



Accelerating Sparse Graph Neural Networks via Dense Tensor Cores on GPUs

Yuke Wang

UC at Santa Barbara


Graphs are everywhere, GNNs are useful hammer...




Social Networks



Financial Services



Molecular chemistry



Point Cloud



Power Grid




Molecular Biology

Background

- Graph Neural Network Basics.

$$a_v^{(k+1)} = \text{Aggregate}^{(k+1)}(h_u^{(k)} | u \in N(v) \cup h_v^{(k)})$$

$$h_v^{(k+1)} = \text{Update}^{(k+1)}(a_v^{(k+1)})$$



- Basic computation in GNNs.

- Neighbor aggregation (SpMM-like).

$$\hat{\mathbf{X}} = (\mathbf{F}_{N \times N} \odot \mathbf{A}_{N \times N}) \cdot \mathbf{X}_{N \times D}$$

- Edge feature computation (SDDMM-like).

$$\mathbf{F} = (\mathbf{X}_{N \times D} \cdot \mathbf{X}_{N \times D}^T) \odot \mathbf{A}_{N \times N}$$

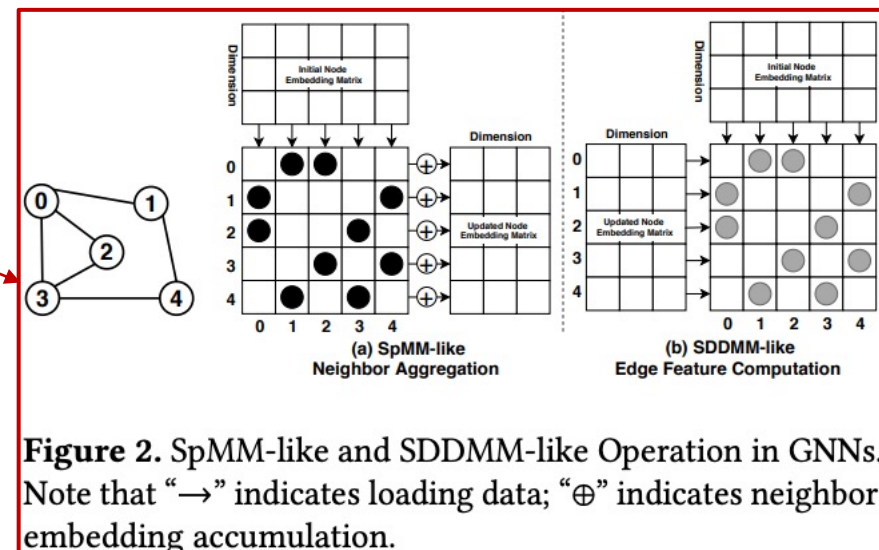


Figure 2. SpMM-like and SDDMM-like Operation in GNNs. Note that “→” indicates loading data; “⊕” indicates neighbor embedding accumulation.

Challenges

- Existing deep-learning frameworks are optimized for dense neural network operations.
- Existing major sparse computation kernels (e.g., cuSPARSE) leverage CUDA cores.
- Existing Tensor-Core based kernels (e.g., Block-SpMM) rely on rigid input sparsity pattern (e.g., block sparsity).

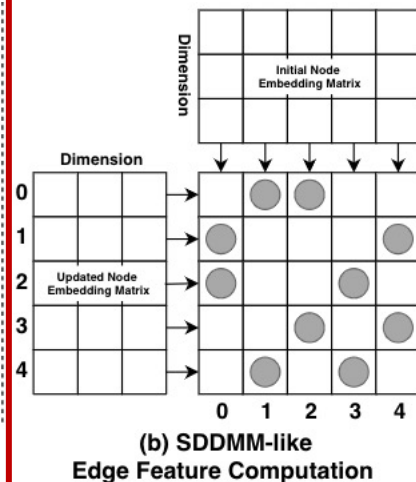



Lack of efficient support for sparse graph neural network computation.

Underutilize the latest GPU with new hardware feature that can offer high-performance computation.

Limits its applicability towards different sparse inputs settings.

Motivations



SpMM-like and SDDMM-like Operation in GNNs.
Note that “ \rightarrow ” indicates loading data; “ \oplus ” indicates neighbor embedding accumulation.

Apply separate optimization on one direction only would hardly work



GPUs fit GNNs

Sparse MM on CUDA core

Profiling of GCN Sparse Operations.

Dataset	Aggr. (%)	Update (%)	Cache(%)	Occ.(%)
Cora	88.56	11.44	37.22	15.06
Citeseer	86.52	13.47	38.18	15.19
Pubmed	94.39	5.55	37.22	16.24

GNNs fit GPUs

Dense MM on CUDA core

Medium-size Graphs in GNNs.

Dataset	# Nodes	# Edges	Memory	Eff.Comp
OVCR-8H	1,890,931	3,946,402	14302.48 GB	0.36%
Yeast	1,714,644	3,636,546	11760.02 GB	0.32%
DD	334,925	1,686,092	448.70 GB	0.03%

Question:

How could we match the sparse GNN workload with GPUs to achieve high computation efficiency and better utilization of GPU resources?



TC-GNN

- The first TC-based GNN acceleration design on GPUs.
- At the input level technique.
- At the kernel level innovation.
- At the framework level design.

