

CS 744: DISTRIBUTED DGL

Shivaram Venkataraman

Fall 2022

ADMINISTRIVIA

- Midterm grades out! →
- Regrade requests (check question numbers)
 - Thu: After class, Roger's OH
 - Mon: Shivaram's OH, Roger's OH
 - Tue: After class
- Course Project: Check in by Nov 23th → Canvas / Piazza today

EXAMPLE: LINK PREDICTION

Task: Predict potential connections in a social network

*Machine learning
on graph data*



[0.25, 0.45, 0.30]

[0.15, 0.85, 0.92]

...

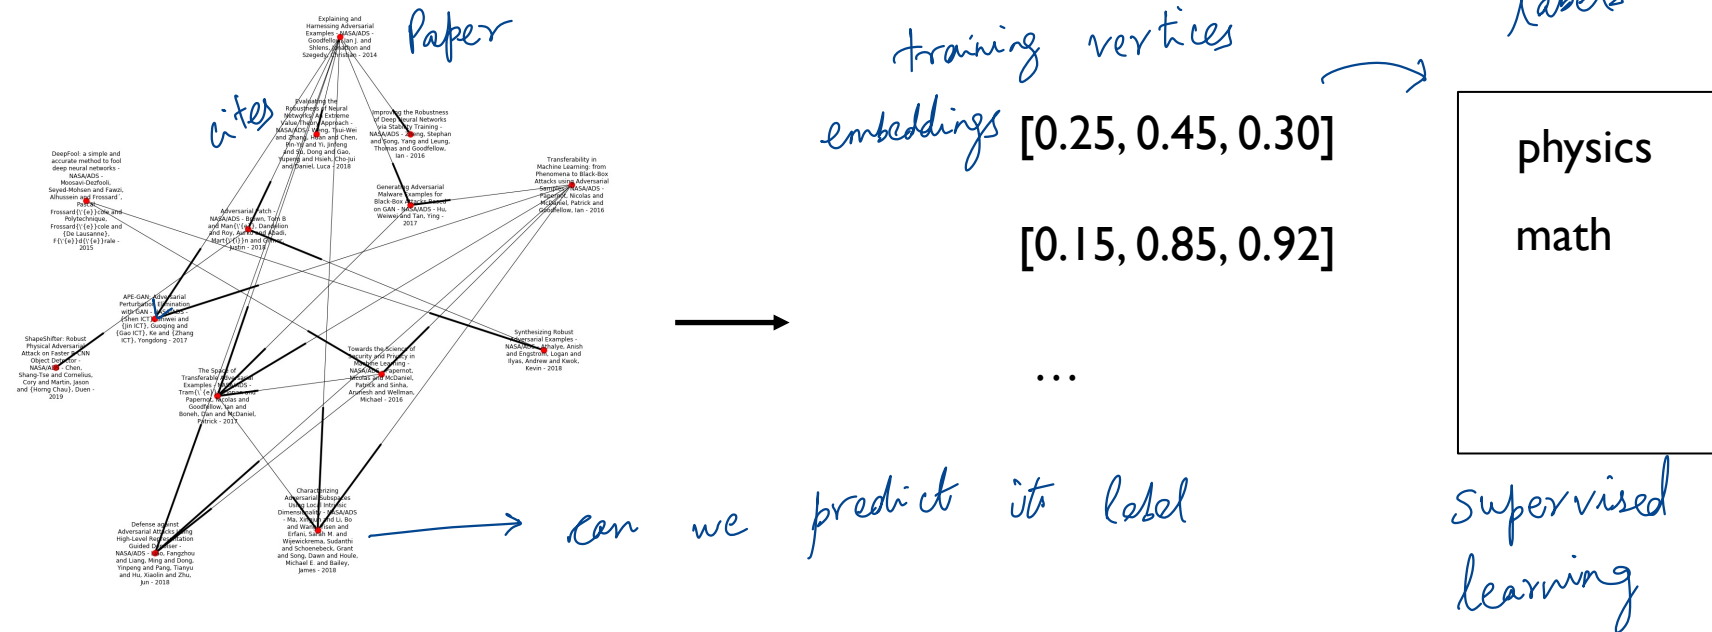


Find K-
nearest
neighbors

EXAMPLE: NODE CLASSIFICATION

Task: Classify papers in a citation graph by subject area

subsets of nodes
have
labels



GRAPH EMBEDDING MODELS: DECODER-ONLY

Loss function

Maximize score for edges in graph

Minimize for others (negative edges)

$$\mathcal{L} = \sum_{e \in G} \sum_{e' \in S'_e} \max(f(e) - f(e') + \lambda, 0)$$

