

TASER: Temporal Addaptive Sampling for Fast and Accurate Dynamic Graph Representation Learning

Gangda Deng*
University of Southern California
Los Angeles, USA
gangdade@usc.edu

Hongkuan Zhou*
University of Southern California
Los Angeles, USA
hongkuaz@usc.edu

Hanqing Zeng
Meta AI
Menlo Park, USA
zengh@meta.com

Yinglong Xia
Meta AI
Menlo Park, USA
yxia@meta.com

Christopher Leung
Meta AI
Menlo Park, USA
chrisleung@meta.com

Jianbo Li
Meta AI
Menlo Park, USA
jianboli@meta.com

Rajgopal Kannan
US Army Research Lab
Los Angeles, USA
rajgopal.kannan.civ@army.mil

Viktor Prasanna
University of Southern California
Los Angeles, USA
prasanna@usc.edu

Abstract—Recently, Temporal Graph Neural Networks (TGNNs) have demonstrated state-of-the-art performance in various high-impact applications, including fraud detection and content recommendation. Despite the success of TGNNs, they are prone to the prevalent noise found in real-world dynamic graphs like time-deprecated links and skewed interaction distribution. The noise causes two critical issues that significantly compromise the accuracy of TGNNs: (1) models are supervised by inferior interactions, and (2) noisy input induces high variance in the aggregated messages. However, current TGNN denoising techniques do not consider the diverse and dynamic noise pattern of each node. In addition, they also suffer from the excessive mini-batch generation overheads caused by traversing more neighbors. We believe the remedy for fast and accurate TGNNs lies in temporal adaptive sampling. In this work, we propose TASER, the first adaptive sampling method for TGNNs optimized for accuracy, efficiency, and scalability. TASER adapts its mini-batch selection based on training dynamics and temporal neighbor selection based on the contextual, structural, and temporal properties of past interactions. To alleviate the bottleneck in mini-batch generation, TASER implements a pure GPU-based temporal neighbor finder and a dedicated GPU feature cache. We evaluate the performance of TASER using two state-of-the-art backbone TGNNs. On five popular datasets, TASER outperforms the corresponding baselines by an average of 2.3% in Mean Reciprocal Rank (MRR) while achieving an average of 5.1× speedup in training time.

Index Terms—Temporal Graph Neural Network, Adaptive Sampling, GPU

I. INTRODUCTION

Dynamic graphs are natural abstractions of time-stamped interactions in many real-world systems. Interacting entities are represented as nodes, while interactions are represented as time-stamped edges. Generating low-dimensional node representations on dynamic graphs (i.e., dynamic graph representation learning) is a fundamental problem for many practical systems, as it allows for monitoring and predicting the evolution of real-world data such as social networks, transportation networks, and financial networks. Researchers

have recently proposed various Temporal Graph Neural Networks (TGNNs) [1]–[6] to learn time-evolving patterns on dynamic graphs. Unlike static graph representation learning approaches [7]–[9] that only accept time-invariant graphs as the input, TGNNs incorporate temporal information jointly with structural and contextual information into low-dimensional embeddings, which have shown superior performance in various real-world applications, including recommendation [10], event prediction [11], and fraud detection [12].

Graph Neural Networks (GNNs), both static and temporal, recursively gather and aggregate information from neighboring nodes to generate node embeddings. To reduce the high computation and memory footprint, neighbor sampling approaches [8], [13]–[15] are widely used to alleviate the exponentially growing neighbor size with respect to the number of GNN layers. However, most sampling methods approximate full neighborhood training using a static distribution, which is agnostic to the node/edge features, model architecture, and task performance. These sampling policies are vulnerable to noise since they cannot distinguish between relevant and irrelevant neighbors, leading to a large sampling variance. To address these issues, researchers have designed adaptive sampling methods [16]–[19], where the sampling distribution is node-dependent and guided by the task performance. On static graphs, these methods, which come with theoretical guarantees for variance reduction and are adaptive to performance, can generate high-quality and robust node embeddings.

Dynamic graphs, much like their static counterparts, are not immune to the presence of noise, which adds false and irrelevant information to the graph signals. Specifically, dynamic graphs introduce two distinctive types of noise: (1) **Deprecated links**. Dynamic graphs accumulate an increasing number of interactions over time. Some old interactions could be irrelevant or even convey information that contradicts the current node status. (2) **Skewed neighborhood distribution**. The distribution of interactions among different nodes in a dynamic graph often exhibits significant disparities or sparsity. Unlike static graphs, temporal graphs have many repeated

*Equal contribution

edges between the same two nodes at different timestamps. A long-standing node may exhibit a skewed distribution of neighbors, while an emerging node may have very few neighbors. For example, deprecated links can be observed in a social network graph when a person relocates to another country, rendering the previous connections gradually less informative or even incorrect. Skewness becomes evident when an individual engages in daily conversations with their best friend while sending only a single message to a car dealer. The noise in dynamic graphs causes two critical issues that significantly impair the accuracy of TGNNs. Firstly, when performing self-supervised training with the link prediction task, inferior interactions are used as positive links. Secondly, it is amplified in the iterative message-passing process, leading to high variance in the output embeddings.

To improve the performance of TGNNs on dynamic graphs with temporal and structural noise, researchers have proposed denoising techniques based on edge dropping [20], [21] and implemented human-defined heuristics [3], [4] to non-adaptive TGNN samplers. However, these approaches require extensive tuning and often achieve worse performance since they assume the whole graph follows the same noise pattern and ignore the differences in different nodes at different timestamps. For example, TGAT [3] employed the inverse timespan sampler to solve the deprecated links problem, which samples past neighbors with probabilities inversely proportional to their time deltas, but found that it performs worse than the original uniform sampler. Adaptive sampling, on the other hand, could learn customized sampling probabilities, encompassing any human-defined heuristics that may exist since the learnable sampler considers not only dynamic graph information but also training dynamics and task performance. Therefore, we believe that adaptive sampling is integral to any approach addressing the noise problem in TGNNs, given its comprehensive consideration of all available information sources when estimating personalized neighborhood sampling probability distributions.

Despite the urgent need for adaptive sampling in TGNNs, it is notably challenging to construct an efficient and reliable solution. We identify the three main challenges as follows: (1) To capture the dynamics in Temporal Graphs, the adaptive sampler should project not only structural and contextual information into sampling probabilities but also the time and frequency of the interactions. (2) Existing adaptive sampling methods only support co-training with simple and static GNN aggregators and cannot be generalized to particularly complex temporal aggregators. (3) Adaptive samplers require traversing a large and time-restricted neighborhood, resulting in enormous training time, especially when scaling to large-scale datasets. Specifically, when the number of traversed neighbors increases, the mini-batch generation overheads (i.e., temporal neighbor finding and feature slicing with CPU-GPU data transfer) lead to an order-of-magnitude increase in the training time. Figure 1 shows the runtime breakdown of TGAT when the receptive field increases. On both datasets, the mini-batch generation time dominates the training time. Besides, adaptive sampling requires encoding node/edge features with learnable




Fig. 1. Runtime (per epoch) breakdown for TGAT with different numbers of neighbors per layer. Prep. refers to the mini-batch generation time (neighbor finding, feature slicing, and CPU-GPU data transferring), while Prop. refers to the propagation time (forward and backward propagation).

weights. Due to this compute-intensive nature, achieving fast adaptive sampling necessitates training on the GPU as well as specific GPU optimizations to alleviate the mini-batch generation bottleneck.

To overcome the aforementioned challenges, we propose TASER, the first efficient and scalable adaptive sampling method for TGNNs. TASER provides a general solution for adaptive sampling in TGNNs and supports most TGNNs designed for Continuous Time Dynamic Graphs (CTDGs). To mitigate the mini-batch generation bottleneck, TASER implements a pure GPU-based temporal neighbor finder and a dedicated GPU feature cache. Our main contributions are:

- We propose a novel two-fold temporal adaptive sampling technique — temporal adaptive mini-batch selection (Section III-A) and temporal adaptive neighbor sampling (Section III-B). Temporal adaptive mini-batch selection selects high-quality training samples, while temporal adaptive neighbor sampling selects high-quality supporting neighbors.
- We implement the first GPU neighbor finder (Section III-C) for dynamic graphs, which is optimized for the massive Single Instruction Multiple Data (SIMD) GPU architecture. Compared with a state-of-the-art CPU neighbor finder, our GPU neighbor finder supports arbitrary training order while achieving an average speedup of $46\times$.
- We design a dynamic GPU cache (Section III-D) to speed up the feature-slicing process for large-scale datasets that cannot be fully stored on the GPU VRAM. Our cache replacing policy achieves near-optimal performance and requires minimal maintenance overhead.
- In the experiments, we implement TASER on two state-of-the-art backbone TGNNs. On five popular datasets, TASER outperforms the corresponding baselines by an average of 2.3% in MRR. Our GPU neighbor finder and 20% GPU feature cache achieve an average of $5.1\times$ speedup in the total training time.

