🚀 Job Offer: C++/HPC PostDoc - MLIR / JAX / Pytorch / Kokkos

🧭 Context

The European HPC landscape is transforming rapidly with the emergence of **Exascale** systems like the upcoming *Alice Recoque* supercomputer at **CEA**. These heterogeneous architectures require new programming paradigms to maintain performance and portability across varied GPU and CPU vendors.

To address this challenge, the **CExA "Moonshot"** initiative contributes to key software infrastructure like **Kokkos** and leverages modern compiler frameworks such as **MLIR** (Multi-level Intermediate Representation) to establish a sustainable approach for domain experts to develop their applications without deep understanding of the modern programming models and architectures. For this purpose, we will explore the coupling technology of kokkos/python with MLIR to allow code developers to maintain kokkos codes only with the python expertise. As a case study, we will port a legacy fortran application for MDFT (Molecular Density Functional Theory) with modern programming approaches including Kokkos, Kokkos/python and python. Then, we will evaluate the performance portability, readability, productivity and mainatinability aspects of each approach to find a reasonable programming model to develop a code for domain scientists.

This position is a 2 years posistion as part of CExA's efforts to integrate advanced techniques, including **JAX-based or PyTorch-based AI models**, into our HPC workflows, fostering hybrid workloads that combine simulation and machine learning.

🎯 Mission

As a core member of the MDFT and CExA teams, your mission will be to **develop, prototype, and evaluate tools and methodologies** for performance optimization of MLIR-Kokkos-based HPC applications. A publication at highly competitive conferences such as Supercomputing is forseen.

Your main responsibilities will include:

- 🔧 **Tooling Development**
  - Introduce a **Kokkos dialect** in **MLIR** as an intermediate representation of Kokkos
  - Implement a translation inflastructure (based on potentially existing intermediate dialects) between Kokkos and JAX / PyTorch
- 📈 **Code Optimization**
  - Profile and optimize MDFT across multiple architectures (Nvidia GPUs, AMD CPUs/GPUs, Intel CPUs)
  - Evaluate integration opportunities for AI-based models via **JAX** or **PyTorch** to support hybrid workflows
- 🤝 **Collaboration & Outreach**
  - Collaborate with other CEA teams and international partners (Europe, US, Japan)
  - Present results in conferences, workshops, and open-source contributions

🧠 Skills

You are a passionate engineer with a background in scientific computing and a drive to push the boundaries of performance portability.

Must-have:

- Experience with LLVM, or compiler infrastructure
- Proficiency in **modern C++ (17 or later)**

Bonus:

- Experience with performance optimization and profiling tools (e.g., VTune, Nsight)
- Familiarity with **Kokkos** or other GPU portability frameworks
- Familiarity with **MLIR**, or compiler infrastructure
- Solid software engineering practice: git, cmake, CI/CD, unit testing
- Knowledge of **JAX** or **PyTorch** and interest in ML or DL for scientific applications

💼 Salary & Benefits

- Competitive salary based on experience and qualifications
- Up to **3 remote days/week**
- **75% reimbursement** of public transport costs
- 5 weeks of vacation + 4 weeks RTT
- **International exposure**: conferences, workshops, and collaborative projects
- Health insurance, retirement savings plans, and more

📩 Apply Now

Send your **CV** and **cover letter** to yuuichi.asahi(at)cea.fr.

Applications will be reviewed continuously until the position is filled.