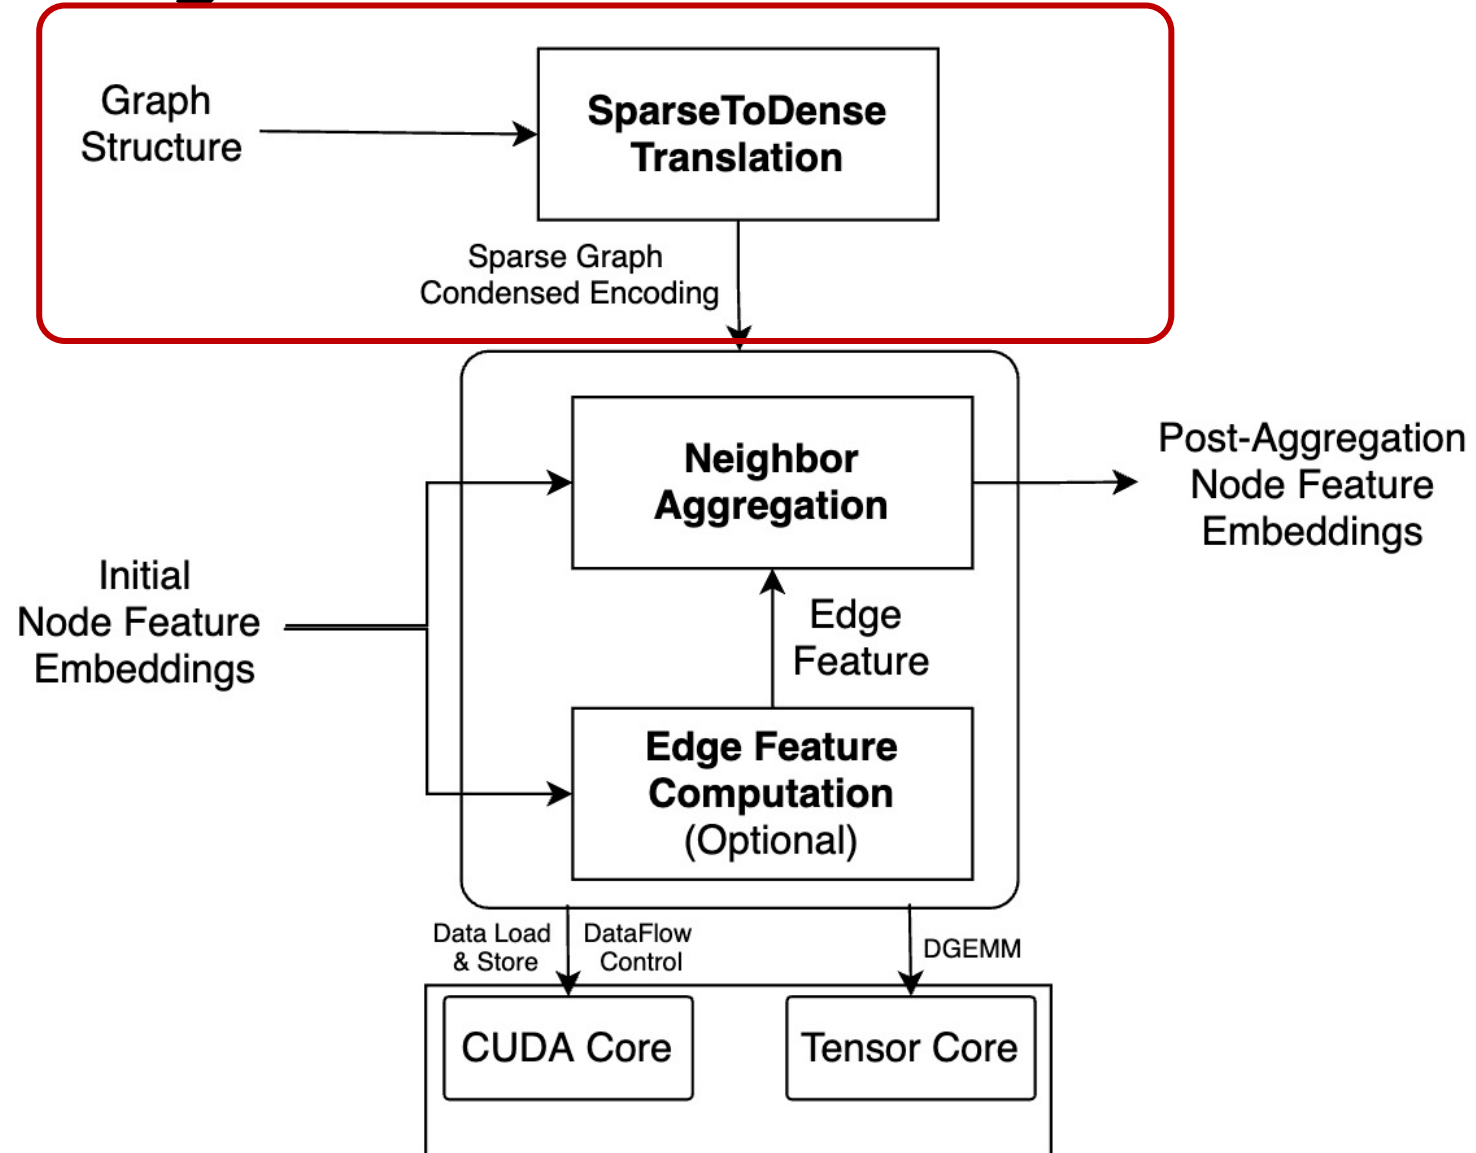
“Let the input sparse graph fit the dense computation of Tensor Core”

Sparse graph translation (SGT) technique condense non-zero elements from sparse tiles into a fewer number of “dense” tiles