II. BACKGROUND

Dynamic graphs can be represented as a series of time-stamped graph events. In this work, we consider the most common dynamic graphs with edges appearing as graph events. Without loss of generality, consider a dynamic graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with events $\{(u, v, \mathbf{x}_{uv}, t)\}$, where each quadruplet represents an edge with edge feature \mathbf{x}_{uv} appearing from node u to node v at time t . The goal of TGNNs is to generate information-rich dynamic embeddings for nodes at given timestamps. The embeddings can be further used in different downstream tasks (e.g., clustering, node classification). Note that TGNNs are usually trained self-supervised with the dynamic link prediction task (i.e., distinguish negative edges from positive edges) [3]–[5]. Let \mathcal{E}_B be a mini-batch of edges sampled in the training set. For each edge $(u, v, t) \in \mathcal{E}_B$, a negative destination node v' is randomly sampled from \mathcal{V} to form a negative edge (u, v', t) . Then, we sample supporting neighbors from the corresponding temporal neighborhood for these root nodes (i.e., u , v , and v') and iteratively apply a series of temporal aggregators to compute their dynamic node embeddings. Lastly, these embeddings are fed into an edge predictor where the binary cross entropy loss is used to perform an iteration of the Stochastic Gradient Descent (SGD) process.

A. Neighbor Finder

For node v at time t , we consider its temporal neighborhood $\mathcal{N}(v, t) = \{(u, t_u) \mid (v, u, t_u) \in \mathcal{E}, t_u < t\}$. To avoid the monotonically increased neighborhood size, TGNNs introduce a neighbor finder to select a subset of neighbors $\mathcal{N}_s(v, t) \subseteq \mathcal{N}(v, t)$ with a fixed size $b_N = |\mathcal{N}_s(v, t)|$. Here, we introduce two neighbor finders that are widely used by existing TGNNs:

- Uniform Neighbor Finder ensures supporting nodes $\mathcal{N}_s(v, t)$ are sampled uniformly, deriving an unbiased approximation of the original neighborhood.
- Most Recent Neighbor Finder only samples the most recent neighbors, ensuring that the latest interactions between nodes are prioritized.

B. Temporal Aggregator

For a target node v at time t , its neighborhood $\mathcal{N}_s(v, t)$, and their input hidden features, a temporal aggregator performs the following two steps: (1) encode time-aware embedding vector for each neighbor $(u, t_u) \in \mathcal{N}_s(v, t)$, and (2) combine these vectors as the output vector. To encode temporal information, existing methods map continuous timestamps to a d_T -dimensional vector space using $\Phi: T \rightarrow \mathbb{R}^{d_T}$. The l -th layer dynamic node embedding $\mathbf{h}_v^{(l)}$ for node v at time t can be computed as follows:

$$\mathbf{m}_u^{(l)} = \left\{ \mathbf{h}_u^{(l-1)} \parallel \mathbf{x}_{uv} \parallel \Phi(\Delta t) \right\} \quad (1)$$

$$\mathbf{h}_v^{(l)} = \text{COMB} \left(\mathbf{m}_v^{(l)}, \left\{ \mathbf{m}_u^{(l)} \mid u \in \mathcal{N}_s(v, t) \right\} \right) \quad (2)$$

where $\Delta t = t - t_u$ and COMB is the combiner that aggregates all the related messages $\mathbf{m}_u^{(l)}$. Note that if node v interacts with

node u at different timestamps, we generate different messages sharing the same node embedding $\mathbf{h}_u^{(l-1)}$. Here, we introduce the two widely-used temporal aggregators:

- TGAT [3] is a self-attention-based aggregator that uses a learnable time encoding Φ :

$$\Phi(\Delta t) = \cos(\Delta t \mathbf{w} + \mathbf{b}), \quad (3)$$

where $\mathbf{w} \in \mathbb{R}^{d_T}$ and $\mathbf{b} \in \mathbb{R}^{d_T}$ are learnable parameters. Let $\mathbf{M}_u^{(l)}$ be the (l) -th layer message matrix for temporal neighborhood $\mathcal{N}_s(v, t)$. The COMB function of TGAT performs

$$\mathbf{q}^{(l)} = \mathbf{W}_q \left\{ \mathbf{h}_v^{(l-1)} \parallel \Phi(0) \right\} + \mathbf{b}_s \quad (4)$$

$$\mathbf{K}^{(l)} = \mathbf{W}_k \mathbf{M}_u^{(l)} + \mathbf{b}_k \quad (5)$$

$$\mathbf{V}^{(l)} = \mathbf{W}_v \mathbf{M}_u^{(l)} + \mathbf{b}_v \quad (6)$$

$$\mathbf{h}_v^{(l)} = \text{Softmax} \left(\frac{\mathbf{q}^{(l)} \mathbf{K}^{(l)\text{T}}}{\sqrt{|\mathcal{N}_s(v, t)|}} \right) \mathbf{V}^{(l)}. \quad (7)$$

- GraphMixer [5] provides a technically simple architecture with comparable performance to RNN-based and self-attention-based methods. It uses a fixed time-encoding

$$\Phi(\Delta t) = \cos(\Delta t \boldsymbol{\omega}), \quad \boldsymbol{\omega} = \left\{ \alpha^{-(i-1)/\beta} \right\}_{i=1}^{d_T}, \quad (8)$$

followed by a 1-layer MLP-Mixer [22] aggregator to combine messages from neighboring nodes:

$$\mathbf{h}_v^{(l)} = \text{Mean} \left(\text{MLP-Mixer} \left(\mathbf{M}_u^{(l)} \right) \right). \quad (9)$$

C. Related Works

The recent success of practical GNN applications is attributed to their ability to quickly and accurately learn graph representations. Techniques such as graph denoising and GPU accelerations enable GNNs to efficiently process noisy real-world data at large scales.

Dynamic Graph Denoising. Existing dynamic graph denoising techniques are mainly based on dynamic graph sparsification. TGAT [3] proposes a heuristics sampling policy based on the probability of inversed timespan. TGN [4] further improves the timespan inversed sampling by sampling the most recent neighbors. To avoid redundancy, TNS [23] proposes to insert learnable spaces in the most recent neighbors. STEP [21] proposes an unsupervised graph pruning method that drops the noisy interactions. TGAC [24] devises dynamic graph augmentation techniques for contrastive TGNN learning, which measures the edge sample probability by computing the PageRank or Eigenvector of nodes in both ends. However, all these methods do not consider the disparity of noise among different nodes at different times.

Adaptive Mini-Batch Selection. Adaptive mini-Batch selection, commonly referred to as adaptive importance sampling, constantly re-evaluates the relative importance of each training sample during training. The main idea behind these methods is to use gradient information to reduce variance in uniformly stochastic gradients in order to improve convergence. GRAD [25] relies on both features and logits for solving

least-squares problems. [26] proposed a general algorithm with efficient computation to speed up coordinate-descent and SGD. MVS [27] extends these methods to GNNs, considering both the variance introduced by mini-batch selection and neighbor sampling. In contrast to reducing variance and speeding up optimization, TGNN training requires avoiding selecting noisy interactions as positive samples.

Adaptive Neighbor Sampling. As one of the graph denoising techniques, adaptive neighbor sampling methods learn a sample probability distribution for each neighboring node of a given target node. AS-GCN [16] minimizes the GCN sampling variance by training a self-dependent function based on node features. Bandit Sampling [17] formulates the variance reduction for adaptive sampling as an adversary bandit problem, and Thanos [19] further proposes a biased reward function to avoid instability. In contrast to variance reduction, PASS [18] directly optimizes task performance by approximating gradient propagation through the non-differentiable sampling operation of GCN. To scale to large graphs, PASS adopts a two-step sampling approach, which first samples a fixed scope and then adaptively samples the neighbors within the scope. However, these adaptive sampling methods can not capture the temporal information of dynamic graphs and are not compatible with temporal aggregators.

Neighbor Finding. Optimized GPU graph neighbor finders could achieve orders-of-magnitude higher throughput compared with CPU neighbor finders by leveraging the massive SIMD architecture and avoiding the data transfer overheads from CPU to GPU. DGL [28] provides an easy-to-use GPU neighbor finder with unified virtual memory access support, demonstrating a speedup of $1.5\times$ to $3.9\times$ in total training time compared to the pipeline, which samples on the CPU and trains on the GPU. Quiver [29] further proposes a workload-based scheduler that dynamically assigns tasks to the CPU and GPU to solve the imbalanced workload problem due to the unpredictable latency when working on sparse nodes. Biased (weighted) neighbor finding based on inverse transformation sampling [30], rejection sampling [31], and alias method [32] are also well studied on GPUs. However, they don't work on dynamic graphs and can not be used for temporal neighborhood sampling. TGL [33] proposes the T-CSR data structure and a parallelized neighbor finder optimized for dynamic graphs on multi-core CPUs. Its key limitation is the reliance on pointer arrays for rapidly locating candidate temporal neighbors, which requires scheduling the training mini-batches chronologically. Besides, Tea [34] is a state-of-the-art general-purpose CPU random walk engine for biased neighbor finding on dynamic graphs. However, it does not support high-dimensional feature transformation, which adaptive sampling requires.

