

Comprehensive Evaluation of GNN Training Systems: A Data Management Perspective

<p>Hao Yuan Northeastern University China yuanhao@stumail.neu.edu.cn</p>	<p>Yajiong Liu Northeastern University China liuyajiong@stumail.neu.edu.cn</p>	<p>Yanfeng Zhang Northeastern University China zhangyf@mail.neu.edu.cn</p>	<p>Xin Ai Northeastern University China aixin0@stumail.neu.edu.cn</p>
<p>Qiang Wang National University of Singapore, Singapore wang.qg@nus.edu.sg</p>	<p>Chaoyi Chen Northeastern University China chenchao@stumail.neu.edu.cn</p>	<p>Yu Gu Northeastern University China guyu@mail.neu.edu.cn</p>	<p>Ge Yu Northeastern University China yuge@mail.neu.edu.cn</p>

ABSTRACT

Many Graph Neural Network (GNN) training systems have emerged recently to support efficient GNN training. Since GNNs embody complex data dependencies between training samples, the training of GNNs should address distinct challenges different from DNN training in data management, such as data partitioning, batch preparation for mini-batch training, and data transferring between CPUs and GPUs. These factors, which take up a large proportion of training time, make data management in GNN training more significant. This paper reviews GNN training from a data management perspective and provides a comprehensive analysis and evaluation of the representative approaches. We conduct extensive experiments on various benchmark datasets and show many interesting and valuable results. We also provide some practical tips learned from these experiments, which are helpful for designing GNN training systems in the future.

PVLDB Reference Format:

Hao Yuan, Yajiong Liu, Yanfeng Zhang, Xin Ai, Qiang Wang, Chaoyi Chen, Yu Gu, and Ge Yu. Comprehensive Evaluation of GNN Training Systems: A Data Management Perspective. PVLDB, 17(6): 1241 - 1254, 2024. doi:10.14778/3648160.3648167

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/iDC-NEU/NeutronBench>.

1 INTRODUCTION

Graph Neural Networks (GNNs) are a class of Deep Neural Networks (DNNs) that can effectively process and analyze graph-structured data [59], and it is widely used in a variety of graph-related tasks [13, 41, 47, 65, 69]. With the increasing size of real-world graph data, it becomes difficult to train a large-scale GNN in a limited memory space. Distributed sample-based mini-batch GNN training emerges as a promising solution [12, 18, 25, 26, 31, 44, 62, 67, 72–75]. In this

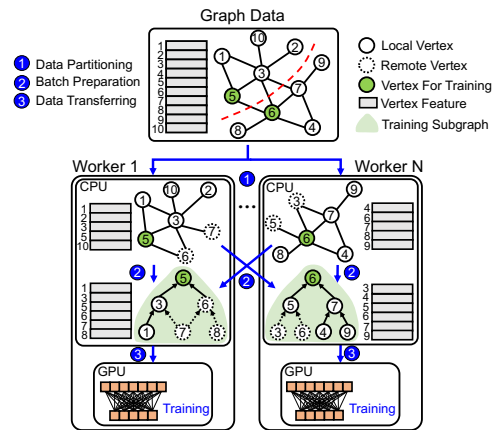


Figure 1: GNN training process.

training approach, the distributed environment offers ample computational and storage resources for training. The sampling reduces the size of the training graph [6, 13, 66], and the mini-batch training increases the frequency of model update to accelerate model convergence [74].

In a multi-node CPU and GPU heterogeneous training scenario, the training process of GNN and DNN can be divided into the following four steps. (1) Data partitioning: The input graph data is partitioned into multiple parts for distributed training. (2) Batch preparation: The data assigned to each worker is divided into multiple batches. (3) Data transferring: The prepared batch data is transferred to the GPU. (4) NN computation: NN computation is performed on the GPU side, including forward and backward propagation and parameter update.

Compared with DNNs, GNNs embody complex data dependencies between training samples. These data dependencies pose key challenges to the management of GNN data. Figure 1 depicts the end-to-end training process of GNNs from the data management perspective. Due to the dependence between data samples, GNN needs to frequently exchange data, which makes parallel GNN training a bottleneck. In the data partitioning step, since the data of the DNN is independent, it is only required to divide the data evenly into multiple partitions. However, for GNN training, the vertices need to access their multi-hop neighbor vertices, and the distribution of these neighbor vertices directly affects GNN training

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 17, No. 6 ISSN 2150-8097.
doi:10.14778/3648160.3648167

performance. In the batch preparation step, DNN only randomly divides the data samples into batches because of no dependency between data samples. GNN generates the training subgraphs by sampling for each batch. Complex data dependencies between data samples increase the sampling overhead. In the data transferring step, the low PCIe bandwidth between CPU and GPU results in high transfer overheads for both GNN and DNN. Compared to DNN, there are a large number of duplicate vertices and edges across different batches [25] due to the complex dependencies between vertices. Transferring such redundant vertices and edges wastes bandwidth resources severely [18, 25, 62], bringing opportunities for data transfer optimizations.

To illustrate the impact of handling data dependencies, Figure 2 shows a step-level time breakdown in both GNN and DNN training. For GNN training, we use a two-layer GCN [21] with a two-layer multi-layer perceptron model (MLP). For DNN training, to ensure fairness, we use a two-layer MLP with the same parameter settings as that in GCN. Since data partitioning is a preprocessing task performed only once before training, its runtime is ignorable. We can observe that the NN computation step consumes the majority of the runtime in DNN training, while the training in GNN only takes up a small portion. Furthermore, the overhead of the data management steps in GNN training (data partitioning, batch preparation, and data transferring) is notably higher than that in DNN training. Compared to DNN training, GNN training involves processing large-scale graph data. The data samples in GNN (i.e., graph vertices) exhibit complex dependencies, making the data management steps of GNN training more intricate and time-consuming than those of DNN training. Therefore, efficient data management is very important for GNN training.

Many GNN training systems have recently emerged to support efficient GNN training [4, 5, 11, 12, 17, 25, 26, 44, 45, 48, 50, 54, 55, 62, 63, 71–73]. These systems differ greatly in their targeted application scenarios and optimization techniques, especially from a data management perspective. The importance of data management in GNN training and the recent emergence of various data management techniques motivate us to study the impact of different optimizations and parameters in GNN training. In this paper, we review the training process of GNN from the data management perspective and provide a comprehensive analysis and evaluation of optimization techniques proposed in GNN training systems. The contributions of this paper can be summarized as follows.

- A taxonomy of data management techniques in GNN training.
- A comprehensive evaluation of data management techniques in GNN training.
- A summarization of lessons learned from our evaluation results.

We believe that our comprehensive analysis and evaluation results should be helpful for researchers and system developers to further improve the existing GNN training systems or design new GNN training systems.

Related Work. Recent surveys [24, 42, 46] provide an overview of the state-of-the-art GNN training systems and related optimization techniques. Unlike these survey papers, we categorize and summarize various optimizations for GNN training from a data management perspective, and provide a comprehensive experimental evaluation of them. Huang et al. investigate the performance gap

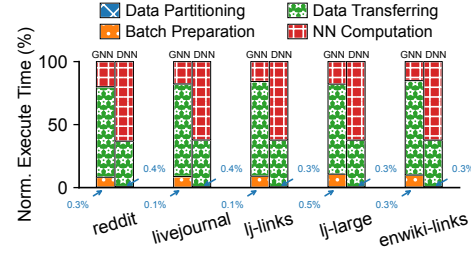


Figure 2: Time portion of different steps in GNN training and DNN training.

of GPU-based GNN systems and provide a set of optimizations to enhance performance [15]. A recent work [30] empirically studies graph partitioning for distributed GNN training. These two works only evaluate a single step of the entire GNN training process and do not separate data management from the entire training process for evaluation. Liu et al. provide a thorough experimental comparison of different sampling algorithms [27]. Wang et al. experimentally evaluate GNN training and inference on a single GPU [57]. These two works focus on single GPU GNN training and do not consider a distributed environment.

2 GNN TRAINING PROCESS

GNNs are a type of neural network that operates on graphs. The key idea behind GNNs is to learn vertex and edge representations that capture the structural information of the graph.

Graph Neural Network (GNN). Similar to traditional neural networks, the training process of GNNs includes forward propagation and backward propagation. In forward propagation, each vertex collects its neighbors’ features to generate the aggregation result using an aggregation function:

$$a_v^{(l)} = \text{AGGREGATE}(\{h_u^{(l-1)} | u \in N(v)\}), \quad (1)$$

where $h_u^{(l-1)}$ denotes the embedding of vertex u at $(l-1)$ -th layer, $N(v)$ denotes the incoming neighbors of vertex v , and $a_v^{(l)}$ denotes the aggregation result of vertex v at l -th layer. Specifically, h_v^0 denotes the input feature of vertex v . The aggregate functions can be sum, average, max/min, etc. Subsequently, the aggregated features are combined with the vertex’s own features using a combine function:

$$h_v^l = \sigma(\mathbf{W}^{(l)} \cdot \text{COMBINE}(h_v^{l-1}, a_v^l)), \quad (2)$$

where σ denotes a non-linear function (e.g., ReLU), $\mathbf{W}^{(l)}$ denotes the weight matrix that transforms the vertex embedding at $(l-1)$ -th layer. The combine function can be concatenation, element-wise multiplication, or summation. Then, a neural network (usually a multi-layer perceptron model, MLP) is used to update the features of each vertex. Finally, the output vertex features are compared to the ground truth labels to compute the loss. In the backward propagation, the loss is propagated through the neural network in the reverse direction, generating gradients that are used to update the model’s parameters.

Distributed Mini-Batch GNN Training. Due to the increasing size of real-world graph data, many emerging GNN training systems [12, 18, 25, 26, 31, 44, 62, 67, 72–75] adopt distributed mini-batch training. This approach splits the training vertices into multiple

mini-batches and conducts GNN training among workers. Figure 1 illustrates an example of the distributed mini-batch training for training a 2-layer GNN on vertices V_5 and V_6 . In the data partitioning step, the graph data are stored in a distributed graph storage, and each worker takes a subset of vertex samples for training. In the batch preparation step, each worker samples the L -hop neighbors of training vertices to generate *training subgraphs*. For example, as shown in Figure 1, the training subgraph of V_5 is sampled 2-hop subgraphs rooted from vertex V_5 (green shaded). The key idea of this step is to let each worker make the remote-dependent neighbors readily prepared locally before training starts. As a L -layer GNN, this approach needs to retrieve not only the vertex’s direct in-neighbors but also its $\{2, \dots, k\}$ -hop in-neighbors. For example, V_6 , V_7 , and V_8 are remote-dependent neighbors of V_5 , and worker 1 needs to fetch their feature vectors before training starts. Note that the sampled vertices may be deduplicated (V_7 in worker 1 is simultaneously sampled by V_3 and V_6). In the data transferring step, the training data, including sampled subgraphs and corresponding feature vectors, are loaded from the CPU to the local GPU through a PCIe interconnect. In the NN computation step, with the L -hop sampled subgraph locally, this approach performs normal forward/backward propagation layer-by-layer within a worker without any communication. Only the gradients need to be synchronized in the parameter update stage.

