



ATPESC 2021 – Tools Track

ROC-profiler and debugger: An Overview of ROCm Tools

Suyash Tandon, Justin Chang, Julio Maia, Noel Chalmers, Paul T. Bauman, Nicholas Curtis,
Nicholas Malaya, Damon McDougall, Rene van Oostrum, Noah Wolfe



Agenda

1

ROC-profiler

2

ROC-debugger

*More information can be found in the official [ROCm documentation](#)

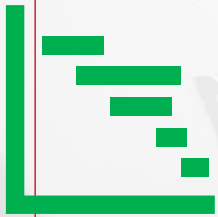
A close-up, low-angle shot of an AMD Radeon Instinct GPU. The card is black with a prominent silver-colored metal grille on the left side. The words "RADEON INSTINCT" are printed in white, bold, sans-serif capital letters on the black surface of the card. The background is dark and out of focus, showing other components of a server or data center environment.

RADEON INSTINCT

Profiling



AMD GPU Profiling



ROC-profiler (or simply rocprof) is the command line front-end for AMD GPU profiling library

- Repo: <https://github.com/ROCm-Developer-Tools/rocprofiler>



rocprof contains the central components allowing the collection of application tracing and counter collection

- Under constant development



Provided in the ROCm releases



The output of rocprof can be visualized using the chrome browser with chrome tracing

Getting started with rocprof

To get help:

- `$ /opt/rocm/bin/rocprof -h`

Useful housekeeping flags:

- `--timestamp <on|off>` : turn on/off gpu kernel timestamps
- `--basenames <on|off>`: turn on/off truncating gpu kernel names (i.e., removing template parameters and argument types)
- `-o <output csv file>`: Direct counter information to a particular file name
- `-d <data directory>`: Send profiling data to a particular directory
- `-t <temporary directory>`: Change the directory where data files typically created in /tmp are placed. This allows you to save these temporary files

Flags directing rocprofiler activity:

- `-i <input.txt|.xml>`: specify an input file (note the output files will be named input.*)
- `--hsa-trace`: to trace GPU Kernels, host HSA events (more later) and HIP memory copies.
- `--hip-trace`: to trace HIP API calls
- `--roctx-trace`: to trace roctx markers

Advanced usage

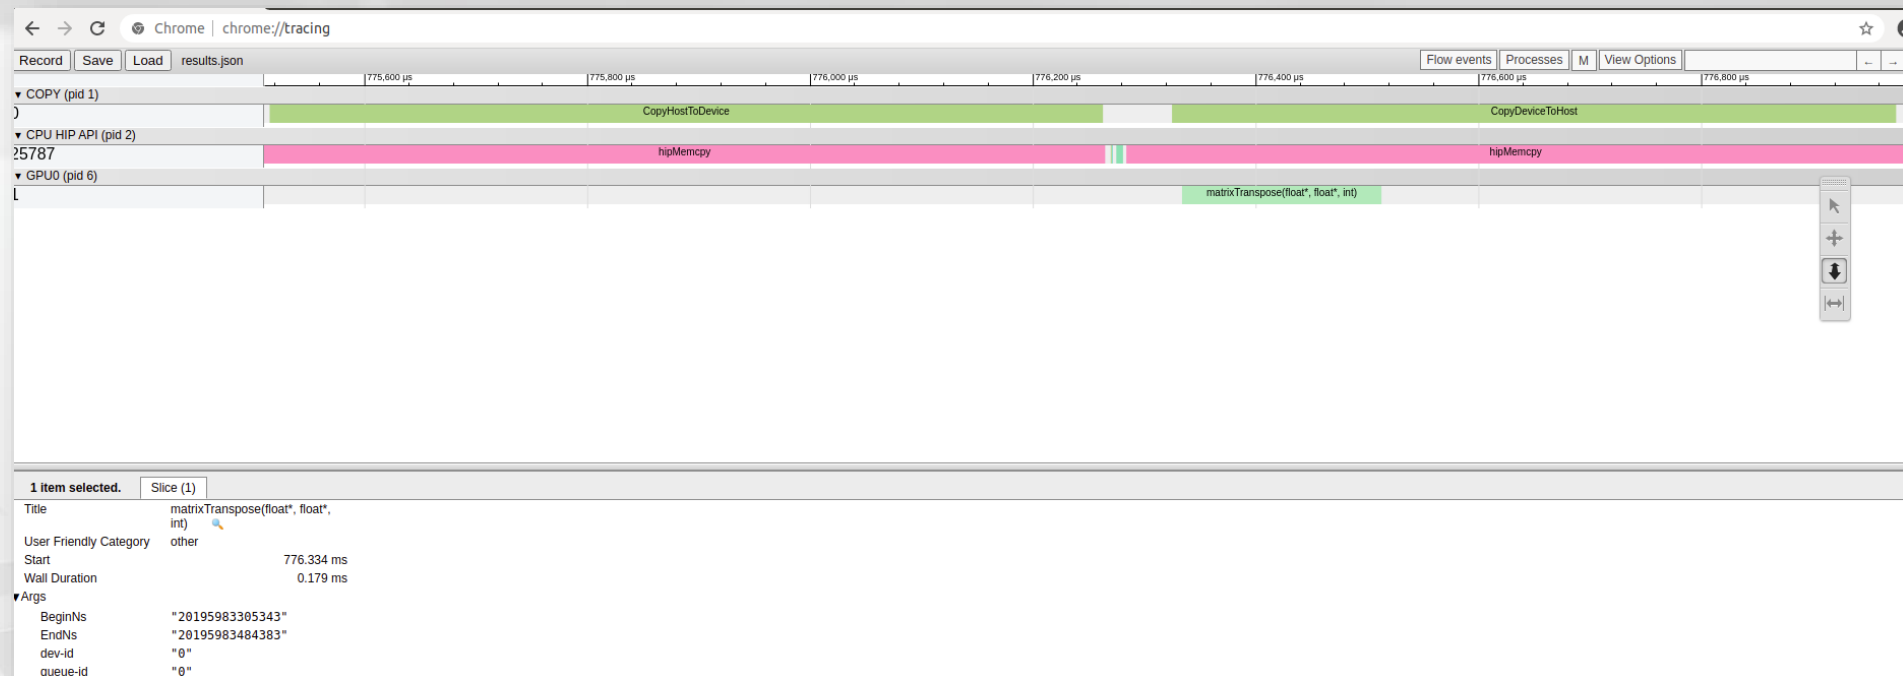
- `-m <metric file>`: Allows the user to define and collect custom metrics. See [rocprofiler/test/tool/*.xml](#) on GitHub for examples

Collecting traces with rocprof

rocprof can collect a variety of trace event types and generate timelines in JSON format for use with chrome-tracing, currently:

Trace Event	rocprof Trace Mode
HIP API call	--hip-trace
GPU Kernels	--hip-trace
Host <-> Device Memory copies	--hip-trace
CPU HSA Calls	--hsa-trace
User code markers	--roctx-trace

Collecting hip-trace with rocprof

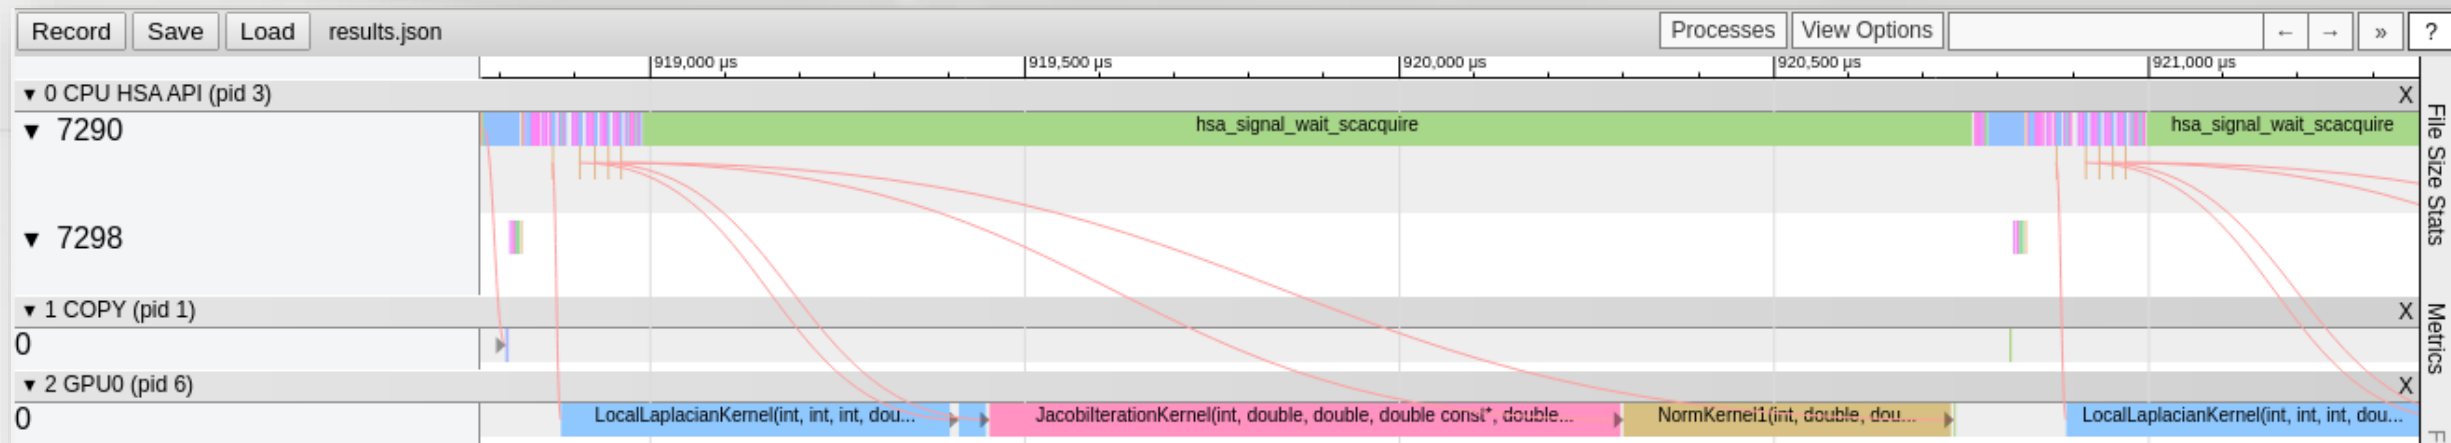


```
$ rocprof --hip-trace <app with arguments>
```

