Demonstration of Federated Learning over Edge-Computing Enabled Metro Optical Networks

B. Shariati*, P. Safari*, A. Mitrovska, N. Hashemi, J. K. Fischer, and R. Freund

Fraunhofer Institute for Telecommunications Heinrich Hertz Institute, Einsteinufer 37, 10587 Berlin, Germany, behnam.shariati@hhi.fraunhofer.de

Abstract: *We demonstrate the benefits of a federated learning framework for (re)training of global ML models over geo-distributed data sources. The demonstration is carried out on a live edge computing enabled optical networking test-bed.*

Introduction

Machine Learning (ML) has recently received a significant attention from every business, which somehow deals with data. One of the very consequences is the massive amount of data that is collected and stored in numerous sites from which the businesses would like to explore insights or train sophisticated ML solutions to improve their product portfolios. In the cloud computing centred paradigm, the main approach has been to move all data to a central location where there is sufficient computing and storage resources to perform the ML model training. However, due to the risk of unauthorized use, which may compromise the privacy of potential customers or put businesses in danger due to the disclosure of confidential information, there is a strong tendency to avoid transporting the data across the telecom infrastructure. Instead, solutions are devised to perform any data exploration or ML training in a distributed fashion at the locations where the data originates. This approach not only addresses the previously mentioned concerns, it can significantly reduce the required bandwidth for data transport over the telecom infrastructure between data sources and cloud computing infrastructures.

Privacy-Preserving Machine Learning (PPML) allows the training of a ML model over privacy-sensitive data by assuring the data owners that their privacy will not be compromised[1]. They have been studied in different applications such as Google Keyboard[2], Apple's QuickType[3], medical screening[4], and disease outbreak discovery[5].

Moreover, many modern ML algorithms are data hungry. These data are usually not gathered in a single location and the mobility of the data is constrained due to privacy concerns, network bandwidth limitations, and resource availability[6]. These issues result in the development of artificial intelligence (AI) on the network edge towards the use of locally hosted data in order to perform model training and contribute to a better global model with partially improved models trained on different remote local sites[7]. These developments together with the use of recently proposed federated aggregation[8] mechanisms can be exploited to use huge amounts of data generated on different devices, which do not necessarily belong to a single owner and are managed by different entities for the realization of ML solutions beneficial for all the consumers[9][10][11][12].

We have developed a federated learning framework, which allows training of ML models over geo-distributed datasets. The framework is developed over PyTorch[13], but can be used to train and validate models based on the TensorFlow[14] library as well. The framework operates in a modular fashion, which allows different downstream tasks (e.g., image recognition, QoT estimation, etc.) to be plugged-in to the framework. The framework is interfaced with a customized performance monitoring dashboard based on Grafana[15], which provides real-time monitoring of traffic flows among different modules of the framework and of the available resources on the distributed sites.

In this demonstration, we perform real-time training of a QoT classifier by exploiting data of three different Domain Managers (DM), representing a multi-vendor ecosystem, without sharing any data with the Network Management System (NMS) in order to avoid transporting any data to a central location and to protect the privacy of different vendors while offering their knowledge to train a global ML model.

Benefits of the Federated Learning Framework for Optical Networks

Optical networks are transforming towards fully autonomous entities and ML is expected to play a significant role[16]. In this transformation phase, there are many technical and regulatory issues to be addressed paving the way for the realization of such solutions and, ultimately, their commercialization and deployment.

Currently, unavailability of real field-collected

[*] These authors have equally contributed to this work.

data is one of the main showstoppers in the development of reliable ML-based solutions. While one of the reasons is the immaturity of monitoring solutions to acquire and process the networking data, a more critical bottleneck emerges due to regulatory issues concerning data sharing among different players (i.e., telecom operators and vendors). In essence, the telecom operators have control over the data, generated from the activity of their networking infrastructure, and usually tend not to share it with third parties due to conflicts of interest and confidentiality of the operation of their infrastructure. However, such data is crucial to develop reliable ML-based solutions. Therefore, it is of great value to develop a solution that enables exploitation of such data, for the purpose of training and validation of ML-based algorithms, without sharing the original data with others.