3 TAXONOMY OF DATA MANAGEMENT TECHNIQUES IN GNN TRAINING

There has been a surge of GNN training systems emerging in academia and industry to support efficient training. A variety of novel data management techniques have been proposed to accelerate the training of GNNs. This section provides a taxonomy of these techniques. As shown in Table 1, the recently emerged systems are listed from four aspects, including deployment platform, data partitioning, batch preparation, and data transferring.

Deployment Platform. From the perspective of deployment platforms, CPU-cluster [29, 67, 72, 73, 75], Multi-GPU [11, 19, 25, 26, 28, 31, 44, 52, 62], and GPU-cluster [3, 12, 17, 18, 39, 49, 55, 74] are the most commonly used. CPU-cluster is a network of multiple computers in which the CPU is the only computing component. Multi-GPU refers to the configuration of multiple GPUs inside a single compute node. GPU-cluster is similar to a CPU-cluster, but each compute node is equipped with one or more GPUs.

Data Partitioning. In terms of data partitioning, we categorize the data partitioning methods of existing GNN training systems into Hash, Metis-extend, and Streaming. Hash is a general graph partitioning method. Different mapping rules are used to meet various task requirements, such as hashing by vertices [3, 10, 17, 44, 45, 55, 67, 75] or hashing by edges [28, 29, 39, 48]. Metis is an iterative data partitioning method that uses community detection and iterative clustering for partitioning to minimize the amount of communication between partitions. Metis has been widely used in GNN training systems [44, 49, 75]. In addition, some GNN training systems [18, 73, 74] combine Metis with sample-based GNN training, and we label these systems as Metis-extend. Streaming data partitioning creates graph partitions in a pass over the sequence of edges, so the decision to assign edges to partitions is made on

the fly. By defining different score functions, streaming can achieve more fine-grained data partitioning [25, 26, 72]. For those systems that do not perform data partitioning, we mark them as N/A.

Batch Preparation. Regarding batch preparation, GNN training systems are categorized in terms of training method, whether they support sampling, and the adopted sampling method. Batch size denotes the number of vertices involved in training within each batch and also determines the frequency of model parameter updates. In terms of batch size, GNN training methods can be categorized into two types: full-batch [3, 17, 28, 29, 39, 45, 49, 51, 55, 70] and mini-batch [18, 26, 43, 60, 62, 72, 73, 75]. With the increase of graph data size in the real world, the full-batch training method suffers from inefficiency and poor scalability [26, 72, 74]. The sample-based mini-batch training method can effectively reduce the size of the training graph [6, 8, 13, 66], thus becoming the mainstream training approach [72]. We categorize the sampling methods into two types: fanout-based and ratio-based, which are used to determine the size of the sampling subgraph. The fanout-based sampling method samples a fixed number of neighbors [13, 73], while the ratio-based sampling method samples neighbors or the whole graph by a ratio [8, 49, 66, 73]. For systems that do not support sampling, we mark it as N/A in the sample method column.

Data Transferring. Regarding data transferring, they are classified in terms of transfer method, whether they support pipeline optimization [31, 44, 48, 74], and whether they support GPU cache optimization [25, 26, 62]. The data transfer between CPU and GPU can be summarized into two categories: extract-load and GPU direct access. In mini-batch training, after the CPU obtains the sampled subgraph through sampling, it first extracts the vertex features from the feature matrix according to the sampled subgraph. It then loads the sampled subgraph and vertex features to the GPU for training. We call this process "Extract-Load" [25, 62]. GPU direct access [18, 31] is another data transfer method that uses implicit data transfer [33, 34] to enable the GPU to access the CPU’s memory directly, avoiding explicit data transfers and additional feature extraction operations. Pipelining decomposes a task into multiple stages and allows these stages to be executed in parallel on different processors. GPU caching is a cache-based data reuse method [19, 25, 26, 39, 44] that effectively reduces the data transfer overhead between CPU and GPU by caching vertex features in GPU memory. For systems that do not support GPU training, we mark it as N/A in the transfer method column.

4 EXPERIMENTAL SETUP

Environments. Our experiments are conducted on the Aliyun ECS cluster with 4 GPU nodes. Each node (ecs.gn6i-c40g1.10xlarge) is equipped with 40 vCPU, 155 GiB DRAM, Intel Xeon (Skylake) Platinum 8163, and an NVIDIA Tesla T4 GPU with 68 SMs, 2560 cores, and 16GB GDDR6 global memory. The host side is running Ubuntu 20.04 LTS OS. The network bandwidth is 10 Gbps/s. Libraries CUDA 11.3, OpenMPI-3.3.2, Pytorch [38] v1.9 backend, and cuDNN 8.4 are used. All the codes are compiled with O3 optimization.

GNN Models. The models used in our experiments are two representative GNN models, Graph Convolutional Network (GCN) [21] and GraphSage [13]. There are also some complex GNN models, e.g., MixHop [1]. Unlike traditional GNN models, it directly aggregates

Table 1: Summary of representative GNN training system and data management techniques.

Year	Systems	Deploy Platform	data partitioning	Batch Preparation			Data Transferring		
			Partition Method	Train Method	Sample	Sample Method	Transfer Method	Pipeline	Cache
2019	DGL [52]	Multi-GPU	N/A	Mini-batch	✓	Fanout-based	Extract-Load	✓	✗
2019	PyG [11]	Multi-GPU	N/A	Mini-batch	✓	Fanout-based	Extract-Load	✗	✗
2019	AliGraph [75]	CPU-cluster	Hash/Metis /Streaming	Mini-batch	✓	Fanout-based/ Ratio-based	N/A	✗	✗
2019	NeuGraph [28]	Multi-GPU	Hash	Full-batch	✗	N/A	Extract-Load	✗	✗
2020	AGL [67]	CPU-cluster	Hash	Mini-batch	✓	Fanout-based	N/A	✗	✗
2020	DistDGL [73]	CPU-cluster	Metis-extend	Mini-batch	✓	Fanout-based/ Ratio-based	N/A	✓	✗
2020	ROC [17]	GPU-cluster	Hash	Full-batch	✗	N/A	Extract-Load	✗	✗
2020	PaGraph [25]	Multi-GPU	Streaming	Mini-batch	✓	Fanout-based	Extract-Load	✗	✓
2021	P3 [12]	GPU-cluster	Hash	Mini-batch	✓	Fanout-based	Extract-Load	✗	✗
2021	DistGNN [29]	CPU-cluster	Hash	Full-batch	✗	N/A	N/A	✗	✗
2021	DGCL [3]	GPU-cluster	Hash	Full-batch	✗	N/A	Extract-Load	✗	✗
2021	Dorylus [45]	Serverless	Hash	Full-batch	✗	N/A	N/A	✓	✗
2021	Pytorch-direct [31]	Multi-GPU	N/A	Mini-batch	✓	Fanout-based	GPU direct access	✓	✗
2022	GNNLab [62]	Multi-GPU	N/A	Mini-batch	✓	Fanout-based	Extract-Load	✓	✓
2022	ByteGNN [72]	CPU-cluster	Streaming	Mini-batch	✓	Fanout-based	N/A	✓	✗
2022	BNS-GCN [49]	GPU-cluster	Metis	Full-batch	✓	Ratio-based	Extract-Load	✗	✗
2022	DistDGLv2 [74]	GPU-cluster	Metis-extend	Mini-batch	✓	Fanout-based	Extract-Load	✓	✗
2022	NeutronStar [55]	GPU-cluster	Hash	Full-batch	✗	N/A	Extract-Load	✗	✗
2022	Sancus [39]	GPU-cluster	Hash	Full-batch	✗	N/A	Extract-Load	✗	✓
2022	SALIENT [19]	Multi-GPU	N/A	Mini-batch	✓	Fanout-based	GPU direct access	✓	✗
2023	Betty [63]	GPU-only	Metis	Mini-batch	✓	Fanout-based	Extract-Load	✗	✗
2023	MariusGNN [48]	GPU-only	Hash	Mini-batch	✓	Fanout-based	Extract-Load	✓	✗
2023	Legion [44]	Multi-GPU	Metis/Hash	Mini-batch	✓	Fanout-based	Extract-Load	✓	✓
2023	SALIENT++ [18]	GPU-cluster	Metis-extend	Mini-batch	✓	Fanout-based	GPU direct access	✓	✓
2023	BGL [26]	Multi-GPU	Streaming	Mini-batch	✓	Fanout-based	Extract-Load	✓	✓

Table 2: Dataset description.

Dataset	V	E	#F	#L	#hidden
Reddit [13]	232.96K	114.85M	602	41	128
OGB-Arxiv [35]	169.34K	2.48M	128	40	128
OGB-Products [37]	2.45M	126.17M	100	47	128
OGB-Papers [36]	111.06M	1.6B	128	172	128
Amazon [66]	1.57M	264.34M	200	107	128
LiveJournal [61]	4.85M	90.55M	600	60	128
Lj-large [32]	7.49M	232.1M	600	60	128
Lj-links [22]	5.2M	205.25M	600	60	128
Enwiki-links [23]	13.59M	1.37B	600	60	128

multiple-hop neighbors at each layer. Despite differences in their computation methods, these models share a common L-hop neighbor access pattern. Therefore, we only use GCN and GraphSage as the default GNN models for experiments. Following the suggestion from prior works [25, 52, 62], the dimension of the hidden layers is set to 128. For all experiments, the default fanout is (25, 10), and the default batch size is 6000 for all experiments.

Datasets. As listed in Table 2, we use nine real-world graph datasets. Among them, OGB-Arxiv, Amazon, LiveJournal, Lj-large, Lj-links, and Enwiki-links have skewed degree distributions. Reddit, Amazon, and Enwiki-links have average degrees relatively higher than other large real-world graphs. OGB-Arxiv, OGB-Products, and OGB-Paper have low feature dimensions. Reddit and OGB-Products are two dense graphs with a high clustering coefficient. In the data partitioning and batch preparation experiments, we use datasets with ground-truth labels, including Reddit, OGB-Arxiv, OGB-Products, and Amazon, because we need to study their effect on the model accuracy. Regarding the data transferring experiments, we use different types of datasets, including a social network (Reddit), a citation network (OGB-Papers), co-purchasing

networks (OGB-Products and Amazon), livejournal communication networks (Livejournal, Lj-large, and Lj-links), and a wikipedia links network (Enwiki-links). These datasets have different characteristics, e.g., average degree, power-law, and feature dimension, which are commonly used by existing GNN training systems in performance evaluation [25, 31, 62].

Evaluation System. For data partitioning experiments, we use DistDGL [73], which is commonly used by existing GNN training systems in performance evaluation experiments, as our evaluation system. For batch preparation and data transferring experiments, we use NeutronStar [55] as our evaluation system. NeutronStar provides a rich programming interface to facilitate the implementation of different performance optimizations.