↓
positive

↓
negative edges



BACKGROUND: GRAPH NEURAL NETWORKS (GNN)

Graph Neural Networks: Use **neural network** to capture neighborhood structure

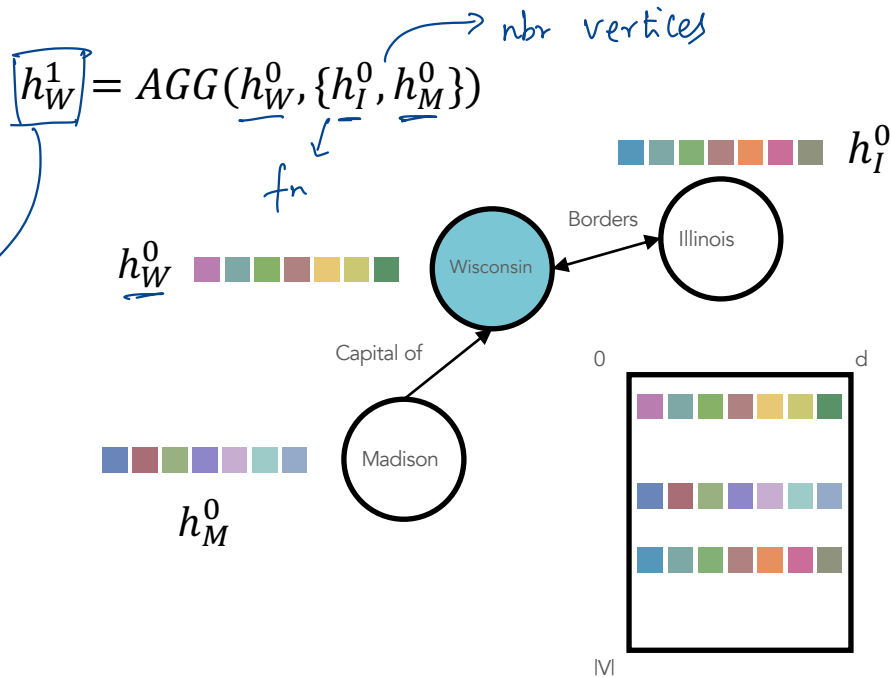
Input: h_i^0 (base node representations)

Model: $h_i^k = \text{AGG}(h_i^{k-1}, \{h_u^{k-1} : u \in N_i\})$

N_i one-hop neighborhood of i

AGG parameterized aggregation func

2-hop nbr hood
 ↳ 1 hop neighborhood of
 nbr
 ↳ AGG on the results
 output



BACKGROUND: GRAPHSAGE

→ example of GNN

$$\mathbf{h}_v^k \leftarrow \sigma(\mathbf{W} \cdot \text{MEAN}(\{\mathbf{h}_v^{k-1}\} \cup \{\mathbf{h}_u^{k-1}, \forall u \in \mathcal{N}(v)\}))$$

