DGCL: An Efficient Communication Library for Distributed GNN Training

Zhenkun CAI

The Chinese University of Hong Kong

Graph Neural Network (GNN)

Graph neural networks in recent years

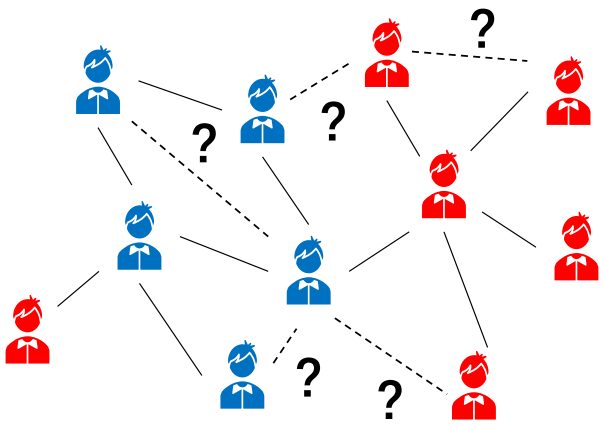

Social network

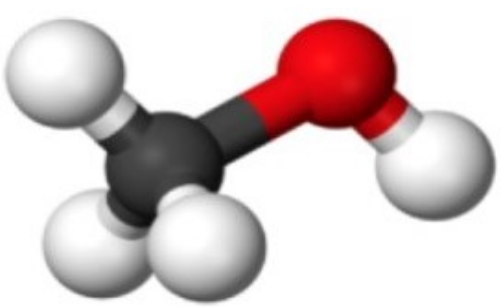

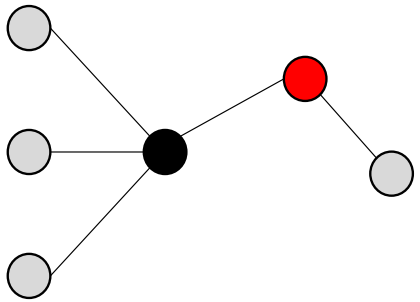

molecule

Graph Neural Networks

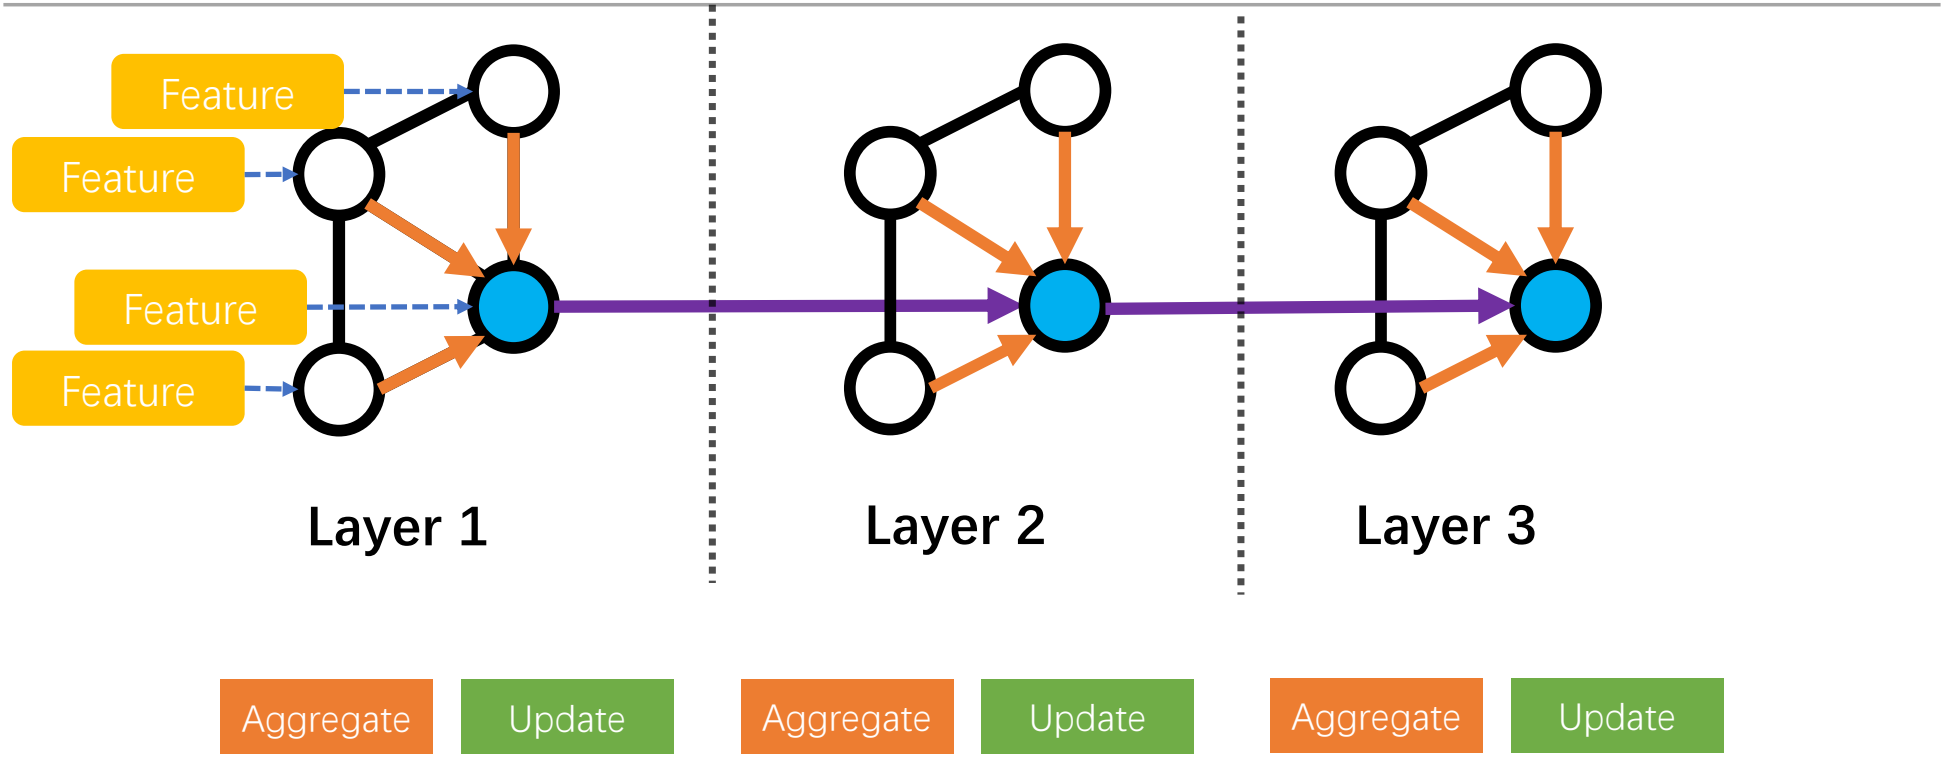

$$a_v^{(k)} = \text{AGGREGATE}^{(k)}(\{h_u^{(k-1)} | u \in \mathcal{N}(v)\})$$
$$h_v^{(k)} = \text{UPDATE}^{(k)}(a_v^{(k)}, h_v^{(k-1)})$$

