# **GPU Profiling:**

#### Performance Timelines with Rocprof and Omnitrace

Suyash Tandon HIP Lecture Series – 04 October 2<sup>nd</sup>, 2023

AMD together we advance\_

### **AMD** Profilers



### Agenda – AMD Profilers with timeline profiling support



# **Background – AMD Profilers**



### **Background – AMD Profilers**



HIP Training Series 05

### Why use a timeline profile

GPU centric

- Visualize the application performance
- Understand the interactions between a program's operations
- Reveal hidden data transfer and other implicit operations
- Understand multiple stream dependencies and performance

CPU and GPU timelines

- Understand interaction between CPU and GPU
- Understand MPI communication process
- Track memory usage
- See hardware temperatures and frequencies

### What is ROC-Profiler?

• ROC-profiler (also referred to as **rocprof**) is the command line front-end for AMD's GPU profiling libraries

- Repo: <u>https://github.com/ROCm-Developer-Tools/rocprofiler</u>
- rocprof contains the central components allowing application traces and counter collection
  - Under constant development
- Distributed with ROCm
- The output of rocprof can be visualized in the Chrome browser with Perfetto (<u>https://ui.perfetto.dev/</u>)

### rocprof: Getting Started + Useful Flags

To get help:

#### \${ROCM\_PATH}/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 now 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
  - --kfd-trace to trace GPU driver calls
- Advanced usage
  - -m <metric file> Allows the user to define and collect custom metrics. See <u>rocprofiler/test/tool/\*.xml</u> on GitHub for examples.

### rocprof: Kernel Information

- rocprof can collect kernel(s) execution stats
  - \$ /opt/rocm/bin/rocprof --stats --basenames on <app with arguments>
- This will output two csv files:
  - results.csv: information per each call of the kernel
  - results.stats.csv: statistics grouped by each kernel
- Content of results.stats.csv provides the list of GPU kernels with their durations and percentage of total GPU time:

"Name","Calls","TotalDurationNs","AverageNs","Percentage" "JacobiIterationKernel",1000,556699359,556699,43.291753895270446 "NormKernel1",1001,430797387,430367,33.500980655394606 "LocalLaplacianKernel",1000,280014065,280014,21.775307970480817 "HaloLaplacianKernel",1000,14635177,14635,1.1381052818810995 "NormKernel2",1001,3770718,3766,0.2932300765671734 "\_\_amd\_rocclr\_fillBufferAligned.kd",1,8000,8000,0.0006221204058583505

• In a spreadsheet viewer, it is easier to read:

	A	В	С	D	E
1	Name	Calls	TotalDurationNs	AverageNs	Percentage
2	JacobiIterationKernel	1000	556699359	556699	43.2917538952704
3	NormKernel1	1001	430797387	430367	33.5009806553946
4	LocalLaplacianKernel	1000	280014065	280014	21.7753079704808
5	HaloLaplacianKernel	1000	14635177	14635	1.1381052818811
6	NormKernel2	1001	3770718	3766	0.293230076567173
7	amd rocclr fillBufferAligned	1	8000	8000	0.000622120405858

# rocprof: Collecting Application Traces

 rocprof can collect a variety of trace event types, and generate timelines in JSON format for use with Perfetto, 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

- You can combine modes like --hip-trace --hsa-trace
- If profiling OpenMP offload code, --hsa-trace is required to show HSA activity

### rocprof + Perfetto: Collecting and Visualizing Application Traces

- rocprof can collect traces
  - \$ /opt/rocm/bin/rocprof --hip-trace <app with arguments>

This will output a .json file that can be visualized using the chrome browser and Perfetto (<u>https://ui.perfetto.dev/</u>)

#### Perfetto – timeline visualization tool

- Both rocprof and Omnitrace currently use Perfetto for timeline presentation
- Perfetto is a built-in timeline visualization tool in the Chrome browser
- Original purpose was to profile android applications
- Open-source application from Google®
- Accessed through <u>https://ui.perfetto.dev/</u> to invoke tool

Works even if off-line No server interaction

• Opens file and reads local profile data in several formats

## rocprof + Perfetto: Collecting and Visualizing Application Traces

#### rocprof can collect traces

\$ /opt/rocm/bin/rocprof --hip-trace <app with arguments>

This will output a .json file that can be visualized using the chrome browser and Perfetto (<u>https://ui.perfetto.dev/</u>)

🎧 Perfetto	≡	8024159.4 s +	0.0 s	0.2 s	0.4 s	
Nevinetien		× <i>≡</i>				
Navigation	Â	∧ CPU HIP API 2				
Dpen trace file			hinMomoot			
Dpen with legacy UI			nipmeniset		HIP API Activity	У
Record new trace		Thread 140096				
	_					
Current Trace	· ^ -					
results.json (152 MB)						
Show timeline		Thread 0			GPU activ	vitv
🛨 Download		Thread 1				
		∧ COPY 1				
Metrics						
Info and stats		Thread 0				
		✔ GPU0 6		Copy activity (H	2D and D2H)	MDZ
Convert trace	^				toge	ether we

### **Perfetto: Visualizing Application Traces**

- Zoom in to see individual events
- Navigate trace using WASD keys

8022896.2 s +		0.268 s	0.269 s	0.270 s	0.271 s	0.272 s	0.273 s	0.274 s
× <i>≡</i>								
▲ CPU HIP API 2								
	hipMemcp	y hipL	aunchK			hipMe hipLaun	hipLa hip	Метсру
Thread 139878						I	1 1	I
▲ GPU2 8								
Thread 0						1	 barr	ier packet>   <
Thread 1			NormK				Lo Jacobil	t NormK Loc
▲ COPY 1								
Thread 0	CopyHostToDev	ice						
✔ GPU0 6								

 $\xrightarrow{\mathsf{D}}$ 





Zoom/

A

# **Perfetto: Kernel Information and Flow Events**

- Zoom and select a kernel, you can see the link to the HIP call launching the kernel
- Try to open the information for the kernel (button at bottom right)

8024159.4 s +		0.272 s		0.273 s		0.274 s	0.275 s		0.276 s
× =									
▲ CPU HIP API 2									
	hipM hipLaun	. hipLau	hipM	етсру		hipMemcpy	h	ірМетсру	h
Thread 140096			1)						
							1		
			/						
▲ GPU2 8									
Thread 0			<barrier< td=""><td>packet&gt;</td><td> ba</td><td>arrier packet&gt;</td><td><ba< td=""><td>rrier packet&gt;</td><td></td></ba<></td></barrier<>	packet>	 ba	arrier packet>	<ba< td=""><td>rrier packet&gt;</td><td></td></ba<>	rrier packet>	
Thread 1		Loc	Jacobilt	NormKe	Loc Ja	acobilt NormKe	Loc Ja	icobilt Nori	mKe L
COPY 1									
Current Selection Flow Events									Ţ ∽



### Perfetto: Kernel Information and Flow Events

Current Selection	Flow Events								Ť	~
Slice Details										
Name	Jacobilte double c	erationKernel(int, o onst*. double*. do	louble, double, double const*, uble*)	Preceding flow Slice	/S	<b>⊿</b> hi	pl aunchKernel			
Category	null	,,,		Delay		6us				
Start time	272ms 5	23us 999ns		Thread		NULL	(CPU HIP API 2)			
Duration	541us 🔺	Durat	ion	Arguments			· · · · · ·			
Thread duration	0s (0.009	%)		args						
Thread	1			BeginNs -	-	80241	59641088210	Kernel nan	ne and ard	s
Process	GPU2 8			Data 👻		NULL				
Slice ID	57			DurationN	S 🔻	54159	19			
				EndNs 👻		80241	59641629809	×		
				Name 👻		Jacob	ilterationKernel(i	nt, double, double,	double const <sup>3</sup>	*,
						14000	e const^, doubler	, double^)		
				pia 👻		14009	10			
				tia →		14009	10	Strea	m where k	erne
				dev-la 👻		1		was I	aunched ir	า
				queue-id	•	1				
				Stream-Iu	•					
Current Selection	Flow Events								Ţ	~
Flow events										
Direction	Duration	Connected Slice ID	Connected Slice Thread Out Name	Thread In	Process (	Dut	Process In	Flow Category	Flow Name	
Incoming	бus	52	hipLaunchKernel NULL	NULL	CPU HIP /	API 2	GPU2 8	DataFlow	dep	



#### rocprof: Collecting Application Traces with rocTX Markers and Regions

- rocprof can collect user defined regions or markers using rocTX •
- Annotate code with roctx regions: #include <roctx.h>

```
• • •
   roctxRangePush("reduce for c");
   reduce_function ();
   roctxRangePop();
```

Annotate code with roctx markers:

• • •

. . .

```
• • •
   roctxMark("start of some code");
   // some code
   roctxMark("end of some code");
```

Add roctx and roctracer libraries to link line: -L\${ROCM PATH}/lib -lroctx64 -lroctracer64

Profile with --roctx-range option:



### rocprof: Collecting Hardware Counters

- 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 rocprof\_counters.txt <app with args>
  - \$ cat rocprof\_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 containing all the requested counters for each invocation of every kernel

# rocprof: Commonly Used GPU Counters

VALUUtilization	The percentage of ALUs active in a wave. Low VALUUtilization is likely due to high divergence or a poorly sized grid
VALUBusy	The percentage of GPUTime vector ALU instructions are processed. Can be thought of as something like compute utilization
FetchSize	The total kilobytes fetched from global memory
WriteSize	The total kilobytes written to global memory
L2CacheHit	The percentage of fetch, write, atomic, and other instructions that hit the data in L2 cache
L2CacheHit MemUnitBusy	The percentage of fetch, write, atomic, and other instructions that hit the data in L2 cache The percentage of GPUTime the memory unit is active. The result includes the stall time
L2CacheHit MemUnitBusy MemUnitStalled	The percentage of fetch, write, atomic, and other instructions that hit the data in L2 cache The percentage of GPUTime the memory unit is active. The result includes the stall time The percentage of GPUTime the memory unit is stalled

Full list at: <u>https://github.com/ROCm-Developer-Tools/rocprofiler/blob/amd-master/test/tool/metrics.xml</u>

### **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 using 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?
  - \$ /opt/rocm/bin/rocprof --timestamp on -i rocprof\_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.stats.csv file will include a kernel duration column
    - Note: the duration is aggregated over repeated calls to the same kernel

### rocprof: 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> ./Jacobi\_hip -g <x> <y>
- Invoke the profiler by executing:

```
mpiexec -np <n> rocprof <rocprof_options> ./Jacobi_hip -g <x> <y>
or
srun --ntasks=n rocprof <rocprof options> ./Jacobi hip -g <x> <y>
```

 By directing output files from each rank to different directories, we can collect traces for each rank separately

Use a helper script for this, and run your program as shown below:
 mpiexec -np <n> helper\_rocprof.sh ./Jacobi\_hip -g <x> <y>

Multi-node profiling currently isn't supported

# **Profiling Multiple MPI Ranks**

\$cat helper\_rocprof.sh

```
#!/bin/bash
set -euo pipefail
# depends on ROCM PATH being set outside; input arguments are the output directory & the name
outdir="$1"
name="$2"
if [[ -n ${OMPI_COMM_WORLD_RANK+z} ]]; then
    # mpich
    export MPI RANK=${OMPI COMM WORLD RANK}
elif [[ -n ${MV2 COMM WORLD RANK+z} ]]; then
    # ompi
    export MPI RANK=${MV2 COMM WORLD RANK}
elif [[ -n ${SLURM PROCID+z} ]]; then
    export MPI RANK=${SLURM PROCID}
else
    echo "Unknown MPI layer detected! Must use OpenMPI, MVAPICH, or SLURM"
    exit 1
fi
rocprof="${ROCM PATH}/bin/rocprof"
                                                               Output directory per rank
pid="$$"
                                                               Filenames annotated with rank as well
outdir="${outdir}/rank_${pid}_${MPI_RANK}"
outfile="${name}_${pid}_${MPI_RANK}.csv"
${rocprof} -d ${outdir} --hsa-trace -o ${outdir}/${outfile} ("${@:3}"
                                                                          Application and its arguments
                                                               Trace mode
```

### rocprof: Profiling Overhead

- As with every profiling tool, there is an overhead
- The percentage of the overhead depends on the profiling options used
  For example, tracing is faster than hardware counter collection
- When collecting many counters, the collection may require multiple passes
- With rocTX markers/regions, tracing can take longer and the output may be large
  - Sometimes too large to visualize
- The more data collected, the more the overhead of profiling
  - Depends on the application and options used

# Summary

- rocprof is the open source, command line AMD GPU profiling tool distributed with ROCm
- Many other tools are built over rocprof
- rocprof provides tracing of GPU kernels, HIP API, HSA API and Copy activity
- rocprof can be used to collect GPU hardware counters with additional overhead
- JSON Traces can be viewed in Perfetto UI
- Other output files are in text/CSV format

### Agenda – AMD Profilers with timeline profiling support



# **Omnitrace: Application Profiling, Tracing, and Analysis**



Refer to current documentation for recent updates



# Installation (if required)



To use pre-built binaries, select the version that matches your operating system, ROCm version, etc.

