CPU-GPU tracing using ROCm

Arnaud Fiorini with Pr. Michel Dagenais

January 8th, 2021

Polytechnique Montreal

DORSAL Laboratory

Agenda

- Background
- Tracing pipeline
- Results
- Future Work

Background

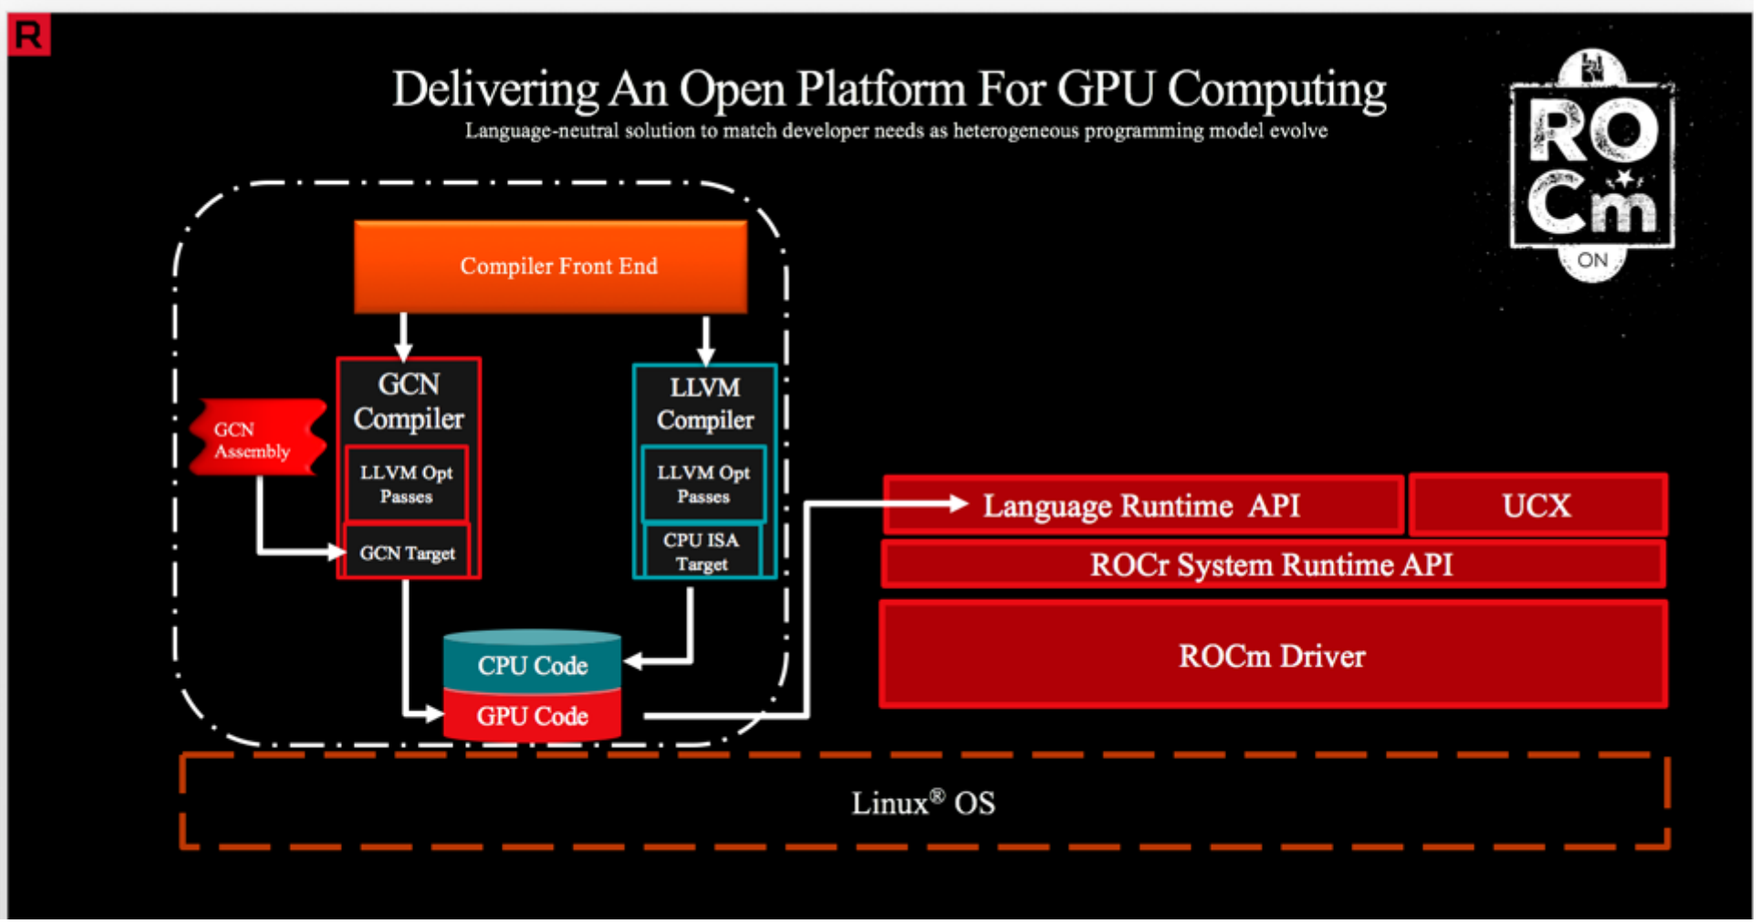

Background

- ROCm works with OpenMP, OpenCL and HIP.
- CUDA code can be easily converted to HIP.
- It is compatible with Deep Learning libraries like Pytorch, Tensorflow and others.
- ROC-profiler and ROC-tracer can trace multiple APIs : HIP, HSA, KFD and code annotation (ROCtx).
- ROC-profiler also provides performance counters.

