CTF traces generation in ROCm and support for OTF2 in Trace Compass

Yoann Heitz

2022/01/12

Polytechnique Montreal

DORSAL Laboratory

Agenda

CTF traces generation in ROCm :

- Context
- Solution and first results
- Improvements

OTF2 support in Trace Compass :

- Background
- OTF2 to CTF converter
- MPI views in Trace Compass

Demo : Tracing a distributed program running on AMD GPUs

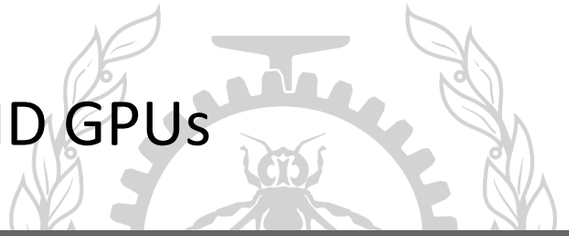

CTF traces generation in ROCm : Context

- ROCm : a platform for GPU computing
- ROC-profiler and ROC-tracer allow to trace several APIs (HIP, HSA) and performance metrics
- Arnaud developed analyzes in Trace Compass

CTF traces generation in ROCm : Context

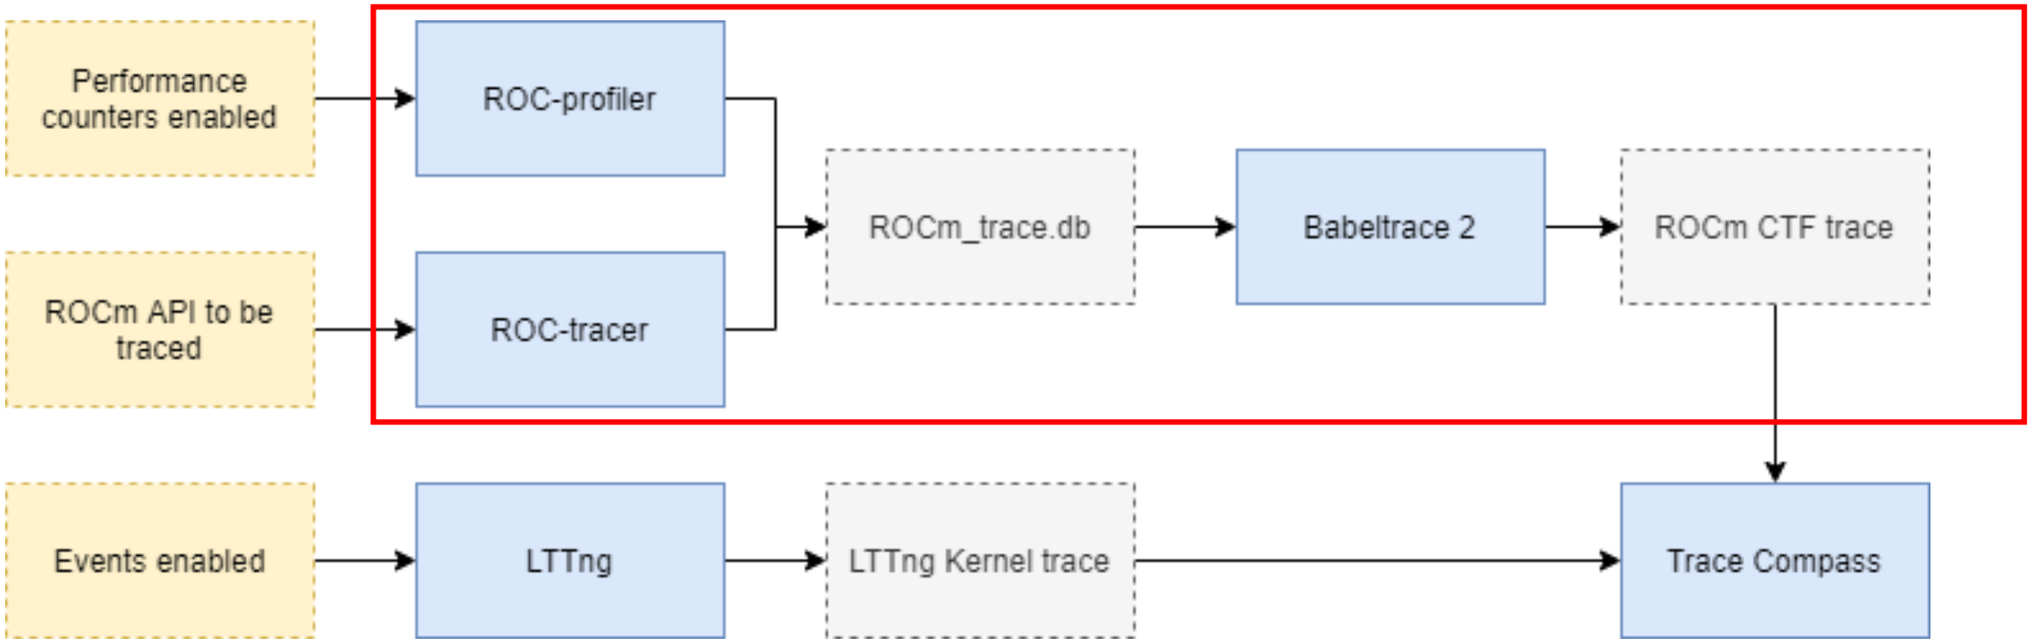

CTF traces generation in ROCm : Context

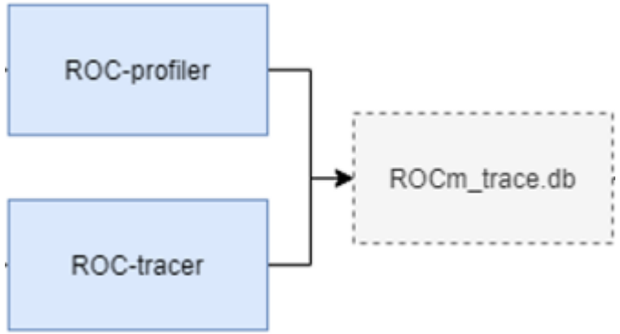

- Overhead due to tracing
- Intermediate text format
- Overhead due to conversion to .db

