Full-time postdoctoral researcher (2 years)

Title : Energy-Aware Distributed AI on Heterogeneous Clusters

Key-words: Distributed Systems, Artificial Intelligence, Sustainable AI, Energy Optimization, Tool

Location: Laboratoire LISTIC, Polytech Annecy-Chambéry, 5 chemin de Bellevue – Annecy-le-Vieux – 74000 Annecy

Context and scientific description:

In 2020, Information and Communication Technologies (ICT) already accounted for 4.4% of global electricity consumption (Malmodin et al, 2024), a share expected to increase especially with the emergence of Artificial Intelligence (AI) tools, such as large language models. Indeed, ADEME (the French Agency for Ecological Transition) anticipates a tripling of greenhouse gas emissions and an 80% increase in electricity consumption for the sector by 2050. Training these models requires huge computational resources, and therefore energy resources, resulting in substantial greenhouse gas emissions (Patterson et al., 2021) and significant water consumption for server cooling (Li et al., 2021)
This increase in energy demand is driven by the rapid expansion of computational requirements: since 2012, the computing power needed to train state-of-the-art AI models has doubled approximately every three to four months (Amodei et al., 2018). This trend has been made possible by the using of parallel computing units, such as GPUs, which significantly accelerate machine-learning processes. GPU-equipped clusters coupled with traditional computing unit allow computations to be parallelized across multiple machines, but they introduce distributed systems issues, including task scheduling, fault tolerance, and workload allocation. Despite their importance, the intersection between energy efficiency, environmental footprint, AI, and distributed systems remains largely underexplored (Orgerie et al., 2014), as most research continues to prioritize performance and scalability over environmental considerations (Guo et al, 2024; Joost Verbraeken et al., 2021).

Recent findings from the ADEME–Arcep report (ADEME, 2025) highlight the huge carbon footprint of ICT sector, with 50% of emissions attributable to terminals, and 46% to data centers. Most device-related emissions come from manufacturing, distribution, and end-of-life processes rather than from operational use. These insights challenge the current paradigm of homogeneous clusters, which rely on identical machines connected via uniform high-speed links. For example, GPT-3 was trained using 10,000 V100 GPUs (Patterson, 2021), consuming energy comparable to that of a European nuclear reactor operating for one hour.

We believe that current distributed AI frameworks are not adapted to heterogeneous computing environments. Algorithms designed for homogeneous platforms treat slower machines as bottlenecks, generating idle time, inefficiencies, and unnecessary energy consumption. This design leads to premature replacement of slower hardware, necessitates investment in large sets of identical machines, and limits the reuse of existing resources. Consequently, the conventional approach to AI training worsens the environmental impact of large-scale computational infrastructures.

The aim of this postdoc is to evaluate, compare, and implement new solutions to reduce the environmental impact and the energy consumption of distributed training in a heterogeneous computing environment.

An open-source tool, DAHL, is currently under development by the PhD student Andrew Mary Huet de Barochez. DAHL is a machine learning framework designed to minimize the energy footprint of AI training on a single machine and to operate efficiently in heterogeneous computing environments. It targets to train a Convolutional Neural Network (CNN) on several datasets.

By leveraging the diversity of hardware resources, the framework enables flexible execution of operations across different computing units and supports energy-aware scheduling and load balancing, contributing to more efficient AI training. Unlike existing frameworks such as PyTorch, TensorFlow, or Candle, which we considered too restrictive for implementing heterogeneity management, DAHL is built on StarPU (Augonnet et al., 2008), a high-performance computing tool providing multiple task schedulers, including energy-aware options. This architecture allows full control over task granularity, distribution across CPUs and GPUs, adaptation of parallelization strategies, and low-level optimizations. Instead of parallelizing entire CNN layers as usual, our approach operates at finer granularity, offering greater flexibility and efficiency.

The main scientific challenge of this postdoc is to extend DAHL **from a single-machine to a multi-machine framework**, heterogeneous platform while optimizing energy efficiency. Then, based on this new version of DAHL, the second scientific challenge lies in exploring different evaluations of the platform and compare with the state of the art the impact of various factors: for example, network latency, network topology, platform carbon footprint, dynamic machine on/off strategies, task postponement, and system resilience/fault tolerance, on overall performance and energy consumption.

Planned experiments will assess how platform support heterogeneity to adapt to a variety of equipment and network topologies (Won et al., 2024). We will also study how the sizing and workload placement affect energy use (Freeh, 2005), (De Langen, 2006), including dynamic strategies such as follow-the-sun or follow-the-wind to exploit greener energy locations (Figuerola, 2009). Additional evaluations will explore the trade-offs between task postponement, result degradation, completion time, and energy consumption (Madon, 2024), providing a comprehensive understanding of the system's efficiency under heterogeneous conditions.