Tracing pipeline

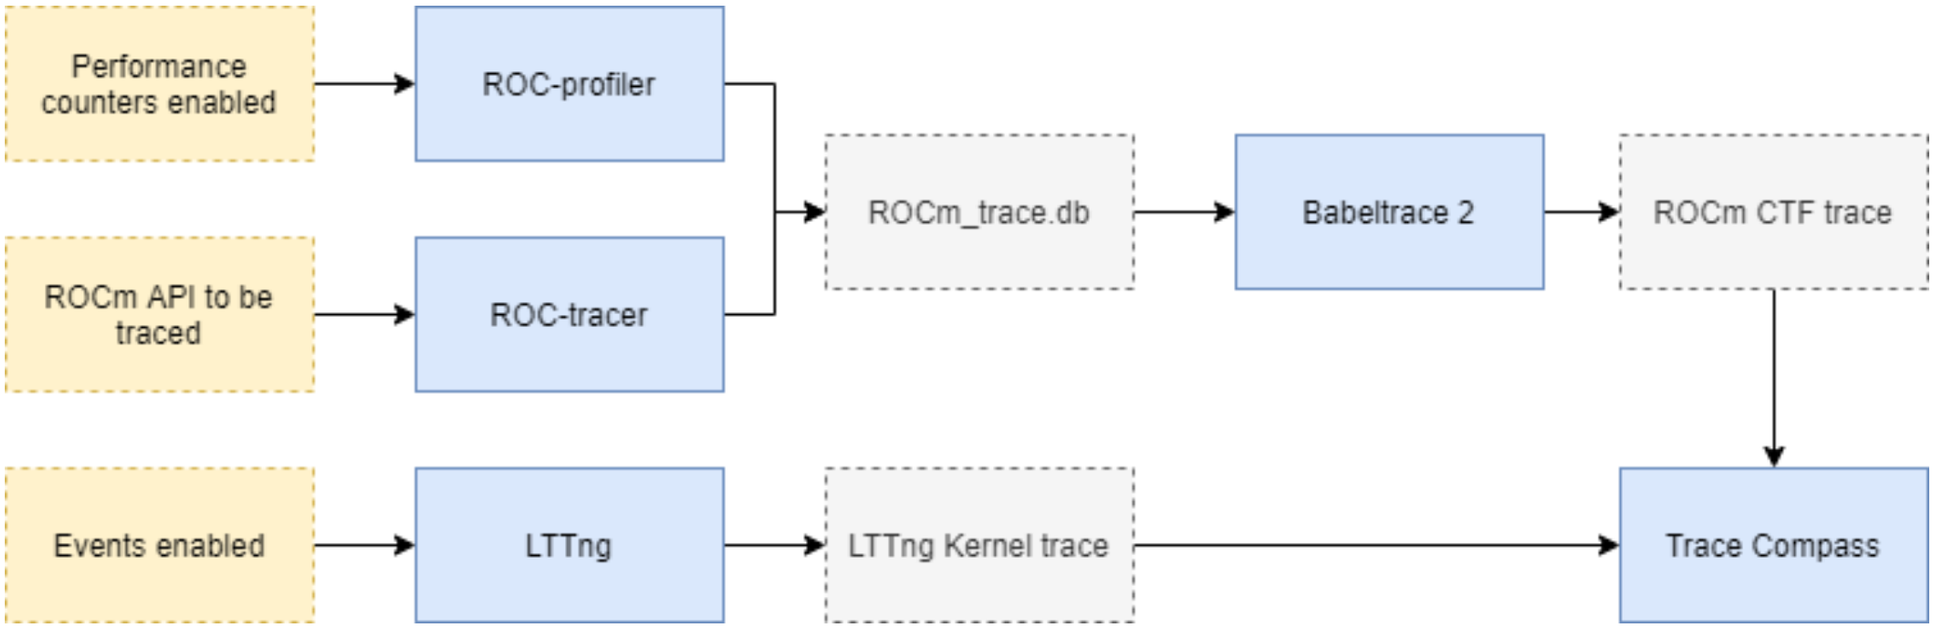

Tracing pipeline

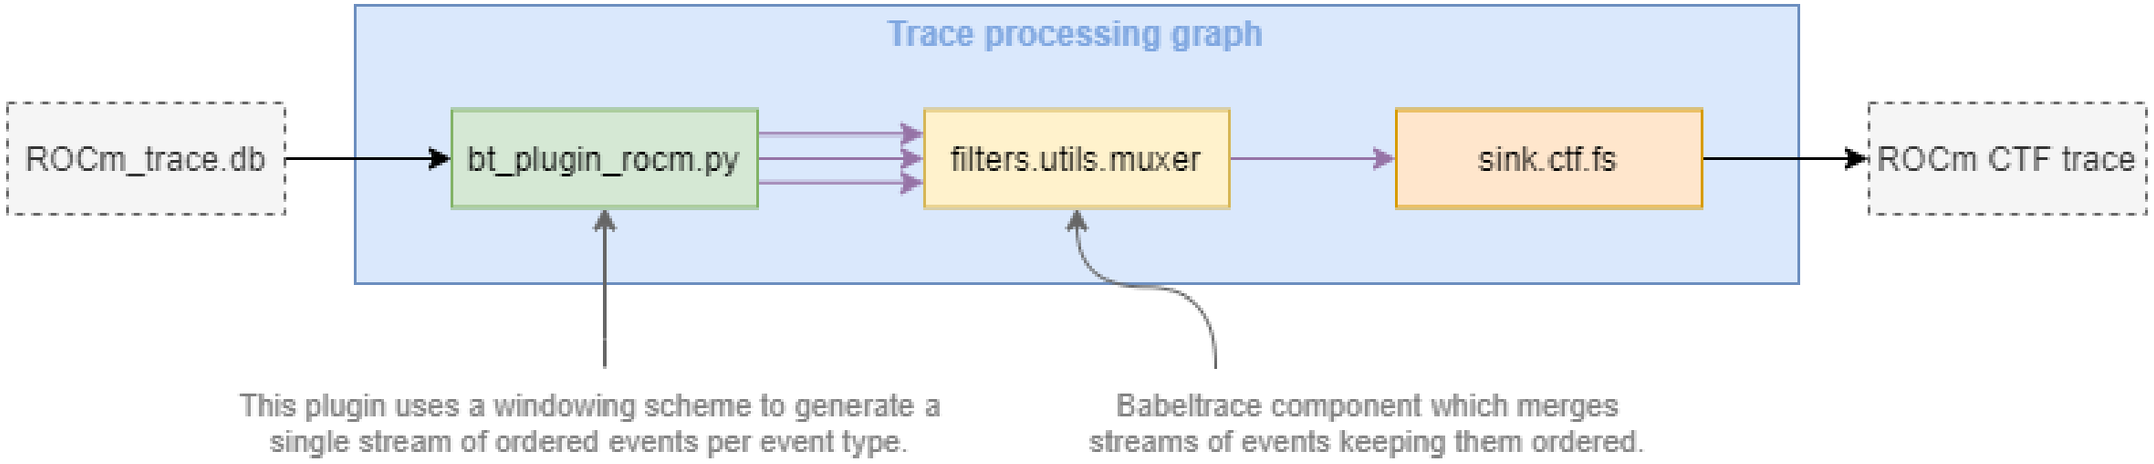

Results

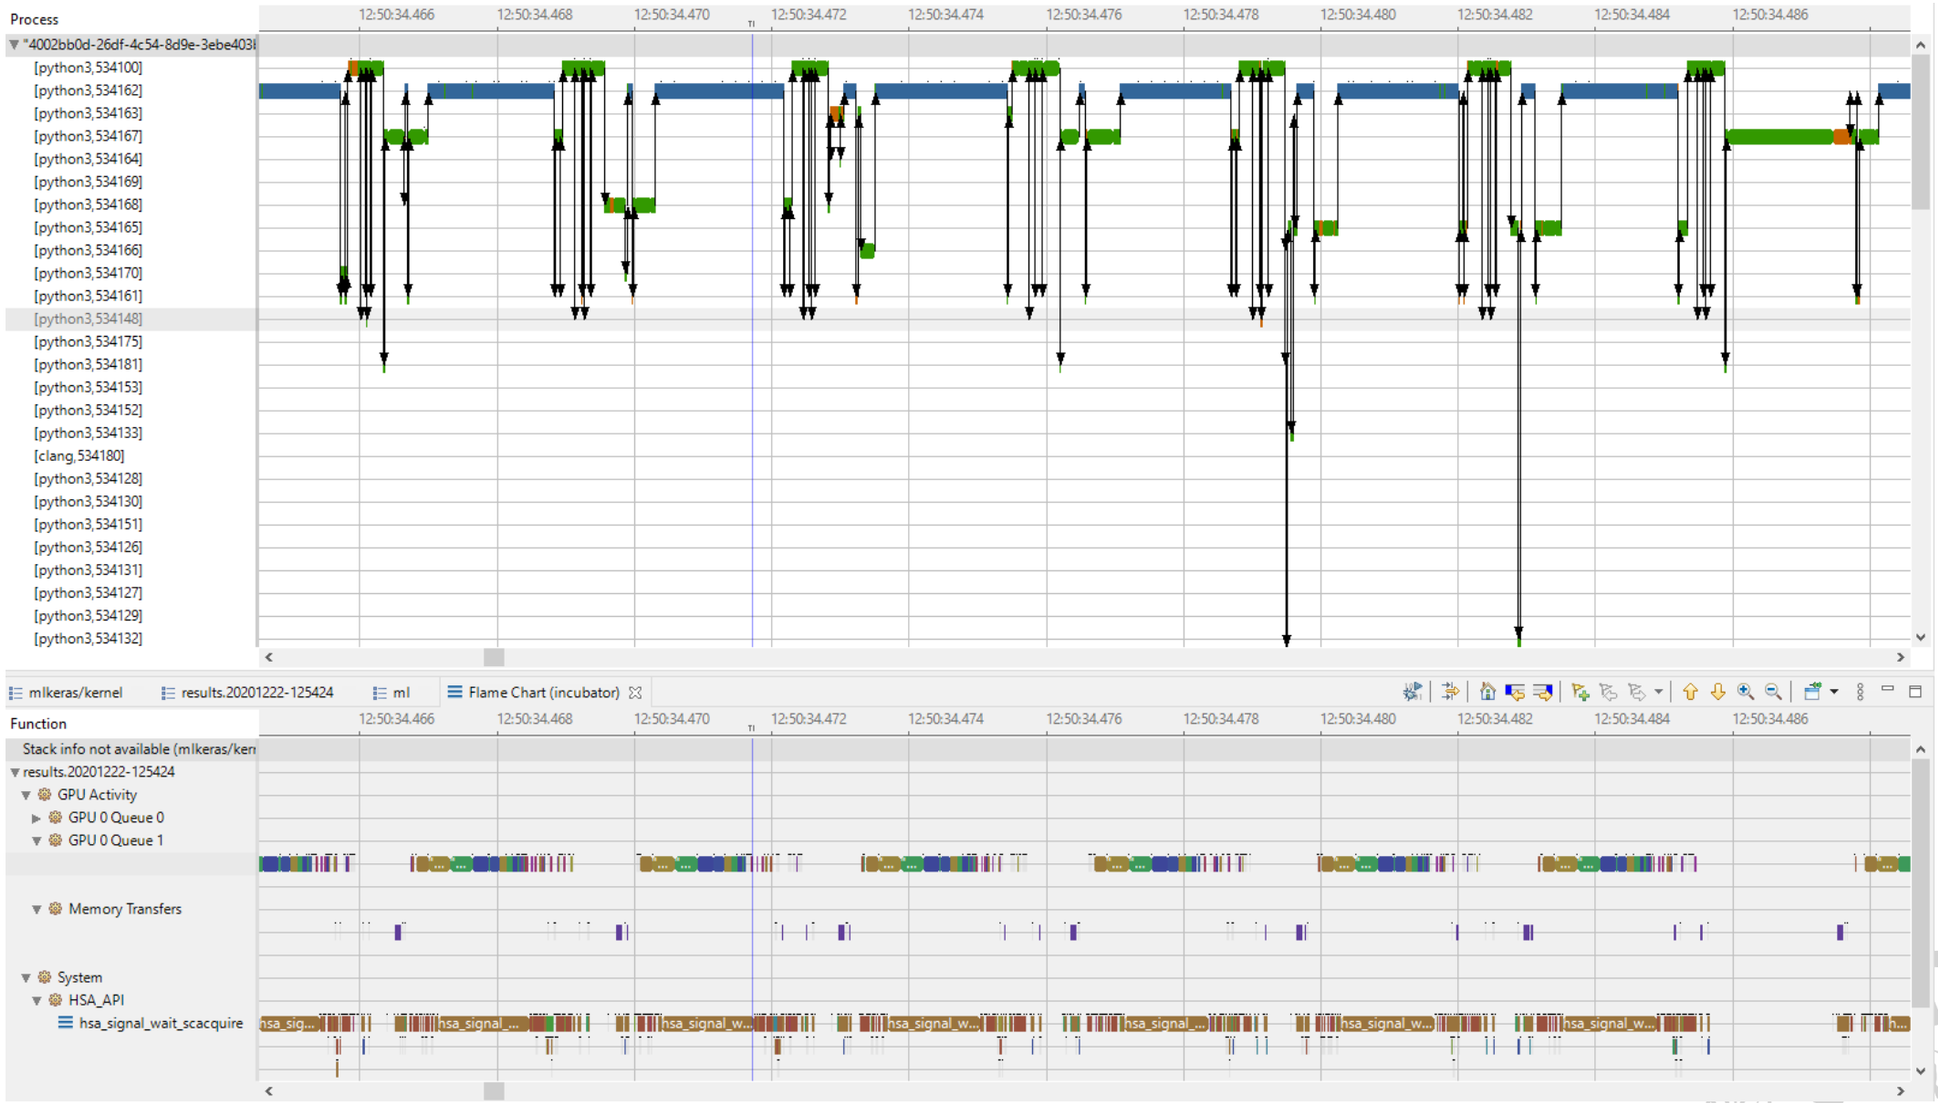
Process
12:50:34.466
12:50:34.468
12:50:34.470
12:50:34.472
12:50:34.474
12:50:34.476
12:50:34.478
12:50:34.480
12:50:34.482
12:50:34.484
12:50:34.486
"4002bb0d-26df-4c54-8d9e-3ebe403
[python3,534100]
[python3,534162]
[python3,534163]
[python3,534167]
[python3,534164]
[python3,534169]
[python3,534168]
[python3,534165]
[python3,534166]
[python3,534170]
[python3,534161]
[python3,534148]
[python3,534175]
[python3,534181]
[python3,534153]
[python3,534152]
[python3,534133]
[clang,534180]
[python3,534128]
[python3,534130]
[python3,534151]
[python3,534126]
[python3,534131]
[python3,534127]
[python3,534129]
[python3,534132]
mlkeras/kernel
results.20201222-125424
ml
Flame Chart (incubator)
Function
Stack info not available (mlkeras/ker
results.20201222-125424
GPU Activity
GPU 0 Queue 0
GPU 0 Queue 1
Memory Transfers
System
HSA_API
hsa_signal_wait_scacquire