- Overhead due to conversion to CTF
- Plugin written in Python

CTF traces generation in ROCm : Solution

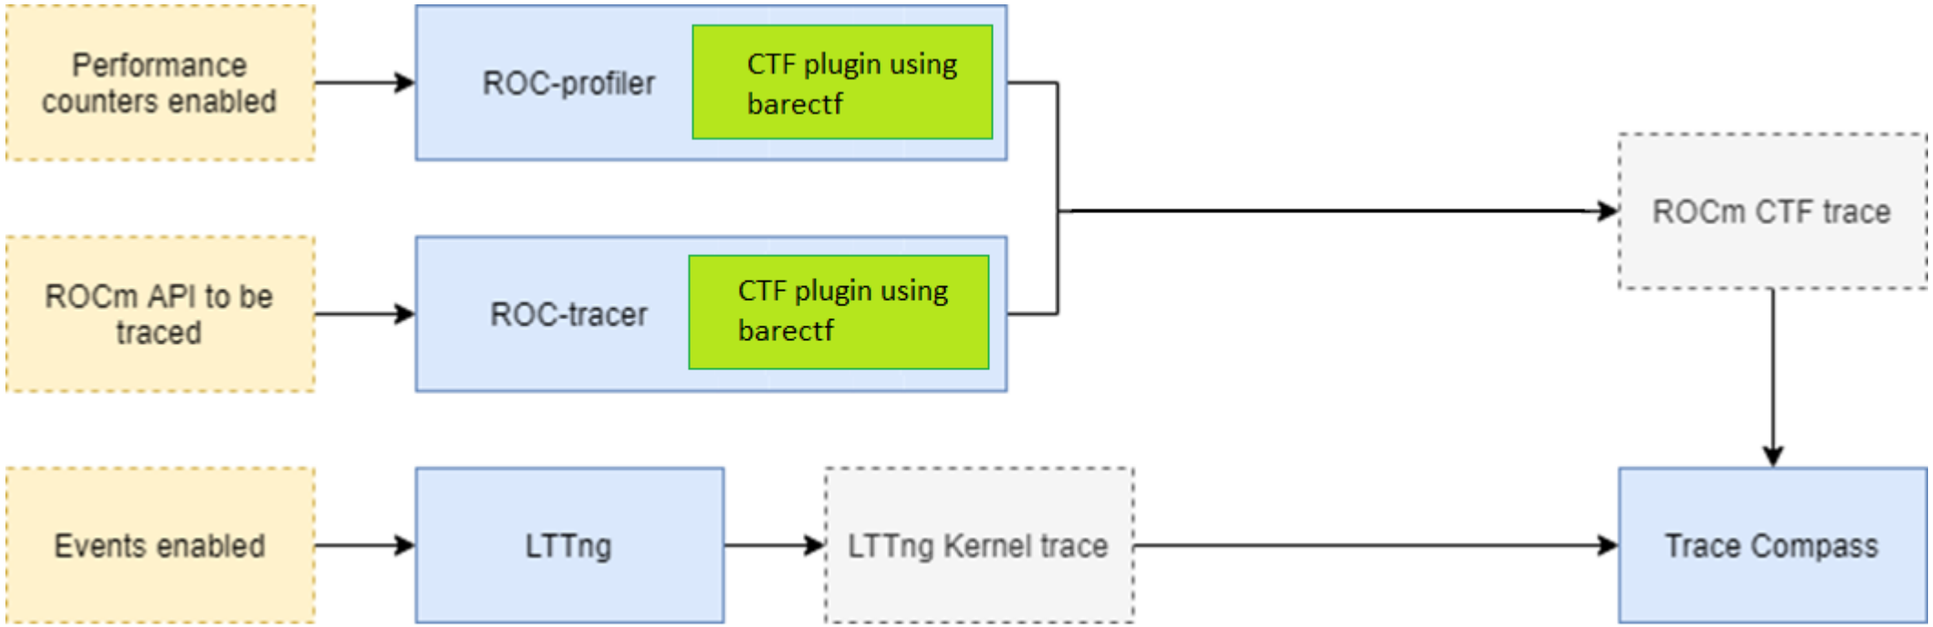

CTF traces generation in ROCm : Solution

Issues :

- events representing intervals
- unordered events
- ROC-profiler and ROC-tracer are multithreaded tools

Initial solutions :

- 2 events per interval
- reordering at runtime based on beginning time of the event
- use of exclusive CTF streams and tracing structures

CTF traces generation in ROCm : Results (2 events per interval)

