OAK RIDGE National Laboratory

OpenMP

CI/CD Efforts for Validation, Verification, and Benchmarking OpenMP Implementations

Authors: Aaron Jarmusch, Felipe Cabarcas, Swaroop Pophale, Andrew Kallai, Johannes Doerfert, Luke Peyralans, Seyong Lee, Joel Denny, and Sunita Chandrasekaran

EXPERIENCE ORNL
MEET. EXPLORE. LEARN.

ORNL is managed by UT-Battelle, LLC for the US Department of Energy

U.S. DEPARTMENT OF ENERGY

UD

OAK RIDGE National Laboratory

UNIVERSITY OF OREGON

Lawrence Livermore National Laboratory

UNIVERSIDAD DE ANTIOQUIA

Background

Concluded in 2023

- All began with the Exascale Computing Project
- SOLLVE: Scaling OpenMP With LLVm for Exascale Performance and Portability
  - LLVM
  - Runtime
  - **Validation and Verification** (originally called the SOLLVE V&V)
  - Specification
  - Outreach and Hackathons

The V&V was used by OLCF for Frontier acceptance.

OpenMP Validation and Verification Testsuite

- A suite of feature centric unit tests for OpenMP features
- **Objectives**
  - Evaluate OpenMP support across systems
  - Keep application programmers informed
  - Additional check for developers and vendors
  - Provide working code examples
  - Where to find it
    - https://github.com/OpenMP-Validation-and-Verification/OpenMP_VV

Currently funded by S4PST: Next Generation Science Software Technologies Project. https://s4pst.org/ (DE-FOA-0003177)

Motivation

- Continuous integration (CI) and continuous delivery (CD):
  - automation
  - increase frequency of testing (hourly runs)
  - catch bugs quicker, especially regression bugs
  - provide results for curated and trusted tests and benchmarks