weights need to be learned

only learn weights

Learn weights and embeddings

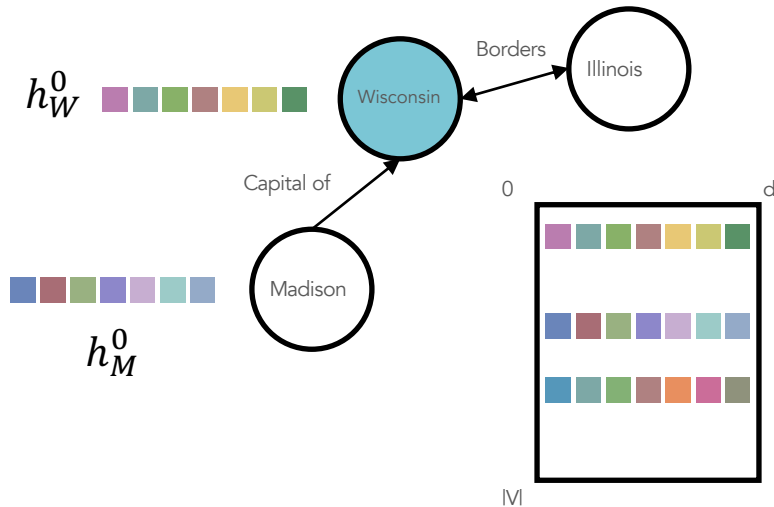
h_v^1

↳ Gather h_u^0 for all $u \in \mathcal{N}(v)$

↳ Mean of all nbr embeddings and h_v^0

weight matrix \sim fully connected layer

h_i^0



DISTDGL: DEEP GRAPH LIBRARY

Distributed system for training GNNs

KVStore

Mini-batch sampler

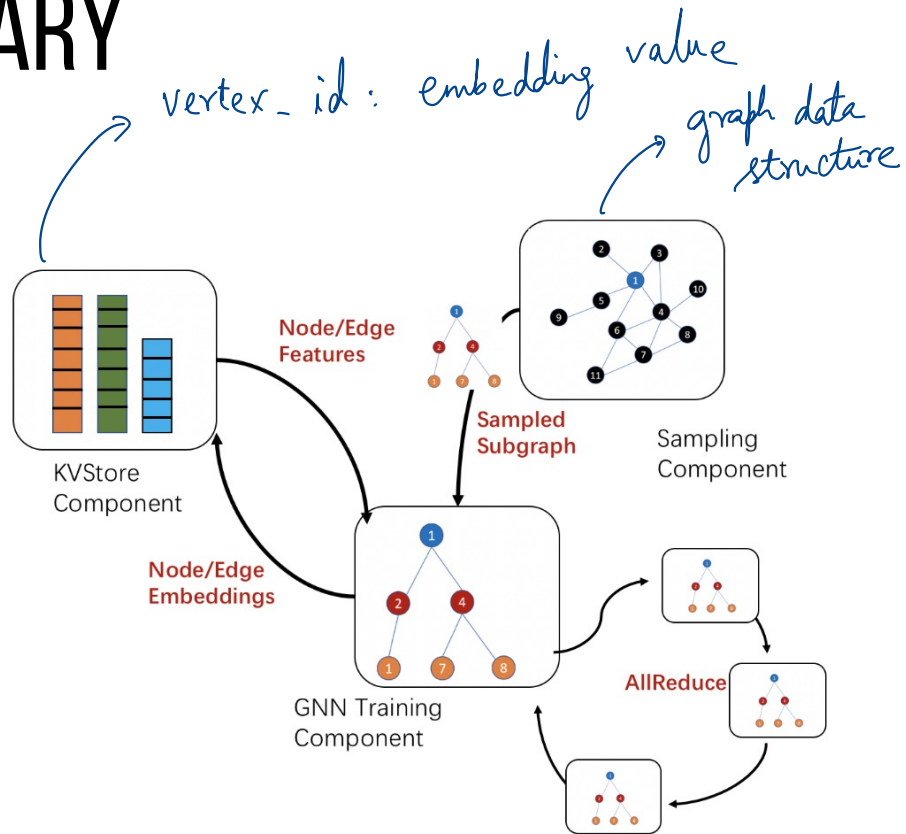
Trainer

System Design

Key techniques

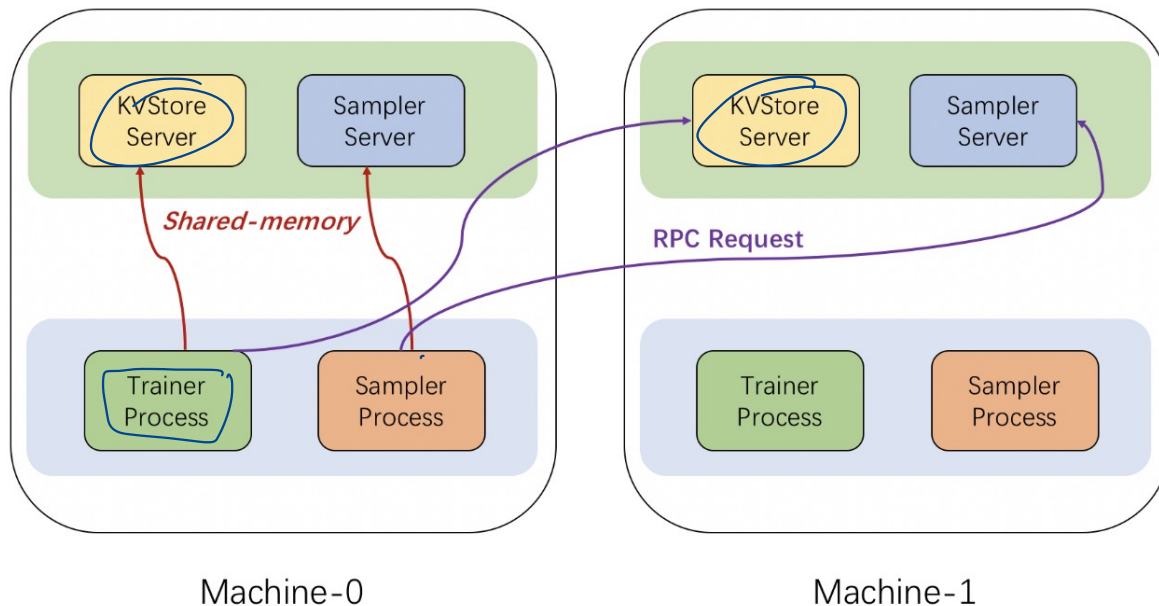
Partitioning heterogeneous graphs

Async mini-batch sampling



DISTDGL SYSTEM SETUP

Sampler:
given the graph,
generate samples
→ list of vertices



① Partition embeddings

② sampler
← Trainer
interact

GRAPH PARTITIONING: METIS → 1999!


Hierarchical METIS

Apply METIS to partition graph
across machines


↪ balance number
of vertices
