

DynaGraph: Dynamic Graph Neural Networks at Scale

Mingyu Guan*, Anand Iyer[^], Taesoo Kim*

*Georgia Institute of Technology [^]Microsoft Research

GRADES-NDA 2022



Graph Neural Networks (GNNs)



- The recent past has seen an increasing interest in GNNs.
- Node embeddings are generated by combining **graph structure** and **feature information**.
- Most GNN models can fit into the **Message Passing Paradigm**.






Initial Features/Embeddings of Each Node

Output Features/Embeddings of Each Node

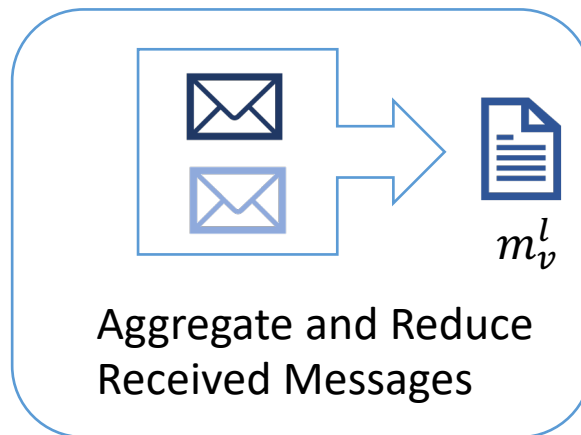
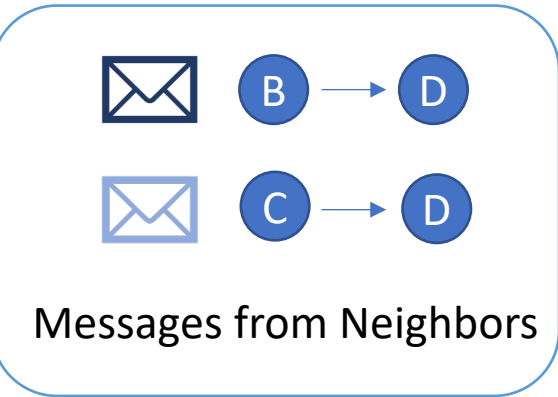
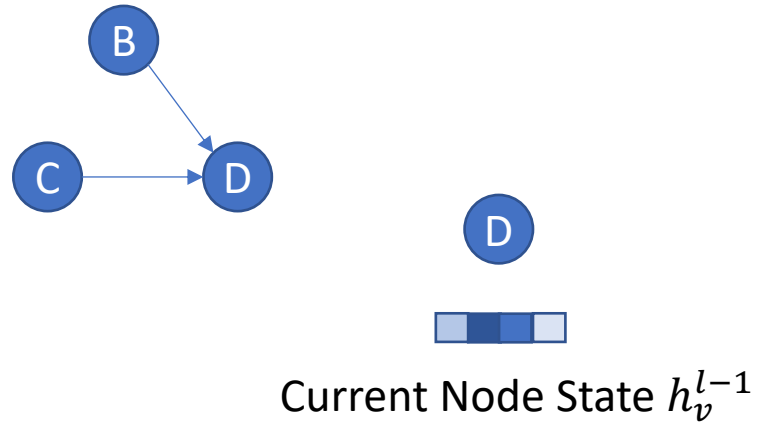

Message Passing Paradigm




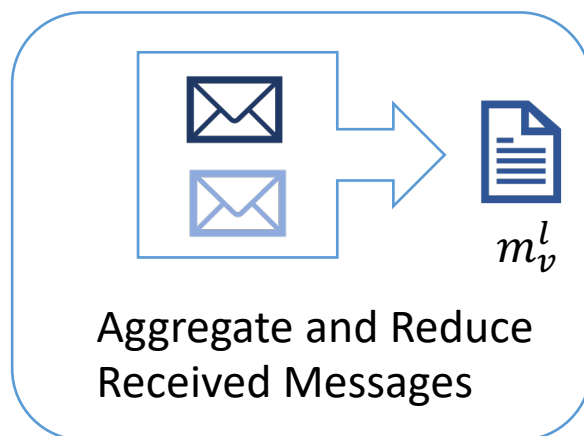


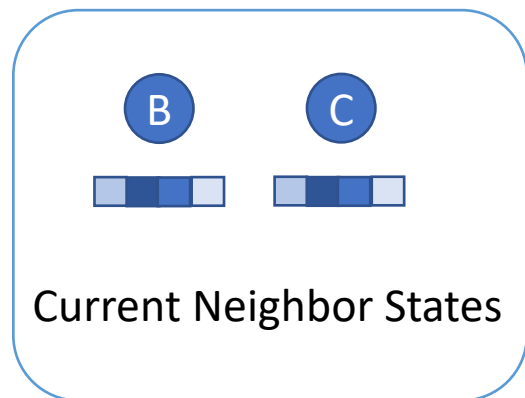
Message Passing Paradigm



Message Passing Paradigm



Message Passing Paradigm

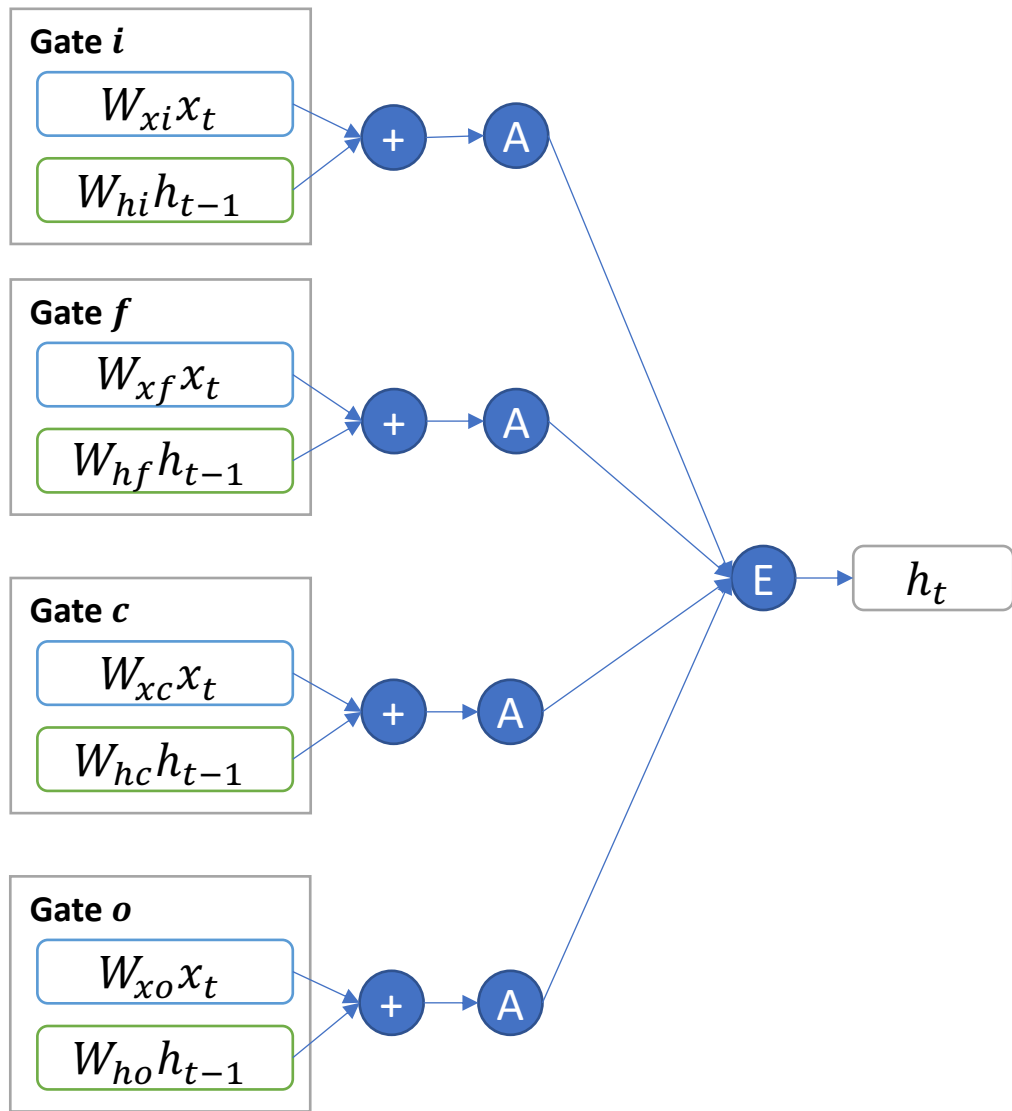


Dynamic GNNs

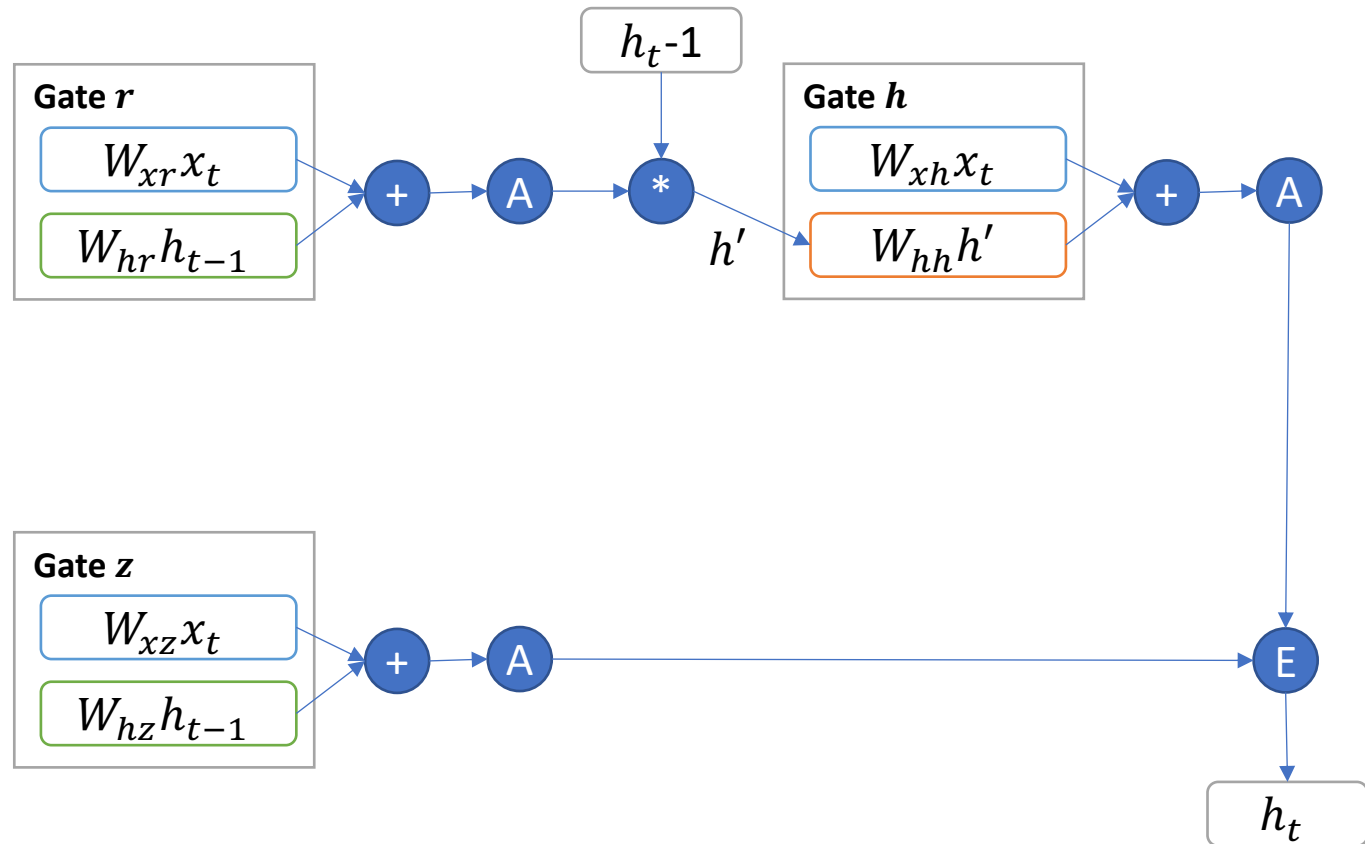
- Most of existing GNN frameworks assume that the input graph is static.
- Real-world graphs are often *dynamic* in nature.
- Representation: a time series of *snapshots* of the graph.
- **Common approach:** Combine GNNs and RNNs.
 - GNNs for encoding spatial information (graph structure)
 - RNNs for encoding temporal information



