



HAL
open science

Deploying Heterogeneity-aware Deep Learning Workloads on the Computing Continuum

Thomas Bouvier, Alexandru Costan, Gabriel Antoniu

► **To cite this version:**

Thomas Bouvier, Alexandru Costan, Gabriel Antoniu. Deploying Heterogeneity-aware Deep Learning Workloads on the Computing Continuum. BDA 2021 - 37e Conférence sur la Gestion de Données - Principes, Technologies et Applications, Oct 2021, Paris, France. hal-03338520v1

HAL Id: hal-03338520

<https://hal.science/hal-03338520v1>

Submitted on 8 Sep 2021 (v1), last revised 22 Oct 2021 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Deploying Heterogeneity-aware Deep Learning Workloads on the Computing Continuum

PhD student: Thomas Bouvier **Advisors:** Alexandru Costan, Gabriel Antoniu
Univ Rennes, INSA Rennes, CNRS, Inria, IRISA — Rennes, France
{thomas.bouvier, alexandru.costan, gabriel.antoniu}@irisa.fr

Abstract—The increasing need for real-time analytics motivated the emergence of new incremental methods to learn representations from continuous flows of data, especially in the context of the Internet of Things. This trend led to the evolution of centralized computing infrastructures towards interconnected processing units spanning from edge devices to cloud data centers. This new paradigm is referred to as the Computing or Edge-to-Cloud Continuum. However, the network and compute heterogeneity across and within clusters may negatively impact Deep Learning (DL) training. We introduce a roadmap for understanding the end-to-end performance of DL workloads in such heterogeneous settings. The goal is to identify key parameters leading to stragglers and devise novel intra- and inter-cluster strategies to address them. We will explore various policies aiming to improve makespan, cost and fairness objectives while ensuring system scalability.

Index Terms—Deep Learning, incremental learning, heterogeneous systems, Computing Continuum.

I. CONTEXT

State-of-the-art **deep learning** (DL) models outperform human experts' capacity in many domains, including image classification, machine translation or gaming. With the recent rise of the Internet of Things (IoT), input data is generated at an increasingly rapid pace by sensors all over the globe. This trend motivated the emergence of the **Computing Continuum** as a means to distribute computation from centralized clouds towards multi-tier processing units (*i.e.* mini-clusters in the fog) or edge devices themselves. Such infrastructure aims to optimize the performance of geo-distributed applications, benefiting from data locality to decrease communication costs and latency. However, **network and compute heterogeneity** should be carefully considered to efficiently leverage this continuum. Besides, deep neural networks (DNNs) need to shift from one-time training (typically used for static data) to approaches capable of learning from incoming flows of data. **Incremental learning** is such an emerging DL method where progressively available data is used to extend the model's knowledge.

II. CHALLENGES

Distributing DL training across fog mini-clusters and cloud data centers poses several challenges:

Network heterogeneity. Distributed DNNs suffer from low-speed and high-latency wide-area networks (WANs) connecting distant nodes. These heterogeneous links lead to stragglers, making it challenging for DL algorithms to perform efficient training.

Compute heterogeneity. As Moore's law came to an end, GPUs and other specialized accelerators emerged alongside traditional CPUs to train DNNs. In practice, clusters provide a different number of nodes equipped with different hardware or virtualized resources, leading to heterogeneous performance across and within clusters.

Scale. By nature, edge-to-cloud computing is geo-distributed. At scale, this characteristic aggravates network heterogeneity, leading to communication bottlenecks on multiple low-bandwidth nodes.

III. PROBLEM STATEMENT

Conciliating both network and compute heterogeneity to distribute the DL training efficiently in large-scale scenarios is therefore a problem gaining more and more interest from both the DL and the HPC – Big Data communities. Strategies at system and algorithmic levels should be explored as users might want to optimize different objectives based on their high-level goals, such as **minimizing the time to complete the job (makespan), costs, or improving fairness.**

IV. PHD OBJECTIVES

During my PhD (started in 2021), I will study the following research questions: 1) *How much can one improve (or degrade) the efficiency of DL training by performing it in the fog (closer to the edge) rather than performing it in the cloud?* and 2) *How to account for the heterogeneous network and compute capabilities of the processing units across the Computing Continuum?*

In order to answer these questions, I plan to devise a set of **strategies and algorithms optimizing DL workloads** in heterogeneous settings. More specifically, in order to assess their impact, I plan to study how the proposed techniques affect the **accuracy** in classifying new data and to what extent they apply to **incremental learning.**

V. STATE OF THE ART

Parallel stochastic gradient descent (PSGD) and its variants have been widely used to train large DNNs. Current approaches allow training model subsets on their local data asynchronously, and exchange gradients via peer-to-peer communication to synchronize local subsets. Techniques like bounded staleness, backup workers and synchronization queues [1] have been devised to account for dynamic network heterogeneity, whereas the selection of the fastest links [2], [3], prevent the emergence of deterministic stragglers. Skipping SGD iterations and optimizing the partitioning of input data [4] help with deterministic compute heterogeneity. At the cluster level, heterogeneity-aware schedulers [5] allow sustaining heavier workloads than agnostic ones.

