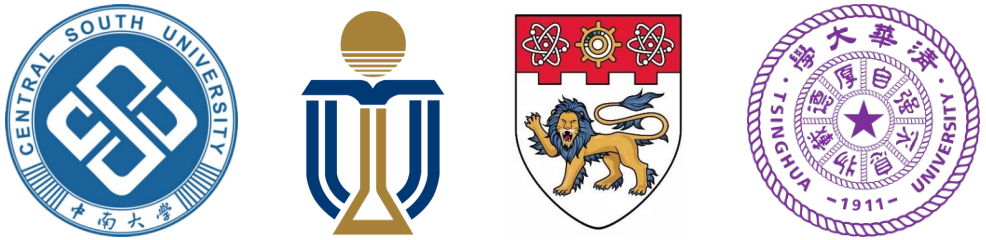

2025 USENIX Annual Technical Conference

Accelerating distributed graph learning by using collaborative in-network multicast and aggregation

Zhaoyi Li[1,3], Jiawei Huang[1], Yijun Li[1], Jingling Liu[1], Junxue Zhang[2], Hui Li[1]
Xiaojun Zhu[1], Shengwen Zhou[1], Jing Shao[1], Xiaojuan Lu[1], Qichen Su[1]
Jianxin Wang[1], **Chee Wei Tan[3]**, Yong Cui[4], Kai Chen[2]

[1]*Central South University,* [2]*Hong Kong University of Science and Technology*
[3]*Nanyang Technological University,* [4]*Tsinghua University*

Introduction

Large-scale Computation of Graph Neural Networks

- Drug discovery, public health surveillance, online social networks

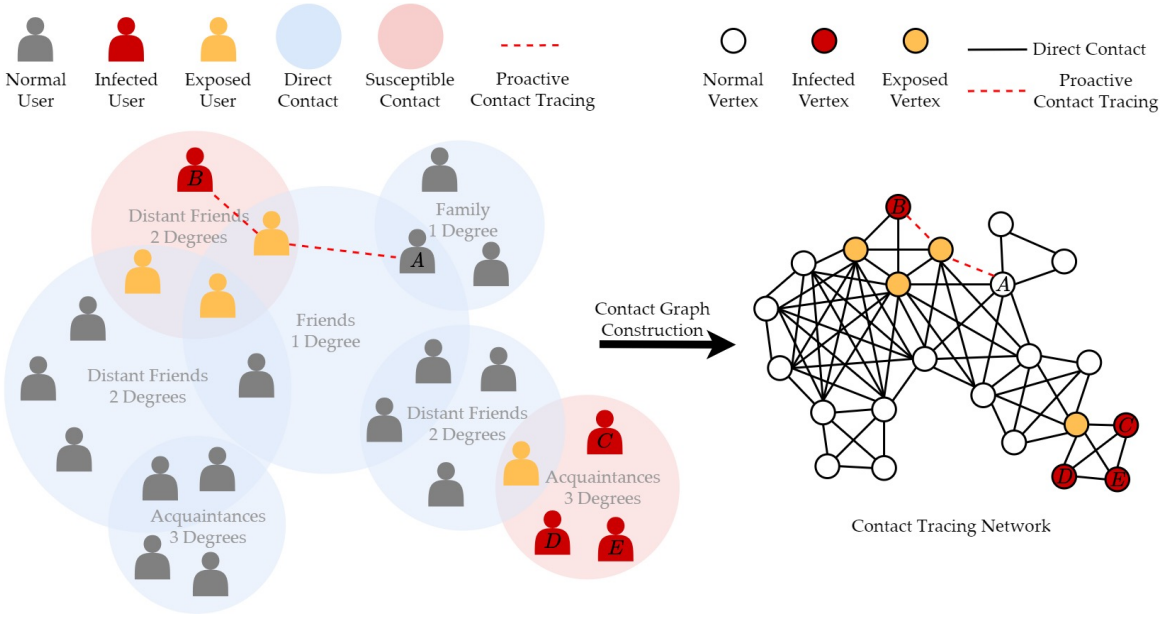

Chee Wei Tan et al, *DeepTrace: Learning to optimize contact tracing in epidemic networks with graph neural networks,* IEEE Trans. on Signal and Information Processing over Networks, 11:97–113, 2025.

Introduction

Graph Neural Network

- GNN can capture structured information from graph-based data

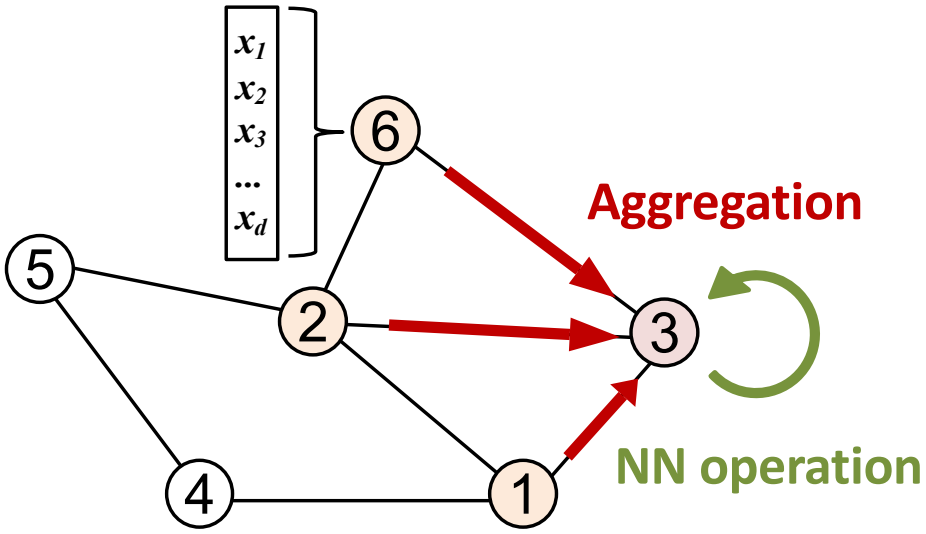

Basic Steps:
(1) Neighbor aggregation
(2) Vertex representation update

Introduction

Communication Bottleneck in Full-Graph Training

- Host-based multicast and aggregation leads to 1-to-N redundant traffic and N-to-1 bandwidth contention