,

Distributed GNN Training on GPUs

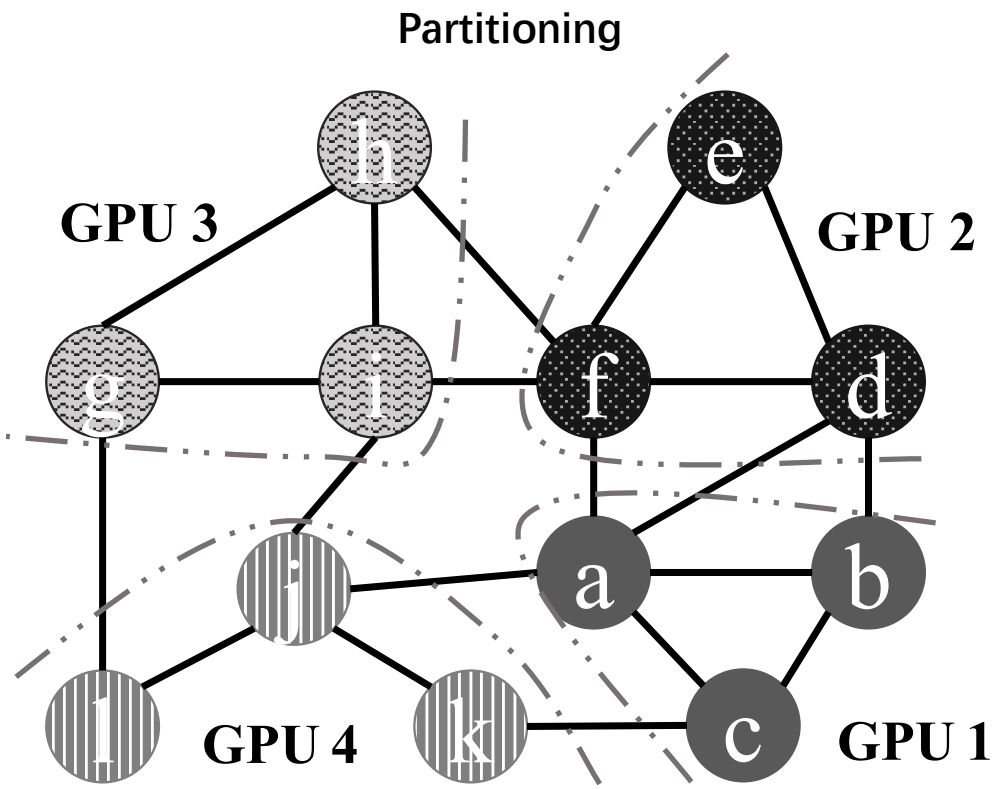

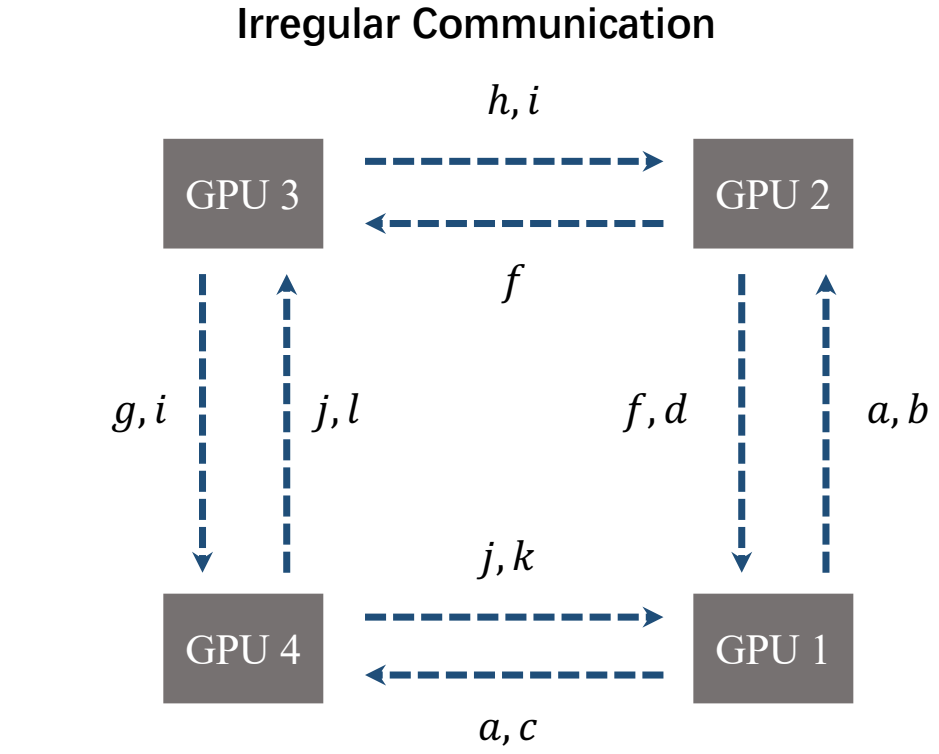

Communication | Aggregate | Update

GNN Systems on GPUs