5 DATA PARTITIONING

5.1 Goals and Challenges

In a distributed setting, the goal of traditional graph partitioning is typically twofold: minimizing communication and balancing the computational load. In traditional iterative graph computations, typically *all vertices* and their *1-hop neighbors* (connected through edges) are involved in the computation (e.g., PageRank [2]). However, in GNN training, only a portion of vertices with ground-truth labels (i.e., training vertices) and their L-hop subgraphs are involved. For example, in a 2-layer GNN, as shown in Figure 3c, A and B represent the labeled vertices. The light blue and light green vertices are their 2-hop neighbors. The dashed lines indicate vertices and edges that are not involved in training.

Due to these differences, traditional graph partitioning methods cannot be directly applied to GNN training. For traditional graph computation, the partitioning in Figure 3a is optimal because it not only balances the number of vertices and edges but also leads to

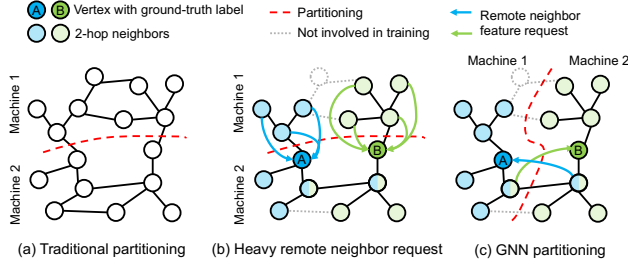


Figure 3: Traditional partitioning vs. GNN partitioning.

the minimum number of cutting edges. However, this partitioning is not optimal for GNN training due to the following two issues: (1) Severe computational load imbalance. The graph partitioner assigns all labeled vertices (A and B) to machine 2, leaving machine 1 idle, thus wasting computational resources. (2) High communication volume. The minimum edge-cut graph partitioning method does not consider the distribution of L-hop neighbors of vertices. The neighbors of labeled vertices are scattered in different partitions, which results in heavy remote neighbor requests. For example, as shown in Figure 3b, for the GNN training, using the minimum edge-cut partitioning method from Figure 3a would result in approximately half of the 2-hop neighbors of training vertices (A and B) being distributed across different partitions. This would lead to significant communication overhead, especially considering that the vertices of a GNN have high-dimensional feature representations. In contrast, the partitioning shown in Figure 3c balances the workload and, at the same time, maximizes the localization of L-hop neighbors of training vertices, thereby reducing inter-partition communication.

Different from traditional data partitioning, GNN data partitioning has new goals. We summarize these goals as follows:

- **Goal 1: Minimize communication.** Partitioning label vertices and their L-hop neighbors into the same partition allows the L-hop sampled subgraphs to be distributed as locally as possible, thus avoiding extensive remote data requests.
- **Goal 2: Balance computational load.** Making labeled vertices and their L-hop neighbors evenly distributed across partitions allows a balanced computational load.
- **Goal 3: Minimize total computational load.** Making the training L-hop subgraphs overlap with each other in a partition allows reusing the NN computation results, which helps reduce the total computational load.
- **Goal 4: Balance communication load.** Communication load imbalance is a common challenge faced by both traditional graph computations and GNNs. However, compared to traditional graph computations, GNNs require frequent access to the L-hop neighbors of vertices, and the vertices of GNNs have high-dimensional feature representations. As a result, the problem of communication load imbalance is even more pronounced in GNNs. Making the remote neighbors of labeled vertices evenly across different partitions can balance the communication load.

5.2 Existing Methods

The graph partitioning methods used in the GNN systems can be categorized into the following three types:

(1) Hash. Hash is a general graph partitioning method that can satisfy various task requirements by defining different mapping rules,

Table 3: Summary of evaluated partitioning methods.

Method	Strategy	Representative System	G1	G2	G3	G4
Hash	Randomly assign vertices or edges.	P3 [12]	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Metis-V	Extend Metis by adding constraints on training vertex masks.	N/A	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Metis-VE	Extend Metis by adding constraints on training vertex masks and vertex degrees.	DistDGL [73]	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Metis-VET	Extend Metis by adding constraints on training/validation/test vertex masks and vertex degrees.	SALIENT++ [18]	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Stream-V	Assign vertex v to a partition P that has the most edges connected to v , while balancing the number of train vertices and caching L-hop neighbors.	PaGraph [25]	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Stream-B	Assign a block B of vertices to a partition P that has the most edges connected to B , while balancing the number of train/val/test vertices.	ByteGNN [72]	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

such as hashing by vertices [3, 10, 17, 44, 45, 55, 67, 75] or hashing by edges [28, 29, 39, 48]. P3 [12] employs a hash-based graph partitioning method that randomly assigns vertices to different partitions, achieving a balance in computational and communication loads (goal 2 and goal 4). However, hash partitioning does not consider the labeled vertices in GNNs and their L-hop neighbors, so it cannot achieve the other two goals of GNN graph partitioning.

(2) Metis-extend. Metis [20] partitions the graph with the goal of minimum edge cuts. In addition, it supports a variety of constraint mechanisms to extend Metis to meet different task requirements. We refer to this constrained Metis partitioning method as Metis-extend, which is used in DistDGL [73], and SALIENT++ [18]. The Metis-extend method effectively achieves the goals of GNN graph partitioning. Firstly, Metis employs a clustering algorithm for partitioning, ensuring that the neighbors of labeled vertices are allocated together, thus minimizing both computational and communication loads (goal 1 and goal 3). Secondly, Metis is extended by adding constraints with vertex masks to balance the number of labeled vertices (goal 2). In addition, Metis can be extended by adding constraints on vertex degrees to balance the number of edges across partitions, alleviating the load imbalance in computation and communication (goal 2 and goal 4).

(3) Streaming. Unlike traditional graph partitioning methods, streaming partitioning does not require storing the entire graph data but dynamically partitions the input vertices or edges. Streaming partitioning can flexibly support various graph partitioning tasks by setting different score functions when assigning vertices or edges. PaGraph [25] and ByteGNN [72] both employ streaming graph partitioning methods. When assigning vertices, they prioritize assigning vertices to the partitions with the highest number of connected edges, aiming to minimize communication overhead (goal 1). Additionally, they use a factor to balance the number of label vertices (goal 2). However, these two streaming graph partitioning methods do not consider the density of partitioned graphs and the distribution of L-hop neighbors and thus suffer from high computational and imbalanced communication loads.

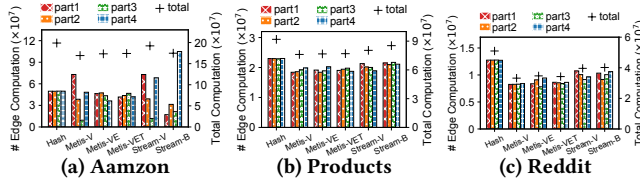


Figure 4: Computational load of different partitionings.

5.3 Evaluation Results

Evaluated Methods. We conduct experimental analysis on the three types of data partitioning, i.e., Hash, Metis-extend, and Streaming. Table 3 summarizes the partitioning methods and the systems that adopt these methods. For the hash partitioning method, P3 [12] is a representative system that randomly partitions the vertices. For the Metis-extend partitioning method, DistDGL [73] and SALIENT++ [18] are two representative systems. DistDGL balances the number of training vertices and edges in partitions, which we refer to as Metis-VE. SALIENT++ not only balances the training vertices and edges but also considers balancing the number of validation and test vertices, ensuring high performance in the inference stage, and we refer to this partitioning method as Metis-VET. In addition, we also compare the Metis-extend data partitioning method that only balances training vertices, which we refer to as Metis-V. Although Metis-V is not employed by any existing GNN system for graph partitioning, we use Metis-V as a control group to comprehensively evaluate the performance of Metis-extend under different constraints. PaGraph [25] and ByteGNN [72] are two representative systems for the streaming partitioning method. PaGraph performs partitioning at the vertex level, while ByteGNN partitions at the block level. We refer to these two partitioning methods as Stream-V and Stream-B, respectively.

5.3.1 Computational Workload Balance. We first conduct an experimental analysis from the perspective of computational workload balancing. The computational workload comprises two parts: sampling and training. Regarding sampling, for remote-dependent vertices, sampling requests are sent to remote machines to be executed, and then the sampling results are returned. Therefore, the sampling workload on each machine includes both processing local sampling requests and processing remote sampling requests. Regarding training, GNN training includes two parts: graph aggregation and NN computation. Since the graph aggregation dominates the overall computational cost, we only count the number of graph aggregations in the training part.

As illustrated in Figure 4, regarding the hash partitioning method, since vertices are randomly assigned to different partitions with equal probability, it exhibits the most balanced computational workload. However, due to the neglect of vertex dependencies during partitioning, hash partitioning results in the highest total computational workload. Regarding the Metis-extend partitioning, Metis-V only balances the number of training vertices, ignoring the distribution of L-hop neighbors of training vertices, resulting in load imbalance issues. By introducing vertex degree constraints to extend Metis, Metis-VE and Metis-VET alleviate the load imbalance problem observed in Metis-V. Metis-extend uses clustering for graph partitioning. When sampling vertices within clusters, different sampled vertices can share a large number of repeated neighbors [68]. Hence, the higher the clustering degree in graph partitioning, the

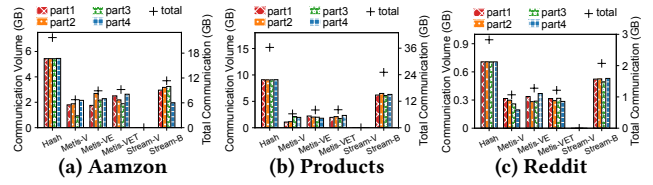


Figure 5: Communication load of different partitionings.

lower the computational load. Metis-VE and Metis-VET introduce additional constraints, which reduce the degree of clustering in the Metis partitioning. Therefore, the total computational load is higher than that of Metis-V. Regarding the streaming partitioning, Stream-V and Stream-B aim to minimize cross-partition edges while balancing the number of labeled vertices when allocating vertices or blocks. This partitioning method performs well on non-power-law graphs but suffers from computational load imbalance on power-law graphs (i.e., Amazon) because it ignores the distribution of L-hop neighbors during partitioning. To illustrate this point, we first compute the clustering coefficient of each partition graph to represent the density of the partitioned graphs. Then, we compute the variance of the clustering coefficients of the partition graph to evaluate their distribution. The results show that the variance of the clustering coefficient [58] of the Hash partition graph is only 3.6×10^{-6} , while the variances of Stream-V and Stream-B are 0.01 and 0.03, respectively. Therefore, the imbalance of partitioned graphs makes Stream-V and Stream-B suffer from computational load imbalance.

5.3.2 Communication Workload Balance. We also conduct an experimental analysis of these graph partitioning methods from the perspective of communication workload balancing. The communication data comprises two parts: remote sample subgraphs and corresponding vertex features. As shown in Figure 5, regarding the hash partitioning, it exhibits the most balanced communication workloads, but the overall communication volume is also the highest. This is because hash partitioning ignores vertex dependencies during partitioning, resulting in a significant number of remote data requests and thus higher network communication overhead. Regarding the Metis-extend partitioning, it can achieve the goal of minimizing communication (i.e., goal 1) by using a cluster-based partitioning method. However, it does not consider how to balance the communication load during the partitioning process and therefore suffers from an imbalanced communication load. Metis-V has the lowest total communication volume because its partitioned graph achieves the best clustering. By introducing vertex degree constraints to balance the number of edges in the partitioned graph, Metis-VE and Metis-VET alleviate the communication load imbalance problem but increase the total communication volume. Regarding streaming partitioning, since Stream-V caches the graph data of L-hop neighbors of training vertices, no communication is required. Stream-B assigns a block to the partition with the most connections to it, which reduces the overall communication but ignores communication load balancing.