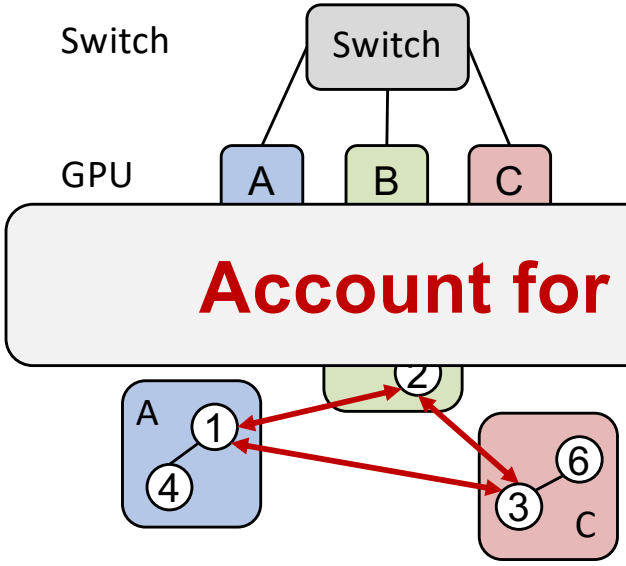

Dataset	Vertices	Edges	Feature	Classes
Ogbn-products	2.4M	62M	100	47
Yelp	0.72M	7M	300	100
Reddit	0.23M	114M	602	41

Account for up to 80% of the epoch time

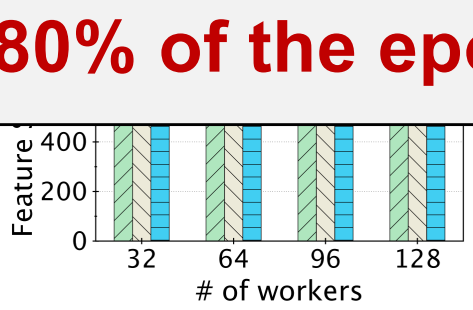

(a) Featrue Size

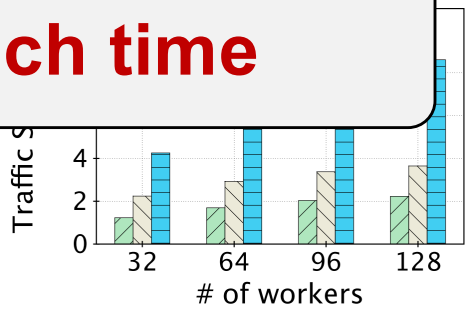

(b) Traffic Size

Introduction

In-Network Aggregation (INA)

- Offloading the gradient aggregation into programmable switch
- Reducing traffic volume for data-parallel training

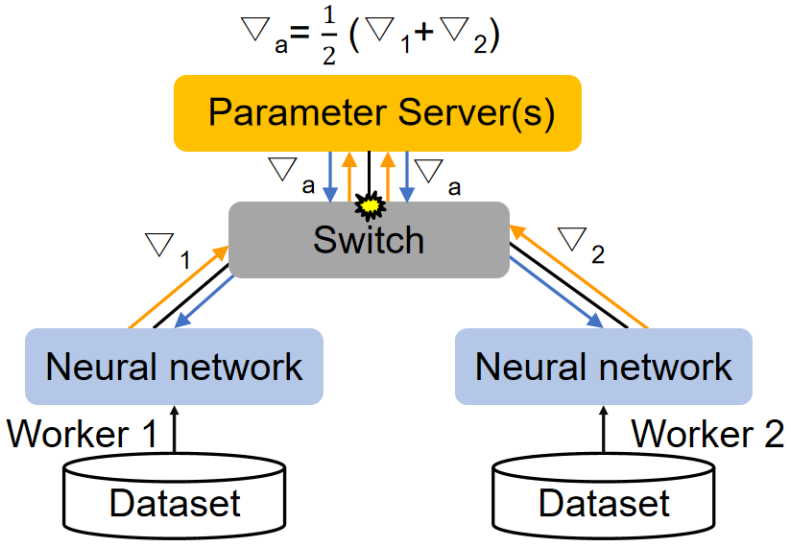

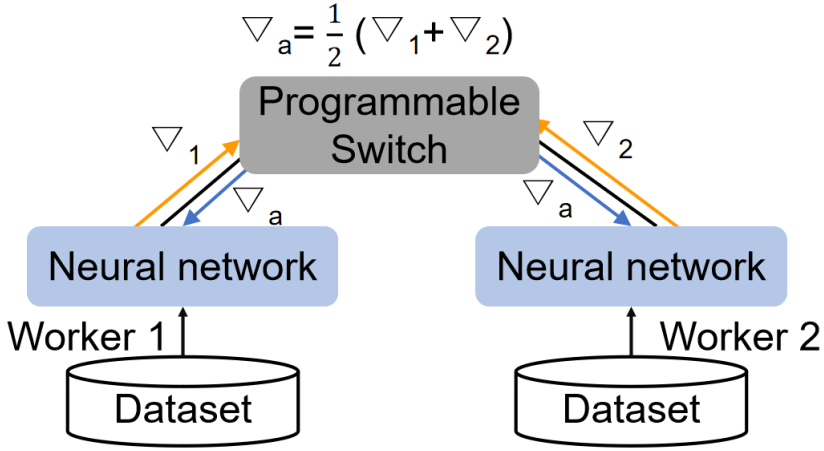

Introduction

- **In-Network Aggregation SwitchML [NSDI'21], ATP [NSDI'21]**
  - Offloading the gradient aggregation into programmable switch
  - Reducing traffic volume for data-parallel training