	Ļ

Select OpenSuse/RHEL operating system for HPE/AMD system: omnitrace-1.10.1-rhel-9.0-ROCm-50500-PAPI-OMPT-Python3.sh



There are .rpm and .deb fpackages along with .sh scripts for installation.



Full documentation: <u>https://amdresearch.github.io/omnitrace/</u>

export OMNITRACE\_VERSION=latest
export ROCM\_VERSION=5.6.0
export OMNITRACE\_INSTALL\_DIR=</path/to/your/omnitrace/install>
wget <u>https://github.com/AMDResearch/omnitrace/releases/\${OMNITRACE\_VERSION}/download/omnitrace-install.py
python3 omnitrace-install.py -p \${OMNITRACE\_INSTALL\_DIR} --rocm \${ROCM\_VERSION}</u>

Set up environment:
source \${OMNITRACE\_INSTALL\_DIR}/share/omnitrace/setup-env.sh

Note: If installing from source, remember to clone the omnitrace repo recursively



# **Omnitrace instrumentation Modes**



Basic command-line syntax:
<pre>\$ omnitrace [omnitrace-options] <cmd> <args></args></cmd></pre>
For more information or help use -h/help/? flags:
<pre>\$ omnitrace -h</pre>
Can also execute on systems using a job scheduler. For example, with SLURM, an interactive session can be used as:
<pre>\$ srun [options] omnitrace [omnitrace-options] <cmd> <args></args></cmd></pre>

For problems, create an issue here: <u>https://github.com/AMDResearch/omnitrace/issues</u> Documentation: <u>https://amdresearch.github.io/omnitrace/</u>

### **Omnitrace Configuration**

#### \$ omnitrace-avail --categories [options]

Get more information about run-time settings, data collection capabilities, and available hardware counters. For more information or help use -h/--help flags:

#### \$ omnitrace-avail -h

Collect information for omnitrace-related settings using shorthand -c for --categories :

\$ omnitrace-avail -c perfetto

\$ omnitrace-avail -c perfetto		
ENVIRONMENT VARIABLE	VALUE	CATEGORIES
OMNITRACE_PERFETTO_BACKEND OMNITRACE_PERFETTO_BUFFER_SIZE_KB OMNITRACE_PERFETTO_FILL_POLICY OMNITRACE_TRACE_DELAY OMNITRACE_TRACE_DURATION OMNITRACE_TRACE_PERIODS OMNITRACE_TRACE_PERIOD_CLOCK_ID OMNITRACE_USE_PERFETTO	inprocess 1024000 discard 0 0 CLOCK_REALTIME true	<pre>custom, libomnitrace, omnitrace, perfetto custom, data, libomnitrace, omnitrace, perfetto custom, data, libomnitrace, omnitrace, perfetto custom, libomnitrace, omnitrace, perfetto, profile, timemory, trace backend, custom, libomnitrace, omnitrace, perfetto</pre>

Shows all runtime settings that may be tuned for perfetto

### **Omnitrace Configuration**

#### \$ omnitrace-avail --categories [options]

Get more information about run-time settings, data collection capabilities, and available hardware counters. For more information or help use -h/--help/? flags:

#### \$ omnitrace-avail -h

Collect information for omnitrace-related settings using shorthand -c for --categories :

\$ omnitrace-avail -c omnitrace

For brief description, use the options:

#### \$ omnitrace-avail -bd

ENVIRONMENT VARIABLE	DESCRIPTION	i I
OMNITRACE_CAUSAL_BINARY_EXCLUDE	Excludes binaries matching the list of provided regexes from causal experiments (separated by tab, sem)	
UMNITRACE_CAUSAL_BINARY_SCOPE	Limits causal experiments to the binaries matching the provided list of regular expressions (separated	
	Length of time to wait (in seconds) before starting the first causal experiment	
OMNITRACE CAUSAL DURATION	Length of time to perform causal experimentation (in seconds) after the first experiment has started	
	Excludes functions matching the tist of provided regress from causal experiments (separated by tab, se	
OMNITRACE_CAUSAL_IONCITON_SCOPE	Seed for random number generator which causal profiling (separated by tab, semi-cotor), and/or quotes (SI)	
OMNITRACE CAUSAL RANDOM SELD	See for random number generator which setects speedups and experiments - prease note that the times [ Excludes source files or source file line on pair (i e files or sfiles discounce) matching the list of	
OMNITRACE CAUSAL SOURCE SCOPE	Limits causal experiments to the source files or source file $+$ lineno pair (i $e$ sfiles or sfiles of the source files)	
OMNITRACE CONFIG FILE	Configuration file for omnitace	1
OMNITRACE CRITICAL TRACE	Enable generation of the critical trace	í.
OMNITRACE ENABLED	Activation state of timemory	í
OMNITRACE OUTPUT PATH	Explicitly specify the output folder for results	i I
OMNITRACE OUTPUT PREFIX	Explicitly specify a prefix for all output files	i I
OMNITRACE PAPI EVENTS	PAPI presets and events to collect (see also: papi avail)	i I
OMNITRACE PERFETTO BACKEND	Specify the perfetto backend to activate. Options are: 'inprocess', 'system', or 'all'	i
OMNITRACE_PERFETTO_BUFFER_SIZE_KB	Size of perfetto buffer (in KB)	Í.
OMNITRACE_PERFETTO_FILL_POLICY	Behavior when perfetto buffer is full. 'discard' will ignore new entries, 'ring_buffer' will overwrite	
OMNITRACE_PROCESS_SAMPLING_DURATION	If > 0.0, time (in seconds) to sample before stopping. If less than zero, uses OMNITRACE_SAMPLING_DURA	1
OMNITRACE_PROCESS_SAMPLING_FREQ	Number of measurements per second when OMNITTRACE_USE_PROCESS_SAMPLING=ON. If set to zero, uses OMNITR	
OMNITRACE_ROCM_EVENTS	ROCm hardware counters. Use ':device=N' syntax to specify collection on device number N, e.g. ':device	
OMNITRACE_SAMPLING_CPUS	CPUs to collect frequency information for. Values should be separated by commas and can be explicit or	
UMNITRACE_SAMPLING_DELAY	lime (in seconds) to wait before the first sampling signal is delivered, increasing this value can fix	
UMNITRACE_SAMPLING_DURATION	11 > 0.0, time (in seconds) to sample before stopping	
UMNITRACE_SAMPLING_FREQ	Number of software interrupts per second when UMNINIACE USE SAMPLING=ON	
UMNITRACE SAMPLING GPUS	, Devices to query when umnifikate use kuth shi=un, values should be separated by commas and can be expli I	

#### Create a config file

Create a config file in \$HOME:

\$ omnitrace-avail -G \$HOME/.omnitrace.cfg

To add description of all variables and settings, use:

#### \$ omnitrace-avail -G \$HOME/.omnitrace.cfg --all

Modify the config file \$HOME/.omnitrace.cfg as desired to enable and change settings:

#### <snip>

MNITRACE_USE_PERFETTO		=	true
MNITRACE_USE_TIMEMORY		=	true
MNITRACE_USE_SAMPLING		=	false
MNITRACE_USE_ROCTRACER		=	true
MNITRACE_USE_ROCM_SMI		=	true
MNITRACE_USE_KOKKOSP		=	false
MNITRACE_USE_CAUSAL		=	false
MNITRACE_USE_MPIP		=	true
MNITRACE_USE_PID	_	=	true
MNITRACE_USE_ROCPROFILER		=	true
MNITRACE_USE_ROCTX	Contents of the config	g i	file

Declare which config file to use by setting the environment:

\$ export OMNITRACE\_CONFIG\_FILE=/path-

to/.omnitrace.cfg

# **Dynamic Instrumentation**

**Runtime Instrumentation** 



### **Dynamic Instrumentation – Jacobi Example**

Clone jacobi example:	Parsing libraries
<pre>\$ git clone <u>https://github.com/amd/HPCTrainingExamples.git</u> \$ cd HPCTrainingExamples/HTP/jacobj</pre>	
Requires ROCm and MPI install, compile:	<pre>[omnitrace][exe] [internal] parsing library: '/usr/lib64/libutil-2.28.so' [omnitrace][exe] [internal] parsing library: '/usr/lib64/libz.so.1.2.11' [omnitrace][exe] [internal] binary info processing required 0.322 sec and 70.724 MB [omnitrace][exe] Processing 72 modules</pre>
\$ make	[omnitrace][exe] Processing 72 modules Done (0.101 sec, 12.084 MB) [ommitrace][exe] Found 'MPT Thit' in '/home/ssitaram/git/HPCTrainingExamples/HIP/jacobi/lacobi hip' Enabling MPT support
Run the non-instrumented code on a single GPU as:	[omnitrace][exe] Finding instrumentation functions [omnitrace][exe] 2 instrumented funcs in//orte/orted/orted_submit.c [omnitrace][exe] 1 instrumented funcs in liband comgr.so.2.4.50403
<pre>\$ time .mpirun -np 1 ./Jacobi_hip -g 1 1</pre>	[omnitrace][exe] 15 Instrumented funcs in Libamon1p64.50.5.4.50403 Functions instrumented
real 0m2.115s	[omnitrace][exe] 10 instrumented funcs in libmpi.so.40.20.3 [omnitrace][exe] 8 instrumented funcs in libopen-pal.so.40.20.3
	[omnitrace][exe] 17 instrumented funcs in libopen-rte.so.40.20.3
Dynamic instrumentation	[omnitrace][exe] Outputting 'omnitrace-Jacobi_hip-output/2023-03-14 17.24/instrumentation/available.json' Done [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip-output/2023-03-14_17.24/instrumentation/available.txt' Done [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip-output/2023-03-14_17.24/instrumentation/instrumented.json' Done [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip-output/2023-03-14_17.24/instrumentation/instrumented.txt' Done [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip-output/2023-03-14_17.24/instrumentation/instrumented.txt' Done
<pre>\$ time mpirun -np 1 omnitrace-instrument/Jacobi_hip -g 1 1</pre>	[omnitrace][exe] Outputting 'omnitrace-Jacobi_hip-output/2023-03-14_17.24/instrumentation/excluded.txt' Done [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip-output/2023-03-14_17.24/instrumentation/overlapping.json' Done [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip-output/2023-03-14_17.24/instrumentation/overlapping.json' Done [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip-output/2023-03-14_17.24/instrumentation/overlapping.txt' Done [omnitrace][exe] Executing
real 1m45.742s	[omnitrace][1649192][omnitrace_init_tooling] Instrumentation mode: Trace Outputs that will be created
Extra time is the overhead of dyninst reading every binary that is loaded, not overhead of omnitrace during app execution	0     0