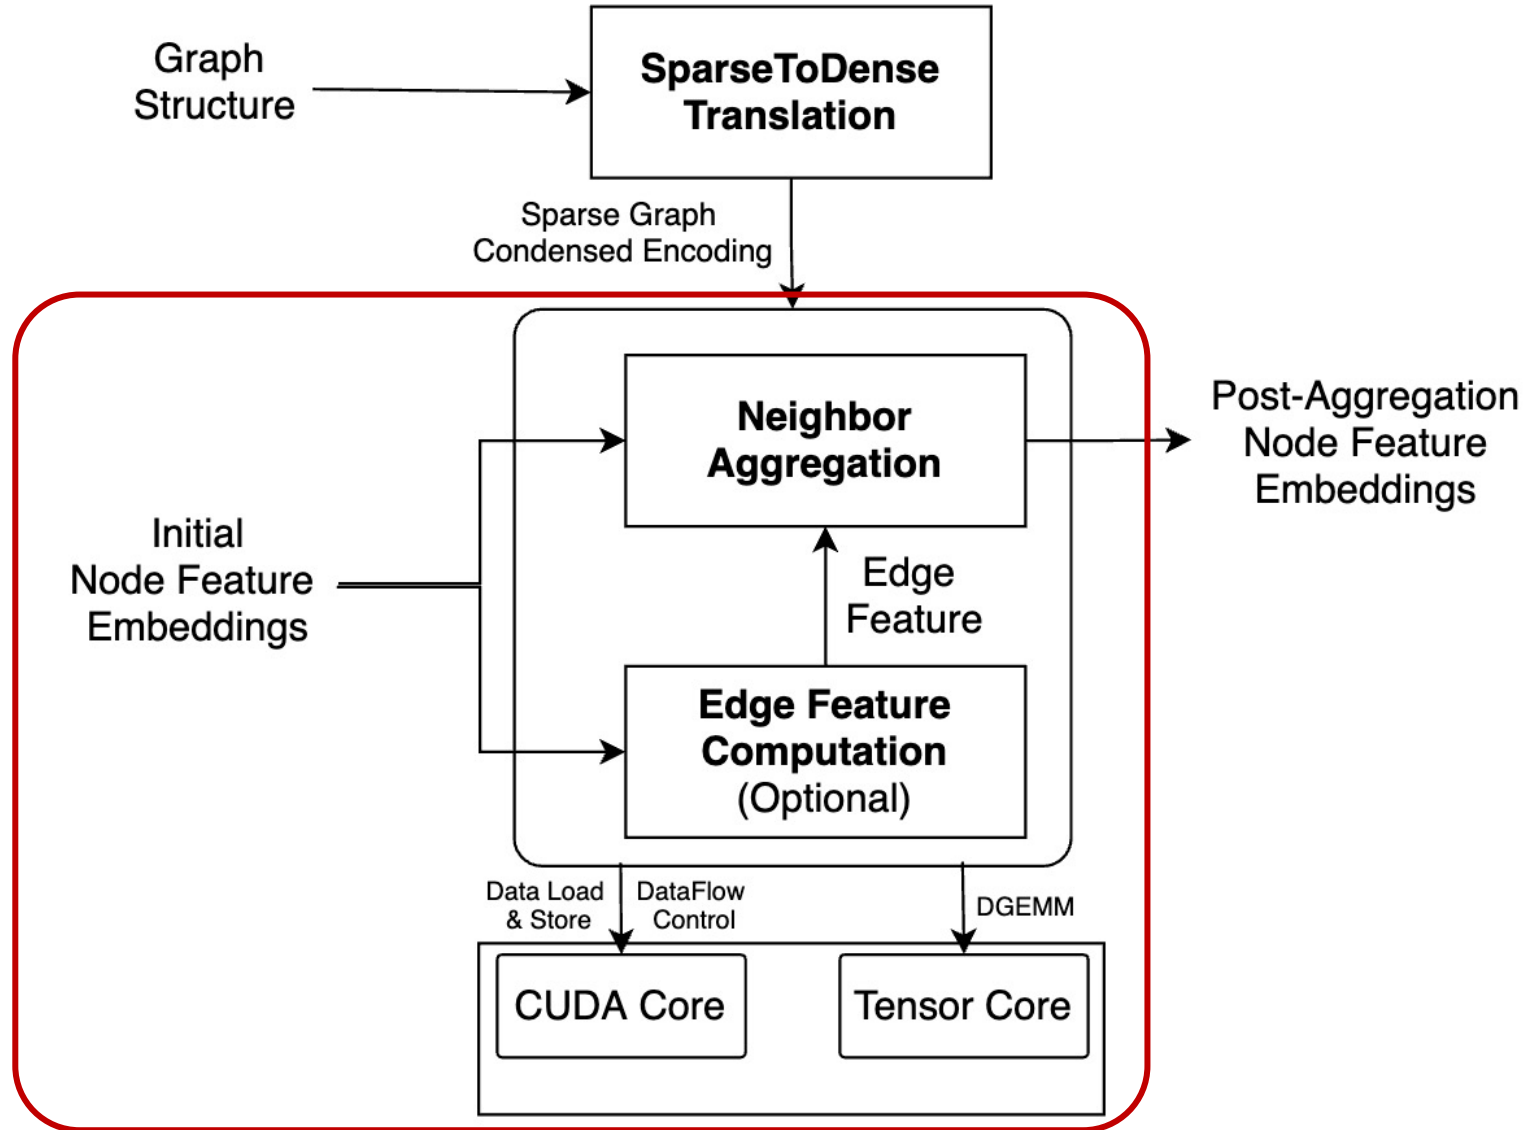
TC-GNN exploits the benefits of CUDA core and tensor core collaboration.

TC-GNN integrates with the popular Pytorch framework.