5.3.3 Partitioning Time vs. Training Time. Although graph partitioning is a preprocessing task that only needs to be performed once before training, in practical applications, the time required cannot be ignored. In order to analyze the cost of graph partitioning, we show the proportion of graph partitioning time and training

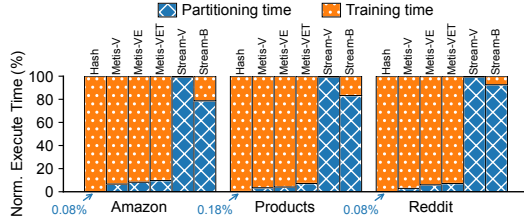


Figure 6: Partitioning time vs. training time.

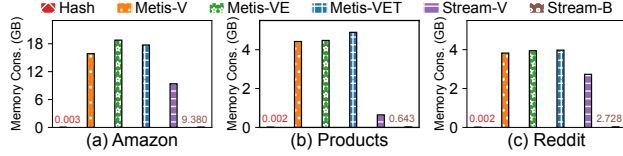


Figure 7: Memory consumption of vary partitioning methods.

time on different datasets. As shown in Figure 6, regarding hash partitioning, it is the fastest and takes the smallest proportion of overall training time (only 0.11% on average), resulting in the lowest partitioning overhead. Regarding Metis-extend partitioning, the graph partitioning time portion for Metis-V, Metis-VE, and Metis-VET is 4.33%, 6.11%, and 7.98%, respectively. As more constraints are introduced, the partitioning overhead of Metis-extend gradually increases. Overall, the partitioning overhead of Metis-extend is deemed acceptable, constituting less than 10% of the total training time. Regarding streaming partitioning, it exhibits the longest graph partitioning time. The graph partitioning time portion for Stream-V and Stream-B is 99.36% and 84.87%, respectively. This is because streaming-based partitioning involves extensive set intersection computations during vertex assignment, incurring substantial computational overhead. Additionally, due to vertex dependencies and the constraints of streaming graph partitioning, it cannot benefit from parallel acceleration.

5.3.4 Partitioning Memory Consumption. Memory consumption is also a factor that should be considered when choosing a graph partitioning method. We conduct experiments to evaluate the memory consumption of partitioning methods. As shown in Figure 7, hash partitioning has the least memory consumption because it only needs to randomly assign vertices to different partitions and does not need to maintain the state of the partitions. Streaming partitioning requires maintaining the state of the partition (e.g., allocated vertices) to guide the allocation of the next vertex or edge, so its memory consumption is higher than hash partitioning. In addition, the memory consumption of Stream-V is much higher than Stream-B. This is because Stream-V needs to obtain the L-hop neighbors of vertices when assigning vertices, which results in large memory consumption. The Metis-extend partition has the highest memory consumption because it requires iteratively performing coarsening and uncoarsening steps, which results in severe memory overhead.

5.3.5 Effect to Accuracy and Convergence Speed. In distributed training, each partition stores a portion of the graph data, and the training subgraphs are constructed only for local training vertices. Different graph partitioning methods result in various data distributions among partitions. For instance, in hash-based partitioning,

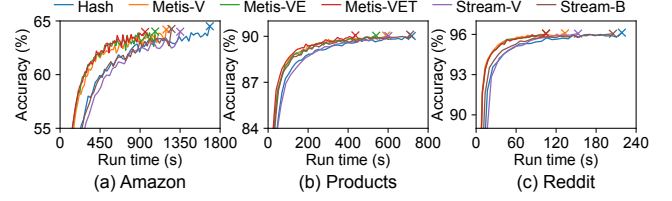


Figure 8: Accuracy and convergence speed of vary partitioning methods.

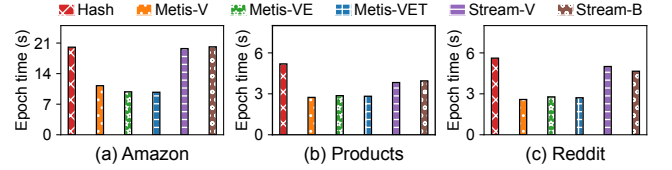


Figure 9: Per-epoch runtime under different partition methods.

Table 4: Model accuracy under different partition methods.

	Hash	Metis-V	Metis-VE	Metis-VET	Stream-V	Stream-B	Diff.
Reddit	96.2%	96.2%	96.2%	96.1%	96.4%	96.2%	$\pm 0.3\%$
Products	90.1%	90.3%	90.2%	90.2%	90.5%	90.2%	$\pm 0.4\%$
Amazon	64.5%	64.2%	64.4%	64.7%	65.1%	64.3%	$\pm 0.9\%$

vertices in each partition are randomly distributed, while in Metis-extend partitioning, vertices within partitions are distributed in a clustered manner. Hence, a question arises: *does graph partitioning impact the accuracy and convergence speed of model training?*

We compare the accuracy and convergence speed of different graph partitioning methods on three datasets. For each graph partitioning method, we train the model to converge and record the highest validation accuracy and run time. Regarding model accuracy, as shown in Table 4, graph partitioning does not affect the highest accuracy. Compared with DNN, there are complex dependencies between GNN data. Even when data is divided into multiple partitions, data exchange is still necessary between partitions due to the dependencies among vertices. Therefore, graph partitioning does not result in the loss of complete graph information and thus does not affect the final model accuracy. Regarding convergence speed, as shown in Figure 8, hash partitioning exhibits the slowest convergence speed, followed by Stream-B. Figure 9 further illustrates that Hash, Stream-V, and Stream-B have the longest per-epoch runtime, indicating that they require more time for the same number of iterations. Regarding the three Metis-extend graph partitioning methods, Metis-VET achieves the fastest convergence speed, followed by Metis-VE and Metis-V. Figure 9 shows that the per-epoch runtime for each graph partitioning method is similar, but the convergence speed of Metis-VET is significantly faster than the other two methods. We think that this is related to the characteristics of Metis. In mini-batch training, the training process needs to select a batch of vertices from local training vertices. Since Metis employs a clustering method to partition densely connected vertices together, this reduces the randomness in batch vertex selection, which in turn affects the model’s convergence speed. Since Metis-VET has the most constraints, it has the least clustering effect, resulting in the fastest convergence speed.

5.4 Lessons Learned

- (1) Existing graph partitioning methods fail to meet the graph partitioning requirements for GNN training, which poses new challenges different from traditional graph computation, such as minimizing total computational load and balancing communication load. (§5.1)
- (2) There is a trade-off between minimizing computational load and achieving load balance when partitioning a graph. Partitioning densely connected vertices together can effectively reduce computation but may disrupt load balancing. (§5.3.1)
- (3) Communication load imbalance is more serious in GNN training than in traditional graph computation due to the high-dimensional vertex features and embeddings. Most existing GNN systems ignore this issue. (§5.3.2)
- (4) Most graph partitioning overhead is acceptable for GNN training. Although streaming partitioning is more flexible, existing methods [25, 72] suffer from high computational costs and inefficient implementation due to low parallelism. (§5.3.3)
- (5) Although Metis-extend has better performance, its graph partitioning requires a large memory consumption, especially for large-scale graphs. This limitation prevents the scalability of the Metis-extend partitioning method. (§5.3.4)
- (6) Shorter per-epoch runtime does not necessarily mean faster convergence speed. Partitioning densely connected vertices into the same partition can significantly reduce inter-partition communication and decrease the per-epoch runtime but may decrease the randomness of model training, thus affecting convergence speed. (§5.3.5)

6 BATCH PREPARATION

6.1 Goals and Challenges

After the graph partitioning step, sample-based mini-batch training need a batch preparation step before training starts, which determines the model quality and training performance. In the batch preparation step, the training process divides the training vertices into multiple batches and then samples the L -hop subgraphs for the training vertices based on a sampling method. The batch size determines the number of batches (i.e., the frequency of model parameter updates) in an epoch. The sampling method determines graph topology information that GNN models can learn.

However, there is a trade-off between accuracy and performance when choosing batch size and sampling method.

- (1) Batch size. Due to vertex dependencies, vertices within a batch usually have a large number of common neighbors. As the batch size increases, the number of these common neighbors also grows. The sampling and NN computation results of these common neighbors can be reused, thus reducing the computational load of an epoch. However, increasing the batch size reduces the frequency of model parameter updates, which may affect accuracy and overall convergence speed.
- (2) Sampling method. The sampling method of GNN aims to reduce the size of the training graph to support large-scale GNN training under limited memory. The training speed of GNN can be significantly improved by reducing the size of the sampled

Table 5: Default settings of batch size and sampling parameters in existing GNN systems.

System	Batch size	Fanout	Sampling rate
P3 [12]	1000	(25, 10)	N/A
DistDGL [73]	2000	(25, 10) (15, 10, 5)	N/A
PaGraph [25]	6000	(2, 2)	N/A
GNNLab [62]	8000	(10, 25) (15, 10, 5)	N/A
ByteGNN [72]	512	(10, 5, 3)	N/A
BNS-GCN [49]	full	N/A	0.1
SALIENT++ [18]	1024	(25, 15) (15, 10, 5)	N/A

subgraph. However, as the size of the sampled subgraph decreases, the information learned by the model each time will also decrease, impacting the accuracy and convergence rate.

This section empirically studies the impact of batch size and sample method on accuracy and performance.

6.2 Existing Methods

Batch Size. GNN training methods can be categorized into two types: full-batch training and mini-batch training. Full-batch training requires all vertices to participate in training simultaneously, it requires more storage and computing resources. In addition, the model parameters are updated only once within an epoch, which results in slower model convergence. In contrast, mini-batch training involves only a subset of vertices in each training iteration, allowing model parameters to be updated multiple times in an epoch.

Sampling Fanout and Sampling Rate. The sampling methods of GNN can be divided into two types: fanout-based [12, 25, 39, 52] and ratio-based [49, 73, 75]. Fanout-based methods generate sampled subgraphs based on a fixed number, while ratio-based methods generate sampled subgraphs based on ratio. These two sampling methods determine the size of the sampled subgraph and are orthogonal to sampling algorithms (e.g., vertex-wise sampling [7, 13, 64], layer-wise sampling [6, 16, 77], and subgraph-wise sampling [8, 66]).