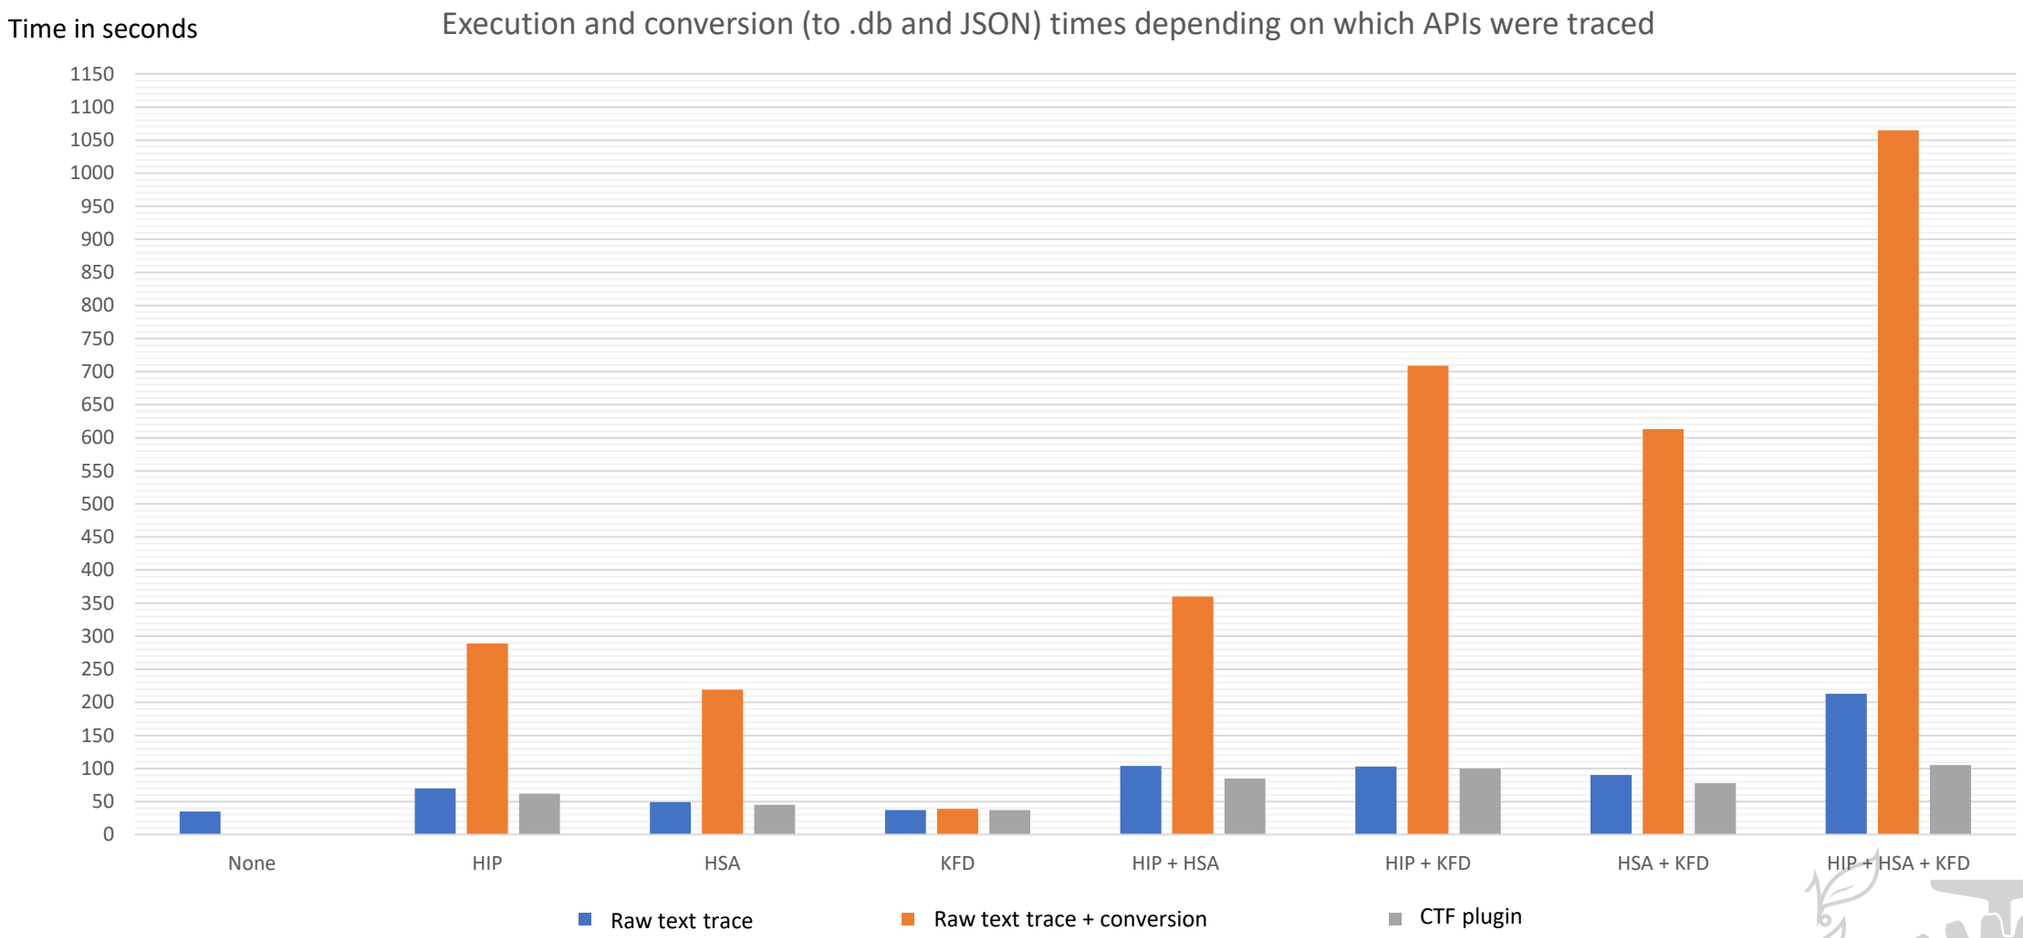

CTF traces generation in ROCm : Results (2 events per interval)