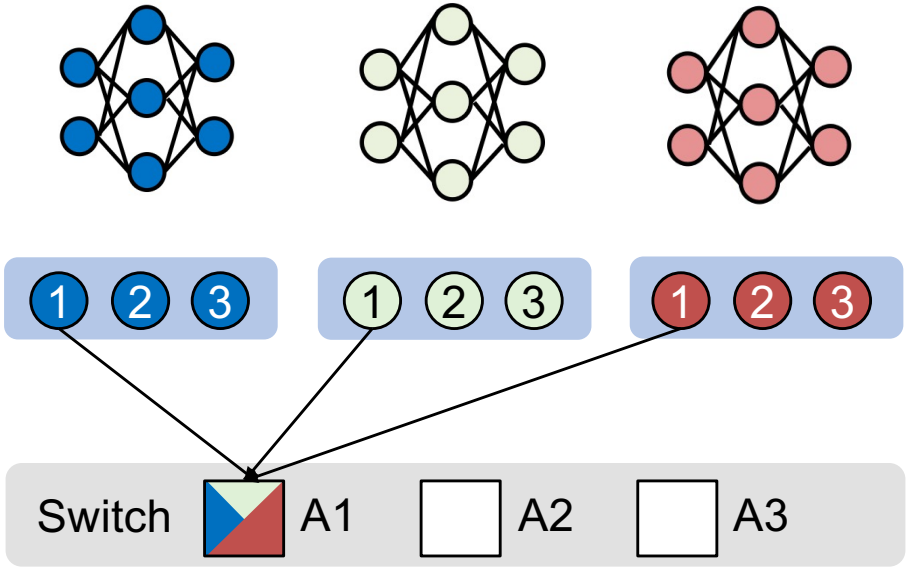

Introduction

- **In-Network Aggregation SwitchML [NSDI'21], ATP [NSDI'21]**
  - Offloading the gradient aggregation into programmable switch
  - Reducing traffic volume for data-parallel training

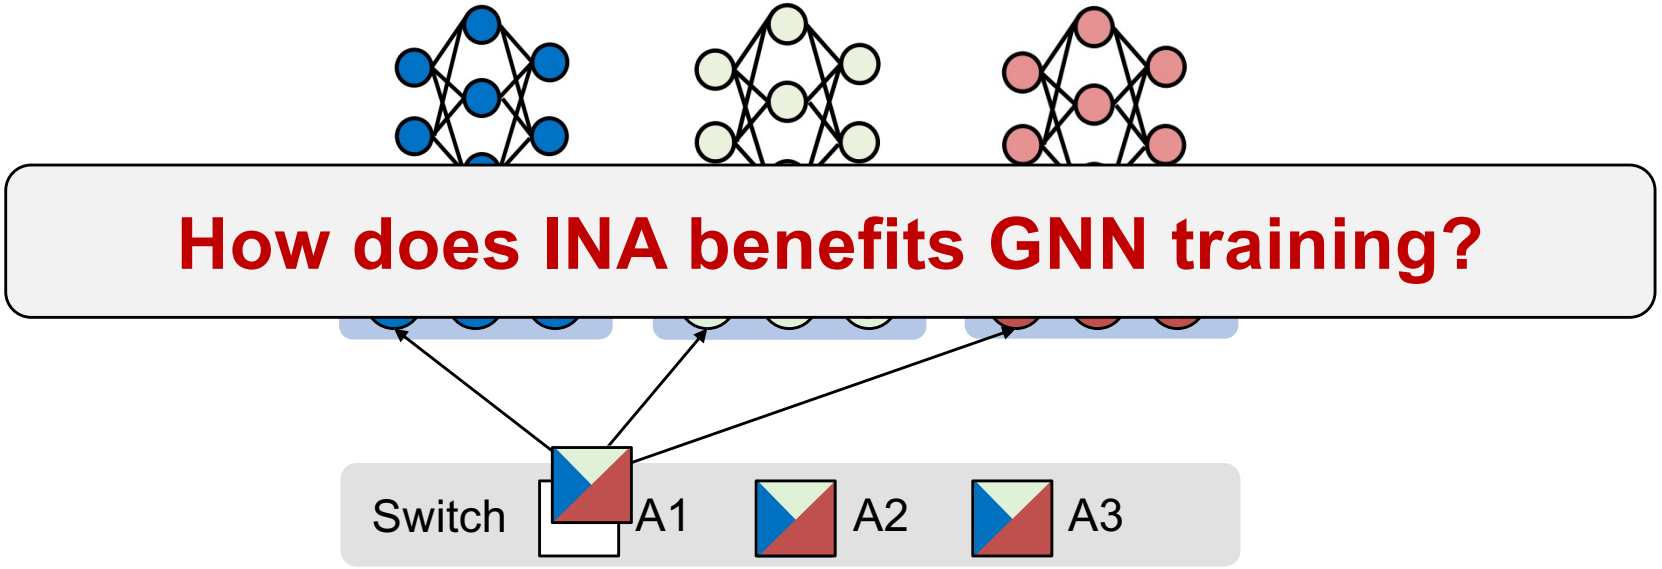

Motivation

Strawman Solution

- Offloading multicast and aggregation of vertices feature into switches
- Eliminating 1-to-N redundant traffic and N-to-1 bandwidth bottleneck