↪ minimize num edge
across partitions

Re-apply METIS to partition
within a machine

→ such that
fits in GPU memory



(a) Assign vertices to graph partitions




A fast and high quality multilevel scheme for partitioning irregular graphs


George Karypis and Vipin Kumar

SIAM Journal on Scientific Computing, Vol. 20, No. 1, pp. 359 - 392, 1999

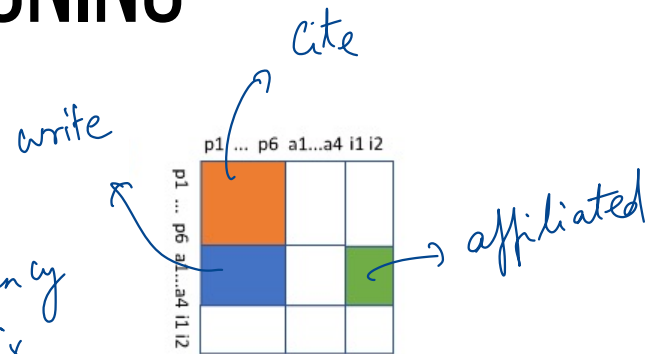
HETEROGENEOUS GRAPH PARTITIONING



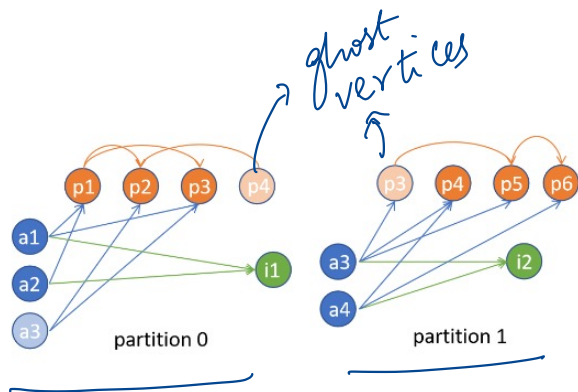
Knowledge graph



adjacency matrix



includes all edge types



METIS

GNN MINI-BATCH PREPARATION

```
for batch in training_examples:
```

```
    sample_neighbors(batch)
```

```
    load_representations(batch)
```

```
    transfer_to_GPU(batch)
```

```
    loss = model(batch)
```

```
    transfer_to_CPU(batch)
```

```
    update_parameters(batch)
```