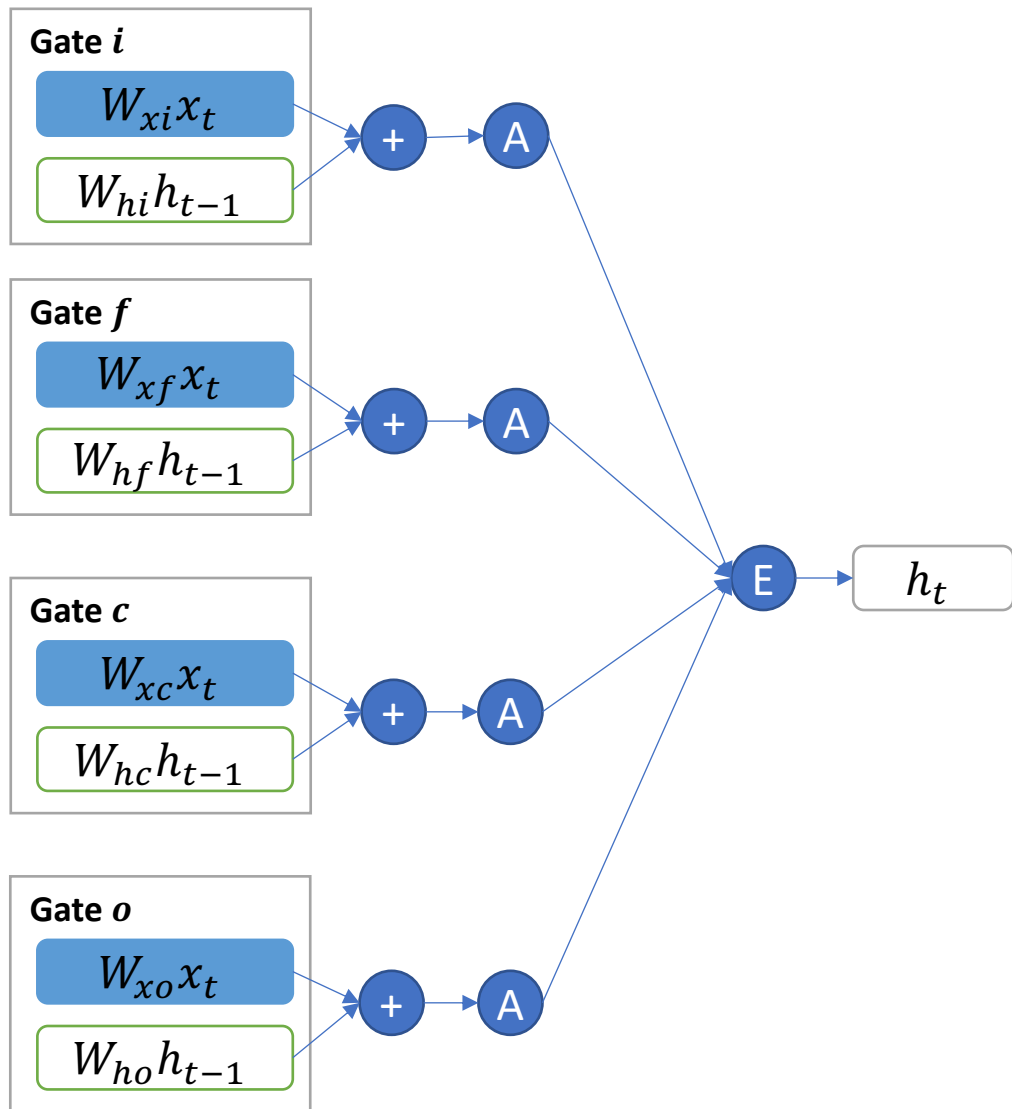
LSTM



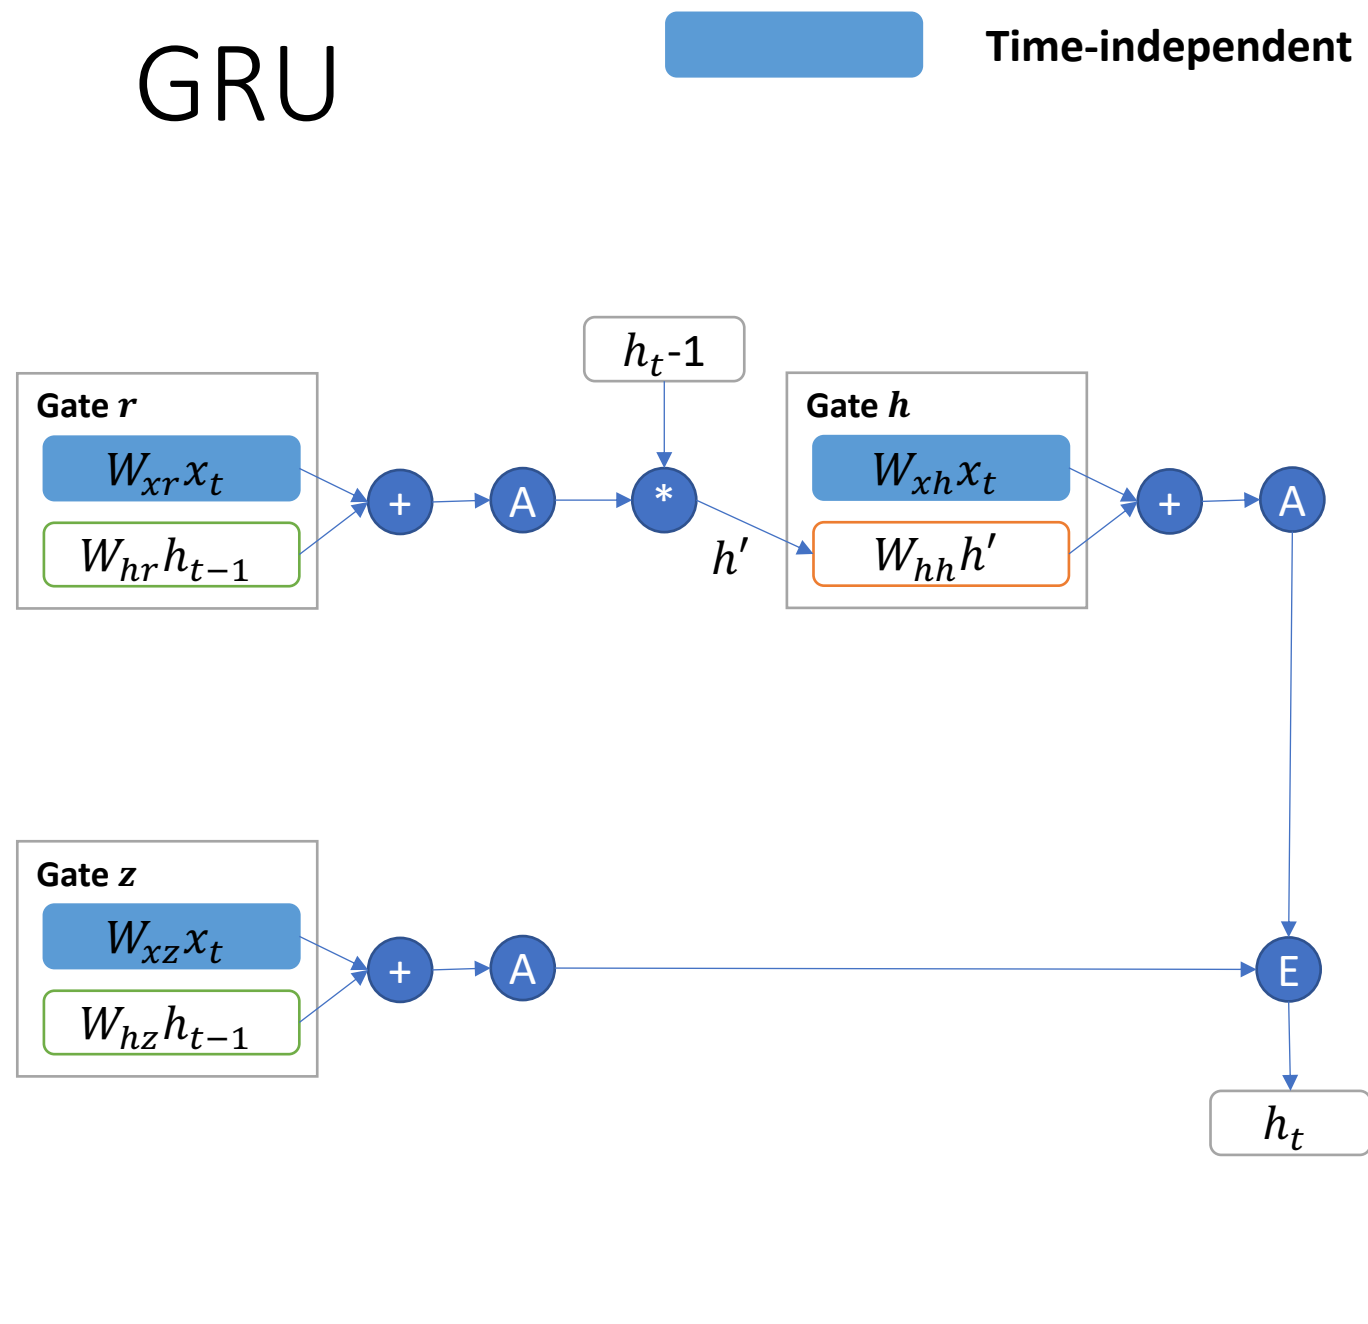
GRU




LSTM



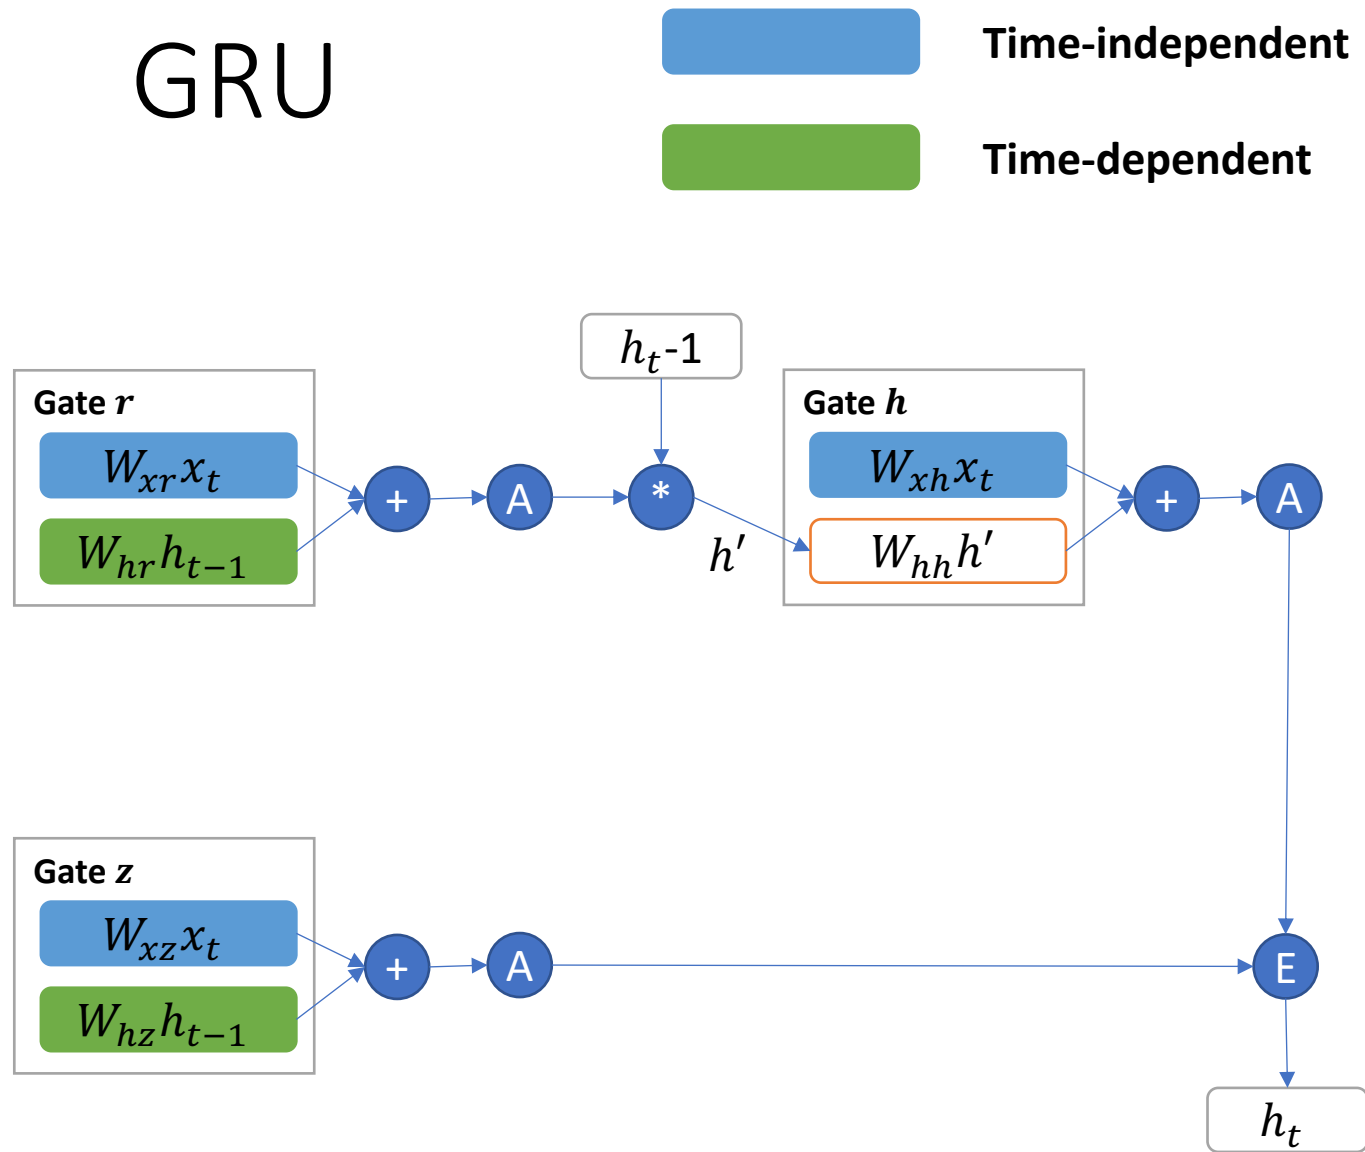
GRU



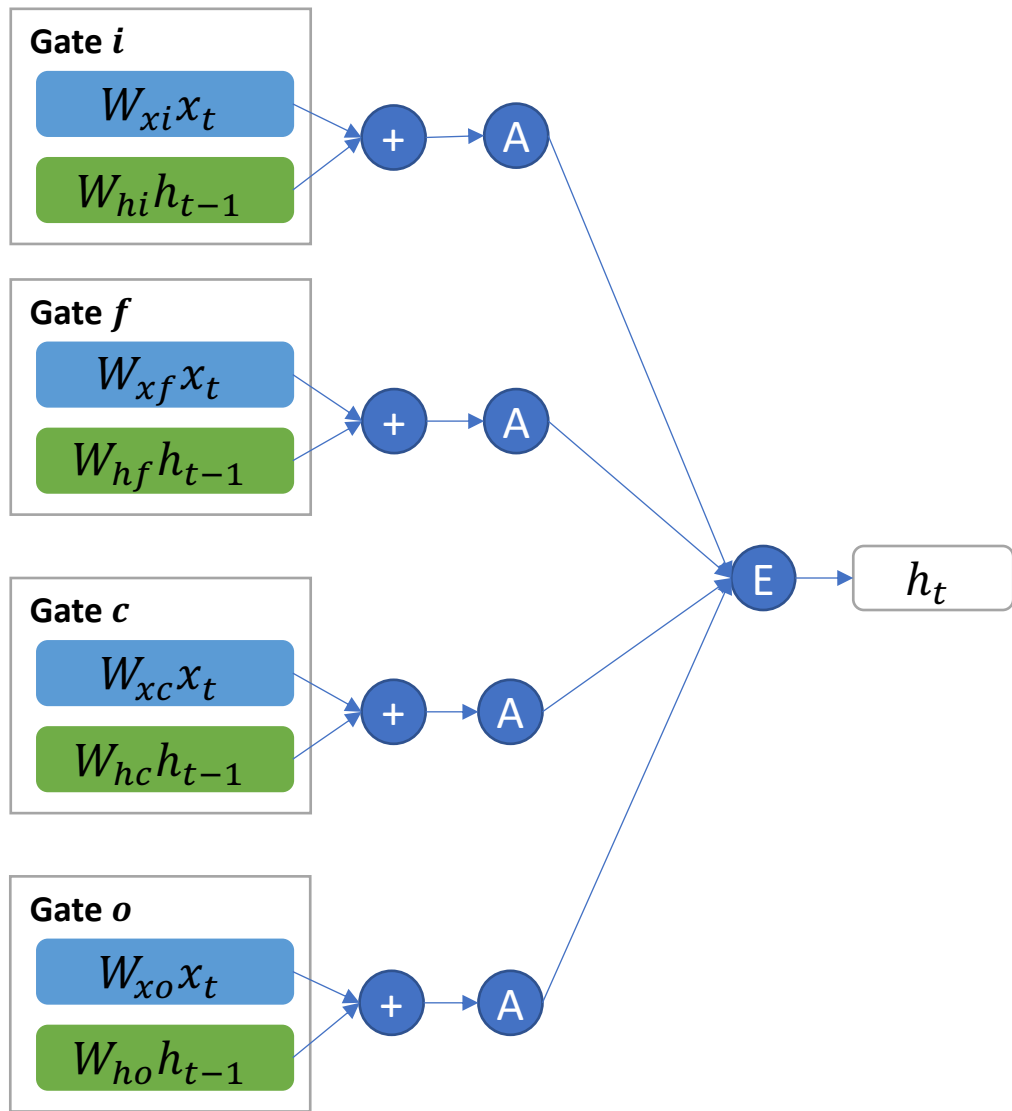
LSTM



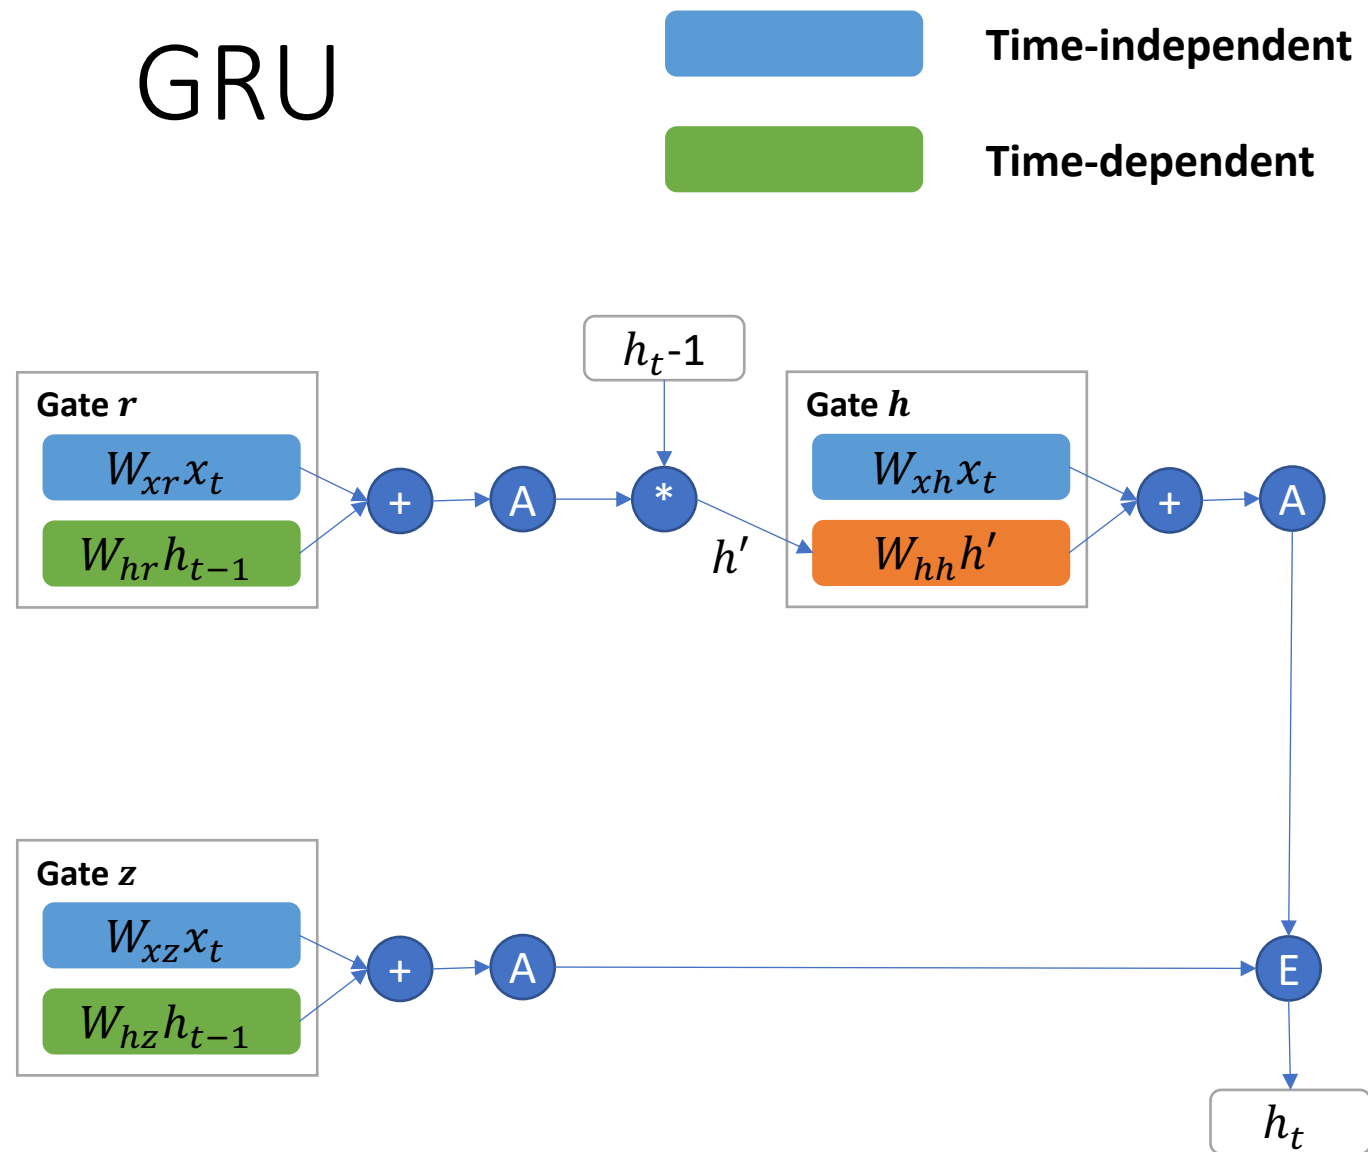
GRU




LSTM



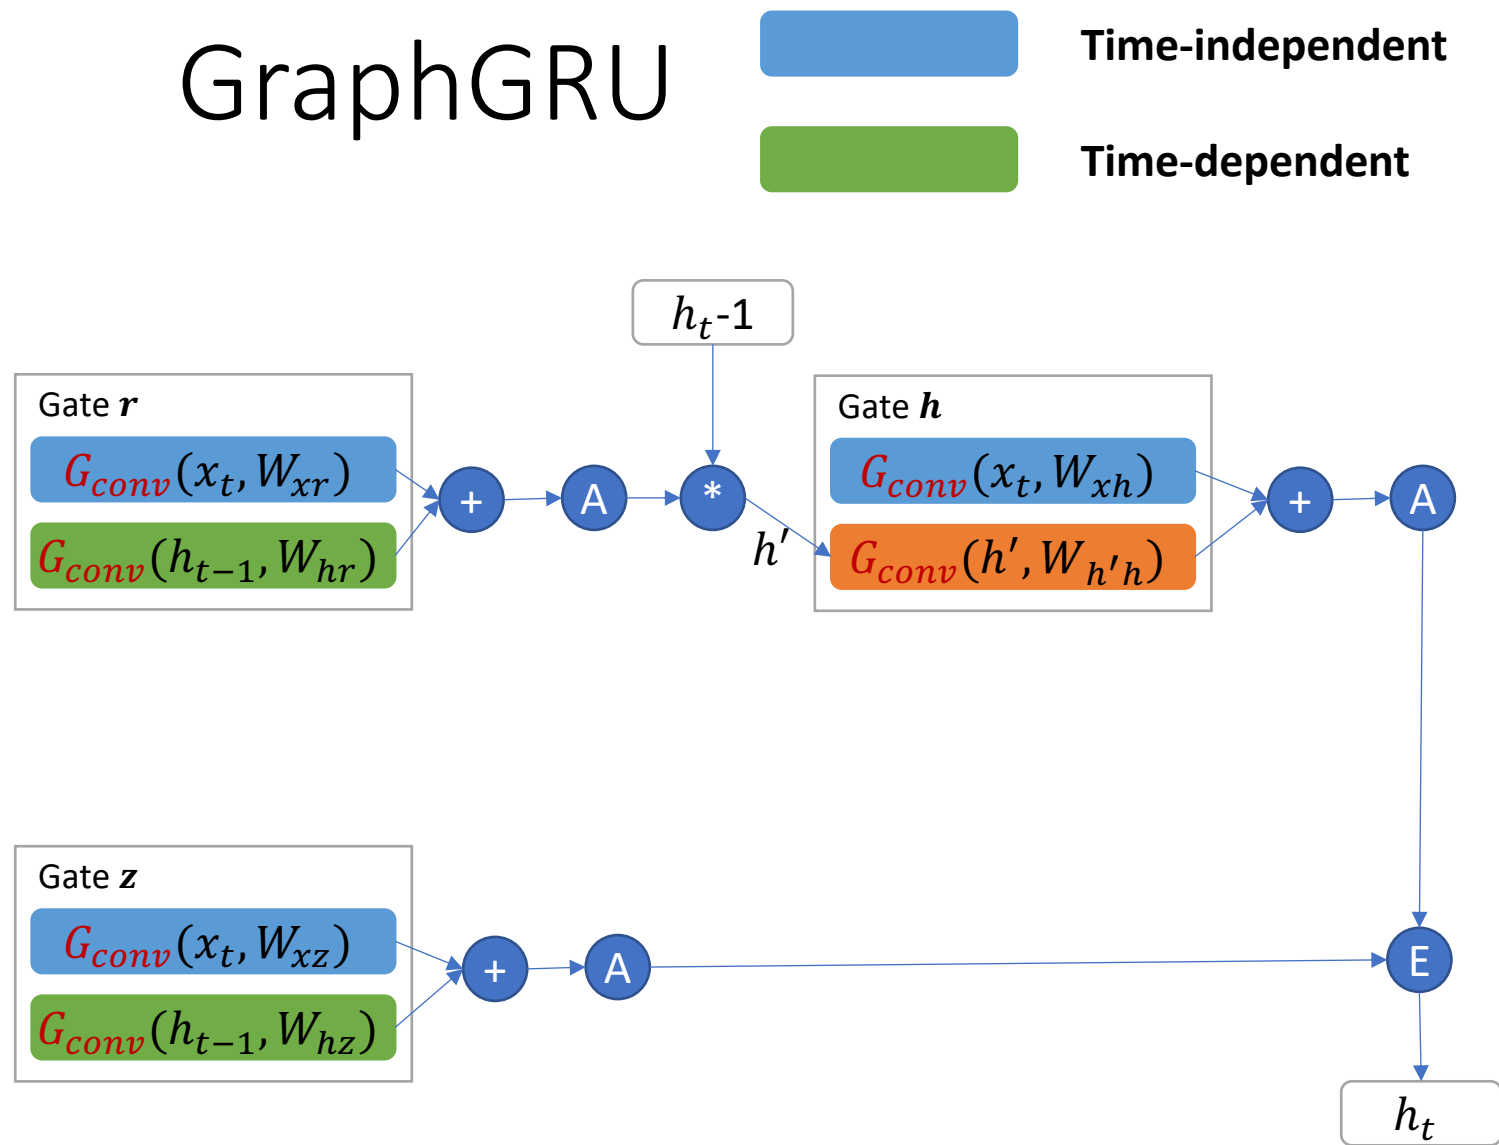
GRU



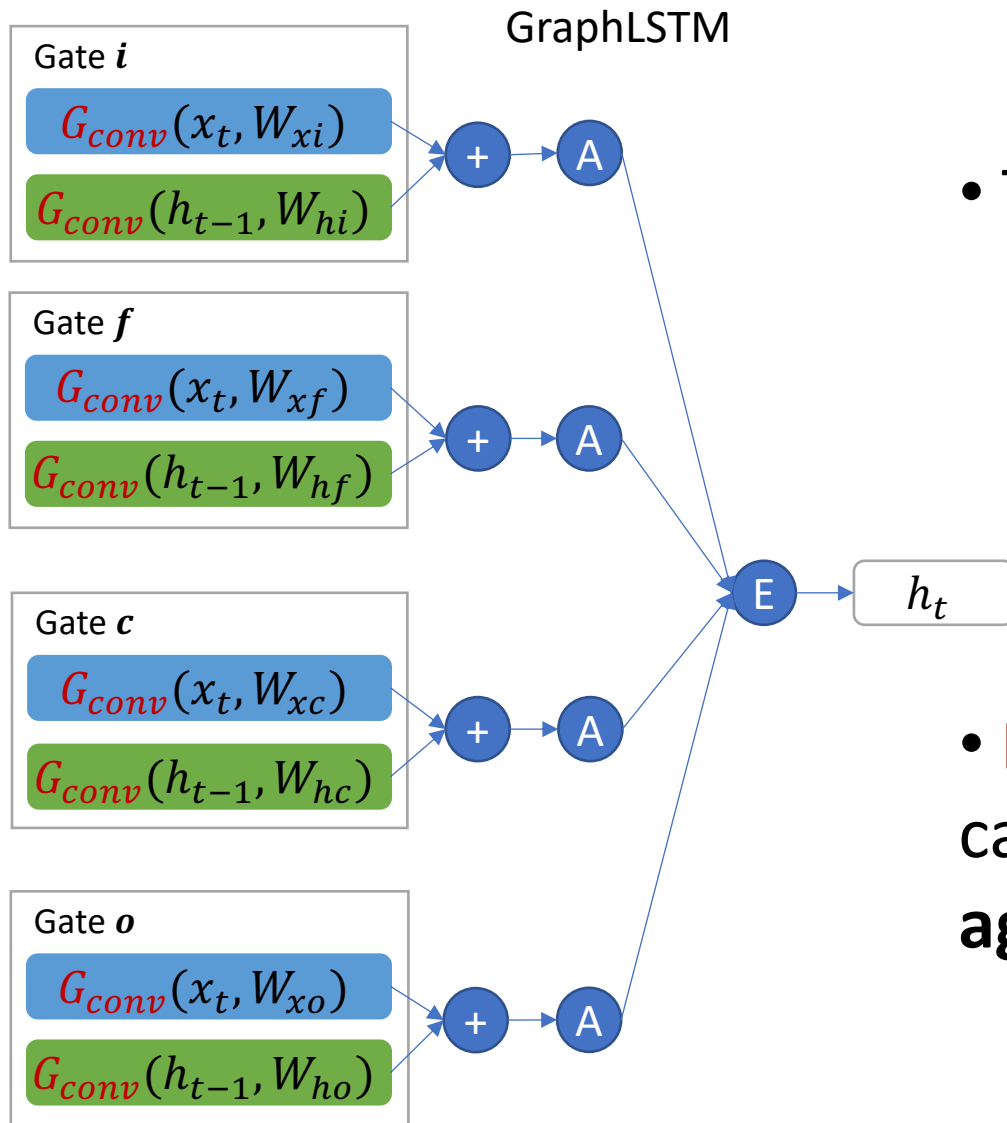
GraphLSTM



GraphGRU



Challenge #1: Redundant Neighborhood Aggregation

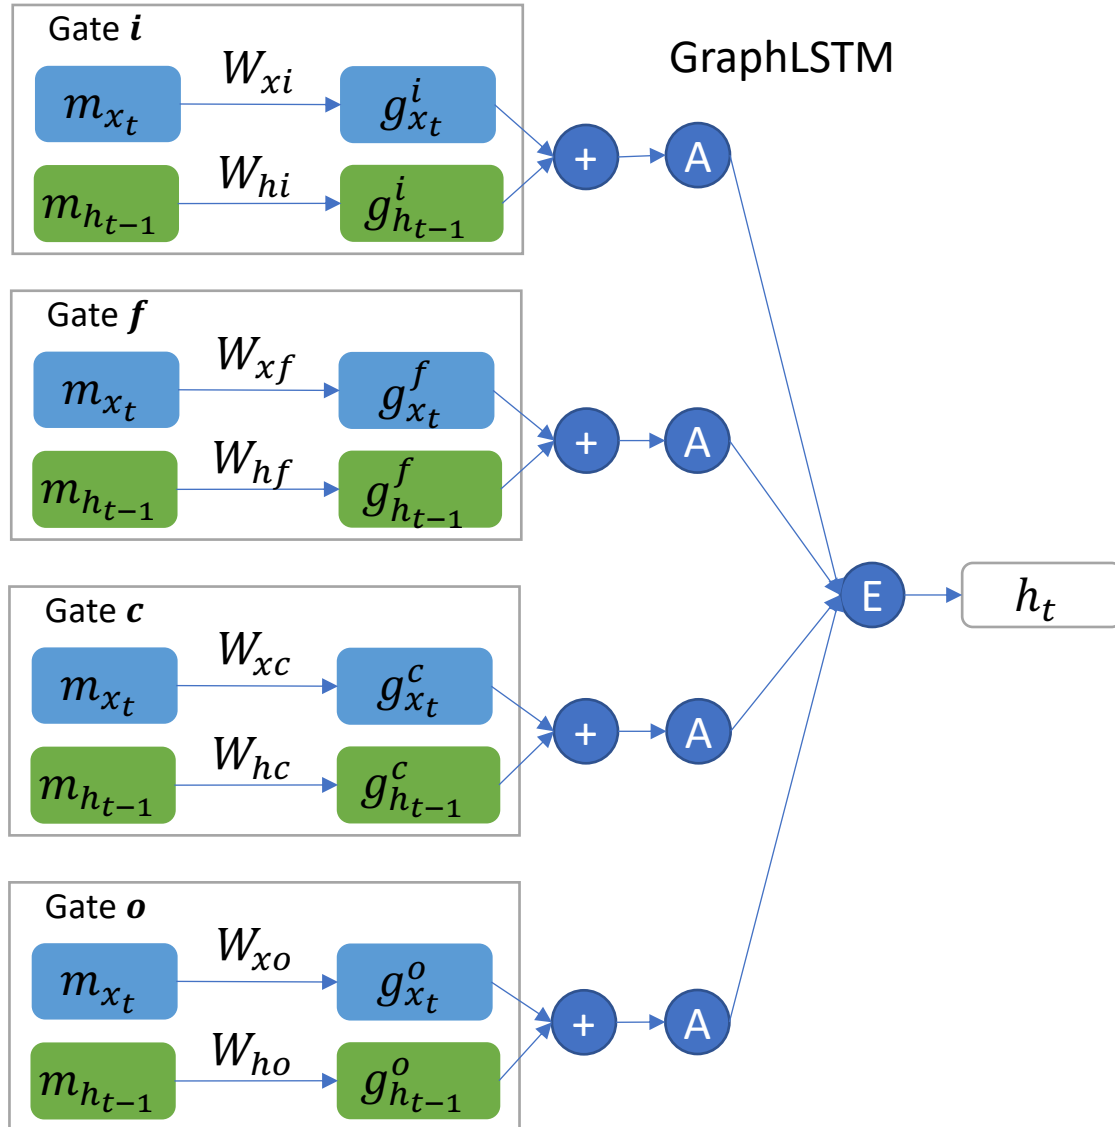


- Two categories of graph convolutions.
 - **Time-independent** graph convolution depends on current representations of nodes.
 - **Time-dependent** graph convolution depends on previous hidden states.
