

AMD GPU development: random topics

A conversation with TImour.

Scalable Tools Workshop 2022

Overview

- ROCm Releases
- New Performance Counters
- Tools with AMD GPU Support
- User Difficulties
- Future

ROCm Releases

- ROCm releases are planned around start of quarters Q1..Q4
- ROCm-5.2.0 will be out this week or next.
- Feedback received from tools groups with release candidates before release.

- Main platform is ubuntu
- Secondary platform is RHEL/CentOS

- Non-AMD effort to make releases available for Fedora by single person.

How to better support (other) Distros

- Try it out and submit bug/issue on repeatable concerns

Release Testing

- AMD testing doesn't involve scaled-up testing, such as larger kernels and/or long runtimes.
- No or little CI style testing at this time. Limits expansion to more distros.
- AMD ROCm group – small staff, looking for more people to hire to expand capabilities.

- Research partners provide testing for release candidates
- Partners also provide testing for tools (Dyninst, HPCToolkit, E4S suite)

Issues / Trouble Reports / Bugs

- Create github issue
- <https://github.com/RadeonOpenCompute/ROCm/issues>
- Email Timour about the issue, including number & URL – and private notes.

External concerns with release process

- AMD doesn't interact with vendors to preview OS releases and updates to ensure they will work seamlessly on new versions.
- AMD could improve communication about new releases – currently non-partners see nothing until after a release occurs.
- Release candidates are public

- Bolo: Are there large or long running kernels, that partners could be testing as part of the RC process? Would it be a better validation of an installation or a RC?

Deprecation of GPU architectures? – Not really

- gfx9xx series (MI25 - MIXX) will remain available as a family.
- New GPU series with different CPU may cause deprecations in the far future.
- ROCm-5.2.0 still supports MI25, MI50 – and will continue to do so.

Open Source?

- Yes, the releases are open source on github
 - Repositories are publicly available, including release candidates.
 - Everything is supposed to be open and public, no API lingering under NDA.
-
- How to improve communication for open source development?

New Performance Counters

- MI200 for now just released, such as on crusher.
- Older GPUs have counters, and will be backfilled as effort permits.
- Difficulty is validating all the counters work as advertised
- Requires extensive testing for validation of counter validity.
- AMD team is small – they are looking for people!
- Miperf and documentation will help with counters
- Miperf and rocprof use same mechanism for counter access
- Counters released are those most demanded by users.
- Users: General concern that we need to understand what the counters really mean and their implications to integrate them into tools.

Tools with AMD GPU Support

- Dyninst
- HPCToolkit
- Score-P – release end of week, waiting for PAPI?
- Tau – roctracer / rocprofiler
- PAPI – GPU metrics under development
- caliber???

User Experience

- Difficulty attributing AMD GPU performance to code.
- Stability of computing environment, especially with larger and/or long duration GPU jobs.
- Performance concerns – slower than expected compared to nVidia
- No way to find which GPU kernels were used with openmp offloading.
- Hip but NOT OpenMP GPU architecture identification
- No PC sampling?
- AMD Fortran versions not compatible with HPC code bases.
- ompt to tie back to openmp – works w/ AMD compiler, callbacks work ... but sampling?
- Everyone have some problems with CI.

User Experience (continued)

- Need more metrics – working with center of excellence people to help understand that in their internal cluster.
- Large apps have high register pressures, don't have instrumentation to understand what is going on. Limited PAPI sampling on counter we have access to
- Apex implemented for rocprof, lacking metrics to understand what is going on.
- Tools need to work on larger apps. Crashes / deadlocks.
- Stability – RocProfiler interface tends to cause issues, not just running apps
- OMPT offloading working for some user apps.
- GPU Aware MPI on crusher allows keeping data on GPUs. However, network size limited to 64 nodes, can't understand more.
- Hard to keep data on GPU-only other than crusher noted above.
- Need interfaces for PC Sampling.

Things that are working in tool chains

- OMPT offloading working for some ECP pseudo apps

Future

- DWARF changes for AMD GPU– Tony Ty – Changes need to be chopped up, DWARF committee won't accept them as a whole.
- Are people giving issues and feedback to AMD / TImour?
- MI200 GPUs on way to tools partners soon.
- MIperf training session later this year
- Complain if something you want doesn't work (issues) AMD wants it to work for you!
-