




Newer Tools: ROCprof Trace Decoder and Omnistat

Presenter: Bob Robey
Oct 21-23, 2025
AMD @ Tsukuba University

AMD 
together we advance_

Role of ROCprof Compute Viewer and Trace Decoder?



- Rocprofv3 enables basic profiling and hardware counters
- Rocprof-sys gives a system level view including CPU, MPI, OpenMP® and many others.




- Rocprof-compute provides a detailed overview of the performance of individual GPU kernels

- But what has been missing is a tool that looks at performance *within* a kernel

- What lines of code in the kernel are taking the most time?
- Where are stalls happening?
- What parts of the kernel are consuming resources such as LDS memory?




ROCprof Compute Viewer – high-level view of kernels




Summary View

From <https://rocm.docs.amd.com/projects/rocprof-compute-viewer/>




Instructions View


ROCprof Compute Viewer – Timeline views




Wave States Timeline



Hotspot Timeline



Compute Unit Timeline





Occupancy Timeline

From <https://rocm.docs.amd.com/projects/rocprof-compute-viewer/>

Instructions view – detail exploration

- We begin with the instructions view in the Rocprof Compute Viewer.
- The lines of code from the kernel are shown on the right.
- Just to the left of the kernel line is a bar indicating the resources used for that line.
- The assembler ISA for the source lines is shown on the left. This can quickly show whether a wide load instruction has been used or what are the compute instructions generated by the compiler. Also shown is the hit count and latency for each line
- Clicking on a source code or ISA line will reorient the other window to the corresponding line and left align the top timeline view of the compute unit.

Color coding for blocks



From <https://rocm.docs.amd.com/projects/rocpof-compute-viewer/>


Quickstart Usage Walkthrough

Generating trace decoder data

- Set up the environment – load the rocprofiler-sdk or rocm module depending on your system
module load rocprofiler-sdk
- Build the HIP code that you want to profile
cd ../HIP/vectorAdd
make vectoradd
- Generate the trace decoder data
rocprofv3 --att --att-activity 16 -d tracedecoder_vectorAdd -- ./vectoradd
- CLI Options
 - --att required to turn on trace decoder
 - --att-activity needed for summary (and maybe source code?). Frequency number is from 1 to 32.
 - Compile debug to get the source code included
- Tar up the directory specified for the data in the previous command
tar -czvf tracedecoder_vectorAdd.tgz tracedecoder_vectorAdd
- Copy the files to your local system
scp <hpc_system>:tracedecoder_vectorAdd.tgz tracedecoder_vectorAdd.tgz
- Clean up afterwards
make clean
rm -rf tracedecoder_vectorAdd

ROCprof Compute Viewer


- Untar the data that you have copied to your system
`tar --xzvf tracedecoder_vectorAdd.tgz`
- Now start up ROCprof Compute viewer on your local system
- Use Import at the very top left and select one of the ui_output_agent* files in the tracedecoder_vectorAdd directory. Center the source code window, select vectorAdd.hip and scroll down to find the kernel.



- Select line 58 `a[i] = b[i] + c[i];` and see what lines get highlighted in the left window.
- There is a very small rectangle on the left of line showing resource usage – we are not using much in this kernel.
- We compiled with `-g -O2`, so the lines may not be exact.


Summary View

- Click on the Summary tab.
- We can see that we are mostly using VALU (Vector Arithmetic Logic Unit) instructions
- We are also mostly seeing Stall/Wait states
- We can conclude that we are memory bandwidth bound and that we could be doing a lot more compute work in this kernel.
- Based on this, let's investigate the wave state timeline to better understand what is happening.




Wave States Timeline

- Click on the Wave States tab.
- We see that it is mostly waits rather than stalls. This tells us that we are likely just waiting on data loads.
- Compute (Exec) is low indicating that we could do more work.
- Let's jump to the Hotspot view to see if we can learn something there.



Hotspot Timeline

- We see an initial memory request and then fairly sparse activity.
- With a more complex kernel, we should see where LDS usage, matrix and other instruction types occur.
- Let's switch to the Compute Unit view to get a big picture of what is going on at that level.



Compute Unit Timeline

- Click on Compute Unit. Then go over to the left control panel and change the WaveView zoom from 10 to 5 so that we can see the whole timeline.
- We clearly see the initial memory load requests in yellow and then scattered VALU instructions in light blue.

Input

Plots

Options

UI Path:

Shader

SIMD

Slot

WaveID

0


0

0

0

GlobalView zoom: 10



WaveView zoom: 5



Installation Instructions

Installing ROCprof Compute Viewer

- Check for releases of HPC Compute Viewer
- <https://github.com/ROCm/rocprof-compute-viewer/releases>
- For Windows
 - , there is a pre-built version for download
 - Uses a Windows setup tool
- To build from source, see
 - <https://rocm.docs.amd.com/projects/rocprof-compute-viewer/en/amd-mainline/install/installation.html#install-viewer>
- Linux®, Mac OS, Windows Subsystem for Linux (WSL) and Windows® native
- Uses Qt 5 or 6



Setting up Trace Decoder

Nothing needed on ROCm 7.0+ !!