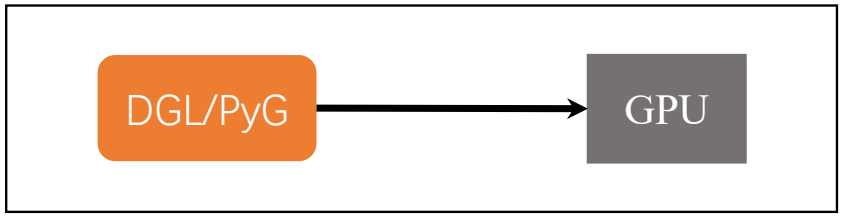

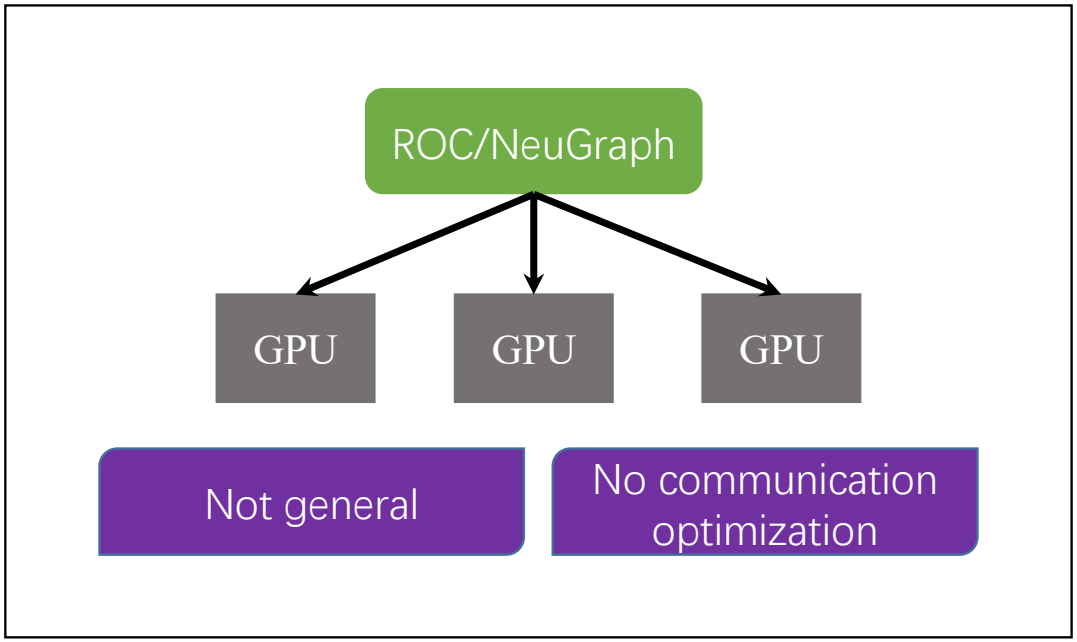

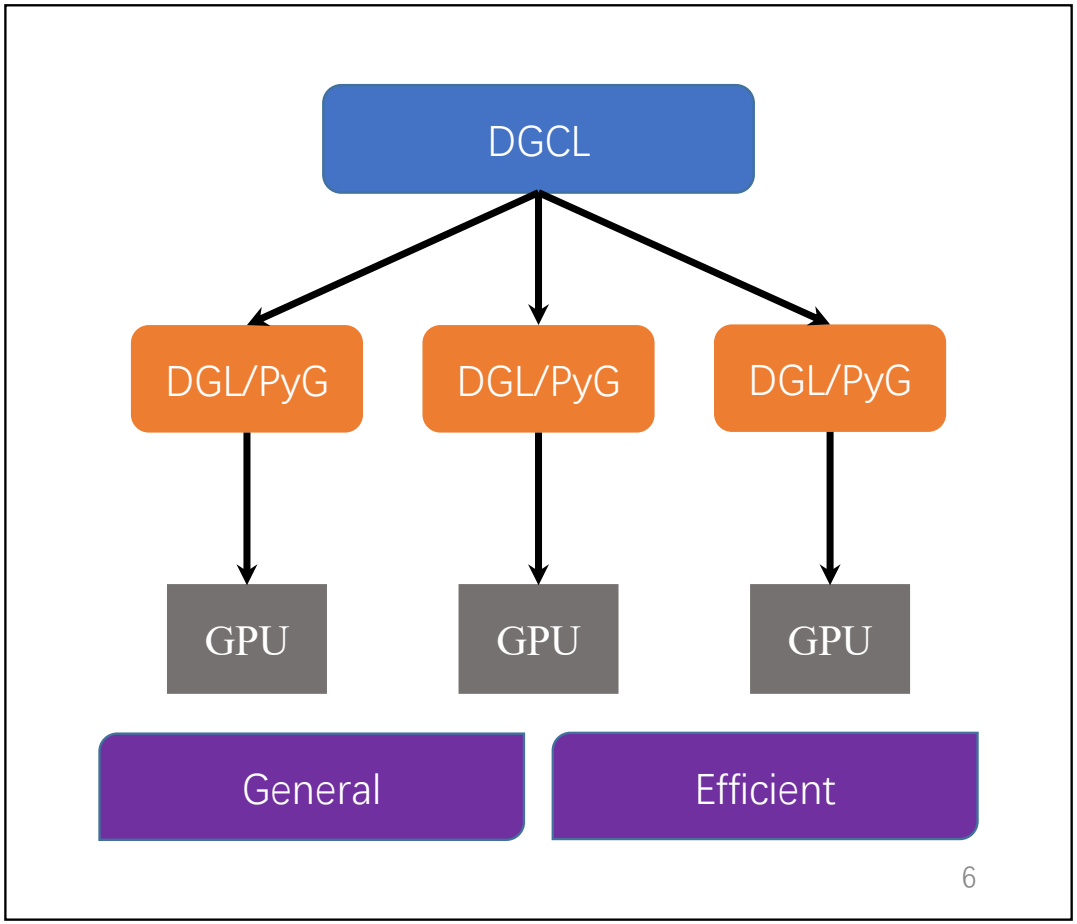

Peer-to-peer communication: each GPU fetch required vertices directly from other GPUs