This will output a .json file that can be visualized using the chrome browser

Go to `chrome://tracing` and then load the .json file

- The trace will display HIP calls, mem copies, kernels



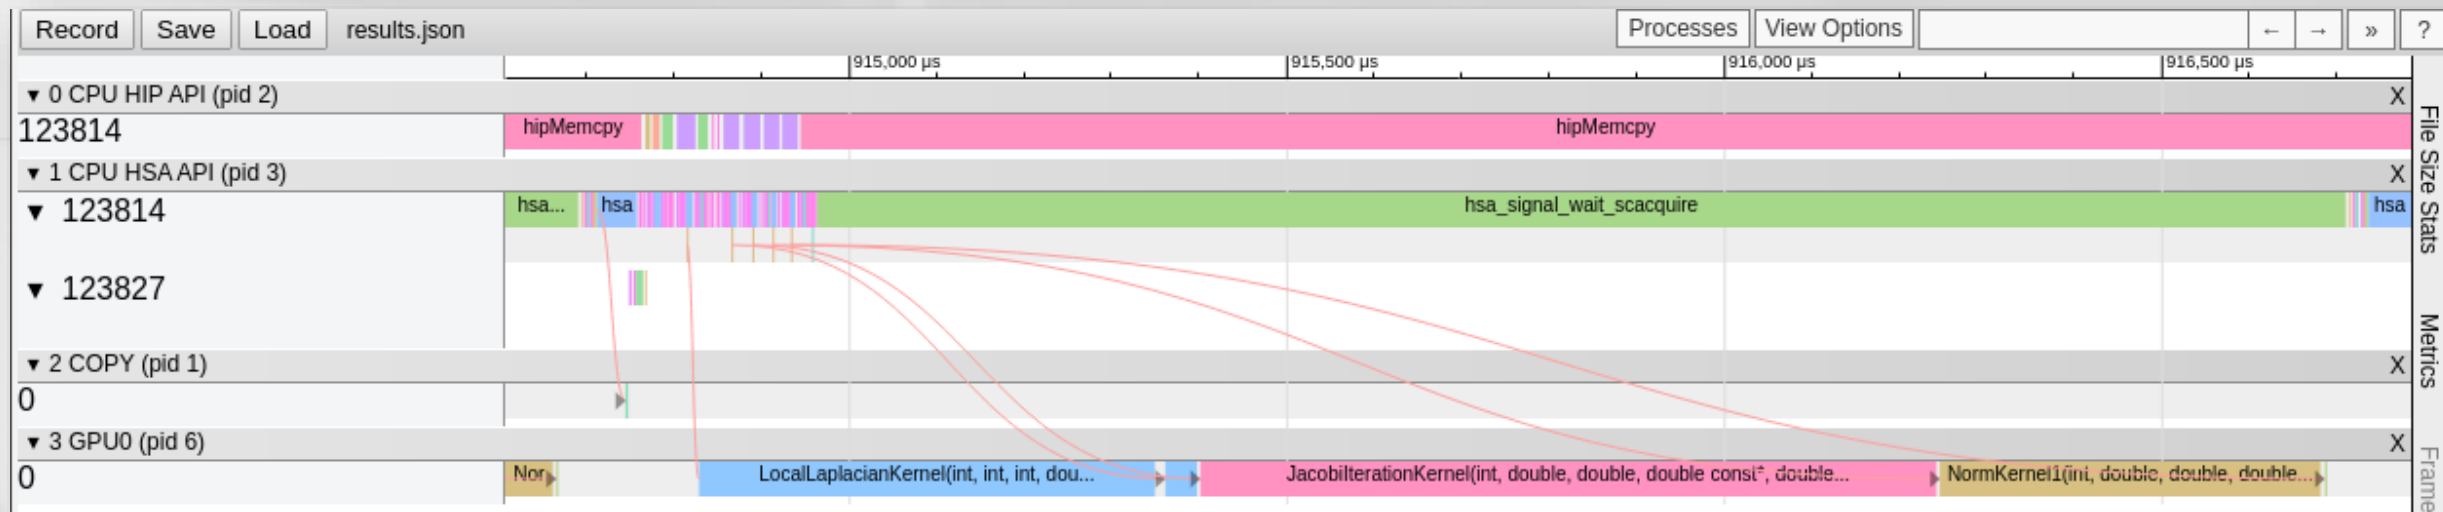
Collecting hsa-trace with rocprof

```
$ rocprof --hsa-trace <app with arguments>
```

This will output a .json file that
can be visualized using the
chrome browser

Go to <chrome://tracing> and
then load the .json file

- The trace will display copies, hsa signals, and kernel calls
- Slowest trace mode – Use with cautions



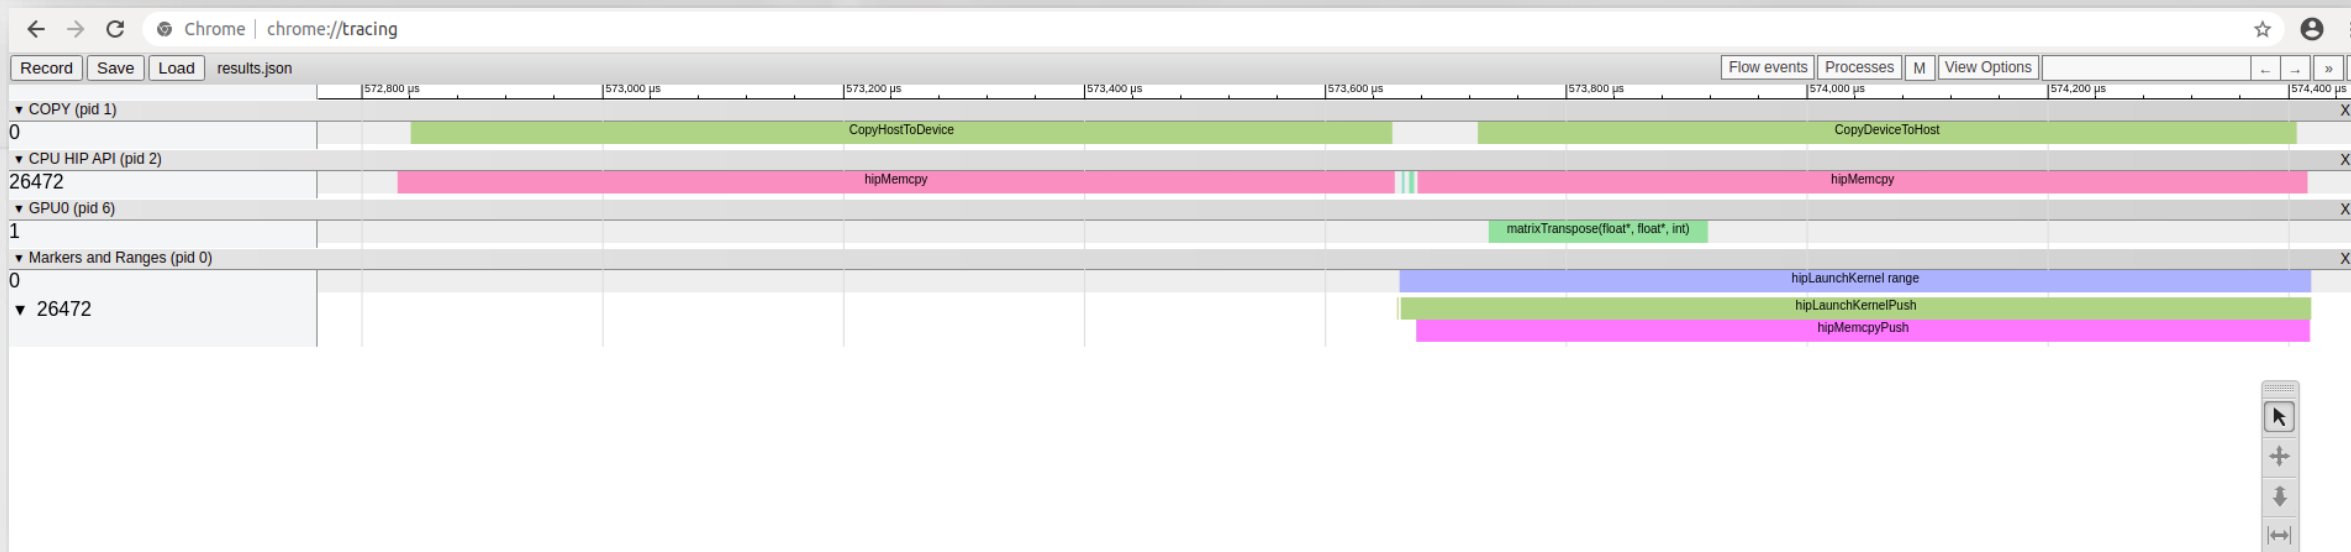
Collecting multiple trace modes with rocprof

```
$ rocprof --hsa-trace --hip-trace <app with arguments>
```

