

# The Frontier Programming Environment at OLCF

**Wael Elwasif**

Frontier Programming Environment Task Lead  
Oak Ridge Leadership Computing Facility

Oak Ridge National Laboratory



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

# Overview



# Contributors to Frontier Programming Environment

## Vendor-Provided

- Cray Programming Environment (CPE)
  - Includes Cray compiler for C, C++, and Fortran plus GCC compiler. All the Cray profiling, tuning, and debugging tools. OpenMP and Cray MPI optimized for AMD GPU direct.
- AMD ROCm programming environment
  - Includes LLVM compiler to generate optimized code for both the AMD Trento CPU and MI250X GPU.
  - Support: C, C++, and Fortran and have GPU offload support. HIP, a CUDA-like direct GPU programming model (with CUDA to HIP conversion utilities).

## Other Sources

- ECP
  - LLVM enhancements: Flang (Fortran front-end), OpenMP, OpenACC
  - Kokkos and RAJA
  - HIP LZ (HIP support for Aurora)
  - OpenMPI, HPCToolkit, PAPI enhancements
  - ...
- ALCF + OLCF
  - Pilot implementation of DPC++/SYCL for Frontier
- OLCF
  - GCC enhancements to better support OpenACC, OpenMP, Fortran on Summit and Frontier

# Programming Environment

- Compilers Offered

- Cray PE (C/C++ [LLVM](#)-based; Cray Fortran)
- AMD ROCm ([LLVM](#)-based)
- [GCC](#)

Items in green are also available on Summit

- Programming Languages & Models Supported (in which compilers)

- [C](#), [C++](#), [Fortran](#) (all)
- [OpenACC](#) (Cray Fortran OpenACC 2.0+ & GCC 2.6 substantially complete, 2.7 planned)
- [OpenMP](#) (all) 5.0-5.2 in progress – most priority features complete, details vary
- [HIP](#) (Cray, AMD)
- [Kokkos/RAJA](#) (all)

- Transition Paths

- CUDA: semi-automatic translation to HIP : `hipify-perl`, `hipify-clang`
- CUDA Fortran: HIP kernels called from Fortran (a more portable approach)
  - CUDA Fortran kernels need to be translated to C++/HIP (manual process)
  - Fortran bindings to HIP and ROCm libraries and HIP runtime available through AMD's [hipfort](#) project

# Programming Tools

## Debuggers and Correctness Tools

Tool	
<i>System-Level Tools</i>	
Linaro DDT	forge/22.1.1
Cray CCDB	cray-ccdb/4.12.13 (D)
Cray ATP	atp/3.14.16 (D)
STAT	cray-stat/4.11.13 (D)
<i>Node-Level Tools</i>	
ROCgdb	rocm/X.Y.Z
Cray GDB4HPC	gdb4hpc/4.14.6 (D)

Items in green are also available on Summit

## Performance Tools

Tool	
<i>System-Level Tools</i>	
Linaro MAP/Performance Reports	forge/22.1.1
CrayPat/Apprentice2 (HPE)	perftools[-lite]
Reveal (HPE)	perftools[-lite]
<i>TAU</i>	
HPCToolkit	ums ums023 hpctoolkit/2023.03
Score-P / VAMPIR	vampir/10.3.0
<i>Node-Level Tools</i>	
PAPI	papi/6.0.0.17
AMD rocprof & libraries	rocm/X.Y.Z
AMD Ominperf	omniperf
AMD Omnitrace	omnitrace

# Scientific Libraries and Tools

Functionality	CPU	GPU	Notes
<b>BLAS</b>	<b>Cray LibSci, AMD BLIS,</b>	<b>Cray LibSci_ACC, AMD roc/hipBLAS, AMD rocAMD ROCm Tensile, MAGMA</b>	MAGMA is open source software led by the UTK Innovative Computing Laboratory
<b>LAPACK</b>	<b>Cray LibSci, AMD libFlame, PLASMA</b>	<b>Cray LibSci_ACC, AMD roc/hipSolver, MAGMA</b>	
<b>ScaLAPACK</b>	Cray LibSci	ECP SLATE, Cray LibSci_ACC	
<b>Sparse</b>		<b>AMD roc/hipSparse, AMD rocALUTION</b>	
Mixed-precision iterative refinement	Cray IRT, MAGMA	MAGMA	
<b>FFTW</b> or similar	<b>Cray, AMD, ECP FFTX, FFT-ECP</b>	<b>AMD rocFFT, ECP FFTX, FFT-ECP</b>	FFT-ECP focuses on 3D FFTs
<b>PETSc, Trilinos, HYPRE, SUNDIALS, SuperLU, ....</b>			Spack recipes from ECP E4S xSDK

Functionality in **green** is also available on Summit

# Useful Documentation Links

- HPE Cray Programming Environment
  - <https://cpe.ext.hpe.com/docs/>
- AMD ROCm docs
  - <https://rocm.docs.amd.com/en/latest/>
- AMD lab notes
  - <https://gpuopen.com/learn/amd-lab-notes/amd-lab-notes-readme/#amd-lab-notes>



# Digging a Little Deeper





# For C/C++ Codes

- Multiple compilers available
  - AMD
  - HPE (Cray)
  - LLVM
- But they're all based on LLVM
  - HPE and AMD are among the many organizations contributing to the development of LLVM
  - Most work is “upstreamed” (contributed to the core LLVM source)
    - But not everything is accepted (immediately), or may be held back as proprietary
  - Capabilities (and bugs) are likely to be generally similar at any point in time...
  - But not identical (***optimizations***) !