# **Dynamic Instrumentation – Jacobi Example**

Clone jacobi example: \$ git clone <u>https://github.com/amd/HPCTrainingExamples.git</u> \$ cd HPCTrainingExamples/HIP/jacobi	<pre>[available] HaloExchange.cpp: [available] [HaloExchange.cold.21][14] [available] [HaloExchange][1267] [available] [ GLOBAL sub I HaloExchange.cpp][8]</pre>
Requires ROCm and MPI install, compile: \$ make Run the non-instrumented code on a single GPU as:	<pre>[available] Input.cpp: [available] [ExtractNumber][19] [available] [FindAndClearArgument][38] [available] [ParseCommandLineArguments][206] [available] [PrintUsage][12]</pre>
<pre>\$ time .mpirun -np 1 ./Jacobi_hip -g 1 1 real 0m2 115s</pre>	<pre>[available] JacobiIteration.cpp: [available] [JacobiIteration][71]</pre>
	[available] JacobiMain.cpp: [available] [main.cold.0][5] [available] [main][35] Functions found in each module
Dynamic instrumentation	[available] JacobiRun.cpp: [available] [Jacobi_t::Run][155]
-g 1 1 real 1m45.742s	<pre>[available] JacobiSetup.cpp: [available] [FormatNumber][53] [available] [Jacobi_t::ApplyTopology][234] [available] [Jacobi_t::CreateMesh][459]</pre>
Available functions to instrument: \$ mpirun -np 1 omnitrace-instrument -v 1simulate print-available functions/Jacobi_hip -g 1 1	<pre>[available] [Jacobi_t::InitializeData][552] [available] [Jacobi_t::Jacobi_t.cold.30][15] [available] [Jacobi_t::Jacobi_t][1043] [available] [Jacobi_t::PrintResults][107] [available] [Jacobi_t::~Jacobi_t][167] [available] [PrintResults][134]</pre>
Here -v gives a verbose output from omnitrace	<pre>[available] [_GLUBALsub_I_JacobiSetup.cpp][8] [available] [std::_cxx11::basic_stringbuf<char, std::char_traits<char="">, std::allocator <char> &gt;::~basic_stringbuf][16] [available] [std::_cxx11::basic_stringbuf<char, std::char_traits<char="">, std::allocator <char> &gt;::~basic_stringbuf][18]</char></char,></char></char,></pre>

The simulate flag does not run the executable, but only demonstrates the available functions

### **Dynamic Instrumentation – Jacobi Example**

Clone jacobi example:	<pre>[omnitrace][exe] [internal] parsing library: '/opt/rocm-5.4.3/lib/librocm_smi64.so.5.0.50403' [omnitrace][exe] [internal] parsing library: '/opt/rocm-5.4.3/lib/librocmTools.so.1.5.0' [ompitrace][eve] [internal] parsing library: '/opt/rocm 5.4.3/lib/librocmTools.so.1.0.50403'</pre>
<b>\$</b> git clone https://github.com/amd/HPCTrainingExamples.git	[omnitrace][exe] [internal] parsing library: '/opt/rocm-5.4.3/lib/libroctracer64.so.4.1.0'
	[omnitrace][exe] [internal] parsing library: '/opt/rocm=5.4.3/lib/libroctx64.so.4.1.0'
\$ cd HPCIrainingExamples/HIP/jacobi	[omnitrace][exe] [internal] parsing library: '/share/contrib-modules/omnitrace/omnitrace18.0/lib/libomnitrace-01.50.18.0' [omnitrace][exe] [internal] parsing library: '/share/contrib-modules/omnitrace/omnitrace18.0/lib/libomnitrace-t so 11.0 1'
	[ommitrace][exe] [internal] parsing library: //share/contrib-modules/ommitrace/ommitrace1.8.0/lib/libommitrace-user.so.1.8.0'
Requires ROCm and MPI install compile:	[omnitrace][exe] [internal] parsing library: '/share/contrib-modules/omnitrace/omnitrace1.8.0/lib/omnitrace/libcommon.so.11.0.1'
Requires Room and with thistail, complete.	[omnitrace][exe] [internal] parsing library: '/share/contrib-modules/omnitrace/annitrace1.8.0/lib/omnitrace/libdw-0.182.so'
¢ make	[ommitrace][exe] [internal] parsing library: '/share/contrib-modules/ommitrace/ommitrace1.8.0/Lib/ommitrace/Libert-0.182.50'
	[ommitrace][exe] [internal] parsing library: '/share/contrib-modules/ommitrace/ambitrace/lib/mitrace/lib/fm.so.4.11.1'
	[omnitrace][exe] [internal] parsing library: '/share/contrib-modules/omnitrace/omnitrace1.8.0/lib/omnitrace/libtbb.so.2'
Run the non-instrumented code on a single GPU as:	[omnitrace][exe] [internal] parsing library: '/share/contrib-modules/omnitrace/omnitrace1.8.0/lib/omnitrace/libtbbmalloc.so.2'
	[omnitrace][exe] [internal] parsing library: '/share/contrib-modules/omnitrace/omnitrace1.8.0/Lib/omnitrace/libLobmatuco_proxy.so.2' [omnitrace][exe] [internal] parsing library: '/share/contrib-modules/omnitrace/omnitrace1.8.0/Lib/omnitrace/lib
\$ time .mpirun -np 1 ./Jacobi hip -g 1 1	[omnitrace][exe] [internal] parsing library: ', on the backstown of the construction of the own of
	[omnitrace][exe] [internal] parsing library: '/usr/lib64/libBrokenLocale-2.28.so'
real 0m2.115s	[omnitrace][exe] [internal] parsing library: '/usr/lib64/libanl-2.28.so'
	[omnitrace][exe] [internal] parsing library: '/usr/libo4/libc-22.28.50'
	[omnitrace][exe] [internal] parsing library: /usr/lib64/libd/-2.28.so'
	[omnitrace][exe] [internal] parsing library: '/usr/lib64/libgcc s-8-20210514.so.1'
	[omnitrace][exe] [internal] parsing library: '/usr/lib64/libnss_compat-2.28.so'
Dynamic instrumentation	[omnitrace][exe] [internal] parsing library: '/usr/lib64/libnss dns-2.28.so'
	[omnitrace][exe] [internal] parsing library: /usr/lib64/libnthread-2.28.so'
<pre>\$ time mninun -nn 1 omnitrace-instrument /lacobi hin</pre>	[omitrace][exe] [internal] parsing library: //usr/lib64/libresolv-2.28.so' Only these two functions
<sup>*</sup> cline inprivation inprivation in the constrained and the const	[omnitrace][exe] [internal] parsing library: '/usr/lib64/librt-2.28.so'
-g 1 1	[omitrace][exe] [internal] parsing library: '/usr/lib64/libstdc++.so.6.0.25' are shown to be
5	[omnitrace][exe] [internal] parsing library: /usr/libo4/libutil_238.so]
	[omnitrace][exe] [internal] parsing library: /usr/lib64/libz.so.1.2.11' Instrumented
real 1m45.742s	[omnitrace][exe] [internal] binary info processing required 0.257 sec and 66.740 MB
	[omnitrace][exe] Processing 72 modules
Available functions to instrument:	[omnitrace][exe] Frocessing /2 modules Done (0.039 sec, 11.030 MB) [omnitrace][exe] Found 'MPI Init' in '(home/scitaram/dit/HPCTrainingExamples/HTP/jacobi/jacobi hin' Fnabling MPI support
	[omitrace][exe] Finding instrumentation functions
<pre>\$ mpirun -np 1 omnitrace-instrument -v 1simulate</pre>	[omnitrace][exe] 1 instrumented funcs in JacobiIteration.cpp
nnint available functions /Jacobi hin g 1 1	[omnitrace][exe] 1 instrumented funcs in JacobiRun.cpp
princ-available functions/Jacobi_nip -g i i	[omnitrace][exe] I instrumented funcs in Jacon np [omnitrace][exe] I instrumented funcs in Jacon here in Jacon her
	[ommitrace][exe] or finitrace lacobi his-output/2023-03-15 12.40/instrumentation/available.ison' Done
Custom include/exclude functions* with -I or -F, resp. For e.g.	[omnitrace][exe] Outputting 'omnitrace-Jacobi_hip-output/2023-03-15_12.40/instrumentation/available.txt' Done
	[omnitrace][exe] Outputting 'omnitrace-Jacobi hip-output/2023-03-15_12.40/instrumentation/instrumented.json' Done
<pre>\$ mpirun -np 1 omnitrace-instrument -v 1 -I</pre>	[omnitrace][exe] outputting 'omnitrace-Jacobi hip-output/2023-03-15 12.40/instrumentation/instrumented.txt' Done
	[omitrace][exe] outputting 'omitrace-Jacobi hip-output/2023-03-12.40/instrumentation/excluded.txt' Done
'Jacobi_t::Run' 'Jacobilteration'/Jacobi_hip -g 1 1	[omnitrace][exe] Outputting 'omnitrace-Jacobi_hip-output/2023-03-15_12.40/instrumentation/overlapping.json' Done
	[omnitrace][exe] Outputting 'omnitrace-Jacobi_hip-output/2023-03-15_12.40/instrumentation/overlapping.txt' Done

together we advance\_



# **Dynamic Instrumentation**

**Binary Rewrite** 



# **Binary Rewrite – Jacobi Example**

Binary Rewrite	<pre>[omnitrace][exe] [internal] parsing library: '/usr/lib64/libgcc_s-8-20210514.so.1' [omnitrace][exe] [internal] parsing library: '/usr/lib64/libnss_compat-2.28.so' [omnitrace][exe] [internal] parsing library: '/usr/lib64/libnss_files-2.28.so' [omnitrace][exe] [internal] parsing library: '/usr/lib64/libnss_files-2.28.so'</pre>
<pre>\$ omnitrace-instrument [omnitrace-options] -o <new-name- of-exec&gt; <cmd> <args></args></cmd></new-name- </pre>	<pre>[omnitrace][exe] [internal] parsing library: '/usr/lib64/libresolv-2.28.so' [omnitrace][exe] [internal] parsing library: '/usr/lib64/librt-2.28.so' [omnitrace][exe] [internal] parsing library: '/usr/lib64/libtdc++.so.6.0.25' [omnitrace][exe] [internal] parsing library: '/usr/lib64/libthread_db-1.0.so' [omnitrace][exe] [internal] parsing library: '/usr/lib64/libtl-2.28.so'</pre>
Generating a new executable/library with instrumentation built-in:	<pre>[omnitrace][exe] [internal] parsing library: '/usr/lib64/libz.so.1.2.11' [omnitrace][exe] [internal] binary info processing required 0.666 sec and 110.500 MB [omnitrace][exe] Processing 9 modules [omnitrace][exe] Processing 9 modules Done (0.001 sec, 0.000 MB) [omnitrace][exe] Found 'MPI nit' in '/omme/csitaram/oit/HPCTrainingExamples/HIP/iacobi/lacobi hin' Enabling MPI support</pre>
This new binary will have instrumented functions	<pre>[omnitrace][exe] Finding instrumentation functions [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip.inst-output/2023-03-15_12.57/instrumentation/available.json' Done [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip.inst-output/2023-03-15_12.57/instrumentation/available.txt' Done [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip.inst-output/2023-03-15_12.57/instrumentation/instrumented.json' Done [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip.inst-output/2023-03-15_12.57/instrumentation/instrumented.json' Done [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip.inst-output/2023-03-15_12.57/instrumentation/instrumented.txt' Done [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip.inst-output/2023-03-15_12.57/instrumentation/excluded.json' Done [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip.inst-output/2023-03-15_12.57/instrumentation/excluded.json' Done</pre>
	<pre>[omnitrace][exe] outputting 'omnitrace-Jacobi_hip.inst-output/2023-05-15_12_57/instrumentation/exctuded.txt bone [omnitrace][exe] Outputting 'omnitrace-Jacobi_hip.inst-output/2023-03-15_12_57/instrumentation/overlapping.txt' Done [omnitrace][exe] [omnitrace][exe] [exe] The instrumented executable image is stored in '/home/ssitaram/git/HPCTrainingExamples/HIP/jacobi/Jacobi_hip.inst' [omnitrace][exe] Getting linked libraries for /home/ssitaram/git/HPCTrainingExamples/HIP/jacobi/Jacobi_hip [omnitrace][exe] Consider instrumenting the relevant libraries [omnitrace][exe] [exe] [exe</pre>
	[omnitrace][exe] /lib64/libpthread.so.0 [omnitrace][exe] /lib64/libm.so.6
Subroutine Instrumentation Default instrumentation is main function and functions of 1024 instructions and more (for CPU)	[omnitrace][exe]/lib64/librt.so.1[omnitrace][exe]/opt/rocm-5.4.3//lib/libroctx64.so.4[omnitrace][exe]/opt/rocm-5.4.3//lib/libroctracer64.so.4[omnitrace][exe]/opt/rocm-5.4.3//lib/libroctracer64.so.4[omnitrace][exe]/opt/rocm-5.4.3//lib/libroctracer64.so.5[omnitrace][exe]/lib64/libstdc++.so.6[omnitrace][exe]/lib64/libc.so.6[omnitrace][exe]/lib64/libc.so.6[omnitrace][exe]/lib64/libc.so.6
To instrument routines with 50 or more cycles, add option "-i 50" (more	

overhead)