Our federated learning framework targets exactly that challenge and allows *shared ownership and governance of ML models* in optical networks, which is the key enabler for the realization of ML-based solutions that can work in real-field scenarios in a robust and reliable way.

Federated Learning Framework Architecture

The Federated Learning (FL) framework[17][18] trains a global model using data hosted on a set of geo-distributed nodes. These nodes are assumed to have computing resources to contribute to the training. We use the term *Edge Contributor Node (ECN)* to refer to those edge nodes. FL mainly relies on the concept of “bring code to the data” rather than “data to the code.” FL i) reduces the amount of transported data significantly, ii) accounts for model inaccuracies by adapting the ML models using local data, and iii) relaxes the constraint of having high-performance computing units in a centralized location. The overall training procedure is orchestrated by the *Training Coordinator Node (TCN)*. In order for the TCN and ECNs to communicate, a secure communication protocol based on WebSocketSecure (WSS) is adopted[19].

In order to realize a FL architecture, we use the well-known Stochastic Gradient Descent (SGD) algorithm[7][8]. In our implementation of so called Federated SGD, each ECN *k* computes the average gradient on its local data at the current model ω_t^k, and the TCN is responsible for gradient aggregation via $\omega_t \leftarrow \sum_{k=1}^{K}[(\frac{n_k}{n}) \times \omega_t^k]$ and then updates the glogal model parameters. This mechanism is called Federated Averaging.

Performance Monitoring Dashboard

We have developed a performance monitoring dashboard based on Grafana[15] and interfaced it to our FL framework. The dashboard provides real-time monitoring of the usage of computational resources (i.e., CPU, Memory, GPU, and filesystem) for each one of the nodes during the training procedure. In addition, the monitoring dashboard allows real-time monitoring of traffic flows among the nodes during the training. The developed dashboard relies on Prometheus as data source and cAdvisor as metrics exporter for providing real-time statistics reporting.

Demo Architecture

The networking architecture envisioned for this demonstration is illustrated in Fig. 1. The metro optical network is composed of three commercial 2-degree reconfigurable optical add-drop

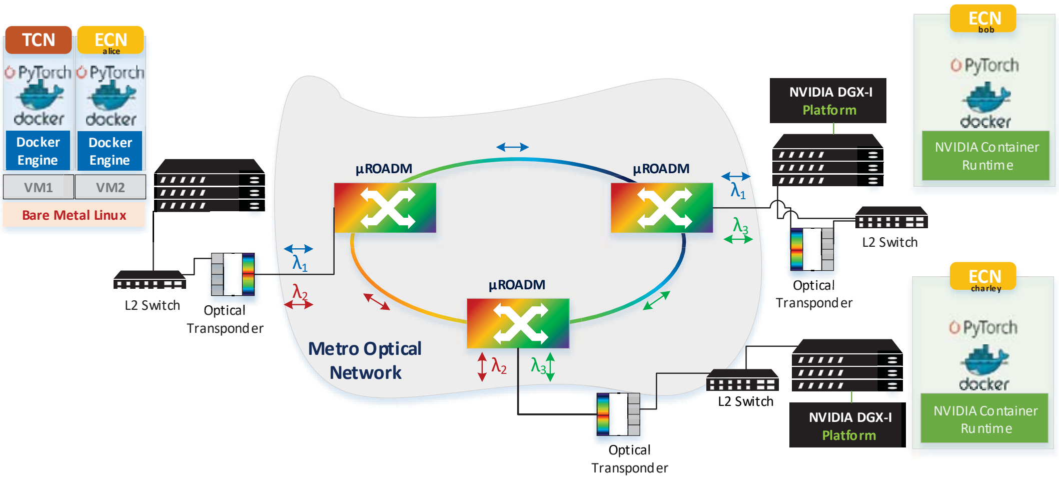

Fig. 1: The 3-node microROADM metro network test-bed used for the demonstration. Each metro node hosts a computing platform, two of them being NVIDIA DGX-1 AI acceleration units. The live traffic exchanges among the three node over three 100G lightpaths shown in blue, red, and green. Each computing node hosts the TED of a particular vendor to which the ECNs have access. The TCN, which moderates the whole training, is hosted on one VM of one of the edge nodes.