# Upstream LLVM @ OLCF

- Summit:
  - OLCF deployed modules (with offloading): latest **llvm/14.0.0**
  - Periodic **main** snapshots :  
`module use /sw/summit/modulefiles/ums/stf010/Core`  
`module load llvm/17.0.0-latest`    *# Also from specific dates*
- Frontier:
  - Periodic **main** snapshots (maintained by the ECP SOLLVE project)  
`module load ums ums012`  
`module load llvm/17.0.0-20230809`    *# Also from other dates*

# For Fortran Codes

- One useful compiler available at present
  - HPE/Cray
    - *Not* based on LLVM
- AMD provides a Fortran implementation, but we don't recommend it
  - It is based on “classic Flang”, in the LLVM ecosystem
  - Support for both the latest language standards and OpenMP offload are limited
- There is extensive work underway in the LLVM community on Flang, but it will be some time before it is production quality

# But What About GCC?

- On this slide “GCC” refers to the whole suite, including gfortran
  - With support for offloading using OpenMP/OpenACC
- OLCF is working with Siemens to implement OpenMP in GCC
- OLCF will provide recent release and development versions of GCC on Frontier
- For various reasons, you should *not* expect gcc-generated executables to be performant for offload at this time
  - Results will vary
  - We are interested in improving the performance of gcc. If you have a troublesome case, reach out to me. (No guarantees, however)
- GCC is also available on Summit

# GCC+offloading

- Summit:

- OLCF deployed modules (with offloading): latest **gcc/12.1.0**
- Periodic development snapshots :  
`module use /sw/summit/modulefiles/ums/stf010/Core`  
`module load gcc/13.2.1-20230727`    *# Also from other dates*

- Crusher:

- Periodic development snapshots  
`module use /sw/crusher/ums/compilers/modulefiles`  
`module load gcc/13.2.1-20230727`    *# Also from other dates*

- Frontier :

- **TBD**

# For HIP (and CUDA) Codes

- HIP runs today on AMD and NVIDIA GPUs
- An ECP project is working on supporting HIP on Intel GPUs
- Recommend a one-time translation of CUDA codes to HIP and make the HIP version primary from then on
- Both Cray and AMD compilers support HIP
  - They both use the underlying AMD compiler & runtime
- [More on HIP available in the OLCF Training Archive](#)
- HIP is also available on Summit

- AMD provides tools to translate CUDA to HIP
  - **hipify-perl** and **hipify-clang**
  - Not fully automatic
- **hipfort**
  - <https://github.com/ROCmSoftwarePlatform/hipfort>
  - Fortran bindings to HIP and ROCm libraries and HIP runtime
  - Build depends on ROCm version & Fortran compiler used
- Related: SYCLomatic - Intel tool to translate CUDA to Sycl
  - <https://github.com/oneapi-src/SYCLomatic>
  - Intel® DPC++ Compatibility Tool
  - Not fully automatic



# For OpenMP Codes

- OpenMP is very much a work in progress in the LLVM community
  - Most of 5.0 is implemented
  - Parts of 5.1, 5.2 are implemented
- We (DOE labs, including ORNL/OLCF) are trying to help prioritize the order of implementation based on what users tell us they need/want
  - So if you could really use features that aren't available yet, please let us know!
- *HPE/Cray and AMD compilers use different OpenMP runtimes*
- Remember that Cray Fortran is not based on LLVM
- OpenMP implementation in GCC is also a work in progress
- More on OpenMP available in the OLCF [Training Archive](#)

# For OpenACC Codes

- Cray Fortran supports OpenACC 2.0+
  - *"CCE supports full OpenACC 2.0 and partial OpenACC 2.x/3.x for Fortran (OpenACC is not supported for C and C++)"*
  - Work is underway to 3.2 (latest)
    - but no timeline has been given
- OLCF provides OpenACC support via GCC
  - 2.6 currently supported --- 2.7 planned
  - 3.x not currently planned – let us know if there are particular features that you could really use
  - Don't expect this to be performant at present
  - Not currently on Frontier
- Work is also underway in the LLVM community on OpenACC : **clacc**
  - **C/C++ only** – supported by the ECP PROTEAS-TUNE project
  - Development snapshot UMS modules : **module load ums ums025 clacc**

```
$> man intro_openacc
```

## **CCE OpenACC 2.x/3.x features – CCE/16:**

- **attach/detach** behavior and clauses
- **default(present)** clause
- Implied present-or behavior for **copy**, **copyin**, **copyout**, and **create** data clauses
- **if\_present** clause on **acc update**
- **if** clause on **acc wait**
- **async** and **wait** clauses on **acc data**
- **acc\_attach** and **acc\_attach\_async** APIs
- **finalize** clause on **exit data**
- **no\_create** clause on structured data and compute constructs
- **if** clause on **host\_data**

# What about SYCL?

- OLCF and ALCF have partnered with Codeplay on a pilot implementation of the Intel DPC++ compiler for AMD GPUs
  - ALCF has also partnered with NERSC on NVIDIA support
- Pilot implementation is complete
  - ~“50%” level of support
  - Tested with a small set of benchmarks and mini-apps
  - Crusher: **module load ums ums015 dpcpp**
  - Frontier: coming soon
- Seeking interested users to try out the pilot implementation
  - Provide feedback
  - Shake out issues
  - Provide motivation to complete the port

# Help Us Help You...

- If you have a liaison, work with them
- **If you encounter an issue, file a ticket with OLCF** – otherwise the facility won't (necessarily) know about it, and can't track it
  - Summit, Frontier...
- Take advantage of training events like this one
  - *Preparing for Frontier* series in the OLCF [Training Archive](#)
- OLCF office hours: <https://docs.olcf.ornl.gov/#olcf-office-hours>