Graph Feature Caching. The neighbor explosion problem [8] causes an enormous number of memory operations to fetch the node and edge features. On large graphs whose entire node and edge feature matrices cannot be stored in GPU VRAM, the CPU-GPU feature slicing and loading process easily becomes the bottleneck during training. GNS [35] addresses this issue

by periodically selecting a global set of nodes for all mini-batches and caching their features on GPU. Data Tiering [36] uses reverse PageRank to predict the access probability of each node. Quiver [29] further proposes a connectivity-aware node feature caching strategy that considers the probability of a node being sampled as a multi-hop neighbor. However, these approaches are designed for the memory access pattern of static GNNs and do not consider temporal information.

III. APPROACH

In this section, we present TASER, a high-performance temporal adaptive sampling method for TGNNs. An overall illustration of one mini-batch training for TASER on a 1-layer TGNN is shown in Fig. 2. First, we maintain an importance score for each training sample, enabling the adaptive selection of a batch of high-quality samples in each step. Next, we adopt the bi-level neighbor sampling scheme used in PASS [18] to improve the performance on large graphs. Initially, a GPU temporal neighbor finder samples a set of candidate neighbors from the temporal neighborhood $\mathcal{N}(v, t)$ using a static policy. Then, we slice features from both the GPU cache and CPU memory, where the GPU cache is updated at the end of every epoch. Following this, a parameterized temporal adaptive neighbor sampler is applied to sample a fixed-size set of informative supporting neighbors from the pre-sampled neighborhood. Finally, the TGNN model is trained on the representative node samples with their denoised supporting neighborhoods. We further update the sample importance score $\mathcal{P}(v)$ and the sampler's parameter θ during forward and backward propagation, respectively.

The rest of the section is arranged as follows. We first propose our two-fold adaptive sampling technique regarding the mini-batch sample selection in Section III-A and the supporting neighbor sampling in Section III-B. Then, we propose the pure-GPU neighbor finder in Section III-C and the dynamic GPU cache in Section III-D.

A. Temporal Adaptive Mini-batch Selection

In order to capture the pattern of node states changing over time, TGNNs are trained on interactions that cover the entire training set. Unlike training GNNs on static graphs, where the models only recover the final states of different nodes, TGNNs need to recover different states for the same node during training. However, learning to recover deprecated or cold-start states may significantly impair the accuracy of TGNNs. To reduce the noise present in the training samples, we propose a temporal adaptive mini-batch selection method that utilizes the dynamic model predictions to sample high-quality training edges.

We first recall the original mini-batch SGD training process of TGNNs. Given a training set \mathcal{E}_{train} , we chronologically sample a subset of edges $\mathcal{E}_B \subseteq \mathcal{E}_{train}$ for each batch. For each training edge $e(v_1, v_2, t) \in \mathcal{E}_B$, we set its label $y_e = 1$ and randomly sample a destination node $v'_2 \in \mathcal{V}$ to form a negative edge $e'(v_1, v'_2, t)$ with label $y'_e = 0$. During the forward propagation, we derive the logits \hat{y}_e and $\hat{y}_{e'}$ for each




Fig. 2. One training iteration of TASER on a one-layer TGNN. (a) Randomly select a set of mini-batch samples based on the pre-computed importance score \mathcal{P} proportional to the logits (temporal adaptive mini-batch selection). (b) Sample a subset of neighbors from the temporal neighborhood using our GPU temporal neighbor finder. (c) Slice the features of sampled neighbors from the VRAM cache and RAM. (d) Apply temporal adaptive neighbor sampling (parameterized by θ) to sub-sample the supporting neighbors for TGNN by encoding timestamps, frequencies, and identities along with features. (e) Perform forward and backward propagation. Update the importance score \mathcal{P} for adaptive mini-batch selection and back-propagate through the model loss and sample loss to train the TGNN model and temporal adaptive sampler.

positive and negative edge, respectively. After that, the model loss is computed as follows:

$$\mathcal{L}_{\text{model}} = -\frac{1}{|\mathcal{E}_{\mathcal{B}}|} \sum_{e \in \mathcal{E}_{\mathcal{B}}} \phi(\hat{y}_e, y_e) + \phi(\hat{y}_{e'}, y_{e'}), \quad (10)$$

where ϕ stands for the cross entropy loss function.

Unlike the original method that chronologically picks mini-batch samples, TASER adaptively selects mini-batch samples based on training dynamics. First, we maintain a list of importance scores $\mathcal{P} \in \mathbb{R}^{|\mathcal{E}_{\text{train}}|}$ to evaluate the noisiness of each edge sample, which is initialized uniformly. Then, we randomly sample a batch of training edges $\mathcal{E}_{\mathcal{B}}$ with the probability proportional to corresponding importance scores. As shown in Figure 2 (e), after the forward propagation, TASER update the importance score $\mathcal{P}(e)$ for every positive sample $e(v_1, v_2, t) \in \mathcal{E}_{\mathcal{B}}$ with

$$\mathcal{P}(e) = \text{sigmoid}(\hat{y}_e) + \gamma, \quad (11)$$

where γ is a hyperparameter representing the magnitude of a uniform distribution mixed with the adaptive sample importance distribution.

Since dynamic graphs tend to be significantly noisier than static graphs, selecting positive samples with high confidence \hat{y}_e can effectively improve accuracy. Given the cross entropy loss function, the gradient of the loss with respect to logits is inversely proportional to the logits. When the gradient update of a sample is large, it indicates that the sample is informative, but it is also more likely to be an outlier in the data distribution. To balance the noisiness and diversity in training samples, we can adjust the value of γ . A larger value of γ makes the selector prone to sample noisier samples to amplify training. In practice, we found that $\gamma = 0.1$ works well on all the datasets.

B. Temporal Adaptive Neighbor Sampling

Existing adaptive neighbor samplers only support attributed static graphs, which do not consider time restrictions and fail to distinguish recurring interactions in the temporal neighborhood. They are specifically designed for a single type of aggregator and cannot achieve high accuracy when extended to other aggregators. To address these issues, we propose a general encoder-decoder scheme that is suitable for various graph data and temporal aggregators. Fig. 2 (d) illustrates the forward propagation of our temporal adaptive neighbor sampler. Formally, given a temporal neighborhood $\mathcal{N}_s(v_i, t_0)$, TASER adaptively computes the sample policy $q_{\theta}(u_j, t_k | v_i, t_0)$ that estimates the probability of sampling neighbors $(u_j, t_k) \in \mathcal{N}_s(v_i, t_0)$ given node v_i at time t_0 .

Neighbor Encoder. Auxiliary information must be incorporated to discriminate the unique noise patterns in dynamic graphs, including outdated or redundant interactions. To generate a time-aware sample policy, we use the fixed time-encoding function $TE_{(\Delta t)}$ proposed by GraphMixer [5] to encode temporal information for each neighbor, as shown in Eq. (8). $TE_{(\Delta t)}$ maps the relative timespan from the continuous time domain to a d_{time} -dimensional vector space. Besides being perceptual of time, the sampler also requires to distinguish the reappearance of neighboring nodes to differentiate redundant neighbors. We propose the frequency encoding by leveraging the sinusoidal encoding [37]:

$$\begin{aligned} FE_{(\text{freq}(u), 2i)} &= \sin\left(\text{freq}(u)/10000^{2i/d_{\text{freq}}}\right) \\ FE_{(\text{freq}(u), 2i-1)} &= \cos\left(\text{freq}(u)/10000^{2i/d_{\text{freq}}}\right) \end{aligned} \quad (12)$$

where $\text{freq}(u)$ denotes the frequency of a neighbor node u appearing in the neighborhood $\mathcal{N}_s(v, t_0)$. Since the frequency

is indeed discrete and has limited values, we choose the positional encoding (i.e. sinusoidal encoding) instead of the time encoding to encode frequency. However, when two nodes exhibit the same appearance frequency, the sampler remains unable to distinguish between them. To address this limitation, we propose the identity encoding $IE_{(u_j)}$. Given a sorted neighbor list $\{(u_1, t_1), (u_2, t_2), \dots, (u_{b_{\mathcal{N}}}, t_{b_{\mathcal{N}}})\}$ of $\mathcal{N}_s(v, t_0)$, where $t_0 > t_1 > \dots > t_{b_{\mathcal{N}}}$, we defined the identity encoding for each neighbor as:

$$IE_{(u_j, i)} = \mathbb{1}_{(u_j = u_i)}, \quad i = 1, 2, \dots, b_{\mathcal{N}}. \quad (13)$$