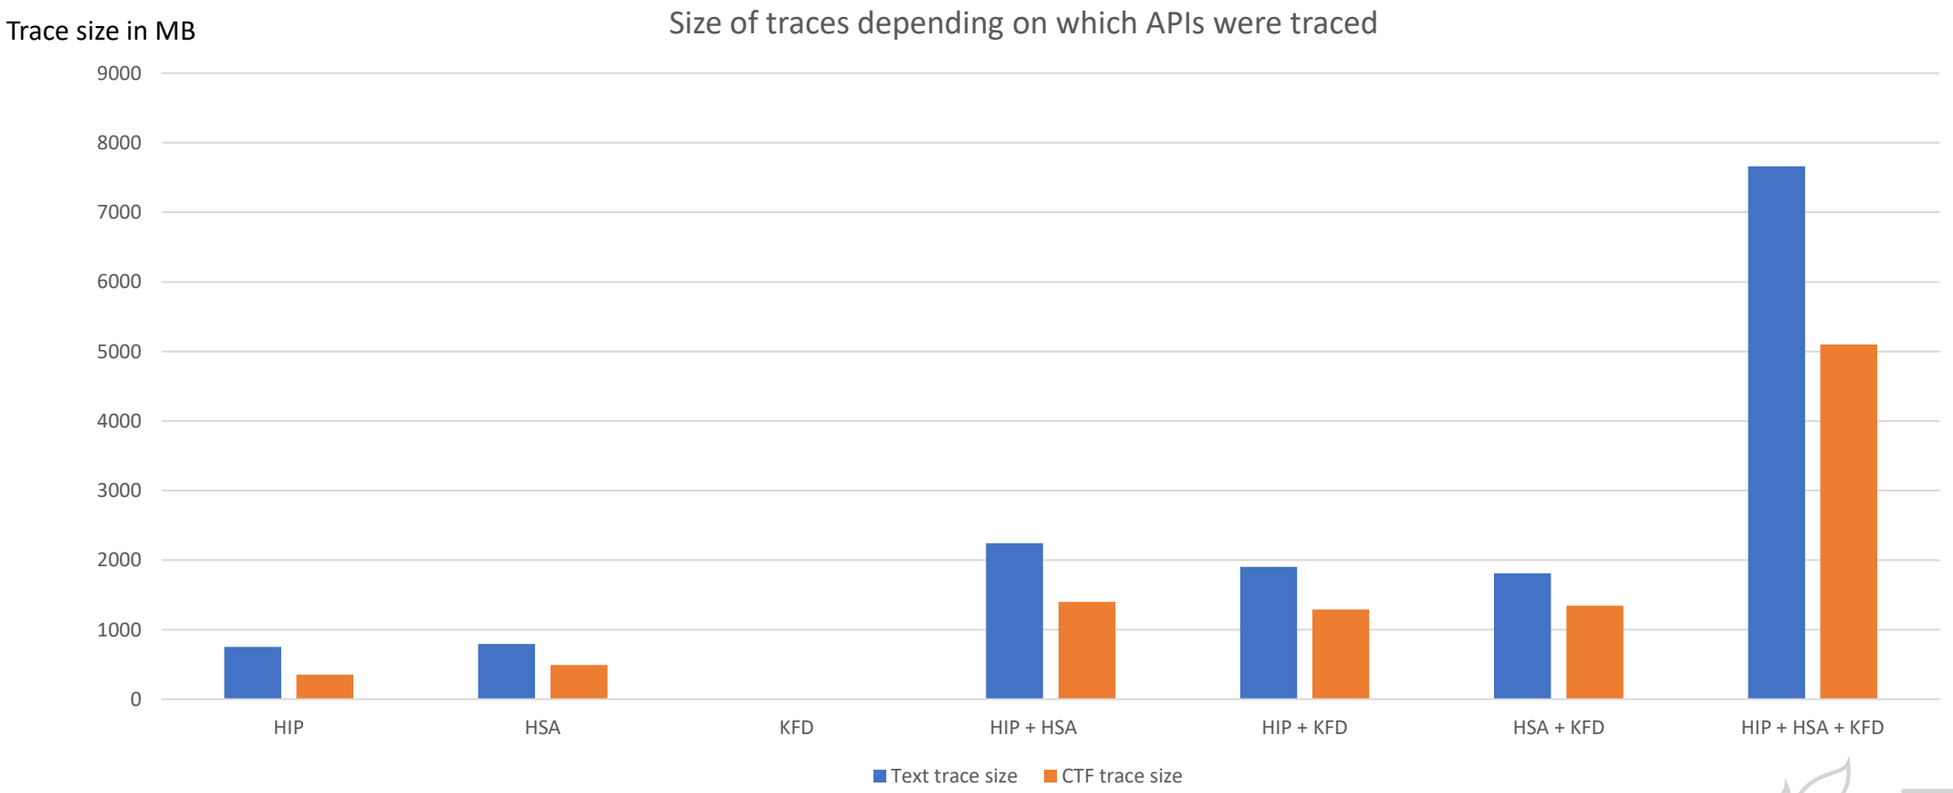

Traces 1,5x lighter in CTF format than in text format

CTF traces generation in ROCm : Improvements

- One event in the form of an interval as in the initial text format
  - Lighter traces
  - Faster processing in TraceCompass
- Better ease of use on the user side
  - Integration of the plugin into ROCm baseline version in progress
- Interface implementation
  - Possibility to implement new plugins for other trace formats

CTF traces generation in ROCm : Improvements

Interval format tracing durations:

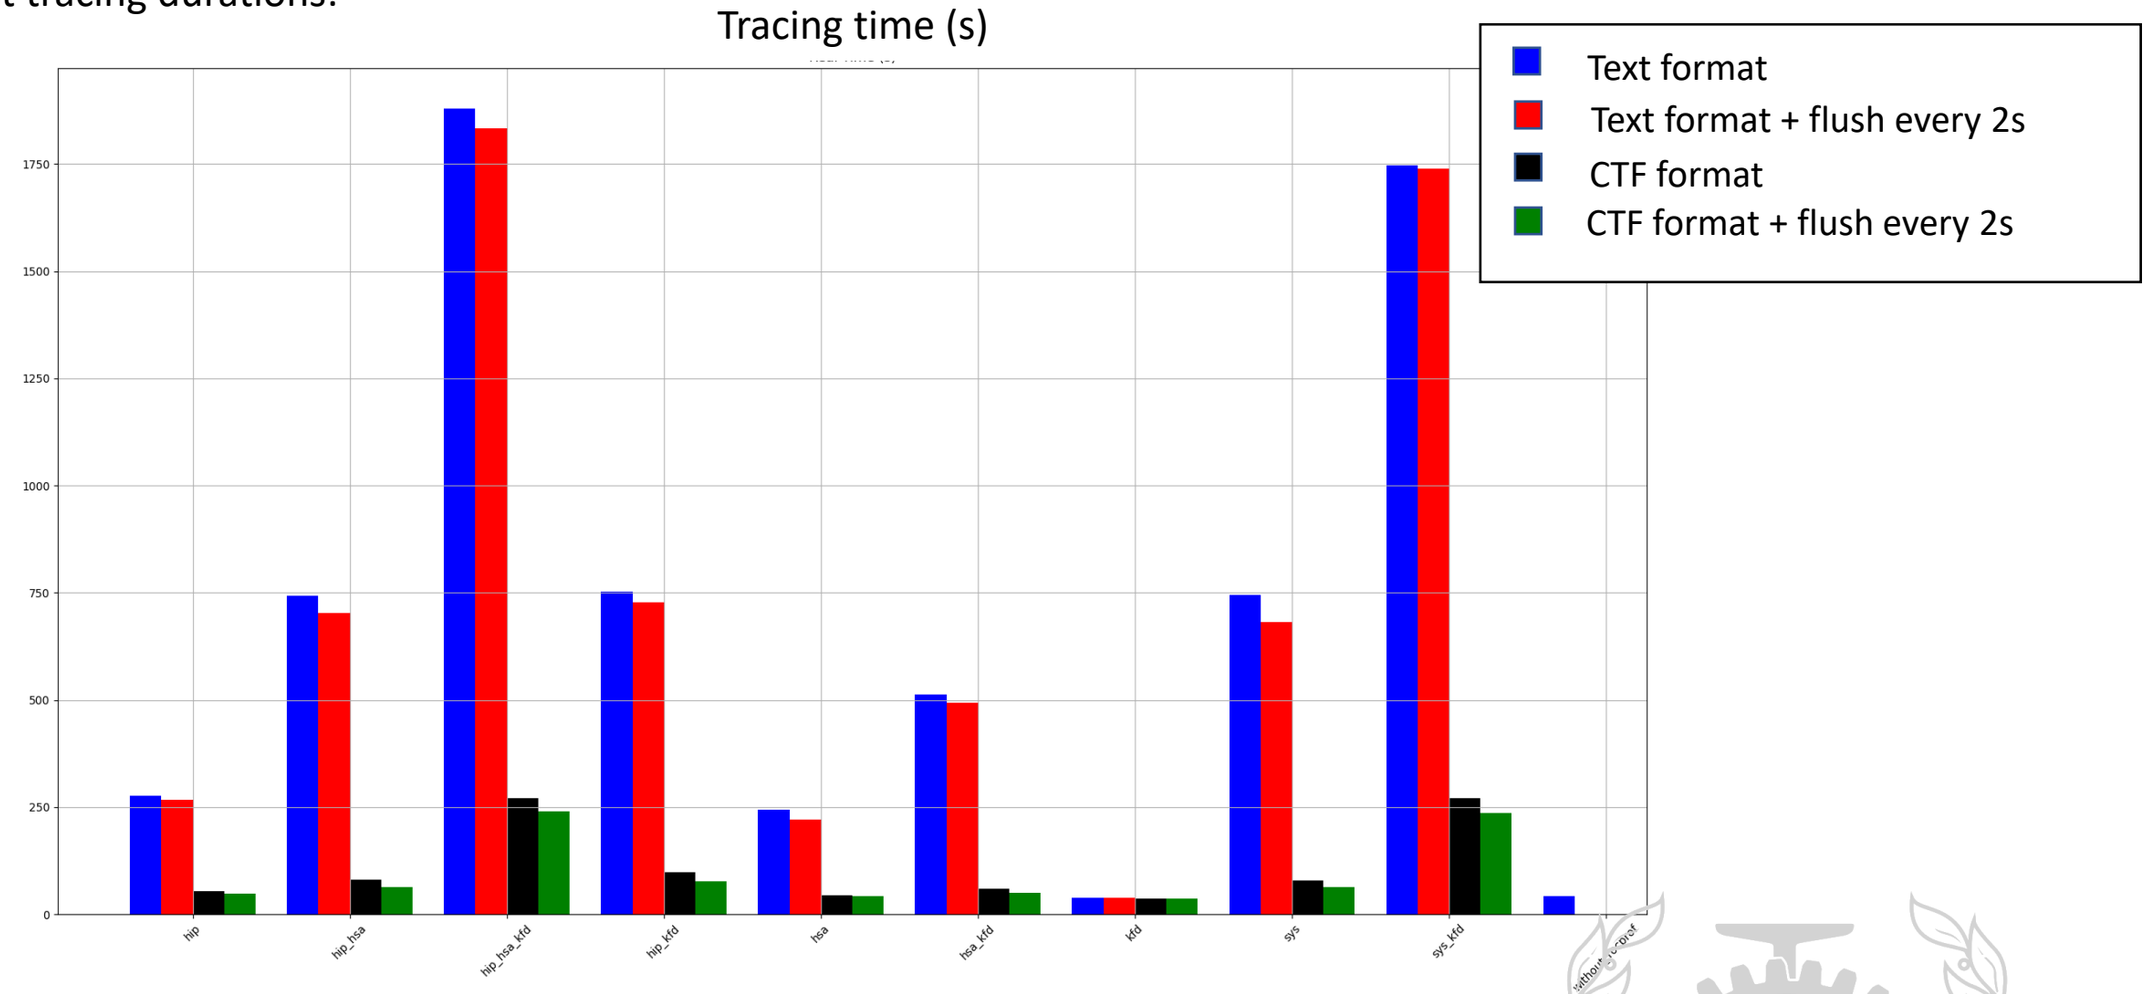

Roughly the same ratios compared with 2 events per interval

CTF traces generation in ROCm : Improvements

Interval format traces size :