Results

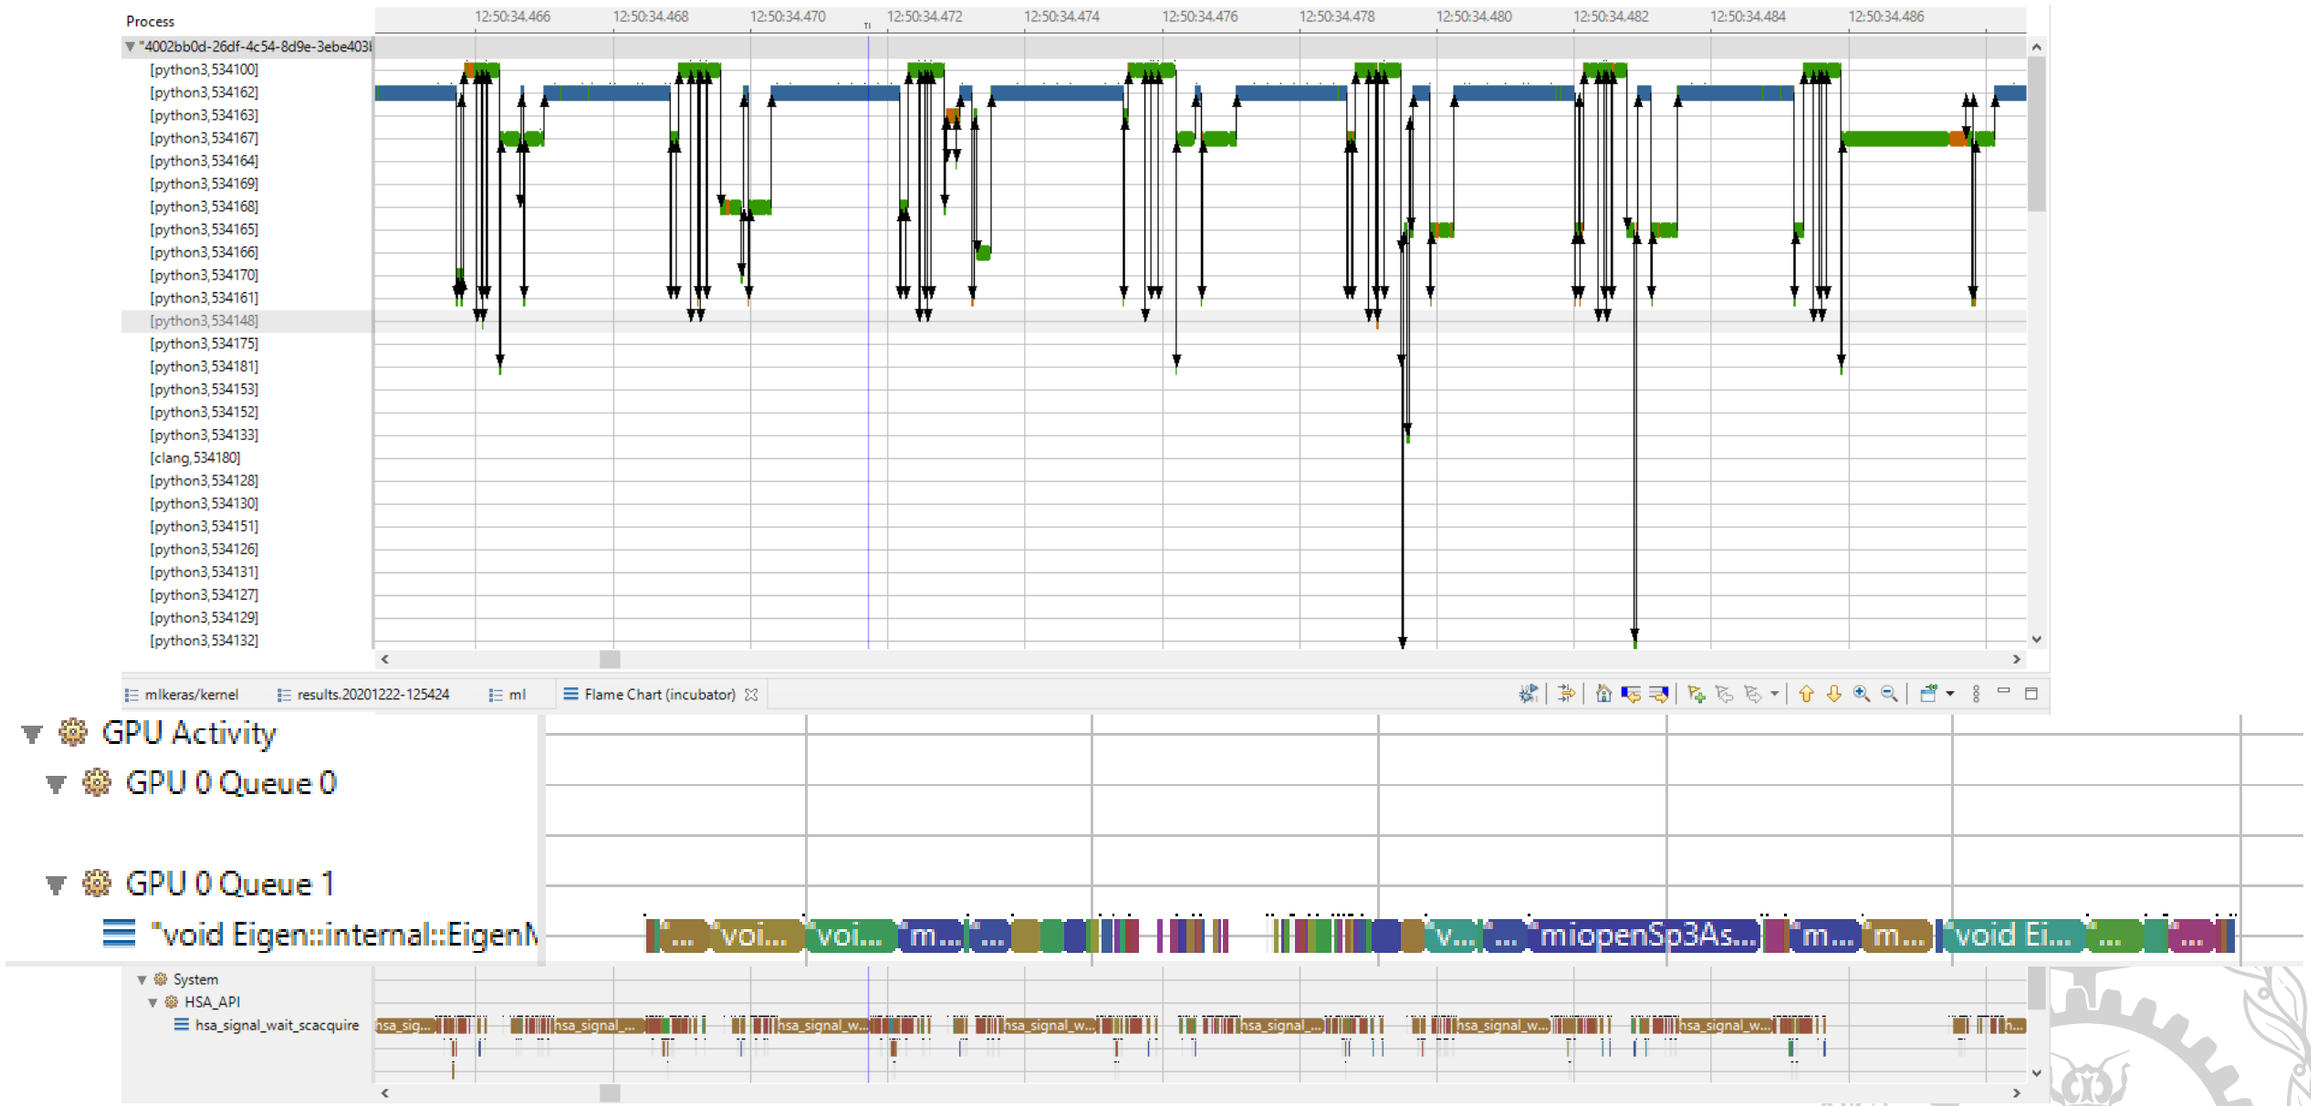
Process
12:50:34.466
12:50:34.468
12:50:34.470
12:50:34.472
12:50:34.474
12:50:34.476
12:50:34.478
12:50:34.480
12:50:34.482
12:50:34.484
12:50:34.486
[python3,534100]
[python3,534162]
[python3,534163]
[python3,534167]
[python3,534164]
[python3,534169]
[python3,534168]
[python3,534165]
[python3,534166]
[python3,534170]
[python3,534161]
[python3,534148]
[python3,534175]
[python3,534181]
[python3,534153]
[python3,534152]
[python3,534133]
[clang,534180]
[python3,534128]
[python3,534130]
[python3,534151]
[python3,534126]
[python3,534131]
[python3,534127]
[python3,534129]
[python3,534132]
mlkeras/kernel
results.20201222-125424
ml
Flame Chart (incubator)
GPU Activity
GPU 0 Queue 0
GPU 0 Queue 1
"void Eigen::internal::EigenM
voi...
voi...
m...
v...
miopenSp3As...
m...
m...
void Ei...
System
HSA_API
hsa_signal_wait_scacquire