CI/CD Workflow

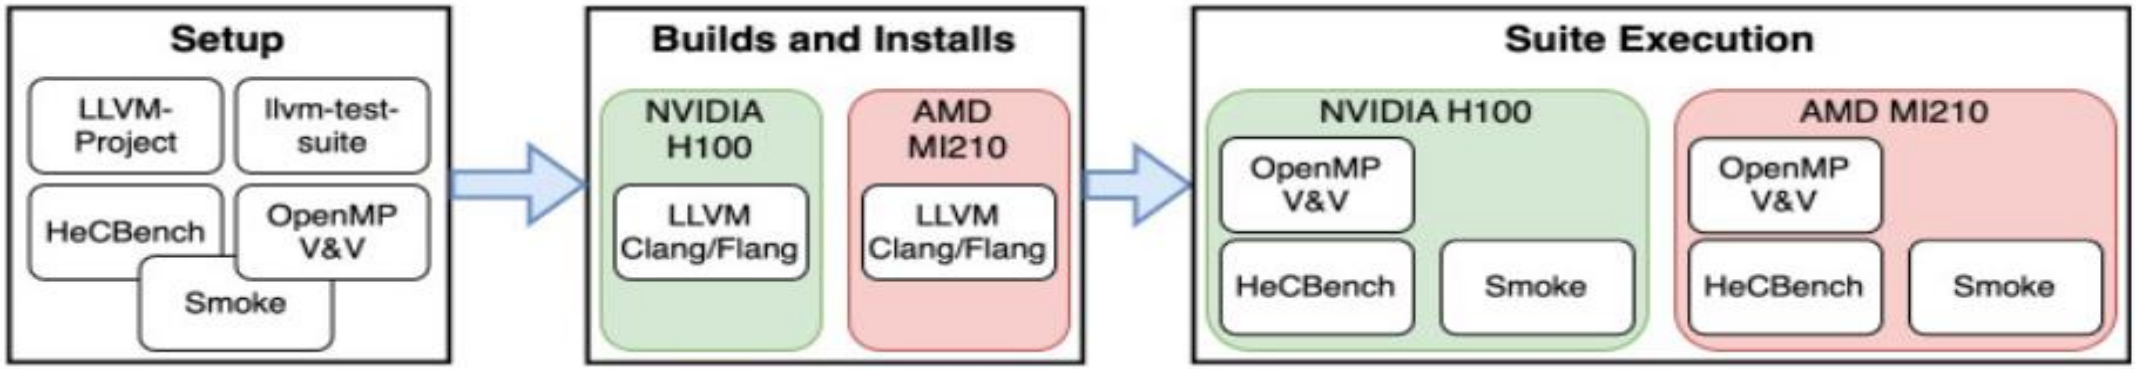

Terminology

- **workflow** has a set of **stages**; here Setup, Build & Install, and Execution
- every stage is a set of **jobs**
- **pipeline** refers to a single run of our workflow

CI/CD Stages

- Setup
  - Clone required resources
    - LLVM (latest from trunk) and LLVM test-suite infrastructure
    - OpenMP Validation and Verification Suite, Smoke, and HeCBench
- Build and Install
  - All mentioned above + SPEChpc
- Execute
  - OMP V&V, Smoke and HecBench - **hourly**
  - SPEChpc - **weekly**

SPEChpc HPG Benchmarking Suite

Benchmark Name	Size: Tiny	Language	Application Area
LBM D2Q37	505.lbm	C	Computational Fluid Dynamics
SOMA	513.soma	C	Polymeric Systems
Tealeaf	518.tealeaf	C	High Energy Physics
Cloverleaf	519.clvleaf	Fortran	High Energy Physics
Minisweep	521.miniswp	C	Nuclear Engineering
POT3D	528.pot3d	Fortran	Solar Physics
SPH-EXA	532.sph_exa	C++14	Astrophysics and Cosmology
HPGMG-FV	534.hpgmgfv	C	Cosmology, Astrophysics
miniWeather	535.weather	Fortran	Weather

- Smoke
  - Developed by AMD, these are primarily used to test AMD GPUs
  - We run them on the AMD and NVIDIA GPUs to compare the results
  - https://github.com/ROCm/aomp
- HecBench
  - A comprehensive collection of benchmarks for heterogeneous computing
  - Include benchmarks beyond OpenMP target offloading
  - https://github.com/zjin-lcf/HeCBench

```
7258  size..got: 40
7259  size..got.plt: 112
7260  size..init: 27
7261  size..init_array: 16
7262  size..interp: 28
7263  size..llvm.offloading: 9856
7264  size..note.ABI-tag: 32
7265  size..note.gnu.property: 32
7266  size..plt: 192
7267  size..plt.got: 8
7268  size..rela.dyn: 600
7269  size..rela.plt: 264
7270  size..rodata: 233
7271  size..text: 1458
7272  size.omp_offloading_entries: 96
7273  **********
7274  Testing Time: 34.43s
7275    Excluded:   7
7276    Passed  : 172
7277  + save amd-results3.json
7278  Running after_script
7279  Running after script...
```

Setup 1: Pipeline Setup for upstream LLVM OpenMP Offloading

University of Oregon - Gilgamesh

- CPU: AMD Epyc Milan 7413
- Accelerators:
  - 2 x AMD MI-200
  - 1 x NVIDIA H100

