

# PhD Candidate in Computer Architecture (HW/SW interaction between OS and microarchitecture)

- **Deadline:** Aug. 31, 2024
- **Career levels:** Master
- **Keywords:** Cloud computing, Computer architecture, Operating systems

Offer 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 than system calls. However, we can envision two scenarios where applications end up spending a significant time in kernel code:

- "Datacenter" applications such as web servers, databases, etc. make significant use of I/O primitives.
- 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 run 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.

Objectives: This thesis will focus on two questions:

- Is system code fundamentally different from user code in terms of e.g., instruction mix, data spatial and temporal locality, code footprint, etc.
- 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.

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

**Personal 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 be mostly 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 language, 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.
- French is not required but will be a plus for day to day interactions in general
- **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 2173.5€.
- **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/>
- **Start date:** October 2024
- **Duration:** 3 years