In addition to these three encodings, we incorporate the contextual information of nodes and edges. For each neighbor $(u, t) \in \mathcal{N}_s(v, t_0)$, we align the dimensions of node feature \mathbf{x}_u and edge feature \mathbf{x}_{vut} to d_{feat} :

$$\mathbf{h}_{(u)} = \text{GeLU}(\mathbf{W}_n \mathbf{x}_u), \quad \mathbf{h}_{(v, u, t)} = \text{GeLU}(\mathbf{W}_e \mathbf{x}_{vut}). \quad (14)$$

Finally, we concatenate all the encodings as well as features to derive the neighbor embedding as the input of the decoder:

$$\mathbf{z}_{(u, t)} = \{\mathbf{h}_{(u)} || \mathbf{h}_{(v, u, t)} || TE_{(\Delta t)} || FE_{(freq(u))} || IE_{(u)}\}. \quad (15)$$

To ensure a balanced impact from various information sources, we set the dimensions to $d_{\text{feat}} = d_{\text{time}} = d_{\text{freq}}$ across all datasets. The dimension of neighbor embedding $\mathbf{z}_{(u, t)}$ is denoted as d_{enc} .

Neighbor Decoder. After we encode the unique characteristics of dynamic graphs into the neighbor embedding $\mathbf{z}_{(u, t)}$, the subsequent processes can be modeled as a general adaptive neighbor sampling problem. For simplicity, we omit the timestamps and also do not differentiate the recurrence of interactions in a neighborhood. The goal of the neighbor decoder is to generate a customized neighborhood importance distribution $q(u|v)$ for each neighborhood $\mathcal{N}_s(v)$. Given that dynamic graphs often lack node features, rather than learning an exact pair-wise importance score for aggregation as GAT [9], we are more interested in estimating the relative importance of a node within the neighborhood $q(u|\{u', u' \in \mathcal{N}_s(v)\})$. Therefore, we use a 1-layer MLP-Mixer [22] to first transform the hidden embedding dimension and then transform the neighbor dimension for each neighborhood:

$$\mathbf{Z}_{\mathcal{N}_s(v)} = \text{MLP-Mixer} \left(\{\mathbf{z}_{u_1}, \mathbf{z}_{u_2}, \dots, \mathbf{z}_{u_{b_{\mathcal{N}}}}\} \right), \quad (16)$$

where $\mathbf{Z}_{\mathcal{N}_s(v)} \in \mathbb{R}^{b_{\mathcal{N}} \times d_{\text{enc}}}$. In doing so, the neighbor embedding not only depends on the global transformation but also captures the neighborhood correlations. To coordinate with different temporal aggregators, our neighbor decoder supports various predictors [9], [37], [38], including

$$q_{\text{linear}}(u|v) = \sigma_u(\mathbf{w}_l \mathbf{Z}_{\mathcal{N}_s(v)}), \quad (17)$$

$$q_{\text{gat}}(u|v) = \sigma_u(\text{LeakyReLU}(\mathbf{a}^T \cdot [\mathbf{W}_g \mathbf{z}_u || \mathbf{W}_g \mathbf{z}_v])) \quad (18)$$

$$q_{\text{gatv2}}(u|v) = \sigma_u(\mathbf{a}^T \text{LeakyReLU}(\mathbf{W}_{g2} \cdot [\mathbf{z}_u || \mathbf{z}_v])) \quad (19)$$

$$q_{\text{trans}}(u|v) = \sigma_u \frac{(\mathbf{W}_t \mathbf{z}_v) (\mathbf{W}'_t \mathbf{Z}_{\mathcal{N}_s(v)})^T}{\sqrt{b_{\mathcal{N}}}}, \quad (20)$$

Algorithm 1: TASER Training: One Iteration

Require: a minibatch of labeled edges $\{e_k, y_k\}_{k=1}^b$,
neighbor finding budget m , neighbor
sampling budget n , L -layer TGNN model f ,
adaptive neighbor sampler $q(j|i)$

Ensure : updated TGNN model and adaptive neighbor
sampler

- 1 $V_{\text{act}} \leftarrow$ all nodes in $\{e_k\}_{k=1}^b$;
- 2 $\mathcal{G} \leftarrow$ empty supporting neighbor set;
- 3 **for** $l \leftarrow L$ **to** 1 **do**
- 4 **for** $(v_i, t) \in V_{\text{act}}$ **do**
- 5 $\mathcal{N}_s \leftarrow \{u_j\}_{j=1}^m$ sampled from $\mathcal{N}(v_i, t)$;
- 6 $\mathcal{N}'_s \leftarrow \{u_j\}_{j=1}^n$ sampled from \mathcal{N}_s with $q(j|i)$;
- 7 update V_{act} and $\mathcal{G}_{[s]}$ with \mathcal{N}'_s ;
- 8 **end for**
- 9 **end for**
- 10 $\mathcal{L}_{\text{model}} \leftarrow$ loss($\{f(e_k, \mathcal{G}_{[k]}), y_k\}_{k=1}^b$);
- 11 update f by back-propagating $\mathcal{L}_{\text{model}}$;
- 12 $\mathcal{L}_{\text{sample}} \leftarrow$ construct loss following Eq.(25) or Eq.(26);
- 13 update $q(j|i)$ by back-propagating $\mathcal{L}_{\text{sample}}$;
- 14 **return** f and $q(j|i)$;

where σ is the softmax function. For the target node embedding \mathbf{z}_v , we concatenate the node feature (if it exists) with zero time encoding and one frequency encoding.

$$\mathbf{z}_v = \{\mathbf{h}_{(v)} || TE_{(0)} || FE_{(1)}\}. \quad (21)$$

Empirically, we observed that this target node embedding works well even without node features.

Co-Training with Temporal Aggregators. Figure 2 (e) demonstrates a temporal aggregator that combines the messages from sampled neighbors during the forward pass and subsequently back-propagates based on $\mathcal{L}_{\text{model}}$. However, the parameter of the sampling policy $q_{\theta}(u|v)$ can not be directly updated through back-propagation since the sampling process is non-differentiable. We need to construct an auxiliary loss $\mathcal{L}_{\text{sample}}$ to update θ . For an arbitrary temporal aggregator with neighbors sampled following $q_{\theta}(\cdot|v)$, we can rewrite the forward propagation in the form of:

$$\mathbf{h}_{v_i}^{(l)} = g^{(l)} \left(\left\{ \mathbb{E}_{q_{\theta}(u_j|v_i)} [f(v_i, u_j)], f \in \mathcal{H}^{(l)} \right\} \right), \quad (22)$$

where $g^{(l)}$ and $\mathcal{H}^{(l)}$ are functions defined by the temporal aggregator at the l -th layer. This implies that every appearance of the expectation should be considered when calculating $\nabla_{\theta} \mathbf{h}_{v_i}^{(l)}$. For simplicity, we concisely denote $q_{\theta}(u_j|v_i)$ as $q_{\theta}(u_j)$ and $f(v_i, u_j)$ as $f(u_j)$ below. We can then approximate each $\nabla_{\theta} \mathbb{E}_{q_{\theta}(u_j)} [f(u_j)]$ using the log-derivative trick [39] with n Monte Carlo samples $\{u_j \sim q_{\theta}(u_j)\}_{j=1}^n$:

$$\nabla_{\theta} \mathbb{E}_{q_{\theta}(u_j)} [f(u_j)] \approx \frac{1}{n} \sum_{j=1}^n \nabla_{\theta} \log q_{\theta}(u_j) f(u_j). \quad (23)$$

Next, we show how to calculate $\nabla_{\theta} \mathcal{L}_{\text{model}}$ when co-training with temporal aggregators. For the TGAT aggregator shown in

