

Flang: Developing an open source Fortran front-end for LLVM

ARM HPC User Group 2016

Doug Miles, PGI Compilers & Tools, NVIDIA Corporation, 14 November, 2016

PGI Fortran , C & C++ Compilers

Optimizing, SIMD Vectorizing, OpenMP

Accelerated Computing Features

OpenACC Directives

CUDA Fortran

Multi-Platform Solution

x86-64 and OpenPOWER CPUs,
Tesla and Radeon GPUs

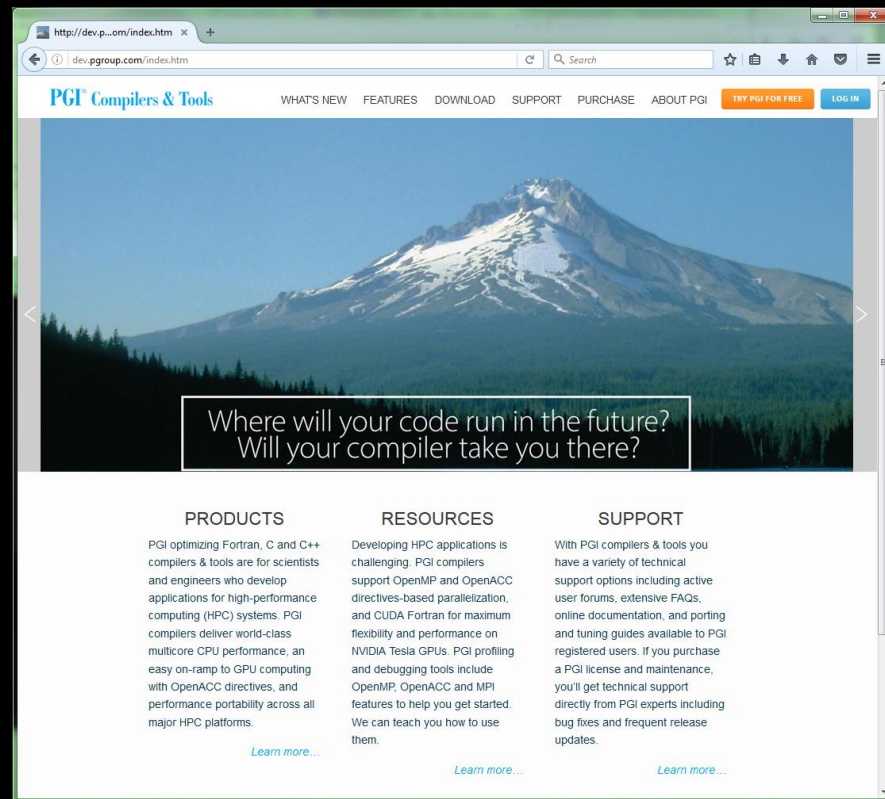
Supported on Linux, macOS, Windows

MPI/OpenMP/OpenACC Tools

PGDBG® debugger

PGPROF® profiler

Interoperable with DDT, Totalview



www.pgroup.com

LLVM: Community Power

Contributing Organizations

Processors

2000



10

15



2005



21

40

Google



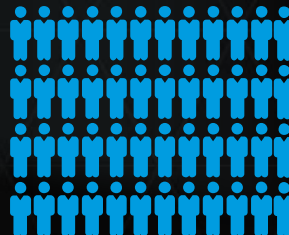
2010



178

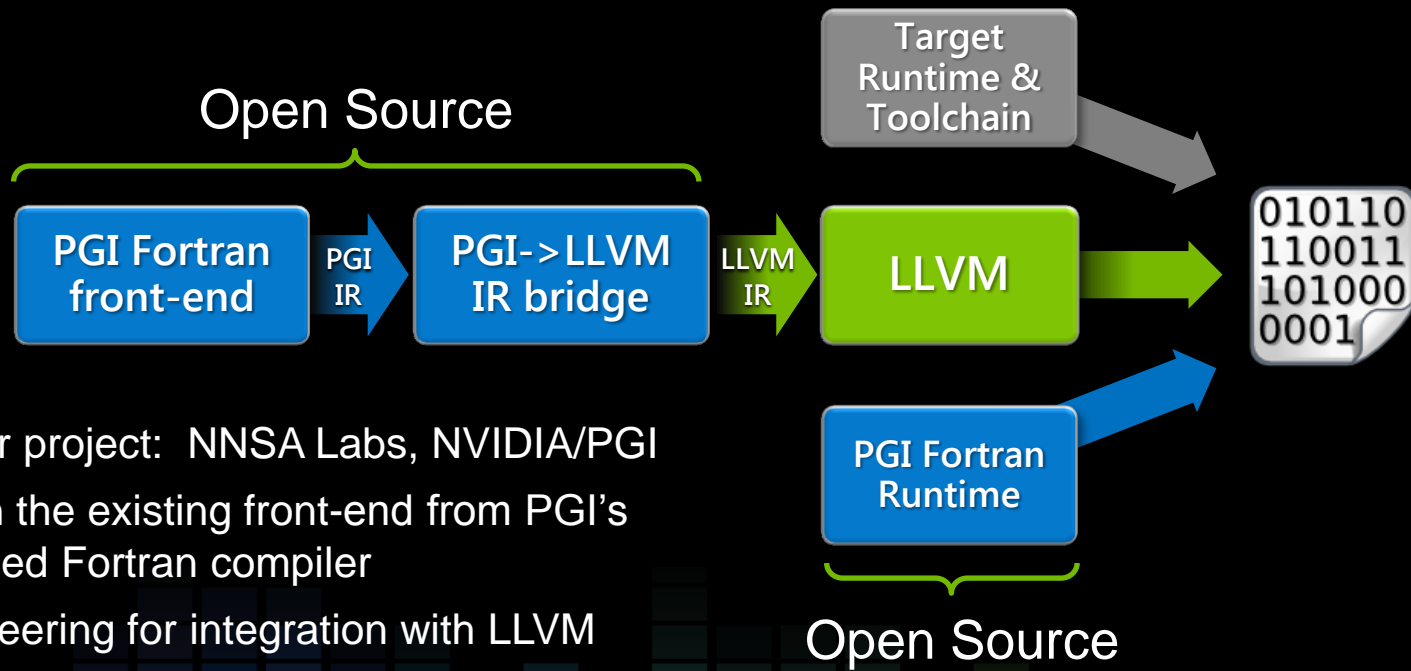
475

SAMSUNG Microsoft



475

An open source Fortran front-end for LLVM a.k.a. the Flang project



- Multi-year project: NNSA Labs, NVIDIA/PGI
- Based on the existing front-end from PGI's widely-used Fortran compiler
- Re-engineering for integration with LLVM
- Develop CLANG-quality Fortran msg facility

