

Characterization of Benchmarks: Memory Characteristics – Capacity and Performance

Bachelor's/Master's Thesis – Open Assignment

Assigning professor:

Prof. Dr. Martin Schulz

Supervisors:

Karlo Kraljic,

Dr. Matthias Maiterth

Contact: karlo.kraljic@tum.de

October 24, 2022

Description:

Benchmarks are an essential tool for performance assessment of HPC systems. During the procurement process of HPC systems both benchmarks and proxy applications are used to assess the system which is to be procured. New generations of HPC systems often serve the current and evolving needs of the applications for which the system is procured. Therefore, with new generations of HPC systems, the selected proxy application and benchmarks to assess the systems' performance are also selected for the specific needs of the system. Only a few of these have stayed persistent over longer time periods. At the same time the quality of benchmarks is typically not questioned as they are seen to only be representatives of specific performance indicators.

This work aims to provide a more systematic approach with the goal of evaluating benchmarks targeting the memory subsystem, looking at capacity latency and bandwidth.

Problem statement:

How can benchmarks used to assess memory performance, including cache usage, be systematically compared amongst each others?

Anticipated Outcome

The objective of this work is to assess a selection of benchmarks used in HPC as indicators of memory performance. This includes small kernels to probe the operation of the memory subsystem as well as applications which are typically limited by the memory performance of the system.

The anticipated outcome is a metric to compare the quality of benchmarks, independent of the used system architecture.

- Quality indicators may include used maximum capacity, read and write latency, bandwidth, random and contiguous access, impact of other hardware components, and representativeness regarding HPC workloads, among others.
- The starting point for the benchmark set are the two prominent examples: STREAM benchmark [1] and the HPCG benchmark [2], evaluating the local memory performance.

- Additionally simple tests for saturating the different levels of caches are in scope. For cross-platform evaluation, the beast system is used [3].
- For collection of performance counters the LIKWID tools are used [4].
- The metric, quality indicators, set of benchmark and architectures to evaluate are defined in the early stages of the thesis.

The possibility of ranking the benchmarks for specific architectures has to be assessed and may be an additional outcome.

References

- [1] John D. McCalpin. *STREAM: Sustainable Memory Bandwidth in High Performance Computers*. URL: <https://www.cs.virginia.edu/stream/ref.html> (visited on 2022-05-16).
- [2] Innovative Computing Laboratory University of Tennessee. *HPCG Benchmark*. URL: <https://www.hpcg-benchmark.org/> (visited on 2022-05-16).
- [3] Leibniz-Rechenzentrum. *Testing HPC Technologies and Shaping the Future of Computers*. URL: https://www.lrz.de/presse/ereignisse/2020-11-06_BEAST/ (visited on 2022-05-16).
- [4] LIKWID Performance Tools git. <https://github.com/RRZE-HPC/likwid>. URL: <https://github.com/RRZE-HPC/likwid> (visited on 2020-03-19).