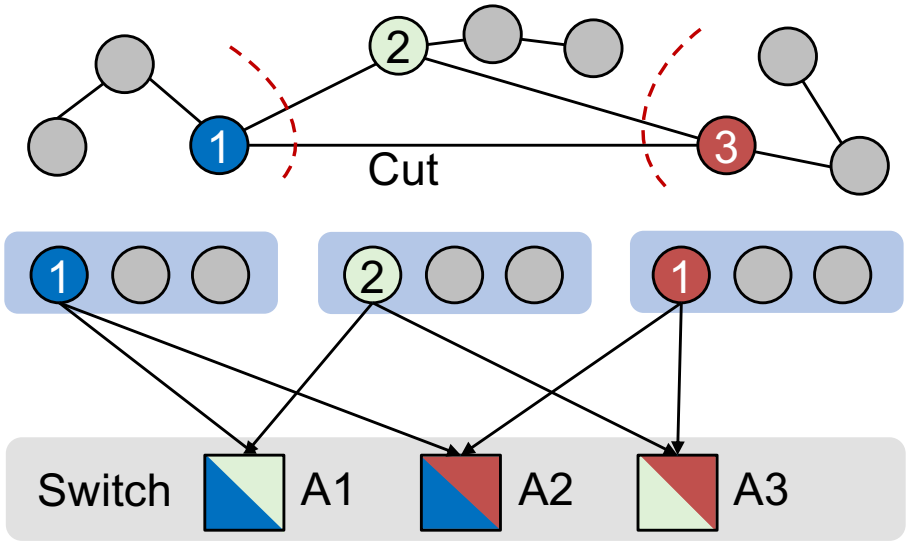

Motivation

Strawman Solution

- Offloading multicast and aggregation of vertices feature into switches
- Eliminating 1-to-N redundant traffic and N-to-1 bandwidth bottleneck

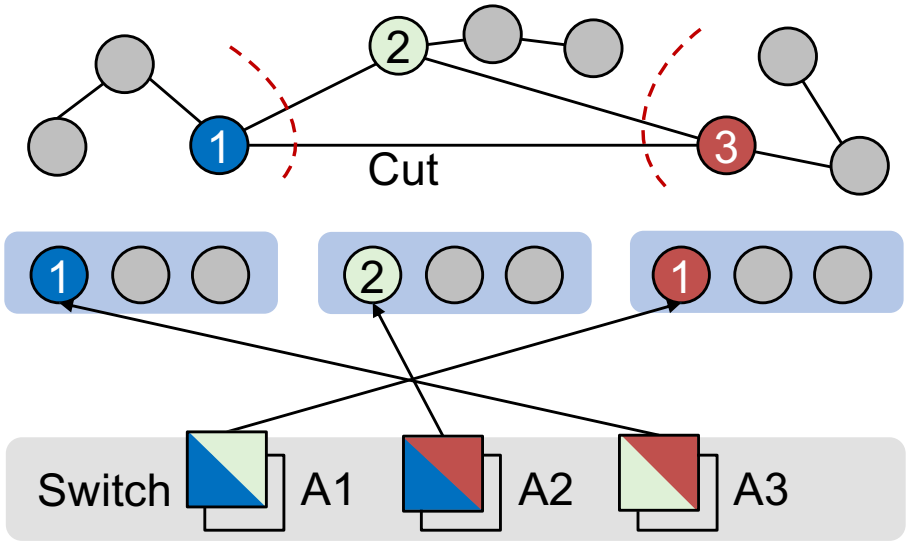

Motivation

Experimental Observations

- Evaluating strawman solution under limited switch memory (100MB)
- Suffering from low aggregation throughput, significant queue backlog, and large traffic volume

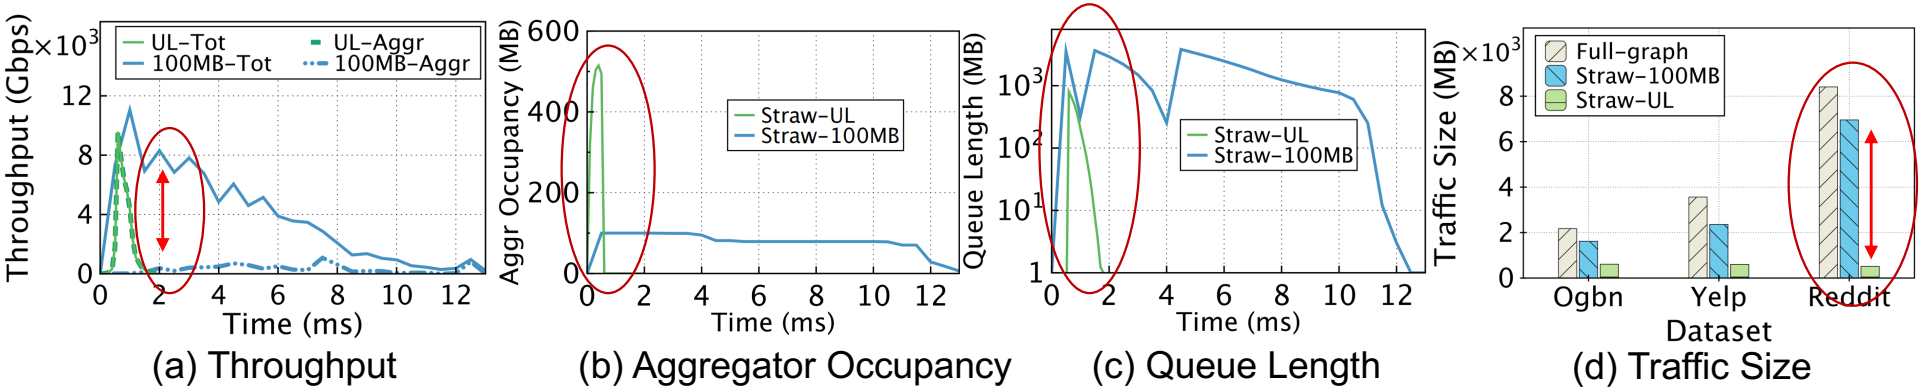

(a) Throughput (b) Aggregator Occupancy (c) Queue Length (d) Traffic Size

Motivation

Deep Dive

- Graph-agnostic multicast order results in inefficient forwarding pipeline

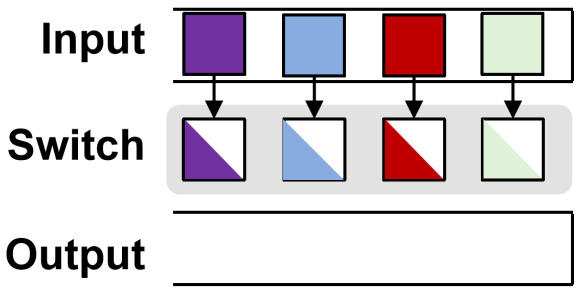

Time 0

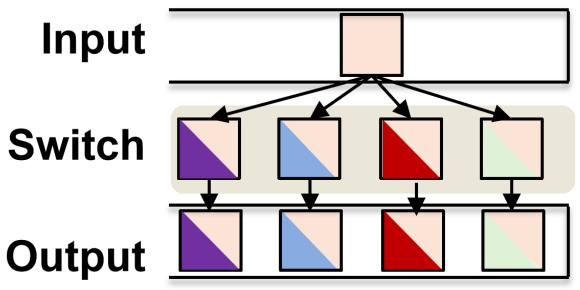

Time 1

Motivation

Deep Dive

- Graph-agnostic multicast order results in inefficient forwarding pipeline
- Huge and interdependent vertex features lead to aggregator overflow