Eq. (7), denote $a(v_i, u_j)$ as the unnormalized attention score and $\tau_{i,j} = e^{a(v_i, u_j)}$. The TGAT aggregator can be transformed in the form of Eq. (22) as:

$$\mathbf{h}_{v_i}^{(l)} = \mathbb{E}_{q_\theta(u_j)} [f_1(u_j)] / \mathbb{E}_{q_\theta(u_j)} [f_2(u_j)], \quad (24)$$

where $f_1(u_j) = \tau_{i,j}[\mathbf{V}^{(l)}]_j$ and $f_2(u_j) = \tau_{i,j}$. According to the chain rule and Eq. (23),

$$\begin{aligned} \nabla_\theta \mathcal{L}_{\text{model}} \approx & \frac{d\mathcal{L}_{\text{model}}}{d\mathbf{h}_{v_i}^{(l)}} \cdot \frac{1}{\lambda^3 n} \sum_{j=1}^n \tau_{i,j} [\mathbf{V}^{(l)}]_j \nabla_\theta \log q_\theta(u_j) \\ & + \frac{d\mathcal{L}_{\text{model}}}{d\mathbf{h}_{v_i}^{(l)}} \cdot \frac{\boldsymbol{\mu}}{\lambda^4 n} \sum_{j=1}^n \tau_{i,j} \nabla_\theta \log q_\theta(u_j), \end{aligned} \quad (25)$$

where $\lambda = \sum_{j=1}^n \tau_{i,j}$, $\boldsymbol{\mu} = \sum_{j=1}^n \tau_{i,j} [\mathbf{V}^{(l)}]_j$. Similarly, for the GraphMixer aggregator,

$$\nabla_\theta \mathcal{L}_{\text{model}} = \frac{d\mathcal{L}_{\text{model}}}{d\mathbf{h}_{v_i}^{(l)}} \cdot \frac{1}{n} \left\{ \sum_{j=1}^n w'_{jk} \mu_{jk} \nabla_\theta \log q_\theta(u_j) \right\}_{k=1}^{d_{\text{enc}}}, \quad (26)$$

where $\mu_{jk} = \mathbf{w}_k^\top \mathbf{h}_{u_j}^{(l-1)}$ and $u_j \sim q_\theta(u_j)$.

Based on Eq. (25) and Eq.(26), we can construct the sample loss $\mathcal{L}_{\text{sample}}$ by freezing the terms except for the log probability $\log q_\theta(u_j)$, and leveraging the autograd mechanism built in deep learning frameworks to update θ . Algorithm 1 shows one iteration of TASER training on a L -layer TGNN.

Remark. (*Adaptive sampling vs. Attention*) *Contrary to the bottom-up approaches of graph attention, our temporal adaptive sampler computes in a top-down manner, which does not require any hidden features $\mathbf{h}_u^{(l')}$ when generating sampling probabilities for nodes at layer l ($l' < l$). Under a fixed-size scope, it reduces the computational complexity exponentially w.r.t. the number of layers.*

C. GPU Temporal Neighbor Finding

Neighbor finding on dynamic Graphs is complex as the neighborhood $\mathcal{N}(v, t)$ for each node varies over time. To rapidly identify the candidate neighbor set, we store dynamic graphs in the T-CSR data structure [33], which sorts the outgoing neighbors according to their timestamps. As shown in Algorithm 2, we employ a block-centric parallel sampling design to leverage the hierarchical GPU architecture. Specifically, each target node is allocated a thread block, and each thread inside the block is assigned to sample a neighbor for the target node. We first identify the pivot pointer in each neighborhood using binary search with a single thread and then use the shared memory bitmap [30] for collision detection in uniform sampling without replacement. After each thread selects a neighbor, an atomic compare-and-update operation is performed to detect whether this neighbor has been selected. This block-centric design has three major benefits. Firstly, self-supervised TGNN training with TASER necessitates a large number of supporting neighbor candidates for thousands of mini-batch samples in each training iteration. The block-wise design can efficiently saturate GPU resources while avoiding intra-warp scheduling overhead. Secondly, the threads in the

Algorithm 2: GPU Temporal Neighbor Finding

Input : target nodes $\{(v_i, t'_i)\}_{i=1}^b$, neighbor budget m , T-CSR graph \mathcal{G}

Output: sampled neighbors $neigh$

```

1 for block  $i \leftarrow 1$  to  $b$  do in parallel
2   for thread  $j \leftarrow 1$  to  $m$  do in parallel
3     if  $j = 1$  then
4        $\{(u_k, t_k)\}_{k=1}^{d_{v_i}} \leftarrow \mathcal{G}_{[v_i]}$ ;
5        $p \leftarrow \text{BinarySearch}(\{t_1, \dots, t_{d_{v_i}}\}, t'_i)$ ;
6     end if
7     SyncThreads();
8     if most recent neighbor finding then
9        $neigh[i][j] \leftarrow (u_{p-j}, t_{p-j})$ ;
10    else if uniform neighbor finding then
11      initialize bitmap  $\mathcal{M}$ ;
12      SyncThreads();
13      keep randomly selecting  $r \in [1, p)$  until
14         $\text{CheckBitmap}(r, \mathcal{M}) = \text{False}$ ;
15       $neigh[i][j] \leftarrow (u_r, t_r)$ ;
16    end
17  end

```

Algorithm 3: GPU Edge Feature Caching

Require: edge features $\{x_e, \forall e \in \mathcal{E}\}$, edge caching budget k , cache replacement threshold ϵ

```

1  $\mathcal{Q} \leftarrow \{0\}_{i=1}^{|\mathcal{E}|}$ ;
2 Randomly cache  $k$  edge features to VRAM;
3 for epoch 1 to  $T$  do
4   for each edge read request  $e$  do
5      $x_e$  serve from VRAM cache or RAM;
6      $\mathcal{Q}[e] \leftarrow \mathcal{Q}[e] + 1$ ;
7   end for
8   if  $|\text{cached edges} \cap \mathcal{Q}_{\text{top}k}| < \epsilon$  then
9     update cache with  $\mathcal{Q}_{\text{top}k}$  edge features
10  end if
11 end for

```

same warp access the same neighbor information, which can be cached in the shared memory. Thirdly, the complexity of the binary search is proportional to the neighbor size while the complexity of the bitmap is inversely proportional to the neighbor size, leading to a balanced workload across different blocks. In experiment IV-C, we verify that our GPU neighbor finder achieves order-of-magnitude speedup compared with existing CPU neighbor finders [33].

D. GPU Feature Caching

To address the issue of the dominant CPU-GPU feature slicing overhead in TGNN training, we propose a GPU cache for the features of nodes and edges with high-access frequency. For dynamic graphs, since edge features are usually tremendously larger than node features, here we demonstrate the more commonly used case of edge feature caching. Due

TABLE I
ACCURACY OF TASER AND BASELINES IN MRR (%). ALL RESULTS ARE AN AVERAGE OF 5 RUNS. (FIRST, SECOND.)

	Wikipedia		Reddit		Flights		MovieLens		GDEL T	
	TGAT	GraphMixer	TGAT	GraphMixer	TGAT	GraphMixer	TGAT	GraphMixer	TGAT	GraphMixer
Baseline	68.76±0.40	74.05±0.21	81.05±0.04	75.11±0.07	80.50±0.08	77.86±0.08	63.11±0.04	69.01±0.06	79.01±0.12	76.26±0.47
w./ Ada. Mini-Batch	72.22±0.41	75.34±0.19	82.57±0.08	76.25±0.07	82.65±0.06	78.89±0.09	63.97±0.08	69.11±0.04	80.34±0.04	76.74±1.09
w./ Ada. Neighbor	73.96±0.51	74.70±1.24	81.66±0.04	75.63±0.17	81.46±0.31	78.94±0.93	65.51±0.42	69.26±0.18	80.22±0.06	76.49±0.15
TASER	75.98±0.35	76.48±0.97	82.59±0.16	76.85±0.56	82.64±0.25	79.39±0.64	65.79±0.13	69.47±0.10	81.04±0.11	76.99±0.72
(Improvement)	(+7.22)	(+2.43)	(+1.54)	(+1.74)	(+2.14)	(+1.53)	(+2.68)	(+0.46)	(+2.03)	(+0.73)

to the temporal adaptive mini-batch selection and neighbor sampling, the access pattern of TASER changes during the training process, which requires a dynamic cache. One naive approach is to maintain an $\mathcal{O}(|\mathcal{E}| \times |\mathcal{E}|)$ matrix to store the access frequency of every supporting neighbor for every training sample. However, this results in unacceptable storage overhead, and the cache update time may even exceed the training time. Although increasing the cache line size can quadratically reduce the memory overhead, the cache hit rate also drops drastically due to the more coarse-grained policy. Empirically, we observe that increasing the cache line size from 1 to 512 leads to more than 20% drop in cache hit rate. On the other hand, since TASER uses the Adam optimizer, the dynamic edge access pattern will eventually stabilize. Therefore, we leverage the historical edge access pattern to update the cache policy. After each epoch, if the overlap between the cached edges and the k most frequently accessed edges $\mathcal{Q}_{\text{top}k}$ of the previous epoch is less than a predefined threshold ϵ , we swap the cached content with features of $\mathcal{Q}_{\text{top}k}$. Note that this lightweight cache replacement policy only requires $\mathcal{O}(|\mathcal{E}|)$ computation, significantly less than the probability-based policy even with a large line size.