- **Redundancy**: Graph convolutions in the same category perform **same neighborhood aggregation**.

Challenge #2: Inefficient Distributed Training

- No existing systems for training static GNNs, for example, DGL, support distributed dynamic GNN training in an efficient way.
- Static GNN training:
 - Partitioning both the graph structure and node features across machines.
 - Using data parallelism to train a static GNN.
- Can we partition each snapshot individually?
 - Partitioning and maintaining a large number of snapshots can be **expensive**.
 - The graph structure and the node features in each snapshot may vary.

Cached Message Passing



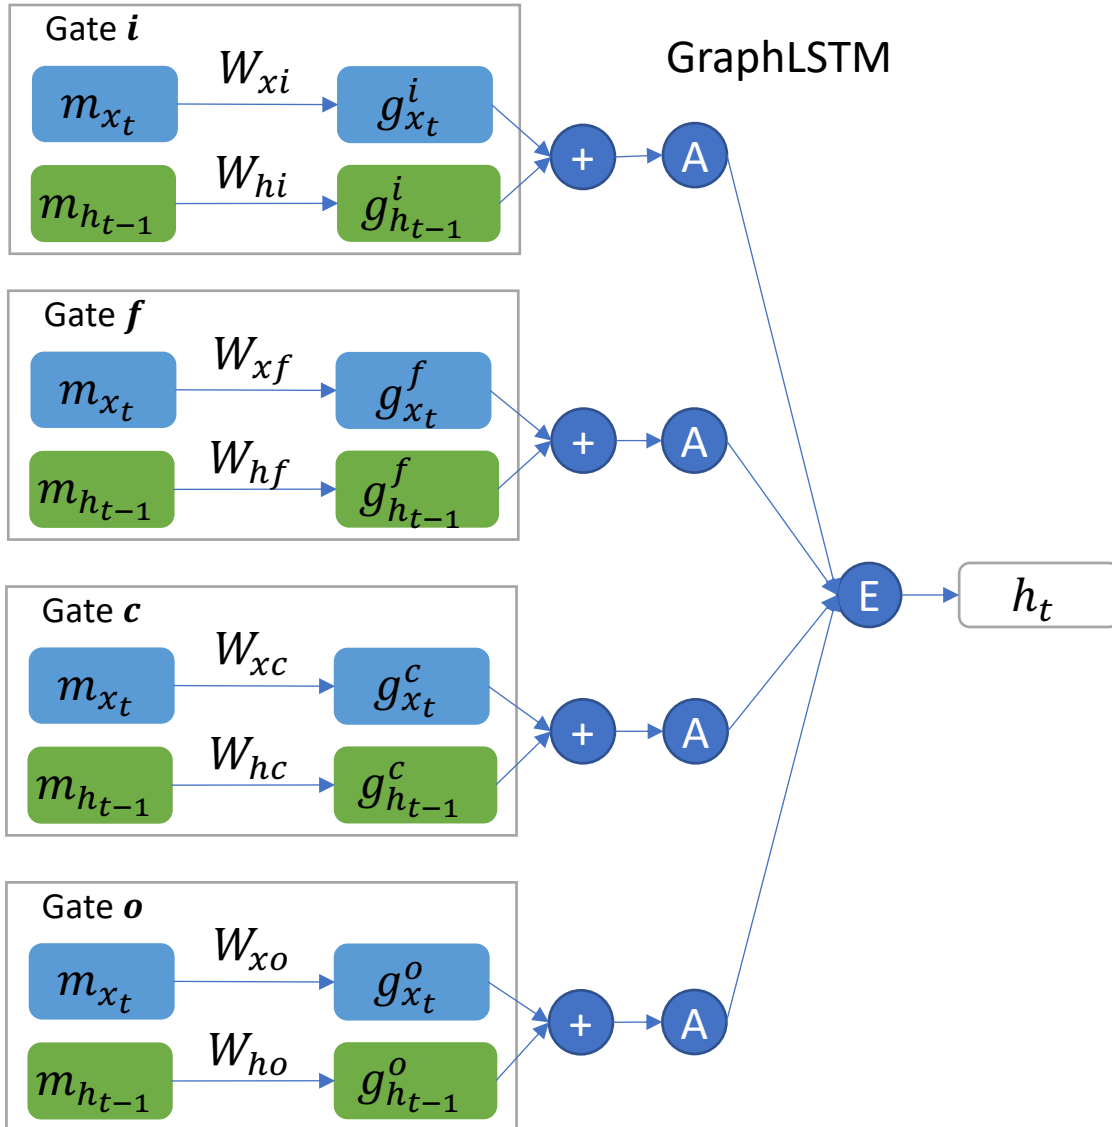
Typical Message Passing Paradigm of GNN:

$$m_{u \rightarrow v}^l = M^l(h_v^{l-1}, h_u^{l-1}, e_{u \rightarrow v}^{l-1})$$

$$m_v^l = \sum_{u \in N(v)} m_{u \rightarrow v}^l$$

$$h_v^l = U^l(h_v^{l-1}, m_v^l)$$

Cached Message Passing



Typical Message Passing Paradigm of GNN:

$$m_{u \rightarrow v}^l = M^l(h_v^{l-1}, h_u^{l-1}, e_{u \rightarrow v}^{l-1})$$

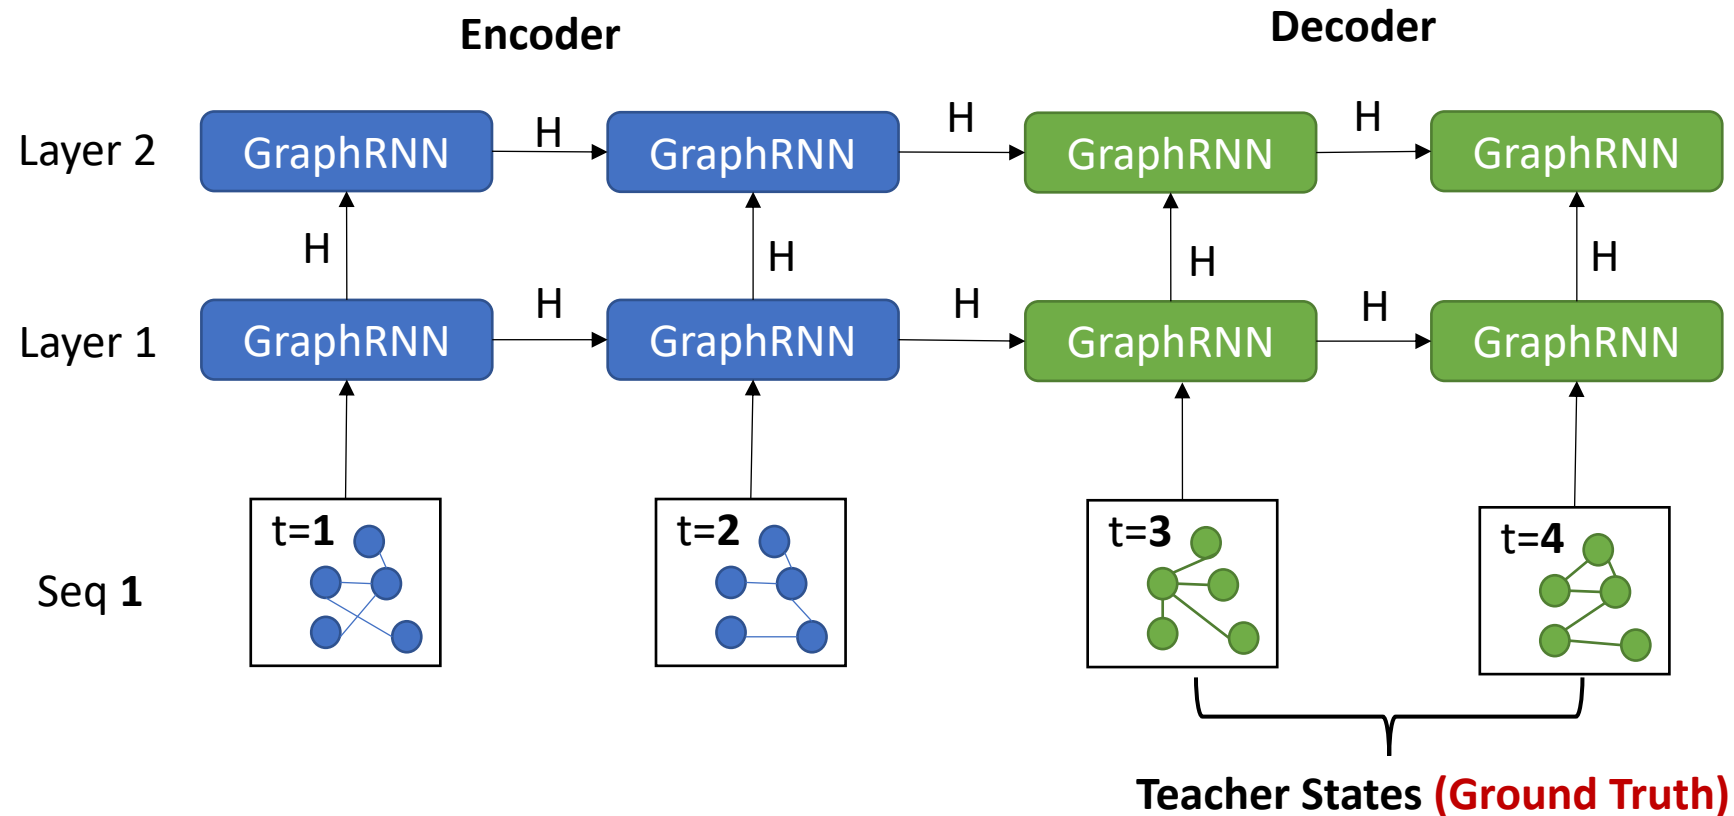
$$m_v^l = \sum_{u \in N(v)} m_{u \rightarrow v}^l$$

$$h_v^l = U^l(h_v^{l-1}, m_v^l)$$

The results after the message passing can be reused for all graph convolution in the same category.