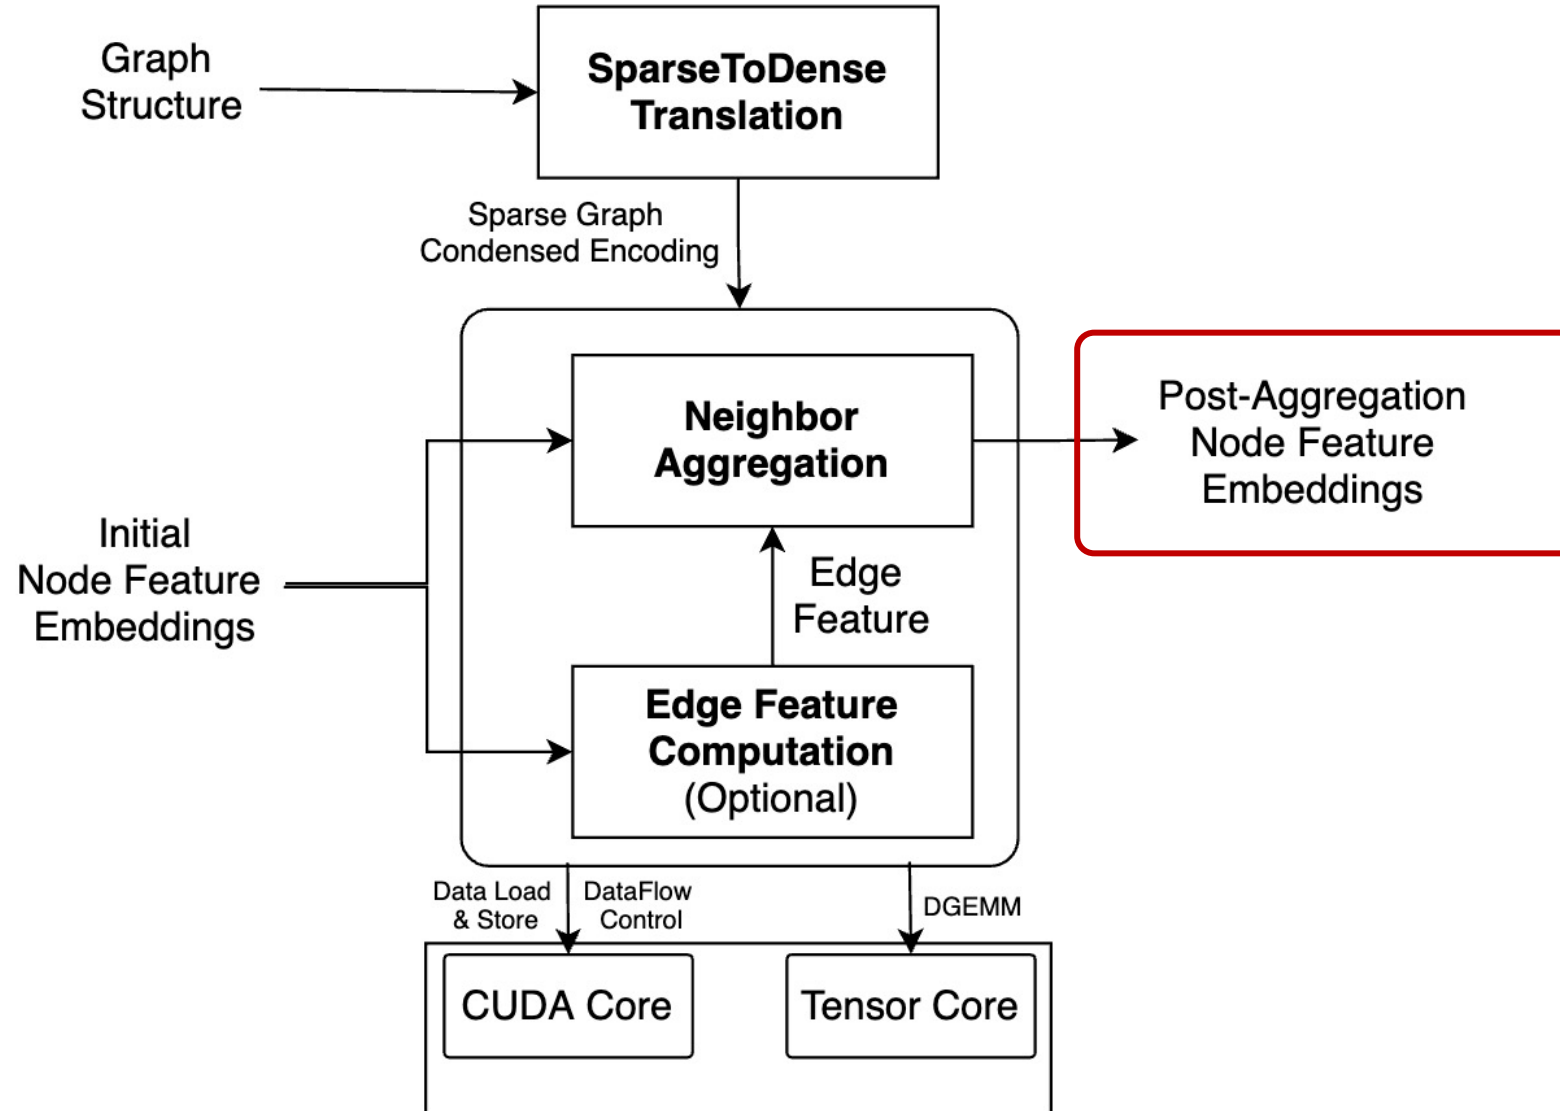
Overall Design



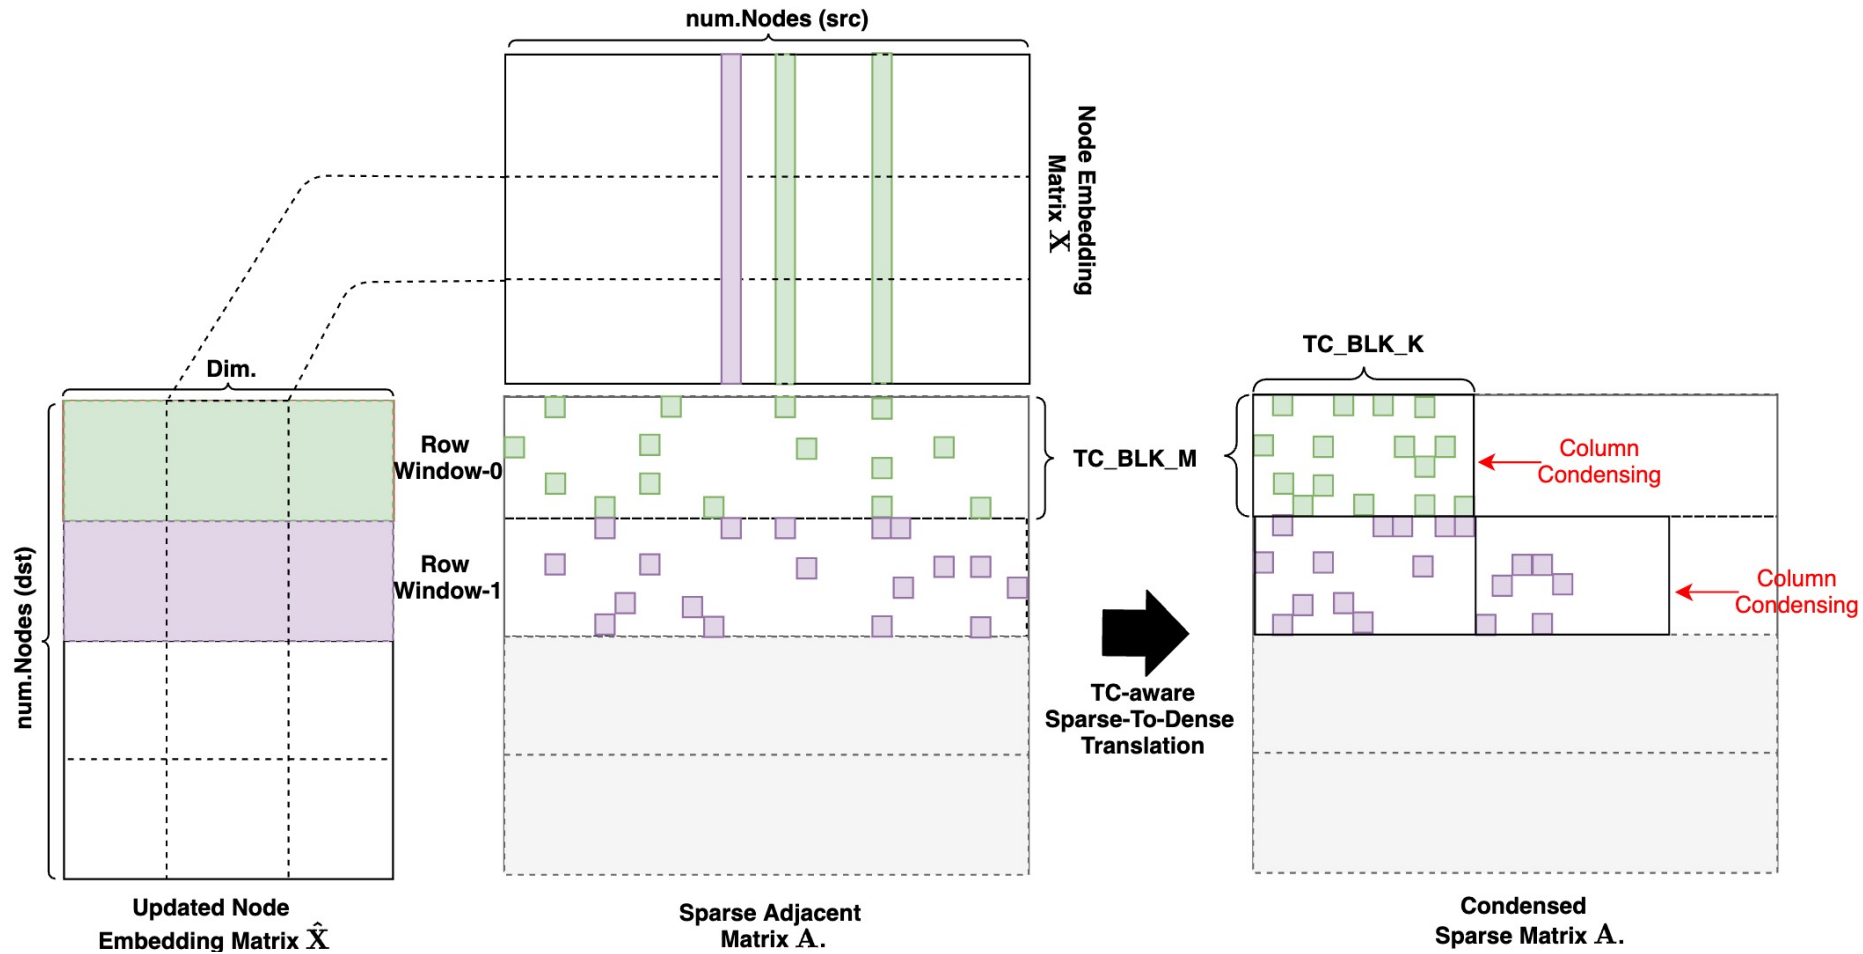
Overall Design



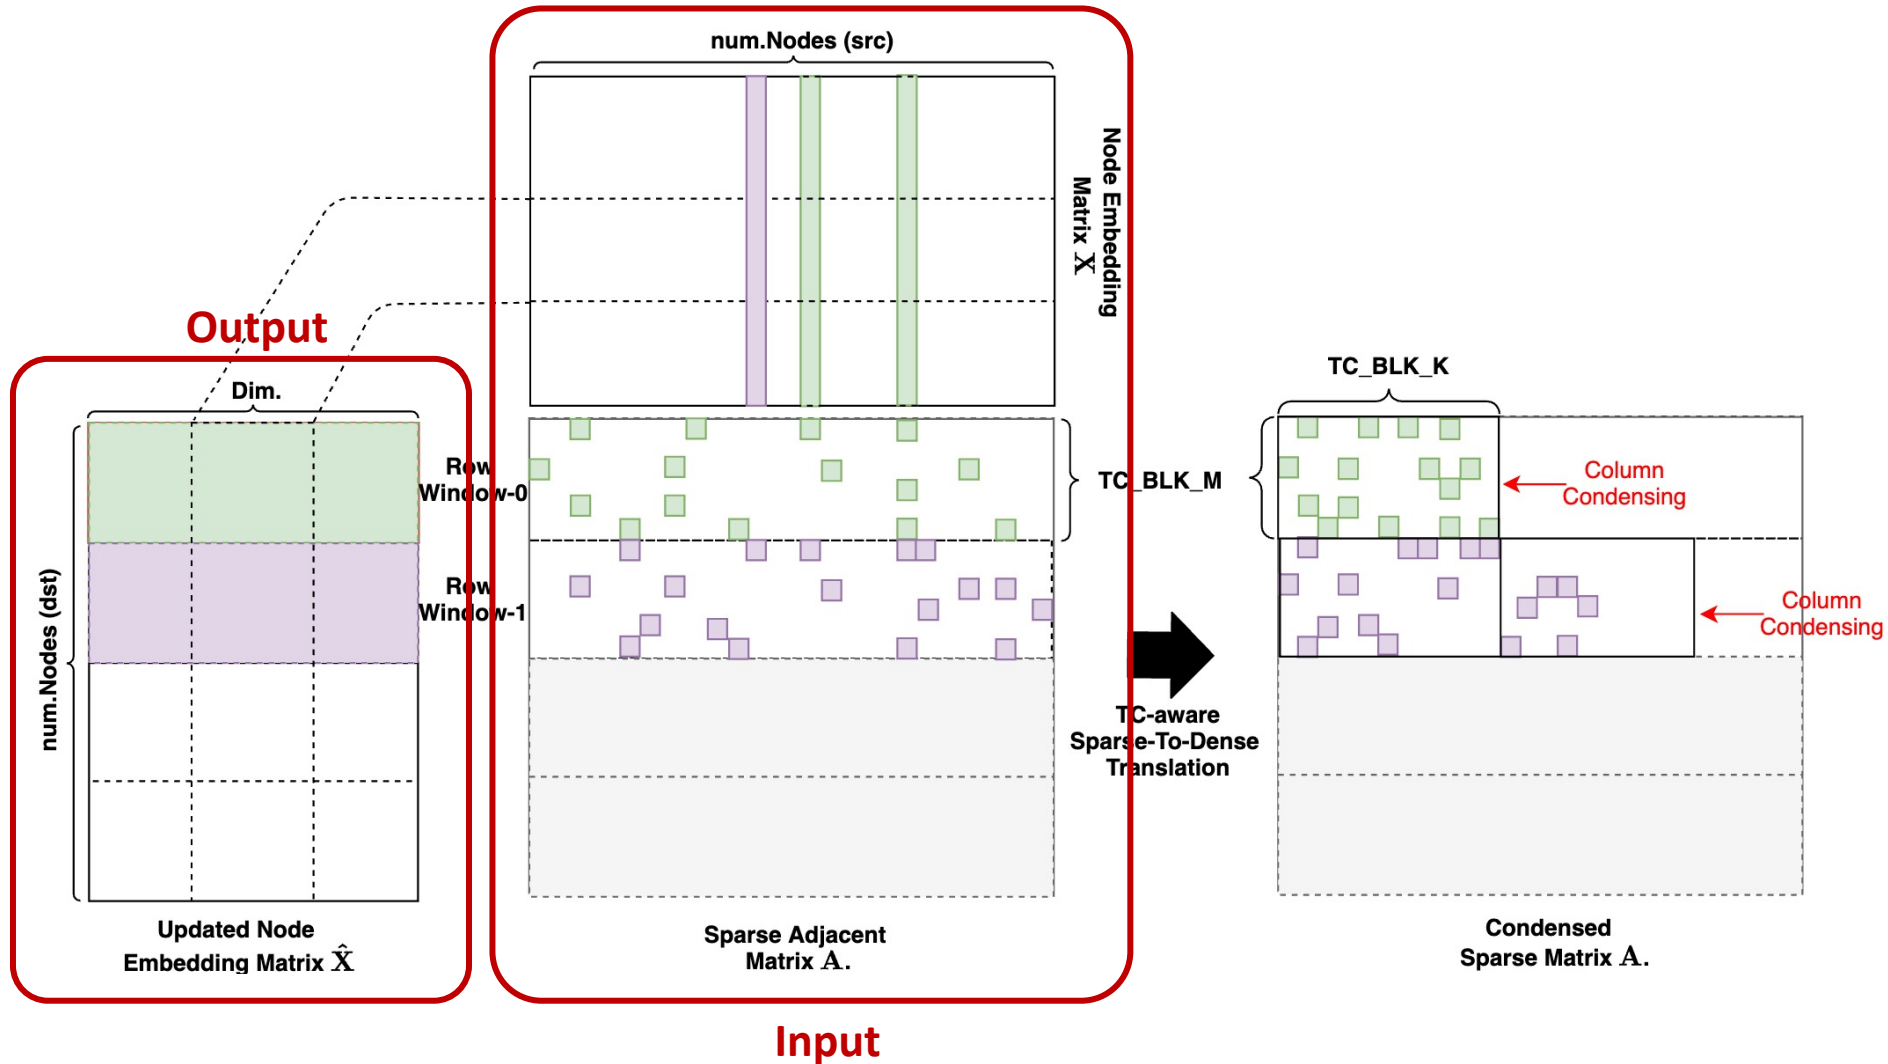
