

# AMD OpenMP offloading

## Lecture 9

Sunita Chandrasekaran  
Associate Professor, University of Delaware  
PDC Summer School, Aug 2023  
Slides courtesy AMD

# Getting Started with OpenMP® Offload Applications on AMD Accelerators

Presenter:  
Event:  
Date:

---

# Agenda

- 
1. AMD Open Source Software Stack: Recap
  2. Building Simple OpenMP® Offload Applications
  3. Hybrid MPI + OpenMP® Offload Applications Support
  4. Runtime Report for Performance Investigation
  5. Summary

A close-up, low-angle shot of an AMD Radeon Instinct graphics card. The card is dark with a prominent blue LED light strip along its edge. The words "RADEON INSTINCT" are printed in white, bold, sans-serif capital letters on the side of the card. The background is dark and out of focus, showing other components of a server or data center environment.

**RADEON INSTINCT**

# **AMD ROCm™ Software Stack: Recap**



# OPEN SOFTWARE PLATFORM FOR GPU COMPUTE



- Unlocked GPU Power To Accelerate Computational Tasks
- Optimized for HPC and Deep Learning Workloads at Scale
- Open Source Enabling Innovation, Differentiation, and Collaboration



## Compilers(1)

- ROCmCC and AOMP
  - ROCmCC provides support for both HIP and OpenMP®
  - AOMP: the AMD OpenMP® research compiler for AMD Instinct™ accelerators. It is used to prototype new OpenMP® features for ROCmCC
  
- GNU Compilers (GNU Compiler Collection - <https://gcc.gnu.org/wiki/Offloading>)
  - Provide offloading support to AMD GPUs (OpenMP®, OpenACC)
    - GCC 11 supports additionally:
      - AMD GCN MI100 (gfx908) GPUs.
    - GCC 13 (under development) adds support for AMD's Instinct MI200 (gfx90a) GPU series.
    - The devel/omp/gcc-12 (OG12) branch augments the GCC 12 branch with OpenMP and offloading features.

# Compilers(2)

- Siemens® Compilers (Sourcery CodeBench Lite – C/C++/Fortran)
  - Siemen's free GCC-based compilers
  - Supports all GCC 12 features, enriched by OpenMP® features from GCC's development branch and AMD GCN improvements such as support for offloading debugging.
  - Still under development.
- If you are on an AMD/HPE HPC system, there are additional options
  - Cray Compilers (HPE compilers)
    - Provide offloading support to AMD GPUs (OpenMP®, HIP, OpenACC)
- List of OpenMP Compilers & Tools :
  - <https://github.com/ROCm-Developer-Tools/aomp>
  - <https://www.openmp.org/resources/openmp-compilers-tools>

# Compilers(3)

	Cray			AMD			
Module	cce			ROCm™			
Language	C	C++	Fortran	C	C++	Fortran	HIP
Compiler	craycc	crayCC	crayftn	amdclang clang clang-cl	amdclang++ clang++	amdflang flang	hipcc

**hipcc**: wrapper for amdclang/amdclang++. It also makes sure to call either amdclang or nvcc based on the platform. Useful portability.

In case of porting CUDA codes to HIP, ROCm™ provides 'HIPification' tools to do the heavy-lifting  
Please consider using HIPify tools like Hipify-perl or Hipify-clang



# ROCm Libraries

ROCm Library	Note/comment
hipBLAS/rocBLAS	Basic Linear Algebra Subroutines
hipFFT/rocFFT	Fast Fourier Transfer Library
hipSPARSE/rocSPARSE	Sparse BLAS + SPMV
hipSolver/rocSolver	Lapack Library
rocALUTION	Sparse iterative solvers & preconditioners with Geometric & Algebraic MultiGrid
hipThrust/rocThrust	C++ parallel algorithms library
rocPRIM	Low Level Optimized Parallel Primitives
MIOpen	Deep learning Solver Library
hipRAND/rocRAND	Random Number Generator Library
EIGEN – HIP port	C++ template library for linear algebra: matrices, vectors, numerical solvers
RCCL	Communications Primitives Library based on the MPI equivalents

Latest status available at: <https://github.com/ROCm-Developer-Tools/HIP>

This talk focuses on simple and hybrid (MPI +) OpenMP Offload Application building and running on AMD GPUs