

SYNC or ASYNC? Time to Fuse for Distributed Graph-parallel Computation

Chenning Xie, Rong Chen, Haibing Guan, Binyu Zang and Haibo Chen

About Graph Algorithm

Collaborative Filtering

□ Alternating Least Squares

Structured Prediction

□ Loopy Belief Propagation

Classification

□ Neural Networks

Graph Analytics

□ PageRank


□ SSSP

□ Triangle-Counting


□ Graph Coloring

"Think as a Vertex"


Dependent Data



Local Accesses




Iterative Computation



Distributed Graph-parallel Systems

Algorithm → Impl. **compute()** for vertex
Computation → Compute all Vertex in **parallel**



G. Malewicz, et al. Pregel: a system for large-scale graph processing. In SIGMOD, 2010.

Y. Low, et al. Distributed GraphLab: a framework for machine learning and data mining in the cloud. VLDB Endowment, 2012.

Deficiency of Existing Schedulings (e.g. SYNC vs. ASYNC)

Properties

→ Communication
→ Convergence

SYNC vs. ASYNC

Regular vs. Irregular
Slow vs. Fast




Favorites

→ Algorithm
→ Parallel Workload
→ Scalability

SYNC vs. ASYNC

I/O Bound vs. CPU bound
Heavyweight vs. Moderate
| Graph | vs. | Machines |

Divergence



Challenges and Opportunities

1. The divergence results **complexity** in mode choosing, and **suboptimal performance** with either single-mode.
2. Diff. in internal structures should be maintained **convertible** from each other.
3. Ideal benefit requires appropriate switch **timing**.


PowerSwitch: Adaptive Mode Switch for Graph-parallel Computation

To support transparent
adaptive mode switch:




- Profiler
- Predictor

Normalized throughput
with heuristic rules.

- Compatible Fault Tolerance
- Convertor:



Experiment



Architecture