This will output a .json file that can be visualized using the chrome browser

Go to <chrome://tracing> and then load the .json file

- The trace will display HIP calls, copies, hsa signals, and kernel calls



Collecting user-code marker using rocTX with rocprof

```
$ rocprof --hip-trace --roctx-trace <app with arguments>
```

See [MatrixTranspose.cpp](#)
example on roctracer GitHub
page for sample in-code usage

This will output a .json file that
can be visualized using the
chrome browser

Collecting hardware counters with rocprof

rocprof can collect a number of hardware counters and derived counters

- `$/opt/rocm/bin/rocprof --list-basic`
- `$/opt/rocm/bin/rocprof --list-derived`

Specify counters in a counter file. For example:

- `$/opt/rocm/bin/rocprof -i my_counters.txt <app with args>`
- `$cat my_counters.txt`
 `pmc : Wavefronts VALUInsts VFetchInsts VWriteInsts VALUUtilization VALUBusy`
 `WriteSize`
 `pmc : SALUInsts SFetchInsts LDSInsts FlatLDSInsts GDSInsts SALUBusy FetchSize`
 `pmc : L2CacheHit MemUnitBusy MemUnitStalled WriteUnitStalled ALUStalledByLDS`
 `LDSBankConflict...`
- A limited number of counters can be collected during a specific pass of code
 - Each line in the counter file will be collected in one pass
 - You will receive an error suggesting alternative counter ordering if you have too many / conflicting counters on one line
- A `.csv` file will be created by this command containing all the requested counters

Commonly used counters

VALUUtilization	<ul style="list-style-type: none">• The percentage of ALUs active in a wave. Low VALUUtilization is likely due to high divergence or a poorly sized grid
VALUBusy	<ul style="list-style-type: none">• The percentage of GPUPTime the vector ALU instructions are processed. Can be thought of as something like compute utilization
FetchSize	<ul style="list-style-type: none">• The total kilobytes fetched from global memory
WriteSize	<ul style="list-style-type: none">• The total kilobytes written to global memory
L2CacheHit	<ul style="list-style-type: none">• The percentage of fetch, write, atomic, and other instructions that hit the data in L2 cache
MemUnitBusy	<ul style="list-style-type: none">• The percentage of GPUPTime the memory unit is active. The result includes the stall time
MemUnitStalled	<ul style="list-style-type: none">• The percentage of GPUPTime the memory unit is stalled
WriteUnitStalled	<ul style="list-style-type: none">• The percentage of GPUPTime the write unit is stalled
Full list at	<ul style="list-style-type: none">• https://github.com/ROCm-Developer-Tools/rocprofiler/blob/amd-master/test/tool/metrics.xml

Performance counters tips and tricks

GPU Hardware counters are global

- Kernel dispatches are serialized to ensure that only one dispatch is ever in flight
- It is recommended that no other applications are running that use the GPU when collecting performance counters

Use “`--basenames on`” which will report only kernel names, leaving off kernel arguments.

How do you time a kernel's duration?

- `$rocprof --timestamps on -i my_counters.txt <app with args>`
- This produces four times: DispatchNs, BeginNs, EndNs, and CompleteNs
- Closest thing to a kernel duration: EndNs - BeginNs
- If you run with `--stats` the resultant results file will automatically include a column that calculates kernel duration
 - Note: the duration is aggregated over repeated calls to the same kernel

Collecting counters and traces

on multiple MPI ranks

- rocprof can collect counters and traces for multiple MPI ranks.
- Say you want to profile an application usually called like this:

```
mpiexec -np <n> <app with args>
```

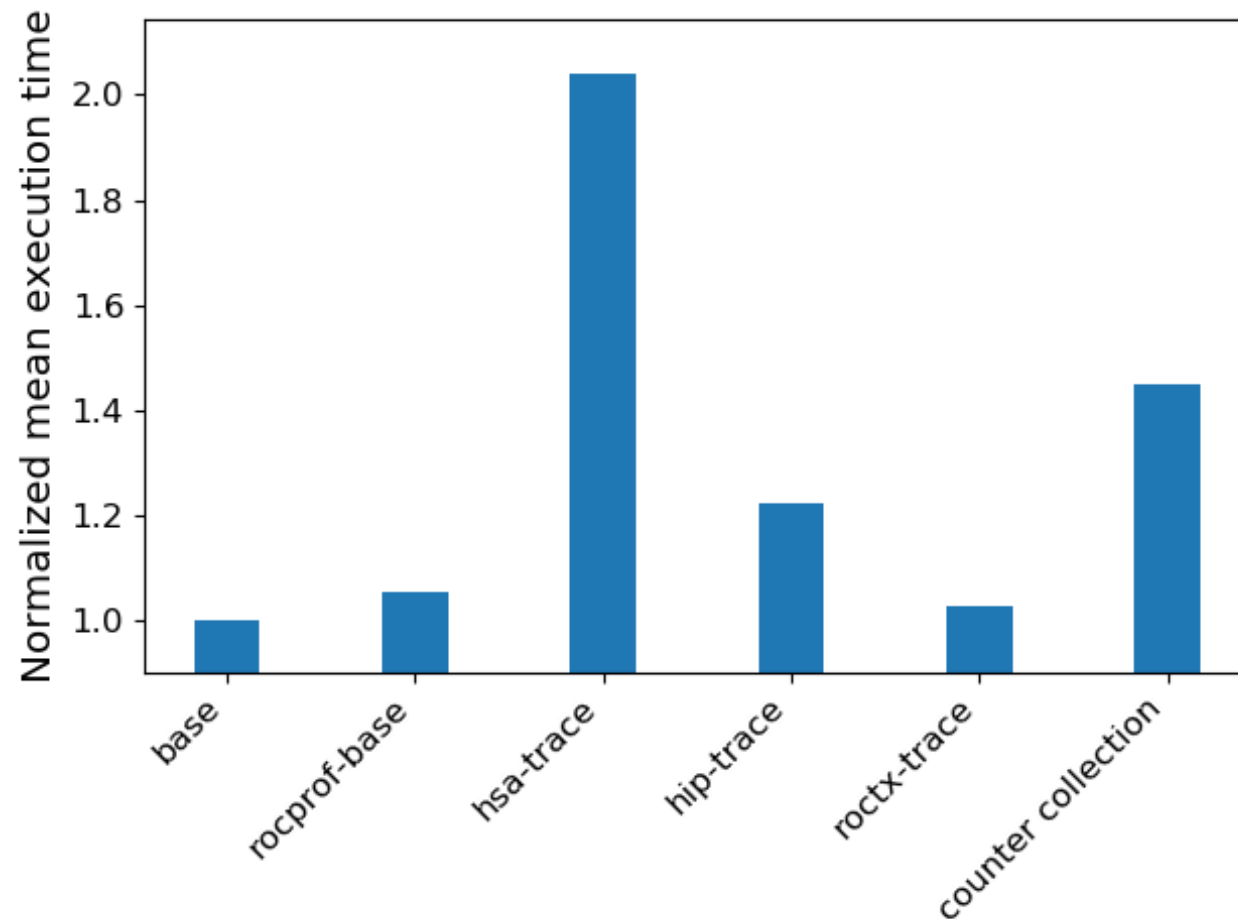
- Then invoke the profiler by executing:

```
rocprof --hip-trace mpiexec -np <n> <app with args>
```

- This will produce a single unified CSV file for all ranks
- Multi-node profiling currently isn't supported

Profiling overhead

Simple estimation of profiling overhead, obtained via wall-clock timing of entire application run via Linux 'time' utility:



A close-up, low-angle shot of an AMD Radeon Instinct GPU. The card is black with a prominent silver-colored metal grille on the left side. The words "RADEON INSTINCT" are printed in white, bold, sans-serif capital letters on the black surface of the card. The background is dark and out of focus, showing other components of a server or data center environment.

RADEON INSTINCT

Debugger



Intro to ROCgdb

What is ROCgdb, from the tin:

The ROCm Debugger (ROCgdb) is the ROCm source-level debugger for Linux, based on the GNU Debugger (GDB). It enables heterogenous debugging on the ROCm platform of an x86-based host architecture along with AMD GPU architectures supported by the AMD Debugger API Library (ROCdbgapi). The AMD Debugger API Library (ROCdbgapi) is included with the ROCm release.

The current ROCm Debugger (ROCgdb) is an initial prototype that focuses on source line debugging and does not provide symbolic variable debugging capabilities. The user guide presents features and commands that may be implemented in future versions.

So... cuda-gdb? Yes, and mostly no -- rocdbg is (or will be) gdb, that is it tracks upstream GDB master.

