

IGNNITION: Bridging the Gap Between Graph Neural Networks and Networking Systems

David Pujol-Perich, José Suárez-Varela, Miquel Ferriol, Shihan Xiao,
Bo Wu, Albert Cabellos-Aparicio and Pere Barlet-Ros

NOTE: Accepted for publication at IEEE Network Magazine. ©2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Abstract—Recent years have seen the vast potential of *Graph Neural Networks* (GNN) in many fields where data is structured as graphs (e.g., chemistry, recommender systems). In particular, GNNs are becoming increasingly popular in the field of networking, as graphs are intrinsically present at many levels (e.g., topology, routing). The main novelty of GNNs is their ability to generalize to other networks unseen during training, which is an essential feature for developing practical Machine Learning (ML) solutions for networking. However, implementing a functional GNN prototype is currently a cumbersome task that requires strong skills in neural network programming. This poses an important barrier to network engineers that often do not have the necessary ML expertise. In this article, we present **IGNNITION**, a novel open-source framework that enables fast prototyping of GNNs for networking systems. **IGNNITION** is based on an intuitive high-level abstraction that hides the complexity behind GNNs, while still offering great flexibility to build custom GNN architectures. To showcase the versatility and performance of this framework, we implement two state-of-the-art GNN models applied to different networking use cases. Our results show that the GNN models produced by **IGNNITION** are equivalent in terms of accuracy and performance to their native implementations in TensorFlow.

INTRODUCTION

Graph Neural Networks (GNN) [1] have recently become a hot topic among the Machine Learning (ML) community. The main novelty behind these models is their unique ability to learn and generalize over graph-structured information. This has enabled the development of groundbreaking applications in different fields where data is fundamentally represented as graphs (e.g., chemistry, physics, biology, recommender systems, social networks) [2].

The last few years have seen an increasing interest among the networking community to exploit the potential of GNN, as many fundamental networking problems involve the use of graphs (e.g., topology, routing). As such, we have already witnessed some successful GNN-based applications for networking (e.g., routing optimization [3], [4], virtual network function placement [5], resource allocation in wireless networks [6], job scheduling in data center networks [7]). In particular, GNN is the only ML technique that demonstrated the capability to effectively operate in other networks unseen




Fig. 1: Overview of IGNNITION.

during the training phase. This is thanks to the generalization power of GNNs over graphs which, in the networking field, represents a cornerstone for the development of a new generation of ML-based solutions that can be trained in a controlled network testbed (e.g., at the vendor’s lab) and, once trained, be ready for deployment in any customer network (topologies of arbitrary size and structure), without the need for re-training on premises [3].

Nowadays, designing and implementing a GNN model involves dealing with complex mathematical formulation and using ML libraries with tensor-based programming, such as TensorFlow, or PyTorch. At the same time, *GNN models* for networking often rely on non-standard GNN architectures (e.g., heterogeneous graphs with sequential dependencies between different element types [3], [4]), which makes their implementation even more difficult. This represents a critical entry barrier for network researchers and engineers that could benefit from the use of GNN, but lack the necessary ML expertise to implement these models.

In this article, we present **IGNNITION**, a TensorFlow-based framework for fast prototyping of GNNs (see Fig. 1). This framework is open source¹, and mainly targets networking experts and researchers with little background on neural network programming. With **IGNNITION**, users can easily design their own GNN models — even for the most complex network scenarios — via an intuitive, human-readable YAML file. Based on this input, the framework automatically generates an efficient TensorFlow implementation of the GNN, without the need for writing a single line of TensorFlow code.

To achieve this, we propose a high-level abstraction we call the Multi-Stage Message Passing graph (*MSMP graph*). This novel abstraction covers a broad definition of GNN, offering great flexibility to design a wide variety of existing GNN variants, and generate new custom architectures combining individual components of them (e.g., message-

David Pujol-Perich, José Suárez-Varela, Miquel Ferriol-Galmés, Albert Cabellos-Aparicio, and Pere Barlet-Ros: Barcelona Neural Networking Center, Universitat Politècnica de Catalunya.
Shihan Xiao and Bo Wu: Network Technology Lab., Huawei Technologies Co.,Ltd.