Unbalanced communication

Bandwidth contention

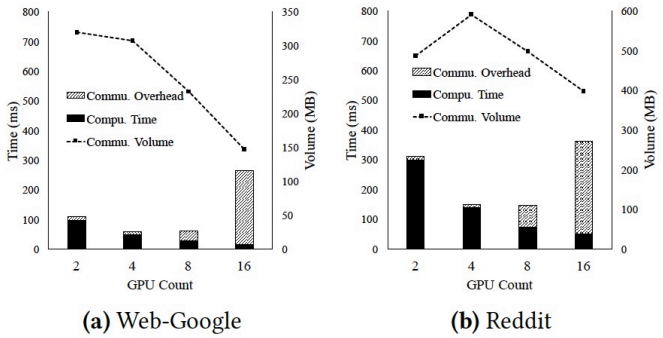

2-layer GCN

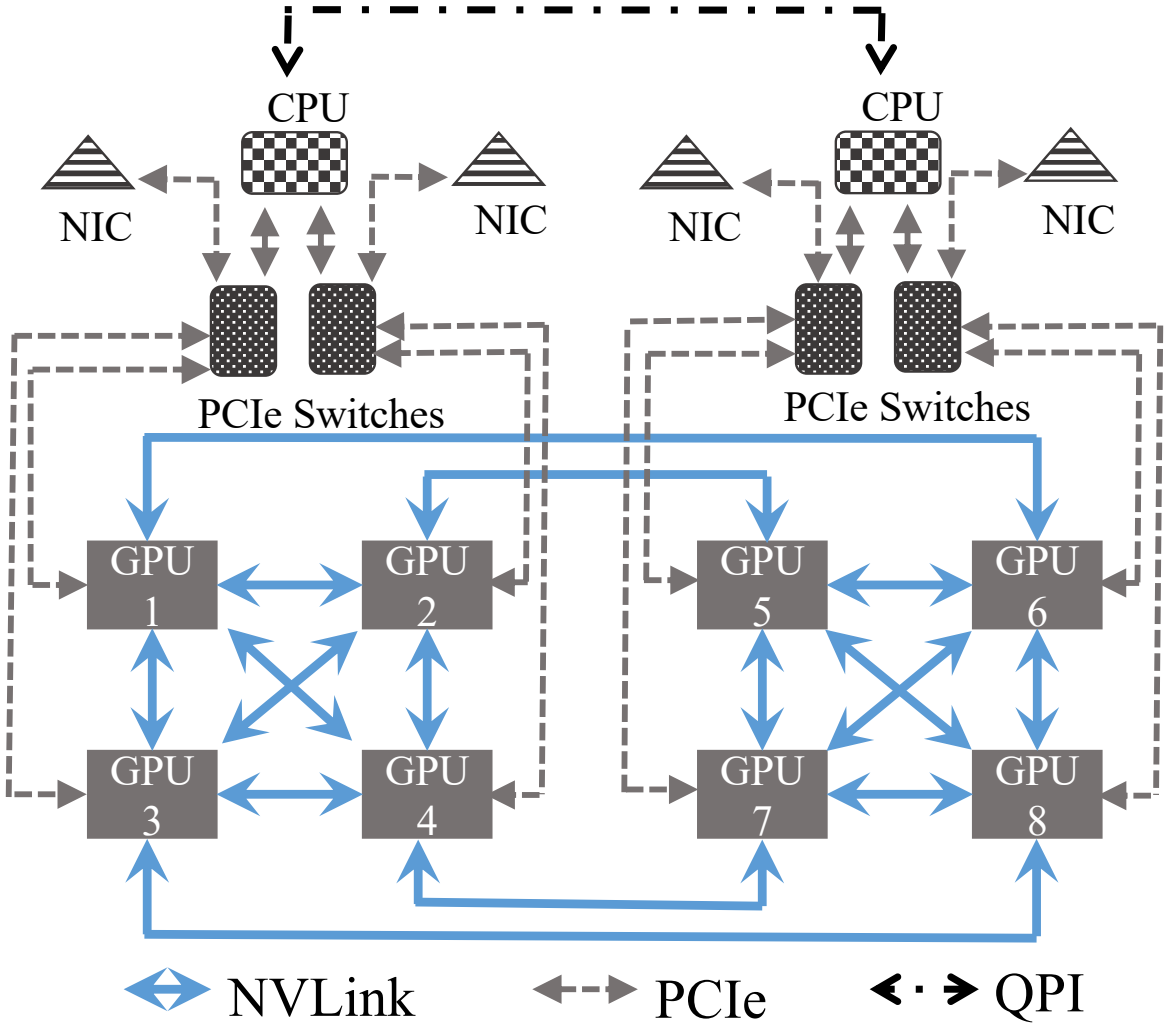

Peer-to-Peer Communication

- Unbalanced communication
- Bandwidth contention

Table 1. The speed (GBps) of common communication links (NV2 and NV1 mean 2 and 1 NVLinks between two GPUs)

Type	NV2	NV1	PCIe	QPI	IB	Ethernet
Speed	48.35	24.22	11.13	9.56	6.37	3.12

Table 2. The time (ms) peer-to-peer communication spends on different links for training a GCN layer with 8 GPUs