Preparing the code for the debugger

Use any optimization level you like,

For example: `-O3`

Have ROCm load code objects at initialization:

```
export HIP_ENABLE_DEFERRED_LOADING=0
```

Add the flags:

`-ggdb`

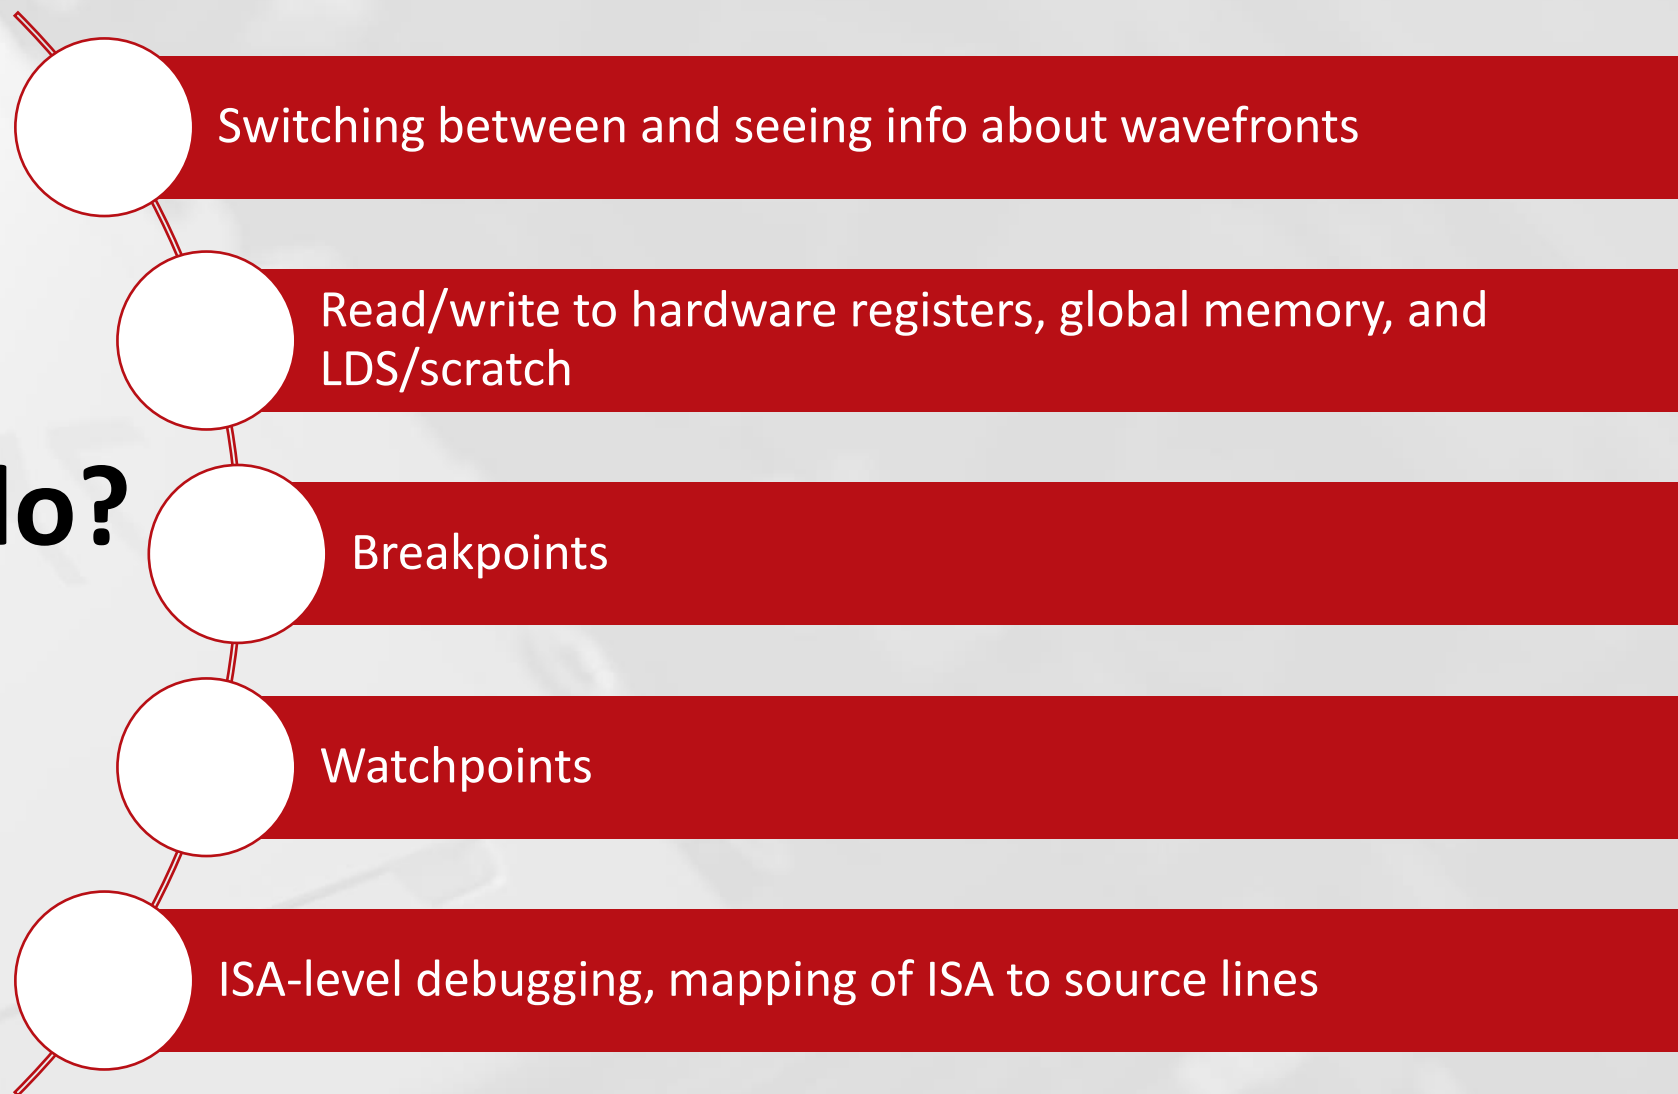
Optionally print even more useful information on API calls

```
export AMD_LOG_LEVEL=3
```

Example of what the compile options may look like...

```
mpic++ -I/usr/lib/x86_64-linux-gnu/openmpi/include/openmpi -L/usr/lib/x86_64-linux-  
gnu/openmpi/include -pthread -O3 -g -ggdb -fPIC -std=c++11 \ -march=native -Wall -  
I/opt/rocm/roctracer/include -I"/opt/rocm-4.2.0/hip/include" -  
I"/opt/rocm/llvm/bin/../lib/clang/12.0.0" -I/opt/rocm/hsa/include -  
I/opt/rocm/roctracer/include -c JacobiSetup.cpp -o JacobiSetup.o
```

What can it do?



Setting a breakpoint in host code

Here we setup a breakpoint in the host code. We can inspect the device pointer and its values:

```

23
24     const int id_l = id - 1;
25     const int id_r = id + 1;
26     const int id_d = id - stride;
27     const int id_u = id + stride;
28
29     AU[id] = (-U[id_l] + 2*U[id] - U[id_r])/(dx*dx) +
30             (-U[id_d] + 2*U[id] - U[id_u])/(dy*dy);
31 }
32 }
33
34 void LocalLaplacian(grid_t& grid, mesh_t& mesh,
35                    hipStream_t stream,
36                    dfloat* d_U,
37                    dfloat* d_AU) {
38
39     //there are (Nx-2)x(Ny-2) node on the interior of the mesh
40     int localNx = mesh.Nx-2;
41     int localNy = mesh.Ny-2;
42
43     int xthreads = 16;
44     int ythreads = 16;
45
46     dim3 threads(xthreads,ythreads,1);
47     dim3 blocks((localNx+xthreads-1)/xthreads,
48                (localNy+ythreads-1)/ythreads, 1);
49
50     hipLaunchKernelGGL(LocalLaplacianKernel,
51                        blocks,
52                        threads,
53                        0, stream,
54                        localNx, localNy, mesh.Nx,
55                        mesh.dx, mesh.dy,
56                        d_U, d_AU);

```

```

For help, type "help".
Type "apropos word" to search for commands related to "word"...
Reading symbols from ./Jacobi_hip...
(gdb) b Laplacian.cpp: 40
Breakpoint 1 at 0x415080: file Laplacian.cpp, line 40.
(gdb) run
Starting program: /home/jychang48/Downloads/hiptutorial/hip/Jacobi_hip -g 1 1
[Thread debugging using libthread_db enabled]
Using host libthread_db library "/lib/x86_64-linux-gnu/libthread_db.so.1".
[Detaching after fork from child process 129946]
[New Thread 0x7f2900108700 (LWP 129953)]
[New Thread 0x7f28ff74a700 (LWP 129954)]
Topology size: 1 x 1
Local domain size (current node): 4096 x 4096
Global domain size (all nodes): 4096 x 4096
[New Thread 0x7f28f7f53700 (LWP 129955)]
Rank 0 selecting device 0 on host jychang48-workstation
[New Thread 0x7f28f71ff700 (LWP 129956)]
[Thread 0x7f28f71ff700 (LWP 129956) exited]
[New Thread 0x7f28fc07f700 (LWP 129957)]
[New Thread 0x7f28f7752700 (LWP 129958)]
[New Thread 0x7f28f757f700 (LWP 129959)]
Starting Jacobi run.
Iteration: 0 - Residual: 0.022108

Thread 1 "Jacobi_hip" hit Breakpoint 1, LocalLaplacian (grid=..., mesh=..., stream=0
x69a6e0, d_U=0x7f27a7e00000, d_AU=0x7f279fc00000) at Laplacian.cpp:40
(gdb) p d_U
$1 = (double *) 0x7f27a7e00000
(gdb) p d_U[0]
$2 = 0
(gdb) p d_U[0]@10
$3 = {0, 0, 0, 0, 0, 0, 0, 0, 0, 0}
(gdb)