# **Binary Rewrite – Jacobi Example**

#### **Binary Rewrite**

\$ omnitrace-instrument [omnitrace-options] -o <new-</pre> name-of-exec> -- <<u>CMD> <ARGS></u>

#### Generating a new /library with instrumentation built-in:

```
$ omnitrace-instrument -o Jacobi hip.inst --
./Jacobi hip
```

#### Run the instrumented binary:

\$ mpirun -np 1 omnitrace-run -- ./Jacobi hip.inst -g

#### subroutine instrumentation

Default instrumentation is main function and functions of 1024 instructions and more (for CPU)

To instrument routines with 50 or more cycles, add option "-i 50" (more overhead)

Binary rewrite is recommended for runs with multiple ranks as omnitrace produces separate output files for each rank

#### omnitrace][3624331][omnitrace init tooling] Instrumentation mode: Trace



#### omnitrace v1.8.0

953.7651 perfetto.cc:58656 Configured tracing session 1, #sources:1, duration:0 ms, #buffers:1, total buffer si e:1024000 KB, total sessions:1, uid:0 session name: "" opology size: 1 x 1 \_ocal domain size (current node): 4096 x 4096 omnitrace][0][pid=3624331] MPI rank: 0 (0), MPI size: 1 (1) Global domain size (all nodes): 4096 x 4096 Rank 0 selecting device 0 on host TheraC60 Starting Jacobi run. [teration: 0 - Residual: 0.022108 [teration: 100 - Residual: 0.000625 teration: 200 - Residual: 0.000371 teration: 300 - Residual: 0.000274 teration: 400 - Residual: 0.000221 teration: 500 - Residual: 0.000187 Generates traces for application run teration: 600 - Residual: 0.000163 teration: 700 - Residual: 0.000145 teration: 800 - Residual: 0.000131 teration: 900 - Residual: 0.000120 [teration: 1000 - Residual: 0.000111 Stopped after 1000 iterations with residue 0.000111 Total Jacobi run time: 1.5470 sec. Measured lattice updates: 10.84 GLU/s (total), 10.84 GLU/s (per process) Measured FLOPS: 184.36 GFLOPS (total), 184.36 GFLOPS (per process) Measured device bandwidth: 1.04 TB/s (total), 1.04 TB/s (per process) omnitrace][3624331][0][omnitrace finalize] finalizing... omnitrace][3624331][0][omnitrace finalize] omnitrace][3624331][0][omnitrace finalize] omnitrace/process/3624331 : 2.364423 sec wall clock, 645.964 MB peak rss, 388.739 MB page rss, 4.330000 sec cpu clock, 183.1 % cpu util [laps: 1] omnitrace][3624331][0][omnitrace finalize] omnitrace/process/3624331/thread/0 : 2.355893 sec wall clock. 1.293230 sec thread cpu clock, 54.9 % thread cpu util, 645.964 MB peak rss [laps: 1] omnitrace][3624331][0][omnitrace finalize] omnitrace/process/3624331/thread/1 : 2.345084 sec wall clock, 0.000261 sec thread cpu clock, 0.0 % thread cpu util, 642.676 MB peak rss [laps: 1]

omnitrace][3624331][0][omnitrace finalize]

omnitrace][3624331][0][omnitrace finalize] Finalizing perfetto...

### **List of Instrumented GPU Functions**

#### \$ cat omnitrace-Jacobi\_hip.inst-output/2023-03-15\_13.57/roctracer-0.txt

ROCM TRACER (ACTIVITY API)							
LABEL	COUNT	DEPTH	METRIC	UNITS	SUM	MEAN	% SELF
10>>> pthread create	i 1	i 0	roctracer	sec	0.000353	0.000353	0.0
>>>   start thread	i 1	i 1	roctracer	sec	2.344864	2.344864	100.0
0>>> hipInit	j 1	i o	roctracer	sec	0.000000	0.000000	0.0
0>>> hipGetDeviceCount	j 1	j O	roctracer	sec	0.000000	0.000000	0.0
0>>> hipSetDevice	j 1	j 0	roctracer	sec	0.000000	0.000000	0.0
0>>> hipHostMalloc	3	j 0	roctracer	sec	0.000000	0.000000	0.0
0>>> hipMalloc	j 7	j 0	roctracer	sec	0.000000	0.000000	0.0
0>>> hipMemset	1	0	roctracer	sec	0.000000	0.000000	0.0
0>>> hipStreamCreate	2	0	roctracer	sec	0.000000	0.000000	0.0
0>>> hipMemcpy	1005	0	roctracer	sec	0.000000	0.000000	0.0
<pre> 0&gt;&gt;&gt;  _LocalLaplacianKernel(int, int, int, double, double, double const*, double*)</pre>	999	1	roctracer	sec	0.279368	0.000280	100.0
<pre> 0&gt;&gt;&gt;  _HaloLaplacianKernel(int, int, int, double, double, double const*, double const*, double*)</pre>	990	1	roctracer	sec	0.014761	0.000015	100.0
<pre> 0&gt;&gt;&gt;  _JacobiIterationKernel(int, double, double, double const*, double const*, double*, double*)</pre>	959	1	roctracer	sec	0.531156	0.000554	100.0
<pre> 0&gt;&gt;&gt;  _NormKernel1(int, double, double, double const*, double*)</pre>	997	1	roctracer	sec	0.430196	0.000431	100.0
0>>>  _NormKernel2(int, double const*, double*)	999	1	roctracer	sec	0.004342	0.000004	100.0
0>>> hipEventCreate	2	0	roctracer	sec	0.000000	0.000000	0.0
0>>> hipLaunchKernel	5002	0	roctracer	sec	0.000000	0.000000	0.0
<pre> 0&gt;&gt;&gt;  _JacobiIterationKernel(int, double, double, double const*, double const*, double*, double*)</pre>	1	1	roctracer	sec	0.000552	0.000552	100.0
<pre> 0&gt;&gt;&gt;  _NormKernel1(int, double, double, double const*, double*)</pre>	1	1	roctracer	sec	0.000425	0.000425	100.0
0>>> hipDeviceSynchronize	1001	0	roctracer	sec	0.000000	0.000000	0.0
<pre> 0&gt;&gt;&gt;  _NormKernel1(int, double, double, double const*, double*)</pre>	2	1	roctracer	sec	0.000850	0.000425	100.0
0>>>  _NormKernel2(int, double const*, double*)	1	1	roctracer	sec	0.000004	0.000004	100.0
<pre> 0&gt;&gt;&gt;  _HaloLaplacianKernel(int, int, int, double, double, double const*, double const*, double*)</pre>	9	1	roctracer	sec	0.000133	0.000015	100.0
<pre> 0&gt;&gt;&gt;  _JacobiIterationKernel(int, double, double, double const*, double const*, double*, double*)</pre>	40	1	roctracer	sec	0.022204	0.000555	100.0
<pre> 0&gt;&gt;&gt;  _LocalLaplacianKernel(int, int, int, double, double, double const*, double*)</pre>	1	1	roctracer	sec	0.000281	0.000281	100.0
0>>> hipEventRecord	2000	0	roctracer	sec	0.000000	0.000000	0.0
0>>> hipStreamSynchronize	2000	0	roctracer	sec	0.000000	0.000000	0.0
0>>> hipEventElapsedTime	1000	0	roctracer	sec	0.000000	0.000000	0.0
<pre> 0&gt;&gt;&gt;  _HaloLaplacianKernel(int, int, int, double, double, double const*, double const*, double*)</pre>	1	1	roctracer	sec	0.000015	0.000015	100.0
0>>> hipFree	4	0	roctracer	sec	0.000000	0.000000	0.0
0>>> hipHostFree Roctracer-0.txt shows duration of	2	0	roctracer	sec	0.000000	0.000000	0.0
HIP APL calls and GPU kernels							

AMD together we advance\_

# **Visualizing Trace**

#### Use Perfetto

Copy perfetto-trace-0.proto to your laptop, go to <u>https://ui.perfetto.dev/</u>, **C**lick "Open trace file", select perfetto-trace-0.proto

4676921.1 s +		0.0 s	0.2 s	0.4 s	0.6 s	0.8 s	1.0 s	1.2 s	1.4 s	1.6 s	1.8 s	2.0 s	2.2 s
× =													
Clock Snapshots metric					Å								A
▲ ./Jacobi_hip.inst 3624331													
Jacobi_hip.inst 3624331		MPi_In	it	Jacobi_t::Jacobi_t CreateMesh::Ini hipMemset	::Top Lev		main						
CPU Context Switches (S)	~	25 K							Tra	ces of CP	U functio	ns	
CPU Frequency [0] (S)	~	5 K										10	
CPU Frequency [1] (S)	~	2.5 K											
CPU Frequency [2] (S)	~	2.5 K											
CPU Frequency [3] (S)	~	2.5 K											
CPU Frequency [4] (S)	~	2.5 K											
CPU Frequency [5] (S)	~	2.5 K											
CPU Frequency [6] (S)	~	2.5 K											
CPU Frequency [7] (S)	~	2.5 K											
CPU Frequency [8] (S)	CPU	metrics											
CPU Frequency [9] (S)	~	2.5 K											
CPU Frequency [10] (S)	$\sim$	2.5 K											

### **Visualizing Trace**

Use Perfetto Zoom in to investigate regions of interest

#### ▲ ./Jacobi\_hip.inst 3624331 main Halo H2D::.. hipDe... MPI\_All pEven. hipEventRecord ipLaunchK. hipEve... Halo D2H: hipLaunc. hipLau. hipLa. hipMemcpy Jacobi\_hip.inst 3624331 hipS.. MPI\_Waita... hipSt. Zoomed in CPU Context Switches (S) 25 K $\sim$ CPU Frequency [0] (S) 5 K $\sim$ 2.5 K CPU Frequency [1] (S) $\sim$ 2.5 K CPU Frequency [2] (S) $\sim$ 2.5 K CPU Frequency [3] (S) $\sim$ 2.5 K CPU Frequency [4] (S) $\sim$ CPU Frequency [5] (S) $\sim$ 2.5 K CPU Frequency [6] (S) $\sim$ 2.5 K 2.5 K CPU Frequency [7] (S) $\sim$ 2.5 K CPU Frequency [8] (S) $\sim$ 2.5 K CPU Frequency [9] (S) $\sim$ 2.5 K CPU Frequency [10] (S) $\sim$



### **Visualizing Trace**

#### Use Perfetto Zoom in to investigate regions of interest

4676921.1 s +			0.75236 s	0.75238 s	0.75240 s	0.75242 s	0.75244 s	0.75246 s	0.75248 s	0.75250 s	0.75252 s	0.75254 s	0.75256 s	0.75258 s	0.75260 s
ž	$\square$														
./Jacobi_hip.inst 3624331	*														
HIP Activity Device 8, Queue 0	*				6	Flow Ev	ents								Marker
HIP Activity Device 8, Queue 1	*								Local	_aplacianKernel(int	, int, int, double, do	ouble, double const <sup>*</sup>	*, double*)		
									main						
Jacobi hip.inst 3624331	*	MPI_AI	l hipDe	hipEven	hipEventRecord	hipLaunc <b>oK</b>	hiptve	Halo D2H: Mi hipS	PI Exchange::Hal MPI_Waita	Halo H2D:: hipSt	hipLaunc h	hipLau hipLau.	hipLa	hipMen	hcpy
	Î	Sel	ect met	rics of	interest t	o view									
CPU Context Switches (S)	~*		se toget	her											
CPU Frequency [0] (S)	/*	5 K													
CPU Kernel Time (S)	/*	5													
CPU Memory Usage (S)	/*	0.75 K													
CPU Page Faults (S)	/*	50 K													
CPU Peak Memory (S)	/*	0.75 K													
CPU User Time (S)	/*	2.5													
CPU Virtual Memory Usage (S)	/*	50 K													
GPU Busy [0] (S)	/*	100													
GPU Memory Usage [0] (S)	/*	1 K													
GPU Temperature [0] (S) GPU	cha	racte	eristics												
GPU Power [0] (S)	/*	0.25 K													