¹Available at: <https://ignnition.net>

passing functions). Likewise, MSMP graphs enable to hide the complex mathematical formulation and the tensor-wise operations behind the implementation of GNNs.

IGNNITION provides a YAML-based *codeless* interface that network engineers and practitioners can leverage to implement custom GNN models applied to networking. In this context, coding these types of models with general-purpose GNN libraries (e.g., [1], [8]–[12]) is arguably more complex for non experts in neural network programming, as they often require using lower-level abstractions, and some of these libraries even lack sufficient flexibility to implement non-standard GNN architectures applied to networking [8], [11].

More in detail, IGNNITION offers the following main features:

- **High-level abstraction:** Codeless interface — based on MSMP graphs and YAML files — that abstracts away both the mathematical formulation of the GNN and its implementation.
- **Flexibility of design:** Broad support for existing GNN variants and custom message-passing architectures via the novel MSMP graph abstraction.
- **High performance:** Efficient GNN implementations equivalent to native TensorFlow code, as performance is crucial for network applications.
- **Easy debugging:** Advanced error-checking mechanisms and enhanced visualizations of the implemented GNNs, to help network engineers troubleshoot their models.
- **Easy integration:** User-friendly dataset interface, based on the well-known NetworkX library, to feed GNN models with data from different sources and in various formats (e.g., network monitoring logs).
- **Networking focus:** IGNNITION is specially designed to facilitate the implementation of non-standard features commonly used in GNN-based networking applications (e.g., support for heterogeneous graphs with possible sequential dependencies [3], [4]).

As a result, with IGNNITION, network researchers and engineers with limited knowledge on neural network programming should be able to quickly produce a fully functional prototype of a GNN model tailored to a particular application.

BACKGROUND ON GNN

The input of a GNN is a graph, which comprises a set of nodes and edges connecting them. Each node v has an associated hidden state vector h_v of predefined size that encodes its internal state. At the beginning of the GNN execution, hidden state vectors are initialized with some node-related features included in the input graph.

After this, a message-passing algorithm is executed according to the connections of the input graph. This message-passing process comprises three main phases (see Fig. 2):

- 1) Message exchange
- 2) Aggregation
- 3) Update

First, each node sends a message to all its neighbors in the graph. This message is the result of combining the hidden states of the sender node h_v and the receiver neighbor h_u




Fig. 2: Phases of a message-passing iteration: Message exchange, Aggregation, and Update

through the so-called *Message* function (M). After this, each node aggregates all the messages collected from its neighbors into a single fixed-size vector m_v (a.k.a., the aggregated message). To do this, nodes use a common *Aggregation* function ($Aggr$), which in standard GNNs is typically an element-wise summation over all the messages collected on the node. Lastly, each node applies an *Update* function (U) that combines its own hidden state h_v with the node’s aggregated message m_v . This message passing is repeated a number of iterations T , until the nodes’ hidden states h_v converge to some fixed values.

After the message-passing phase, a *Readout* function is used to produce the output of the GNN model. This function takes as input the nodes’ hidden states and converts them into the final predictions of the model. Note that the output of the Readout function can be either a set of global graph-level properties or specific node-level features.

An essential aspect of GNN is that the internal functions that shape its architecture are modeled by neural networks (NN). Particularly, the *Message*, *Update*, and *Readout* functions are typically approximated by three independent NNs (e.g., feed-forward NN, Recurrent NN). These NNs are then dynamically assembled within the GNN architecture according to the nodes and connections of the input graphs. During training, the aforementioned NNs (e.g., the *Message* NN) optimize jointly their internal parameters, effectively learning the target function from the perspective of each individual node in the graph. As a result, after training, the GNN is able to make accurate predictions on graphs with different sizes and structures not seen before.

Note that beyond the generic GNN description provided in this section, there is a rich body of literature with different GNN architectural variants (e.g., Graph Convolutional Networks, Gated Neural Networks, Graph Attention Networks, Graph Recurrent Networks, GraphSAGE) [2], while all of them share the basic principle of an initial message-passing phase followed by a final readout. We refer the reader to [1], [2] for more generic background on GNN.

