# 2-year Post-doctoral Position - Open hardware compilation for near-memory dataflow computing

Keywords: `High-Performance computing`, `Near-memory computing`, `Compilation`, `Dataflow models of computation`, `Stream processing`, `Open hardware`, `RISC-V`, `FPGA`, `Computer architecture`.

**Laboratory: Institut d'Electronique et des Technologies du numéRique**
**UMR CNRS 6164, Rennes, France**

## Context

**Near memory computing (NMC) consists of adding computing capacities (typically processors) as close as possible to the main memory of the computer, thus bypassing the hierarchy of caches, and limiting its numerous data transfers** [Singh19]. Memory is a vital part of any digital system. It is used to store the programs of the system and the data on which these programs operate. The speed at which a program runs directly depends on the speed at which instructions and data can be transferred between memory and processor. The memory must therefore be (very) fast, but ideally also (very) large to store all programs and data, and inexpensive. The impossible meeting between these three requirements has led to proposing solutions based on a layered memory hierarchy, relying on various memory technologies. Through the numerous hardware solutions implemented from generation to generation, more and more transistors and energy in a circuit are allocated for the sole purpose of improving memory access [Horowitz14]. Cache memories, embedded as close as possible to the processor on the same chip gives the programmer the illusion of both fast and large memory. The cache memory exploits the assumptions of both spatial and temporal data access locality, assumptions that do not hold when processing data streams [Ghasemi21], as data undergo a *process-and-forget* management. This post-doctoral contract is financed by the **RIDIM project, led by Lab-STICC and involving Irisa and IETR laboratories, in collaboration with the University of Maryland, USA**. The RIDIM project aims to integrate stream processing capabilities all along the data path from the main memory to the processor.

## Objectives

This post-doctoral contract is primarily intended towards compilation objectives but the objectives of the post-doc will be adapted w.r.t. the experience of the recruited candidate. One goal of the post-doc is to invent new compilation methods to efficiently make use of near-memory computing facilities, and to test these methods in simulation and on a functional FPGA prototype. These methods will be integrated in the PREESM dataflow rapid prototyping tool https://preesm.github.io/, a compiler based on Eclipse and written in Java and for which a new version based on LLVM/MLIR is currently considered [Lattner21]. Signal processing from the Square Kilometer Array (SKA) telescope [Miomandre20] is among the first applications to be considered, already coded. Other applications such as the SqueezeNet convolutional neural network, a Versatile Video Coding (VVC) decoder, a swimmer detector in underwater images, and a stereo matching depth map computer are available on the Preesm Apps repository as PiSDF dataflow functional descriptions [Desnos13].

## Environment and Expected Impact

The post-doc will benefit from a close collaboration between the three RIDIM research teams: the project is coordinated by Kevin Martin at Lab-STICC in Lorient and involves Steven Derrien from the TARAN team of IRISA and the VAADER team of IETR in Rennes. **The project will build on the momentum of the Open Hardware and RISC-V communities that are eager for novel architecture propositions. RIDIM teams are involved in this community and participate to the organising committee of the yearly RISC-V week, largest event on RISC-V in Europe** . The RIDIM post-doc will be closely linked to the Cominlabs International Chaire at INSA/IETR of Prof. Shuvra S. Bhattacharyya, Professor of Electrical and Computer Engineering at the

University of Maryland, College Park, USA. The post-doc will be located in Rennes, France, a very dynamic city with nearly 700,000 inhabitants under its urban umbrella, and a fast growing population (the city has the second largest population growth in France).

This post-doc and the RIDIM project are expected to set the foundations of a programming model and its supporting hardware for building smarter distributed systems with stream processing facilities. The main objective of project dissemination is through high quality publications at events such as DATE, DAC, Embedded Systems Week, ACM TACO, as well as in highly ranked journals on system design.

## Candidate

The candidate shall hold a PhD in computer science or electrical engineering, with experience in the following areas:

| Required experience | Experience to be gained in the post-doc |
| --- | --- |
| • Software design (required), Compilation (appreciated)<br>  • C, C++, and (Java or LLVM/MLIR) (required)<br>  • Python (appreciated)<br>• Digital hardware design (appreciated)<br>• Parallel programming (appreciated) | • Open source hardware and RISC-V systems<br>• Dataflow models of computation<br>• Project-level task lead<br>• Astronomy computing workloads |

## Characteristics

- location: Vaader team, IETR laboratory - INSA Rennes, 20 Avenue des Buttes de Coësmes , 35708 Rennes, France
- Duration: 2 years, Start: ASAP
- Salary per month: depends on experience, minimum 2950€ gross / 2350€ netto
- Supervisors
  - Maxime Pelcat (IETR, INSA Rennes) - maxime.pelcat@insa-rennes.fr
  - Jean-François Nezan (IETR, INSA Rennes) - jnezan@insa-rennes.fr
  - Karol Desnos (IETR, INSA Rennes) - kdesnos@insa-rennes.fr

## Applications

You may request details on the subject, and send your resume and application letter to Maxime Pelcat

## References

[Desnos13] Karol Desnos, et al. "Pimm: Parameterized and interfaced dataflow meta-model for mpsocs runtime reconfiguration." SAMOS XIII. IEEE, 2013.
[Ghasemi21] Alemeh Ghasemi, et al. "On cache limits for dataflow applications and related efficient memory management strategies". DASIP'21, ACM, 2021.
[Horowitz14] Mark Horowitz. "Computing's energy problem (and what we can do about it)"". In 2014 IEEE ISSCC. IEEE, 2014.
[Kermarrec19] Florent Kermarrec, et al. "LiteX: an open-source SoC builder and library based on Migen Python DSL." OSDA 2019.
[Lattner21] Chris Lattner, et al. "Mlir: Scaling compiler infrastructure for domain specific computation." 2021 IEEE/ACM CGO. IEEE, 2021.
[Miomandre20] Hugo Miomandre, et al. "Approximate Buffers for Reducing Memory Requirements: Case study on SKA". IEEE SiPS 2020.
[Singh19] Gagandeep Singh, et al. "Near-memory computing: Past, present, and future". Microprocessors and Microsystems, 2019.