an earlier paper [5]. Based on that work, we implemented a scalable and high-performance implementation of an optimized Graph500 benchmark for large distributed environments [6][7]. In contrast to the computation-intensive benchmark used by TOP500, Graph500 is a data-intensive benchmark. It does breadth-first searches in undirected large graphs generated by a scalable data generator based on a Kronecker graph. There are six problem classes: toy, mini, small, medium, large, and huge. Each problem solves a different size graph defined by a Scale parameter, which is the base 2 logarithm of the number of vertices. For example, the level Scale 26 for toy means 226 and corresponds to 1010 bytes occupying 17 GB of memory. The six Scale values are 26, 29, 32, 36, 39, and 42 for the six classes. The largest problem, huge (Scale 42), needs to handle around 1.1 PB of memory. As of this writing, Scale 38 is the largest that has been solved by a top-ranked supercomputer. Our work [6] proposed an optimized method based on 2D partitioning and other methods such as communication compression and vertex sorting. Our optimized implementation can handle BFS (Breadth First Search) of a large graph with Scale 35 with 462.25 GE/s while using 1366 nodes and 16,392 CPU cores. This competition is greatly challenging since new scalable algorithms have been proposed rapidly. We have been continuously enhancing the scalable algorithm and implementation on various supercomputers.

4. Billion-Scale Social Simulation

We introduce billion-scale social simulation [10][11][12] in this section. Towards the contribution to the human society, global economy, ecology, the analysis of human brain characteristics and our daily life, the research in multi-agent simulation is entering into the era of simulating billion-scale agents. Although prior arts tackle distributed agent simulation platform to achieve this goal, it is not sufficient to simulation billion-scale agent behaviors. Based on this observation, we report the first effort for building such an infrastructure platform that handles billion-scale agent simulation platform. In our previous work, we introduce X10-based agent simulation platform for such a purpose and presents its application to traffic simulation. We were able to handle only at maximum 10 millions of agents, but the performance was not scalable due to various reasons such as work load imbalance, global synchronization. Our work in [12] present the work of purely implementing the whole simulation stack including both the simulation runtime and the application layer such as traffic simulation [11] by the use of the state-of-the-art PGAS language. By implementing the system in such a manner and evaluating the system in highly distributed systems, it is observed that the system can be close to handle billion-scale agents in near real-time. The first experimental result is that the performance scalability is greatly achieved by simulating 1 millions of agents on 1536 CPU cores and 256 nodes. By compiling fully X10-based agent simulation system into C++ and MPI, it only takes 77 seconds for 600 simulation steps which is nearly 10 times faster than real-time. Moreover, by using the entire whole country-wide network of India as the agents' underlying infrastructure, we successfully simulated 1 billion agents with 400 nodes in TSUBAME 2.0. This is the first attempt to deal with such a gigantic number of agents and we believe that this infrastructure would be the basis of large-scale agent simulation in various fields.

5. REFERENCES

[1] Masaru Watanabe and Toyotaro Suzumura, "How Social Network is Evolving ? A Preliminary Study on Billion-scale Twitter Network", LSNA 2013 Workshop in conjunction with WWW 2013

[2] Miyuru Dayarathna, Charuwat Houngkaew, Toyotaro Suzumura, "Scalable Performance of ScaleGraph for Large Scale Graph Analysis", HiPC 2012 (IEEE International Conference on High Performance Computing)

[3] ScaleGraph Library: http://www.scalegraph.org/

[4] Graph500 : http://www.graph500.org/.

[5] Toyotaro Suzumura, Koji Ueno, Hitoshi Sato, Katsuki Fujisawa and Satoshi Matsuoka, "Performance Evaluation of Graph500 on Large-Scale Distributed Environment", IEEE IISWC 2011

[6] Koji Ueno and Toyotaro Suzumura "Highly Scalable Graph Search for the Graph500 Benchmark", the 21st International ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC), 2012/06, Delft, Netherland.

[7] Koji Ueno and Toyotaro Suzumura, "Parallel Distributed Breadth First Search on GPU", HiPC 2013 (IEEE International Conference on High Performance Computing), India, 2013/12

[8] Charuwat Houngkaew and Toyotaro Suzumura, "X10-Based Distributed and Parallel Betweenness Centrality and Its Application to Social Analytics", HiPC 2013 (IEEE International Conference on High Performance Computing), India, 2013/12

[9] Koichi Shirahata, Hitoshi Sato, Toyotaro Suzumura and Satoshi Matsuoka, "A Scalable Implementation of a MapReduce-based Graph Processing Algorithm for Largescale Heterogeneous Supercomputers", 13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), Delft, Netherlands, 2013/5

[10] Toyotaro Suzumura and Hiroki Kanezashi, "Highly Scalable X10-based Agent Simulation Platform and its Application to Large-scale Traffic Simulation", 2012 IEEE/ACM 16th International Symposium on Distributed Simulation and Real Time Application

[11] Takayuki Osogami, Takashi Imamichi, Hideyuki Mizuta, Toyotaro Suzumura, and Tsuyoshi Ide, "Toward simulating entire cities with behavioral models of traffic," *IBM Journal of Research and Development*, **57**(5): Paper 6, 2013

[12] Toyotaro Suzumura and Hiroki Kanezashi, "A Holistic Architecture for Super Real-Time Multiagent Simulation Platform", Winter Simulation Conference 2013, 2013/12, Washington D.C., US

[13] Miyuru Dararathna and Toyotaro Suzumura, "XGDBench: A Benchmarking Platform for Graph Stores in Exascale Clouds", IEEE CloudCom 2012 conference, Taipei, Taiwan, 2012/12

[14] Bao Nguyen and Toyotaro Suzumura, "Towards Highly Scalable Pregel-based Graph Processing Platform with X10", The 2nd International Workshop on Large Scale Network Analysis (LSNA 2013) In conjunction with WWW 2013, 2013/05, Rio de janeiro, Brazil.