THE MSMP GRAPH ABSTRACTION

In the networking field, GNN models need to be highly customized to accommodate specific networking use cases. Particularly, it is common to find non-standard message-passing architectures [3], [4], [6] that stem from the need to explicitly model the circular dependencies between various network elements, based on the actual behavior of real network infrastructures. As a result, state-of-the-art GNN models for networking often consider heterogeneous graphs as input

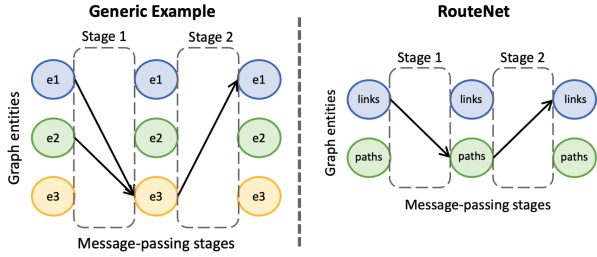


Fig. 3: MSMP graphs — Generic example and RouteNet [3].

— with different element types (e.g., forwarding devices, links, paths) —, and custom message-passing definitions divided in multiple stages [3], [13], where in each stage a specific set of elements (e.g., paths) share their hidden states with some other elements (e.g., devices).

Existing general-purpose GNN libraries offer high-level abstractions to simplify the implementation of custom GNN models. Among the most popular libraries, DGL [10] provides support for heterogeneous graphs, offering flexibility to define GNNs with differentiated message, aggregation, and update functions for each element type of input graphs, as well as combining components from well-known GNN architectures (e.g., Graph Convolutional Networks, Graph Attention Networks). Likewise, PyTorch Geometric [9] provides pre-implemented GNN layers that can be combined to create complex GNN architectures. However, the higher-level abstractions of these libraries do not provide direct support for implementing multi-stage message passing definitions, with flexibility to define message exchanges between different element types at different execution stages, which is a common architectural framework among state-of-the-art GNN models for networking [3], [13]. This eventually requires the use of lower-level abstractions to implement these types of models.

In light of the above, IGNNITION introduces a novel high-level abstraction called the *Multi-Stage Message Passing graph (MSMP graph)*, which is carefully designed to facilitate the implementation of GNN models applied to networking. With MSMP graphs, users can define their custom GNN architectures through a visual graph representation, abstracting them from the complex mathematical formulation behind their designs. Particularly, a GNN can be defined by a set of element types (hereafter, *graph entities*) and how they relate to each other in a sequential order, allowing the definition of multi-stage message passing schemes in a natural and intuitive way.

Figure 3 (left) illustrates a generic example of an MSMP graph with three graph entities (e_1 , e_2 and e_3). In this scheme, two message-passing stages can be differentiated: First, elements of type e_1 and e_2 send their hidden states to their neighbors of type e_3 , according to their connections in the input graph. Then, in a second stage e_3 elements share their states with their linked elements of type e_1 and e_2 . Likewise, Figure 3 (right) depicts the MSMP graph of RouteNet [3]. In this case, the GNN model considers two graph entities (*links* and *paths*), and implements a two-stage message-passing scheme that explicitly represent the circular dependencies between these two entities. The implementation of the *Message*, *Aggregation*, *Update*, and *Readout* functions is later defined

by the user via an intuitive model description file in YAML, as described in the next section of this article. An MSMP graph description eventually describes a message-passing iteration of the GNN. Then, IGNNITION automatically unrolls this definition a number of iterations T defined by the user.

At the same time, the MSMP graph abstraction offers a flexible modular design, which supports a wide range of existing GNN variants (e.g., Message-Passing NNs, Graph Convolutional Networks, Gated NNs, Graph Attention Networks, Graph Recurrent Networks) [2], as well as custom combinations with individual components of them (e.g., message, aggregation, update, normalization). To the best of our knowledge, this abstraction supports all the state-of-the-art GNN models applied to networking at the time of this writing.