mini-batch preparation



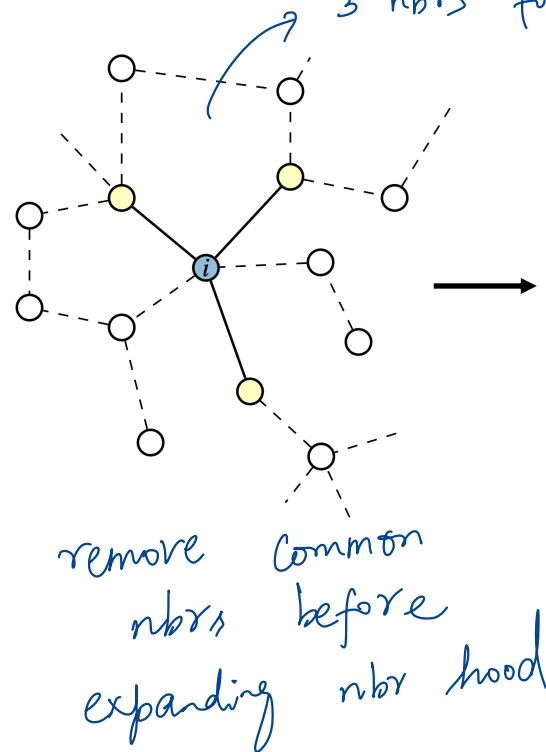

figure out which representations to load

MINIBATCH SAMPLING FOR GNNs

$b=8 \Rightarrow$ start with 8 target vertices

up to 24 vertices

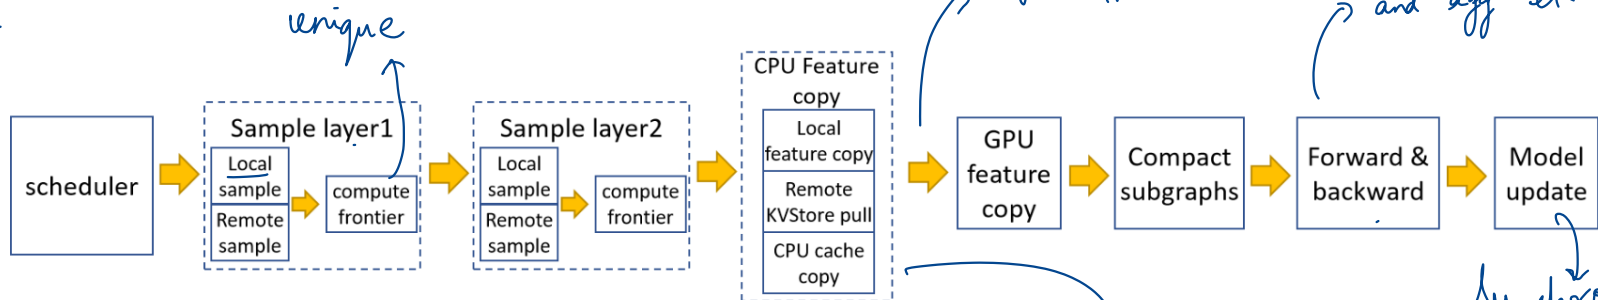
3 nbrs for each vertex \equiv



Access all of these for training

DISTDGL PIPELINE

2 layer GNN



copy features to GPU

Apply nn functions and agg etc.

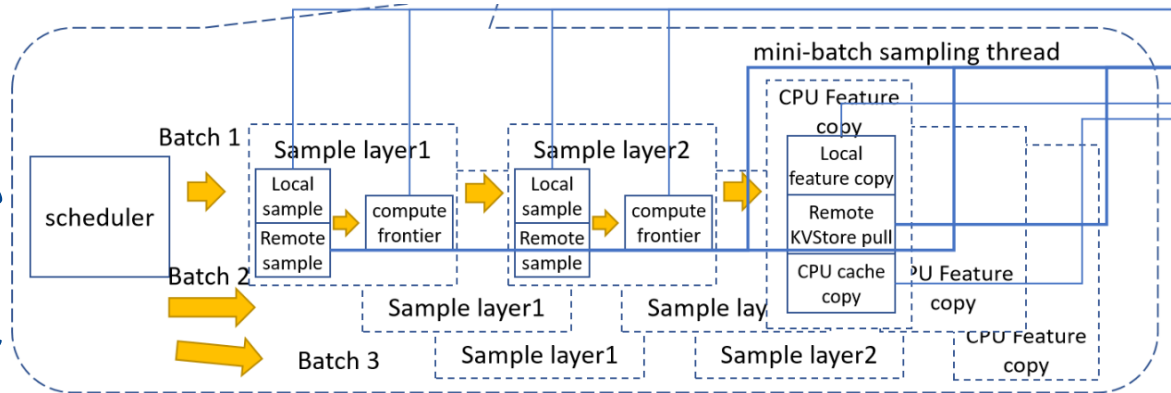
Async mini-batch sampling with sync training

Async if not updating emb.