W + D,



S -

A ✦

# **Hardware Counters**



### Hardware Counters – List All

#### \$ mpirun -np 1 omnitrace-avail --all

#### Components, Categories

	1	1	1	1	1	
COMPONENT	AVAILABLE	VALUE_TYPE	STRING_IDS	FILENAME	DESCRIPTION	CATEGORY
allinea_map   caliper_marker	false false	void void	   "allinea", "allinea_map", "forge"   "cali". "caliper". "caliper marker"		Controls the AllineaMAP sampler.	<pre>category::external, os::supports_linux, t   category::external, os::supports_unix, tp</pre>
caliper_config	false	void	"caliper_config"		Caliper configuration manager.	category::external, os::supports_unix, tp
cpu_clock	true	l long	"cpu_clock"	   cpu_clock	Total CPU time spent in both user- and ke	project::timemory, category::timing, os::
cpu_util	true	<pre>std::pair<long, long=""></long,></pre>	"cpu_util", "cpu_utilization"	cpu_util	Percentage of CPU-clock time divided by w	<pre>project::timemory, category::timing, os::  </pre>
craypat counters	false	<pre>std::vector<unsigned long,="" pre="" std::allocato<=""></unsigned></pre>	"craypat counters"	craypat counters	Names and value of any counter events tha	<pre>  category::external, os::supports linux, t  </pre>

· · · · · · · · · · · · · · · · · · ·					
ENVIRONMENT VARIABLE	VALUE	DATA TYPE	DESCRIPTION	CATEGORIES	
OMNITRACE_CAUSAL_BINARY_EXCLUDE OMNITRACE_CAUSAL_BINARY_SCOPE OMNITRACE_CAUSAL_DELAY OMNITRACE_CAUSAL_DURATION OMNITRACE_CAUSAL_FUNCTION_EXCLUDE OMNITRACE_CAUSAL_FUNCTION_SCOPE OMNITRACE_CAUSAL_RANDOM_SEED OMNITRACE_CAUSAL_SOURCE_EXCLUDE OMNITRACE_CAUSAL_SOURCE_SCOPE	%MAIN% 0 0 0	string   string   double   double   string   string   unsigned long   string   string	Excludes binaries matching the list of pr Limits causal experiments to the binaries Length of time to wait (in seconds) befor Excludes functions matching the list of p List of <function> regex entries for caus Seed for random number generator which se Excludes source files or source file + li Limits causal experiments to the source f</function>	<pre>analysis, causal, custom, libomnitrace, o analysis, causal, custom, libomnitrace, o</pre>	Environment Variables

HARDWARE COUNTER	AVAILABLE	DESCRIPTION
СРИ		
PAPI_L1_DCM PAPI_L1_ICM PAPI_L2_DCM PAPI_L2_ICM PAPI_L3_DCM PAPI_L3_ICM PAPI_L1_TCM	true false true true false false	Level 1 data cache misses Level 1 instruction cache misses Level 2 data cache misses Level 2 instruction cache misses Level 3 data cache misses Level 3 instruction cache misses Level 1 cache misses
CPU Hardware Cou	inters	
perf::CYCLES	true	PERF COUNT HW CPU CYCLES
perf::CYCLES:u=0	true	perf::CYCLES + monitor at user level
perf::CYCLES:k=0	true	perf::CYCLES + monitor at kernel level
perf::CYCLES:h=0	true	perf::CYCLES + monitor at hypervisor level
perf::CYCLES:period=0	true	perf::CYCLES + sampling period
perf::CYCLES:freq=0	true	perf::CYCLES + sampling frequency (Hz)
perf::CYCLES:precise=0	true	perf::CYCLES + precise event sampling
perf::CYCLES:excl=0	true	perf::CYCLES + exclusive access

TCC NORMAL WRITEBACK sum:device=0	true	Number of writebacks due to requests that
TCC_ALL_TC_OP_WB_WRITEBACK_sum:device=0	true	Number of writebacks due to all TC_OP wri
TCC_NORMAL_EVICT_sum:device=0	true	Number of evictions due to requests that
TCC_ALL_TC_OP_INV_EVICT_sum:device=0	true	Number of evictions due to all TC_OP inva
TCC EA RDREQ DRAM sum:device=0	true	Number of TCC/EA read requests (either 32
TCC_EA_WRREQ_DRAM_sum:device=0	true	Number of TCC/EA write requests (either 3
FETCH_SIZE:device=0	true	The total kilobytes fetched from the vide
WRITE_SIZE:device=0	true	The total kilobytes written to the video
WRITE REQ 32B:device=0	true	The total number of 32-byte effective mem
GPUBusy:device=0	i true	The percentage of time GPU was busy.
GPUBusy:device=0 Wavefronts:device=0 GPULHardware	Counters	The percentage of time GPU was busy. Total wavefronts.
GPUBusy:device=0 Wavefronts:device=0 VALUInsts:device=0 GPU Hardware	Counters	The percentage of time GPU was busy. Total wavefronts. The average number of vector ALU instruct
GPUBusy:device=0 Wavefronts:device=0 VALUInsts:device=0 SALUInsts:device=0	Counters	The percentage of time GPU was busy. Total wavefronts. The average number of vector ALU instruct The average number of scalar ALU instruct
GPUBusy:device=0 Wavefronts:device=0 VALUInsts:device=0 SALUInsts:device=0 SFetchInsts:device=0	Counters true true true	The percentage of time GPU was busy. Total wavefronts. The average number of vector ALU instruct The average number of scalar ALU instruct The average number of scalar fetch instru
GPUBusy:device=0 Wavefronts:device=0 VALUInsts:device=0 SALUInsts:device=0 SFetchInsts:device=0 GDSInsts:device=0	true Counters true true true	The percentage of time GPU was busy. Total wavefronts. The average number of vector ALU instruct The average number of scalar ALU instruct The average number of scalar fetch instru The average number of GDS read or GDS wri
GPUBusy:device=0 Wavefronts:device=0 VALUInsts:device=0 SALUInsts:device=0 GDSInsts:device=0 MemUnitBusy:device=0	Counters true true true true true	The percentage of time GPU was busy. Total wavefronts. The average number of vector ALU instruct The average number of scalar ALU instruct The average number of scalar fetch instru The average number of GDS read or GDS wri The percentage of GPUTime the memory unit

A very small subset of the counters shown here

# **Commonly Used GPU Counters**

VALUUtilization	The percentage of ALUs active in a wave. Low VALUUtilization is likely due to high divergence or a poorly sized grid
VALUBusy	The percentage of GPUTime vector ALU instructions are processed. Can be thought of as something like compute utilization
FetchSize	The total kilobytes fetched from global memory
WriteSize	The total kilobytes written to global memory
L2CacheHit	The percentage of fetch, write, atomic, and other instructions that hit the data in L2 cache
MemUnitBusy	The percentage of GPUTime the memory unit is active. The result includes the stall time
MemUnitStalled	The percentage of GPUTime the memory unit is stalled
WriteUnitStalled	The percentage of GPUTime the write unit is stalled

Full list at: https://github.com/ROCm-Developer-Tools/rocprofiler/blob/amd-master/test/tool/metrics.xml

#### **Execution with Hardware Counters**

(after modifying cfg file to set up OMNITRACE\_ROCM\_EVENTS with GPU metrics) \$ mpirun -np 1 omnitrace-run -- ./Jacobi hip.inst -g 1 1

[omnitrace][501266][0][omnitrace\_finalize] Finalizing perfetto...

[omnitrace][501266][perfetto]> Outputting '/shared/prod/home/ssitaram/HPCTrainingExamples/HIP/jacobi/omnitrace-Jacobi hip-output/2023-03-15 22.57/perfetto-trace-0.proto' (1] .. Done [omnitrace][501266][rocprof-device-0-GPUBusy]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/rocprof-device-0-GPUBusy-0.json' GPU hardware [omnitrace][501266][rocprof-device-0-GPUBusy]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/rocprof-device-0-GPUBusy-0.txt' [omnitrace][501266][rocprof-device-0-Wavefronts]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/rocprof-device-0-Wavefronts-0.json' counters [omnitrace][501266][rocprof-device-0-Wavefronts]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/rocprof-device-0-Wavefronts-0.txt' [omnitrace][501266][rocprof-device-0-MemUnitBusy]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/rocprof-device-0-MemUnitBusy-0.json' [omnitrace][501266][rocprof-device-0-MemUnitBusy]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/rocprof-device-0-MemUnitBusy-0.txt' [omnitrace][501266][trip count]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/trip count-0.ison' [omnitrace][501266][trip\_count]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/trip\_count-0.txt' [omnitrace][501266][wall\_clock]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15\_22.57/wall\_clock-0.json' [omnitrace][501266][wall\_clock]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/wall\_clock-0.txt' [omnitrace][501266][roctracer]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/roctracer-0.json' [omnitrace][501266][roctracer]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/roctracer-0.txt' [omnitrace][501266][sampling percent]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/sampling percent-0.json' [omnitrace][501266][sampling percent]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/sampling percent-0.txt' [omnitrace][501266][sampling\_cpu\_clock]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/sampling\_cpu\_clock-0.json' [omnitrace][501266][sampling\_cpu\_clock]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15\_22.57/sampling\_cpu\_clock-0.txt' [omnitrace][501266][sampling wall clock]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/sampling wall clock-0.json' [omnitrace][501266][sampling\_wall\_clock]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15\_22.57/sampling\_wall\_clock-0.txt' [omnitrace][501266][sampling gpu memory usage]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/sampling gpu memory usage-0.json' [omnitrace][501266][sampling gpu memory usage]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/sampling gpu memory usage-0.txt' [omnitrace][501266][sampling gpu power]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/sampling gpu power-0.json' [omnitrace][501266][sampling gpu power]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/sampling gpu power-0.txt' [omnitrace][501266][sampling gpu temperature]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/sampling gpu temperature-0.json' [omnitrace][501266][sampling\_gpu\_temperature]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15\_22.57/sampling\_gpu\_temperature-0.txt' [omnitrace][501266][sampling gpu busy percent]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/sampling gpu busy percent-0.json' [omnitrace][501266][sampling\_gpu\_busy\_percent]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/sampling\_gpu\_busy\_percent-0.txt' [omnitrace][501266][metadata]> Outputting 'omnitrace-Jacobi hip-output/2023-03-15 22.57/metadata-0.json' and 'omnitrace-Jacobi hip-output/2023-03-15 22.57/functions-0.json' [omnitrace][501266][0][omnitrace finalize] Finalized: 31.657272 sec wall clock, 0.000 MB peak rss, 179.700 MB page rss, 29.950000 sec cpu clock, 94.6 % cpu util [889.832] perfetto.cc:60129 Tracing session 1 ended. total sessions:0



### **Visualization with Hardware Counters**



ROCTX Regions

# **Tracing Multiple Ranks**



# Profiling Multiple MPI Ranks – Jacobi Example

Binary Rewrite Generating a new /library with instrumentation built-in:
\$ omnitrace-instrument -o Jacobi_hip.inst ./Jacobi_hip
Run the instrumented binary with 2 ranks:
\$ mpirun -np 2 omnitrace-run/Jacobi_hip.inst -g 2 1

[omnitrace][3628199][perfetto]> Outputting '/home/ssitaram/git/HPCTrainingExamples/HIP/jacobi/omnitrace-Jacobi\_hip.inst-output/2023-03-15\_18.02/perfetto-trace-1.proto' [perfetto]> Outputting '/home/ssitaram/git/HPCTrainingExamples/HIP/jacobi/omnitrace-Jacobi\_hip.inst-output/2023-03-15\_18.02/perfetto-trace-0.proto' (7856.71 KB / 7.86 M

[omnitrace][3628199][wall\_clock]> Outputting 'omnitrace-Jacobi\_hip.inst-output/2023-03-15\_18.02/wall\_clock-1.json' [omnitrace][3628196][wall\_clock]> Outputting 'omnitrace-Jacobi\_hip.inst-output/2023-03-15\_18.02/wall\_clock-0.json' [omnitrace][3628199][wall\_clock]> Outputting 'omnitrace-Jacobi\_hip.inst-output/2023-03-15\_18.02/wall\_clock-1.txt' [omnitrace][3628196][wall\_clock]> Outputting 'omnitrace-Jacobi\_hip.inst-output/2023-03-15\_18.02/wall\_clock-0.txt'

