Efficient Distribution for Deep Learning on Large Graphs

Loc Hoang[1], Xuhao Chen[2], Hochan Lee[1], Roshan Dathathri[3], Gurbinder Gill[3], and Keshav Pingali[1,3]

[1]The University of Texas at Austin
[2]Massachusetts Institute of Technology
[3]Katana Graph

DeepGalois Overview

- Scalable distributed GNN framework
- Generalize GNN to **vertex programming model**
  - *Topology driven*: All vertices are active
  - *Operators*
    - Aggregate features from 1-hop neighbors
    - Local linear transformation
  - *Termination condition*: number of layers
- **Distributed Graph Engine**
  - CuSP [1]: Graph partitioner
  - Gluon [2]: Communication substrate
  - Galois [3]: Computation engine
- Outperform the state-of-the art
  - 4x speedup over DistDGL [4]

Synchronization

- Support arbitrary use-defined partitioning policy
  - Optimize communication based on the partitioning policy

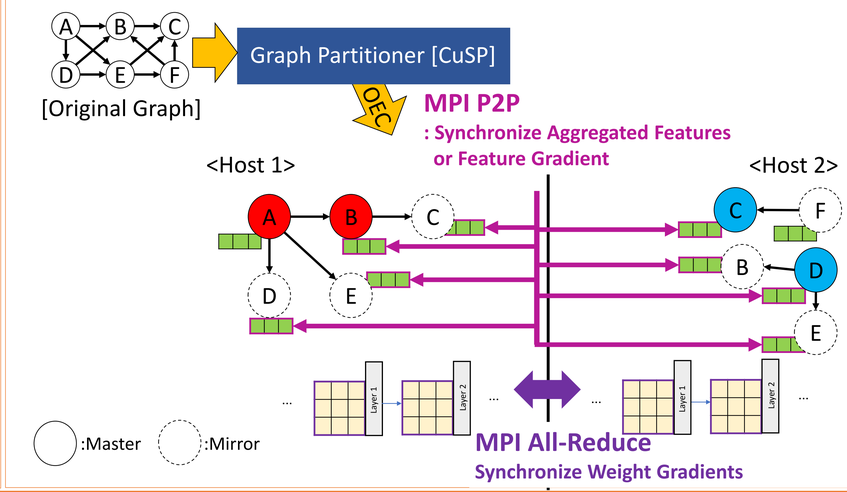

Evaluation

GraphSAGE average epoch time

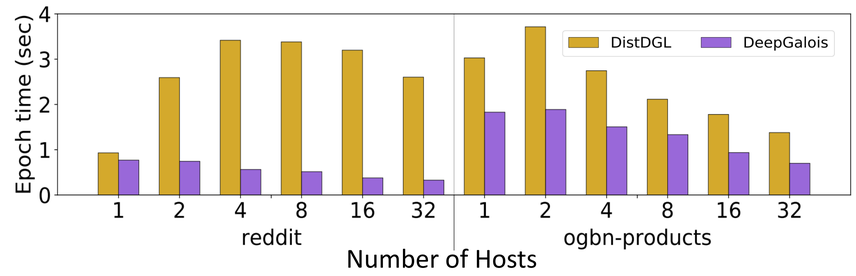

- DeepGalois 4x faster than DistDGL
- 32 hosts vs 1 host for reddit
  - DeepGalois: 2.4x speedup
  - DistDGL: 2.8x slowdown