Algorithm 3 shows the GPU edge feature caching during training. For each iteration of mini-batch training, TASER concurrently slices a batch of edge features layer by layer and updates the frequency of accessed edges in parallel. For edge features that are not stored in the VRAM cache, we directly slice the feature through the unified virtual memory with zero-copy access over PCI-e.

IV. EXPERIMENTS

A. Experimental Setup

Datasets. We evaluate the performance of TASER on five dynamic graph datasets, whose statistics are shown in Ta-

TABLE II
DATASET STATISTIC. $|d_v|$ AND $|d_e|$ SHOW THE DIMENSIONS OF NODE AND EDGE FEATURES, RESPECTIVELY.

	$ \mathcal{V} $	$ \mathcal{E} $	$ d_v $	$ d_e $	train/val/test
Wikipedia	9,227	157,474	-	172	110k/23k/23k
Reddit	10,984	672,447	-	172	470k/101k/101k
Flights	13,169	1,927,145	100	-	600k/200k/200k
MovieLens	371,715	48,990,832	-	266	600k/200k/200k
GDEL T	16,682	191,290,882	413	130	600k/200k/200k

ble II. Among them, *Wikipedia* [40], *Reddit*¹ [40], and *MovieLens* [41] are bipartite graphs without node features. *Flights* [42] is a traffic graph without edge features, and *GDEL T* [33] is a large-scale knowledge graph including both node and edge features. The tasks are to predict user posts (*Wikipedia*, *Reddit*, *MovieLens*), flight schedules (*Flights*), and news (*GDEL T*). To simulate the use cases in real-world applications, for large-scale datasets with more than one million temporal edges, we use the latest one million edges with 60%, 20%, and 20% chronological splits as the training, validation, and test sets, respectively.

TGNN Models. We build TASER on two state-of-the-art TGNN models introduced in Section II-B. TGAT [3] uses a 2-layer attention-based temporal aggregator with supporting nodes uniformly sampled from the historical neighbors. GraphMixer [5] uses a single-layer MLP-Mixer temporal aggregator with the most recent neighbors as supporting nodes. To ensure a fair comparison, we keep the number of supporting neighbors to 10, the default value in both baselines. Note that TASER does not provide any additional input to the TGNN models other than selecting high-quality supporting neighbors of the same size.

Configurations. For the TGNN models, we follow the default parameters used in the TGL framework [33] for a fair comparison. In particular, we use the 0.0001 learning rate, 600 batch size, 200 training epochs, and $n = 10$ supporting neighbors per node for all the datasets and all the models. We set the dimension of all the hidden embeddings and encodings to 100. For methods with adaptive neighbor sampling, we set $m = 25$ as the budget of the neighbor finder for all the datasets, except for the ablation study in Section IV-F. We follow DistTGL [43] to evaluate the performance of transductive temporal link prediction using Mean Reciprocal Rank (MRR) with 49 randomly sampled negative destination nodes. Please refer to our open-sourced code² for more details on the hyper-parameters.

Hardware and Software. We implement TASER using Python 3.11, PyTorch 2.0.1, DGL 1.1, and CUDA 12.2. All the experiments are conducted on a machine with dual 96-Core AMD EPYC 9654 CPUs paired with 1.5TB ECC-DDR5

¹The Reddit dataset used in this paper is obtained exclusively from the work [40], and no data is directly scraped from the Reddit website.

²<https://github.com/facebookresearch/taser-tgmn>

TABLE III

TOTAL RUNTIME BREAKDOWN PER EPOCH (SEC). NF, AS, FS, AND PP DENOTE NEIGHBOR FINDING, ADAPTIVE NEIGHBOR SAMPLING, FEATURE SLICING, AND PROPAGATION, RESPECTIVELY. THE PERCENTAGES (%) REPRESENT THE RUNTIME RATIOS OF A PARTICULAR STEP RELATIVE TO THE TOTAL EPOCH. THE ARROWS (\uparrow) REFER TO THE SAME RUNTIME AS THE ONE IT POINTS TO.

		TGAT					GraphMixer				
		NF (%)	AS	FS (%)	PP	Total (Impr.)	NF (%)	AS	FS (%)	PP	Total (Impr.)
Wikipedia	Baseline	40.27 (70%)	2.55	11.26 (19%)	3.73	57.81 (1.00 \times)	0.75 (23%)	0.46	0.61 (19%)	1.45	3.28 (1.00 \times)
	+GPU NF	0.07 (0%)	\uparrow	\uparrow (64%)	\uparrow	17.61 (3.28 \times)	0.04 (2%)	\uparrow	\uparrow (24%)	\uparrow	2.56 (1.27 \times)
	+10% Cache	\uparrow (1%)		0.99 (13%)		7.35 (7.86 \times)	\uparrow (2%)		0.18 (8%)		2.13 (1.53 \times)
	+20% Cache	\uparrow (1%)		0.71 (10%)		7.07 (8.17 \times)	\uparrow (2%)		0.16 (8%)		2.11 (1.55 \times)
	+30% Cache	\uparrow (1%)		0.54 (8%)		6.90 (8.38 \times)	\uparrow (2%)		0.13 (6%)		2.08 (1.57 \times)
Reddit	Baseline	218.56 (77%)	10.18	41.56 (15%)	12.62	282.93 (1.00 \times)	3.23 (23%)	1.98	2.36 (17%)	6.23	13.79 (1.00 \times)
	+GPU NF	0.37 (1%)	\uparrow	\uparrow (64%)	\uparrow	64.73 (4.37 \times)	0.19 (2%)	\uparrow	\uparrow (22%)	\uparrow	10.75 (1.28 \times)
	+10% Cache	\uparrow (1%)		4.41 (16%)		27.58 (10.25 \times)	\uparrow (2%)		0.81 (9%)		9.19 (1.50 \times)
	+20% Cache	\uparrow (1%)		2.95 (11%)		26.12 (10.82 \times)	\uparrow (2%)		0.71 (8%)		9.09 (1.51 \times)
	+30% Cache	\uparrow (2%)		2.36 (9%)		25.53 (11.08 \times)	\uparrow (2%)		0.60 (7%)		8.98 (1.53 \times)
MovieLens	Baseline	276.28 (69%)	17.62	79.15 (20%)	27.60	400.66 (1.00 \times)	5.57 (13%)	4.80	19.61 (52%)	12.71	42.68 (1.00 \times)
	+GPU NF	0.54 (0%)	\uparrow	\uparrow (63%)	\uparrow	124.92 (3.20 \times)	0.32 (1%)	\uparrow	\uparrow (24%)	\uparrow	37.45 (1.14 \times)
	+10% Cache	\uparrow (1%)		12.05 (21%)		57.81 (6.93 \times)	\uparrow (2%)		0.46 (2.5%)		18.30 (2.33 \times)
	+20% Cache	\uparrow (1%)		9.67 (17%)		55.43 (7.22 \times)	\uparrow (2%)		0.45 (2.5%)		18.29 (2.33 \times)
	+30% Cache	\uparrow (1%)		7.99 (15%)		53.75 (7.45 \times)	\uparrow (2%)		0.46 (2.5%)		18.30 (2.33 \times)
GDELT	Baseline	322.40 (83%)	17.08	17.84 (5%)	29.52	386.84 (1.00 \times)	6.5 (12%)	6.19	15.37 (29%)	25.33	53.33 (1.00 \times)
	+GPU NF	0.56 (1%)	\uparrow	\uparrow (27%)	\uparrow	65.00 (5.95 \times)	0.36 (1%)	\uparrow	\uparrow (32%)	\uparrow	47.24 (1.12 \times)
	+10% Cache	\uparrow (1%)		3.08 (6%)		50.25 (7.69 \times)	\uparrow (1%)		0.52 (2%)		32.40 (1.64 \times)
	+20% Cache	\uparrow (1%)		2.17 (4%)		49.34 (7.83 \times)	\uparrow (1%)		0.54 (2%)		32.42 (1.64 \times)
	+30% Cache	\uparrow (1%)		2.39 (5%)		49.56 (7.80 \times)	\uparrow (1%)		0.54 (2%)		32.41 (1.64 \times)

RAM and a single NVIDIA RTX 6000 Ada GPU with 48GB ECC-GDDR6 VRAM.

B. Accuracy

Table I shows the accuracy of TASER on the five datasets. We create two variants to better evaluate the effectiveness of each of these two adaptive sampling methods in TASER, where *w./ Ada. Mini-Batch* denotes baseline methods with adaptive mini-batch selection and *w./ Ada. Neighbor* is the one with adaptive neighbor sampling. With both adaptive mini-batch selection and neighbor sampling, TASER achieves an average of 2.3% MRR improvements over the baselines. TGAT gets an average of 3.1% improvements with TASER, while GraphMixer only gets 1.4% improvements. Intuitively, this is because TGAT takes 2-hop neighbors as the input, which benefits more from the adaptive neighbor sampler compared to the 1-hop neighbors of GraphMixer. On the one hand, each variant of TASER consistently outperforms the baseline TGNNs by a large margin, revealing the effectiveness of TASER both in denoising training samples and supporting neighbors. On the other hand, our results suggest that these two orthogonal adaptive sampling techniques can be employed collectively to either enhance or, at least, maintain accuracy.