The provisional postdoc schedule starts with the extension of DAHL to multiple machines, followed by systematic benchmarking against the state-of-the-art. Initial experiments will leverage existing infrastructure, including GPUs and CPUs already available within the LISTIC laboratory, ensuring immediate access to data and hardware at the start of the postdoctoral fellowship.

Subsequent phases will focus on the development and testing of energy-aware scheduling and load-balancing algorithms across heterogeneous platforms. This work addresses a largely unexplored area at the intersection of distributed AI, energy efficiency, and heterogeneous computing.

The postdoctoral position will include the following tasks:

- **Months 1–6:** Learn the state of the art in the field, with a particular focus on understanding the DAHL tool. Assist the PhD student with the simulation work.
- **Months 7–15:** Propose a new prototype that generalizes the PhD student's solution. The current solution involves only a single machine; the postdoctoral researcher will design an extension of DAHL to support multiple machines.
- **Months 10–20:** Conduct experiments to evaluate the upgraded tools, for example, measuring the impact of latency and data transfer between networked machines.
- **Months 20–24:** Disseminate the results by publishing the updated tool on GitHub and writing scientific papers.

Profile required: The postdoctoral researcher should have an expertise in distributed systems and networking. The postdoctoral researcher will also supervise an intern to assist him with the simulations once the multi-machine prototype is operational.

Eligibility Requirements: Applications are open to individuals who hold a doctorate awarded by a French university, or a degree recognized as equivalent by the university, including a doctorate or PhD awarded by a foreign university.

Documents to be submitted with the application:

- cover letter,
- detailed curriculum vitae,
- copy(ies) of the degree(s),
- thesis defense report.

Contact: Stéphan Plassart, Sébastien Monnet (firstname.lastname@univ-smb.fr)

References

J. Malmodin, N. Lövehagen, P. Bergmark, and D. Lundén, ICT sector electricity consumption and greenhouse gas emissions, 2020 outcome, Telecommunications Policy, vol. 48, no. 3, p. 102701, 2024.

A.-C. Orgerie, M. D. d. Assuncao, and L. Lefevre, A survey on techniques for improving the energy efficiency of large-scale distributed systems, ACM Computing Surveys (CSUR), vol. 46, no. 4, pp. 1–31, 2014.

Z. Guo et al., A Survey on Performance Modeling and Prediction for Distributed DNN Training, IEEE Transactions on Parallel and Distributed Systems, vol. 35, no. 12, p. 2463 – 2478, 2024

D. Patterson et al., Carbon emissions and large neural network training, arXiv preprint arXiv:2104.10350, 2021.

P. Li, J. Yang, M. A. Islam, and S. Ren, Making ai less" thirsty": Uncovering and addressing the secret water footprint of ai models, arXiv preprint arXiv:2304.03271, 2023.
Joost Verbraeken, Matthijs Wolting, Jonathan Katzy, Jeroen Kloppenburg, Tim Verbelen, Jan S. Rellermeyer: A Survey on Distributed Machine Learning. ACM Comput. Surv. 53(2): 30:1-30:33, 2021

D. Amodei, D. Hernandez, G. Sastry, J. Clark, G. Brockman, and I. Sutskever, AI and Compute, 2018.

ADEME 2025, Évaluation de l'impact environnemental du numérique en France.
C. Augonnet and R. Namyst, A unified runtime system for heterogeneous multi-core architecures, in European Conference on Parallel Processing, 2008, pp. 174–183.

Andrew Mary Huet de Barochez, Stéphan Plassart, Sébastien Monnet, Falafels: A Tool for Estimating Federated Learning Energy Consumption via Discrete Simulation. AINA (3) 2025: 141-152

DAHL: Distributed, Arbitrary and Heterogeneous Learning, https://github.com/PhoqueEberlue/dahl

FaLaFEℓS: Federated Learning Frugality and Efficiency via Simulation https://github.com/PhoqueEberlue/falafels

W. Won, M. Elavazhagan, S. Srinivasan, S. Gupta, and T. Krishna, TACOS: Topology-Aware Collective Algorithm Synthesizer for Distributed Machine Learning, 57th IEEE/ACM International Symposium on Microarchitecture (MICRO), 2024, pp. 856–870.

S. Figuerola, M. Lemay, V. Reijs, M. Savoie, and B. St. Arnaud, Converged optical network infrastructures in support of future internet and grid services using IaaS to reduce GHG emissions, Journal of Lightwave Technology, vol. 27, no. 12, pp. 1941–1946, 2009.

V. W. Freeh, F. Pan, N. Kappiah, D. K. Lowenthal, and R. Springer, Exploring the energy-time tradeoff in MPI programs on a power-scalable cluster, in 19th IEEE International Parallel and Distributed Processing Symposium, 2005, p. 10–pp.

P. De Langen and B. Juurlink, Leakage-aware multiprocessor scheduling for low power, in Proceedings 20th IEEE International Parallel & Distributed Processing Symposium, 2006, p. 8–pp.

M. Madon, Digital Sufficiency in Data Centers: Studying the Impact of User Behaviors, 2024.