Overall Design



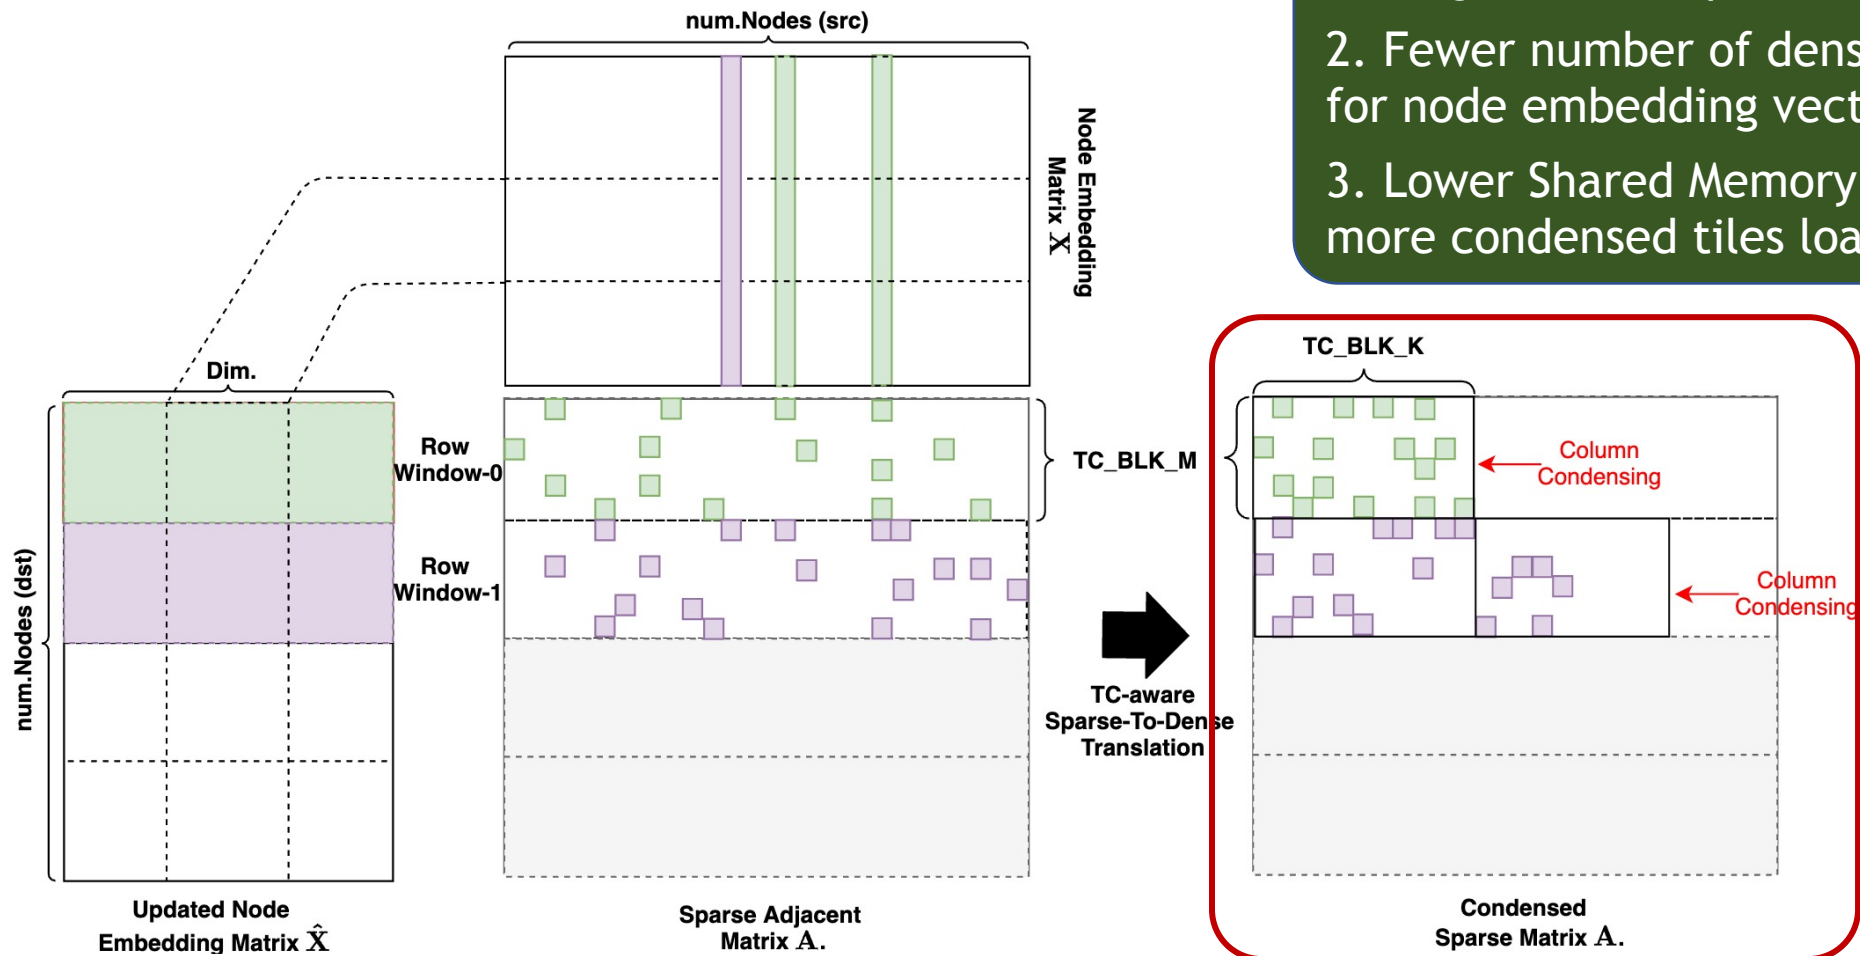
Sparse Graph Translation



Sparse Graph Translation



Sparse Graph Translation



1. Fewer number of iterations for Calling TC WMMA primitives.
2. Fewer number of dense row access for node embedding vector.
3. Lower Shared Memory Usage due to more condensed tiles loading.