To the best of our knowledge, the settings of these parameters have not yet been fully discussed. Users typically configure batch size and sampling methods based on their experience and specific task requirements. Table 5 summarizes parameter settings for several systems. Regarding batch size, common choices include 512, 1024, 2000, 6000, and 8000. Regarding fanout, in 2-layer GNN models, it is typically set to (25, 10), while for 3-layer models, (15, 10, 5) is commonly chosen. Regarding sampling rate, BNS-GCN [49] recommends using a 0.1 ratio to sample boundary vertices (i.e., vertices involved in cross-partition communication). In this section, we conduct extensive experiments to analyze the training performance of the GNN under different parameter settings and also propose a hybrid training method to accelerate the convergence of the model.

Other Optimizations. We note that some work to optimize batch preparation from a system optimization perspective. For example, shared-memory parallel batch preparation [19] uses zero-copy communication to access the remote feature across training processes. However, these optimizations are orthogonal to our concerns, so we do not evaluate them experimentally.

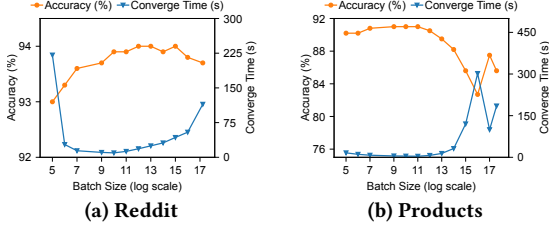


Figure 10: Accuracy and convergence speed with varying batch sizes.

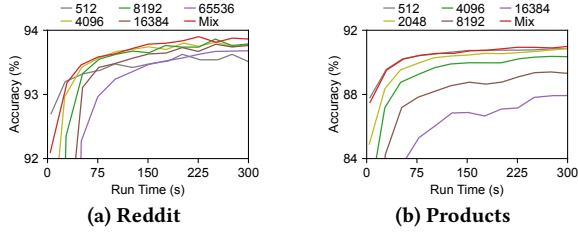


Figure 11: Performance with adaptive batch size.

6.3 Evaluation Results

6.3.1 Batch Size. To investigate the impact of batch size on training, we compare the training accuracy and convergence speed of GNN models under different batch sizes. As shown in Figure 10, we find that different batch sizes have a significant impact on both the accuracy and convergence speed. The specific phenomenon can be summarized in two points:

- (1) Reducing the batch size will speed up convergence. However, after exceeding a lower batch size, the convergence speed of the model decreases.
- (2) Increasing the batch size increases the accuracy of the model. However, after exceeding a higher batch size, the model’s accuracy decreases.

Taking the Reddit dataset as an example, regarding convergence speed, the model converges faster as the batch size decreases. For example, when the batch size is reduced from 32,768 to 128, the model’s convergence speed increases by 3.02 times. However, if the batch size continues to decrease, e.g., from 128 to 64, the model converges 1.95 times slower. Regarding accuracy, the accuracy of the model increases with batch size. For example, when the batch size increases from 32 to 4096, the accuracy of the model improves by 1%. However, if the batch size continues to increase, such as from 4096 to 196615, the model’s accuracy decreases by 0.3%.

In addition, we find that small batch sizes have larger gradient magnitudes, and large batch sizes have smaller gradient magnitudes. Therefore, we try to explain these two phenomena by the training characteristics of Mini-Batch Gradient Descent [14] (MGD). For phenomenon (1), we think that a larger gradient magnitude is good for finding the direction of the model’s optimal point faster, but too large a gradient magnitude is unfavorable for model convergence. For phenomenon (2), we think that a smaller gradient magnitude is good for finding the optimal point of the model, but too small gradient magnitudes tend to fall into the local optimal point of the model. To ensure both accuracy and convergence speed, we can first use a large gradient magnitude (small batch size) to find the optimal point direction and then use a small gradient magnitude (large batch size) to close the optimal point.

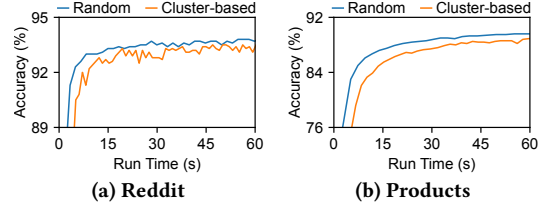


Figure 12: Accuracy and convergence speed of random and cluster-based batch selections.

Table 6: Per-epoch runtime of different batch selection methods.

Dataset	Method	Per-epoch runtime (s)	Involved #V	Involved #E
Products	random	2.4	24.8 M	45.9 M
	cluster-based	1.0	15.7 M	38.3 M
Reddit	random	1.7	6.8 M	25.2 M
	cluster-based	0.6	4.2 M	16.7 M

Adaptive Batch Size. Previous work used fixed batch sizes for GNN training. Through experiments, we find that smaller batch sizes have faster convergence and larger batch sizes have higher accuracy. Based on this, we propose an adaptive batch size training method, which dynamically combines the advantages of different batch sizes during training. The strategy for switching batch sizes is to use a small batch size at the beginning of training so that the model can quickly converge to a higher accuracy. Then, the batch size is gradually increased to converge to the highest accuracy. For example, on the Reddit dataset, we first start training with a batch size of 512 (because it has the fastest convergence rate) and then gradually increase the batch size until the batch size reaches 8,192. As illustrated in Figure 11, adopting this adaptive training method significantly accelerates the model’s convergence rate. On the Reddit and Products datasets, the convergence speed increased by 1.64 and 1.52 times, respectively.

6.3.2 Batch Selection. Batch selection determines the vertices that each batch contains in training. Selecting batch vertices with bias increases the variance in model training, thereby affecting the model’s convergence. Furthermore, the density of connections between vertices within the same batch can impact the overall computational load. To investigate the influence of different batch selection methods on accuracy and convergence speed, we compare two existing batch selection methods. 1) *Random* [11, 18, 25, 62, 73]. It randomly selects a batch of vertices. 2) *Cluster-based* [68]. It leverages graph clustering algorithm (e.g., Metis [20]) to arrange batches.

Table 6 shows the average per-epoch runtime and computational load for each batch selection method, with computational load measured by the number of vertices and edges involved in the training subgraph. The cluster-based method results in the shortest per-epoch runtime. This method selects a batch of vertices from the cluster graph. The vertices in the cluster graph have denser connections, which increases the chance of reusing computational results. The random batch selection method does not consider the dependencies between vertices when selecting batched vertices, it involves the highest number of vertices and edges in computations.

Then, we compare the accuracy of these two batch selection methods on the two datasets of Reddit and Ogbn-products. As shown in Figure 12, the random selection method has the highest accuracy because it has no bias when selecting samples. Since cluster-based batch selection is limited to the internal structure

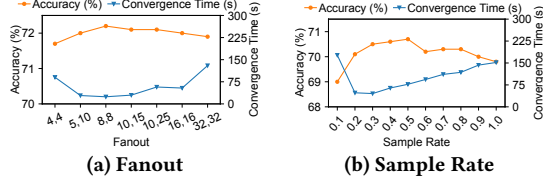


Figure 13: Accuracy and convergence speed of different fanout settings and sample rate settings (Arxiv).

Table 7: Accuracy of high and low degree vertices (Arxiv).

Vertex type	Fanout			
	4, 4	8, 8	16, 16	32, 32
Low-degree vertices	0.718	0.722	0.720	0.718
High-degree vertices	0.786	0.794	0.810	0.817

of the clustered graph, it introduces data bias, resulting in lower accuracy. In addition, we observe that the training process of cluster-based is unstable. We think this phenomenon is related to the distribution of batched subgraphs. Since each cluster graph has a different density, this may lead to an imbalanced distribution of the subgraphs involved in training. To illustrate this point, we calculate the clustering coefficient for each batched subgraph as a measure of its density. Subsequently, we calculate the variance of the clustering coefficients for the batched subgraphs within an epoch to assess their distribution. The experimental results show that the variance of the cluster-based method (2×10^{-4}) is much larger than that of the random selection method (1.1×10^{-6}). Therefore, the imbalanced distribution of batch subgraphs leads to the training process being unstable.

6.3.3 Fanout. The fanout is another key parameter in GNN training. To investigate the impact of fanout on training, we compare the accuracy and convergence speed of GNN models under different fanouts on the Ogbn-arxiv dataset. As shown in Figure 13a, the accuracy of the GNN model shows a trend of "first increase and then decrease" as the fanout increases, while the convergence speed shows the opposite. We think this phenomenon is related to the degree of vertices. Most real-world graphs are skewed (e.g., social networks, e-commerce networks, and citation networks), which means that only a few vertices have very large degrees, while most vertices have small degrees. Hence, a fixed fanout is not conducive to these power-law graphs. We find that low-degree vertices are better suited for smaller fanouts, while high-degree vertices are more appropriate for larger fanouts.

To illustrate this point, we select four typical fanout based on the results in Figure 13a: the smallest fanout (4, 4), the best-performing fanout (8, 8), and two larger fanout (16, 16) and (32, 32). Then, we measure their prediction accuracy on low-degree and high-degree vertices using the Ogbn-arxiv dataset. As shown in Table 7, as fanout increases, the prediction accuracy on low-degree vertices decreases while the accuracy on high-degree vertices increases. This is because, for low-degree vertices, using smaller fanouts can enhance the randomness of sampling while preserving the integrity of graph information. However, using smaller fanouts for high-degree vertices may result in too few neighbor samples, making it challenging to learn the complete graph structure during training.

6.3.4 Sampling Rate. Through the above experimental analysis on fanout, it is not optimal to use a fixed fanout for vertices of different degrees. In addition, we find that fanout is not flexible

Table 8: Accuracy and performance comparison of fanout-based sampling and fanout-rate hybrid sampling (Arxiv).

	4, 4	8, 8	10, 15	10, 25	32, 32	hybrid
Accuracy (%)	71.5	72.1	72.1	72.1	71.5	72.1
Time (s)	300	172	165	237	300	99

enough to express sampling, especially when sampling high-degree vertices. Sampling rate [8, 49, 66] seems to be a good choice, and sampling vertices of different degrees proportionally is more fair than fanout. As shown in Figure 13b, we compare the training accuracy and convergence speed under different sampling rates. Similar to the experimental results of fanout, as the sampling rate increases, the accuracy shows a trend of "first increase and then decrease". It should be noted that the overall accuracy of the sampling rate is lower than that of fanout. There are two main reasons, firstly, smaller sampling rates disadvantage low-degree vertices. For example, for a vertex with degree 20, if the sampling rate is set to 0.1, only 2 vertices can be sampled each time, which is far less than the number of fanout samples (10, 25). Secondly, increasing the sampling rate will sacrifice the randomness of sampling, thus affecting the accuracy and convergence speed of the model. Therefore, using only a sampling rate is not optimal.

Fanout-Rate Hybrid Sampling. Previous work used fixed fanout or sampling rate for GNN training. Through experiments, we found that low-degree vertices are better suited for smaller fanouts, while high-degree vertices are more appropriate for larger fanouts. In addition, as the sampling rate increases, the model accuracy shows a trend of "first increase and then decrease". Based on this, we propose a hybrid training method that combines fanout and sampling rate. The fundamental strategy of the hybrid training method is to employ fanout sampling for low-degree vertices and sampling rate sampling for high-degree vertices. As shown in Table 8, using a hybrid training method can significantly speed up the convergence of the model. Compared to using only fanout (8, 8), the convergence speed is improved by 1.74 times.