IGNNITION ARCHITECTURE

This section provides an overview of IGNNITION. In particular, we describe below the four main building blocks that shape this framework:

- 1) Model Description Interface
- 2) Dataset Interface
- 3) Core Engine
- 4) Debugging Assistant

Model Description Interface

The first step to build a GNN with IGNNITION is to describe the model architecture using the MSMP graph abstraction, introduced in the previous section. This can be done through a codeless interface, by filling a short YAML file. More specifically, this file contains the following information:

Entities Definition: First, the user defines the different entities (i.e., element types) to consider in the networking scenario and, consequently, in the corresponding MSMP graph. In the networking context, an entity represents a set of network elements, which can be physical (e.g., routers, links), or logical (e.g., paths, virtual network functions). As an example, Fig. 4 shows the definition of an entity in the YAML model description file. There, the user indicates the entity name (e.g., *path*), the size of the hidden state vectors (e.g., 32 elements), and the method and features used to initialize these vectors. The initialization can be defined as a flexible pipeline of operations that supports elaborate methods used in state-of-the-art models (e.g., normalization, initial feature embedding). Feature names (e.g., *traffic*) are used as unique identifiers to then feed the model with any input dataset, as shown later in this section.

Message Passing: After defining the entities, the user completes the MSMP graph by describing the message-passing operations in the GNN model. These are based on the relationships between the graph entities previously defined. To do so, the user defines a single message-passing iteration, which can be divided into multiple stages. In each stage, a set of entities share their hidden states with other specific entities, thus offering flexibility to define the source and destination entities participating in each message-passing stage.

For example, in the generic MSMP graph of Fig. 3 (left), the message passing is divided in two stages. In the first one,

```

# Entity definition
entities:
- name: path
  state_dimension: 32
  initial_state:
  - type: build_state
    input: [traffic]

...

# Message-passing definition
message_passing:
  num_iterations: N
  stages:
    # Stage 1
    - stage_message_passings:
      - destination_entity: path
        source_entities:
          - name: link
            message:
              - type: direct_assignment
        aggregation:
          - type: ordered
        update:
          type: neural_network
          nn_name: recurrent1

```

Fig. 4: Example of a YAML model description file.

the following message exchanges are executed: $e1 \rightarrow e3$, and $e2 \rightarrow e3$; while the second stage performs the reverse message-passing operations: $e3 \rightarrow e1$, and $e3 \rightarrow e2$.

In the model description file, this can be defined as a series of YAML objects describing the several stages. For each stage, the user defines the names of the *source* and *destination* entities, and the *Message*, *Aggregation* and *Update* functions to be applied. IGNNITION provides great flexibility to implement these functions, using any type of NN (e.g., multilayer perceptron, recurrent, convolutional), or other operations such as element-wise sum, average, max, min, product, or even direct variable assignments, which are often used in state-of-the-art GNN models for networking [3], [13]. A full description of all the options available can be found at [14]. As an example, Fig. 4 shows the definition of *Stage 1* in the message passing of RouteNet [3] (see the corresponding MSMP graph in Fig. 3-right). Note that, in case of using NNs (e.g., in the *Update* function), the user can define a name (e.g., *recurrent_1*) that is then used to describe the NN details in a separate section of the file, as detailed later in this section.

Readout: After the message passing, the user defines the *Readout* function, producing either per-node or global graph-level outputs. In IGNNITION, the *Readout* is defined via a flexible pipeline of instructions that may combine *pooling* operations, *neural networks*, and other less common operations used in state-of-the-art GNN models applied to networking. *Pooling* operations (e.g., element-wise sum, mean) transform a variable-size set of nodes’ hidden states into a single fixed-size embedding vector, which is often used to make global graph-level predictions. Likewise, *neural network* operations take as input nodes’ hidden states — or the embeddings resulting from previous operations — and pass them through a NN. More information about the *Readout* definition can be found at [14].

Neural Networks Definition: Lastly, the user can define all the neural networks previously referenced in a separate section

(e.g., *Message*, *Update*, *Readout* NNs). For this purpose, IGNNITION supports all the native functions of the well-known Keras library [15]. Thus, to define for instance a multilayer perceptron, the user should specify in the YAML file at least the fields required by Keras, which are the number and type of layers and some basic parameters (e.g., neurons per layer). For other undefined parameters, IGNNITION uses the Keras default values, thus considerably simplifying the definition of the NN. This YAML-based interface provides flexibility to implement any NN as long as it is supported by Keras, which is widely considered as one of the most complete low-level NN programming libraries.

