Hardware Engineers – CPU development
From Micro-architecture to Physical Implementation

Arm Sophia Antipolis
Le Paros – 25 allée Pierre Ziller
Sophia-Antipolis
06560 VALBONNE
Tel: 04.97.23.51.00

[2018-D-1] CPU frontend simple instructions elimination
[2018-D-2] Deep Analysis of state-of-the-art Branch Prediction Algorithm
[2018-D-3] Design of a Safety Module allowing CPU functional mode testing
[2018-D-4] Design of CPU power limitation module using machine learning
[2018-D-5] L2 Cache performance and power optimization techniques
[2018-D-6] Optimization techniques for address translation cache
[2018-D-7] Performance and feasibility of alternative store queue structures
[2018-D-8] Prioritization of critical instructions in a superscalar CPU
[2018-D-9] Scalable issue structure for superscalar processor

[2018-I-1] High Performance Clock distribution
[2018-I-2] Processor Dynamic Power Profiling
[2018-I-3] Register file custom implementation
[2018-I-4] RTL benchmark library development

Essential skills:
Different sets of skills and knowledge are required depending on the selected internship topic. These should desirably include:

- Microprocessor, ASIC systems
- HDL Design/Synthesis/Place & Route knowledge
- Use of UNIX and shell programming
- Strong analytical skills
- Ability to express ideas and communicate effectively

Interested in the latest processor technologies? Join us at: **http://www.arm.com/about/careers/students/student-vacancies.php**

Micro-Architecture & RTL Design
goo.gl/Mu9ubo

[2018-D-1] CPU front-end simple instructions elimination

In a superscalar CPU, the front-end stages of the pipeline are responsible for preparing instructions for execution in the out-of-order back-end. Some instructions are simple enough to be eliminated in the front-end, reducing the use of the execution stages. The student will study and identify, by mean of a CPU model, the instructions that can be eliminated; the RTL module corresponding to the identified optimization will be implemented in a state-of-the-art CPU and its impact on performance/power/area will be assessed.

[2018-D-2] Deep Analysis of state-of-the-art Branch Prediction Algorithm

Branch prediction is a key element for superscalar processors performance. The goal of this internship is to study different high-accuracy branch prediction algorithms in a cycle accurate CPU model, provide a precise comparison of their characteristics and implement them in a state-of-the-art CPU instruction fetch pipeline.

[2018-D-3] Design of a Safety Module allowing CPU functional mode testing

With the growing demand of Safety compliance, this internship aims at designing a module that could allow automatic testing of a CPU in normal operation. The student will work on a modern CPU and cover all phases of the design of the safety module in order to assess its feasibility

[2018-D-4] Design of CPU power limitation module using machine learning

Power consumption is a concern in a CPU and must be limited to a maximum threshold. The goal of this internship is to define a model to correlate internal events with power consumption and design a hardware module (RTL) to limit the maximum power of the CPU. Machine learning will be used to model and estimate the power consumption of the processor.

[2018-D-5] L2 Cache performance and power optimization techniques

Associativity and power consumption of the L2 cache are important parameters in a processor. The goal of this internship is to investigate new solutions to improve these parameters in a state-of-the-art L2 cache of a superscalar CPU. The student will study the current cache architecture, get familiar with Arm design and verification environments, propose and benchmark the new solution.

[2018-D-6] Optimization techniques for address translation cache

The TLB speeds up mapping of virtual address to physical address in MMU-based translation systems. The goal of this internship is to study optimization techniques such as TLB prefetching in order to increase the efficiency of this structure. The student will discover state of the art of TLB, identify new opportunities for more efficient translation caching, and design an improved structure on a modern CPU.

[2018-D-7] Performance and feasibility of alternative store queue structures

The store queue is a key structure for performances in a processor memory system. The goal of this internship is to study alternative (more scalable and more efficient) structures for this component in the context of a superscalar CPU. The student will study several solutions in a cycle-accurate CPU model and then design the best ones in RTL to compare its characteristics with state-of-the-art implementations.

[2018-D-8] Prioritization of critical instructions in a superscalar CPU

Long latency instruction can limit overall IPC in a superscalar processor. The goal of this internship is studying a predictor which is capable of identifying critical long latency instructions to be prioritized using iterative backward dependency analysis. The student will propose a solution to be implemented in a state-of-the-art CPU architecture.

[2018-D-9] Scalable issue structure for superscalar processor

Superscalar CPUs need out-of-order schedulers to feed the increasing number of execution units. The goal of this internship is to study a new scalable structure for next generation Arm processors. This solution will be first investigated in a cycle-accurate CPU model and then designed in RTL to compare its characteristics with state of the art implementations.

Physical Implementation
goo.gl/6AHF9P

[2018-I-1] High Performance Clock distribution

Clock is the most important signal in a synchronous design, pacing all operations. Distributing it efficiently in larger and larger design is a growing challenge. The intern choosing this subject will first review different clock distribution styles.
He/She will have to put in place scripted flow to implement with various clock distribution styles in the most efficient way and analyze impact on frequency/power and area of a high-end application microprocessor.
If needed, clock tree micro-architecture modification will be proposed.

[2018-I-2] Processor Dynamic Power Profiling

Power is a major concern during microprocessor design. Not only average power, but also real-time power and real-time current as it can swing above and below the average value.
The intern will have to set up a flow based on existing tools and scripts to monitor dynamic power profile on an existing microprocessor.
Then the intern will study power changes on typical test and benchmarks and develop more focused code to activate specific part of the micro-architecture.
The intern will focus not only on overall dynamic power evolution through time but also on physical hot spot map across a microprocessor layout and its evolution through test execution.
Additional reverse analysis will be done to see how processor activity can be deduced from power profiling.

[2018-I-3] Register file custom implementation

Modern microprocessor development relies heavily on automated placement through EDA tool.
For some structured parts of the design, a more hand-crafted approach can bring power, area and frequency benefits.
The intern will first explore the various steps of a usual ASIC flow, and possible automated customization of it, to improve physical implementation of a regular structure commonly found in high performance microprocessors.

[2018-I-4] RTL benchmark library development

Modern microprocessor development relies heavily on EDA tool usage and HDL description of functionality.
This internship aims at studying interactions between these two fields. It should lead to validation of coding rules used to get the best frequency possible, power and area on a HDL IP.
By extracting relevant structures and gathering them in a library of test case, it will improve new EDA tool evaluation.