```


Setting a breakpoint in device kernel

Invoke 'b' or 'break' to the device kernel of interest:

```

1  //*****
2  /* Copyright (c) 2019, Advanced Micro Devices, Inc. All rights reserved.
3  //*****
4
5  #include "Jacobi.hpp"
6
7  // AU_i,j = (-U_i+1,j + 2U_i,j - U_i-1,j)/dx^2 +
8  //          (-U_i,j+1 + 2U_i,j - U_i,j-1)/dy^2
9  >__global__ void LocalLaplacianKernel(const int localNx,
10                                     const int localNy,
11                                     const int stride,
12                                     const dfloat dx,
13                                     const dfloat dy,
14                                     const dfloat *__restrict__ U,
15                                     dfloat *__restrict__ AU) {
16
17     const int i = threadIdx.x+blockIdx.x*blockDim.x;
18     const int j = threadIdx.y+blockIdx.y*blockDim.y;
19
20     if ((i<localNx) && (j<localNy)) {
21
22         const int id = (i+1) + (j+1)*stride;
23
24         const int id_l = id - 1;
25         const int id_r = id + 1;
26         const int id_d = id - stride;
27         const int id_u = id + stride;
28
29         AU[id] = (-U[id_l] + 2*U[id] - U[id_r])/(dx*dx) +
30                 (-U[id_d] + 2*U[id] - U[id_u])/(dy*dy);
31     }
32 }
33
34 void LocalLaplacian(grid t& grid, mesh t& mesh,
/home/jychang48/Downloads/hiptutorial/hip/Laplacian.cpp

```

```

Topology size: 1 x 1
Local domain size (current node): 4096 x 4096
Global domain size (all nodes): 4096 x 4096
[New Thread 0x7f28f7f53700 (LWP 129955)]
Rank 0 selecting device 0 on host jychang48-workstation
[New Thread 0x7f28f71ff700 (LWP 129956)]
[Thread 0x7f28f71ff700 (LWP 129956) exited]
[New Thread 0x7f28fc07f700 (LWP 129957)]
[New Thread 0x7f28f7752700 (LWP 129958)]
[New Thread 0x7f28f757f700 (LWP 129959)]
Starting Jacobi run.
Iteration: 0 - Residual: 0.022108

Thread 1 "Jacobi hip" hit Breakpoint 1, LocalLaplacian (grid=..., mesh=..., stream=0
x69a6e0, d U=0x7f27a7e00000, d_AU=0x7f279fc00000) at Laplacian.cpp:40
(gdb) p d U
$1 = (double *) 0x7f27a7e00000
(gdb) p d_U[0]
$2 = 0
(gdb) p d_U[0]@10
$3 = {0, 0, 0, 0, 0, 0, 0, 0, 0, 0}
(gdb) step
(gdb) step
(gdb) b LocalLaplacianKernel
Function "LocalLaplacianKernel" not defined.
Make breakpoint pending on future shared library load? (y or [n]) y

Breakpoint 2 (LocalLaplacianKernel) pending.(gdb) continue
Continuing.
[Switching to AMDGPU Thread 1:5:1:1 (0,0,0)/0]

Thread 9 "Jacobi hip" hit Breakpoint 2, LocalLaplacianKernel (localNx=<optimized out
>, localNy=<optimized out>, stride=<optimized out>, dx=<optimized out>, dy=<optimize
d out>, U=<optimized out>, AU=<optimized out>) at Laplacian.cpp:9
(gdb)

```

Examine switching context to new thread

What happens when you type 'step'? Another thread hit the same breakpoint!

GDB will switch context to the new thread:

```
jychang48@jychang48-workstation: ~/Downloads/hiptutorial/hip

1 //*****
2 /** Copyright (c) 2019, Advanced Micro Devices, Inc. All rights reserved.
3 //*****
4
5 #include "Jacobi.hpp"
6
7 // AU_i,j = (-U_i+1,j + 2U_i,j - U_i-1,j)/dx^2 +
8 //          (-U_i,j+1 + 2U_i,j - U_i,j-1)/dy^2
9 >_global__ void LocalLaplacianKernel(const int localNx,
10                                     const int localNy,
11                                     const int stride,
12                                     const dfloat dx,
13                                     const dfloat dy,
14                                     const dfloat *__restrict__ U,
15                                     dfloat *__restrict__ AU) {
16
17     const int i = threadIdx.x+blockIdx.x*blockDim.x;
18     const int j = threadIdx.y+blockIdx.y*blockDim.y;
19
20     if ((i<localNx) && (j<localNy)) {
21
22         const int id = (i+1) + (j+1)*stride;
23
24         const int id_l = id - 1;
25         const int id_r = id + 1;
26         const int id_d = id - stride;
27         const int id_u = id + stride;
28
29         AU[id] = (-U[id_l] + 2*U[id] - U[id_r])/(dx*dx) +
30                 (-U[id_d] + 2*U[id] - U[id_u])/(dy*dy);
31     }
32 }
33
34 void LocalLaplacian(grid t& grid, mesh t& mesh,
```

```
$2 = 0
(gdb) p d_U[0]@10
$3 = {0, 0, 0, 0, 0, 0, 0, 0, 0, 0}
(gdb) step
(gdb) step
(gdb) b LocalLaplacianKernel
Function "LocalLaplacianKernel" not defined.
Make breakpoint pending on future shared library load? (y or [n]) y

Breakpoint 2 (LocalLaplacianKernel) pending.(gdb) continue
Continuing.
[Switching to AMDGPU Thread 1:5:1:1 (0,0,0)/0]

Thread 9 "Jacobi_hip" hit Breakpoint 2, LocalLaplacianKernel (localNx=<optimized out>, localNy=<optimized out>, stride=<optimized out>, dx=<optimized out>, dy=<optimized out>, U=<optimized out>, AU=<optimized out>) at Laplacian.cpp:9
(gdb) step
[Switching to AMDGPU Thread 1:5:1:1287 (65,1,0)/2]

Thread 1295 "Jacobi_hip" hit Breakpoint 2, LocalLaplacianKernel (localNx=<optimized out>, localNy=<optimized out>, stride=<optimized out>, dx=<optimized out>, dy=<optimized out>, U=<optimized out>, AU=<optimized out>) at Laplacian.cpp:9
(gdb) step
[Switching to AMDGPU Thread 1:5:1:1269 (61,1,0)/0]

Thread 1277 "Jacobi_hip" hit Breakpoint 2, LocalLaplacianKernel (localNx=<optimized out>, localNy=<optimized out>, stride=<optimized out>, dx=<optimized out>, dy=<optimized out>, U=<optimized out>, AU=<optimized out>) at Laplacian.cpp:9
(gdb) step
[Switching to AMDGPU Thread 1:5:1:363 (90,0,0)/2]

Thread 371 "Jacobi_hip" hit Breakpoint 2, LocalLaplacianKernel (localNx=<optimized out>, localNy=<optimized out>, stride=<optimized out>, dx=<optimized out>, dy=<optimized out>, U=<optimized out>, AU=<optimized out>) at Laplacian.cpp:9
(gdb)
```

/home/jychang48/Downloads/hiptutorial/hip/Laplacian.cpp

Examine switching context to new thread

AMDGPU Thread agent-id:queue-id:dispatch-num:wave-id (work-group-z,work-group-y,work-group-x)/work-group-thread-index

```
[Switching to AMDGPU Thread 1:5:1:1 (0,0,0)/0]
Thread 9 "Jacobi_hip" hit Breakpoint 2, LocalLaplacianKernel (localNx=<optimized out>, localNy=<optimized out>, stride=<optimized out>, dx=<optimized out>, dy=<optimize
d out>, U=<optimized out>, AU=<optimized out>) at Laplacian.cpp:9
```