Dataset Interface

The *Dataset Interface* is intended to decouple the GNN model description from the input data. Similar to other GNN libraries, such as DGL [10] or PyTorch Geometric [9], IGNNITION enables to directly process data with the well-known NetworkX library. This Python library already implements a plethora of functions that automatize the definition of graphs from datasets, which enables to easily feed the GNN model with datasets from different sources and in various formats (e.g., NetFlow/IPFIX measurement reports, router configuration logs).

Core Engine

This module implements the main logic behind IGNNITION. Once the model has been designed and the datasets have been properly formatted — as explained above — the user can call different functionalities from the *Core Engine*:

- Generate a TensorFlow implementation of the GNN
- Run the training/validation of the model
- Make predictions with a trained model
- Generate visual representations of the model (for debugging purposes)

The *Core Engine* internally implements a generic definition of the MSMP graph abstraction, thus covering the broad spectrum of GNN architectures supported by such abstraction. To generate the model implementation, the *Core Engine* proceeds in a similar way to traditional compilers, first performing a lexical analysis, then a syntactical analysis, and finally a semantic analysis. These three steps enable to detect unexpected structures in the model definition (in YAML files) and trigger numerous error-checking messages to help users debug their models.

After validating all the input information, IGNNITION automatically generates the GNN model implementation directly in TensorFlow. This enables to work internally with the efficient computational graph produced by this ML library. As a result, IGNNITION implementations have comparable performance to models directly coded in TensorFlow, as shown later in the evaluation of this article.

Debugging Assistant

One of the biggest challenges when designing a GNN model with traditional ML libraries is to identify potential bugs and fix them. Typically, having a clear picture of the resulting GNN model is a cumbersome and time-consuming task. For

this reason, IGNNITION incorporates an advanced debugging system that assists the user in different ways.

First, it automatically produces an interactive visual representation of the internal GNN architecture. To this end, it relies on *Tensorboard* — a TensorFlow-based visualization toolkit — allowing the user to dynamically visualize its GNN model at different levels of granularity. Standard visualizations of TensorBoard show the resulting computational graph of the model. However, it is difficult to identify the different building blocks of the GNN from the tensor operations shown in this graph (e.g., message, aggregation, update functions, message-passing iterations). IGNNITION implements a layer on top of this computational graph (using scope variables) that differentiates the main building blocks of the GNN at different levels of hierarchy, which makes it easier to identify the internal functions. For example, after implementing RouteNet [3] with IGNNITION the user can navigate through the resulting TensorBoard graph and easily identify the two message-passing stages of this model (see Fig. 3-right): *links*→*paths*, and *paths*→*links*. Moreover, by zooming in on the visualization, the user can easily find the internal functions and how they are implemented (e.g., *Message*, *Aggregation*, *Update*).

Additionally, the debugging assistant incorporates several advanced error-checking mechanisms to ensure the correct definition of the model and to provide guidance to fix badly defined fields in YAML files (e.g., wrong NN descriptions, entities definitions, missing features in the datasets).

CASE STUDIES

GNNs have been successfully applied to a plethora of networking use cases, such as routing optimization [3], [4], virtual network function placement [5], resource allocation in wireless networks [6], job scheduling in data center networks [7], MPLS configuration, or Multipath TCP. A more comprehensive list of relevant GNN-based applications for communication networks can be found at [13].

Figure 5 illustrates the general workflow for producing with IGNNITION a GNN model adapted to a particular network problem. First, the user — e.g., a network engineer — needs to identify the main elements involved in the networking problem to be addressed (e.g., devices, links, apps, paths) and, according to the purpose of the GNN model (e.g., Traffic Engineering), define potential dependencies between these different elements. In the networking context, we can often find circular dependencies between different elements that can be naturally encoded in GNNs as a sequence of message-passing stages. These elements and dependencies shape the input graphs of the GNN, thus defining the structure of the samples in the train, validate and/or test datasets used for the model. Afterward, the user can define the architecture of the GNN. In IGNNITION, this can be done by filling the YAML model description file, leveraging the novel MSMP graph abstraction. Lastly, IGNNITION generates an efficient implementation of the model in TensorFlow that can be trained with the generated datasets.