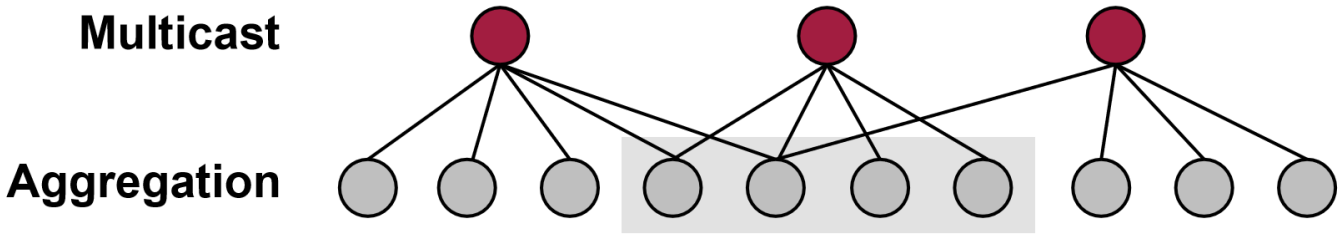

SwitchGNN: Overview

- Graph-aware multicast reordering optimizes the forwarding pipeline
- Multi-level graph partitioning mitigates aggregator overflow
- Reliability and congestion control avoids network congestion

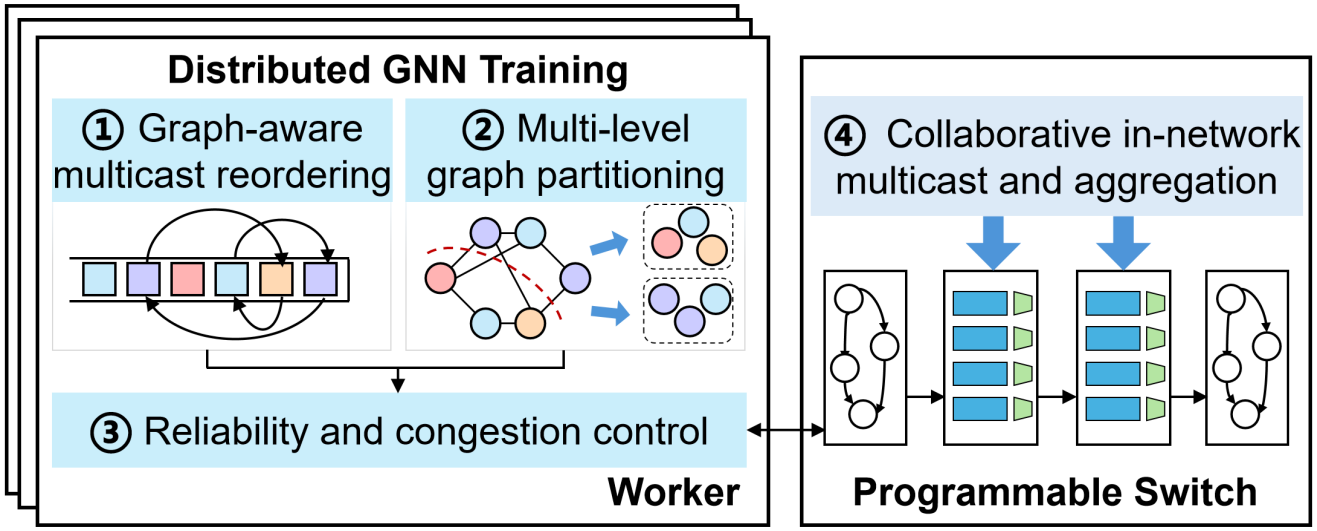

SwitchGNN: Design Details

- **Graph-Aware Multicast Reordering**
  - Prioritizing transmission of high-outdegree vertex features improves forwarding efficiency

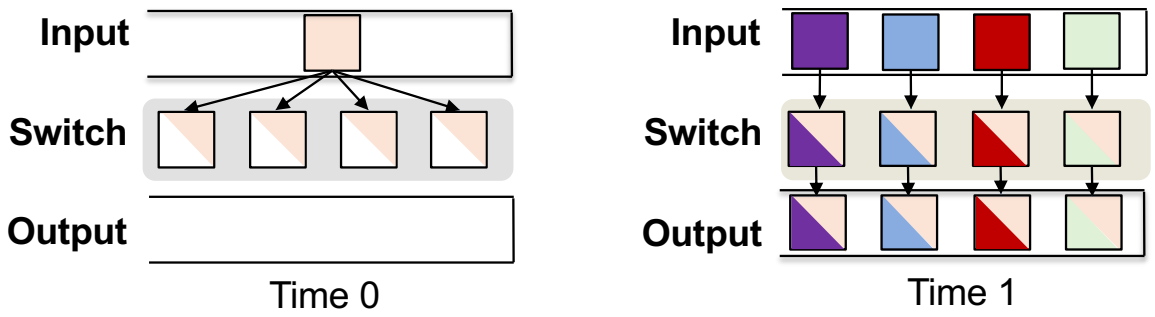

SwitchGNN: Design Details

Graph-Aware Multicast Reordering

- Prioritizing transmission of high-outdegree vertex features improves forwarding efficiency
- Uploading features of neighboring vertices in close time enables faster aggregation at the switch