Pipeline Setup

- Hourly
- https://gitlab.e4s.io/uo-public/llvm-openmp-offloading/-/pipelines

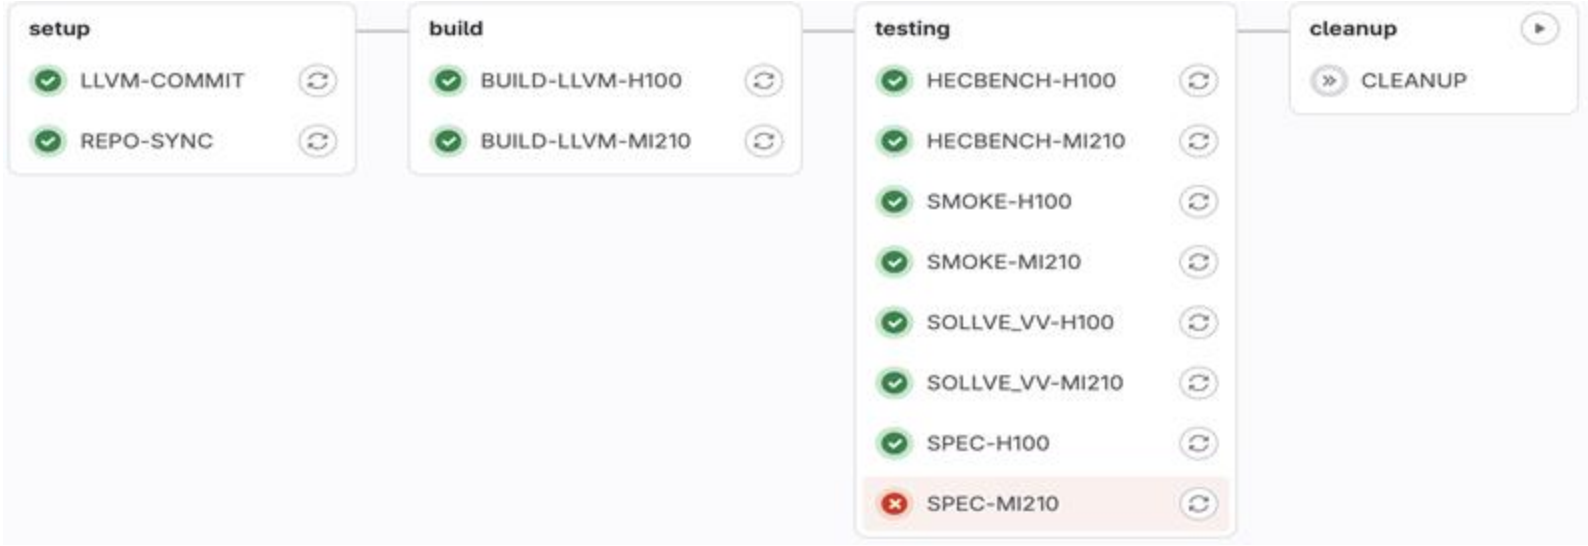

9

Community Reporting - LLVM

- Reporting
  - Failures are reported via e-mail and slack
- Additional analysis
  - investigation may be needed to rule out system issues
  - failures due to a new commit, once confirmed, are reported on the PR

Other Systems Tested

System	Vendor	Accelerator	Compiler and Versions
Perlmutter	HPE	NVIDIA A100	NVIDIA 24.5, Cray 17.0.0, LLVM 19.0.0 commit (18ec885a), and GNU 14.1
Frontier	HPE	AMD MI250X	AMD's ROCm 6.0.0, Cray 17.0.0, LLVM 19.0.0 commit (18ec885a), and GNU 14.1
Sunspot	Intel	Intel GPU Max Series	OneAPI 18.0.0