Many Stakeholders, Many Goals

LANL

New developer productive in source base in 4 – 8 weeks

Sandia

Single-thread/SIMD and OpenMP 3.1 performance

LLNL

OpenMP 4.x features, GPU and OpenPOWER support

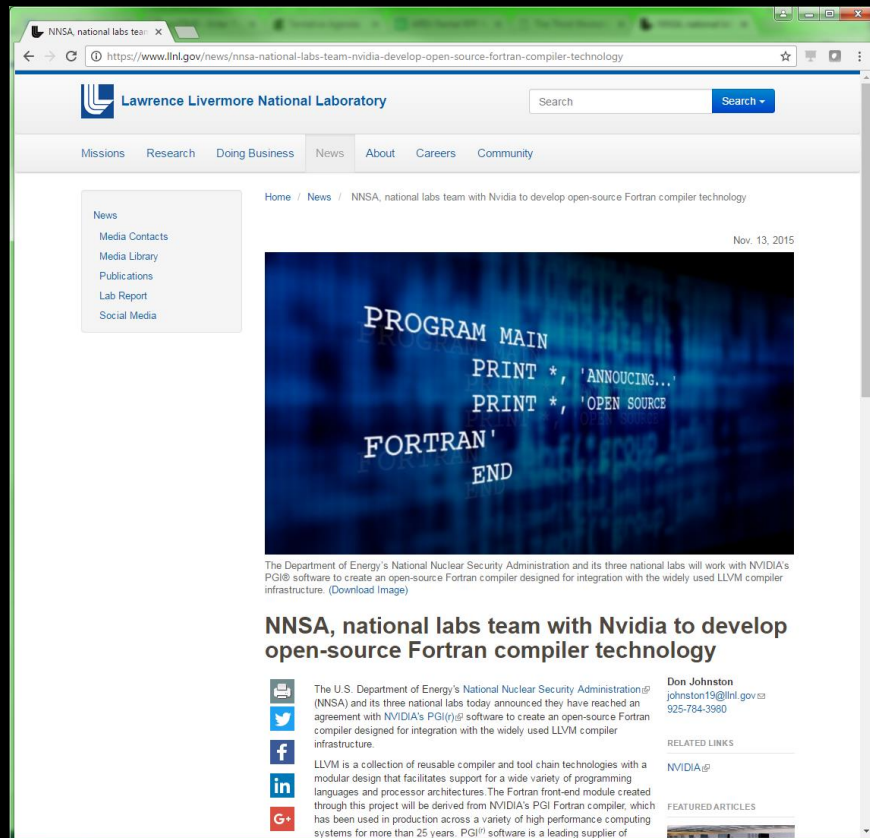
NVIDIA

Accelerate Fortran features support, PGI interoperability

Everyone

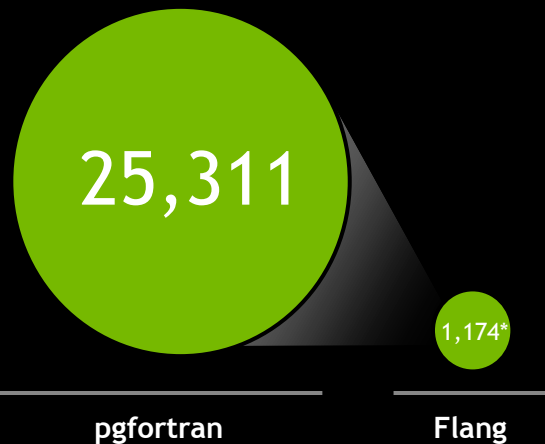
Adoption by both the HPC and LLVM communities

ANL, IBM, ARM Ltd, ORNL, Codethink, ...