As an example, we briefly describe below how we can implement with IGNNITION two state-of-the-art GNN models that introduce some non-standard architectural patterns.

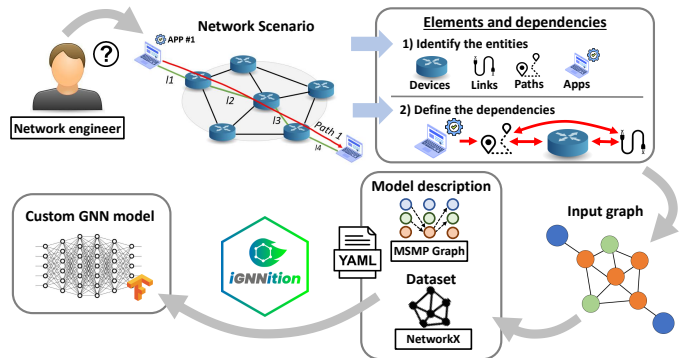


Fig. 5: General workflow to produce GNN models for networking with IGNNITION.

RouteNet

RouteNet [3] was proposed as a network modeling tool that predicts path-level performance metrics (e.g., delay, jitter) given a network state snapshot as input, which is defined by: a network topology, a routing configuration, and a traffic matrix. This GNN is then used to optimize the routing configuration in networks, by combining the model with an optimizer. To this end, RouteNet considers input graphs with two different entities (*links* and *paths*) that have circular dependencies between them, and encodes its input parameters as features within the initial hidden states of these elements. As a result, this GNN model implements a two-stage message-passing scheme combining the hidden states of these entities, which can be naturally defined by an MSMP graph (see Fig. 3-right).

Graph-Query Neural Network (GQNN)

The GQNN model [4] addresses a different problem: supervised learning of traditional routing protocols with GNN, such as shortest path or max-min routing. To this end, this GNN model uses a novel architecture with two graph entities: *routers* and *interfaces*, being the latter the several network interfaces of each router in the network. This model considers a single-stage message-passing scheme where routers and interfaces share their hidden states simultaneously. As output, it determines whether the interfaces are used to transmit traffic or not (i.e., $[0,1]$), which eventually defines the routing configuration of the network. GQNN introduces a particularity in the *Readout* definition, which uses a non-standard operation pipeline with an element-wise product and then a NN to produce the final prediction. With IGNNITION, this can be easily defined in the YAML model description file as a pipeline with a *product* operation and then a NN function.

For more information on how to implement a GNN prototype with IGNNITION, we refer the reader to [14]. There, the user can find an easy-to-follow quick-start tutorial, where a basic GNN model is built to compute the shortest path routing configuration given a network topology with weights on links. Also, IGNNITION includes a library with state-of-the-art GNN models applied to different networking use cases, which is under continuous development by the community of this open-source project.

EVALUATION: ACCURACY AND COST

As mentioned earlier, IGNNITION internally implements GNN models in TensorFlow, thus leveraging the efficient computational graphs produced by this ML library. This section aims to showcase this aspect by making a direct comparison of IGNNITION implementations with respect to implementations directly coded in TensorFlow. As an example, we consider the two state-of-the-art GNN models described in the previous section (RouteNet and GQNN).

For the evaluation, we use the implementations of these models in IGNNITION — publicly available at [14] — and, as a reference, their native implementations in TensorFlow, and reproduce some of the experiments made in their corresponding papers [3], [4]. Note that IGNNITION does not provide additional compiling optimizations. Instead, in the evaluation we reproduce the same scheme used in the original TensorFlow implementation. We use the datasets of the original evaluations of RouteNet, and GQNN (with 300,000, and 40,000 samples respectively), where 80 percent of the samples are randomly selected for training and 20 percent for evaluation.

We first evaluate the accuracy achieved by the models, to check that both implementations achieve the same results under equal training conditions. Our experimental results show that the IGNNITION implementation of RouteNet yields a Mean Relative Error of 2.62 percent, and the implementation of GQNN an accuracy of 98.07 percent. Both numbers are in line with the results obtained by the original TensorFlow implementations.