Trace Decoder is a part of ROCprofv3

- We have installed on ROCm version 6.4.1 with some effort.
 - Install recent rocprofiler-sdk from source to a different location (/opt/rocmplus-6.4.1/rocprofiler-sdk)
 - Install aqlprofile to same location
 - Install tracedecoder binary to same location
 - Set paths to the new rocprofiler-sdk directory before the standard ROCm location
 - Need to override default gfx codes since some do not exist before ROCm 7.0
 - Install to prior ROCm versions may break in the future
 - Cannot install prior to ROCm 6.2
- May consider using a ROCm 7.0 container to get early access as well

Setup pre-7.0 ROCm version step 1

- Requires system have libdw installed
- Setup script with these steps at
 - `git clone https://github.com/AMD/HPCTrainingDock`
 - File: `HPCTrainingDock/tools/scripts/rocprofiler-sdk_setup.sh`
 - Check script for latest instructions
- Set `INSTALL_PATH` to `/opt/rocmplus-6.4.1/rocprofiler-sdk` or other directory
- Install trace decoder

```
wget https://github.com/ROCm/rocprof-trace-decoder/releases/download/0.1.2/rocprof-trace-decoder-manylinux-2.28-0.1.2-Linux.tar.gz
tar -xzf rocprof-trace-decoder-manylinux-2.28-0.1.2-Linux.tar.gz
mv rocprof-trace-decoder-manylinux-2.28-0.1.2-Linux/opt/rocm/lib/librocprof-trace-decoder.so
$INSTALL_PATH/lib
rm -rf rocprof-trace-decoder-manylinux-2.28-0.1.2-Linux
```

Setup pre-7.0 ROCm version step 2

- Install rocm-profiler-sdk
- AMDGPU_GFXMODEL is set to “gfx90a;gfx942” to override default gfx model list. Use what you need here.

```
git clone -b develop https://github.com/ROCm/rocm-systems.git rocm-systems-source
```

```
cmake -B rocm-profiler-sdk-build \
      -DCMAKE_INSTALL_PREFIX=${INSTALL_PATH} \
      -DGPU_TARGETS=\"${AMDGPU_GFXMODEL}\" \
      -DOPENMP_GPU_TARGETS=\"${AMDGPU_GFXMODEL}\" \
      -DCMAKE_PREFIX_PATH=/opt/rocm-${ROCM_VERSION} \
      rocm-systems-source/projects/rocm-profiler-sdk
```

```
cmake --build rocm-profiler-sdk-build --target all --parallel 8
cmake --build rocm-profiler-sdk-build --target install
```

```
rm -rf rocm-profiler-sdk-source rocm-profiler-sdk-build
```

Setup pre-7.0 ROCm version step 3

- Install aqlprofile

```
cmake -B aqlprofile-build \
      -DGPU_TARGETS="${AMDGPU_GFXMODEL}" \
      -DCMAKE_PREFIX_PATH=/opt/rocm-${ROCM_VERSION}/lib:/opt/rocm-${ROCM_VERSION}/include/hsa \
      -DCMAKE_INSTALL_PREFIX=$INSTALL_PATH \
      rocm-systems-source/projects/aqlprofile
```

```
cmake --build aqlprofile-build --target all --parallel 8
```

```
cmake --build aqlprofile-build --target install
```

```
rm -rf rocm-systems-source aqlprofile-build
```

Omnistat

Omnistat

- Omnistat is primarily for HPC cluster administrators
- Has both a system-wide and a user-mode
- Metrics include:
 - GPU utilization
 - High-bandwidth memory (HBM) usage
 - GPU power
 - GPU temperature
 - GPU clock frequency
 - GPU memory clock frequency
 - Inventory information:
 - ROCm driver version
 - GPU type
 - GPU vBIOS version
- Optional metrics
 - RAS information (error counts per GPU block)
 - GPU power caps
 - GPU throttling events
 - Host network traffic (received/transmitted)

Sample output – GPU Usage Summary



Omnistat: Where to get more information

- AMD Research project and is not part of the ROCm software stack.
- Online documentation
 - <https://rocm.github.io/omnistat>
- Source code
 - <https://github.com/ROCm/omnistat>
- Uses either Prometheus client and server or VictoriaMetrics server for scalable collection of metrics
- Uses Grafana for visualizing data

Disclaimer

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED 'AS IS.' AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

AMD, the AMD Arrow logo, ROCm and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

© 2025 Advanced Micro Devices, Inc. All rights reserved.

Git and the Git logo are either registered trademarks or trademarks of Software Freedom Conservancy, Inc., corporate home of the Git Project, in the United States and/or other countries

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

The OpenMP name and the OpenMP logo are registered trademarks of the OpenMP Architecture Review Board

Windows is a registered trademark of Microsoft Corporation in the US and/or other countries.