6.4 Lessons Learned

- (1) There is a trade-off between accuracy and performance when setting a batch size. By adaptively setting a batch size (using a smaller size at the beginning and then gradually increasing it), it can effectively accelerate the convergence. (§6.3.1)
- (2) Although the complex cluster-based batch selection method can reduce the per-epoch runtime, simple batch selection methods (e.g., random selection) perform better, because no bias is introduced when selecting data samples, resulting in higher accuracy and faster convergence. (§6.3.2)
- (3) It is nonsense to perform sampling with a fixed fanout on all vertices. It is suggested to set a smaller fanout for low-degree vertices and a larger fanout for high-degree vertices. (§6.3.3)
- (4) Dynamically adjusting the batch size during training and combining fanout and sampling rate during sampling can significantly improve the convergence speed without sacrificing accuracy. (§6.3.1 and §6.3.4)

7 DATA TRANSFERRING

7.1 Goals and Challenges

After the batch preparation step, we need to extract the features of the selected vertices in a batch and then load the batch data to

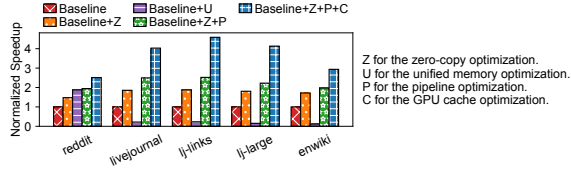


Figure 14: Performance gain of data transfer optimizations.

the GPU for NN computation. The data transferring of GNN has the following special challenges.

- (1) **Irregular memory access.** Compared with DNN, the random memory access pattern of GNN causes a large amount of random I/O overhead in feature extraction.
- (2) **Redundant communication.** There are a large number of repeated vertices and edges between sampled subgraphs due to vertex dependencies. Unlike graph computation, vertices and edges in GNNs are represented as feature vectors, which usually have thousands of dimensions. Transferring these large and redundant vertex/edge features increases the burden of data transferring.

7.2 Existing Methods

The three following data transfer optimization methods have been widely employed in GNN systems [18, 25, 26, 39, 44, 45, 62, 73], including UVA-based data direct access, task pipelining, and cache-based data reusing.

Data transferring method. The data transfer methods between CPU and GPU can be categorized as explicit and implicit transfer. In *explicit* transfers, developers need to explicitly call the CUDA [9] API to transfer data between the CPU and GPU memory. Explicit transfer is suitable for transferring large blocks of data, which can fully utilize the bandwidth. Therefore, for GNN training, scattered vertex features need to be extracted into a contiguous memory space and then transferred to the GPU. In contrast, zero-copy [34] is an *implicit* transfer method based on Unified Virtual Addressing (UVA), which allows GPU threads to access CPU memory directly, thus avoiding the overhead of extracting vertex features. However, kernel performance suffers from high latency accesses to zero-copy memory over the PCIe. Therefore, it requires fine-grained data access orchestration to mitigate the performance impact of latency access [31]. Unified memory [33] is another *implicit* memory access method that provides a single memory address space for GPU and CPU. The requested data memory pages are automatically migrated and cached to the device, and subsequent accesses to the same pages will read the data from GPU memory without additional data transfer. When the memory consumed by cached pages exceeds the GPU memory, some pages may need to be evicted from the GPU to make space for new pages. However, due to the high cost of page migrations and the irregular access pattern of graphs, unified memory is not an efficient way of handling graph algorithms [53].

Task Pipelining. Pipelining decomposes a task into multiple stages and allows these stages to be executed in parallel on different processors or resources. In a CPU and GPU heterogeneous training scenario, a complete batch training process consists of three steps: batch preparation, data transferring, and NN computation. These three steps are executed on three different devices, i.e., CPU, PCIe, and GPU. Due to the data dependency between these steps, the traditional training process performs these three steps sequentially

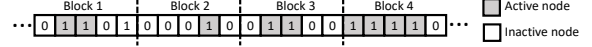


Figure 15: An example of active vertices distribution in a batch.

[3, 12, 25]. Since there is no obvious dependency between training in different batches, we can use pipeline optimization to improve hardware utilization. Specifically, when the CPU finishes sampling for batch b , it transfers the batch data to the GPU via PCIe, while the CPU starts sampling for batch $b + 1$. Once the GPU receives the data for batch b , it begins NN computation. In this way, all hardware resources can be scheduled simultaneously, maximizing the use of hardware resources.

Cache-based Data Reusing. The training process needs to sample labeled vertices to generate L-hop training subgraphs. Due to the intricate dependencies between vertices, there are a large number of duplicate vertices between batch training subgraphs [25]. This results in redundant data transfers between the CPU and GPU, which seriously wastes bandwidth resources. GPU caching is a cache-based data reuse method [19, 25, 26, 39, 44]. Unlike zero-copy and pipeline optimization, by caching vertex features in GPU memory, the data transfer volume between the CPU and GPU can be fundamentally reduced.

7.3 Evaluation Results

7.3.1 Data transferring method. As shown in Figure 2, the data transferring step dominates the entire training process (taking up 73.4% in total), with feature extraction accounting for 31.2% and data loading accounting for 42.2%. As shown in Figure 14, compared to explicit data transfer (e.g., Baseline), zero-copy optimization (e.g., Baseline+Z) has 1.74 times the performance gain on average. Using unified memory can achieve an improvement of 1.88 times on a small-scale dataset (Reddit). However, on large-scale datasets, using unified memory does not bring positive performance gains due to frequent page migration overhead.

Does Hybrid Data Transferring Help? Regarding data transfer optimization, a hybrid data transfer optimization method [53] was recently proposed. This approach combines explicit and implicit transferring methods, dynamically selecting the most suitable transfer method for each data block. In order to explore whether hybrid data transfer optimization can accelerate GNN training, we analyze the distribution of active vertices (sampled vertices) within each batch. As shown in Figure 15, we count the active vertices in the data block in units of 256KB [31]. For each data block, we set a threshold. If the number of active vertices exceeds this threshold, the block will be considered suitable for explicit transfer. Otherwise, it is suitable for implicit transfer (e.g., when the threshold is 0.5, only block 4 is suitable for explicit transfer).

Figure 16 shows the ratio of data blocks suitable for explicit transfer at different thresholds. We can observe that as the threshold increases, the ratio of data blocks suitable for explicit transfer decreases significantly. Since the Reddit dataset has a high average degree, which makes the distribution of active vertices denser, it is more suitable for the explicit transfer of data blocks. However, after applying GPU cache optimization, the number of data blocks suitable for explicit transfer is significantly reduced (orange line in Figure 16). For example, when the threshold is set to 80%, only 2% of the data blocks are suitable for explicit transfer even for the

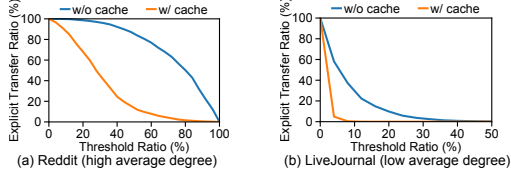


Figure 16: Ratio of active blocks with different thresholds.

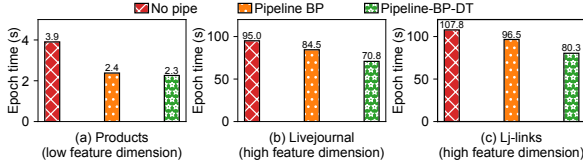


Figure 17: Pipeline training ablation study.

Reddit dataset. Based on the above analysis, We find that the demand for hybrid data transferring in GNN training is not significant. This is due to the unique computational patterns of GNNs. Traditional graph computations have different vertex access patterns in iterative computations [53, 76]. However, in GNN training, the distribution of vertices in the sampled subgraph is often random and fragmented, especially after implementing a caching optimization. Therefore, though the hybrid data transfer shows promising results in graph computation, it does not help GNN training.

7.3.2 Task Pipelining. In the CPU-GPU heterogeneous training scenarios, a complete batch training process can be divided into three steps: batch preparation (BP), data transfer (DT), and NN computation (NN). As shown in Figure 14, Baseline+Z+P represents the simultaneous use of zero-copy transfer optimization and pipelining optimization. It offers a performance improvement of 2.26 times compared to Baseline (i.e., explicit transfer). When compared to Baseline+Z (i.e., zero-copy optimization), the pipeline optimization brings an additional performance gain of 1.30 times.

Figure 17 is a pipeline ablation experiment that reports the per-epoch runtime when adding each step to the pipeline. No Pipe is the slowest one due to its totally sequential execution. Pipeline BP can significantly reduce the running time by pipelining the BP step. The results show that the acceleration effect of pipelining is not significant because data transfer is still the main bottleneck. On the LiveJournal and Lj-links datasets, the time proportion of data transfer is 58.8% and 53.1%, respectively, which dominate the training process. Some research works [40, 50, 56] have demonstrated that improving the performance of the pipeline is possible by balancing GNN workloads between CPUs and GPUs and employing fine-grained task scheduling.

7.3.3 Cache-based Data Reusing. There are two main caching strategies: (1) degree-based caching (degree), (2) pre-sampling-based caching (sample). The degree-based caching strategy prioritizes caching vertices with large out-degrees. The sampling-based caching strategy first obtains the access frequency of each vertex by pre-sampling and then prioritizes the caching of frequently accessed vertex features. We compare the performance of these two caching strategies. As shown in Figure 18, on the Amazon dataset, the performance of these two caching strategies is comparable. However, on the OGB-Paper dataset, the pre-sampling-based caching strategy significantly outperforms the degree-based caching strategy. This is because degree-based caching is based on the assumption

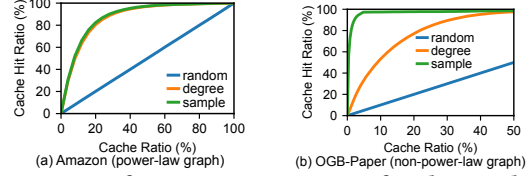


Figure 18: Performance comparison of caching policies.

that higher-degree vertices are more likely to be sampled. Therefore, degree-based caching strategies perform well on power-law graphs but perform poorly on non-power-law graphs. In addition, the degree-based caching strategy is only applicable to the uniform vertex sampling algorithm [13]. For special sampling algorithms (such as importance sampling [6]), the degree-based assumption is no longer valid. In contrast, the pre-sampling-based caching strategy has better robustness and performs well on both power-law graphs and non-power-law graphs.