	Web-Google	Reddit	Wiki-Talk
NVLink	0.99	1.70	1.39
Others	6.20	18.1	6.13

Table 3. Attainable bandwidth (Gbps) of a GPU when there are different number of GPUs using the QPI link

Number of GPUs	1	2	3
Attainable bandwidth	9.50	5.12	3.34

System

- Read the graph and communication topology
- Graph partitioning
- Communication planning
- Training and communication on each device

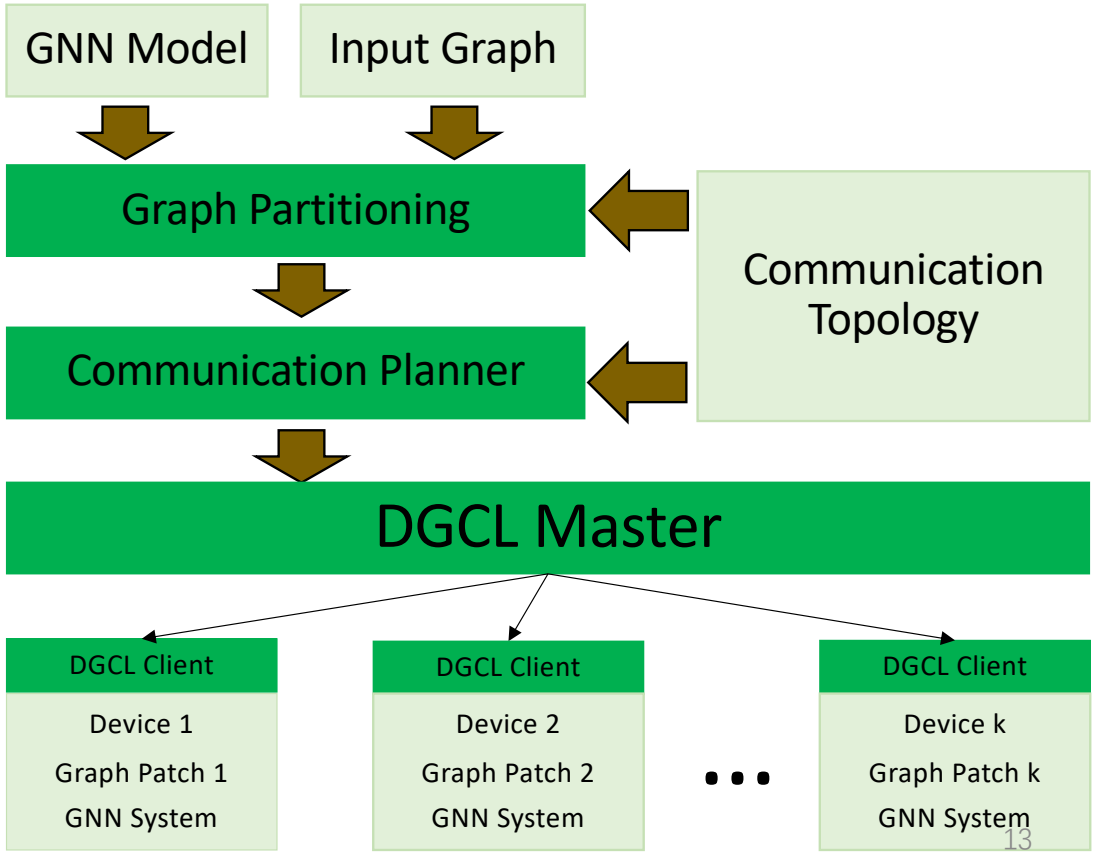

Design

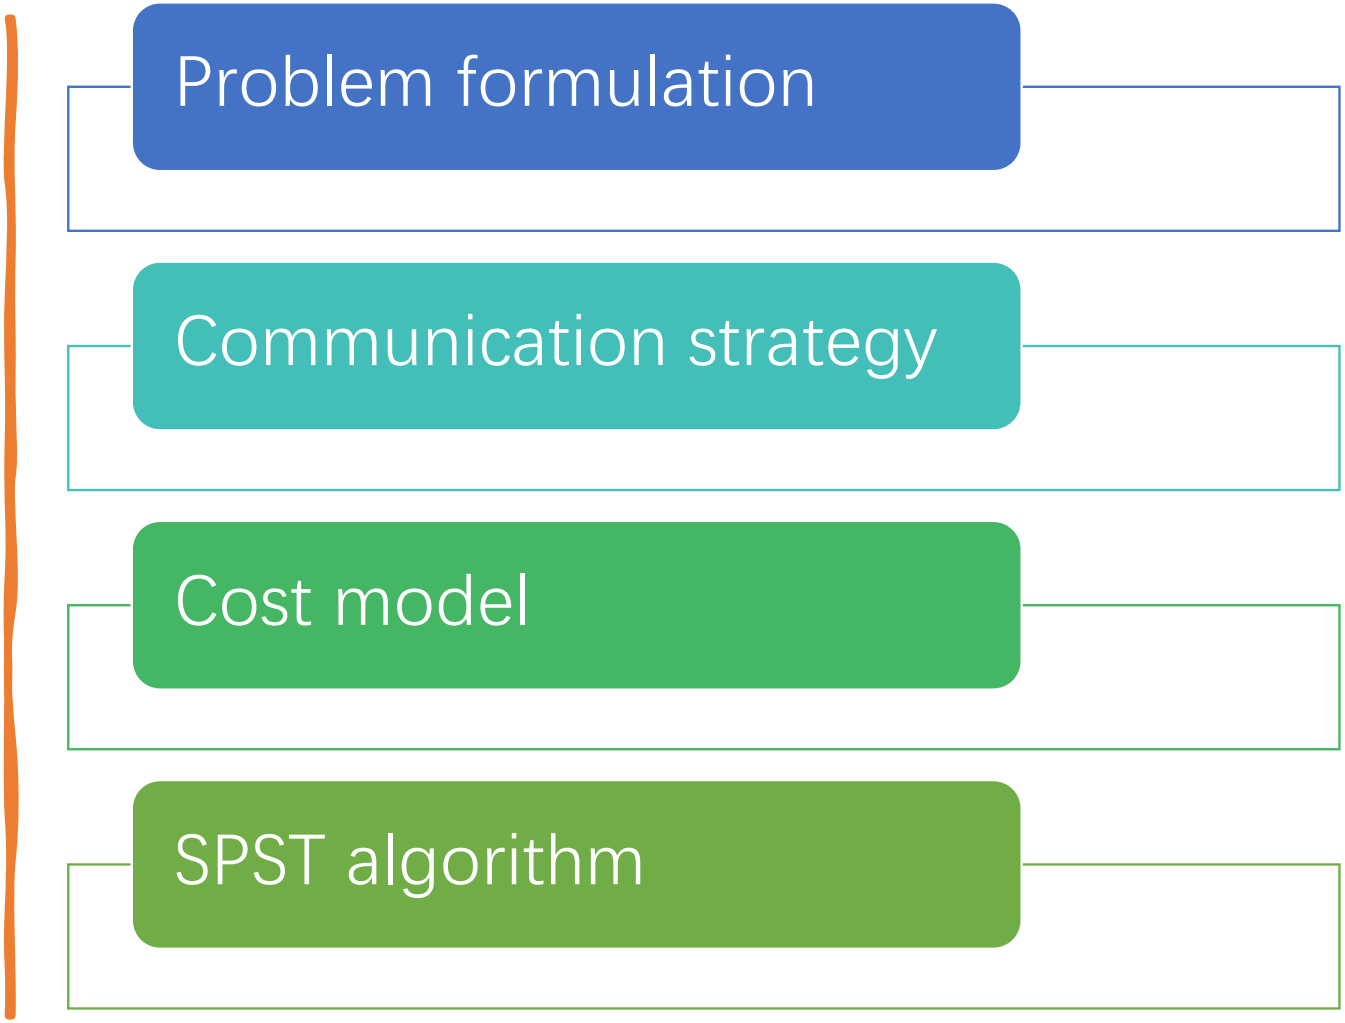
Problem formulation
Communication strategy
Cost model
SPST algorithm