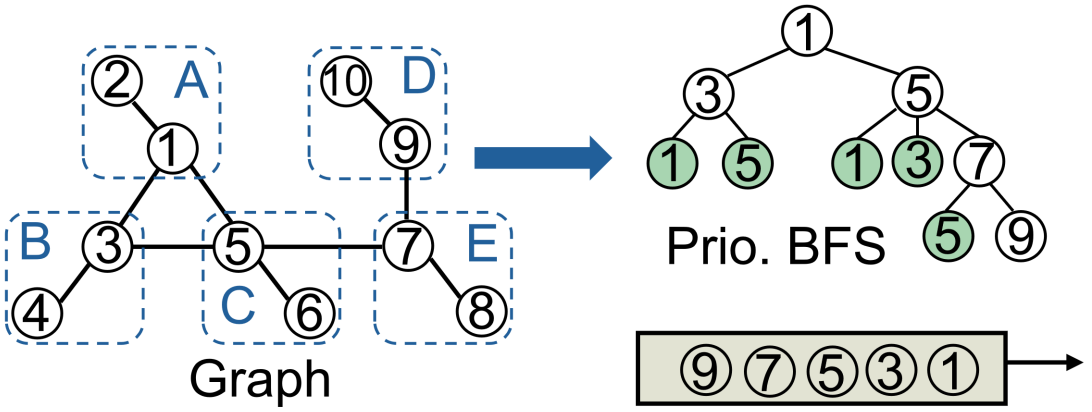

SwitchGNN: Design Details

Multi-level Graph Partitioning

- Partitioning the graph (METIS) to fit switch memory and aggregating them in batches
- Subgraphs formed by cut edges should be still aggregated

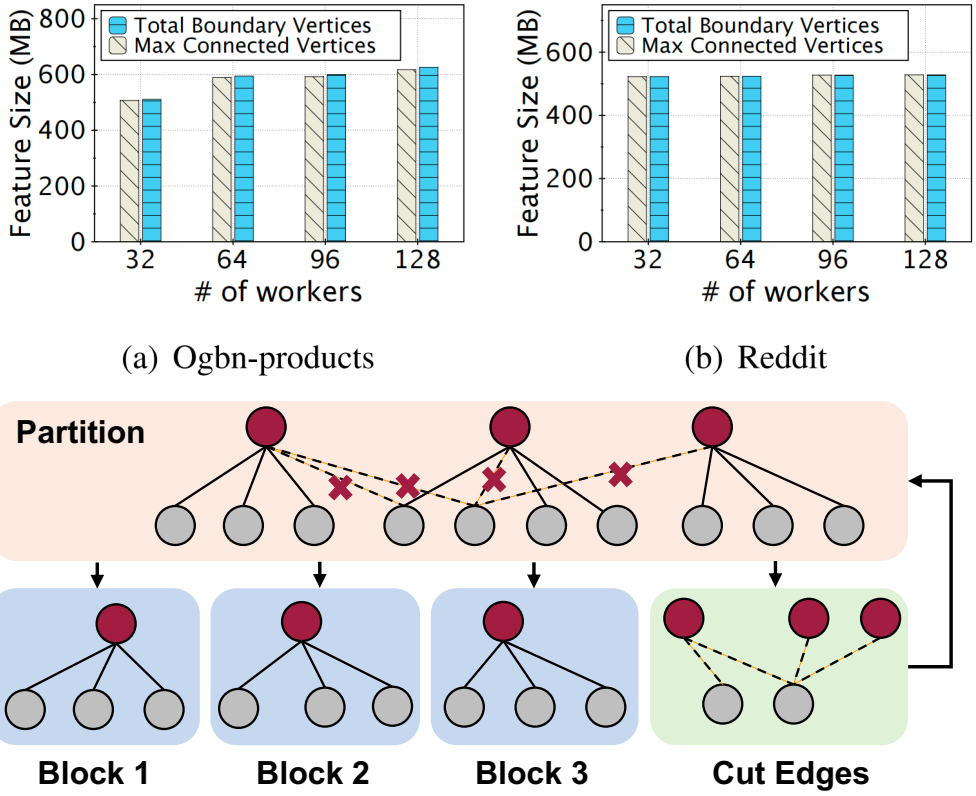

SwitchGNN: Design Details

Multi-level Graph Partitioning

- Partitioning the graph (METIS) to fit switch memory and aggregating them in batches
- Subgraphs formed by cut edges should be still aggregated

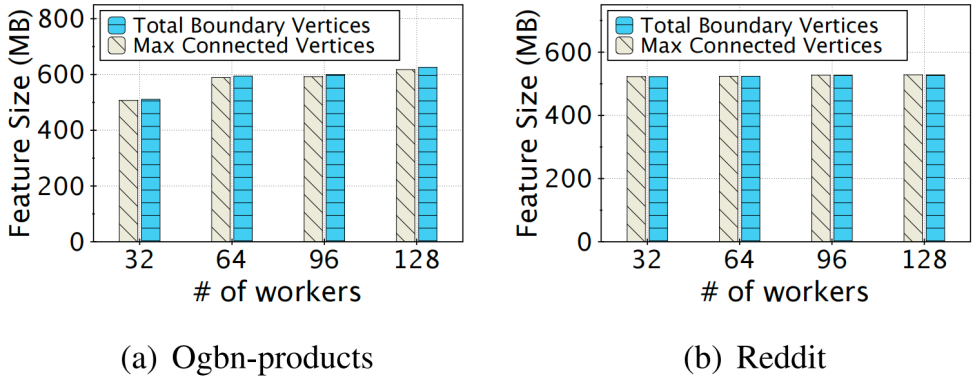

Reliability and Congestion Control

- Guaranteeing correct aggregation and avoid network congestion in all-to-all transmission