7.4 Lessons Learned

- (1) Due to GNN's irregular data access pattern, the vertex features required for heterogeneous training are often scattered in memory. Through UVA-based data direct access (zero-copy) optimization, the GPU device can access CPU memory directly, avoiding expensive feature extraction operations. (§7.3.1)
- (2) Vertices do not exhibit an obvious active or deactivated state in GNN training, the hybrid data transfer optimization will not help improve data transferring. (§7.3.1)
- (3) The effect of pipeline optimization is not outstanding (less than 50% improvement in most cases) under our experimental settings because the data transferring step is the main bottleneck, overwhelming the other tasks. (§7.3.2)
- (4) GPU caching is the most significant data transfer optimization method as it fundamentally reduces data transfer volume. Compared to degree-based caching, pre-sampling-based caching exhibits better robustness. (§7.3.3)

8 CONCLUSION

In this paper, we conduct a comprehensive evaluation of the optimization techniques adopted by existing GNN systems from a data management perspective. We show that graph partitioning of GNN is a complex task. Achieving the goals of GNN graph partitioning presents significant challenges. In addition, we find that there are significant differences in the accuracy and convergence speed of GNN under different parameter settings. We propose adaptive batch size training and hybrid fanout and sample rate training methods to accelerate convergence while maintaining accuracy. Lastly, we conduct evaluations of three data transfer optimizations, summarizing their advantages, disadvantages, and suitable scenarios. Based on these results, we provide insights and lessons learned for enhancing the training performance of GNNs.

ACKNOWLEDGMENTS

This work is supported by the National Key R&D Program of China (2018YFB1003400), the National Natural Science Foundation of China (U2241212, 62072082, 62202088, 62072083, and 62372097), and the Fundamental Research Funds for the Central Universities (N2216015 and N2216012). Yanfeng Zhang and Qiange Wang are the corresponding authors.

REFERENCES

- [1] Sami Abu-El-Hajja, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. 2019. MixHop: Higher-Order Graph Convolutional Architectures via Sparsified Neighborhood Mixing. In *Proceedings of the 36th International Conference on Machine Learning, ICML'19, Long Beach, California, USA*. 21–29.
- [2] Sergey Brin and Lawrence Page. 1998. The Anatomy of a Large-Scale Hypertextual Web Search Engine. *Comput. Networks* 30, 1-7 (1998), 107–117.
- [3] Zhenkun Cai, Xiao Yan, Yidi Wu, Kaihao Ma, James Cheng, and Fan Yu. 2021. DGCL: an efficient communication library for distributed GNN training. In *Sixteenth European Conference on Computer Systems, EuroSys'21, Online Event, United Kingdom*. 130–144.
- [4] Zhenkun Cai, Qihui Zhou, Xiao Yan, Da Zheng, Xiang Song, Chenguang Zheng, James Cheng, and George Karypis. 2023. DSP: Efficient GNN Training with Multiple GPUs. In *Proceedings of the 28th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming, PPoPP'23, Montreal, QC, Canada*. 392–404.
- [5] Chaoyi Chen, Dechao Gao, Yanfeng Zhang, Qiange Wang, Zhenbo Fu, Xuechang Zhang, Junhua Zhu, Yu Gu, and Ge Yu. 2023. NeutronStream: A Dynamic GNN Training Framework with Sliding Window for Graph Streams. *Proc. VLDB Endow.* 17, 3 (2023), 455–468.
- [6] Jie Chen, Tengfei Ma, and Cao Xiao. 2018. FastGCN: Fast Learning with Graph Convolutional Networks via Importance Sampling. In *6th International Conference on Learning Representations, ICLR'18, Vancouver, BC, Canada*.
- [7] Jianfei Chen, Jun Zhu, and Le Song. 2018. Stochastic Training of Graph Convolutional Networks with Variance Reduction. In *Proceedings of the 35th International Conference on Machine Learning, ICML'18, Stockholm, Sweden*. 941–949.
- [8] Wei-Lin Chiang, Xuanqing Liu, Si Si, Yang Li, Samy Bengio, and Cho-Jui Hsieh. 2019. Cluster-GCN: An Efficient Algorithm for Training Deep and Large Graph Convolutional Networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD'19, Anchorage, AK, USA*. 257–266.
- [9] CUDA 2023. CUDA. <https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html>.
- [10] Euler 2019. Euler. <https://github.com/alibaba/euler/wiki/System-Introduction>.
- [11] Matthias Fey and Jan E. Lenssen. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- [12] Swapnil Gandhi and Anand Padmanabha Iyer. 2021. P3: Distributed Deep Graph Learning at Scale. In *15th USENIX Symposium on Operating Systems Design and Implementation, OSDI'21, 2021*. 551–568.
- [13] William L. Hamilton, Zitao Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, NeurIPS'17 Long Beach, CA, USA*. 1024–1034.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR'16, Las Vegas, NV, USA*. 770–778.
- [15] Kezhao Huang, Jidong Zhai, Zhen Zheng, Youngmin Yi, and Xipeng Shen. 2021. Understanding and bridging the gaps in current GNN performance optimizations. In *26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPoPP'21, Virtual Event, Republic of Korea*. 119–132.
- [16] Wen-bing Huang, Tong Zhang, Yu Rong, and Junzhou Huang. 2018. Adaptive Sampling Towards Fast Graph Representation Learning. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS'18, Montréal, Canada*. 4563–4572.
- [17] Zhihao Jia, Sina Lin, Mingyu Gao, Matei Zaharia, and Alex Aiken. 2020. Improving the Accuracy, Scalability, and Performance of Graph Neural Networks with Roc. In *Proceedings of Machine Learning and Systems 2020, MLSys'20, Austin, TX, USA*.
- [18] Tim Kaler, Alexandros-Stavros Iliopoulos, Philip Murzynowski, Tao B. Schardl, Charles E. Leiserson, and Jie Chen. 2023. Communication-Efficient Graph Neural Networks with Probabilistic Neighborhood Expansion Analysis and Caching. *CoRR* abs/2305.03152 (2023).
- [19] Tim Kaler, Nickolas Stathas, Anne Ouyang, Alexandros-Stavros Iliopoulos, Tao B. Schardl, Charles E. Leiserson, and Jie Chen. 2022. Accelerating Training and Inference of Graph Neural Networks with Fast Sampling and Pipelining. In *Proceedings of Machine Learning and Systems 2022, MLSys'22, Santa Clara, CA, USA*.
- [20] George Karypis and Vipin Kumar. 1998. A Fast and High Quality Multilevel Scheme for Partitioning Irregular Graphs. *SIAM J. Sci. Comput.* 20, 1 (1998), 359–392.
- [21] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *5th International Conference on Learning Representations, ICLR'17, Toulon, France, Conference Track Proceedings*. OpenReview.net.
- [22] KONECT 2017. LiveJournal links network dataset. <http://konect.uni-koblenz.de/networks/livejournal-links>.
- [23] KONECT 2017. Wikipedia links, English network dataset. http://konect.cc/networks/wikipedia_link_en.
- [24] Haiyang Lin, Mingyu Yan, Xiaochun Ye, Dongrui Fan, Shirui Pan, Wenguang Chen, and Yuan Xie. 2022. A Comprehensive Survey on Distributed Training of Graph Neural Networks. *CoRR* abs/2211.05368 (2022).
- [25] Zhiqi Lin, Cheng Li, Youshan Miao, Yunxin Liu, and Yinlong Xu. 2020. PaGraph: Scaling GNN training on large graphs via computation-aware caching. In *ACM Symposium on Cloud Computing, SoCC'20, Virtual Event, USA*. 401–415.
- [26] Tianfeng Liu, Yangrui Chen, Dan Li, Chuan Wu, Yibo Zhu, Jun He, Yanghua Peng, Hongzheng Chen, Hongzhi Chen, and Chuanxiong Guo. 2023. BGL: GPU-Efficient GNN Training by Optimizing Graph Data I/O and Preprocessing. In *20th USENIX Symposium on Networked Systems Design and Implementation, NSDI'23, Boston, MA*. 103–118.
- [27] Xin Liu, Mingyu Yan, Lei Deng, Guoqi Li, Xiaochun Ye, and Dongrui Fan. 2022. Sampling Methods for Efficient Training of Graph Convolutional Networks: A Survey. *IEEE CAA J. Autom. Sinica* 9, 2 (2022), 205–234.
- [28] Lingxiao Ma, Zhi Yang, Youshan Miao, Jilong Xue, Ming Wu, Lidong Zhou, and Yafei Dai. 2019. NeuGraph: Parallel Deep Neural Network Computation on Large Graphs. In *2019 USENIX Annual Technical Conference, ATC'19, Renton, WA, USA*. 443–458.
- [29] Vasimuddin Md, Sanchit Misra, Guixiang Ma, Ramanarayan Mohanty, Evangelos Georganas, Alexander Heinecke, Dhiraj D. Kalamkar, Nesreen K. Ahmed, and Sasikanth Avancha. 2021. DistGNN: scalable distributed training for large-scale graph neural networks. In *International Conference for High Performance Computing, Networking, Storage and Analysis, SC'21, St. Louis, Missouri, USA*. 76.
- [30] Nikolai Merkel, Daniel Stoll, Ruben Mayer, and Hans-Arno Jacobsen. 2023. An Experimental Comparison of Partitioning Strategies for Distributed Graph Neural Network Training. *CoRR* abs/2308.15602 (2023).
- [31] Seungwon Min, Kun Wu, Sitao Huang, Mert Hidayetoglu, Jinjun Xiong, Eiman Ebrahimi, Deming Chen, and Wen-mei W. Hwu. 2021. Large Graph Convolutional Network Training with GPU-Oriented Data Communication Architecture. *Proc. VLDB Endow.* 14, 11 (2021), 2087–2100.
- [32] Alan Mislove, Massimiliano Marcon, P. Krishna Gummadi, Peter Druschel, and Bobby Bhattacharjee. 2007. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM Internet Measurement Conference, IMC'07, San Diego, California, USA*. 29–42.
- [33] NVIDIA. 2017. Mark Harris. Unified Memory for CUDA Beginners – nvidia developer blog. <https://developer.nvidia.com/blog/unified-memory-cuda-beginners>.
- [34] NVIDIA. 2022. Zero Copy. <https://docs.nvidia.com/cuda/cuda-c-best-practices-guide/index.html#zero-copy>.
- [35] Open Graph Benchmark 2021. The ogbn-arxiv dataset. <https://ogb.stanford.edu/docs/nodeprop/#ogbn-arxiv>.
- [36] Open Graph Benchmark 2021. The ogbn-papers100M dataset. <https://ogb.stanford.edu/docs/nodeprop/#ogbn-papers100M>.
- [37] Open Graph Benchmark 2021. The ogbn-products dataset. <https://ogb.stanford.edu/docs/nodeprop/#ogbn-products>.
- [38] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 8024–8035. <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [39] Jingshu Peng, Zhao Chen, Yingxia Shao, Yanyan Shen, Lei Chen, and Jiannong Cao. 2022. SANCUS: Staleness-Aware Communication-Avoiding Full-Graph Decentralized Training in Large-Scale Graph Neural Networks. *Proc. VLDB Endow.* 15, 9 (2022), 1937–1950.
- [40] Morteza Ramezani, Weilin Cong, Mehrdad Mahdavi, Anand Sivasubramaniam, and Mahmut T. Kandemir. 2020. GCN meets GPU: Decoupling “When to Sample” from “How to Sample”. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS'20, virtual, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.)*.
- [41] Victor Garcia Satorras and Joan Bruna Estrach. 2018. Few-Shot Learning with Graph Neural Networks. In *6th International Conference on Learning Representations, ICLR'18, Vancouver, BC, Canada*.
- [42] Yingxia Shao, Hongzheng Li, Xizhi Gu, Hongbo Yin, Yawen Li, Xupeng Miao, Wentao Zhang, Bin Cui, and Lei Chen. 2022. Distributed Graph Neural Network Training: A Survey. *CoRR* abs/2211.00216 (2022).
- [43] Shihui Song and Peng Jiang. 2022. Rethinking graph data placement for graph neural network training on multiple GPUs. In *27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPoPP'22, Seoul, Republic of Korea*. 455–456.
- [44] Jie Sun, Li Su, Zuocheng Shi, Wenting Shen, Zeke Wang, Lei Wang, Jie Zhang, Yong Li, Wenyan Yu, Jingren Zhou, and Fei Wu. 2023. Legion: Automatically

