

DiskGNN: Bridging I/O Efficiency and Model Accuracy for Out-of-Core GNN Training




Renjie Liu*, Yichuan Wang*, Xiao Yan, Haitian Jiang, Zhenkun Cai, Minjie Wang, Bo Tang, Jinyang Li (* indicates equal contributions)



Out-of-core GNN Training

- In resource-constrained environments, large graphs (topology & features) can not be loaded from disk to memory.
- A training system is needed that can fetch data directly from disk while minimizing I/O overhead.



Challenge: I/O Efficiency & Model Acc

- **Issues of existing systems:**
 - **High I/O volume** due to read amplification
 - **Degrade GNN model accuracy** due to loss of edges


Metrics	Ginex [39]	MariusGNN [52]	DiskGNN (ours)
End-to-end time (hrs)	9.72	3.66	1.09
Pre-processing time (hrs)	1.66	0.81	0.03
Training time (hrs)	8.06	2.85	1.06
Avg. epoch time (sec)	580	205	76.3
Disk access time (sec)	412	27.1	51.2
Disk access volume (GB)	484	6.46	73.9
Final test accuracy (%)	65.9	64.0	65.9

- (1) **Ginex (VLDB2022)**: uses random reads to fetch cache-missed features but has read amplification and high I/O volume.
- (2) **MariusGNN (Eurosys2023)**: uses partition-based accesses but ignores cross-partition edges during in-memory graph sampling.

Existing systems either have **poor I/O efficiency** or **degrade the model accuracy**.


Node Feature Orchestration

- **Four-level Feature Store:**
 - (1) Store popular features in GPU and CPU cache.
 - (2) Deduplicate on-disk features by disk cache and reordering. (Segmented Disk Cache with MinHash perfectly balance tradeoff).
 - (3) Pack cache-missed features in contiguous disk location.




Efficient Feature Packing

- **Batched Feature Packing:** Change the packing iteration order from mini-batches to node features, using only sequential reads and thus eliminate read replication and amplification.



Evaluation Results

- **Hardware:** AWS G5 instance, default cache size: 10%.
- **Datasets:** Ogbn-papers100M, Friendster, MAG240M, IGB.



Offline Sampling-based Training

Offline Sampling: Pre-sample mini-batches in advance and reuse them training. Packed feature files are used to reduce read amplification.

