



Contribution ID: 124

Type: **Poster**

## A Microbenchmark Framework for Performance Evaluation of OpenMP Target Offloading

*Thursday, 14 March 2024 16:10 (30 minutes)*

We present a framework based on Catch2 to evaluate performance of OpenMP's target offload model via micro-benchmarks. The compilers supporting OpenMP's target offload model for heterogeneous architectures are currently undergoing rapid development. These developments influence performance of various physics applications in different ways. This framework can be employed to track the impact of compiler upgrades and compare their performance with the native programming models. We use the framework to benchmark performance of a few commonly used operations on leadership class supercomputers such as Perlmutter at National Energy Research Scientific Computing (NERSC) Center and Frontier at Oak Ridge Leadership Computing Facility (OLCF). Such a framework will be useful for compiler developers to gain insights into the overall impact of many small changes, as well as for users to decide which compilers and versions are expected to yield best performance for their applications.

### Significance

Several portable programming models have been developed in the last decade as a solution to avoid code duplication and diversion for different GPU backends. This work focuses on OpenMP's target offload model which has undergone rapid development in the past few years and gained increasing vendor support. Simplified algorithmic performance benchmarks provide a good overview of the current compiler support for OpenMP target offload, however, they often lack the granularity to evaluate the performance of specific, often specialized, operations, such as atomic, memset or scan. We have encountered several issues related to these operations when porting application codes to OpenMP, which motivated us to develop necessary tools for the better evaluation of specific operations in OpenMP target offload.

### References

### Experiment context, if any

**Primary authors:** LIN, Meifeng; ATIF, Mohammad (Brookhaven National Laboratory); WANG, Tianle (Brookhaven National Lab); DONG, Zhihua

**Presenter:** ATIF, Mohammad (Brookhaven National Laboratory)

**Session Classification:** Poster session with coffee break

**Track Classification:** Track 1: Computing Technology for Physics Research