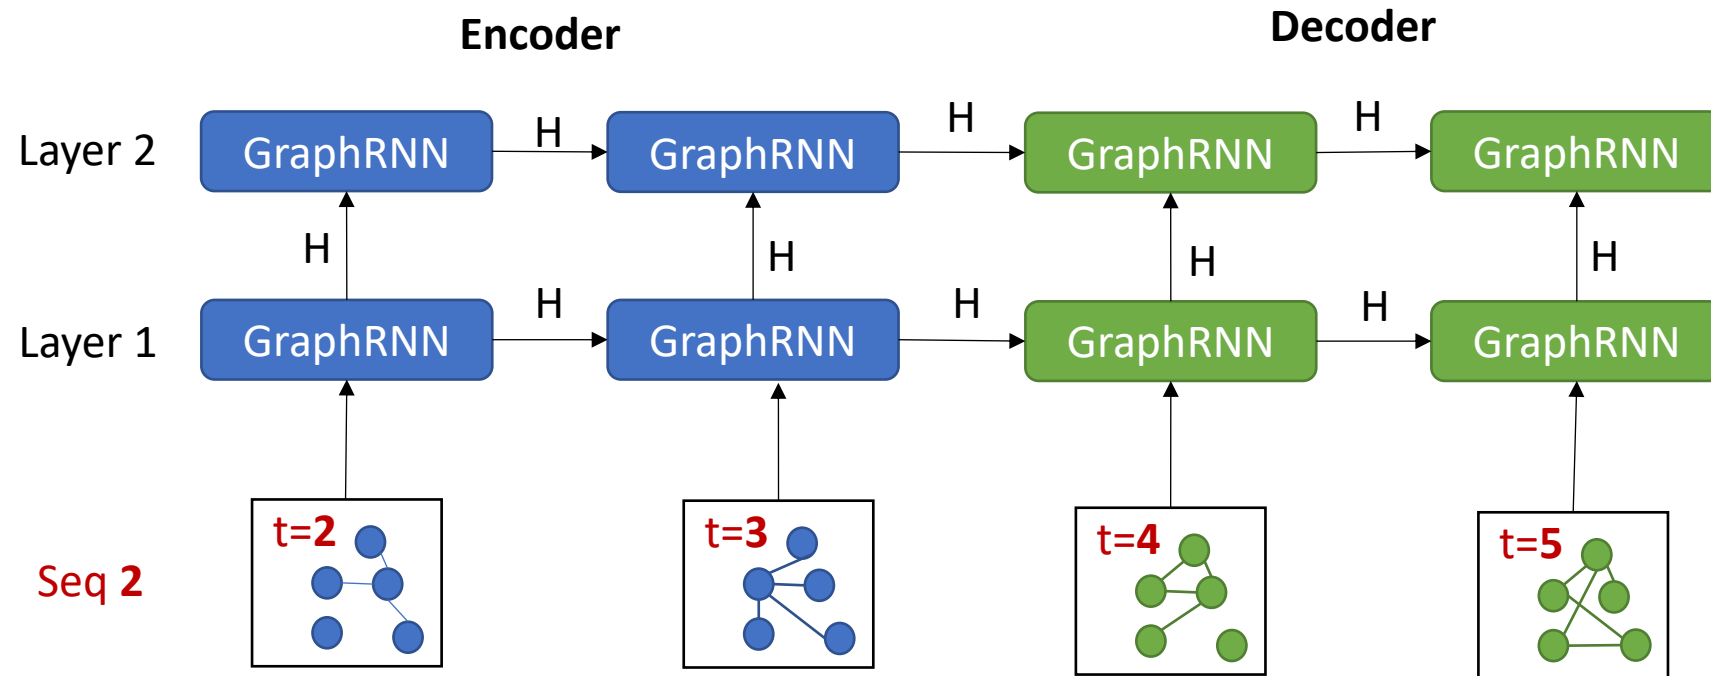
Cached Message Passing

- Dynamic graphs are often trained using **sequence-to-sequence** models in a **sliding-window** fashion.



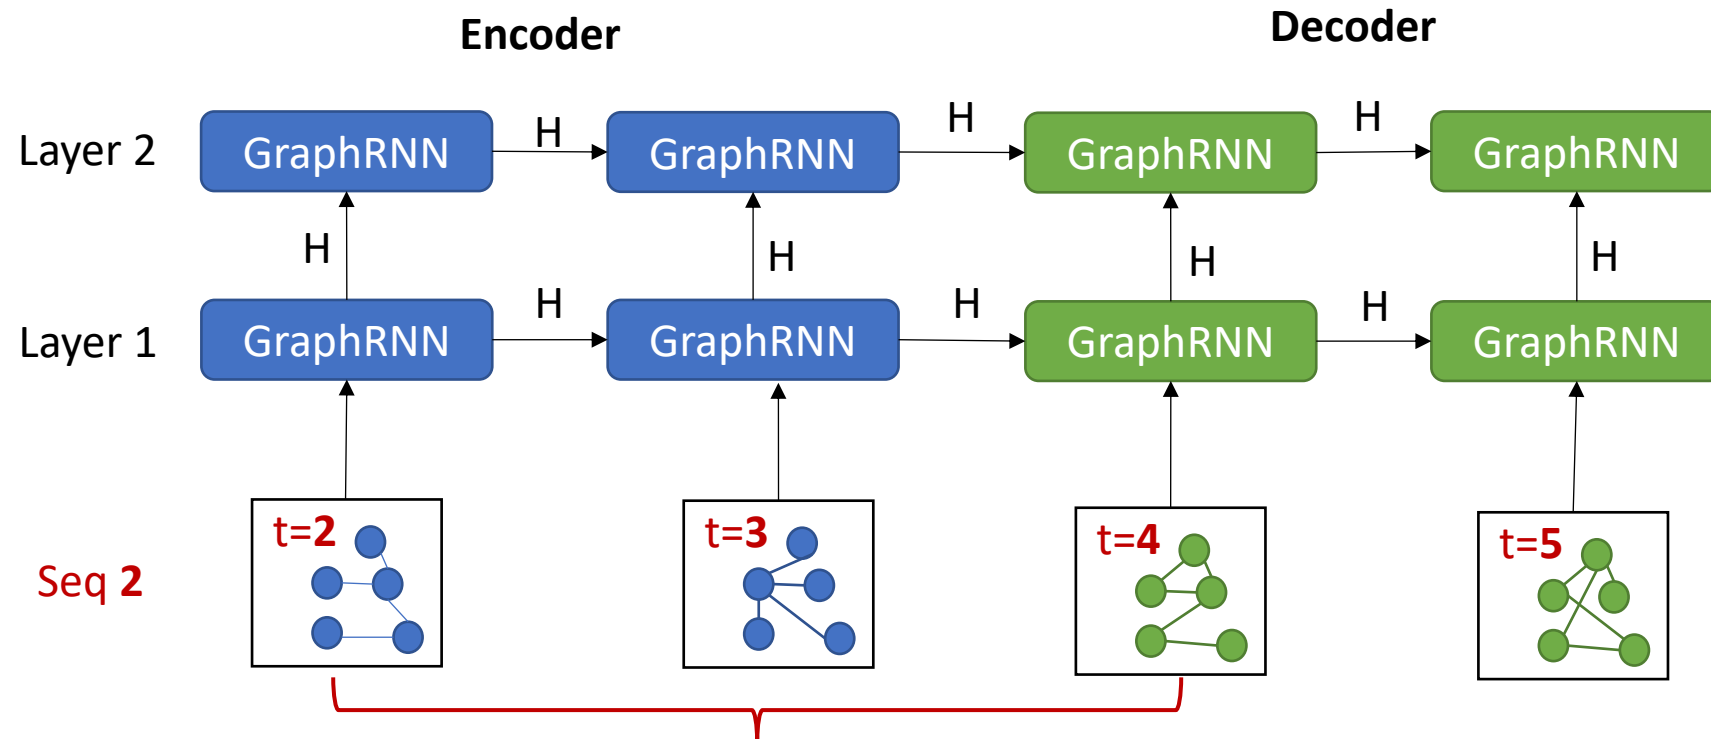
Cached Message Passing

- Dynamic graphs are often trained using **sequence-to-sequence** models in a **sliding-window** fashion.



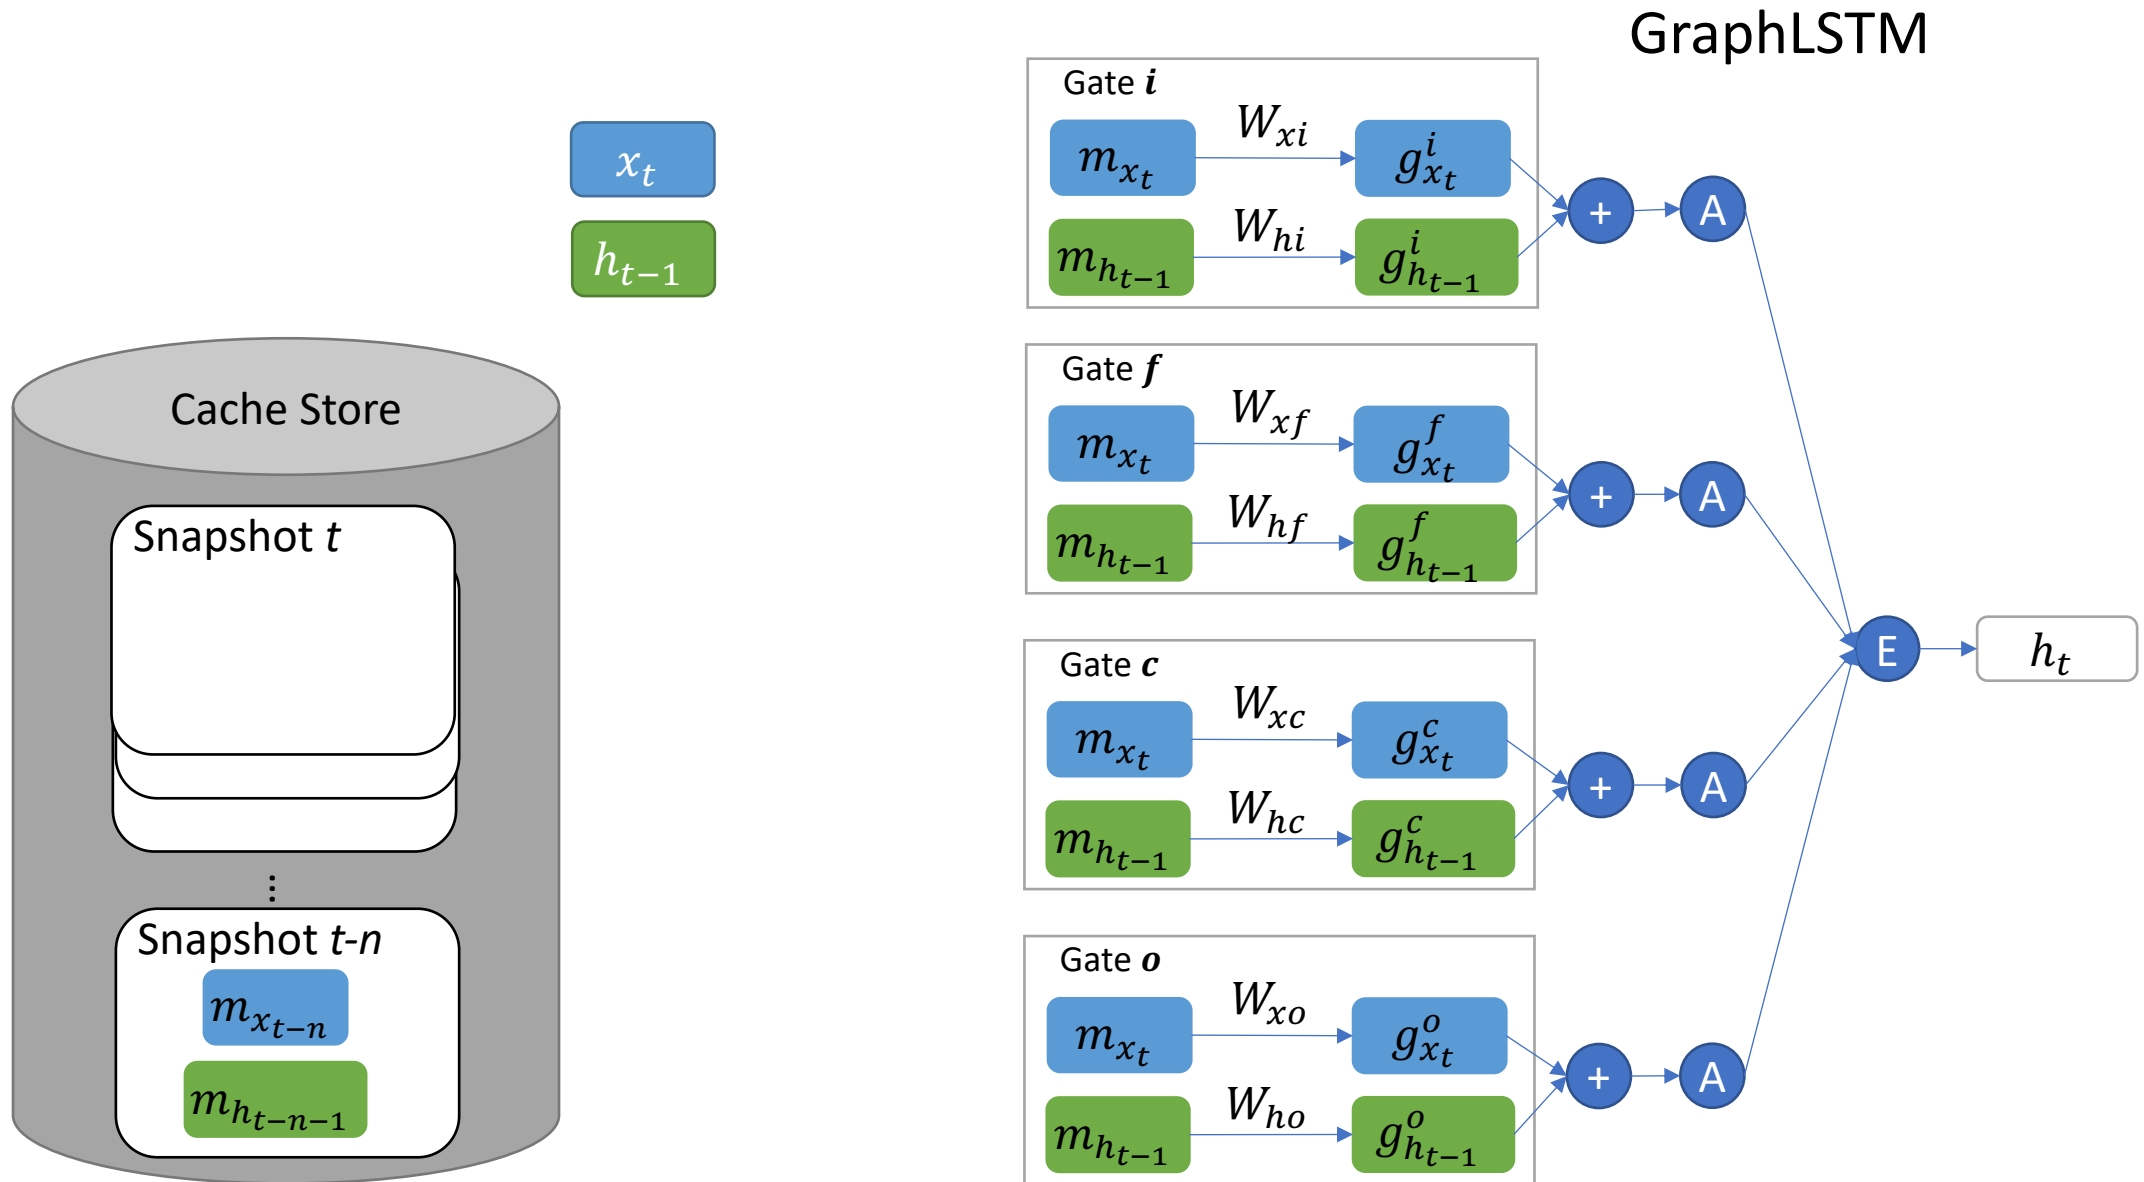
Cached Message Passing

- Dynamic graphs are often trained using **sequence-to-sequence** models in a **sliding-window** fashion.