Examine the ISA with rocgdb using cgdb

Using cgdb, type

ESC -> :set dis -> ENTER

```

3 9  __global__ void LocalLaplacianKernel(const int localNx,
4    0x00007f28f722f000 <+0>:    s_load_dwordx4 s[0:3], s[6:7], 0x0
5
6 /opt/rocm-4.2.0/hip/include/hip/amd_detail/hip_runtime.h:
7 232 return __ockl_get_local_size(x);
8    0x00007f28f722f008 <+8>:    s_load_dword s10, s[4:5], 0x4
9    0x00007f28f722f010 <+16>:   s_waitcnt lgkmcnt(0)
10   0x00007f28f722f014 <+20>:   s_lshr_b32 s3, s10, 16
11   0x00007f28f722f018 <+24>:   s_and_b32 s4, s10, 0xffff
12   0x00007f28f722f020 <+32>:   s_mul_i32 s8, s8, s4
13   0x00007f28f722f024 <+36>:   s_mul_i32 s9, s9, s3
14
15 Laplacian.cpp:
16 17  const int i = threadIdx.x+blockIdx.x*blockDim.x;
17   0x00007f28f722f028 <+40>:   v_add_u32_e32 v0, s8, v0
18
19 18  const int j = threadIdx.y+blockIdx.y*blockDim.y;
20  -> 0x00007f28f722f02c <+44>:   v_add_u32_e32 v1, s9, v1
21
22 19
23 20  if ((i<localNx) && (j<localNy)) {
24   0x00007f28f722f030 <+48>:   v_cmp_gt_i32_e32 vcc, s0, v0
25   0x00007f28f722f034 <+52>:   v_cmp_gt_i32_e64 s[0:1], s1, v1
26   0x00007f28f722f03c <+60>:   s_and_b64 s[0:1], vcc, s[0:1]
27   0x00007f28f722f040 <+64>:   s_and_saveexec b64 s[4:5], s[0:1]
28   0x00007f28f722f044 <+68>:   s_cbranch_execz 104
29
30 21
31 22  const int id = (i+1) + (j+1)*stride;
32   0x00007f28f722f048 <+72>:   v_add_u32_e32 v1, 1, v1
33   0x00007f28f722f04c <+76>:   v_mul_lo_u32 v1, v1, s2
34   0x00007f28f722f054 <+84>:   s_load_dwordx2 s[0:1], s[6:7], 0x20
35
36 9  __global__ void LocalLaplacianKernel(const int localNx,
>de for function __Z20LocalLaplacianKerneliiiddPKdPd: (7f28f722f000 - 7f28f722f1e8) **

```

Make breakpoint pending on future shared library load? (y or [n]) y

Breakpoint 2 (LocalLaplacianKernel) pending.(gdb) continue
Continuing.
[Switching to AMDGPU Thread 1:5:1:1 (0,0,0)/0]

Thread 9 "Jacobi_hip" hit Breakpoint 2, LocalLaplacianKernel (localNx=<optimized out>, localNy=<optimized out>, stride=<optimized out>, dx=<optimized out>, dy=<optimized out>, U=<optimized out>, AU=<optimized out>) at Laplacian.cpp:9
(gdb) step
[Switching to AMDGPU Thread 1:5:1:1287 (65,1,0)/2]

Thread 1295 "Jacobi_hip" hit Breakpoint 2, LocalLaplacianKernel (localNx=<optimized out>, localNy=<optimized out>, stride=<optimized out>, dx=<optimized out>, dy=<optimized out>, U=<optimized out>, AU=<optimized out>) at Laplacian.cpp:9
(gdb) step
[Switching to AMDGPU Thread 1:5:1:1269 (61,1,0)/0]

Thread 1277 "Jacobi_hip" hit Breakpoint 2, LocalLaplacianKernel (localNx=<optimized out>, localNy=<optimized out>, stride=<optimized out>, dx=<optimized out>, dy=<optimized out>, U=<optimized out>, AU=<optimized out>) at Laplacian.cpp:9
(gdb) step
[Switching to AMDGPU Thread 1:5:1:363 (90,0,0)/2]

Thread 371 "Jacobi_hip" hit Breakpoint 2, LocalLaplacianKernel (localNx=<optimized out>, localNy=<optimized out>, stride=<optimized out>, dx=<optimized out>, dy=<optimized out>, U=<optimized out>, AU=<optimized out>) at Laplacian.cpp:9
(gdb) disable 2
(gdb) step
(gdb) stepi
0x00007f28f722f010 in __HIP_BlockDim::operator() (this=<optimized out>, x=<optimized out>) at /opt/rocm-4.2.0/hip/include/hip/amd_detail/hip_runtime.h:232
(gdb) step
(gdb) step
(gdb)

Switching wavefronts

Use info threads to see the location of both host threads and GPU wavefronts

The screenshot shows a debugger window with two panes. The left pane displays C++ code for a Laplacian kernel, and the right pane shows the output of the 'info threads' command, listing GPU threads and their execution context.

```

15 Laplacian.cpp:
16 17      const int i = threadIdx.x+blockIdx.x*blockDim.x;
17      0x00007f28f722f028 <+40>:    v_add_u32_e32 v0, s8, v0
18
19 18      const int j = threadIdx.y+blockIdx.y*blockDim.y;
20      0x00007f28f722f02c <+44>:    v_add_u32_e32 v1, s9, v1
21
22 19
23 20      if ((i<localNx) && (j<localNy)) {
24      0x00007f28f722f030 <+48>:    v_cmp_gt_i32_e32 vcc, s0, v0
25      0x00007f28f722f034 <+52>:    v_cmp_gt_i32_e64 s[0:1], s1, v1
26      0x00007f28f722f03c <+60>:    s_and_b64 s[0:1], vcc, s[0:1]
27      0x00007f28f722f040 <+64>:    s_and_saveexec_b64 s[4:5], s[0:1]
28      0x00007f28f722f044 <+68>:    s_cbranch_execz 104
29
30 21
31 22      const int id = (i+1) + (j+1)*stride;
32  → 0x00007f28f722f048 <+72>:    v_add_u32_e32 v1, 1, v1
33      0x00007f28f722f04c <+76>:    v_mul_lo_u32 v1, v1, s2
34      0x00007f28f722f054 <+84>:    s_load_dwordx2 s[0:1], s[6:7], 0x20
35
36 9      __global__ void LocalLaplacianKernel(const int localNx,
37      0x00007f28f722f05c <+92>:    s_load_dwordx4 s[8:11], s[6:7], 0x10
38
39 21
40 22      const int id = (i+1) + (j+1)*stride;
41      0x00007f28f722f064 <+100>:   v_add_u32_e32 v4, v1, v0
42
43 28
44 29      AU[id] = (-U[id_l] + 2*U[id] - U[id_r])/(dx*dx) +
45      0x00007f28f722f068 <+104>:   v_ashrrev_i32_e32 v5, 31, v4
46      0x00007f28f722f06c <+108>:   v_lshlrev_b64 v[5:6], 3, v[4:5]
47      0x00007f28f722f074 <+116>:   s_waitcnt lgkmcnt(0)
48      0x00007f28f722f078 <+120>:   v_mov_b32_e32 v15, s1
49  }
50 }
51
52 de for function Z20LocalLaplacianKerneliiddPKdPd: (7f28f722f000 - 7f28f722f1e8) **
  
```

The right pane shows the output of the 'info threads' command, listing GPU threads and their execution context. The threads are listed in a table with columns: Id, Target Id, and Frame. The threads are grouped by their target ID and frame, showing multiple instances of the 'Jacobi_hip' thread.

```

(gdb) info threads
Id      Target Id      Frame
1      Thread 0x7f291814d8c0 (LWP 129937) "Jacobi_hip" 0x00007f2918b7a89b in sched_y
ield () from /lib/x86_64-linux-gnu/libc.so.6
2      Thread 0x7f2900108700 (LWP 129953) "Jacobi_hip" 0x00007f2918b8aaff in poll ()
from /lib/x86_64-linux-gnu/libc.so.6
3      Thread 0x7f28ff74a700 (LWP 129954) "Jacobi_hip" 0x00007f2918b975ce in epoll_w
ait () from /lib/x86_64-linux-gnu/libc.so.6
4      Thread 0x7f28f7f53700 (LWP 129955) "Jacobi_hip" 0x00007f2918b8c50b in ioctl (
) from /lib/x86_64-linux-gnu/libc.so.6
5      Thread 0x7f28fc07f700 (LWP 129957) "Jacobi_hip" 0x00007f2918b7a89b in sched_y
ield () from /lib/x86_64-linux-gnu/libc.so.6
6      Thread 0x7f28f7752700 (LWP 129958) "Jacobi_hip" 0x00007f2918b8c50b in ioctl (
) from /lib/x86_64-linux-gnu/libc.so.6
7      Thread 0x7f28f757f700 (LWP 129959) "Jacobi_hip" 0x00007f291a2f4678 in do_fute
x_wait.constprop () from /lib/x86_64-linux-gnu/libpthread.so.0
* 371 AMDGPU Thread 1:5:1:363 (90,0,0)/2 "Jacobi_hip" LocalLaplacianKernel (localNx
=<optimized out>, localNy=<optimized out>, stride=<optimized out>, dx=<optimized out
>, dy=<optimized out>, U=<optimized out>, AU=<optimized out>) at Laplacian.cpp:22
(gdb)
  