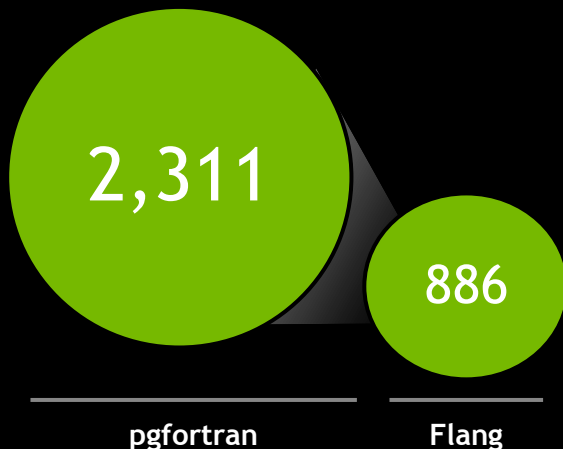
Creating the initial Flang source base

- Front-end 85%
- Runtime Libraries 15%



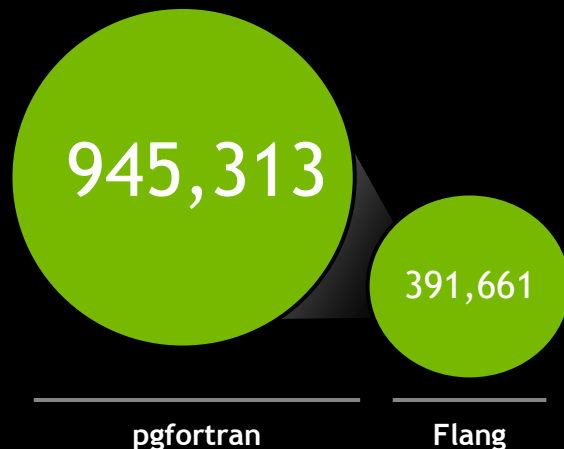
95% fewer #ifdefs

- Front-end 85%
- Runtime Libraries 15%



62% fewer files

- Front-end 80%
- Runtime Libraries 20%



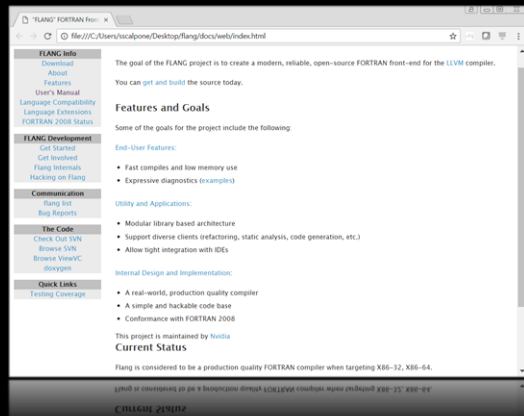
59% fewer LOC

*Clang has 212 #ifdefs in lib, include, tools

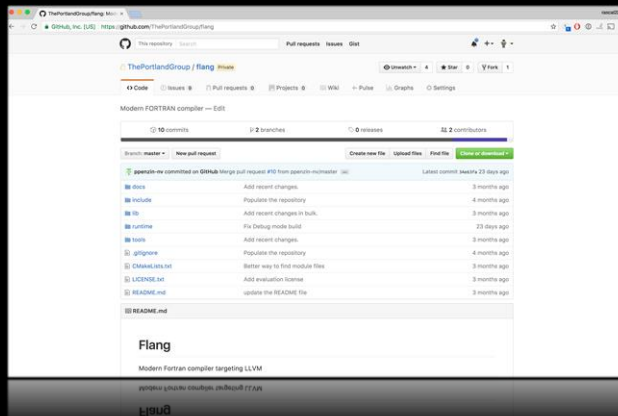
Flang Development Status

- Source code clean-up, refactoring & documentation ongoing
- Vendor neutrality nearly complete
- Frequent source and Flang binary updates to partners
- Passes most PGI Fortran Linux/x86 QA tests
- SIMD vectorization via the LLVM vectorizer, tuning ongoing
- Most of OpenMP 4.5 is implemented (CPU-side only)

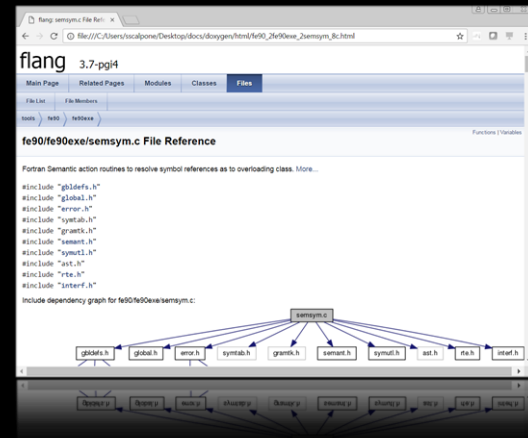
Flang Source Code



Home page



Github



Doxygen

Flang Single-core Performance

SPEC CPU 2006 Fortran codes, all times in seconds, 1 Haswell core

	PGI FORTRAN 16.10	GFORTTRAN 6.1	FLANG DEV LLVM 3.9
410.bwaves	182s	220s	251s
416.gamess	507s	Fails	475s
434.zeusmp	183s	221s	240s
436.cactusADM	165s	194s	208s
437.leslie3d	179s	209s	435s
454.calculix	171s	297s	608s
459.GemsFDTD	261s	286s	391s
465.tonto	295s	373s	Fails
481.wrf	157s	271s	247s

PGI Fortran: -fast -Mfprelaxed -Mstack_arrays gfortran: -O3 -funroll-loops -fpeel-loops -ffast-math Flang: -O3 -march=core-avx2 -ffp-contract=fast -Knoieee