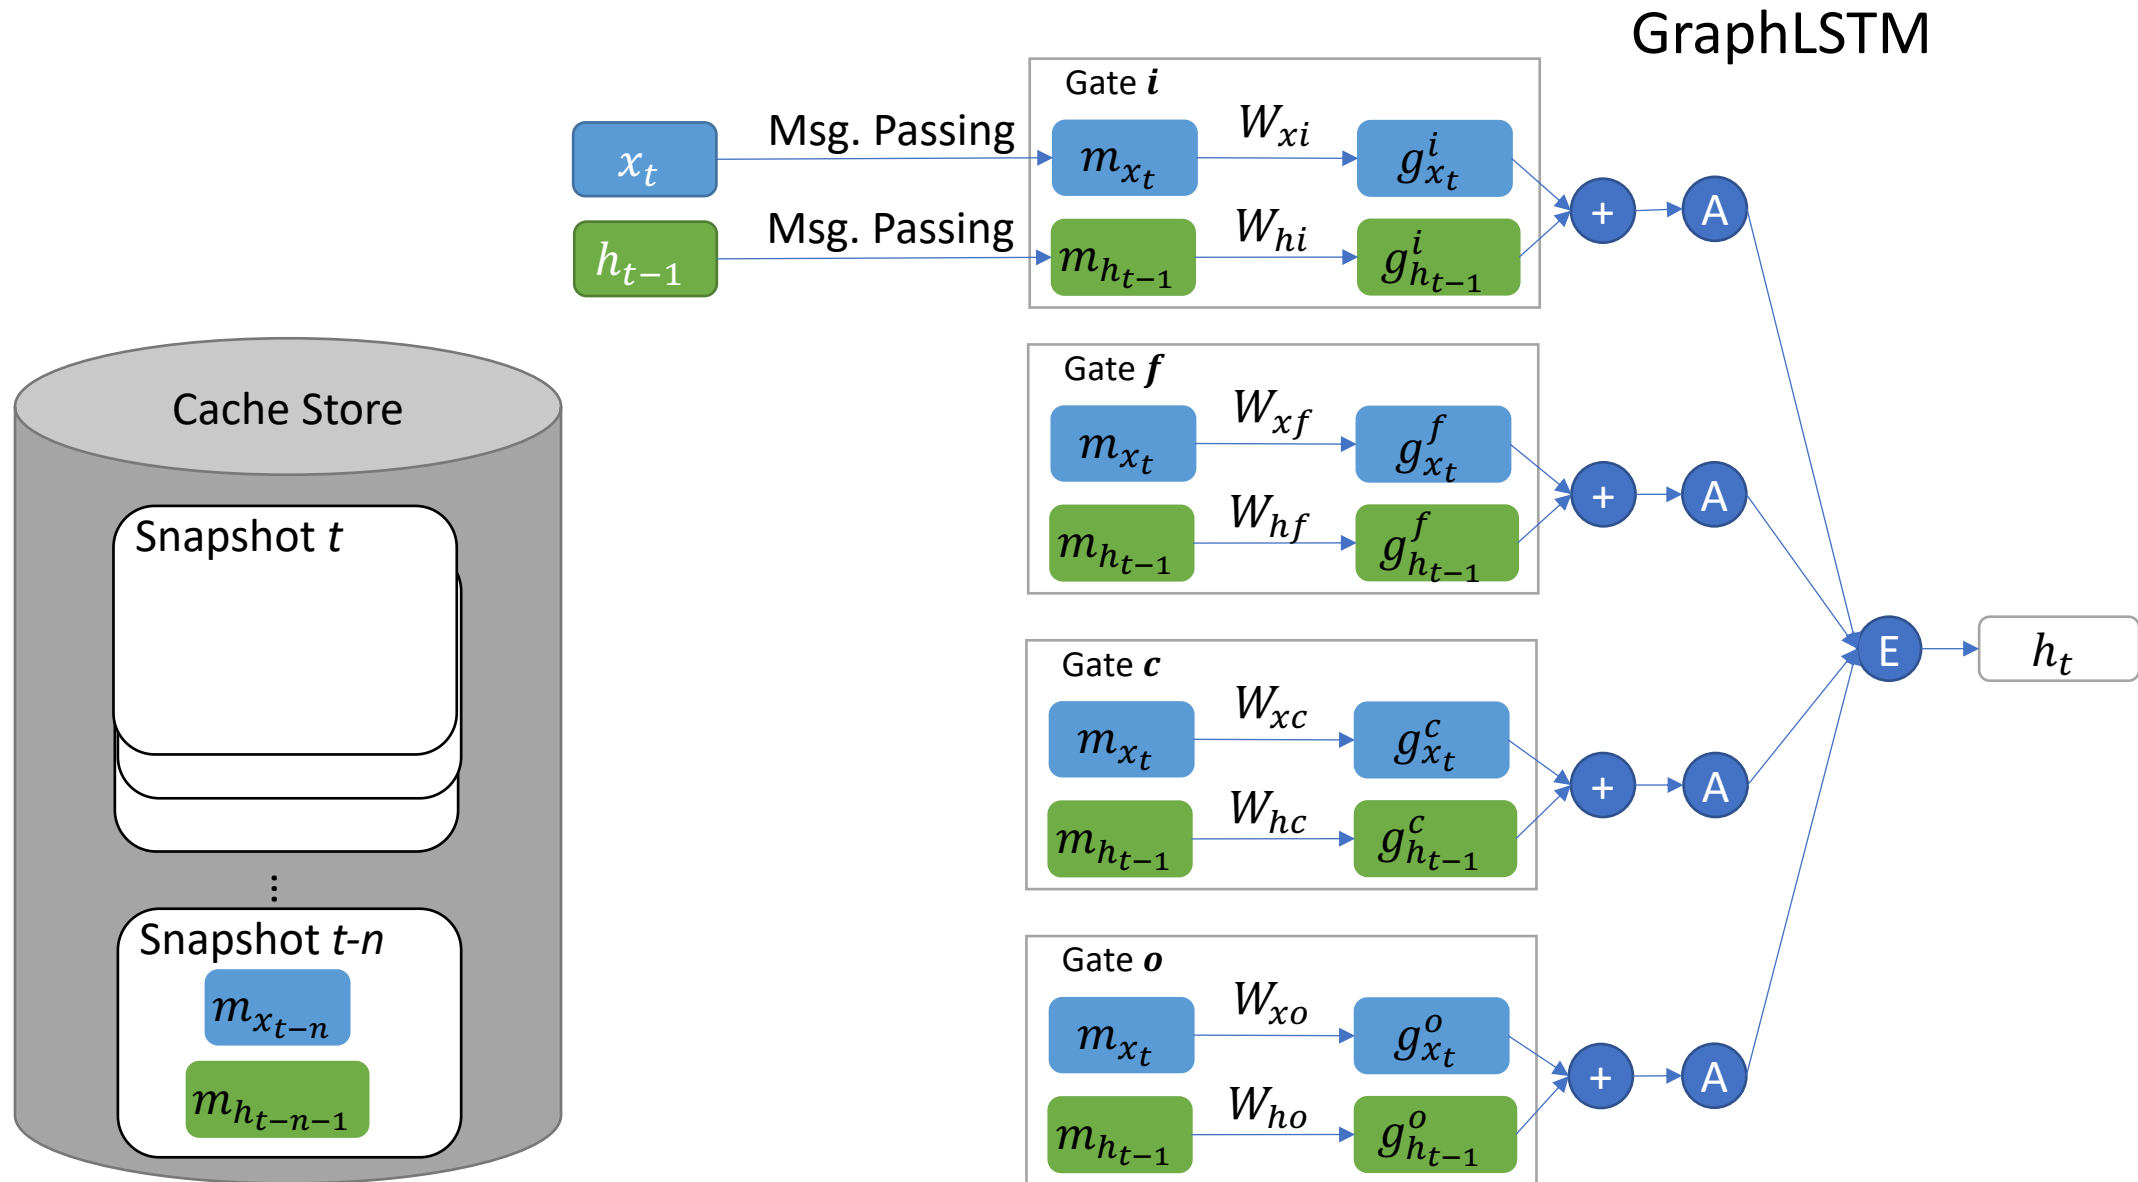


Neighborhood aggregation has already been performed in previous sequence(s)!