Avoid async if updating

Synchronize weights in Dist setup

↳ Create nbr hood and feature for future mini batches
 → Sample layers with affecting sync training



SUMMARY

Graph NN: capture the structure of *graph* in *creating embeddings*

DistDGL: Distributed GNN training

Partition graphs using METIS , *hierarchical*


Pipelining to use CPU, GPU for sampling

DISCUSSION

<https://forms.gle/Dp8qtqdpWuoVeys67>

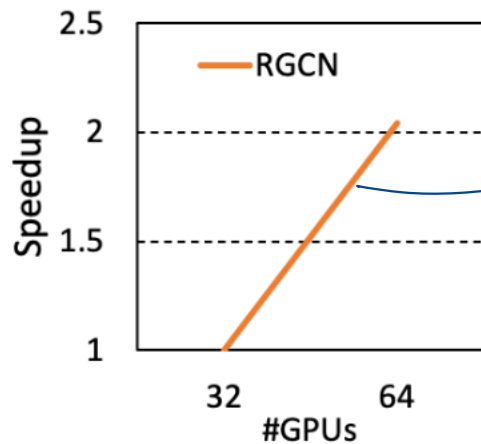
Speedup is
sub-linear

↳ CPU saturation
mini batches



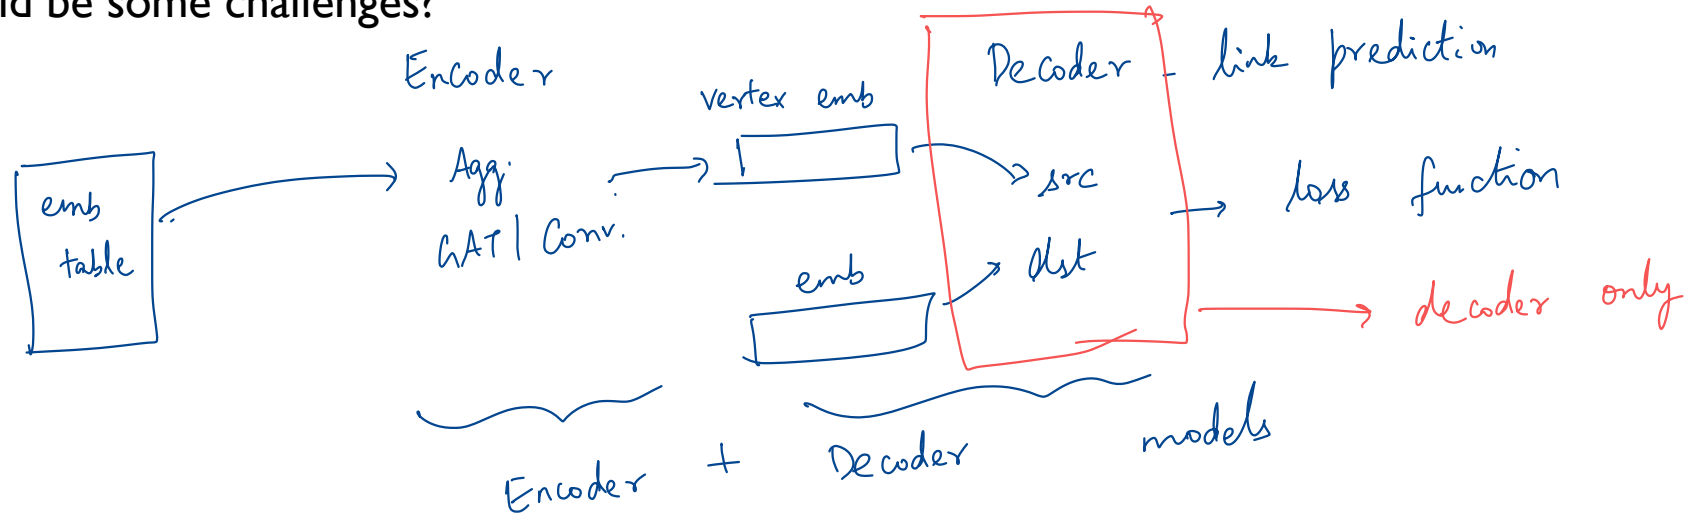
Graph Attention → more compute intensive
easier to scale

log scale



seems to
be
linear

If you wished to extend the design of Marius to support GNNs, how would you do that? What would be some challenges?



If Marius is distributed → BETA ordering. Avoid Staleness → Change Pipelining

Need to sample nbrs of given vertex → Partitioning → sample nbrs which are in cache

↳ METIS or smarter partitioning

NEXT STEPS

Next class: Serverless computing

Project check-ins by Nov 23th