

Workshop Proposal

The First Workshop on LLVM Compiler and Tools for HPC (LLVM-CTH 20)

Dr. Johannes Doerfert, jdoerfert@anl.gov
Argonne National Laboratory, USA

Anja Gerbes, gerbes@fias.uni-frankfurt.de
Goethe University Frankfurt, Germany

Dr. Sameer Shende, sameer@cs.uoregon.edu
University of Oregon, USA

Dr. Jeremy Bennett, jeremy.benneet@embecosm.com
Embecosm, United Kingdom

February 13, 2020

1 Workshop Organizers

Johannes Doerfert is a researcher in the Argonne Leadership Computing Facility at the Argonne National Laboratory. He develops LLVM and Clang enhancements that enable compiler optimization for parallel programs. Johannes is also part of several ongoing efforts to make compiler software ready for exascale computing. He is an active member of the OpenMP Language Committee and already organized various LLVM related workshops and conferences, including the LLVM Performance Workshops @ CGO, and the EuroLLVM in 2017. Johannes received his Ph.D. in Computer Science from Saarland University in 2018.

Anja Gerbes works at the Center for Scientific Computing and is a member of Hessian Competence Center for High Performance Computing located at Goethe University in Frankfurt/Main. A considerable part of her job role is to develop a range of courses and resources to enable users to work with the cluster. In addition, she is doing a PhD at the German Climate Research Center in Hamburg as an external member. The main topic is Compiler Optimization in High-Performance Computing with an aim to improve weather forecasting and climate modeling. The goal of her PhD is to study the compiler for deficits in terms of performance when translating HPC applications and to understand the limitations of compilers in making the necessary optimizations. These insights can then be incorporated into the compiler for future automatic compiler optimization. Automatic program transformation using source-to-source instrumentation of parallel programs will prepare HPC applications for future performance analysis.

Sameer Shende serves as the Director of the Performance Research Laboratory at the University of Oregon and the President and Director of ParaTools, Inc. He serves as the lead developer of the TAU Performance System, Program Database Toolkit (PDT), HPCLinux, and Extreme-Scale Scientific SoftwareStack (E4S.io). His research interests include performance instrumentation, compiler optimizations, measurement, and analysis tools for HPC. He leads the SDK project for the Exascale Computing Project (ECP), in the Programming Models and Runtime (PMR) area at the University of Oregon to provide ECP Software Technology (ST) products in a container environment for HPC. He has served

as the chair of the Performance Measurement, Modeling, and Tools track at SC17 and the co-chair for the technical program at the ICPP 2018 conferences.

Jeremy Bennett founded in 2008 Embecosm. He is an expert on hardware modeling and embedded software development. Previously Dr. Bennett was Vice President of ARC International plc, following their acquisition of Tenison Design where he had been CEO and CTO. Dr. Bennett is author of the popular textbook, Introduction to Compiling Techniques (McGraw-Hill 1990, 1995, 2003) and holds an MA and PhD in Computer Science from Cambridge University.

2 Relevance and Impact of Workshop for ISC

The LLVM framework is a vast ecosystem surrounding a compiler core which enabled various advances in source-code tools, debuggers, linkers, and a whole host of programming-language and toolchain-related components. Similar to the very successful LLVM-HPC workshop at SuperComputing (SC), we want to provide a forum for LLVM-related discussions and talks from academia and industry at ISC.

3 List of tentative program/steering committee members

We do not have a dedicated program or steering committee. As this is the first time the workshop happening, the organizers send out invitations to potential speakers that work on new LLVM-based techniques and tools in the HPC field. As of this writing, ten talks are tentatively confirmed by their respective authors. As shown in the schedule below we have contributions from almost all HPC vendors and other topics relevant to HPC and LLVM.

4 Workshop Scope

The workshop is centered around HPC compiler vendors, HPC compiler research, and HPC application needs. HPC vendors have accepted our invitation to present updates on their LLVM-based compilers and toolchain-related products. Compiler researchers and users are talking about opportunities and are given a platform to communicate their needs and pain points directly to the (industry) compiler developers.

5 Format of Workshop

This workshop will feature invited talks from industry and academia as well as a panel discussions focusing on recent developments under the LLVM umbrella related to HPC.

6 Expected Outcome from the Workshop

We expect interesting talks from both academia and industry on novel HPC-related topics. These will spark discussions and collaborations not only during the Q/A sections of the workshop.

7 Strategies for advertise and attract attendees

The workshop, as all LLVM related events, will be advertised on the LLVM developers mailing lists (llvm-dev, cfe-dev, openmp-dev, etc.) and at preceding LLVM events, e.g., CGO'20 in San Diego and EuroLLVM'20 in Paris. Furthermore, we commonly utilize established mailing lists that distribute CFPs for compiler venues, e.g., `compilation-news@lists.gforge.inria.fr`.

8 Estimated Number of Attendees

Between 20 and 50. This is an estimate based on the number of attendees for LLVM-related workshops at similar venues, e.g., at CGO and SC.

9 Targeted audience

We expect compiler and tool developers, compiler and tools users, to attendees interested in updates on vendor compilers (based on LLVM).

10 Workshop Website

It will follow the common scheme used for other events, see <https://llvm.org/devmtg/>

11 Workshop Length

Full-day workshop (9:00 am - 6:00 pm)

12 Tentative Schedule

Time	Topic	Speaker	Affiliation
09:00 - 09:30	Welcome	Anja Gerbes & Johannes Doerfert	Goethe University / ANL
09:30 - 10:30	Keynote	tbd	tbd
10:30 - 11:00	Compiling three generations of Fortran with LLVM	Steve Scalpone	NVIDIA (PGI)
11:00 - 11:30	break		
11:30 - 12:00	The CLANG/LLVM based Cray Compiling Environment (CCE)	Luiz DeRose	HPE (Cray)
12:00 - 12:30	Intel(r) LLVM-based Compilers for Xe Accelerators	Xinmin Tian	Intel
12:30 - 13:00	Overview of the LLVM Loop Optimization Pipeline and How it Compares to the IBM XL Compiler	Kit Barton	IBM
13:00 - 14:00	break		
14:00 - 14:30	Hidden Gems: Exposing and exploiting hidden compiler optimization opportunities	Zbigniew Chamski	University of Bristol
14:30 - 15:00	LLVM OpenMP on AMDGPU	Jon Chesterfield	AMD
15:00 - 15:30	OpenMP in MLIR and Flang	Kiran Chandramohan	ARM
15:30 - 16:00	Benchmarking compiler optimizations on OpenMP performance	Giorgis Georgakoudis	LLNL
16:00 - 16:30	break		
16:30 - 17:00	ARM Scalable Vector Extension in LLVM	Will Lovett	ARM
17:00 - 17:30	Clang Tooling for Kokkos	David Hollman	SNL
17:30 - 18:00	Panel: Ask Us Anything	Panel	various
18:00 - 18:05	Closing Remarks	Anja Gerbes & Johannes Doerfert	Goethe University / ANL