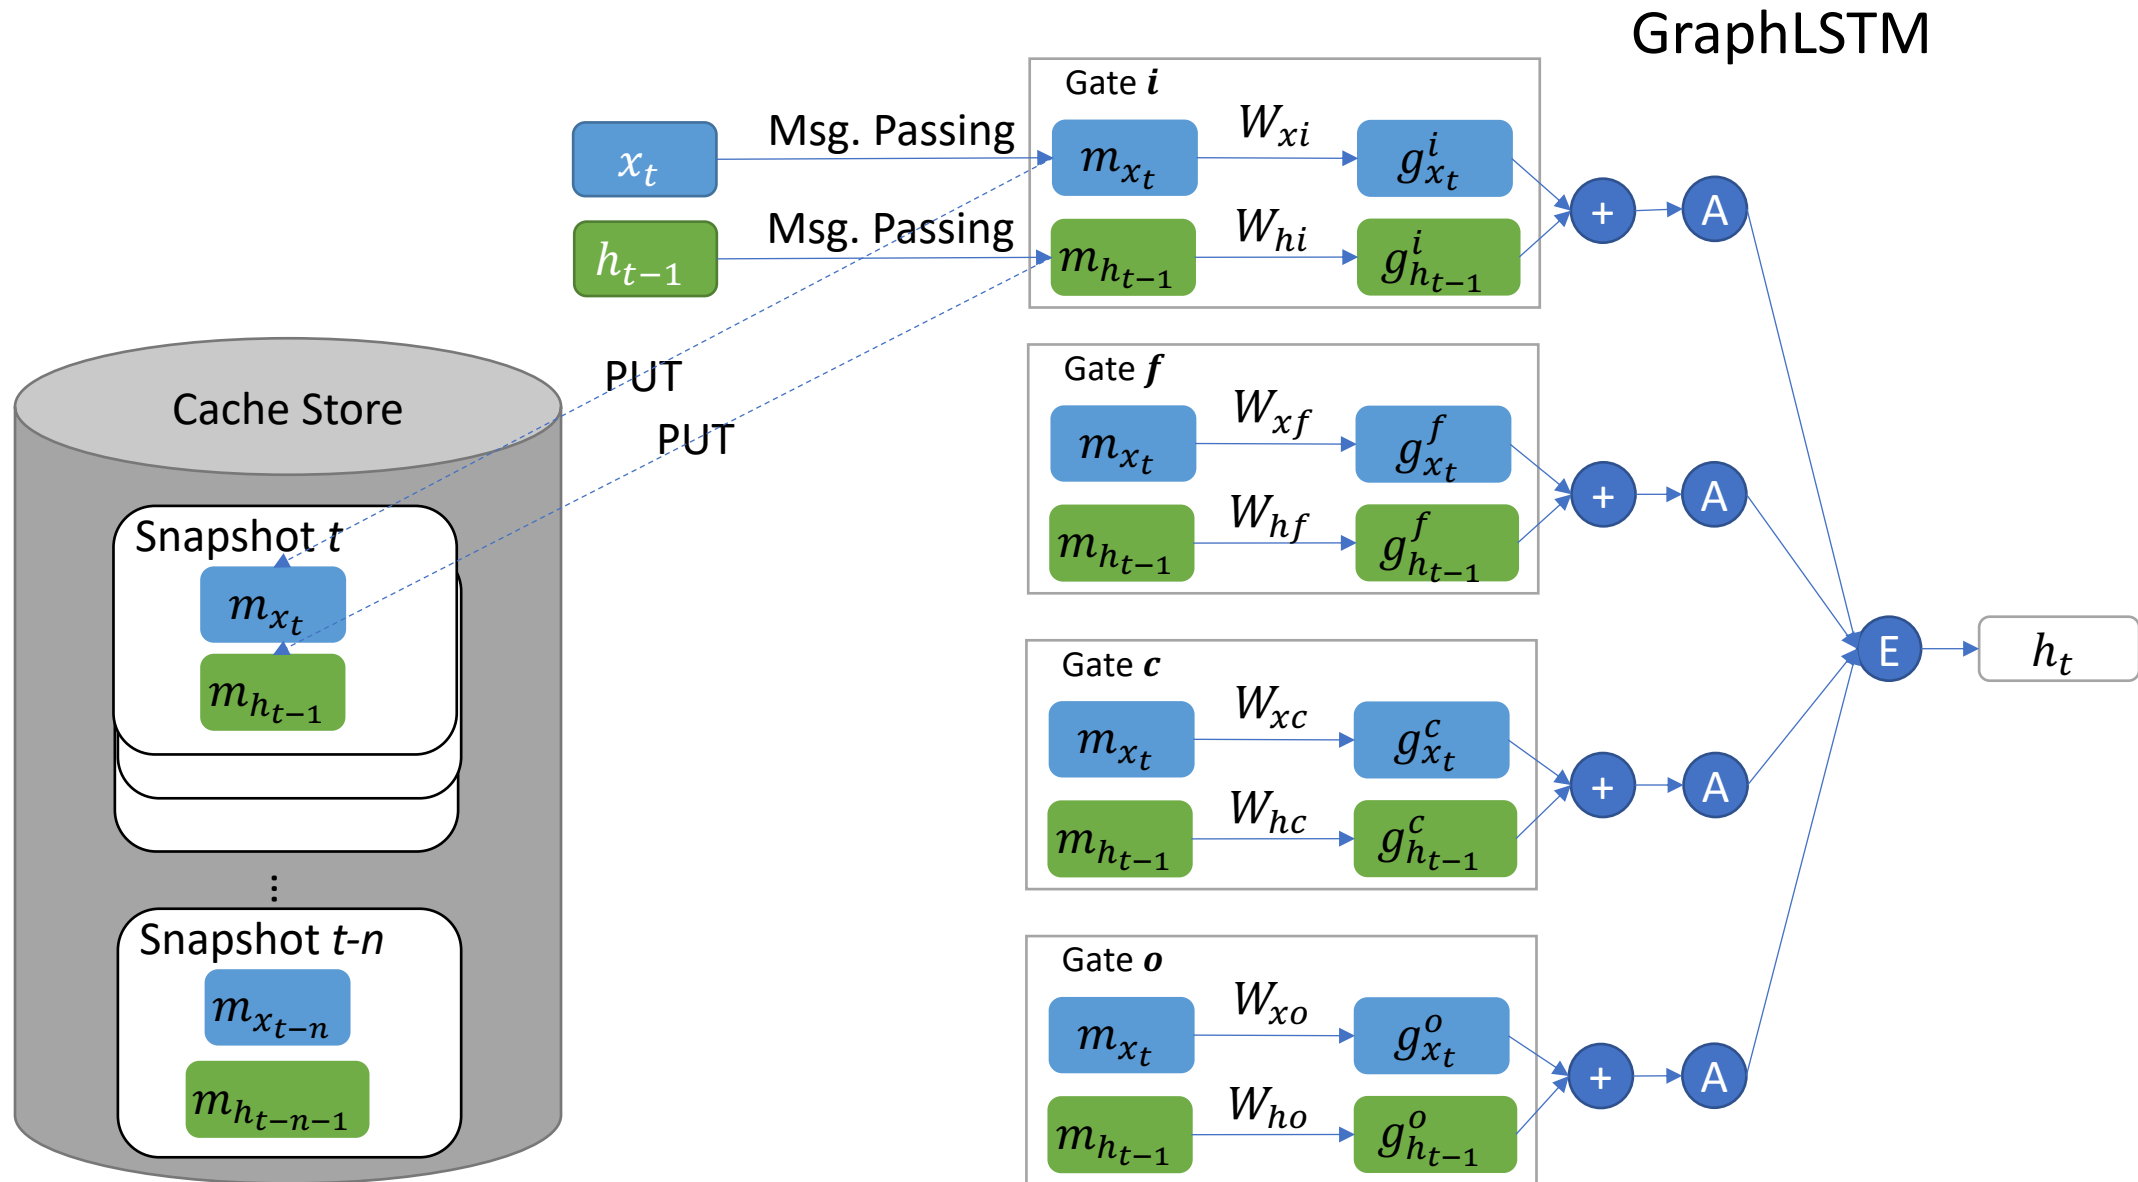
Cached Message Passing



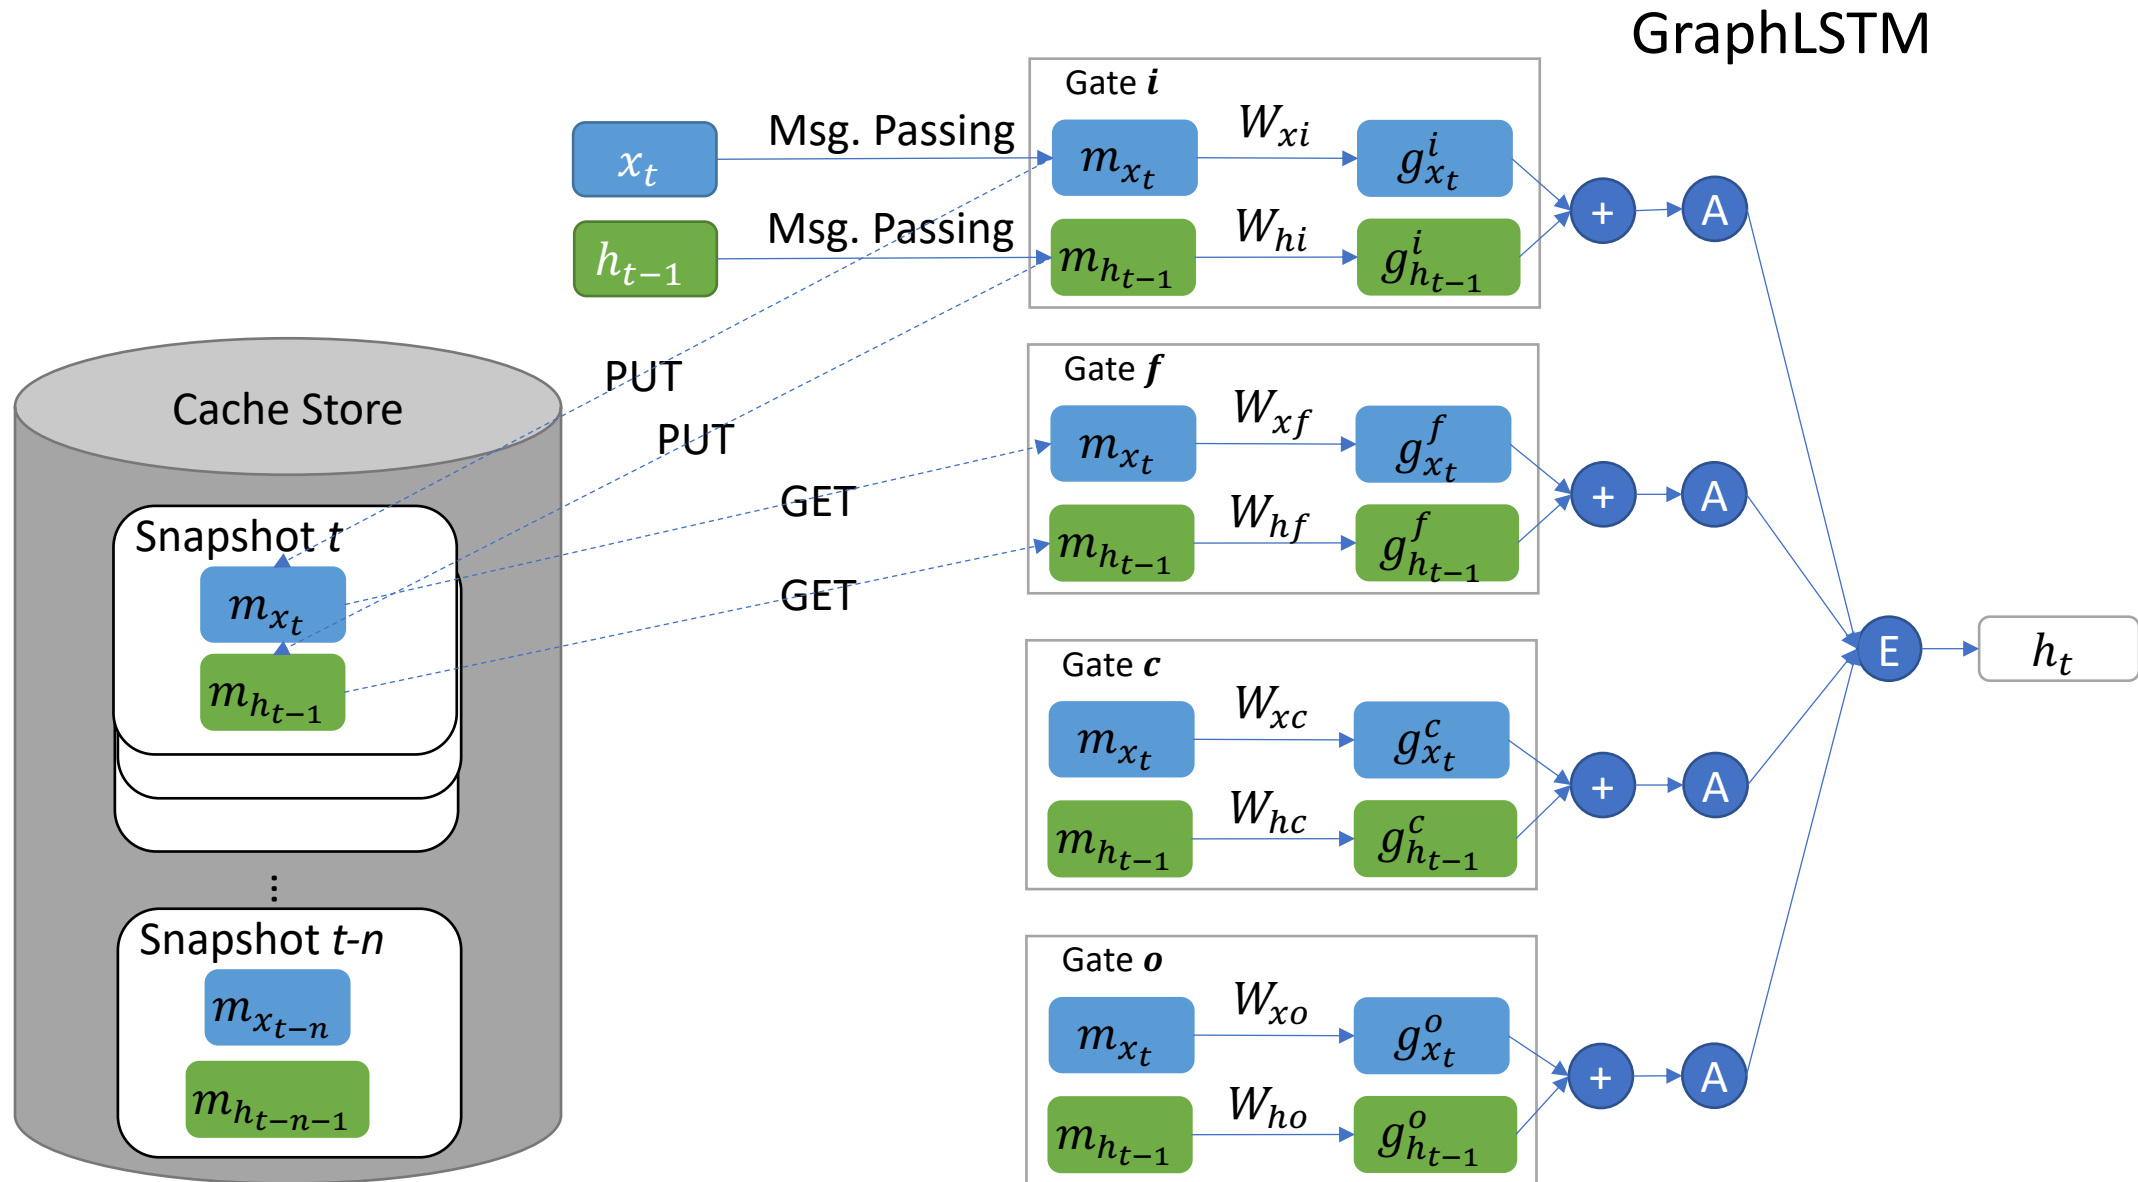
Cached Message Passing



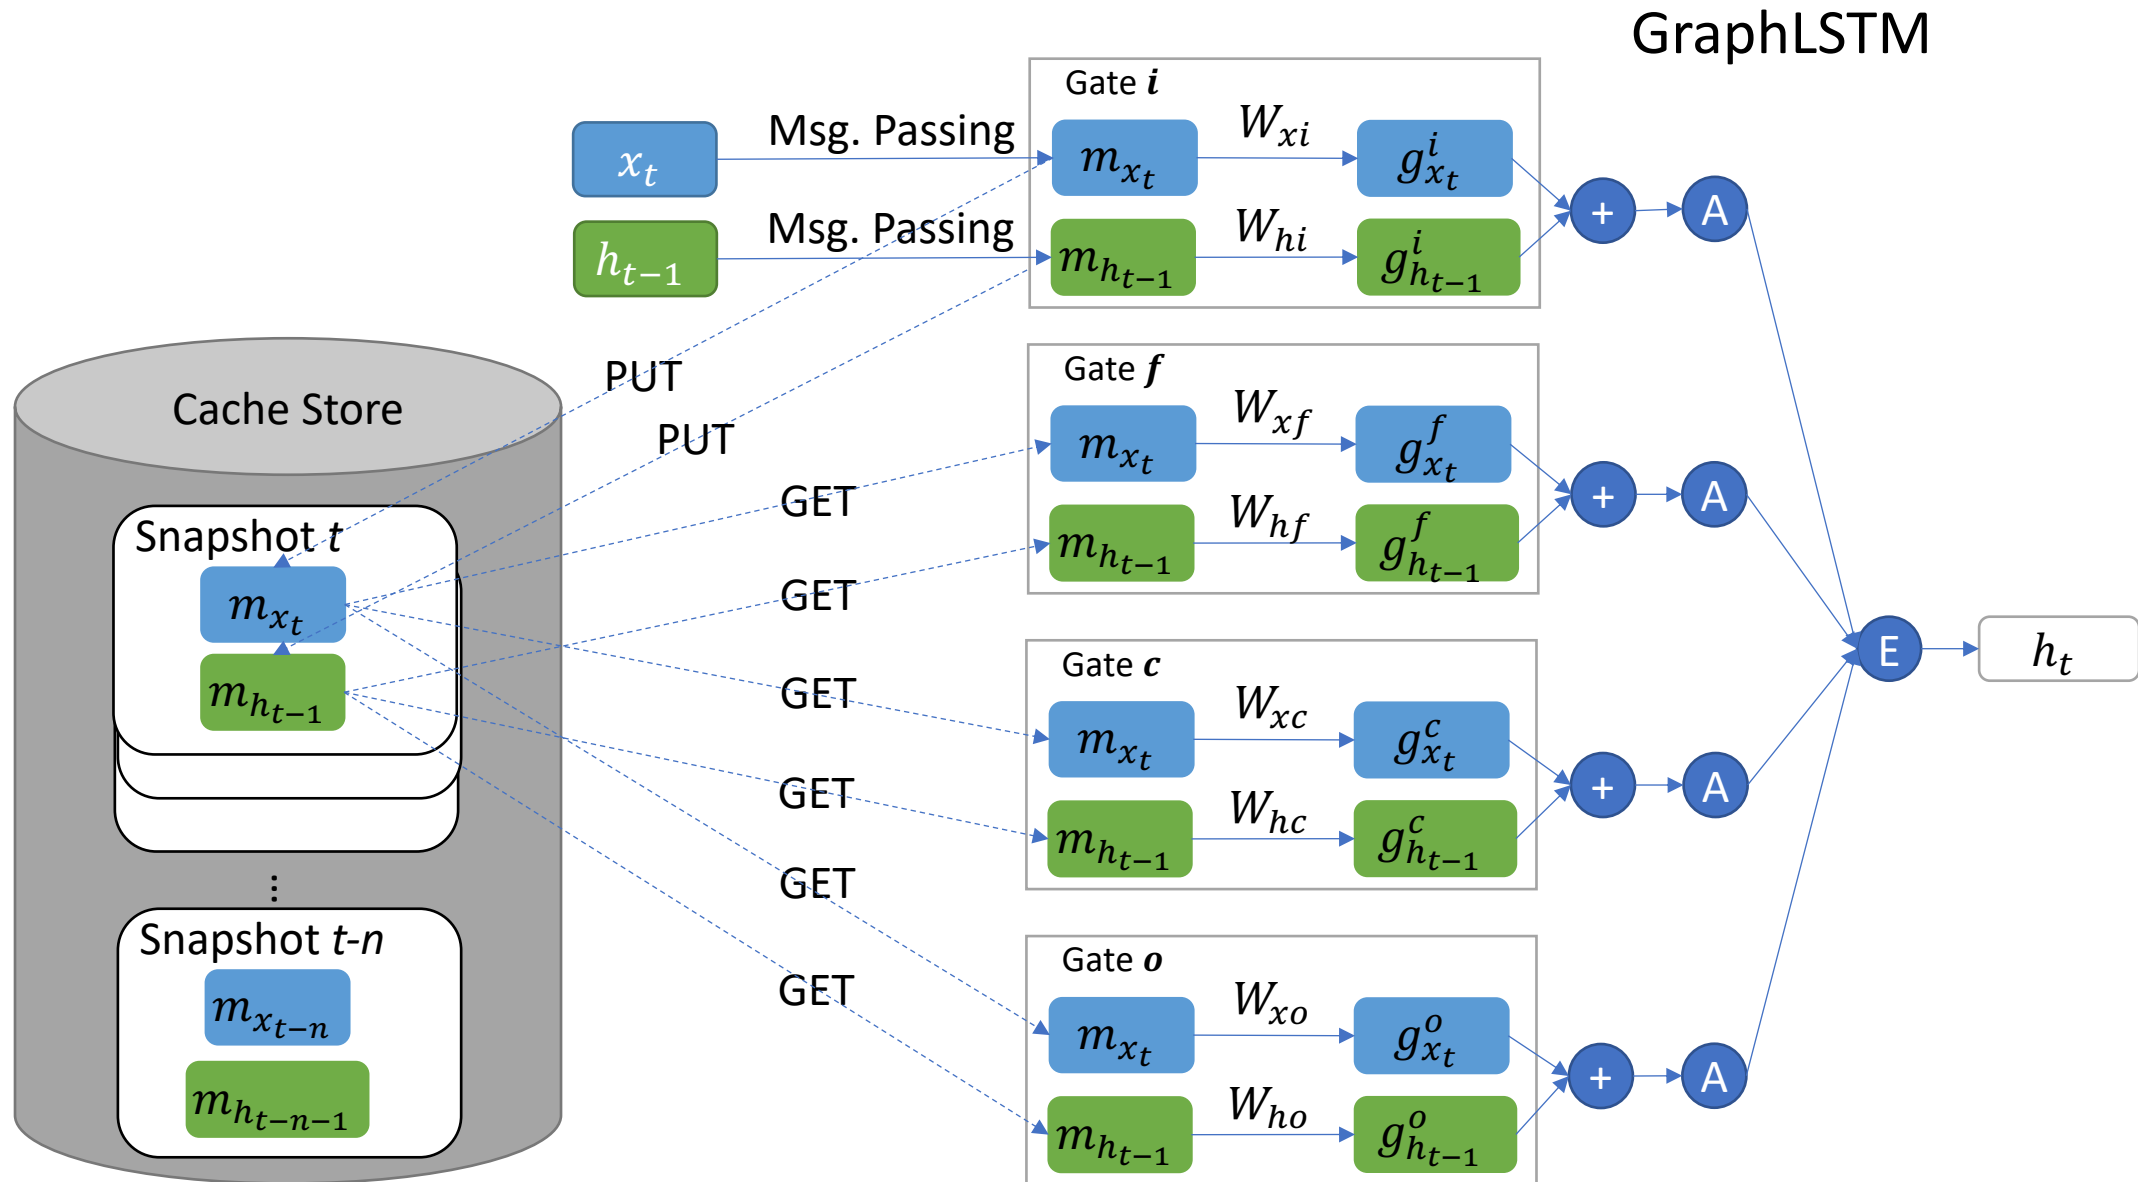
Cached Message Passing



Cached Message Passing

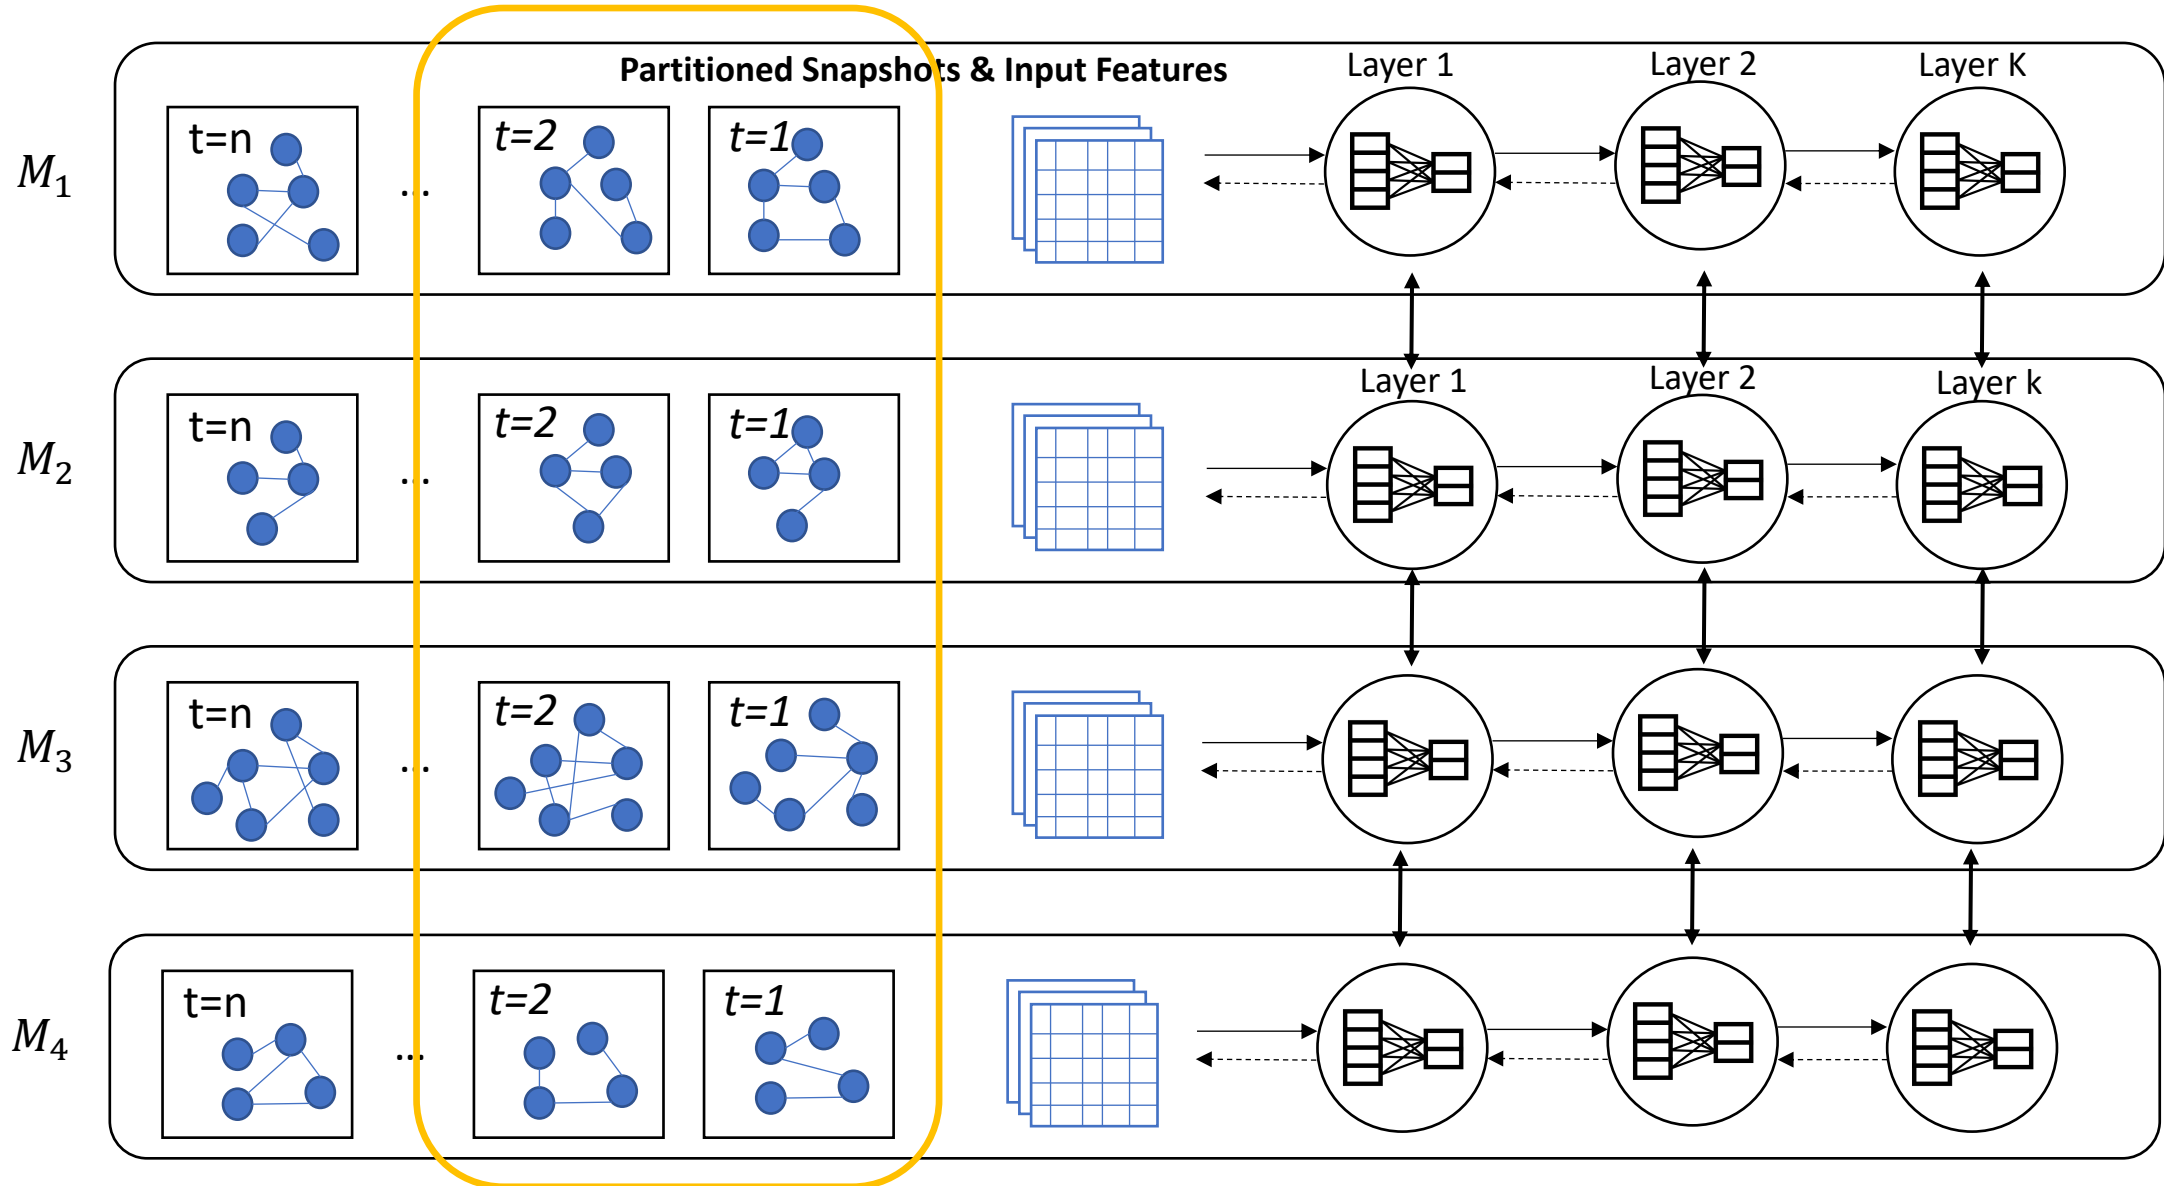


Cached Message Passing



Distributed DGNN Training

Sliding Window




DynaGraph API

- cache()** Cache caller function outputs; do nothing if already cached.
- msg_pass()** Computes intermediate message passing results.
- update()** Computes output representation from intermediate message passing results.
- integrate()** Integrates a GNN into a GraphRNN to create a dynamic GNN.
- stack_seq_model()** Stacks dynamic GNN layers to an encoder-decoder structure.


Implementation & Evaluation

- Implemented on Deep Graph Library (DGL) v0.7
- Evaluated using 8 machines, each with 2 NVIDIA Tesla V100 GPUs
 - **METR-LA**: 207 nodes/snapshots, $|F|=2$, $|S|=34K$