TC-aware Sparse Graph Translation

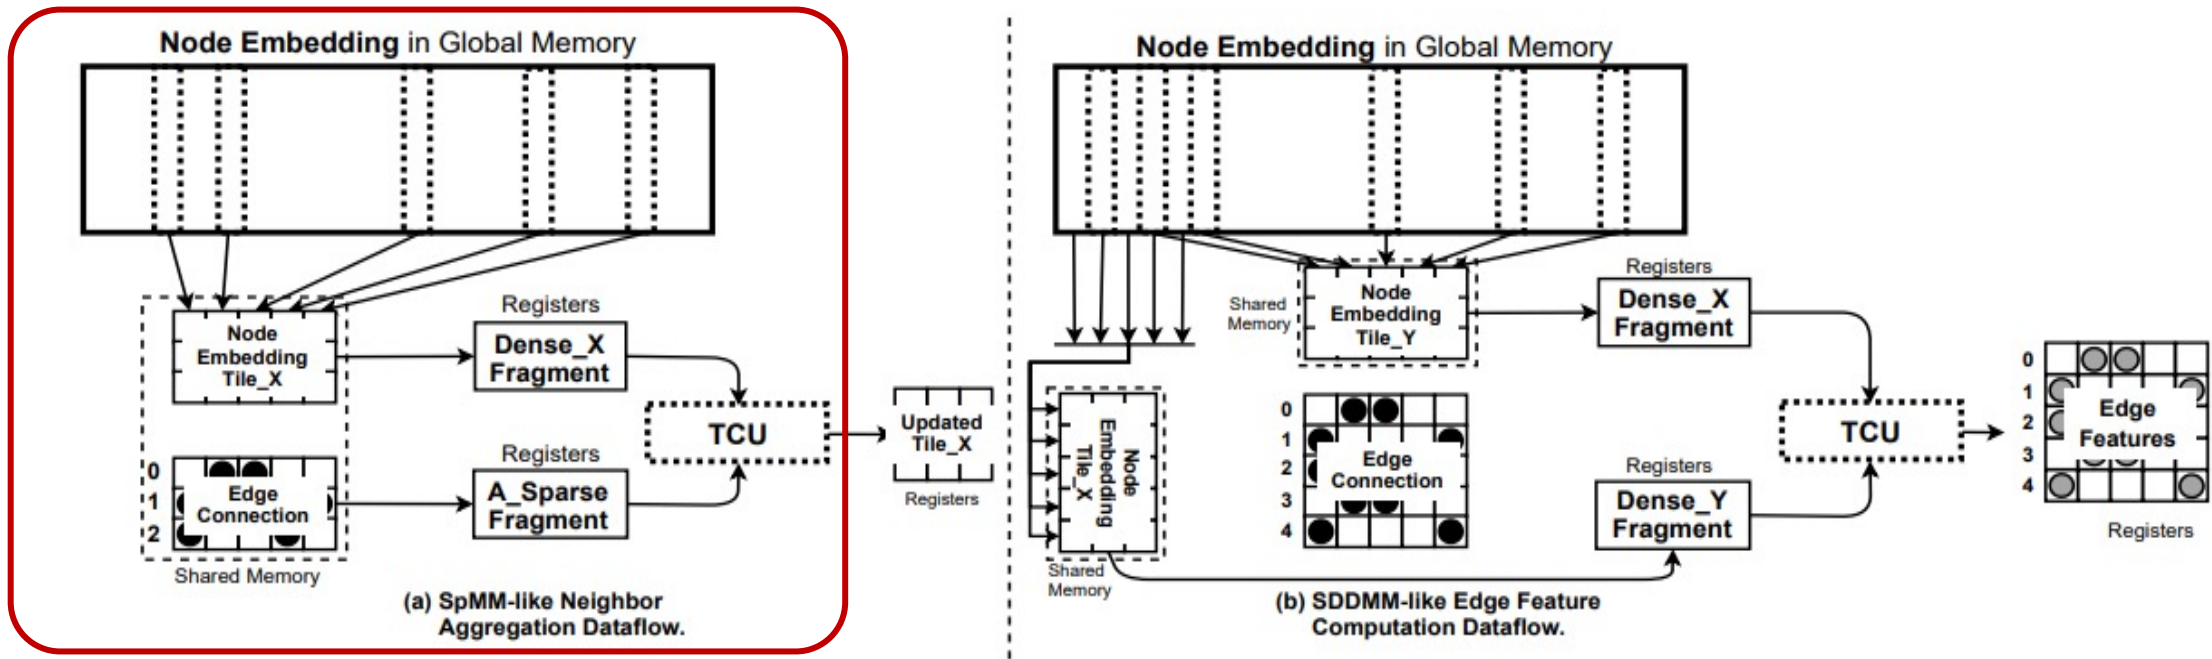
Algorithm 1: TCU-aware Sparse Graph Translation.

input : Graph adjacent matrix A (*nodePointer*, *edgeList*).

output : Result of *winPartition* and *edgeToCol*


```
1  /* Compute the total number of row windows. */
1  numRowsWin = ceil(numNodes/winSize);
2  for winId in numRowsWin do
3      /* EdgeIndex range of the current rowWindow. */
4      winStart = nodePointer[winId * winSize];
5      winEnd = nodePointer[(winId + 1) * winSize];
6      /* Sort the edges of the current rowWindow. */
7      eArray = Sort(winStart, winEnd, edgeList);
8      /* Deduplicate edges of the current rowWindow. */
9      eArrClean = Deduplication(nIdArray);
10     /* #TC blocks in the current rowWindow. */
11     winPartition[winId] = ceil(eArrClean.size/TC_BLK_w);
12     /* Edges-to-columnID mapping in TC Blocks. */
13     for eIndex in [winStart, winEnd] do
14         eid = edgeList[eIndex];
15         edgeToCol[eIndex] = eArrClean[eid];
16     end
17 end
```

TC-optimized Dataflow



TC-Optimized Dataflow Design for (a) Neighbor Aggregation and (b) Edge Feature Computing in GNNs

TC-optimized Dataflow



TC-Optimized Dataflow Design for (a) Neighbor Aggregation and (b) Edge Feature Computing in GNNs

