Grenoble INP - UGA is a member of **international** engineering and management education and research **networks.** It is widely recognized in national and international rankings**.**

8 schools + **39** laboratories

8300 students

1 300 teaching, research, administrative and technical staff

Grenoble INP - UGA is a renowned public institution of higher education and research, and a major player in the Grenoble ecosystem. It is the engineering and management institute of Grenoble Alpes University, and plays a leading role in the scientific and industrial community.

PhD Candidate in Computer Design

Job reference number	2024-PHDCOMPUDESIGN-TIMA
Research field	Computer Architecture and Design
Host laboratory	TIMA (UMR 5159 Grenoble-INP, UGA and CNRS) / Website : https://tima.univ-grenoble-alpes.fr
Researcher profile	PhD / Doctorate
Location	Grenoble, France
Date of recruitment / contract term	1/10/2024 (36 months)
Contacts	Arthur Perais (arthur.perais@univ-grenoble-alpes.fr)

Grenoble INP - UGA is a leading public institution accredited with the French label "Initiative d'excellence". It offers innovative engineering and management programs, with an increasing internationalization of its course offers. The courses are grounded in sound scientific knowledge and linked to digital, industrial, organizational, environmental and energy transitions. The Engineering and Management Institute of Grenoble Alpes brings together more than 1300 staff members (teacher-researchers, lecturers, administrative and technical staff) and 8300 students, located on 8 sites (Grenoble INP - Ense3, Grenoble INP - Ensimag, Grenoble INP - Esisar, Grenoble INP - Génie industriel GI, Grenoble INP - Pagora, Grenoble INP - Phelma, Polytech Grenoble, Grenoble IAE and the INP Prepa). Grenoble INP is also a highly-ranked institution of higher education and research, leading the way in the fields of engineering and management on an international scale. It is a member of a large number of international academic and research networks. It is part of the European University UNITE!.

As part of Grenoble Alpes University, Grenoble INP has associated guardianship of 39 national and international research laboratories and of technological platforms. The research conducted there benefits both its socio-economic partners and its students. Grenoble INP is at the heart of the following scientific fields: physics, energy, mechanics and materials; digital; micronanoelectronics, embedded systems; industry of the future, production systems, environment; management and business sciences.

Grenoble INP - UGA is an equal opportunity employer committed to sustainability. Grenoble INP-UGA celebrates diversity and equity and is committed to creating an inclusive environment for all employees. All qualified applications will be considered without discrimination of any kind.

Research

The research topics of the TIMA laboratory cover the specification, design, verification, test, CAD tools and design methods for integrated systems, from analog and digital components on one end of the spectrum, to multiprocessor Systems-on-Chip together with their basic operating system on the other end.

Within TIMA and in the context of this offer, the SLS team focuses on general purpose microarchitecture, that is, how to arrange the billions of transistors available thanks to Moore's Law to accelerate the execution of program instructions. Microarchitecture is at the interface between circuit design and instruction set architecture design (ISA), as it in charge of building a coherent system that executes the instructions defined in the ISA knowing that this system should have good performance when built as a circuit (max frequency, power consumption, silicium area). The team is known internationally in the field of microarchitecture with publications in flagsip conferences (e.g., IEEE/ACM ISCA, IEEE/ACM MICRO).

Position description:

Modern general purpose processors leverage speculation at the hardware level to accelerate program execution. Typical examples include hardware branch prediction, cache replacement policies, data and instruction prefetching, etc. Those mechanisms are typically designed with user code in mind, as many embedded and desktop-level applications spend most of their time executing their own code rather system calls. However, we can envision two scenarios where applications end up spending a significant time in kernel code:

1. "Datacenter" applications such as web servers, databases, etc. make significant use of I/O primitives.
2. Following the end of Dennard scaling, system-on-chips now feature accelerators dedicated to specific applications (e.g., GPU for graphics processing, NPU/TPU for machine learning applications, etc.). In the future, one can imagine that most of the code will be ran on accelerators, when the general purpose processor will mostly be used to run the system.

Given this, it appears that efficiently executing system code may become comparatively more important than it was for embedded and desktop-level applications.

