
Automated performance analysis of software stacks for distributed systems

- ▶ Author: Fabian Schröder
- ▶ Type: Bachelor's Thesis
- ▶ Date: 2023-09-13
- ▶ Reviewers: Prof. Dr. Michael Kuhn, Michael Blesel
- ▶ Download: PDF

The performance of clusters is influenced by both the hardware and the software it uses. For key functionalities such as a compiler and libraries such as MPI and BLAS, there are a large number of different implementations that might achieve different performances depending on the exact cluster setup. Finding well-performing implementations is therefore important to maximise the efficiency with which the available resources are used in terms of program runtime, throughput, etc. HPC-Benchmarks-2, an extension of HPC-Benchmarks, is a configurable tool designed to find a software stack consisting of a compiler and an implementation of MPI and BLAS on a SLURM-based cluster, using spack to manage the different software configurations. The stack is built during in a slurm job by iteratively running benchmarks, analysing the results obtained and selecting the best software implementation of each of the three types. The comparison is based on an estimate of the performance improvement each implementation could offer, using only selected performance metrics and weights to quantify their assumed importance. The program largely works as intended, although there are some software incompatibility issues. The results suggest that a combination of the specifications gcc@13.2.0 (compiler), mpich@4.1.2 netmod=tcp (MPI) and openblas@0.3.23 (BLAS) form the best software stack for the cluster on which the tests were performed. The choice of weights and metrics was mixed in the extent to which they accurately matched and emphasized the performance differences between the software implementations.
