

1 REPRODUCIBILITY DESCRIPTION APPENDIX

1.1 Description

1.1.1 Check-List (Meta Information).

- **Program:** The md_hom code generator implemented in C++; the Auto-Tuning-Framework (ATF) implemented in C++; the Lift compiler implemented in Scala; Benchmark programs implemented in C/C++ and OpenCL.
- **Compilation:** C++ compiler, Intel OpenCL SDK, NVIDIA OpenCL SDK
- **Runtime environment:** Linux with OpenCL
- **Hardware:** Intel Xeon E5 CPU, NVIDIA Tesla K20 GPU
- **Profiling:** OpenCL Events, C++ chrono library, NVIDIA profiler (nvprof)
- **Output:** Runtime in CSV files and plot as PDF
- **Experiment workflow:** Git clone; run build scripts; run test scripts; run plotting script; observe performance results

1.2 Hardware Dependencies

An OpenCL enabled multi-core CPU and GPU is required. For our results, we use an Intel Xeon E5 multi-core CPU and an NVIDIA Tesla K20 GPU.

1.3 Software Dependencies

- The md_hom code generator requires a working C++ compiler.
- Lift requires Java 8, OpenCL, and a working C++ compiler as its main dependencies.
- ATF requires a working C++ compiler and the OpenTuner framework¹.

1.4 Evaluation and Expected Result

The main results of the artifact evaluation is to reproduce the performance comparison given in Figure ?? of the paper. Depending on the precise multi-core CPU and GPU used for evaluation, we expect similar performance trend as reported in the paper between the md_hom OpenCL kernels compared to the reference implementations.

¹<https://github.com/jansel/opentuner>