All output files are generated for each rank



### Visualizing Traces from Multiple Ranks - Separately





AMD together we advance\_

# **Statistical Sampling**



# Sampling Call-Stack (I)

#### OMNITRACE\_USE\_SAMPLING = false



#### OMNITRACE\_USE\_SAMPLING = true; OMNITRACE\_SAMPLING\_FREQ = 100 (100 samples per second)

						samnle	مما مع	nitracel	1							_					
laashi tu laashi t(grid t° mash t°)								miracej							1 1 1	1 1 1					
Jacobi_tJacobi_t(ghu_ta, mesh_ta)	J	J J J J	J J J J	5 5 5	J .	JJJJ	JJ	J J J J J J .	JJJJ	J J J J J	JJJ		JJJ		JJJ	J J J	JJ	J J J	J J J .	JJJJJ	
Jacobi_t::CreateMesh()	 J	NNN	<b>N</b> NN		N		ΝN	JHLHNF		NNNN	NLN		NNP	NNN	n n N		NN	NNN	NNN	NNN	
hipMemset	h	h h h	h h h h	h h h	h t	h h h	s h	h h h h h s	s h h	h h h h	h h h	i h h	hhł	h h h	s s h	h h h	h h ł	h h h	h h h	rhhh	
hipApiName	h	h s h	0 h h s	s h h	h	s h h	0 h	sshhhh	s h h	h h h h	s h h	ı h h	h h ł	ıhh	0 0 h	s h h	h h ł	h h h	h h h	) h h h	
hipDeviceGetByPCIBusId	h	h O h	h h 🕻	0 h h	h C	D h h	0 h	00 <b>h</b> 0h	h h	hhhh	h h	h h	hhł	n <mark>h</mark> h h	s O h	0 h h	h h ł	h h h	h h h l	0 0 h h	
hipExtStreamGetCUMask	h	h <mark>s</mark> h	h h 🕻	0 h h	h C	D h h	h	0 h 0 h	h h	h h h h	h h	h h	hhł	n <mark>h</mark> h h	s_h	s h h	h h ł	h h h	h h h	Ohh	
hipExtStreamGetCUMask	h	h s h	h h 🕻	0 h h	h C	) h h	h	_ <b>O</b> h	h h	hhhh	h t	h h	hhł	n h h h	s _ h	h h	h h ł	h h h	h h h	Oh s	
hiprtcLinkAddData	h	h s h	h h	h h	h s	s h h	h	h h	h h	hhhh	ŀ	h h	hhł	n hh	h	h h	h h ł	h h h	O h h	s h	
hiprtcLinkAddData	h	h h	h h	h h	hs	s h h	h	h	h h	hhhh	ŀ	h h	hhł	n hr	h	h h	h h ł	h h h	r h h	k h	
hiprtcLinkAddData	h	h h	h h	h h	h	h h	h	h	h h	hhhh	ŀ	h h	hhł	h h	h	h h	h h ł	h h h	h h	h	
hiprtcLinkAddData	h	h h	h h	h h	h	h h	h	h	h h	hhhh	ŀ	h h	hhł	n h	h	h h	h h ł	h h h	h h		
hiprtcLinkAddData	h	h h	h h	h h	h	h h	h	h	h h	hhhh	ŀ	h h	hhh	n r	h	h h	h h ł	h h h	h h		
hiprtcLinkAddData	h	r r	r r	r r	r	r r	r	r	r r		ſ		r r i		r	r r		r r r	r r		
hiprtcLinkAddData	h	h h	h h	h h	h	h h	h	h	h h	h h h h	ŀ	h h	h h ł	1 I	h	h h	h h ł	h h h	h h		
hiprtcLinkAddData	r							_						-							
hiprtcLinkAddData	h																				
hiprtcLinkAddData	h																				
amd_comgr_do_action	h															a a va L			ما ا		
amd_comgr_data_set_remove	h													Eacu	sai	npie	e si	IOW	s ine	2	
amd_comgr_data_set_remove														call s	tacl	< at	tha	it tin	ne		
amd_comgr_data_set_remove																· at					
amd comgr data set remove																					
amd_comgr_data_set_remove																					
and_congr_ddtd_sct_tenove	_										_		_	_	_			_	_		
																				7AN M I	

Scroll down all the way in Perfetto to see the sampling output!

together we advance\_

# Sampling Call-Stack (II)

#### Zoom in call-stack sampling

					samples [omnitrace	]				
Jacobi	Jacobi_t::Run()	Jacobi_t::Run()	Jacobi_t::Run()	Jacobi_t::Run()	Jacobi_t::Run()	Jacobi_t::Run()	Jacobi_t::Run()	Jacobi_t::Run()	Jacobi_t::Run()	Jacobi_t::Ru
Norm(gr	LocalLaplacian(gri	Norm(grid_t&, me	Norm(grid_t&, me	hipEventRecord	Norm(grid_t&, me	Jacobilteration(	HaloExchange(gri	LocalLaplacian(g	HaloExchange(grid	Norm(grid_t&
hipMemc	hipLaunchKernel	hipMemcpy	hipMemcpy	std::basic_string<	hipMemcpy	hipLaunchKernel	hipStreamSynchro	hipLaunchKernel	hipStreamSynchroni.	. hipMemcpy
hipApiN	std::basic_string<	hipApiName	hipApiName	OnUnload	hipApiName	std::basic_strin	std::basic_strin	hipMemPoolGetAtt	hipLaunchHostFunc	hipApiName
hiprtcL	OnUnload	hiprtcLinkAddData	hiprtcLinkAddData	OnUnload	hiprtcLinkAddData	OnUnload	OnUnload	hip_impl::hipLau	OnUnload	hiprtcLinkAd
hiprtcL	OnUnload	hiprtcLinkAddData	hiprtcLinkAddData		hiprtcLinkAddData		OnUnload	hipGetCmdName	OnUnload	hiprtcLinkAd
hiprtcL	OnUnload	hiprtcLinkAddData	hiprtcLinkAddData		hiprtcLinkAddData			hipGetPCH	OnUnload	hiprtcLinkAd
hiprtcL	std::ostream& std:	hiprtcLinkAddData	hiprtcLinkAddData		hiprtcLinkAddData			hiplpcGetEventHa		hiprtcLinkAd
hiprtcL	std::ostreambuf_it	hiprtcLinkAddData	hiprtcLinkAddData		hiprtcLinkAddData					hiprtcLinkAd
hiprtcL		hiprtcLinkAddData	hiprtcLinkAddData		hiprtcLinkAddData					hiprtcLinkAd
hiprtcL		hiprtcLinkAddData	hiprtcLinkAddData		hiprtcLinkAddData					hiprtcLinkAd
hiprtcL		hiprtcLinkAddData	hiprtcLinkAddData		hiprtcLinkAddData					hiprtcLinkAd
roctrac		roctracer_disabl	roctracer_disabl		roctracer_disabl					roctracer_di
hsa_amd		hsa_amd_image_ge	hsa_amd_image_ge		hsa_amd_image_ge					hsa_amd_imag

Thread 0 (S) 3625610

Sampling data is annotated with (S)



# **Other Features**



## **Kernel Durations**

#### \$ cat omnitrace-Jacobi\_hip.inst-output/2023-03-15\_13.57/wall\_clock-0.txt

If you do not see a wall\_clock.txt dumped by omnitrace, try modify the config file \$HOME/.omnitrace.cfg and enable OMNITRACE\_USE\_TIMEMORY:

···	
OMNITRACE USE PERFETTO	= tru
OMNITRACE_USE_TIMEMORY	= tru
OMNITRACE_USE_SAMPLING	= fal

0>>>	MPI_Allreduce	1	5	wall_clock	sec	0.000012	0.000012	0.000012	0.000012	0.000000	0.000000	100.0
0>>>	_hipDeviceSynchronize	1	5	wall_clock	sec	0.000019	0.000019	0.000019	0.000019	0.000000	0.000000	94.4
0>>>	NormKernel1(int, double, double const*, double*)	1	6	wall clock	sec	0.000001	0.000001	0.000001	0.000001	0.000000	0.000000	100.0
0>>>	NormKernel2(int, double const*, double*)	1	6	wall_clock	sec	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	100.0
0>>>	_MPI_Barrier	1	5	wall_clock	sec	0.000001	0.000001	0.000001	0.000001	0.000000	0.000000	100.0
0>>>	_hipEventRecord	2	5	wall_clock	sec	0.000027	0.000014	0.000011	0.000016	0.000000	0.000003	100.0
0>>>	Halo D2H::Halo Exchange	1	5	wall_clock	sec	1.628420	1.628420	1.628420	1.628420	0.000000	0.000000	0.0
0>>>	hipStreamSynchronize Call Stack	1	6	wall_clock	sec	0.000003	0.000003	0.000003	0.000003	0.000000	0.000000	100.0
0>>>	MPI Exchange::Halo Exchange	1	6	wall_clock	sec	1.628395	1.628395	1.628395	1.628395	0.000000	0.000000	0.0
0>>>	MPI Waitall	1	7	wall clock	sec	0.000002	0.000002	0.000002	0.000002	0.000000	0.000000	100.0
0>>>	_Halo H2D::Halo Exchange	1	7	wall_clock	sec	1.628104	1.628104	1.628104	1.628104	0.000000	0.000000	0.0
0>>>	_hipStreamSynchronize	1	8	wall_clock	sec	0.000003	0.000003	0.000003	0.000003	0.000000	0.000000	100.0
0>>>	_hipLaunchKernel	5	8	wall_clock	sec	0.000615	0.000123	0.000005	0.000578	0.000000	0.000254	99.6
0>>>	_mbind	1	9	wall_clock	sec	0.000003	0.000003	0.000003	0.000003	0.000000	0.000000	100.0
0>>>	hipMemcpy	1	8	wall_clock	sec	0.001122	0.001122	0.001122	0.001122	0.000000	0.000000	99.9
0>>>	<pre>[_LocalLaplacianKernel(int, int, int, double, double, double const*, double*)</pre>	1	9	wall_clock	sec	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	100.0
0>>>	_HaloLaplacianKernel(int, int, int, double, double, double const*, double const*, double*)	1	9	wall_clock	sec	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	100.0
0>>>	JacobiIterationKernel(int, double, double, double const*, double const*, double*, double*)	1	9	wall clock	sec	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	100.0

Text file is for quick reference. JSON output is easy to script for and can be read by Hatchet, a Python package (<u>https://hatchet.readthedocs.io/en/latest/</u>)