Summary of OMP V&V Results

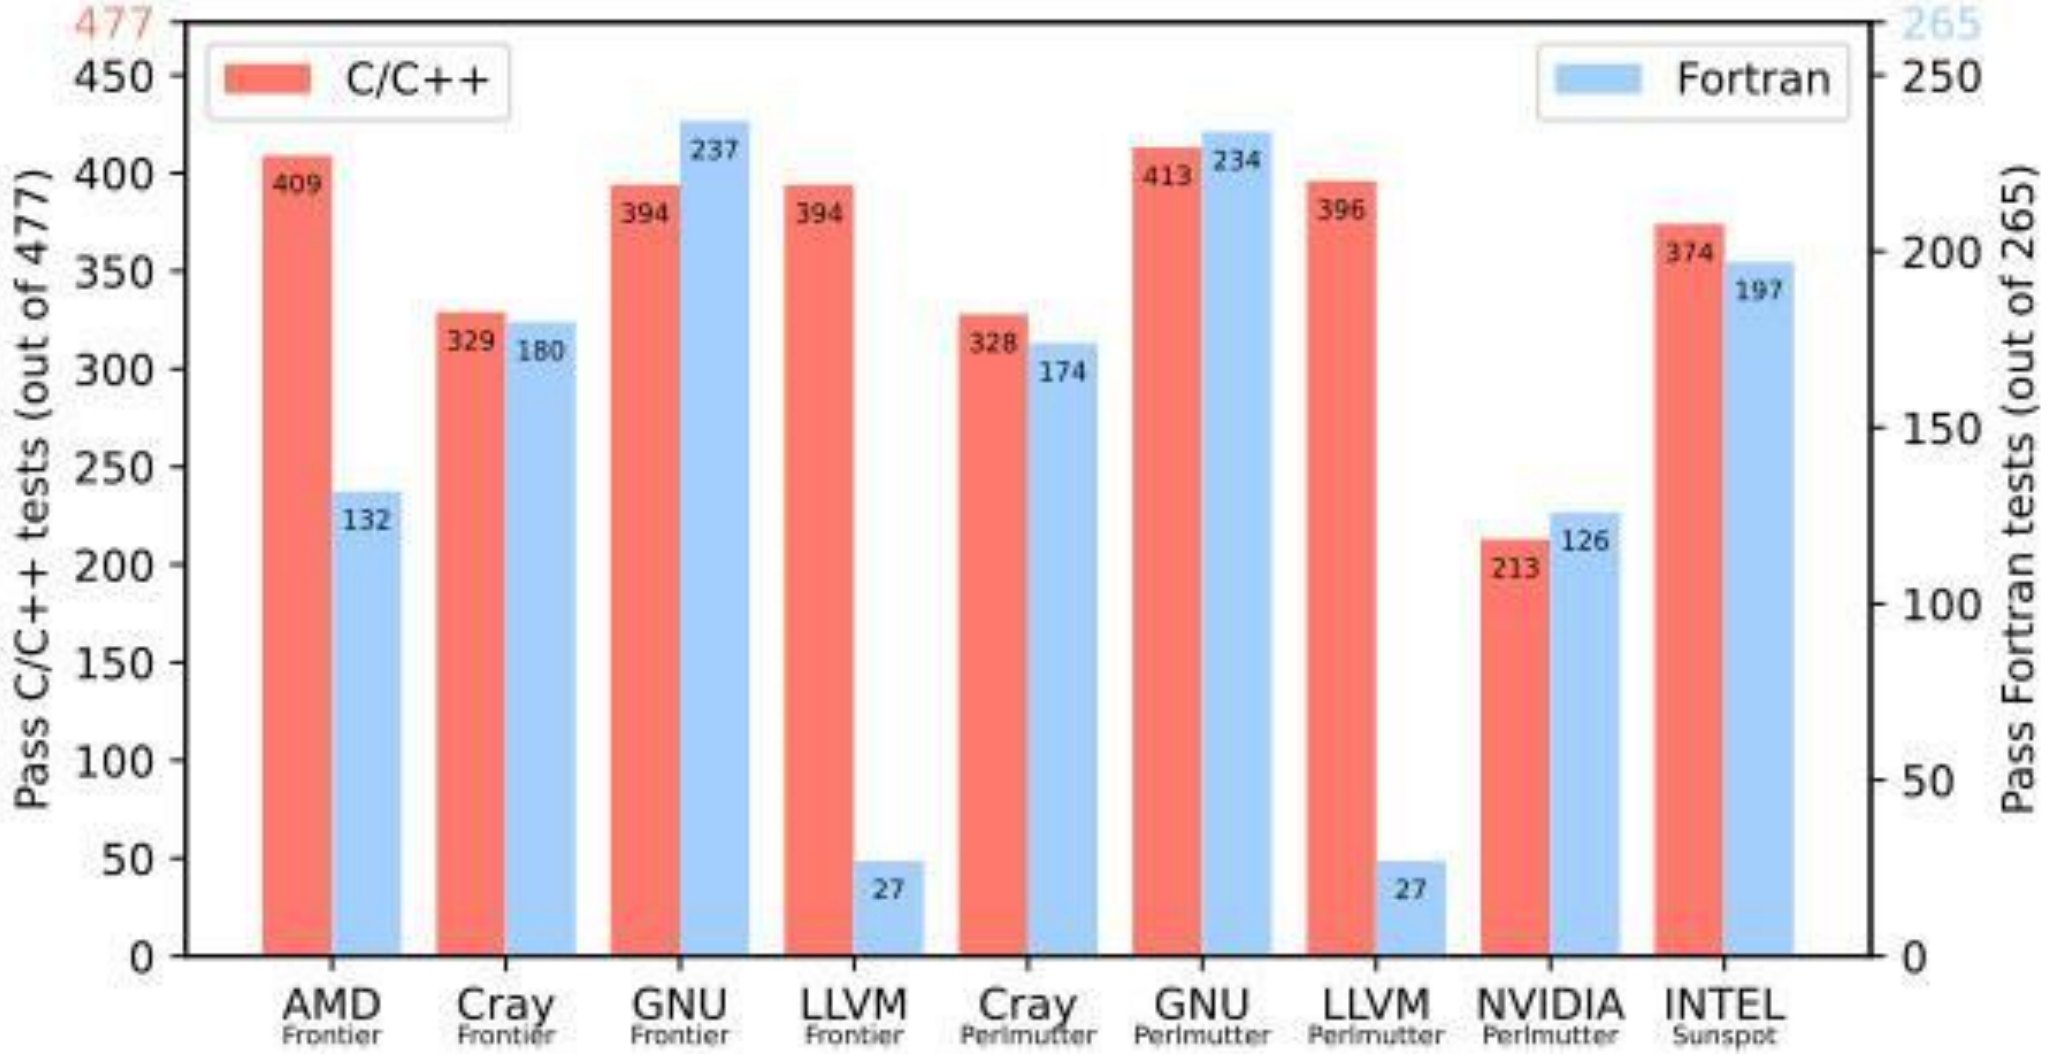

SPEChpc HPG Benchmarking Suite Results

SPEChpc Results Estimate Base Time (seconds)									
	OpenMP								OpenACC
	Frontier				Perlmutter				
Compiler	**GNU**	**LLVM**	**Cray cce**	**ROCm**	**GNU**	**LLVM**	**Cray cce**	**Nvidia nvc**	**Nvidia nvc**
Commit/Version	**14.1.0**	**20/4eb97802**	**17.0.0**	**6.0.0**	**14.1.0**	**20/4eb97802**	**17.0.0**	**nvc 24.5**	**nvc 24.5**
505.lbm_t	2813.46446	43.117335	40.819769	54.642662	484.888265	39.047365	28.342833	35.895276	28.482348
513.soma_t	BE 2	88.75995	78.540314	70.046082	855.051911	69.972128	56.745078	65.643469	45.82468
518.tealeaf_t	337.123638	41.875657	40.706717	48.505303	2200.952798	91.707654	49.092477	40.491684	48.230906
519.clvleaf_t	BE 1	BE 4	58.754946	72.73099	BE 1	BE 4	EE 3	45.536804	35.691745
521.miniswp_t	EE 1	160.877619	93.366468	142.614606	EE 1	209.451968	96.755836	573.088305	52.377006
528.pot3d_t	BE 3	BE 4	46.423345	92.606593	926.238377	BE 4	55.335936	61.541239	53.575591
532.sph_exa_t	BE 2	203.335568	224.179829	207.403858	1454.46486	849.597139	EE 4	491.411615	129.084128
532.sph_exa_t_mod	BE 2	145.435232	163.263467	144.82898	5973.459466	165.879632	128.356885	EE 5	Not modified
534.hpgmgfv_t	BE 3	102.658218	99.793166	95.589992	EE 2	151.425221	71.19688	163.330139	64.270558
535.weather_t	2569.959735	BE 4	32.512782	53.189687	1391.838681	BE 4	38.505173	42.717578	37.233478