 - **PEMS-BAY**: 325 nodes/snapshots, $|F|=2$, $|S|=52K$
 - **METR-LA-Large**: 0.4m nodes/snapshots, $|F|=128$, $|S|=34k$
 - **PEMS-BAY-Large**: 0.7m nodes/snapshots, $|F|=128$, $|S|=52k$
- Several Dynamic GNN architectures
 - GCRN-GRU, GCRN-LSTM [ICONIP '18]
 - DCRNN [ICLR '18]


DynaGraph Single-Machine Performance



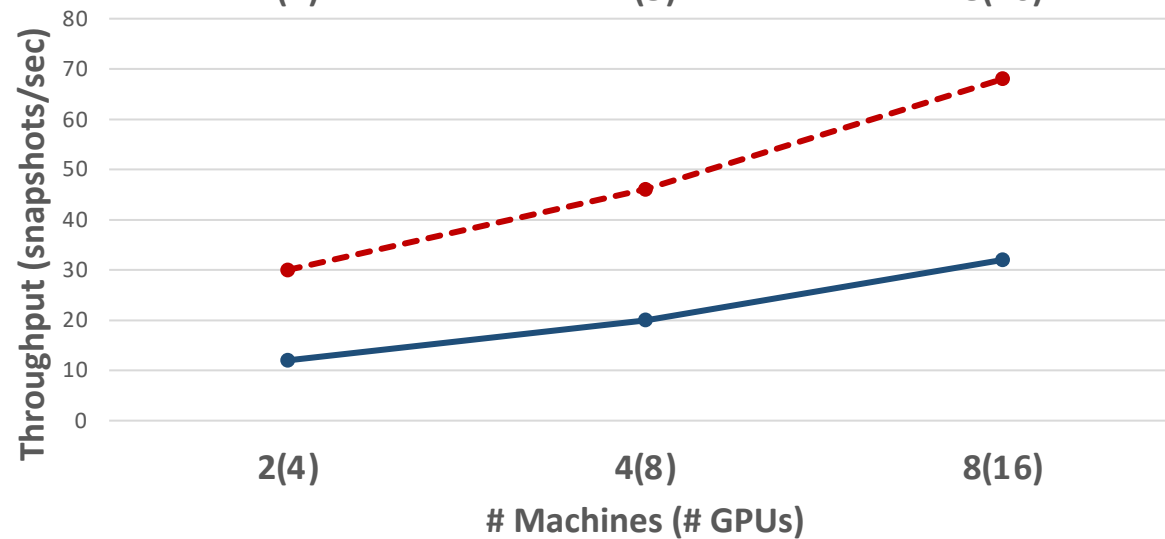
DynaGraph Distributed Performance



DynaGraph Scaling



GCRN-GRU



GCRN-LSTM

—●— DGL - - - ● - - - DynaGraph

Summary

- Supporting dynamic graphs is increasingly important for enabling many GNN applications.
 - Existing GNN systems mainly focus on static graphs and static GNNs.
 - Dynamic GNN architectures combine GNN techniques and temporal embedding techniques like RNNs.
- DynaGraph enables dynamic GNN training at scale.
 - Several techniques to reuse intermediate results.
 - Efficient distributed training.
 - Outperforms state-of-the-art solutions.

Thank you!

Contact: mingyu.guan@gatech.edu