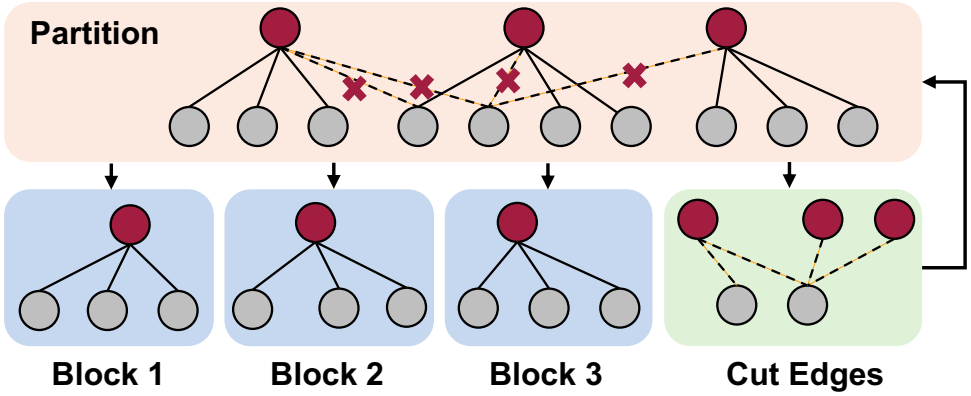

Evaluation

- **Testbed** **DPDK+P4; 8 servers; Star topology; GCN Model**

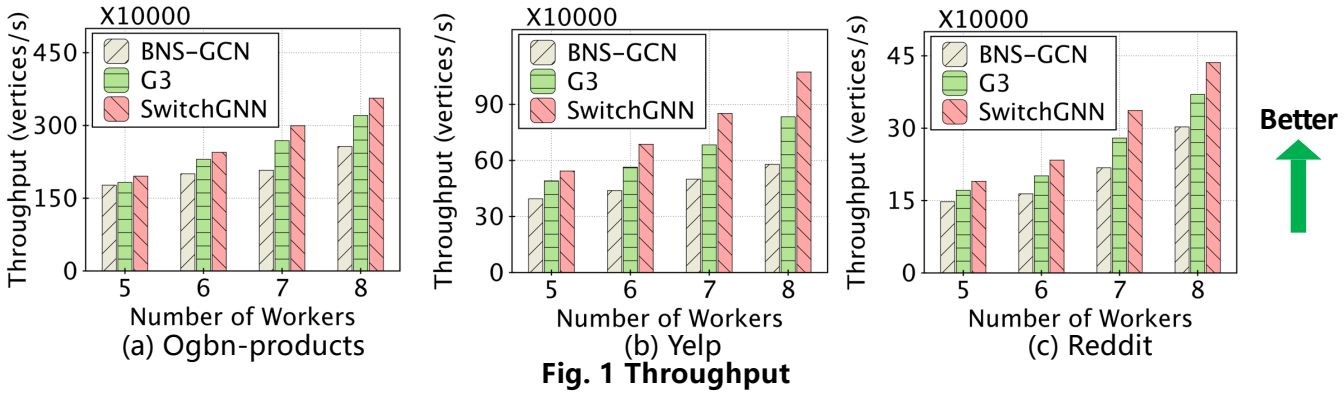

Fig. 1 Throughput

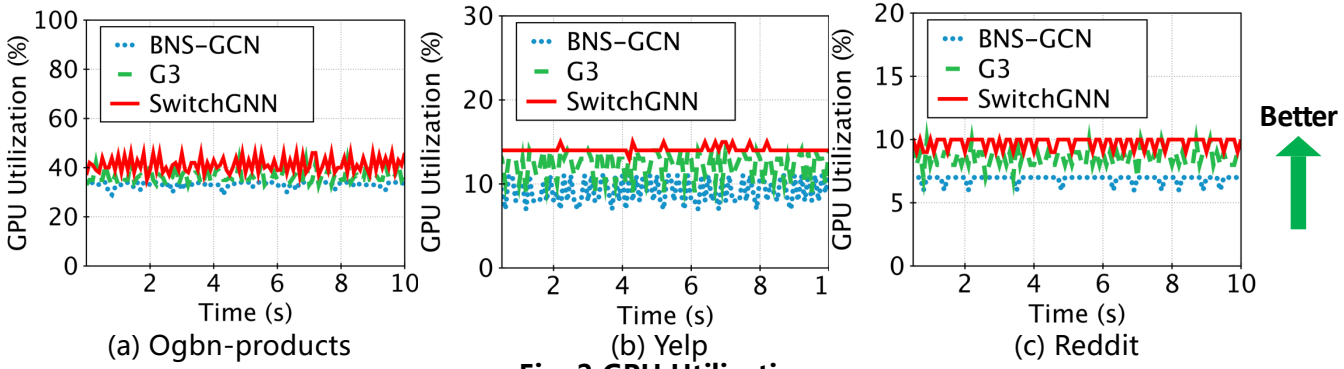

Fig. 2 GPU Utilization

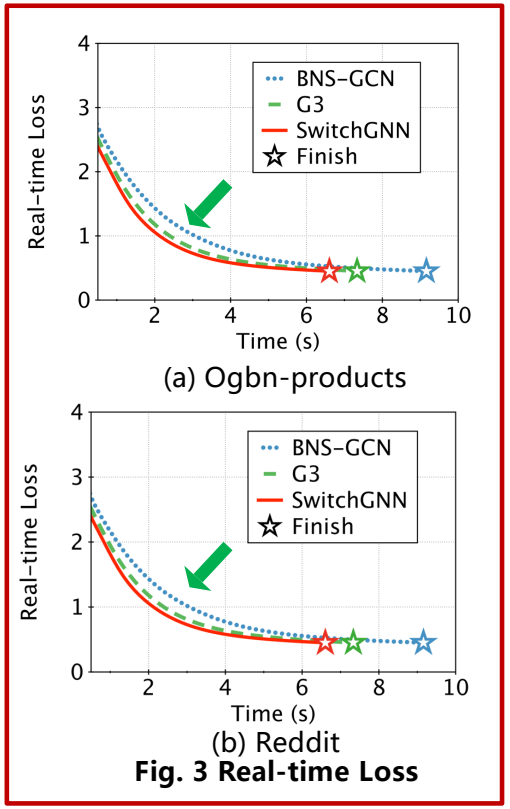

Fig. 3 Real-time Loss

Evaluation

- **NS3 Simulation** **128 workers; Leaf-spine topology; GCN model**

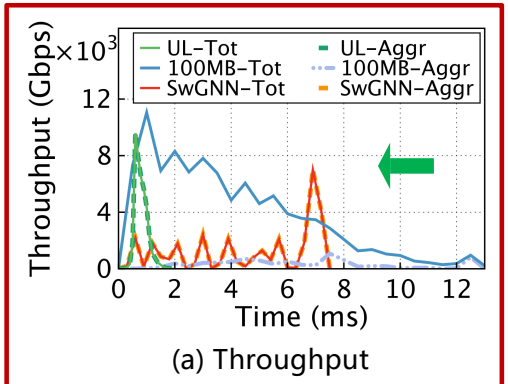

Fig. 4 Basic Performance

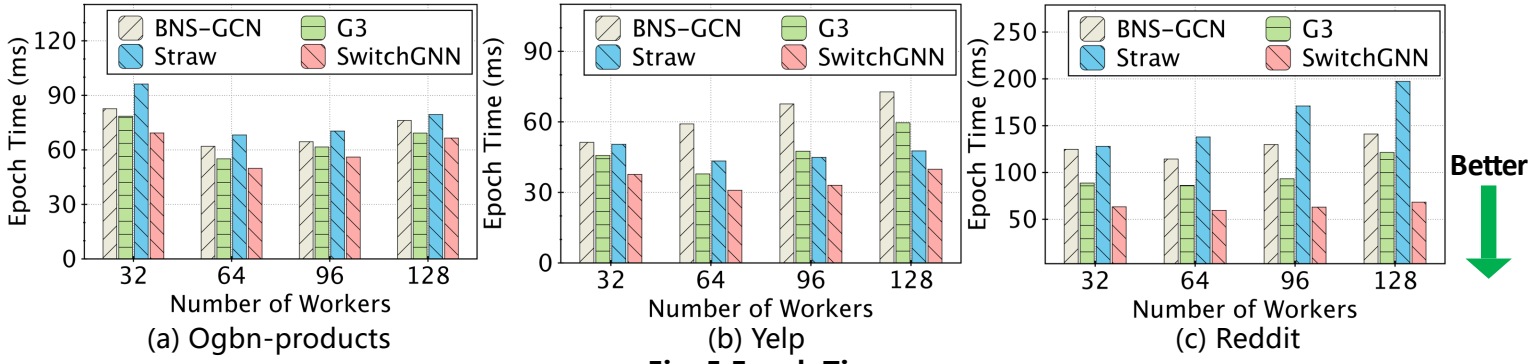

Fig. 5 Epoch Time

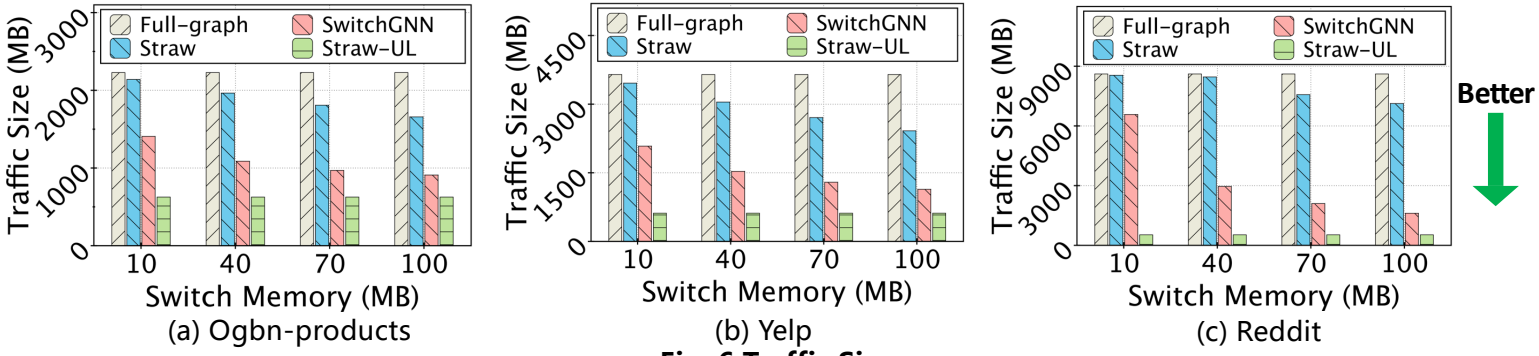

Fig. 6 Traffic Size

Summary

- SwitchGNN exploits in-network multicast and aggregation to accelerate GNN training in data centers
- Graph-aware multicast reordering to optimize the vertex transmission sequence for efficient in-network pipelining
- Multi-level graph partitioning to prevent aggregator overflow and minimize switch traffic
- Key Takeaway: Harness ML workload characteristics to optimize both training efficiency and data center resource usage



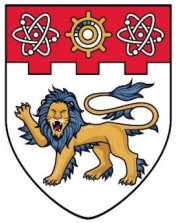


Thank You

cheewei.tan@ntu.edu.sg

Accelerating distributed graph learning by using collaborative in-network multicast and aggregation

Zhaoyi Li[1,3], Jiawei Huang[1], Yijun Li[1], Jingling Liu[1], Junxue Zhang[2], Hui Li[1]
Xiaojun Zhu[1], Shengwen Zhou[1], Jing Shao[1], Xiaojuan Lu[1], Qichen Su[1]
Jianxin Wang[1], **Chee Wei Tan[3]**, Yong Cui[4], Kai Chen[2]

[1]*Central South University,* [2]*Hong Kong University of Science and Technology*
[3]*Nanyang Technological University,* [4]*Tsinghua University*