We notice that the same neighbor decoder, when paired with different temporal aggregators, leads to remarkably different performances. This justifies the need for a general encoder-decoder scheme in TASER. In addition, increasing the integrity of the whole model by using a neighbor decoder with a similar architecture as the temporal aggregator can reduce training difficulties and thus improve accuracy. We note substantial accuracy gains (up to 6%) when training MLP-Mixer with GraphMixer, yet observe minimal improvements with TGAT,

whereas TGAT exhibits a preference for the GATv2 neighbor decoder. For the neighbor encoder, our proposed frequency encoding and identity encoding consistently work well with any neighbor decoders, reducing the variance of test accuracy and improving the MRR by 0.6% \sim 1.8%.

C. Runtime

In this section, we evaluate the speedup of our proposed optimizations in TASER. The training time of TASER can be broken down into the four dominant steps: neighbor finding, adaptive neighbor sampling, feature slicing, and forward and backward propagation. We build TASER using the optimized temporal aggregators as proposed in TGL [33]. For the baseline, we slice all the features from RAM to GPU in each training iteration and use the original neighbor finder implementation in TGAT [3] and GraphMixer [5]. Note that although TGL provides a high-performance parallel CPU neighbor finder, it maintains a pointer array for efficient temporal neighborhood searching that only supports training in chronological order, which does not work in TASER since our mini-batch selection is randomly sampled from a dynamic distribution.

As shown in Table III, the bottlenecks of the baseline are neighbor finding and feature slicing. After applying our GPU neighbor finder and GPU feature caching with 20% of total edge features, the ratio of mini-batch generation time (i.e., neighbor finding time plus feature slicing time) to the total runtime drops significantly from 40% \sim 92% to 3% \sim 18%. The rest of the runtime mainly lies in the neural network computation, which is proportional to the computational complexity. Since the *Flights* dataset does not contain edge features and the node features can be entirely stored on




Fig. 3. (a) Total sampling time per epoch of a 2-layer TGAT with different neighbor finders and different numbers of neighbors per layer. (b) Cache Hit Rate of TASER caching strategy and Oracle caching strategy with different training epochs.

GPU, we do not demonstrate its runtime. TASER achieves an average of $8.68\times$ speedup on TGAT and $1.77\times$ speedup on GraphMixer. TGAT is a 2-layer TGNN and requires a squared number of supporting neighbors, suffering a greater impact from the inefficiency of neighbor finding and feature slicing. On *GDELT*, since we use the latest one million temporal edges for training and evaluation, caching 20% of edge features is already sufficient for the training set.

D. GPU Neighbor Finder

Fig. 3(a) compares the runtime of different uniform neighbor finders, including the original Python-implemented neighbor finder [3], the high-performance CPU parallel neighbor finder from TGL [33], and our TASER GPU neighbor finder. Since the TGL neighbor finder only supports chronological training order, we use chronological order on all three neighbor finders for a fair comparison. To better reflect the actual runtime of CPU neighbor finders, we also include the CPU-GPU data loading time for the sampled neighbor indices. Note that the TGL neighbor finder is built on a pointer array that leverages the chronological training order for fast memory access. Although we do not specifically optimize for the chronological training order, our GPU neighbor sampler is still orders of magnitude faster than the TGL neighbor finder. As shown in Fig. 3(a), when the number of neighbors per layer is set to 25, our TASER neighbor finder achieves a speedup of more than three orders of magnitude compared to the original neighbor finder, and a $37 \sim 56\times$ speedup compared to the TGL neighbor finder, across all five datasets.

E. GPU Cache

We compare our GPU caching strategy with the Oracle caching strategy, which assumes the access frequency of each

edge is known in advance. Both our caching strategy and the Oracle caching strategy are updated at the end of each epoch. Fig. 3 (b) shows that our GPU caching strategy achieves a near-optimal cache hit rate, close to the Oracle cache with the same size. We choose the 10%, 20%, and 30% cache ratio as they fit mainstream GPUs with 8GB, 16GB, and 40GB VRAM on the *GDELT* dataset, respectively. The cache hit rates increase proportionally with the cache ratio until the Oracle cache is able to include all the accessed features. We observe that, as the entire model’s weights progressively stabilize, our GPU cache rarely necessitates an update after 20 epochs, further reaffirming the low maintenance of our strategy. Note that the cache hit rate of the Oracle caching strategy can also reflect the explore-and-exploit strategy of the adaptive samplers. For instance, on the *Wikipedia* dataset, the cache hit rate of the 10% Oracle cache first increased to 73% and then gradually decreased to 69%, illustrating that

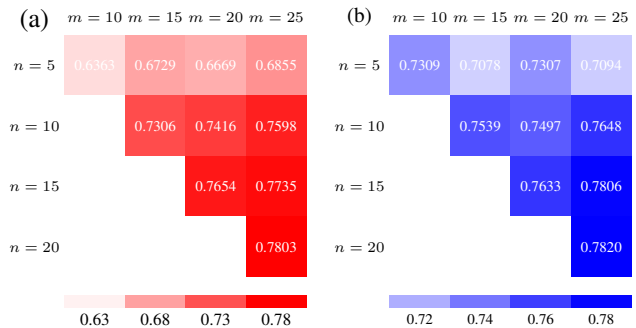


Fig. 4. Test MRR of (a) TGAT and (b) GraphMixer with TASER on the *Wikipedia* dataset. m and n denote the numbers of neighbors selected by the neighbor finder and the adaptive neighbor sampler, respectively.

the adaptive samplers initially exploit high-reward edges and subsequently explore other training samples and supporting neighbors to improve the accuracy.

F. Ablation Study

We evaluate the performance of TASER with different neighbor budgets. Fig. 4 demonstrates that TASER is versatile to various numbers of neighbor candidates m and sampled supporting neighbors n . We note that the accuracy does not improve when increasing m for GraphMixer with $n = 5$. Since GraphMixer is a one-layer TGNN model, it has only 5 supporting nodes per root node when $n = 5$, while a 2-layer TGAT has $5 + 5 \times 5 = 30$ supporting nodes. Selecting $n = 5$ as the hyper-parameter choice for GraphMixer is suboptimal for real-world applications, leading to inaccurate supervision for the adaptive sampler from the TGNN model. The results validate our hypothesis that, with a larger number of neighbor candidates m , the adaptive neighbor sampler is capable of selecting supporting neighbors that provide more pivotal information for task prediction. It also shows that TASER consistently performs well when TGNNs prefer a larger number of supporting neighbors n .

V. CONCLUSION

In this work, we proposed TASER, a novel temporal neighbor sampling method for fast and accurate representation learning on dynamic graphs. With a two-fold adaptive sampling method, temporal adaptive mini-batch selection, and temporal adaptive neighbor sampling, TASER endowed TGNNs with the ability to handle distinctive noise in dynamic graphs. However, the introduced adaptive sampling increases the number of neighborhood traversals, prolonging the runtime of existing bottlenecks in TGNNs, specifically neighbor finding and CPU-GPU feature slicing. We proposed two system optimizations to address these bottlenecks: an efficient GPU neighbor finder and a GPU feature caching strategy. On two state-of-the-art backbone TGNNs and five real-world datasets, TASER not only improved the accuracy by an average of 2.3% in MRR but also achieved an average speedup of $5.1\times$ on a single GPU.

ACKNOWLEDGMENT

This work is supported by Meta Platforms Inc. under grant number INB2675366, National Science Foundation (NSF) under grant OAC-2209563, and DEVCOM Army Research Lab (ARL) under grant W911NF2220159. Distribution Statement A: Approved for public release. Distribution is unlimited.