Durations

### **Kernel Durations (flat profile)**

#### Edit in your omnitrace.cfg:

OMNITRACE\_USE\_TIMEMORY OMNITRACE FLAT PROFILE

= true

= true

# Use flat profile to see aggregate duration of kernels and functions

REAL-CLOCI	K TIMER (]	[.E. WALL	-CLOCK TIMER)								
LABEL	COUNT	DEPTH	METRIC	UNITS	SUM	MEAN	MIN	MAX	VAR	STDDEV	% SELF
0>>> main	1	0	wall_clock	sec	82.739099	82.739099	82.739099	82.739099	0.000000	0.000000	100.0
0>>> MPI_Init	1	0	wall_clock	sec	34.056610	34.056610	34.056610	34.056610	0.000000	0.000000	100.0
0>>> pthread_create	3	0	Wall_clock	sec			0.001169	0.011974	0.000038		
Initia Initia	205	0	wall_clock						0.000000		
10>>> NPT Comm rank	1	0	wall_clock	l sec	0.000212	0.000212	0.000212	0.000212	0.000000	0.000000	100.0
0>>> MPI Comm size	ī	Õ	wall clock	sec	0.000004	0.000004	0.000004	0.000004	0.000000	0.000000	100.0
0>>> hipInit	1	0	wall clock	sec	0.000372	0.000372	0.000372	0.000372	0.000000	0.000000	100.0
0>>> hipGetDeviceCount	1	0	wall_clock	sec	0.000017	0.000017	0.000017	0.000017	0.000000	0.000000	100.0
0>>> MPI_Allgather	1	0	wall_clock	sec	0.000009	0.000009	0.00009	0.000009	0.000000	0.000000	100.0
0>>> hipSetDevice	1	0	wall_clock	sec	0.000024	0.000024	0.000024	0.000024	0.000000	0.000000	100.0
0>>> hipHostMalloc	3	0	wall_clock	sec	0.126827	0.042276	0.000176	0.126453	0.005314	0.072900	100.0
0>>> hipMalloc	7	0	wall_clock	sec	0.000458	0.000065	0.000024	0.000178	0.000000	0.000052	100.0
0>>> hipMemset	1	0	Wall_clock	sec	35.770403	35.770403	35.770403	35.//0403	0.000000	0.000000	
USSS hiptoreur	1005	0	wall_clock	sec					0.000018		
UA>>> hipEventCreate	1005	0	wall_clock								
10->> hiptontcreate	∠   5002	0	wall_clock								100.0
10>>> MPT All reduce	1003	0	wall clock	sec	0.101301				0.000000		100.0
10>>> hipDeviceSynchronize	1001	0	wall clock	sec	0.016813	0.000017	0.000015	0.000043	0.000000	0.000004	100.0
0>>> MPI Barrier	3	0	wall clock	sec	0.000007	0.000002	0.000001	0.000004	0.000000	0.000001	100.0
0>>> hipEventRecord	2000	0	wall_clock	sec	0.046701	0.000023	0.000020	0.000225	0.000000	0.000006	100.0
0>>> hipStreamSynchronize	2000	0	wall_clock	sec	0.030366	0.000015	0.000013	0.000382	0.000000	0.000009	100.0
0>>> MPI_Waitall	1000	0	wall_clock	sec	0.001665	0.000002	0.000002	0.000007	0.000000	0.000000	100.0
<pre> 0&gt;&gt;&gt; NormKernel1(int, double, double, double const*, double*)</pre>	1001	0	wall_clock	sec	0.001502	0.000002	0.000001	0.000006	0.000000	0.000000	100.0
0>>> NormKernel2(int, double const*, double*)	1000	0	wall_clock	sec	0.001972	0.000002	0.000001	0.000003	0.000000	0.000001	100.0
0>>> LocaLaplacianKernel(int, int, int, double, double, double const*, double*)	1000	0	wall_clock	sec	0.001488	0.000001	0.000001	0.000007	0.000000	0.000000	100.0
U>>> HaloLaplaClankernel(int, int, int, double, double, double const*, double const*, double*)		0	Wall_clock	sec					0.000000		
US>> hipeventetabseutime		0	wall_clock	sec							
Joss sthread ioin	1000	0	wall_clock	Sec	0.002598	0.000003	0.000001	0.000000	0.000000		100.0
loss hinFree	4	0	wall clock	sec	0.000526	0.000131	0.000021	0.000243	0.000000	0.000091	100.0
0>>> hipHostFree	2	0	wall clock	sec	0.000637	0.000318	0.000287	0.000350	0.000000	0.000044	100.0
3>>> start thread	1	0	wall clock	sec	0.004802	0.004802	0.004802	0.004802	0.000000	0.000000	100.0
1>>> start_thread	1	0	wall clock	sec	81.987779	81.987779	81.987779	81.987779	0.000000	0.000000	100.0
2>>> start_thread	-	0		i -	i -	-	-	-	-	i -	-

AMD together we advance\_

### **User API**

#### Omnitrace provides an API to control the instrumentation

API Call	Description
int <b>omnitrace_user_start_trace(void)</b>	Enable tracing on this thread and all subsequently created threads
int <b>omnitrace_user_stop_trace(void)</b>	Disable tracing on this thread and all subsequently created threads
int <b>omnitrace_user_start_thread_trace(void)</b>	Enable tracing on this specific thread. Does not apply to subsequently created threads
int <b>omnitrace_user_stop_thread_trace(void)</b>	Disable tracing on this specific thread. Does not apply to subsequently created threads
int <b>omnitrace_user_push_region(void)</b>	Start user defined region
int <b>omnitrace_user_pop_region(void)</b>	End user defined region, FILO (first in last out) is expected

All the API calls: https://amdresearch.github.io/omnitrace/user\_api.html

# **OpenMP**<sup>®</sup>

We use the example omnitrace/examples/openmp/
Build the code with CMake:
\$ cmake -B build
Use the openmp-lu binary, which can be executed with:
\$ export OMP_NUM_THREADS=4 \$ srun -n 1 -c 4 ./openmp-lu
Create a new instrumented binary:
\$ srun -n 1 omnitrace-instrument -o openmp-lu.inst ./openmp-lu
Execute the new binary:
\$ srun -n 1 -c 4 omnitrace-run/openmp-lu.inst

	REAL-CLOCK TIMER (I.E. WALL-CLOCK TIMER)													
	LABEL	COUNT	DEPTH	METRIC	UNITS	   SUM 	MEAN	MIN	 MAX	VAR	STDDEV 	% SELF		
0>>>	main	1		wall_clock	sec	1.096702	1.096702	1.096702	1.096702	0.000000	0.000000	9.2		
0>>>	_pthread_create	3	1	wall_clock	sec	0.002931	0.000977	0.000733	0.001420	0.000000	0.000385	0.0		
3>>>		1	2	wall_clock	sec	2.451520	2.451520	2.451520	2.451520	0.000000	0.000000	57.7		
3>>>	_erhs	1	3	wall_clock	sec	0.001906	0.001906	0.001906	0.001906	0.000000	0.00000	100.0		
3>>>	_rhs	153	3	wall_clock	sec	0.229893	0.001503	0.001410	0.001893	0.000000	0.000116	100.0		
3>>>	_jacld	3473	3	wall_clock	sec	0.170568	0.000049	0.000047	0.000135	0.000000	0.000005	100.0		
3>>>	_blts	3473	3	wall_clock	sec	0.232512	0.000067	0.000040	0.000959	0.000000	0.000034	100.0		
3>>>	_jacu	3473	3	wall_clock	sec	0.166229	0.000048	0.000046	0.000148	0.000000	0.000005	100.0		
3>>>	_buts	3473	3	wall_clock	sec	0.236484	0.00068	0.000041	0.000391	0.000000	0.000031	100.0		
2>>>	_start_thread	1	2	wall_clock	sec	2.452309	2.452309	2.452309	2.452309	0.000000	0.00000	58.1		
2>>>	_erhs	1	3	wall_clock	sec	0.001895	0.001895	0.001895	0.001895	0.000000	0.000000	100.0		
2>>>	_rhs	153	3	wall_clock	sec	0.229776	0.001502	0.001410	0.001893	0.000000	0.000115	100.0		
2>>>	_jacld	3473	3	wall_clock	sec	0.204609	0.000059	0.000057	0.000152	0.000000	0.000006	100.0		
2>>>	_blts	3473	3	wall_clock	sec	0.192986	0.000056	0.000047	0.000358	0.000000	0.000026	100.0		
2>>>	_jacu	3473	3	wall_clock	sec	0.199029	0.000057	0.000055	0.000188	0.000000	0.000007	100.0		
2>>>	_buts	3473	3	wall_clock	sec	0.198972	0.000057	0.000048	0.000372	0.000000	0.000026	100.0		
1>>>	_start_thread	1	2	wall_clock	sec	2.453072	2.453072	2.453072	2.453072	0.000000	0.000000	58.6		
1>>>	_erhs	1	3	wall_clock	sec	0.001905	0.001905	0.001905	0.001905	0.000000	0.000000	100.0		
1>>>	_rhs	153	3	wall_clock	sec	0.229742	0.001502	0.001410	0.001894	0.000000	0.000115	100.0		
1>>>	_jacld	3473	3	wall_clock	sec	0.206418	0.000059	0.000057	0.000934	0.000000	0.000016	100.0		
1>>>	_blts	3473	3	wall_clock	sec	0.186097	0.000054	0.000047	0.000344	0.000000	0.000023	100.0		
1>>>	_jacu	3473	3	wall_clock	sec	0.198689	0.000057	0.000055	0.000186	0.000000	0.000006	100.0		
1>>>	_buts	3473	3	wall_clock	sec	0.192470	0.000055	0.000048	0.000356	0.000000	0.000022	100.0		
0>>>	_erhs	1	1	wall_clock	sec	0.001961	0.001961	0.001961	0.001961	0.000000	0.000000	100.0		
0>>>	_rhs	153	1	wall_clock	sec	0.229889	0.001503	0.001410	0.001891	0.000000	0.000116	100.0		
0>>>	_jacld	3473	1	wall_clock	sec	0.208903	0.000060	0.000057	0.000359	0.000000	0.000017	100.0		
0>>>	_blts	3473	1	wall_clock	sec	0.172646	0.000050	0.000047	0.000822	0.000000	0.000020	100.0		
0>>>	_jacu	3473	1	wall_clock	sec	0.202130	0.000058	0.000055	0.000350	0.000000	0.000016	100.0		
0>>>	_buts	3473	1	wall_clock	sec	0.176975	0.000051	0.000048	0.000377	0.000000	0.000016	100.0		
0>>>	_pintgr	1	1	wall_clock	sec	0.000054	0.000054	0.000054	0.000054	0.00000	0.00000	100.0		

# **OpenMP**<sup>®</sup> Visualization

Clock Snapshots metric														
▲ openmp-lu.inst 117836														
openmp-lu.inst 117836	main juckd bits juckd bits	acid bits												
Thread 1 117844	the blag and blag blag blag blag blag blag blag blag	jacid												
Thread 2 117846	rits juck bits j	bits jac												
Thread 3 117848		cid bits												
	sundes (environ)													
	societi CAMA navilai													
	ssor(nt) clone_emp_fn.4													
	no unwind info found													
Thread 0 (S) 117857														
	sandis (amitrad													
	no united info fourd													
	omititace:.component.pthread_create_ptotawapper.coperator()) const	/												
	omp_fulfikerent ssonteni [closeemp_fulfi													
Thread 1 (S) 117858	saturing under company no unividad for found													
Thread 1 (S) 117858														
	anges extenses no unifiel info fond													
	omititace:.component:ptitead_create_ptota:.wrapper:coperator()) const	/												
	ong fulfil (vent research faile													
	anticity generation and anticity generation and anticity of the second anticity of the seco													
Thread 2 (S) 11/859														
	samples (anahase)													
	no uminidi ilini found													
Thread 3 (S) 117860	south) folceemp.fh.4													
	no uneited info found													



# Python<sup>™</sup>

The omnitrace Python package is installed in /path/omnitrace\_install/lib/pythonX.Y/site-packages/omnitrace

Setup the environment:

\$ export PYTHONPATH=/path/omnitrace/lib/python/sitepackages/:\${PYTHONPATH}

We use the Fibonacci example in omnitrace/examples/python/source.py

Execute the python program with:

\$ omnitrace-python ./external.py

Profiled data is dumped in output directory:

