RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

ONERA
THE FRENCH AEROSPACE LAB

www.onera.fr

PROPOSITION DE STAGE EN COURS D'ETUDES

Référence : **DAAA-2026-34**
(à rappeler dans toute correspondance)

Lieu : Châtillon

Département/Dir./Serv. : DAAA

Tél. :

Responsable(s) du stage : Alexandre Suss

Email. : alexandre.suss@onera.fr

DESCRIPTION DU STAGE

Thématique(s) : 09 - Méthodes et outils numériques pour les simulations hautes-performances

Type de stage : ☒ Fin d'études bac+5 ☒ Master 2 ☐ Bac+2 à bac+4 ☐ Autres

Title: Performance & portability of lattice Boltzmann methods on exascale computing platforms

Subject:

The simulation of turbulent flows is a central challenge in the aerospace industry, with applications ranging from drag prediction on aircraft components to characterizing the noise generated by landing gear and high-lift devices. These flows are inherently complex, spanning a wide range of spatial and temporal scale, and high-fidelity simulation approaches such as Direct Numerical Simulation (DNS) or Large Eddy Simulation (LES) are often considered out of reach for industrial-scale problems given current computational resources. Nevertheless, the advent of exascale supercomputers, capable of performing over 10^{18} operations per second; opens unprecedented opportunities to perform advanced numerical simulations and push the boundaries of our understanding of turbulence [1]. Yet fully harnessing this massive computational power is far from straightforward, and remains a major technical hurdle.

Two major trends in modern High-Performance Computing (HPC) help explain these challenges. First, most exascale machines rely on heterogeneous architectures, combining CPUs with GPU accelerators. While such configurations deliver impressive performance and energy efficiency, they introduce new layers of complexity. In particular, the multiplicity of GPU vendors makes performance portability a critical concern: achieving sustained efficiency across CPUs and GPUs from different vendor (Intel, AMD, NVIDIA, among others) requires careful tuning of memory layouts, parallelization strategies, and cross-platform programming models, which vary widely in terms of intrusiveness, abstraction level, and complexity (e.g., OpenMP, OpenACC, SYCL) [2]. Second, the growing demand for AI-oriented workloads has driven the adoption of hardware optimized for lower floating-point precision. While this offers higher performance, improved energy efficiency, and reduced memory footprint, it also requires careful algorithmic design to maintain numerical stability and accuracy, particularly in the context of high-fidelity flow simulations [3].

Characterizing the impact of these trends on computational fluid dynamics (CFD) solvers is therefore a crucial challenge. Existing large-scale Navier-Stokes solvers, developed over decades in both industry and academia, are notoriously difficult to port to modern heterogeneous architectures due to their complex data structures and intricate algorithmic patterns. By contrast, the lattice Boltzmann Method (LBM) [3] offers a more tractable alternative. Based on mesoscopic kinetic theory, the LBM evolves particle distribution functions over a lattice rather than directly solving macroscopic conservation laws. This results in a simple evolution algorithm consisting of a local collision step followed by a streaming step with minimal data dependencies. Its high computational efficiency not only makes the LBM suitable for high-fidelity turbulent flow simulations [4] (see Figure 1), but also provides an ideal testbed for exploring performance optimization strategies and adapting solvers to emerging HPC architectures [5,6].

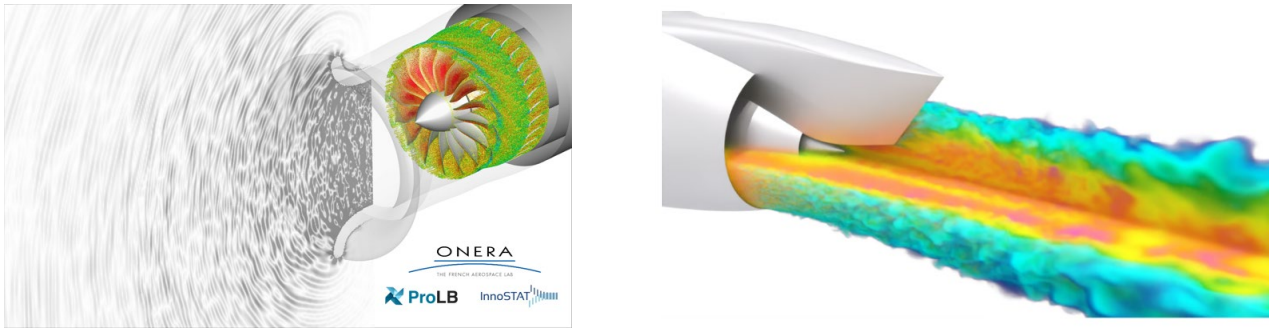

Figure 1 - Examples of high-fidelity flow simulations performed with the lattice Boltzmann method.