Epoch time breakdown

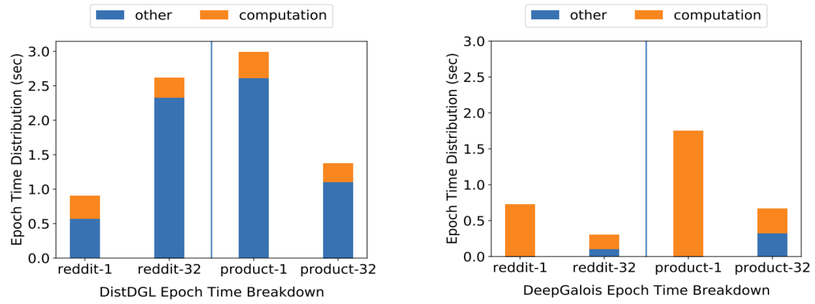

Speed of Convergence

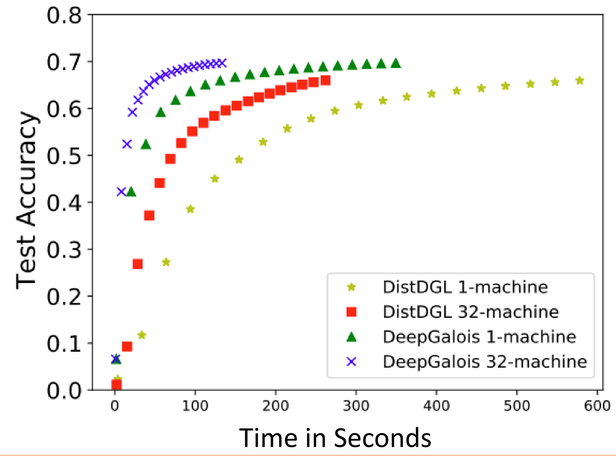

Graph Partitioning

- Edges are uniquely assigned among hosts
  - Proxies for end points
- One proxy is designated the *master* proxy
- Different policies have different edge/master assignment
- Policies trade-off
  - Computation load-balance
  - Communication overhead

- Each vertex proxy has feature vectors
- Hosts share weight matrices for layers
- No accuracy loss

[Partitioning Policy Examples]

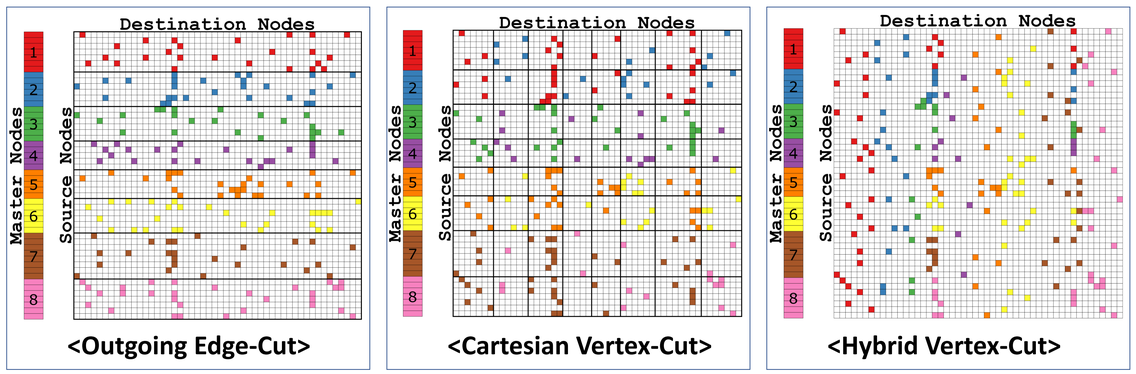

References

[1] **"Cusp: A customizable streaming edge partitioner for distributed graph analytics"** Loc Hoang, et el. In *International Parallel and Distributed Processing Symposium (IPDPS)*, 2019
[2] **"Gluon: A communication optimizing framework for distributed heterogeneous graph analytics"** Roshan Dathathri, et el. In *Proceedings of ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI 2018
[3] **"A lightweight infrastructure for graph analytics"** Donald Nguyen, et el. In *Proceedings of ACM Symposium on Operating Systems Principles, SOSP 2013*
[4] **"DistDGL: Distributed Graph Neural Network Training for Billion-Scale Graphs"** Da Zheng, et el. In *arXiv 2020*