multiplexers (ROADM)[20] connected to each other in a ring topology. We consider 20 km of standard single mode fibre (SSMF) to interconnect each two ROADMs. Each ROADM is connected to a commercial optical transponder card[21], which provides two wavelength channels. We use each transponder to establish two distinct 100G connections, as shown with different colours in the figure. On the client side, we have two 100G interfaces connected to a L2 switch and, from there, to the edge compute nodes. Two of the compute nodes are NVIDIA DGX-1 platforms[22], which host 8 Tesla V100 Graphic Processing Units (GPU) and are used for acceleration of model training.

In terms of software stack, both DGX-1 platforms run a Linux 18.04.2 LTS over which the NVIDIA Container Runtime is running to execute Docker containers. On the GPU-less edge node, we run a bare metal Linux 18.04 LTS over which we host two virtual machines (VM): one dedicated to run an ECN and one to run the TCN over a Docker engine. The framework runs in a containerized fashion as a set of micro-services, which are deployed on each dedicated machine.

Demonstration Workflow

In order to show the benefits of our proposed federated learning framework, we showcase a use-case of Quality of Transmission (QoT) estimation in the context of network data sharing based on mutual trust for network automation. In this regard, we consider that each edge node hosts the Domain Manager (DM) (see Fig.2b) of a corresponding vendor with a dedicated Traffic Engineering Database (TED) of the domain.The demonstration will showcase a distributed version of training a QoT classifier using datasets that are hosted on three different machines.

The demonstration is initiated by a supervisor. In the first step, the Docker containers, which include the framework and all its dependencies, will be deployed on the machines. Once the images are deployed, the micro-services start. At this stage, the TCN performs handshake with all the ECNs to establish the WSS connection. When the connection is established, the TCN will query the eligibility of the ECNs. The ECNs will respond whether they have an available dataset, computation power, and willingness to contribute to the training. Once the TCN receives all the responses from the ECNs, it distributes the training configuration to all the ECNs and the training on the ECNs starts. When the local training is completed, each ECN returns the updated model parameters and the TCN proceeds to perform Federated Averaging to obtain the global model. The training will continue until the global model achieves a certain accuracy, which is measured during the validation procedure. Otherwise, it stops when the maximum number of training rounds set by the TCN is reached. The performance monitoring dashboard reports all the interactions and resource usage during the training procedure.

Conclusions

This demonstration shows the possibility of training a global ML model over a set of datasets that are located on different machines, interfaced with GPU-enabled AI acceleration units, and interconnected with three nodes in a metro optical networking test-bed that carries live traffic.

Acknowledgment

This work has been partially developed in the project OTB-5G+ funded by the German ministry of education and research (BMBF) (reference number: 16KIS0979K)