The main goal of the internship is (i) to explore the performance of the lattice Boltzmann method on modern heterogeneous HPC architectures, with a particular focus on GPU accelerators, and (ii) to identify best practices for developing efficient and portable solvers adapted to such heterogeneous systems. To achieve this, the internship will include the following tasks:

1. **Familiarization with GPU programming and cross-platform models.** Benchmark different programming models (such as OpenMP, OpenACC, SYCL, …) to identify the most efficient & portable strategies.
2. **Understanding the lattice Boltzmann method and identifying performance-critical kernels.** Perform a baseline CPU performance analysis and profile the code to identify the most computationally intensive kernels.
3. **Porting and optimization of critical kernels on GPUs.** Implement the key kernels on GPU using the chosen programming model and evaluate the achieved performance against the CPU baseline.
4. **Exploration of reduced-precision arithmetic.** Assess the impact of lower floating-point precision on the computation time, numerical stability, and solution accuracy, using concepts introduced in [7].

The intern will have access to ONERA's internal computing facilities and may also perform tests on national HPC facilities. Code development will be tracked and managed using Git to ensure reproducibility and collaborative workflow. The internship will follow an open science approach, making use of ONERA's publicly available codes such as Cassiopee and Fast where relevant. The resulting insights and best practices will directly contribute to ongoing efforts at ONERA to port and optimize its simulation tools for next-generation computing platforms.

Bibliography:

[1] Slotnick, J. P., Khodadoust, A., Alonso, J. J., Darmofal, D. L., Gropp, W. D., Lurie, E. A., Mavriplis, D. J., & Venkatakrishnan, V. (2014). Enabling the environmentally clean air transportation of the future: A vision of computational fluid dynamics in 2030. Philosophical Transactions of the Royal Society A, 372(2022).

[2] Holzer, M. (2024). Code generation in a lattice Boltzmann framework for exascale computing. PhD thesis. Friedrich-Alexander-Universität Erlangen-Nürnberg.

[3] Karp, M., Stanly, R., Mukha, T., Galimberti, L., Toosi, S., Song, H., Dalcin, L., Rezaeiravesh, S., Jansson, N., Markidis, S., Parsani, M., Bose, S., Lele, S., & Schlatter, P. (2025). Effects of lower floating-point precision on scale-resolving numerical simulations of turbulence. https://arxiv.org/abs/2506.05150v1

[3] Krüger, T., Kusumaatmaja, H., Kuzmin, A., Shardt, O., Silva, G., & Viggen, E. M. (2017). The Lattice Boltzmann Method: Principles and Practice. Springer International Publishing.

[4] Suss, A., Mary, I., Le Garrec, T., & Marié, S. (2023). Comprehensive comparison between the lattice Boltzmann and Navier–Stokes methods for aerodynamic and aeroacoustic applications. Computers & Fluids, 257, 105881.

[5] Coreixas, C., & Latt, J. (2025). GPU-based compressible lattice Boltzmann simulations on non-uniform grids using standard C++ parallelism: From best practices to aerodynamics, aeroacoustics and supersonic flow simulations. https://arxiv.org/abs/2504.04465v2

[6] Latt, J., Coreixas, C., & Beny, J. (2021). Cross-platform programming model for many-core lattice Boltzmann simulations. PLOS ONE, 16(4), e0250306.

[7] Lehmann, M., Krause, M.J., Amati, G., Sega, M., Harting, J., Gekle, S. (2022). Accuracy and performance of the lattice Boltzmann method with 64-bit, 32-bit, and customized 16-bit number formats. Physical Review E, 106, 015308.

Est-il possible d'envisager un travail en binôme ? Non

Méthodes à mettre en œuvre :

☐ Recherche théorique ☐ Travail de synthèse

☒ Recherche appliquée ☒ Travail de documentation

☐ Recherche expérimentale ☐ Participation à une réalisation

Possibilité de prolongation en thèse : A renseigner

Durée du stage : Minimum : 5 mois Maximum : 5 mois

Période souhaitée : 1st semester of 2026 (between January and September)

PROFIL DU STAGIAIRE

Connaissances et niveau requis :	Ecoles ou établissements souhaités :
Computer science, programming (Python, C++, Fortran), high-performane computing, applied mathematics, numerical methods	Master of science majoring in computer science, applied mathematics or fluid mechanics French Engineering school