Performance measured November, 2016 and are considered estimates per SPEC run and reporting rules. SPEC® and SPEC CPU® are registered trademarks of the
Standard Performance Evaluation Corporation (www.spec.org).

Flang OpenMP Performance

SPEC OMP 2012 Fortran codes, all times in seconds, 32 Haswell cores (64 threads)

	PGI FORTRAN 16.10	GFORTRAN 6.1	FLANG DEV LLVM 3.9
350.md	517s	3460s	459s
351.bwaves	469s	519s	805s
357.bt331	449s	492s	474s
360.ilbdc	541s	6846s	539s
362.fma3d	575s	504s	656s
363.swim	633s	634s	632s
370.mgrid	693s	697s	690s
371.applu	451s	414s	514s

PGI Fortran: -fast -mp -Mfprelaxed -Mstack_arrays gfortran: -O3 -funroll-loops -fpeel-loops -ffast-math -fopenmp

Flang: -O3 -mp -march=core-avx2 -ffp-contract=fast -Knoieee All: OMP_NUM_THREADS=64 OMP_PROC_BIND=true

Performance measured November, 2016 and are considered estimates per SPEC run and reporting rules. SPEC® and SPEC OMP® are registered trademarks of the Standard Performance Evaluation Corporation (www.spec.org).

Flang Year 2 Development Plans

- Source code

- Continue source clean-up, refactoring, documentation
- Create repository and release as open source
- Deploy an open source testing infrastructure

- Features

- Enhance compile-time Fortran error/warning messages
- Incremental F08 and OpenMP 4.5 features
- LLVM enhancements to enable Fortran DWARF generation

- Performance

- Incremental, likely to be reactive after initial pass is done