Problem Formulation

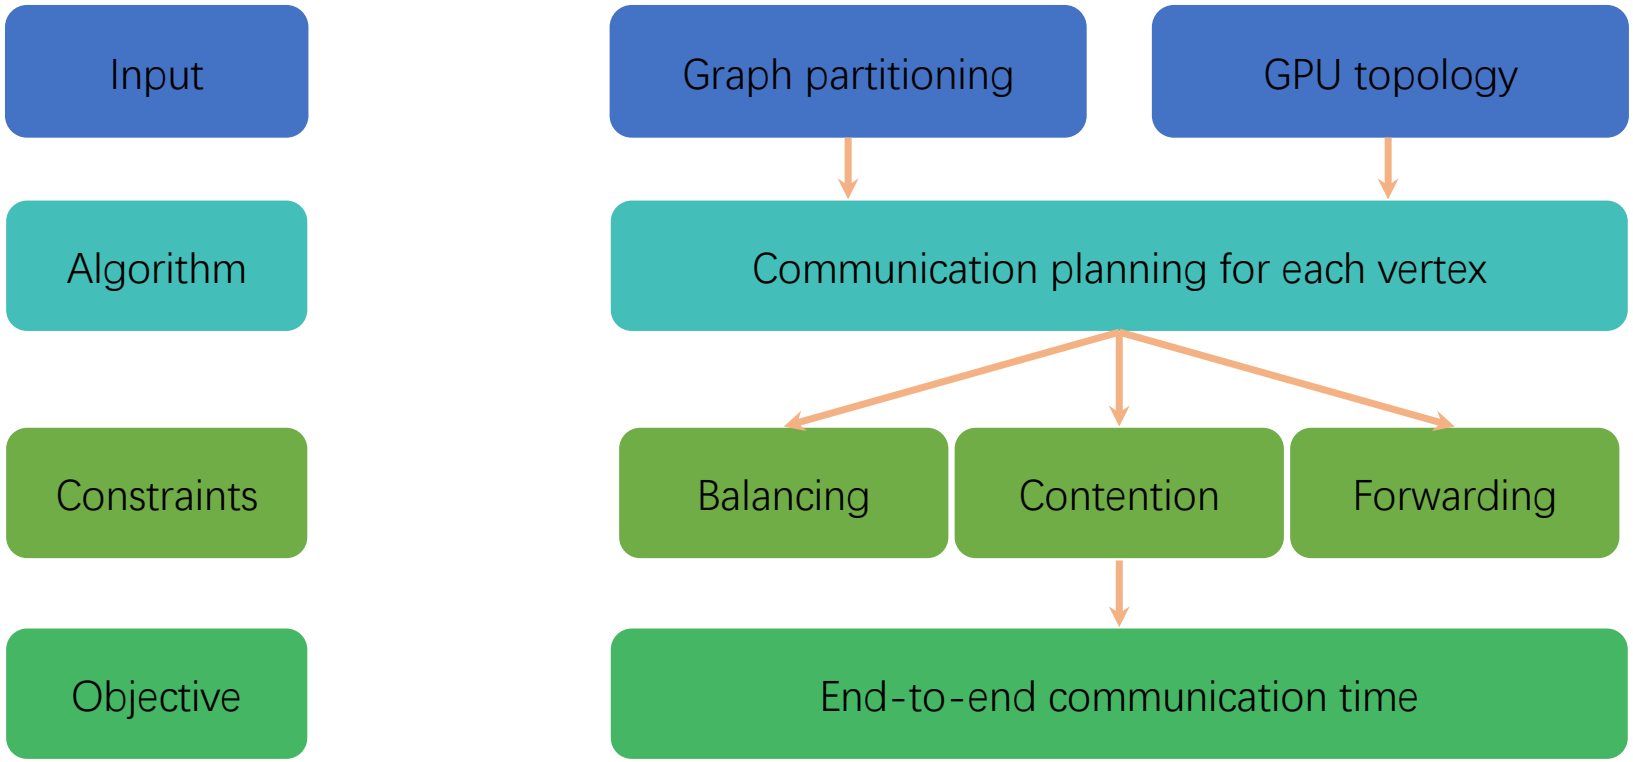

Communication Strategy

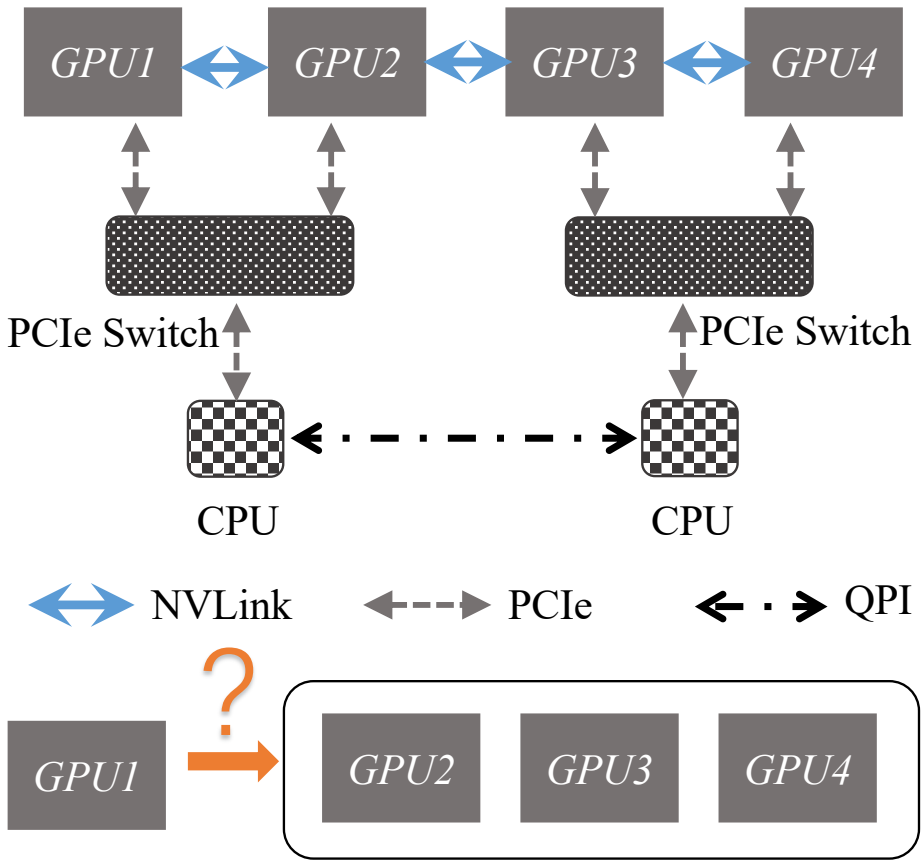

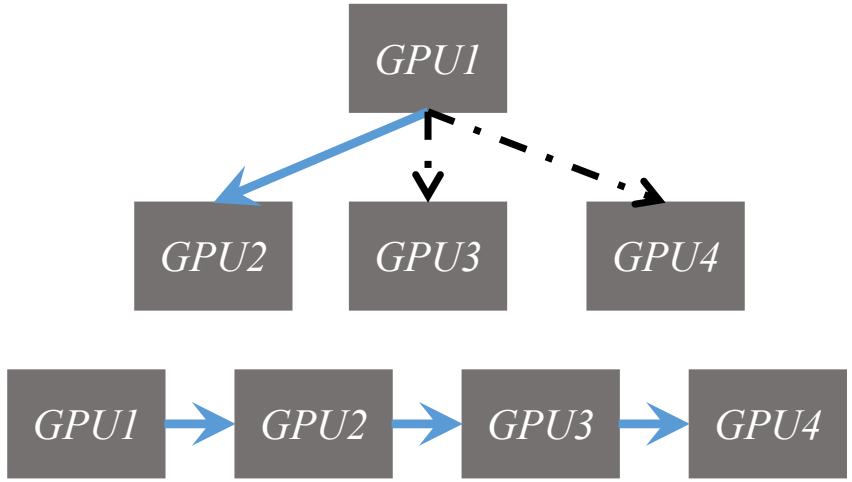

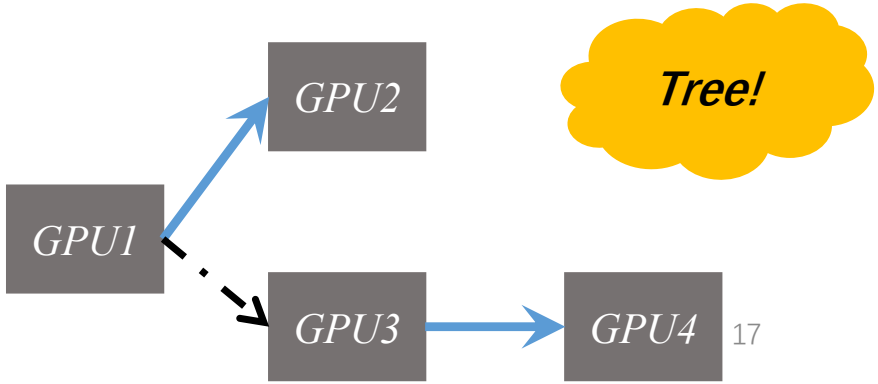

Cost Model

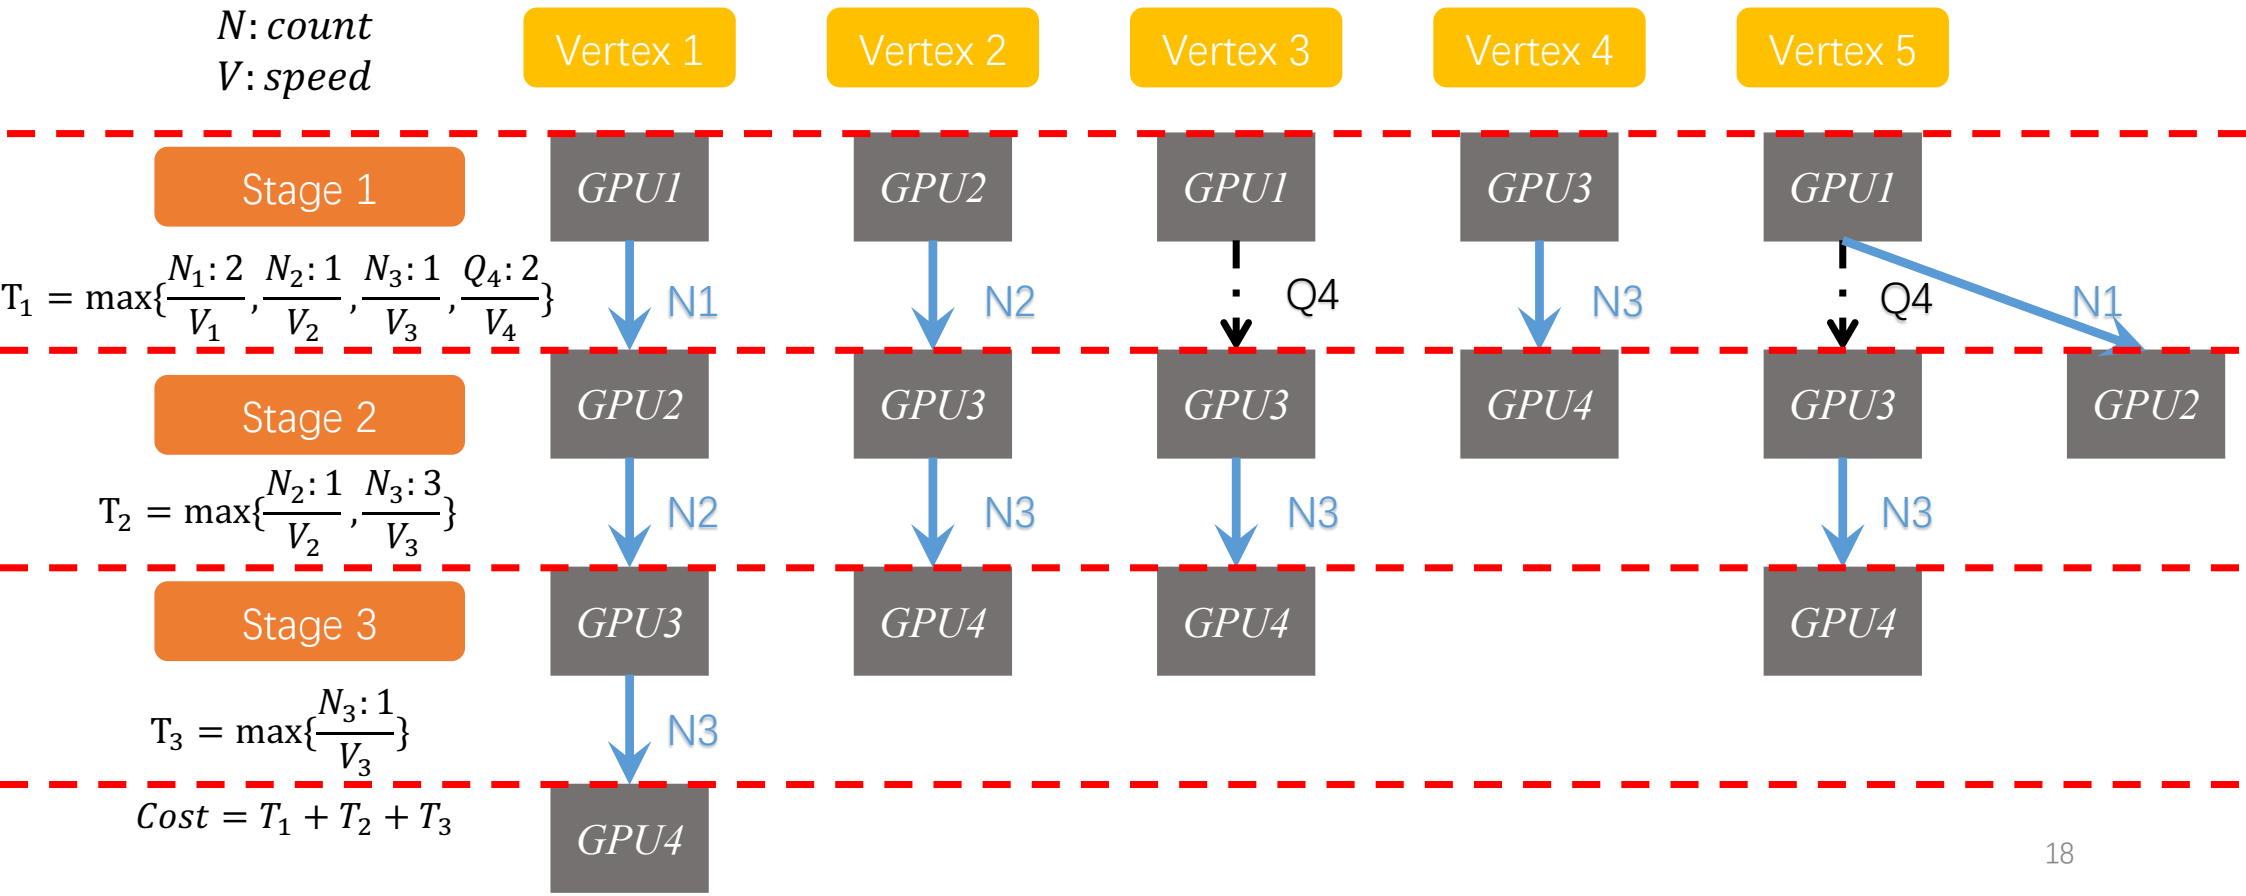
$N: count$
$V: speed$
Vertex 1
Vertex 2
Vertex 3
Vertex 4
Vertex 5
Stage 1
$T_1 = \max\{\frac{N_1:2}{V_1}, \frac{N_2:1}{V_2}, \frac{N_3:1}{V_3}, \frac{Q_4:2}{V_4}\}$
Stage 2
$T_2 = \max\{\frac{N_2:1}{V_2}, \frac{N_3:3}{V_3}\}$
Stage 3
$T_3 = \max\{\frac{N_3:1}{V_3}\}$
$Cost = T_1 + T_2 + T_3$
GPU1
GPU2
GPU1
GPU3
GPU1
N1
N2
Q4
N3
Q4
N1
GPU2
GPU3
GPU3
GPU4
GPU3
GPU2
N2
N3
N3
N3
GPU3
GPU4
GPU4
GPU4
N3
GPU4

