

Foreword

LLVM, winner of the 2012 ACM Software System Award, has become an integral part of the software-development ecosystem for optimizing compilers, dynamic-language execution engines, source-code analysis and transformation tools, debuggers and linkers, and a whole host of programming-language and toolchain-related components. Now heavily used in both academia and industry, where it allows for rapid development of production-quality tools, LLVM is increasingly used in work targeted at high-performance computing. Research in, and implementation of, programming-language analysis, compilation, execution, and profiling has clearly benefited from the availability of a high-quality, freely-available infrastructure on which to build.

LLVM-HPC2016 is the third workshop specifically focusing on research making use of the LLVM compiler infrastructure in High-Performance Computing (HPC). The call for papers attracted five submissions, and after peer review by the program committee, four of these were accepted. The accepted papers cover a variety of topics focused around compiling for modern HPC architectures and extending LLVM's intermediate representation to further express parallelism. These are important topics for the growing application of LLVM to HPC.

We thank the authors of our selected papers, our program committee, our invited speakers, the ACM Special Interest Group on High Performance Computing, and all of the other people and organizations that have made this workshop possible.

<p>Hal Finkel LLVM-HPC2016 Chair Argonne National Laboratory</p>
