technische universität dortmund

computer science 12

Bachelor Thesis

Generating ISA microbenchmarks for modern RISC microarchitectures

Nils Hölscher
Otto-Hahn Str. 16
Technische Universität Dortmund
Email: Email: nils.hoelscher@tu-dortmund.de
13.01.2022

Modern Microarchitectures are very complex man-made systems. To predict the performance or optimize such systems, a profound knowledge about the underlying models is needed. However, such models are often incorrect, incomplete or not available at all. Therefore, automated tools are desirable to measure and benchmark microarchitectures.

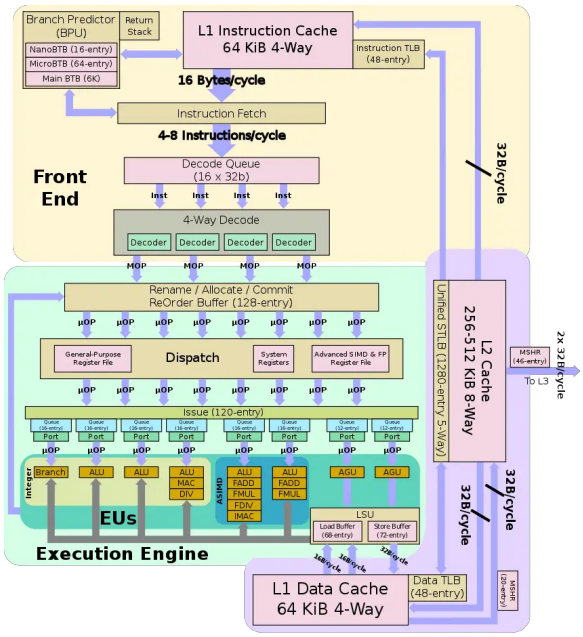

Figure 1: Pipeline of the Cortex A76, taken from its WikiChip.

Modern compilers like LLVM and GCC use information about instruction latency, throughput and functional unit usage for code optimization. Latency is the number of clock cycles that are required for the execution core to complete the execution of the instruction. Throughput is the number of clock cycles required to wait before the functional units are free to accept the same instruction again.
Functional unit are small sections of a CPU or microcontroller, that are able to process a subset of instructions from the ISA. For example, most CPUs have a functional unit for floating point calculation, another example would be a vector unit and the most common ALU (Arithmetic-Logic-Unit). While the usage of functional units can be easily guessed per instruction, it is still necessary to evaluate such claims.
Another integral part for measuring and evaluating the documentation are the provided performance measurement units (PRUs). These are counters embedded in the CPUs hardware and allow measuring many different events, happening in the pipeline. Most PRUs posses some counting registers and way more events to count. This means that the events have to be mapped to counting registers and not all can be measured at once. The most common events are cycles, cache hits/misses and many more.
Only if these measurements are accurate, compilers like LLVM or GCC are able to perform yielding optimizations. Not only program optimizations needs this information also analysis tools like llvm-mca or WCET-analyzers (Worst Case Execution Time) will gain from correct and faithful information.

In this thesis, students should investigate the given hardware properties of RISC microarchitecture and get familiar with the specific hardware performance counters. Students afterwards should implement simple microbenchmarks to acquire the before stated metrics, such as latency, throughput and functional unit usage. After that the measured metrics need to be compared with official and unofficial documentation of the studied microarchitecture.

Other suggestions and related topics are also welcome. Please do not hesitate to make an appointment.

Required Skills:

- Knowledge of computer architecture
- Knowledge of RISC ISAs like ARM or RISC-V

Acquired Skills after the thesis:

- Deep Knowledge about the targeted ISA
- Understanding of microarchitecture
- Experience with hardware performance counters