Likewise, we evaluate the execution cost of both types of implementations. Figure 6 depicts the average execution time per sample during training and inference. In line with the previous results, we can observe that the execution times of IGNNITION implementations are equivalent to those of the original models implemented in TensorFlow. Note that all these experiments were made in a controlled testbed, considering the same computing resources for all cases.

These results reveal the capability of IGNNITION to produce efficient GNN models, while offering substantial time savings for non-ML experts to implement them. Moreover, users can leverage additional functionalities such as the visual representations of the debugging system and the advanced error-checking mechanisms to easily fix potential bugs.

RELATED WORK

Several GNN libraries have emerged in recent years, motivated by the outstanding applications of GNNs to a plethora of relevant real-world problems. Among the most popular open-source libraries, we can find: DGL [10], PyTorch Geometric [9], or Graph Nets [1]. Also, NeuGraph [12] has recently become a popular GNN library that focuses on parallelizing the execution of GNN models, although it is not open source. These libraries provide high-level abstractions to code standard GNN layers (e.g., Graph Convolutional Networks, Graph Attention Networks) that can be concatenated. However, these higher-level abstractions do not provide direct support for implementing multi-stage message passing iterations, where in each stage specific graph entities share their hidden states




Fig. 6: Evaluation of the execution time of IGNNITION vs. native TensorFlow implementations (training and inference).

with some other entities, which is a common architectural pattern for modeling circular dependencies in GNN models for networking [3], [13]. As an example, DGL provides several built-in functionalities to deal with heterogeneous graphs. Nevertheless, these functionalities do not handle directly the implementation of message-passing definitions where different graph entities share their hidden states sequentially. As a result, it is eventually needed to use lower-level abstractions to implement these types of models. Lastly, some other popular libraries, such as GraphGym [8] or Spektral [11], provide higher-level abstractions compared to the previous libraries. However, these libraries do not support many non-standard GNN architectures used in networking (e.g., [3], [4]).

In contrast, IGNNITION provides a *codeless interface* that users can leverage to easily implement GNN prototypes. This interface is based on the novel MSMP graph abstraction proposed along with this framework, which is inspired by the architectural patterns commonly used in GNN models for networking. In particular, this abstraction naturally introduces the definition of non-standard message-passing architectures considering heterogeneous graphs with sequential dependencies between different sets of elements.

CONCLUSION

In this article, we have introduced IGNNITION, an open-source framework based on TensorFlow that enables fast prototyping of *Graph Neural Networks* for networking applications. IGNNITION works over a novel high-level abstraction called the *Multi-Stage Message Passing graph (MSMP graph)*, which isolates users from the complex mathematical formulation and implementation in traditional ML libraries (e.g., TensorFlow, PyTorch). MSMP graphs are flexible enough to support state-of-the-art GNN architectures with non-standard message-passing strategies. We implemented with IGNNITION two state-of-the-art GNNs for networking, to showcase the versatility of this framework. Likewise, we validated that the performance of the GNN models produced by IGNNITION is equivalent to that of native TensorFlow implementations, as efficiency is a must for networking applications.

ACKNOWLEDGMENT

This open-source project has received funding from the European Union’s Horizon 2020 research and innovation programme within the framework of the NGI-POINTER Project funded under grant agreement No. 871528. This article reflects only the author’s view; the European Commission is not responsible for any use that may be made of the information it contains. The work was also supported by the Spanish MINECO under contract TEC2017-90034-C2-1-R (ALLIANCE) and the Catalan Institution for Research and Advanced Studies (ICREA).