However, these approaches do not comprehensively consider the challenges posed by the continuum. In particular, they enable DL deployments on distributed platforms by addressing the issues of network and compute heterogeneity separately.

VI. CONTRIBUTIONS

A first step towards executing efficient DL training in heterogeneous settings is to get a **holistic understanding of performance**. For this purpose, we will leverage E2Clab [6], a framework supporting the complete experimental cycle across the continuum, including configuration, deployment and execution in a reproducible way. We will extend it to **identify inefficient gradients computation and propagation** across nodes.

The next step is to devise novel intra- and inter- cluster strategies to address the above observations. In particular, the following facets should be considered:

- **Coarse-grained compute capacity.** The compute capacity of clusters should be estimated to adjust the initial data parallelism batch size accordingly. Input data migration could be performed at runtime.
- **Fine-grained compute capacity.** DNN training proceeds in iterations. Within a cluster, heterogeneity-aware scheduling strategies should improve utilization and optimize fairness by time-multiplexing jobs over CPUs and GPUs.
- **Deterministic network capacity.** Fast links should be favored to exchange data. Limiting communication to certain links requires ensuring the convergence of the SGD algorithm across nodes.
- **Dynamic network capacity.** Temporary network slowdowns should be detected in real time and countered by decentralized asynchronous SGD methods.

The main contribution of my PhD will be **the design of a prototype framework helping to distribute DL training** (Figure 1), accounting for deterministic and dynamic of both network and compute heterogeneity.

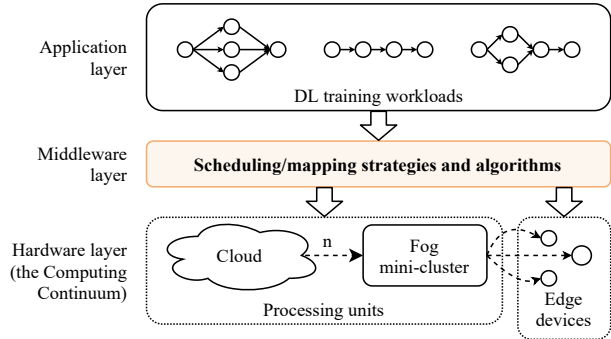


Fig. 1. Processing DL workloads on the Computing Continuum.

VII. METHODOLOGY & ROADMAP

During the first year, I plan to conduct experiments on Grid’5000, a large-scale testbed for distributed computing. The methodology implemented in E2Clab will help reproduce both relevant behaviors of the given DL workloads and representative settings of the physical infrastructure underlying the continuum. The goal is to identify the main factors leading to stragglers slowing down the whole training process. The architectural differences of DNNs should be compared.

The outcomes will drive the design of a framework mitigating the identified infrastructure bottlenecks in DL scenarios. By explicitly considering heterogeneity, various policies should improve the following objectives:

- **Makespan** refers to completing the training process as soon as possible.
- **Cost** is relevant when using elastic resources from public clouds. It specifically aims to minimize communication and resource usage costs.
- **Fairness** aims to optimize the fair sharing of processing units across the continuum.

Eventually, it will be interesting to study the extent to which these strategies apply to incremental learning. The makespan should be largely improved by deferring the global model synchronization across more distant nodes.

REFERENCES

- [1] Q. Luo *et al.*, “Hop: Heterogeneity-aware decentralized training,” in *24th Intl. Conf. on Architectural Support for Programming Languages and OS*, ser. ASPLOS 19, 2019, p. 893–907.
- [2] K. Hsieh *et al.*, “Gaia: Geo-distributed machine learning approaching LAN speeds,” in *14th USENIX NSDI 17*, 2017, pp. 629–647.
- [3] Q. Luo *et al.*, “Prague: High-performance heterogeneity-aware asynchronous decentralized training,” in *25th Intl. Conf. on Arch. Support for Prog. Lang and OS*, ASPLOS 20, 2020, p. 401–416.
- [4] R. Hong *et al.*, “Dlion: Decentralized distributed deep learning in micro-clouds,” in *11th USENIX HotCloud 19*, Renton, WA, 2019.
- [5] D. Narayanan *et al.*, “Heterogeneity-aware cluster scheduling policies for deep learning workloads,” in *14th USENIX OSDI 20*, 2020, pp. 481–498.
- [6] D. Rosendo *et al.*, “E2clab: Exploring the computing continuum through repeatable, replicable and reproducible edge-to-cloud experiments,” in *2020 IEEE CLUSTER*, 2020, pp. 176–186.