```

Switching wavefronts

Or use `thread <tid>` to examine one particular thread

```

15 Laplacian.cpp:
16 17      const int i = threadIdx.x+blockIdx.x*blockDim.x;
17      0x00007f28f722f028 <+40>:    v_add_u32_e32 v0, s8, v0
18
19 18      const int j = threadIdx.y+blockIdx.y*blockDim.y;
20      0x00007f28f722f02c <+44>:    v_add_u32_e32 v1, s9, v1
21
22 19
23 20      if ((i<localNx) && (j<localNy)) {
24      0x00007f28f722f030 <+48>:    v_cmp_gt_i32_e32 vcc, s0, v0
25      0x00007f28f722f034 <+52>:    v_cmp_gt_i32_e64 s[0:1], s1, v1
26      0x00007f28f722f03c <+60>:    s_and_b64 s[0:1], vcc, s[0:1]
27      0x00007f28f722f040 <+64>:    s_and_saveexec_b64 s[4:5], s[0:1]
28      0x00007f28f722f044 <+68>:    s_cbranch_execz 104
29
30 21
31 22      const int id = (i+1) + (j+1)*stride;
32  → 0x00007f28f722f048 <+72>:    v_add_u32_e32 v1, 1, v1
33      0x00007f28f722f04c <+76>:    v_mul_lo_u32 v1, v1, s2
34      0x00007f28f722f054 <+84>:    s_load_dwordx2 s[0:1], s[6:7], 0x20
35
36 9      __global__ void LocalLaplacianKernel(const int localNx,
37      0x00007f28f722f05c <+92>:    s_load_dwordx4 s[8:11], s[6:7], 0x10
38
39 21
40 22      const int id = (i+1) + (j+1)*stride;
41      0x00007f28f722f064 <+100>:   v_add_u32_e32 v4, v1, v0
42
43 28
44 29      AU[id] = (-U[id_l] + 2*U[id] - U[id_r])/(dx*dx) +
45      0x00007f28f722f068 <+104>:   v_ashrrev_i32_e32 v5, 31, v4
46      0x00007f28f722f06c <+108>:   v_lshlrev_b64 v[5:6], 3, v[4:5]
47      0x00007f28f722f074 <+116>:   s_waitcnt lgkmcnt(0)
48      0x00007f28f722f078 <+120>:   v_mov_b32_e32 v15, s1
de for function Z20LocalLaplacianKerneliiddPKdPd: (7f28f722f000 - 7f28f722f1e8) **

(gdb) info threads
[110/112]
  Id   Target Id                                     Frame
  1     Thread 0x7f291814d8c0 (LWP 129937) "Jacobi_hip" 0x00007f2918b7a89b in sched_y
ield () from /lib/x86_64-linux-gnu/libc.so.6
  2     Thread 0x7f2900108700 (LWP 129953) "Jacobi_hip" 0x00007f2918b8aaff in poll ()
from /lib/x86_64-linux-gnu/libc.so.6
  3     Thread 0x7f28ff74a700 (LWP 129954) "Jacobi_hip" 0x00007f2918b975ce in epoll_w
ait () from /lib/x86_64-linux-gnu/libc.so.6
  4     Thread 0x7f28f7f53700 (LWP 129955) "Jacobi_hip" 0x00007f2918b8c50b in ioctl (
) from /lib/x86_64-linux-gnu/libc.so.6
  6     Thread 0x7f28fc07f700 (LWP 129957) "Jacobi_hip" 0x00007f2918b7a89b in sched_y
ield () from /lib/x86_64-linux-gnu/libc.so.6
  7     Thread 0x7f28f7752700 (LWP 129958) "Jacobi_hip" 0x00007f2918b8c50b in ioctl (
) from /lib/x86_64-linux-gnu/libc.so.6
  8     Thread 0x7f28f757f700 (LWP 129959) "Jacobi_hip" 0x00007f291a2f4678 in do_fute
x_wait.constprop () from /lib/x86_64-linux-gnu/libpthread.so.0
* 371 AMDGPU Thread 1:5:1:363 (90,0,0)/2 "Jacobi_hip" LocalLaplacianKernel (localNx
=<optimized out>, localNy=<optimized out>, stride=<optimized out>, dx=<optimized out
>, dy=<optimized out>, U=<optimized out>, AU=<optimized out>) at Laplacian.cpp:22
(gdb)
  Id   Target Id                                     Frame
  1     Thread 0x7f291814d8c0 (LWP 129937) "Jacobi_hip" 0x00007f2918b7a89b in sched_y
ield () from /lib/x86_64-linux-gnu/libc.so.6
  2     Thread 0x7f2900108700 (LWP 129953) "Jacobi_hip" 0x00007f2918b8aaff in poll ()
from /lib/x86_64-linux-gnu/libc.so.6
  3     Thread 0x7f28ff74a700 (LWP 129954) "Jacobi_hip" 0x00007f2918b975ce in epoll_w
ait () from /lib/x86_64-linux-gnu/libc.so.6
  4     Thread 0x7f28f7f53700 (LWP 129955) "Jacobi_hip" 0x00007f2918b8c50b in ioctl (
) from /lib/x86_64-linux-gnu/libc.so.6
  6     Thread 0x7f28fc07f700 (LWP 129957) "Jacobi_hip" 0x00007f2918b7a89b in sched_y
ield () from /lib/x86_64-linux-gnu/libc.so.6
  7     Thread 0x7f28f7752700 (LWP 129958) "Jacobi_hip" 0x00007f2918b8c50b in ioctl (
) from /lib/x86_64-linux-gnu/libc.so.6
  8     Thread 0x7f28f757f700 (LWP 129959) "Jacobi_hip" 0x00007f291a2f4678 in do_fute
x_wait.constprop () from /lib/x86_64-linux-gnu/libpthread.so.0

```

Other tips and tricks with rocgdb

Use `export AMD_LOG_LEVEL=3` to print all API calls and more happening

```

jychang48@jychang48-workstation: ~/Downloads/hiptutorial/hip
23
24     const int id_l = id - 1;
25     const int id_r = id + 1;
26     const int id_d = id - stride;
27     const int id_u = id + stride;
28
29     AU[id] = (-U[id_l] + 2*U[id] - U[id_r])/(dx*dx) +
30             (-U[id_d] + 2*U[id] - U[id_u])/(dy*dy);
31 }
32 }
33
34 void LocalLaplacian(grid_t& grid, mesh_t& mesh,
35                    hipStream_t stream,
36                    dfloat* d_U,
37                    dfloat* d_AU) {
38
39     //there are (Nx-2)x(Ny-2) node on the interior of the mesh
40 → int localNx = mesh.Nx-2;
41    int localNy = mesh.Ny-2;
42
43    int xthreads = 16;
44    int ythreads = 16;
45
46    dim3 threads(xthreads,ythreads,1);
47    dim3 blocks((localNx+xthreads-1)/xthreads,
48               (localNy+ythreads-1)/ythreads, 1);
49
50    hipLaunchKernelGGL(LocalLaplacianKernel,
51                       blocks,
52                       threads,
53                       0, stream,
54                       localNx, localNy, mesh.Nx,
55                       mesh.dx, mesh.dy,
56                       d_U, d_AU);

```

```

Reading symbols from ./Jacobi_hip... [41/144]
(gdb) b Laplacian.cpp: 40

Breakpoint 1 at 0x415080: file Laplacian.cpp, line 40.(gdb) run
Starting program: /home/jychang48/Downloads/hiptutorial/hip/Jacobi_hip -g 1 1
[Thread debugging using libthread_db enabled]
Using host libthread_db library "/lib/x86_64-linux-gnu/libthread_db.so.1".
[Detaching after fork from child process 151584]:3:rocdevice.cpp :457 : 4
3543864078 us: Initializing HSA stack.
:3:comgrctx.cpp :33 : 43543879647 us: Loading COMGR library.
:3:rocdevice.cpp :200 : 43543882744 us: Numa selects cpu agent[0]=0x55e1a
d860740(fine=0x55e1ad8490e0,coarse=0x55e1ad7e32c0, kern_arg=0x55e1ad8d7b20) for gpu
agent=0x7ffa27a13769

[New Thread 0x7f39eab37700 (LWP 151591)]
[New Thread 0x7f39ea179700 (LWP 151592)]
Topology size: 1 x 1
Local domain size (current node): 4096 x 4096
Global domain size (all nodes): 4096 x 4096
:3:rocdevice.cpp :457 : 43544254545 us: Initializing HSA stack.
[New Thread 0x7f39e253e700 (LWP 151593)]
:3:comgrctx.cpp :33 : 43544267424 us: Loading COMGR library.
:3:rocdevice.cpp :200 : 43544292744 us: Numa selects cpu agent[0]=0x88ab4
0(fine=0x84b9c0,coarse=0x8a39a0, kern_arg=0x8a2900) for gpu agent=0x7f3a03a431d9
:3:hip_context.cpp :124 : 43544298171 us: 151575: [7f3a02b7c8c0] hipInit: R
eturned hipSuccess :
:3:hip_device_runtime.cpp :497 : 43544298229 us: 151575: [7f3a02b7c8c0] hipGetDevi
ceCount ( 0x7ffd1902c1a4 )
:3:hip_device_runtime.cpp :499 : 43544298236 us: 151575: [7f3a02b7c8c0] hipGetDevi
ceCount: Returned hipSuccess :
:3:hip_device_runtime.cpp :512 : 43544298264 us: 151575: [7f3a02b7c8c0] hipSetDevi
ce ( 0 )
:3:hip_device_runtime.cpp :517 : 43544298269 us: 151575: [7f3a02b7c8c0] hipSetDevi
ce: Returned hipSuccess :
Rank 0 selecting device 0 on host jychang48-workstation

```


Other tips and tricks with rocgdb

Use `export AMD_LOG_LEVEL=3` to print all API calls and more happening

```

jychang48@jychang48-workstation: ~/Downloads/hiptutorial/hip
23
24     const int id_l = id - 1;
25     const int id_r = id + 1;
26     const int id_d = id - stride;
27     const int id_u = id + stride;
28
29     AU[id] = (-U[id_l] + 2*U[id] - U[id_r])/(dx*dx) +
30             (-U[id_d] + 2*U[id] - U[id_u])/(dy*dy);
31 }
32 }
33
34 void LocalLaplacian(grid_t& grid, mesh_t& mesh,
35                    hipStream_t stream,
36                    dfloat* d_U,
37                    dfloat* d_AU) {
38
39     //there are (Nx-2)x(Ny-2) node on the interior of the mesh
40 → int localNx = mesh.Nx-2;
41    int localNy = mesh.Ny-2;
42
43    int xthreads = 16;
44    int ythreads = 16;
45
46    dim3 threads(xthreads,ythreads,1);
47    dim3 blocks((localNx+xthreads-1)/xthreads,
48               (localNy+ythreads-1)/ythreads, 1);
49
50    hipLaunchKernelGGL(LocalLaplacianKernel,
51                       blocks,
52                       threads,
53                       0, stream,
54                       localNx, localNy, mesh.Nx,
55                       mesh.dx, mesh.dy,
56                       d_U, d_AU);