Shortest Path Spanning Tree (SPST) Algorithm

Process vertices one by one

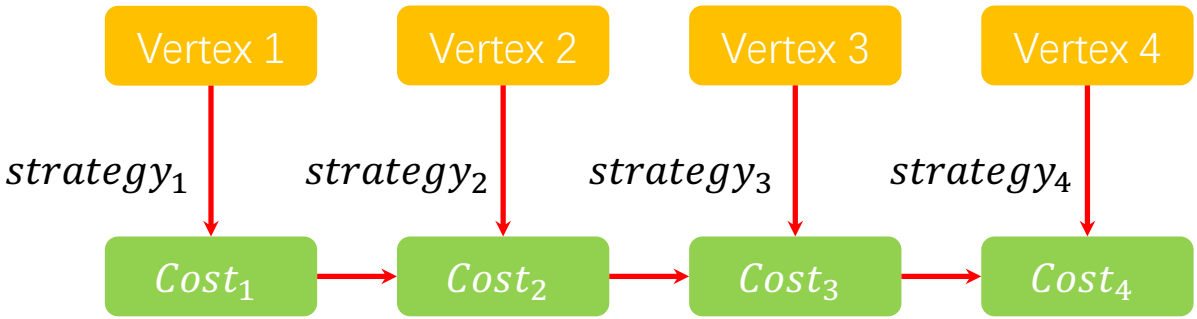

Heuristic: minimize $Cost_i$ when determining $strategy_i$

Stage 1

$N_1: c_1, N_2: c_2, N_3: c_3, \ldots$

Stage 2

$N_1: c_1, N_2: c_2, N_3: c_3, \ldots$

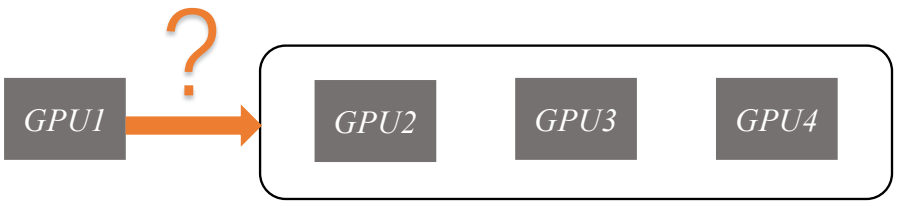

Shortest path w.r.t. cost increase

Shortest shortest path

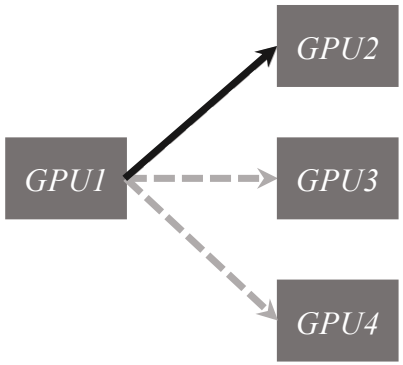

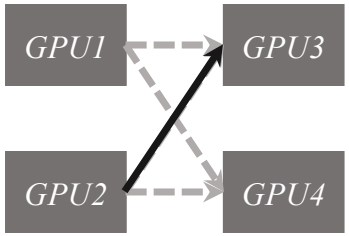

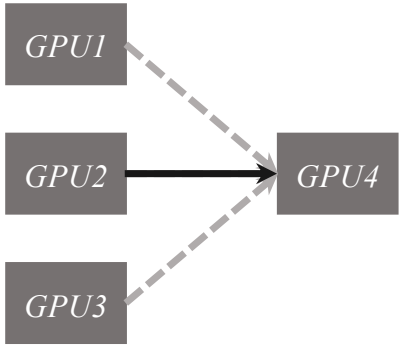

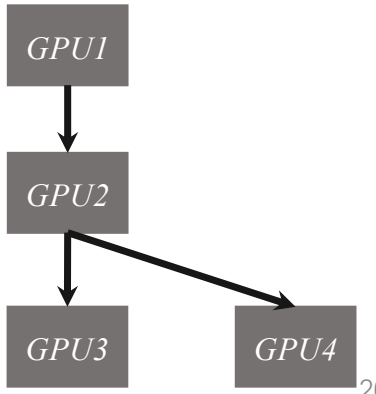

Experiment Setup

Hardware

- Two machines with 8 V100 GPU (NVLink)
- One machine with 8 1080ti GPU (No NVLink)

Baseline

- Peer-to-peer
- Replication: replicate k-hop vertices to eliminate communication
- Swap: using the main memory as shared memory

Dataset

- Reddit, Webgoogle, Wiki-Talk, Com-Orkut

GNN Model

- GCN, CommNet, GIN

End-to-End Performance on 8 V100 GPUs

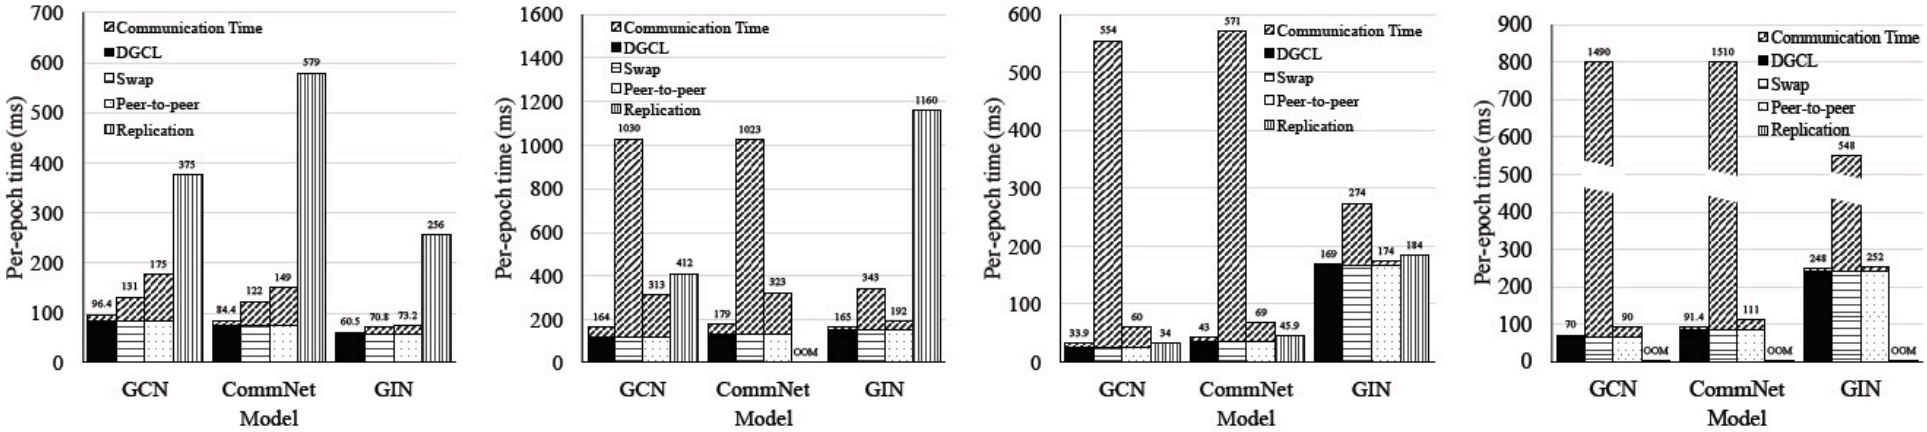

Figure 7. The per-epoch time and communication time for training the 3 GNN models on 4 datasets with 8 GPUs

	Peer-to-peer	Swap
Average communication time	4.5x	60x
Average per-epoch time	1.5x	7.4x

Performance on 8 1080ti GPUs

Table 6. Time (ms) for one *graphAllgather* operation in a hardware configuration without NVLink

	Reddit	Com-Orkut	Web-Google	Wiki-Talk
DGCL	14.3	128	7.84	5.86
Swap	14.5	1220	116	317
Peer-to-peer	17.9	179	8.72	8.51

Cost Model Accuracy

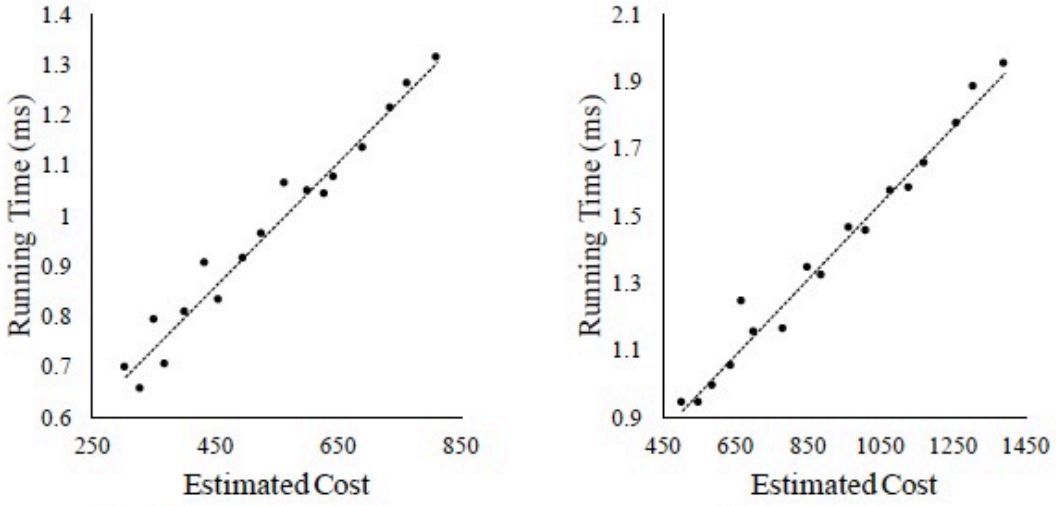

Figure 10. Relation between the model estimated communication cost and the actual communication time for one *graphAllgather* operation with 8 GPUs

Breakdown Analysis

Table 7. The breakdown of the communication time (ms) of one *graphAllgather* operation for DGCL with 8 GPUs

	NVLink	Others	Relative difference
Web-Google	0.787	0.821	4.32%
Reddit	1.16	1.07	7.41%
Com-Orkut	7.43	7.30	1.78%
Wiki-Talk	0.783	0.882	12.6%

Conclusions

- Start from a communication problem
- Analysis of peer-to-peer communications
- General communication library DGCL with cost model and planning algorithm
- DGCL outperforms other communications on different datasets, GNN models and GPU types.

Q & A