Evolution of compilers on Frontier and Perlmutter systems

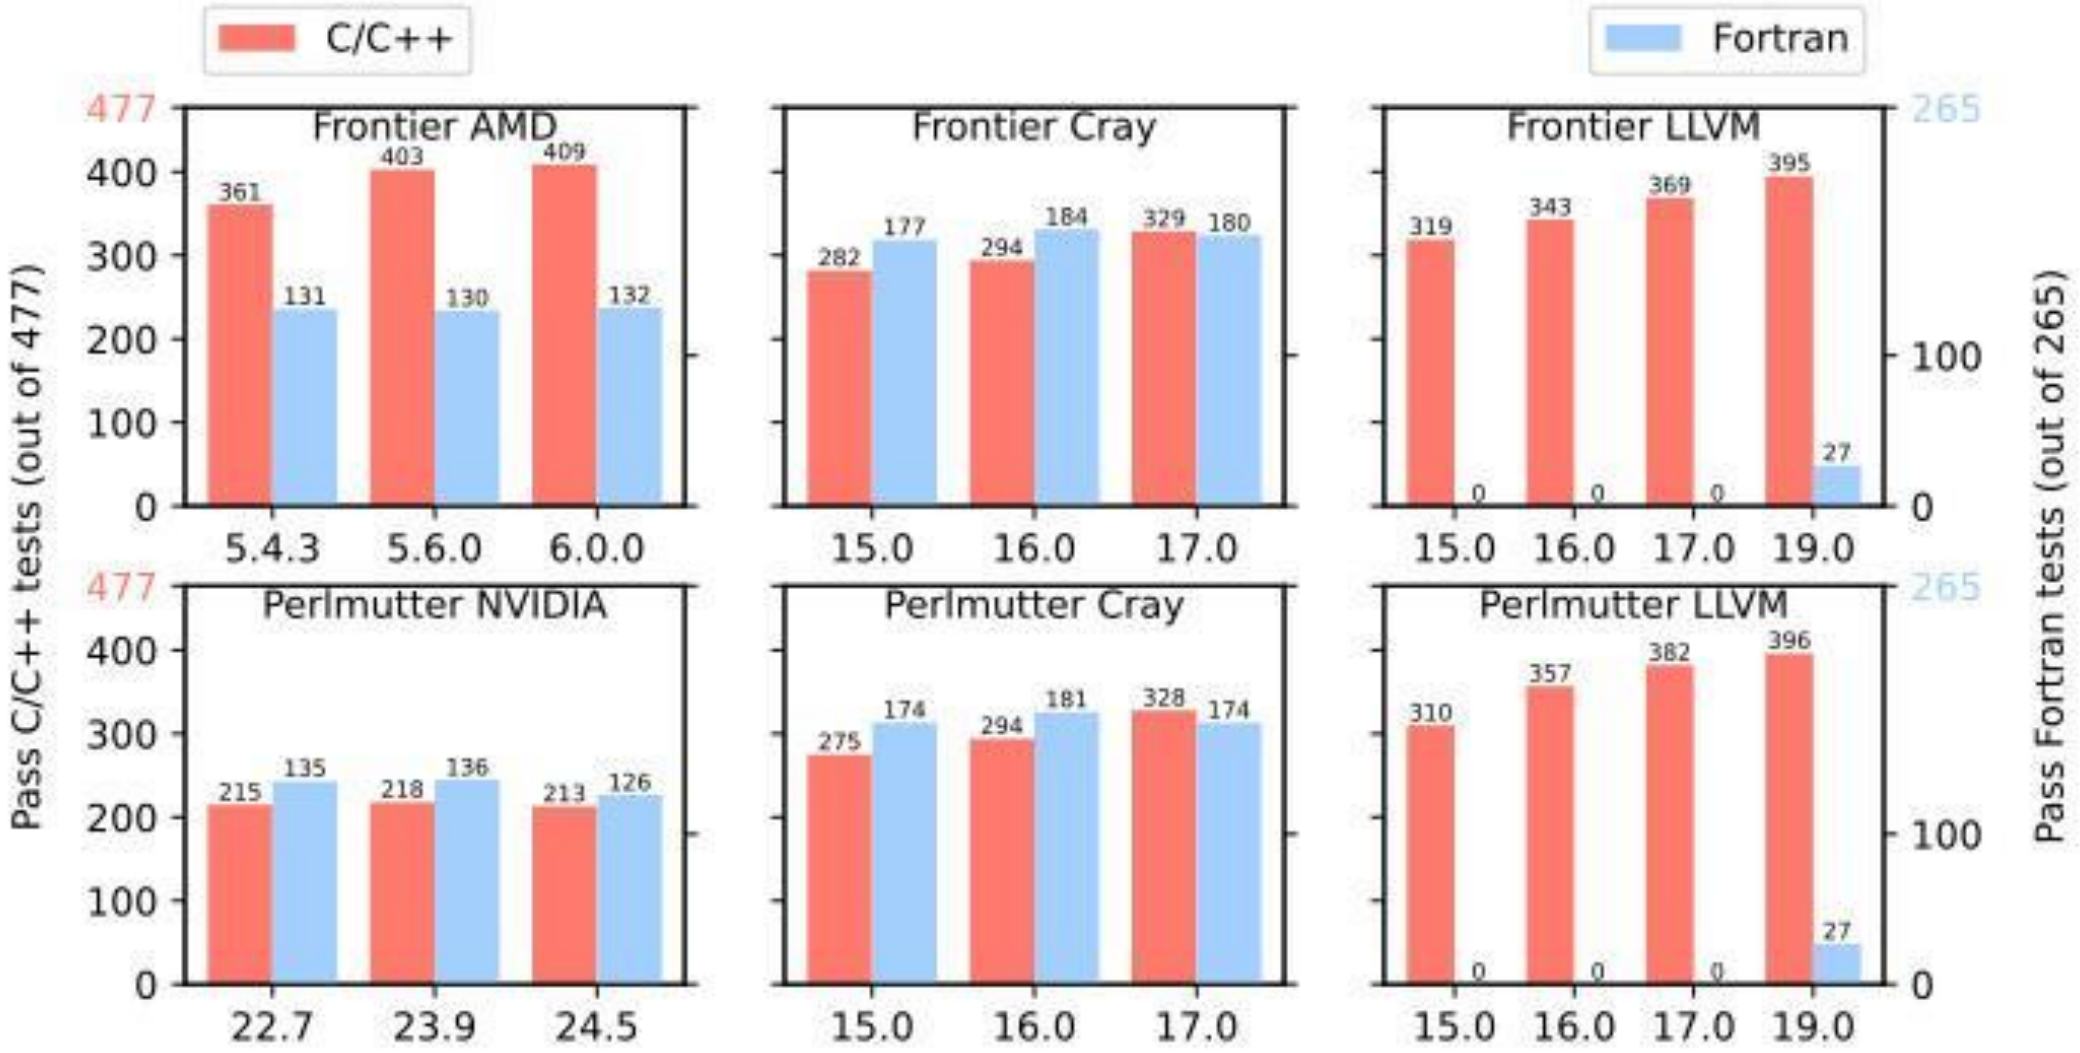

Takeaways

CI/CD Workflow

- We implemented a CI/CD workflow to capture bugs and accelerate feedback for compiler development.

Comprehensive Compiler Evaluation

- Our study evaluates OpenMP offloading support across compilers from AMD, HPE, GNU, LLVM, and Intel.

Goal

- Provide insights into the state of OpenMP offloading support and compiler quality for HPC.

Acknowledgment

The S4PST project is a collaboration across the US Department National Laboratories (ORNL, LANL, LLNL, LBNL, ANL), the University of Delaware and the Massachusetts Institute of Technology. The project is focused on the stewardship and advancement of programming systems for DOE. The project is funded by the U.S. DOE Office of Science, Advanced Scientific Computing Research (ASCR) program.

Work supported by the U.S. Department of Energy, Office of Science, the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration under contract number DE-AC05-00OR22725.