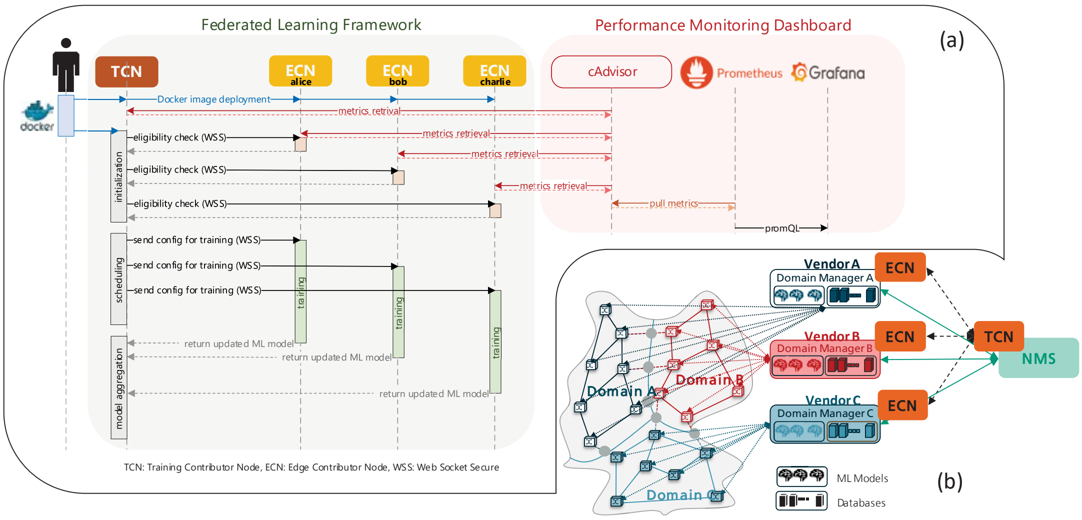

Fig. 2: (a) Training workflow of the federated learning framework considering three ECNs. (b) Illustration of a multi-vendor ecosystem considered for use-case demonstration over the framework.

References

[1] K. Bonawitz et al., “Practical Secure Aggregation for Privacy-Preserving Machine Learning”, in Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 2017.

[2] T. Yang et al. “Applied federated learning: Improving google keyboard query suggestions,” arXiv preprint arXiv:1812.02903, 2018.

[3] ADP Team et al. “Learning with privacy at scale,” Apple Machine Learning Journal, 1(8), 2017.

[4] J. Paparrizos, et al., “Screening for pancreatic adenocarcinoma using signals from web search logs: Feasibility study and results,” Journal of Oncology Practice, 12(8):737–744, 2016.

[5] V. Lampos, Andrew C Miller, Steve Crossan, and Christian Stefansen. Advances in nowcasting influenza-like illness rates using search query logs. Scientific reports, 5:12760, 2015.

[6] H. Wang, et al., “Federated learning with matched averaging,” In International Conference on Learning Representations, 2020.

[7] K. Bonawitz et al, “Towards Federated Learning at Scale: System Design”, in arXiv preprint arXiv:1902.01046.

[8] H. B. McMahan, et al., “Communication-Efficient Learning of Deep Networks from Decentralized Data”, in Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), 2017.”

[9] T. Li, et al., “Federated learning: Challenges, methods, and future directions,” arXiv preprint arXiv:1908.07873, 2019. 9.

[10] V. Smith, et al., “Federated multi-task learning,” In Advances in Neural Information Processing Systems, pp. 4424–4434, 2017.

[11] S. Caldas, et al., “A benchmark for federated settings,” arXiv preprint arXiv:1812.01097, 2018.

[12] P. Kairouz, et al., “Advances and open problems in federated learning,” arXiv preprint arXiv:1912.04977, 2019.

[13] https://pytorch.org/ [accessed in Aug 2020]

[14] https://www.tensorflow.org/ [accessed in Aug 2020]

[15] https://grafana.com/ [accessed in Aug 2020]

[16] L. Velasco, B. Shariati, F. Boitier, P. Layec, and M. Ruiz, “Learning life cycle to speed up autonomic optical transmission and networking adoption,” IEEE/OSA JOCN, vol. 11, pp. 226-237, 2019.

[17] B. Shariati, P. Safari, and J. K. Fischer, “Applications of distributed learning for optical communication networks,” presented at OSA APC, Montreal, Canada, Jul 2020.

[18] P. Safari, B. Shariati, and J. K. Fischer, “Privacy-preserving distributed learning framework for 6G telecom ecosystems,” arXiv:2008.07225, Aug 2020.

[19] [RFC 6455] The WebSocket Protocol.

[20] ADVA FSP3000 Open Line System with microROADMs (https://www.adva.com/) [accessed in Aug 2020]

[21] ADVA Quadflex Transponder (https://www.adva.com/) [accessed in Aug 2020]

[22] NVIDIA DGX-1 Deep Learning Platform (https://www.nvidia.com/en-us/data-center/dgx-1/) [accessed in Aug 2020]