\$ cat omnitrace-source-output/timestamp/wall\_clock.txt

			REAL-CLOO	K TIMER (I.E	. WALL-CL	DCK TIMER)						
	LABEL	COUNT	DEPTH	METRIC	UNITS 	SUM 	MEAN	MIN 	MAX	VAR 	STDDEV	% SELF   
0>>> m	ain_loop	3	0	wall_clock	sec	2.786075	0.928692	0.926350	0.932130	0.000009	0.003042	0.0
0>>>	_run	3	1	wall_clock	sec	2.785799	0.928600	0.926250	0.932037	0.00009	0.003043	0.0
0>>>		3	2	wall_clock	sec	2.750104	0.916701	0.914454	0.919577	0.000007	0.002619	0.0
0>>>	_fib	6	3	wall_clock	sec	2.749901	0.458317	0.348962	0.567074	0.013958	0.118145	0.0
<del>0</del> >>>	_fib	12	4	wall_clock	sec	2.749511	0.229126	0.133382	0.350765	0.006504	0.080650	0.0
<del>0</del> >>>	_fib	24	5	wall_clock	sec	2.748734	0.114531	0.050867	0.217030	0.002399	0.048977	0.1
<del>0</del> >>>	_fib	48	6	wall_clock	sec	2.747118	0.057232	0.019302	0.134596	0.000806	0.028396	0.1
<del>0</del> >>>	_fib	96	7	wall_clock	sec	2.743922	0.028583	0.007181	0.083350	0.000257	0.016026	0.2
<del>0</del> >>>	_fib	192	8	wall_clock	sec	2.737564	0.014258	0.002690	0.051524	0.000079	0.008887	0.5
<del>0</del> >>>	_fib	384	9	wall_clock	sec	2.724966	0.007096	0.000973	0.031798	0.000024	0.004865	0.9
<del>0</del> >>>	_fib	768	10	wall_clock	sec	2.699251	0.003515	0.000336	0.019670	0.00007	0.002637	1.9
<del>0</del> >>>	_fib	1536	11	wall_clock	sec	2.648006	0.001724	0.000096	0.012081	0.000002	0.001417	3.9
θ>>>	_fib	3072	12	wall_clock	sec	2.545260	0.000829	0.000016	0.007461	0.00001	0.000758	8.0
0>>>	_fib	6078	13	wall_clock	sec	2.342276	0.000385	0.000016	0.004669	0.00000	0.000404	16.0
0>>>	_fib	10896	14	wall_clock	sec	1.967475	0.000181	0.000015	0.002752	0.00000	0.000218	28.6
<del>0</del> >>>	_fib	15060	15	wall_clock	sec	1.404069	0.00093	0.000015	0.001704	0.00000	0.000123	43.6
<del>0</del> >>>	_fib	14280	16	wall_clock	sec	0.791873	0.000055	0.000015	0.001044	0.00000	0.000076	58.3
<del>0</del> >>>	_fib	8826	17	wall_clock	sec	0.330189	0.000037	0.000015	0.000620	0.00000	0.000050	70.9
θ>>>	_fib	3456	18	wall_clock	sec	0.096120	0.000028	0.000015	0.000380	0.00000	0.000034	81.0
0>>>	_fib	822	19	wall_clock	sec	0.018294	0.000022	0.000015	0.000209	0.00000	0.000024	88.9
0>>>	fib	108	20	wall_clock	sec	0.002037	0.000019	0.000016	0.000107	0.00000	0.000015	94.9
0>>>	_fib	6	21	wall_clock	sec	0.000104	0.000017	0.000016	0.000019	0.00000	0.00001	100.0
0>>>	_inefficient	3	2	wall_clock	sec	0.035450	0.011817	0.010096	0.012972	0.000002	0.001519	95.8
<del>0</del> >>>	sum	3	3	wall_clock	sec	0.001494	0.000498	0.000440	0.000537	0.000000	0.000051	100.0

Python documentation: https://amdresearch.github.io/omnitrace/python.html

# Visualizing Python<sup>™</sup> Perfetto Tracing



#### Kokkos

Omnitrace can instrument Kokkos applications too.

Edit the \$HOME/.omnitrace.cfg file and enable omnitrace:

• • •

OMNITRACE\_USE\_KOKKOSP = true

• • •

Profiling with omnitrace produces \*kokkos\*.txt files:

#### \$ cat kokkos\_memory0.txt

0>>>	[_[kokkos][dev0] Kokkos::deep_copy: copy between contiguous views, post deep copy fence	1	3	kokkos_memory	MB	0	0	Θ
0>>>	_[kokkos] Kokkos::deep_copy: copy between contiguous views, post deep copy fence	1	3	kokkos_memory	MB	0	0	Θ
0>>>	_[kokkos][deep_copy] Host=DataBlock_A2_mirror HIP=DataBlock_A2	1	2	kokkos_memory	MB	142	142	100
0>>>	_[kokkos][dev0] Kokkos::deep_copy: copy between contiguous views, pre view equality check	1	3	kokkos_memory	MB	0	0	Θ
0>>>	_[kokkos] Kokkos::deep_copy: copy between contiguous views, pre view equality check	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	_[kokkos][dev0] Kokkos::deep_copy: copy between contiguous views, post deep copy fence	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	_[kokkos] Kokkos::deep_copy: copy between contiguous views, post deep copy fence	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	_[kokkos][deep_copy] Host=DataBlock_dV_mirror HIP=DataBlock_dV	1	2	kokkos_memory	MB	140	140	100
0>>>	_[kokkos][dev0] Kokkos::deep_copy: copy between contiguous views, pre view equality check	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	_[kokkos] Kokkos::deep_copy: copy between contiguous views, pre view equality check	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	_[kokkos][dev0] Kokkos::deep_copy: copy between contiguous views, post deep copy fence	1	3	kokkos_memory	MB	0	0	Θ
0>>>	_[kokkos] Kokkos::deep_copy: copy between contiguous views, post deep copy fence	1	3	kokkos_memory	MB	0	0	Θ
0>>>	_DataBlockHost::SyncToDevice()	1	1	kokkos_memory	MB	0	Θ	Θ
0>>>	_[kokkos][deep_copy] HIP=Hydro_Vc Host=Hydro_Vc_mirror	1	2	kokkos_memory	MB	1124	1124	100
0>>>	_[kokkos][dev0] Kokkos::deep_copy: copy between contiguous views, pre view equality check	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	[_[kokkos] Kokkos::deep_copy: copy between contiguous views, pre view equality check	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	_[kokkos][dev0] Kokkos::deep_copy: copy between contiguous views, post deep copy fence	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	_[kokkos] Kokkos::deep_copy: copy between contiguous views, post deep copy fence	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	_[kokkos][deep_copy] HIP=Hydro_InvDt Host=Hydro_InvDt_mirror	1	2	kokkos_memory	MB	140	140	100
0>>>	_[kokkos][dev0] Kokkos::deep_copy: copy between contiguous views, pre view equality check	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	[_[kokkos] Kokkos::deep_copy: copy between contiguous views, pre view equality check	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	_[kokkos][dev0] Kokkos::deep_copy: copy between contiguous views, post deep copy fence	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	_[kokkos] Kokkos::deep_copy: copy between contiguous views, post deep copy fence	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	_[kokkos][deep_copy] HIP=Hydro_Vs Host=Hydro_Vs_mirror	1	2	kokkos_memory	MB	426	426	100
0>>>	_[kokkos][dev0] Kokkos::deep_copy: copy between contiguous views, pre view equality check	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	[_[kokkos] Kokkos::deep_copy: copy between contiguous views, pre view equality check	1	3	kokkos_memory	MB	0	Θ	Θ
0>>>	[_[kokkos][dev0] Kokkos::deep_copy: copy between contiguous views, post deep copy fence	1	3	kokkos_memory	MB	0	0	Θ
0>>>	[_[kokkos] Kokkos::deep_copy: copy between contiguous views, post deep copy fence	1	3	kokkos_memory	MB	0	0	Θ

together we advance\_

# Visualizing Kokkos with Perfetto Trace

Visualize perfetto-trace-0.proto (with sampling enabled)

5.9	S	+1.6 ms	+3.6 ms	+5.6 ms	+7.6 ms	+9.6 ms	+11.6 ms	+13.6 ms	+15.6 ms	+17.6 ms	+19.6 ms	+21.6 ms	+23.6 ms	+25.6 ms	+27.6 ms	+29.6 ms	+31.6 ms	+33.6 ms	+35.6 ms	+37.6 ms	+39.6 ms	+41.6 ms	+43.6 ms	+45.6 ms	+47.6 ms	+49.6 ms	+51.6 ms	+53.6 ms	+55.6 ms	+57.
																									9					
	8			223															2							-			9	
																TimeIntegra	ator::Cycle													
																TimeIntegra	ator::Cycle													
	StateContainer::CopyFrom						DataBlock::EvolveStage								[kokkos] Timestep_reduction				ly DataBlock::EvolveStage											
	Kokkos::d	deep_copy<	Kok Hydro	::CalcRightHan	dSide	Hydro::CalcR	RiemannFlux	Hydro::	CalcRightHandS	Side<1>	Hydro::CalcF	RiemannFlux		Hydro::CalcRig	ghtHandSide<2	>	ElectroMotive	Force::CalcCorn	erEMF	[kokkos] Kok	kos::Impl::Paral	lelReduce <mdr< th=""><th>an</th><th></th><th>Hy</th><th>/dro::CalcRight</th><th>HandSide&lt;0&gt;</th><th></th><th></th><th>Hydr</th></mdr<>	an		Hy	/dro::CalcRight	HandSide<0>			Hydr
	hipMe	[kokkos][d	[k Hydro	o::CalcRightHar	ndSide	Hydro::HL	LLD_MHD	Hydro	o::CalcRightHand	dSide	Hydro::HL	LLD_MHD		Hydro::CalcF	RightHandSide		ElectroMotiveF	Force::CalcCont	actAv		hipStreamSyncl	hronize			)	Hydro::CalcRigh	htHandSide			H
		hipMemcpy	hi [kokk	os] CalcRightHa	andSi	[kokkos] Calc	RiemannFlux	[kokko	os] CalcRightHan	ndSide	[kokkos] Calc	RiemannFlux		[kokkos] Calc	:RightHandSide		ElectroMotive	Force::CalcCont	actAv						þ	kokkos] CalcRig	phtHandSide			[kokk
	hipEventSynchronize		ize	hipEventSy	nchronize	nize hipEventSynchronize			hipEventSynchronize hipEventSynchroniz		Synchronize	[kokkos] EMF_Integrate_to_Corner			Corner							hipEventSync	chronize			hipl				
																	hipEve	ntSynchronize												
																								1						
1																														

### **Other Executables**

#### • omnitrace-sample

- For sampling with low overhead, use omnitrace-sample
- Use omnitrace-sample --help to get relevant options
- Settings in the OmniTrace config file will be used by omnitrace-sample
- Example invocation to get a flat tracing profile on Host and Device (-PTHD), excluding all components (-E all) and including only rocm-smi, roctracer, rocprofiler and roctx components (-I ...)
   mpirun -np 1 omnitrace-sample -PTHD -E all -I rocm-smi -I roctracer -I rocprofiler -I roctx -- ./Jacobi hip -g 1 1

#### • omnitrace-causal

- Invokes causal profiling
- omnitrace-critical-trace
  - Post-processing tool for critical-trace data output by omnitrace

Current documentation: https://amdresearch.github.io/omnitrace/development.html#executables

# **Tips & Tricks**

- My Perfetto timeline seems weird how can I check the clock skew?
  - Set OMNITRACE\_VERBOSE=1 or higher for verbose mode and it will print the timestamp skew
- It takes too long to map rocm-smi samples to kernels.
  - Temporarily set OMNITRACE\_USE\_ROCM\_SMI=OFF
- What is the best way to profile multi-process runs?
  - Use OmniTrace's binary rewrite (-o) option to instrument the binary first, run the instrumented binary with mpirun/srun
- If you are doing binary rewrite and you do not get information about kernels, set:
  - HSA\_TOOLS\_LIB=libomnitrace.so in the env. and set OMNITRACE\_USE\_ROCTRACER=ON in the cfg file
- My HIP application hangs in different points, what do I do?
  - Try to set HSA\_ENABLE\_INTERRUPT=0 in the environment, this changes how HIP runtime is notified when GPU kernels complete
- My Perfetto trace is too big, can I decrease it?
  - Yes, with v1.7.3 and later declare OMNITRACE\_PERFETTO\_ANNOTATIONS to false
- I want to remove the many rows of CPU frequency lines from the Perfetto trace
  - Declare the OMNITRACE\_USE\_PROCESS\_SAMPLING = false

# Summary

- OmniTrace is a powerful tool to understand CPU + GPU activity
  - Ideal for an initial look at how an application runs
- Leverages several other tools and combines their data into a comprehensive output file
  - Some tools used are AMD uProf, rocprof, rocm-smi, roctracer, perf, etc.
- Easy to visualize traces in Perfetto
- Includes several features:
  - Dynamic Instrumentation either at Runtime or using Binary Rewrite
  - Statistical Sampling for call-stack info
  - Process sampling, monitoring of system metrics during application run
  - Causal Profiling
  - Critical Path Tracing

# **Questions?**

# **DISCLAIMERS AND ATTRIBUTIONS**

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18

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.

© 2023 Advanced Micro Devices, Inc. All rights reserved.

AMD, the AMD Arrow logo, Radeon<sup>™</sup>, Instinct<sup>™</sup>, EPYC, Infinity Fabric, ROCm<sup>™</sup>, 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.

#