

Towards continuous improvement of code-generation for RISC-V

Tuesday, 21 September 2021 10:30 (30 minutes)

Architectures competing with RISC-V have expended considerable time and resources on optimizing their development tools for improved performance on industry-standard benchmarks. For the future growth of the RISC-V ecosystem, a concerted effort to optimize the generated code for performance will be required. This effort will in a large part be independent of the underlying microarchitecture and can be distributed across our entire ecosystem, if we develop the necessary tools and infrastructure to assess for gaps, distribute the work and cooperate.

We propose a data-driven methodology, based on the gathering and comparison of hot-block information, instruction-type histograms and dynamic instructions counts, to evaluate the performance of compilers for RISC-V using qemu. Based on example findings and data, we will illustrate the proposed workflow and how it can allow the prioritisation of potential optimizations based on an expected gain.

We aim to motivate increased cooperation and the creation of a data-driven workflow built around standard tools (primarily plugins for qemu and analysis tools) to continuously monitor and improve the quality of the RISC-V compilers.

I agree to abide by the anti-harassment policy

I agree

Primary authors: TOMSICH, Philipp (VRULL GmbH); MÜLLNER, Christoph (SBA Research)

Presenters: TOMSICH, Philipp (VRULL GmbH); MÜLLNER, Christoph (SBA Research)

Session Classification: RISC-V MC

Track Classification: RISC-V MC