This thesis will therefore focus on two questions:

1. Is system code fundamentally different from user code in terms of e.g., instruction mix, data spatial and temporal locality, code footprint, etc.
2. Are microarchitectural speculation techniques that were designed with user code in mind adapted to system code?

A first requirement is to be able to analyze (e.g., trace) both userspace and kernel space code within the same application. This can be achieved through binary translation and instrumentation tools that execute the system (e.g., QEMU) and not only the application (e.g., Pin). This may also be obtainable through hardware extensions (e.g., Intel Processor Trace) or by adding tracing directly at the RTL level in an open-source processor (e.g., Boom, Xiangshan). Regardless, the goal is to obtain representative traces of execution and to replay them offline to obtain metrics of interest and find interesting differences between user code and kernel code that could be leveraged at the hardware level.

Therefore, as a first step, the candidate will be tasked with understanding the different options available to obtain those traces and to weigh their pros and cons. They will then setup an infrastructure to obtain traces knowing that target applications are large (e.g., DeathStarBench), potentially multithreaded, and potentially client-server, thus are non trivial to setup for tracing.

As a second step, and based on the analysis performed in the first step, the candidate will propose hardware schemes that have potential to improve system code performance. Those schemes may have to work for both user code and system code, despite the two potentially having different characteristics.

For instance, it is not unlikely that the dataset used by a system call will have temporal locality within the system call execution, but will not be reused once the system call has terminated. That is, the next instance of that system call will use different data. Therefore, it may be desirable to tailor the cache replacement policy based on who (user or kernel) brought the data into the cache.

Specific requirements or conditions

The candidate should hold a Master of Science in Computer Science or Computer Engineering, or any equivalent degree.

Personnal Skills:

The candidate should be aware that a PhD programme is vastly different from a BS or MS programme. Pursuing a PhD requires strong motivation and the ability to focus on a specific topic for three years.

- The candidate should expect to be autonomous in developing software, experiments, and analyzing results.
- The candidate should be able to clearly express their ideas and conclusions, and to motivate their research directions.
- The candidate should be open to constructive criticism from their peers and supervisors.

Technical Skills:

- **Programming:** C/C++. Strong knowledge and understanding of the langage, data structures, algorithms, testing and debugging tools.
- **Linux scripting**: Python, bash or other.
- **Computer Architecture:** At least undergrad level (caches, virtual memory, pipelining) and Instruction Set Architecture concepts. Advanced level (branch prediction, out of order execution) is a plus.
- **Operating Systems (Linux):** At least undergrad level (virtual memory, privilege levels, system calls), advanced level is a plus.
- **English:** Strong level is required as scientific articles are written and presented in english.

Specifics of the position:

Direction: The thesis will be directed by Prof.~Frédéric Pétrot (Grenoble INP) and co-supervised by Arthur Perais (CNRS).

Institution: The PhD will be hosted in the TIMA laboratory in the SLS team: https://tima.univ-grenoble-alpes.fr/research/sls.

Doctoral School: The PhD programme is tied to the MSTII Doctoral School of Université of Grenoble: https://edmstii.univ-grenoble-alpes.fr/mstii-doctoral-school/

Funding: The PhD is fully funded for 3 years by project ARCHI-Cesam which is part of France 2030 PEPR Cloud, and the employer is the National Polytechnic Institute of Grenoble (Grenoble INP, https://www.grenoble-inp.fr/). The monthly gross salary is 2200€ (after 1/01/2025) and will increase to 2300€ after 1/01/2026.

Teaching: During the 3 years, teaching (labs and practicals) is possible but english classes are the exception, so if the candidate wishes to teach, fluency in french is highly preferable. Teaching is remunerated at around 40€ gross per hour in front of students.

Mandatory Training: During the 3 years, the candidate will be expected to attend 120 hours of training, split in three topical buckets of 40h: Scientific, Transversal and Professional Project. More information: https://edmstii.univ-grenoble-alpes.fr/formation/

Position assigned to a restricted area: NO

(Device for the protection of the scientific and technical potential of the nation, conditioning the appointment of the researcher to the authorization of the Defense Security Officer).

How to apply

Applications (resume and cover letter) must be sent to : arthur.perais@univ-grenoble-alpes.fr

Application deadline : 27/06/2024