Results

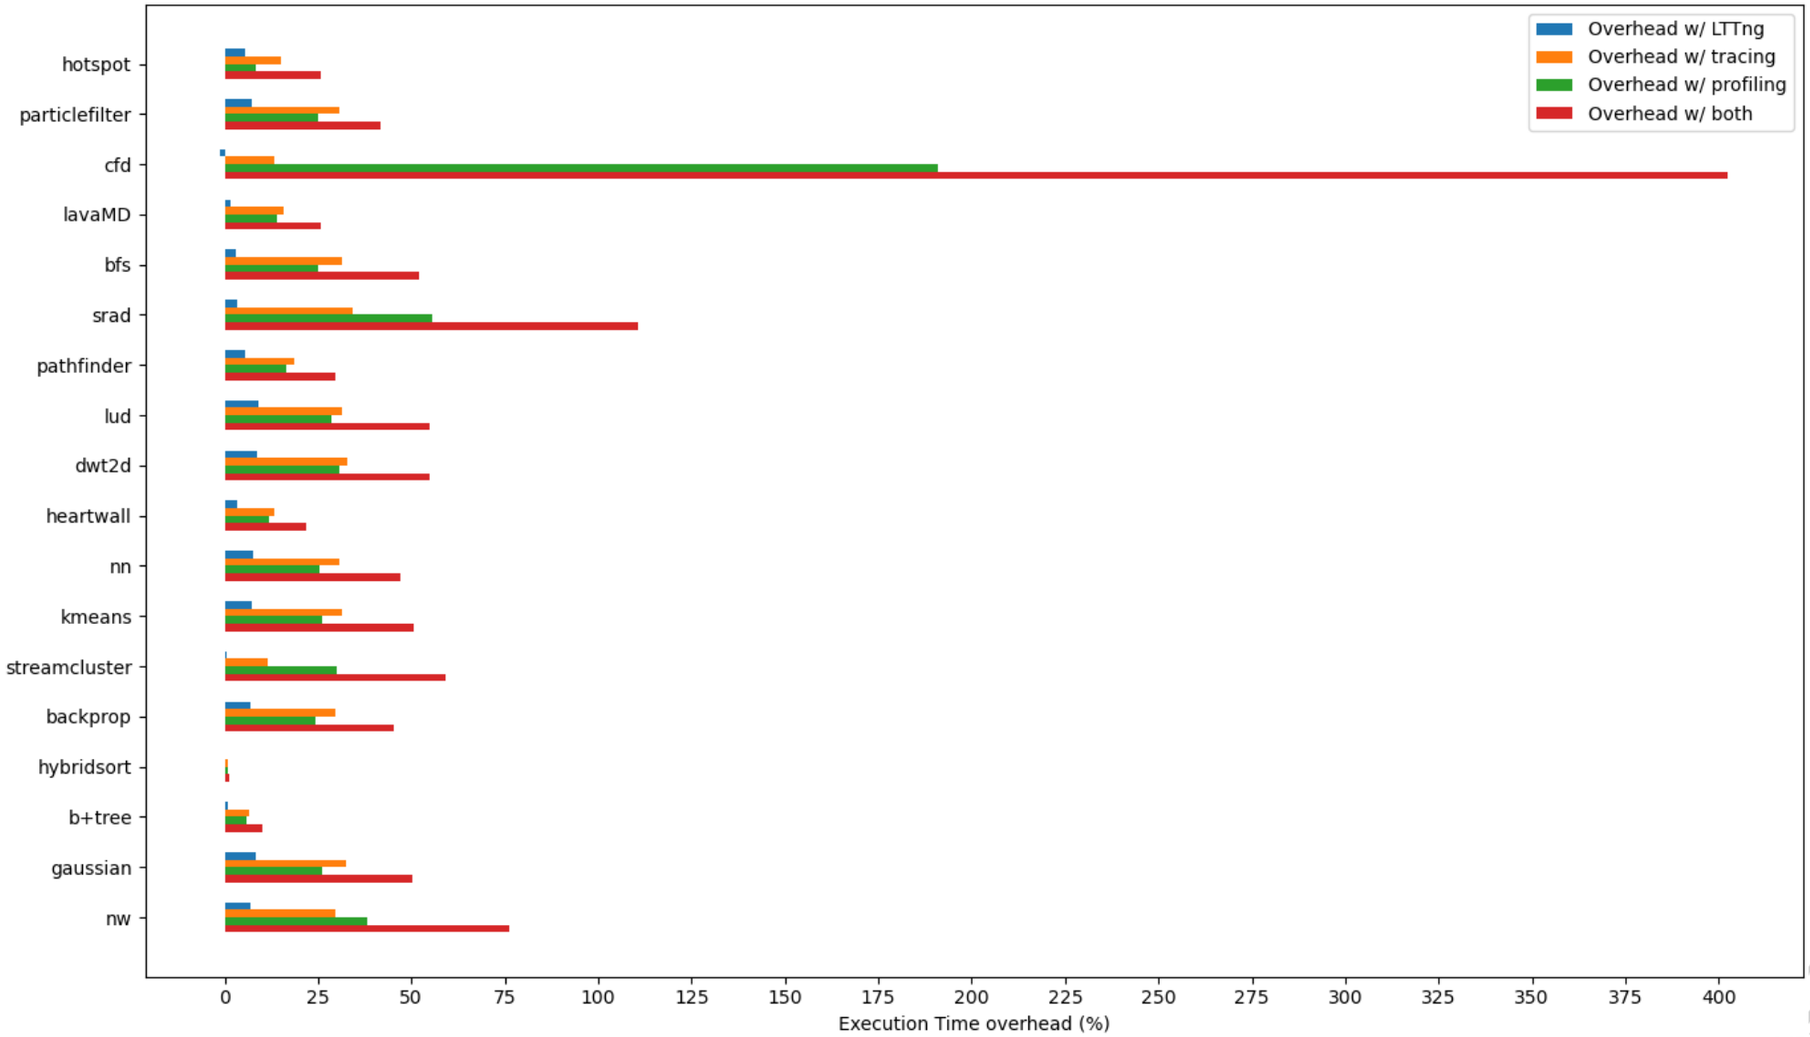

The experiences are run with LTTng and only enabling the HSA API.

Results

Experience	Number of kernels	Number of events	Trace Size
nn	1	8377	1.23 MB
hotspot	1	8413	1.24 MB
lavaMD	1	8434	1.24 MB
pathfinder	5	8641	1.24 MB
gaussian	4	8632	1.25 MB
backprop	2	8548	1.25 MB
b+tree	2	8725	1.27 MB
lud	46	10946	1.28 MB
dwt2d	27	10045	1.29 MB
kmeans	9	9220	1.29 MB
bfs	16	9603	1.30 MB
particlefilter	9	9894	1.38 MB
hybridsort	120	15934	1.49 MB
nw	255	22878	1.54 MB
heartwall	20	11265	1.54 MB
srad	502	40497	2.32 MB
cfd	16004	932665	23.25 MB
streamcluster	5549	456819	26.34 MB

The number of events is recorded for the ROCm trace whereas trace size takes into account the ROCm trace as well as the kernel trace.

Results

- The trace shown had a size of 1.72 GB.
- The trace size varies a lot depending on what API is traced.
- In particular, KFD events are very frequent

Future Work

- Synchronization of ROCm traces between different nodes
- Support for HIP streams
- Support for MPI programs
- And more…