REFERENCES

- [1] P. W. Battaglia *et al.*, “Relational inductive biases, deep learning, and graph networks,” *arXiv preprint arXiv:1806.01261*, 2018.
- [2] J. Zhou *et al.*, “Graph neural networks: A review of methods and applications,” *AI Open*, vol. 1, pp. 57–81, 2020.
- [3] K. Rusek *et al.*, “Unveiling the potential of graph neural networks for network modeling and optimization in SDN,” in *Proceedings of ACM SOSR*, 2019, pp. 140–151.
- [4] F. Geyer and G. Carle, “Learning and generating distributed routing protocols using graph-based deep learning,” in *Proceedings of the ACM SIGCOMM Big-DAMA Workshop*, 2018, pp. 40–45.
- [5] P. T. A. Quang, Y. Hadjadj-Aoul, and A. Outtagarts, “A deep reinforcement learning approach for VNF forwarding graph embedding,” *IEEE Transactions on Network and Service Management*, vol. 16, no. 4, pp. 1318–1331, 2019.
- [6] M. Eisen and A. Ribeiro, “Optimal wireless resource allocation with random edge graph neural networks,” *IEEE Transactions on Signal Processing*, vol. 68, pp. 2977–2991, 2020.
- [7] H. Mao, M. Schwarzkopf, S. B. Venkatakrisnan, Z. Meng, and M. Alizadeh, “Learning scheduling algorithms for data processing clusters,” in *Proceedings of ACM SIGCOMM*, 2019, pp. 270–288.
- [8] J. You, Z. Ying, and J. Leskovec, “Design space for graph neural networks,” *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, 2020.
- [9] M. Fey and J. E. Lenssen, “Fast graph representation learning with pytorch geometric,” in *Proceedings of the ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- [10] M. Wang *et al.*, “Deep graph library: A graph-centric, highly-performant package for graph neural networks,” *arXiv preprint arXiv:1909.01315*, 2019.
- [11] D. Grattarola and C. Alippi, “Graph neural networks in TensorFlow and Keras with spektral,” in *Proceedings of the ICML workshop on Graph Representation Learning and Beyond*, 2020.
- [12] L. Ma, Z. Yang, Y. Miao, J. Xue, M. Wu, L. Zhou, and Y. Dai, “Neugraph: parallel deep neural network computation on large graphs,” in *USENIX Annual Technical Conference (ATC)*, 2019, pp. 443–458.
- [13] Barcelona Neural Networking center. (2021) Must-read papers on GNN for communication networks. Accessed: 04-08-2021. [Online]. Available: <https://github.com/BNN-UPC/GNNPapersCommNets>
- [14] —. (2021) IGNNITION-Documentation. Accessed: 15-04-2021. [Online]. Available: <https://ignnition.net/doc/>
- [15] F. Chollet *et al.* (2015) Keras. Accessed: 07-04-2021. [Online]. Available: <https://github.com/fchollet/keras>

BIOGRAPHY

David Pujol-Perich is an MSc student in advanced computing from the Universitat Politècnica de Catalunya (2020-2022), where he also received his BSc in Computer Science (2020). In 2019-2020, he did an academic exchange in ETH Zurich. He is currently a machine learning researcher at the Barcelona Neural Networking Center (BNN-UPC).

José Suárez-Varela received his PhD in Computer Science from the Universitat Politècnica de Catalunya (2020). He is currently a postdoctoral researcher at the Barcelona Neural Networking Center (BNN-UPC), and co-Principal Investigator of the EU-funded project IGNNITION (NGI Pointer). His main research interests are in Machine Learning applied to networking.

Miquel Ferriol-Galmés received his BSc in Computer science (2018) and MSc in Data Science (2020) from the Universitat Politècnica de Catalunya. He is currently pursuing a PhD at the Barcelona Neural Networking center (BNN-UPC). His main research interests are in the application of Graph Neural Networks to computer networks.

Shihan Xiao received a B.Eng. degree in electronic and information engineering from the Beijing University of Posts and Telecommunications, in 2012, and a PhD degree from the Department of Computer Science and Technology, Tsinghua University. He is currently a senior engineer with Huawei 2012 NetLab.

Bo Wu received his Bachelor’s degree from the School of Software at Shandong University, China, in 2014, and his PhD degree from the Department of Computer Science and Technology at Tsinghua University in 2019. Currently, he works in the Network Technology Laboratory at Huawei Technologies.

Albert Cabellos-Aparicio is an assistant professor at Universitat Politècnica de Catalunya (UPC), where he obtained his PhD in computer science in 2008. He is director of the Barcelona Neural Networking Center (BNN-UPC) and scientific director of the NaNoNetworking Center in Catalunya.

Pere Barlet-Ros is an associate professor at Universitat Politècnica de Catalunya (UPC) and scientific director of the Barcelona Neural Networking Center (BNN-UPC). From 2013 to 2018, he was co-founder and chairman of the machine learning startup Talaia Networks, which was acquired by Auvik Networks in 2018.