```

```

Rank 0 selecting device 0 on host jychang48-workstation [61/144]
:3:hip_memory.cpp:289: 43544298319 us: 151575: [7f3a02b7c8c0] hipHostMal
loc ( 0x7ffd1902c6f8, 131072, 0 )
:3:hip_memory.cpp:318: 43544298463 us: 151575: [7f3a02b7c8c0] hipHostMal
loc: Returned hipSuccess : 0x7f39e85c0000: duration: 144 us
:3:hip_memory.cpp:289: 43544298473 us: 151575: [7f3a02b7c8c0] hipHostMal
loc ( 0x7ffd1902c700, 131072, 0 )
:3:hip_memory.cpp:318: 43544298589 us: 151575: [7f3a02b7c8c0] hipHostMal
loc: Returned hipSuccess : 0x7f39e8580000: duration: 116 us
:3:hip_memory.cpp:283: 43544298607 us: 151575: [7f3a02b7c8c0] hipMalloc
( 0x7ffd1902c708, 131072 )
:3:rocdevice.cpp:2065: 43544298694 us: device=0x8d2890, freeMem_ = 0xfef
e0000
:3:hip_memory.cpp:285: 43544298702 us: 151575: [7f3a02b7c8c0] hipMalloc:
Returned hipSuccess : 0x7f39e1a00000: duration: 95 us
:3:hip_memory.cpp:1877: 43544298712 us: 151575: [7f3a02b7c8c0] hipMemset
( 0x7f39e1a00000, 0, 131072 )
:3:rocdevice.cpp:2566: 43544299081 us: number of allocated hardware que
es with low priority: 0, with normal priority: 0, with high priority: 0, maximum per
priority is: 4
[New Thread 0x7f39e19ff700 (LWP 151594)]
[Thread 0x7f39e19ff700 (LWP 151594) exited]
[New Thread 0x7f39e857f700 (LWP 151595)]
:3:rocdevice.cpp:2638: 43544316233 us: created hardware queue 0x7f39e88c
6000 with size 1024 with priority 1, cooperative: 0
:3:devprogram.cpp:2463: 43544489762 us: Using Code Object V4.
:3:rocvirtual.cpp:572: 43544495580 us: ! arg1: uint* bufUInt = ptr:0x
7f39e1a00000 obj:[0x7f39e1a00000-0x7f39e1a20000] threadId : 7f39e857f700
:3:rocvirtual.cpp:572: 43544495588 us: ! arg2: uchar* pattern = ptr:0
x7f39e876e080 obj:[0x7f39e876e000-0x7f39e876f000] threadId : 7f39e857f700
:3:rocvirtual.cpp:2521: 43544495591 us: [7f39e857f700]! ShaderName :
__amd_rocclr_fillBuffer

```

Other tips and tricks with rocgdb

Use `i th` to see a list of all active host threads. Currently viewing thread 1 (default)

```

jychang48@jychang48-workstation: ~/Downloads/hiptutorial/hip
23
24     const int id_l = id - 1;
25     const int id_r = id + 1;
26     const int id_d = id - stride;
27     const int id_u = id + stride;
28
29     AU[id] = (-U[id_l] + 2*U[id] - U[id_r])/(dx*dx) +
30             (-U[id_d] + 2*U[id] - U[id_u])/(dy*dy);
31 }
32 }
33
34 void LocalLaplacian(grid_t& grid, mesh_t& mesh,
35                    hipStream_t stream,
36                    dfloat* d_U,
37                    dfloat* d_AU) {
38
39     //there are (Nx-2)x(Ny-2) node on the interior of the mesh
40     int localNx = mesh.Nx-2;
41     int localNy = mesh.Ny-2;
42
43     int xthreads = 16;
44     int ythreads = 16;
45
46     dim3 threads(xthreads,ythreads,1);
47     dim3 blocks((localNx+xthreads-1)/xthreads,
48                (localNy+ythreads-1)/ythreads, 1);
49
50     hipLaunchKernelGGL(LocalLaplacianKernel,
51                        blocks,
52                        threads,
53                        0, stream,
54                        localNx, localNy, mesh.Nx,
55                        mesh.dx, mesh.dy,
56                        d_U, d_AU);
/home/jychang48/Downloads/hiptutorial/hip/Laplacian.cpp

:3:rocvirtual.cpp      :2521: 43544718546 us: [7f39e8446700]!      ShaderName :
      _Z11NormKernel2iPKdPd

:3:hip_memory.cpp      :331 : 43544719486 us: 151575: [7f3a02b7c8c0] hipMemcpy:
      Returned hipSuccess : : duration: 957 us
      Iteration: 0 - Residual: 0.022108
:3:hip_device_runtime.cpp :458 : 43544719507 us: 151575: [7f3a02b7c8c0] hipDeviceS
      ynchronize ( )
:3:hip_device_runtime.cpp :470 : 43544719512 us: 151575: [7f3a02b7c8c0] hipDeviceS
      ynchronize: Returned hipSuccess :
:3:hip_event.cpp        :310 : 43544719523 us: 151575: [7f3a02b7c8c0] hipEventRe
      cord ( event:0x948d00, stream:0x789710 )
:3:hip_event.cpp        :352 : 43544719530 us: 151575: [7f3a02b7c8c0] hipEventRe
      cord: Returned hipSuccess :

Thread 1 "Jacobi hip" hit Breakpoint 1, LocalLaplacian (grid=..., mesh=..., stream=0
x789710, d_U=0x7f3893e00000, d_AU=0x7f388bc00000) at Laplacian.cpp:40
(gdb) i th
      Id      Target Id                                     Frame
* 1      Thread 0x7f3a02b7c8c0 (LWP 151575) "Jacobi hip" LocalLaplacian (grid=..., mes
h=..., stream=0x789710, d_U=0x7f3893e00000, d_AU=0x7f388bc00000) at Laplacian.cpp:40
      2      Thread 0x7f39eab37700 (LWP 151591) "Jacobi hip" 0x00007f3a035b9aff in poll ()
      from /lib/x86_64-linux-gnu/libc.so.6
      3      Thread 0x7f39ea179700 (LWP 151592) "Jacobi hip" 0x00007f3a035c65ce in epoll_w
      ait () from /lib/x86_64-linux-gnu/libc.so.6
      4      Thread 0x7f39e253e700 (LWP 151593) "Jacobi hip" 0x00007f3a035bb50b in ioctl (
      ) from /lib/x86_64-linux-gnu/libc.so.6
      6      Thread 0x7f39e857f700 (LWP 151595) "Jacobi hip" 0x00007f3a04d23678 in do_fute
      x_wait.constprop () from /lib/x86_64-linux-gnu/libpthread.so.0
      7      Thread 0x7f39e8446700 (LWP 151596) "Jacobi hip" 0x00007f3a04d23678 in do_fute
      x_wait.constprop () from /lib/x86_64-linux-gnu/libpthread.so.0
      8      Thread 0x7f39e827f700 (LWP 151597) "Jacobi hip" 0x00007f3a04d23678 in do_fute
      x_wait.constprop () from /lib/x86_64-linux-gnu/libpthread.so.0
(gdb)

```

And more...

ROCgdb has several other features and capabilities not covered in this presentation.

See the following for much more:

[https://github.com/RadeonOpenCompute/ROCm/blob/master/ROCm Debugger User Guide v4.2.pdf](https://github.com/RadeonOpenCompute/ROCm/blob/master/ROCm%20Debugger%20User%20Guide%20v4.2.pdf)

[/opt/rocm-4.2.0/share/doc/rocgdb/rocannotate.pdf](#)

[/opt/rocm-4.2.0/share/doc/rocgdb/rocgdb.pdf](#)

[/opt/rocm-4.2.0/share/doc/rocgdb/rocrefcard.pdf](#)

[/opt/rocm-4.2.0/share/doc/rocgdb/rocstabs.pdf](#)

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.

Third-party content is licensed to you directly by the third party that owns the content and is not licensed to you by AMD. ALL LINKED THIRD-PARTY CONTENT IS PROVIDED “AS IS” WITHOUT A WARRANTY OF ANY KIND. USE OF SUCH THIRD-PARTY CONTENT IS DONE AT YOUR SOLE DISCRETION AND UNDER NO CIRCUMSTANCES WILL AMD BE LIABLE TO YOU FOR ANY THIRD-PARTY CONTENT. YOU ASSUME ALL RISK AND ARE SOLELY RESPONSIBLE FOR ANY DAMAGES THAT MAY ARISE FROM YOUR USE OF THIRD-PARTY CONTENT.

© 2021 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, ROCm, Radeon, Radeon Instinct and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. Other names are for informational purposes only and may be trademarks of their respective owners.

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