- Pushing the Envelope of Multi-GPU System for Billion-Scale GNN Training. In *Annual Technical Conference, ATC'23, Boston, MA, USA*. 165–179.
- [45] John Thorpe, Yifan Qiao, Jonathan EyoIfson, Shen Teng, Guanzhou Hu, Zhihao Jia, Jinliang Wei, Keval Vora, Ravi Netravali, Miryung Kim, and Guoqing Harry Xu. 2021. Dorylus: Affordable, Scalable, and Accurate GNN Training with Distributed CPU Servers and Serverless Threads. In *15th USENIX Symposium on Operating Systems Design and Implementation, OSDI'21*. 495–514.
- [46] Jana Vatter, Ruben Mayer, and Hans-Arno Jacobsen. 2024. The Evolution of Distributed Systems for Graph Neural Networks and Their Origin in Graph Processing and Deep Learning: A Survey. *ACM Comput. Surv.* 56, 1 (2024), 6:1–6:37.
- [47] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *6th International Conference on Learning Representations, ICLR'18, Vancouver, BC, Canada*.
- [48] Roger Waleffe, Jason Mohoney, Theodoros Rekatsinas, and Shivaram Venkataraman. 2023. MariusGNN: Resource-Efficient Out-of-Core Training of Graph Neural Networks. In *Eighteenth European Conference on Computer Systems, EuroSys'23, Rome, Italy*. 144–161.
- [49] Cheng Wan, Youjie Li, Ang Li, Nam Sung Kim, and Yingyan Lin. 2022. BNS-GCN: Efficient Full-Graph Training of Graph Convolutional Networks with Partition-Parallelism and Random Boundary Node Sampling. In *Proceedings of Machine Learning and Systems 2022, MLSys'22, Santa Clara, CA, USA*.
- [50] Xinchen Wan, Kaiqiang Xu, Xudong Liao, Yilun Jin, Kai Chen, and Xin Jin. 2023. Scalable and Efficient Full-Graph GNN Training for Large Graphs. *Proc. ACM Manag. Data* 1, 2 (2023), 143:1–143:23.
- [51] Lei Wang, Qiang Yin, Chao Tian, Jianbang Yang, Rong Chen, Wenyuan Yu, Zihang Yao, and Jingren Zhou. 2021. FlexGraph: a flexible and efficient distributed framework for GNN training. In *Sixteenth European Conference on Computer Systems, EuroSys'21, Online Event, United Kingdom*. 67–82.
- [52] Minjie Wang, Da Zheng, Zihao Ye, Quan Gan, Mufei Li, Xiang Song, Jinjing Zhou, Chao Ma, Lingfan Yu, Yu Gai, Tianjun Xiao, Tong He, George Karypis, Jinyang Li, and Zheng Zhang. 2019. Deep Graph Library: A Graph-Centric, Highly-Performant Package for Graph Neural Networks. *arXiv preprint arXiv:1909.01315* (2019).
- [53] Qiang Wang, Xin Ai, Yanfeng Zhang, Jing Chen, and Ge Yu. 2023. HyTGraph: GPU-Accelerated Graph Processing with Hybrid Transfer Management. In *39th IEEE International Conference on Data Engineering, ICDE'23, Anaheim, CA, USA*. 558–571.
- [54] Qiang Wang, Yao Chen, Weng-Fai Wong, and Bingsheng He. 2023. HongTu: Scalable Full-Graph GNN Training on Multiple GPUs (via communication-optimized CPU data offloading). *CoRR abs/2311.14898* (2023).
- [55] Qiang Wang, Yanfeng Zhang, Hao Wang, Chaoyi Chen, Xiaodong Zhang, and Ge Yu. 2022. NeutronStar: Distributed GNN Training with Hybrid Dependency Management. In *International Conference on Management of Data, SIGMOD'22, Philadelphia, PA, USA*. 1301–1315.
- [56] Yuke Wang, Boyuan Feng, Zheng Wang, Tong Geng, Kevin J. Barker, Ang Li, and Yufei Ding. 2023. MGG: Accelerating Graph Neural Networks with Fine-Grained Intra-Kernel Communication-Computation Pipelining on Multi-GPU Platforms. In *17th USENIX Symposium on Operating Systems Design and Implementation, OSDI'23, Boston, MA, USA*.
- [57] Zhaokang Wang, Yunpan Wang, Chunfeng Yuan, Rong Gu, and Yihua Huang. 2021. Empirical analysis of performance bottlenecks in graph neural network training and inference with GPUs. *Neurocomputing* (2021), 165–191.
- [58] Duncan J Watts and Steven H Strogatz. 1998. Collective dynamics of 'small-world' networks. *nature* 393, 6684 (1998), 440–442.
- [59] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. 2021. A Comprehensive Survey on Graph Neural Networks. *IEEE Trans. Neural Networks Learn. Syst.* 32, 1 (2021), 4–24.
- [60] Zihui Xue, Yuedong Yang, Mengtian Yang, and Radu Marculescu. 2022. SUGAR: Efficient Subgraph-level Training via Resource-aware Graph Partitioning. *CoRR abs/2202.00075* (2022).
- [61] Jaewon Yang and Jure Leskovec. 2012. Defining and Evaluating Network Communities Based on Ground-Truth. In *12th IEEE International Conference on Data Mining, ICDM'12, Brussels, Belgium*. 745–754.
- [62] Jianbang Yang, Dahai Tang, Xiaoniu Song, Lei Wang, Qiang Yin, Rong Chen, Wenyuan Yu, and Jingren Zhou. 2022. GNNLab: a factored system for sample-based GNN training over GPUs. In *Seventeenth European Conference on Computer Systems, EuroSys'22, Rennes, France*. 417–434.
- [63] Shuangyan Yang, Minjia Zhang, Wenqian Dong, and Dong Li. 2023. Betty: Enabling Large-Scale GNN Training with Batch-Level Graph Partitioning. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2, ASPLOS'23, Vancouver, BC, Canada*. 103–117.
- [64] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L. Hamilton, and Jure Leskovec. 2018. Graph Convolutional Neural Networks for Web-Scale Recommender Systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD'18, London, UK*. 974–983.
- [65] Zhitao Ying, Jiaxuan You, Christopher Morris, Xiang Ren, William L. Hamilton, and Jure Leskovec. 2018. Hierarchical Graph Representation Learning with Differentiable Pooling. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS'18, Montréal, Canada*. 4805–4815.
- [66] Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor K. Prasanna. 2020. GraphSAINT: Graph Sampling Based Inductive Learning Method. In *8th International Conference on Learning Representations, ICLR'20, Addis Ababa, Ethiopia*.
- [67] Dalong Zhang, Xin Huang, Ziqi Liu, Jun Zhou, Zhiyang Hu, Xianzheng Song, Zhibang Ge, Lin Wang, Zhiqiang Zhang, and Yuan Qi. 2020. AGL: A Scalable System for Industrial-purpose Graph Machine Learning. *Proc. VLDB Endow.* 13, 12 (2020), 3125–3137.
- [68] Lizhi Zhang, Kai Lu, Zhiqian Lai, Yongquan Fu, Yu Tang, and Dongsheng Li. 2023. Accelerating GNN Training by Adapting Large Graphs to Distributed Heterogeneous Architectures. *IEEE Trans. Comput.* (2023), 1–14. <https://doi.org/10.1109/TC.2023.3305077>
- [69] Muhan Zhang and Yixin Chen. 2018. Link Prediction Based on Graph Neural Networks. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS'18, Montréal, Canada*.
- [70] Wentao Zhang, Yu Shen, Zheyu Lin, Yang Li, Xiaosen Li, Wen Ouyang, Yangyu Tao, Zhi Yang, and Bin Cui. 2022. PaSca: A Graph Neural Architecture Search System under the Scalable Paradigm. In *The ACM Web Conference 2022, WWW'22, Virtual Event, Lyon, France*.
- [71] Xin Zhang, Yanyan Shen, Yingxia Shao, and Lei Chen. 2023. DUCATI: A Dual-Cache Training System for Graph Neural Networks on Giant Graphs with the GPU. *Proc. ACM Manag. Data* 1, 2 (2023), 166:1–166:24.
- [72] Chenguang Zheng, Hongzhi Chen, Yuxuan Cheng, Zhezhen Song, Yifan Wu, Changji Li, James Cheng, Hao Yang, and Shuai Zhang. 2022. ByteGNN: Efficient Graph Neural Network Training at Large Scale. *Proc. VLDB Endow.* 15, 6 (2022), 1228–1242.
- [73] Da Zheng, Chao Ma, Minjie Wang, Jinjing Zhou, Qidong Su, Xiang Song, Quan Gan, Zheng Zhang, and George Karypis. 2020. DistDGL: Distributed Graph Neural Network Training for Billion-Scale Graphs. In *10th IEEE/ACM Workshop on Irregular Applications: Architectures and Algorithms, IA3'20, Atlanta, GA, USA*. 36–44.
- [74] Da Zheng, Xiang Song, Chengru Yang, Dominique LaSalle, and George Karypis. 2022. Distributed Hybrid CPU and GPU training for Graph Neural Networks on Billion-Scale Heterogeneous Graphs. In *The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD'22, Washington, DC, USA*. 4582–4591.
- [75] Rong Zhu, Kun Zhao, Hongxia Yang, Wei Lin, Chang Zhou, Baole Ai, Yong Li, and Jingren Zhou. 2019. AliGraph: A Comprehensive Graph Neural Network Platform. *Proc. VLDB Endow.* 12, 12 (2019), 2094–2105.
- [76] Xiaowei Zhu, Wenguang Chen, Weimin Zheng, and Xiaosong Ma. 2016. Gemini: A Computation-Centric Distributed Graph Processing System. In *12th USENIX Symposium on Operating Systems Design and Implementation, OSDI' 2016, Savannah, GA, USA, November 2-4, 2016*, Kimberly Keeton and Timothy Roscoe (Eds.). 301–316.
- [77] Difan Zou, Ziniu Hu, Yewen Wang, Song Jiang, Yizhou Sun, and Quanquan Gu. 2019. Layer-Dependent Importance Sampling for Training Deep and Large Graph Convolutional Networks. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS'19, Vancouver, BC, Canada*. 11247–11256.