TC-tailored SpMM

Algorithm 2: TC-GNN Neighbor Aggregation.

input : Condensed graph structural information (*nodePointer*, *edgeList*, *edgeToCol*, *winPartition*) and node embedding matrix (*X*).

output: Updated node embedding matrix (\hat{X}).

```

1  /* Traverse through all row windows.          Block */
1  for winId in numRowsWindows do
2      /* #TC blocks of the row window.          */
2      numTCblocks = winPartition[winId];
3      /* Edge range of TC blocks of the row window. */
3      edgeRan = GetEdgeRange(nodePointer, winId);
4      for TCblkId in numTCblocks do
5          /* The edgeList chunk in current TC block. */
5          edgeChunk = GetChunk(edgeList, edgeRan, TCblkId);
6          /* Neighbor node Ids in current TC block. */
6          colToNId = GetNeighbors(edgeChunk, edgeToCol);
7          /* Initiate a dense tile (ATile).          */
7          ATile = InitSparse(edgeChunk, winId);
8          /* Initiate a dense tile (XTile).          */
8          XTile, colId = FetchDense(colToNId, X);
9          /* Compute XnewTile via TCU GEMM.          */
9          XnewTile = TCcompute(ATile, XTile);
10         /* Store XnewTile of  $\hat{X}$ .                  */
10          $\hat{X}$  = StoreDense(XnewTile, winId, colId);
11     end
12 end

```

Warp

TC-tailored SDDMM

Algorithm 3: TC-GNN Edge Feature Computation.

input : Condensed graph structural information (*nodePointer*, *edgeList*, *edgeToCol*, *winPartition*) and node embedding matrix (*X*).

output: Edge Feature List (*edgeValList*).

```

1  /* Traverse through all row windows.          */
1  for winId in numRowsWin do
2      /* #TC blocks in the row window.          */
2      numTCblocks = winPartition[winId];
3      /* Edge range of TC blocks of the row window. */
3      edgeRan = GetEdgeRange(nodePointer, winId);
4      for TCblkId in numTCblocks do
5          /* EdgeList chunk in current TC block. */
5          edgeChunk = GetChunk(edgeList, edgeRan, TCblkId);
6          /* Neighbor node Ids in current TC block. */
6          colToNId = GetNeighbors(edgeChunk, edgeToCol);
7          /* Fetch a dense tile (XTileA).          */
7          XTileA = FetchDenseRow(winId, TCblkId, X);
8          /* Fetch a dense tile (XTileB).          */
8          XTileB = FetchDenseCol(colToNId, edgeToCol, X);
9          /* Compute edgeValTile via TCU GEMM.          */
9          edgeValTile = TCcompute(XTileA, XTileB);
10         /* Store edgeValTile to edgeValList.          */
10         StoreSparse(edgeValList, edgeValTile,
11                     edgeList, edgeToCol);
12     end
13 end

```



Evaluation

- **Baseline:**
 - Deep Graph Library (DGL)
 - PyTorch Geometric (PyG)
- **GNN model:**
 - GCN (Graph Convolutional Network)
 - AGNN (Attention-based GNN)
- **Dataset:** see right-side table.
- **Platform:**
 - A desktop server with 8-core 16-thread Intel Xeon Silver 4110 CPU (64GB host memory) and NVIDIA RTX3090 GPU (24GB device memory)

Table 4. Datasets for Evaluation.

Type	Dataset	#Vertex	#Edge	Dim.	#Class
I	Citeseer	3,327	9,464	3703	6
	Cora	2,708	10,858	1433	7
	Pubmed	19,717	88,676	500	3
	PPI	56,944	818,716	50	121
II	PROTEINS_full	43,471	162,088	29	2
	OVCAR-8H	1,890,931	3,946,402	66	2
	Yeast	1,714,644	3,636,546	74	2
	DD	334,925	1,686,092	89	2
	YeastH	3,139,988	6,487,230	75	2
III	amazon0505	410,236	4,878,875	96	22
	artist	50,515	1,638,396	100	12
	com-amazon	334,863	1,851,744	96	22
	soc-BlogCatalog	88,784	2,093,195	128	39
	amazon0601	403,394	3,387,388	96	22

End-to-end Performance



Speedup over (a) DGL and (b) PyG on GCN and AGNN.

Avg: 1.70X

Operator Performance (dgl.op)

- SpMM (dgl.op.copy_u_sum)

	dgl.op (ms)	TC-GNN (ms)
PROTEINS_full	0.088	0.044
OVCAR-8H	1.295	1.018
Yeast	1.183	0.862
DD	0.454	0.287
SW-620H	1.291	1.018

Avg: 1.50X

- SDDMM (dgl.op.u_dot_v)

	dgl.op (ms)	TC-GNN (ms)
PROTEINS_full	0.062	0.019
OVCAR-8H	0.466	0.054
Yeast	0.401	0.051
DD	0.170	0.026
SW-620H	0.476	0.055

Avg: 6.98X

Kernel Performance (cuSPARSE)

- SpMM w.r.t cuSPARSE with different embedding dimension. (GFLOPS)

	D (16)		D (32)		D (64)	
	cuSPARSE	TC-GNN	cuSPARSE	TC-GNN	cuSPARSE	TC-GNN
PROTEINS_full	90.13	130.89	170.55	226.63	276.46	348.73
OVCAR-8H	135.26	143.54	237.81	239.05	237.96	340.02
Yeast	135.42	157.97	238.12	261.76	230.25	366.61
DD	156.17	207.04	309.67	350.57	467.94	498.02
SW-620H	135.22	143.56	237.72	239.17	238.13	340.04

Avg: 1.23X

Future Works

- GPU-accelerated Preprocessing.

- Current version is based on CPU + OpenMP parallel.
- Intra-warp/block sorting for variable length edge list is needed (may use CUB library for fixed-length array sorting + padding).

- Support/optimization for multiple precision TC.

- Current version is using TF32 on Ampere with WMMA shape of 16x8x16.
- Adaptive optimization for different inputs settings (graph/dimension) when multiple WMMA shape available (e.g., FP16 with 16x8x8 and 16x8x16).

- Kernel Fusion with other layers.

- Current version focuses on training.
- More fusion operation in inference, such as Graphconv+BatchNorm.

$$\text{BN}(x_{i,j}) = \left(\frac{x_{i,j} - \mathbb{E}[x_{*,j}]}{\sqrt{\text{Var}[x_{*,j}] + \epsilon}} \right) \cdot \gamma_j + \beta_j$$

Thank You

Email: yuke_wang@cs.ucsb.edu

Paper: <https://arxiv.org/pdf/2112.02052.pdf>

Code: <https://github.com/YukeWang96/TCGNN-Pytorch.git>