REFERENCES

- [1] A. Pareja, G. Domeniconi, J. Chen, T. Ma, T. Suzumura, H. Kanezashi, T. Kaler, T. Schardl, and C. Leiserson, "Evolvegcn: Evolving graph convolutional networks for dynamic graphs," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, no. 04, 2020, pp. 5363–5370.
- [2] S. Yan, Y. Xiong, and D. Lin, "Spatial temporal graph convolutional networks for skeleton-based action recognition," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 32, no. 1, 2018.
- [3] D. Xu, C. Ruan, E. Körpeoglu, S. Kumar, and K. Achan, "Inductive representation learning on temporal graphs," in *ICLR*, 2020.
- [4] E. Rossi, B. Chamberlain, F. Frasca, D. Eynard, F. Monti, and M. Bronstein, "Temporal graph networks for deep learning on dynamic graphs," in *ICML 2020 Workshop on Graph Representation Learning*, 2020.
- [5] W. Cong, S. Zhang, J. Kang, B. Yuan, H. Wu, X. Zhou, H. Tong, and M. Mahdavi, "Do we really need complicated model architectures for temporal networks?" in *ICLR*, 2023.
- [6] Y. Luo and P. Li, "Neighborhood-aware scalable temporal network representation learning," in *Learning on Graphs Conference*. PMLR, 2022, pp. 1–1.
- [7] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *ICLR (Poster)*, 2017.
- [8] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," *Advances in neural information processing systems*, vol. 30, 2017.
- [9] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," in *ICLR (Poster)*, 2018.
- [10] M. Zhang, S. Wu, X. Yu, Q. Liu, and L. Wang, "Dynamic graph neural networks for sequential recommendation," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 5, pp. 4741–4753, 2022.
- [11] H. Zhou, J. Orme-Rogers, R. Kannan, and V. Prasanna, "Sedyt: A general framework for multi-step event forecasting via sequence modeling on dynamic entity embeddings," in *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 2021, pp. 3667–3671.
- [12] D. Cheng, X. Wang, Y. Zhang, and L. Zhang, "Graph neural network for fraud detection via spatial-temporal attention," *IEEE Transactions on Knowledge Data Engineering*, vol. 34, no. 08, pp. 3800–3813, aug 2022.
- [13] H. Zeng, H. Zhou, A. Srivastava, R. Kannan, and V. K. Prasanna, "Graphsaint: Graph sampling based inductive learning method," in *ICLR*, 2020.
- [14] J. Chen, T. Ma, and C. Xiao, "Fastgcn: Fast learning with graph convolutional networks via importance sampling," in *ICLR (Poster)*, 2018.
- [15] D. Zou, Z. Hu, Y. Wang, S. Jiang, Y. Sun, and Q. Gu, "Layer-dependent importance sampling for training deep and large graph convolutional networks," *Advances in neural information processing systems*, vol. 32, 2019.
- [16] W. Huang, T. Zhang, Y. Rong, and J. Huang, "Adaptive sampling towards fast graph representation learning," *Advances in neural information processing systems*, vol. 31, 2018.
- [17] Z. Liu, Z. Wu, Z. Zhang, J. Zhou, S. Yang, L. Song, and Y. Qi, "Bandit samplers for training graph neural networks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 6878–6888, 2020.
- [18] M. Yoon, T. Gervet, B. Shi, S. Niu, Q. He, and J. Yang, "Performance-adaptive sampling strategy towards fast and accurate graph neural networks," in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, pp. 2046–2056.
- [19] Q. Zhang, D. Wipf, Q. Gan, and L. Song, "A biased graph neural network sampler with near-optimal regret," *Advances in Neural Information Processing Systems*, vol. 34, pp. 8833–8844, 2021.
- [20] S. Tian, R. Wu, L. Shi, L. Zhu, and T. Xiong, "Self-supervised representation learning on dynamic graphs," in *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, ser. CIKM '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 1814–1823. [Online]. Available: <https://doi.org/10.1145/3459637.3482389>
- [21] J. Li, S. Tian, R. Wu, L. Zhu, W. Zhao, C. Meng, L. Chen, Z. Zheng, and H. Yin, "Less can be more: Unsupervised graph pruning for large-scale dynamic graphs," *arXiv preprint arXiv:2305.10673*, 2023.
- [22] I. O. Tolstikhin, N. Houlsby, A. Kolesnikov, L. Beyer, X. Zhai, T. Unterthiner, J. Yung, A. Steiner, D. Keysers, J. Uszkoreit *et al.*, "Mlp-mixer: An all-mlp architecture for vision," *Advances in neural information processing systems*, vol. 34, pp. 24 261–24 272, 2021.
- [23] Y. Wang, Y. Cai, Y. Liang, H. Ding, C. Wang, and B. Hooi, "Time-aware neighbor sampling for temporal graph networks," 2021.
- [24] H. Chen, P. Jiao, H. Tang, and H. Wu, "Temporal graph representation learning with adaptive augmentation contrastive," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2023, pp. 683–699.
- [25] R. Zhu, "Gradient-based sampling: An adaptive importance sampling for least-squares," *Advances in neural information processing systems*, vol. 29, 2016.

- [26] S. U. Stich, A. Raj, and M. Jaggi, “Safe adaptive importance sampling,” *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [27] W. Cong, R. Forsati, M. Kandemir, and M. Mahdavi, “Minimal variance sampling with provable guarantees for fast training of graph neural networks,” in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, ser. KDD ’20. New York, NY, USA: Association for Computing Machinery, 2020, p. 1393–1403. [Online]. Available: <https://doi.org/10.1145/3394486.3403192>
- [28] M. Wang, D. Zheng, Z. Ye, Q. Gan, M. Li, X. Song, J. Zhou, C. Ma, L. Yu, Y. Gai, T. Xiao, T. He, G. Karypis, J. Li, and Z. Zhang, “Deep graph library: A graph-centric, highly-performant package for graph neural networks,” *arXiv preprint arXiv:1909.01315*, 2019.
- [29] Z. Tan, X. Yuan, C. He, M.-K. Sit, G. Li, X. Liu, B. Ai, K. Zeng, P. Pietzuch, and L. Mai, “Quiver: Supporting gpus for low-latency, high-throughput gnn serving with workload awareness,” 2023.
- [30] S. Pandey, L. Li, A. Hoisie, X. S. Li, and H. Liu, “C-saw: A framework for graph sampling and random walk on gpus,” in *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2020, pp. 1–15.
- [31] A. Jangda, S. Polisetty, A. Guha, and M. Serafini, “Accelerating graph sampling for graph machine learning using gpus,” in *Proceedings of the Sixteenth European Conference on Computer Systems*, ser. EuroSys ’21. New York, NY, USA: Association for Computing Machinery, 2021, p. 311–326. [Online]. Available: <https://doi.org/10.1145/3447786.3456244>
- [32] P. Wang, C. Li, J. Wang, T. Wang, L. Zhang, J. Leng, Q. Chen, and M. Guo, “Skywalker: Efficient alias-method-based graph sampling and random walk on gpus,” in *2021 30th International Conference on Parallel Architectures and Compilation Techniques (PACT)*, 2021, pp. 304–317.
- [33] H. Zhou, D. Zheng, I. Nisa, V. Ioannidis, X. Song, and G. Karypis, “Tgl: A general framework for temporal gnn training on billion-scale graphs,” *Proc. VLDB Endow.*, vol. 15, no. 8, p. 1572–1580, apr 2022. [Online]. Available: <https://doi.org/10.14778/3529337.3529342>
- [34] C. Huan, S. L. Song, S. Pandey, H. Liu, Y. Liu, B. Lepers, C. He, K. Chen, J. Jiang, and Y. Wu, “Tea: A general-purpose temporal graph random walk engine,” in *Proceedings of the Eighteenth European Conference on Computer Systems*, ser. EuroSys ’23. New York, NY, USA: Association for Computing Machinery, 2023, p. 182–198. [Online]. Available: <https://doi.org/10.1145/3552326.3567491>
- [35] J. Dong, D. Zheng, L. F. Yang, and G. Karypis, “Global neighbor sampling for mixed cpu-gpu training on giant graphs,” in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, pp. 289–299.
- [36] S. W. Min, K. Wu, M. Hidayetoglu, J. Xiong, X. Song, and W.-m. Hwu, “Graph neural network training and data tiering,” in *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022, pp. 3555–3565.
- [37] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” 2017, pp. 5998–6008.
- [38] S. Brody, U. Alon, and E. Yahav, “How attentive are graph attention networks?” in *ICLR*, 2022.
- [39] R. J. Williams, “Simple statistical gradient-following algorithms for connectionist reinforcement learning,” *Machine learning*, vol. 8, pp. 229–256, 1992.
- [40] S. Kumar, X. Zhang, and J. Leskovec, “Predicting dynamic embedding trajectory in temporal interaction networks,” in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 2019, pp. 1269–1278.
- [41] J. Vig, S. Sen, and J. Riedl, “The tag genome: Encoding community knowledge to support novel interaction,” *ACM Trans. Interact. Intell. Syst.*, vol. 2, no. 3, sep 2012. [Online]. Available: <https://doi.org/10.1145/2362394.2362395>
- [42] F. Poursafaei, S. Huang, K. Pelrine, and R. Rabbany, “Towards better evaluation for dynamic link prediction,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 32928–32941, 2022.
- [43] H. Zhou, D. Zheng, X. Song, G. Karypis, and V. Prasanna, “Disttgl: Distributed memory-based temporal graph neural network training,” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2023.