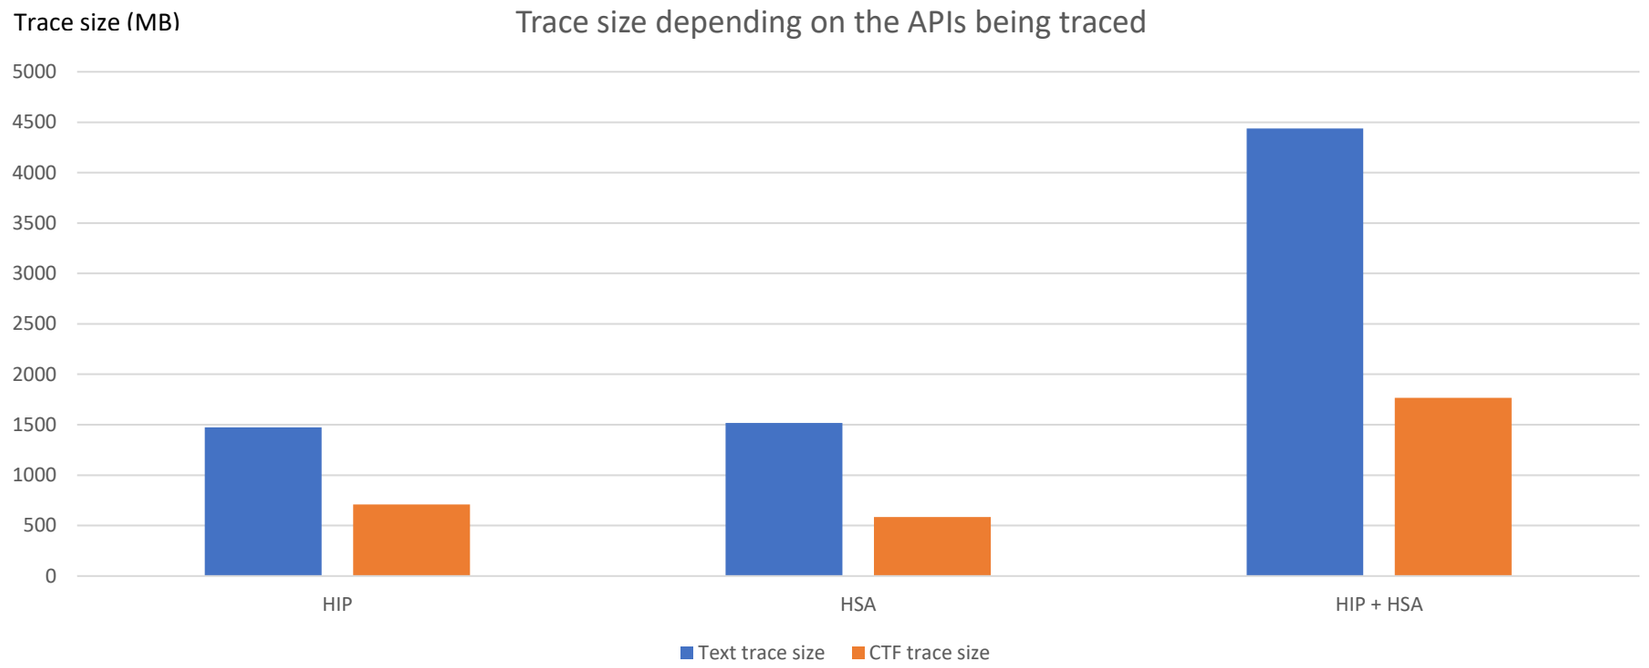

Traces 2x lighter in CTF format (with intervals) than in text format

CTF traces generation in ROCm : Improvements

Demo : using the plugin

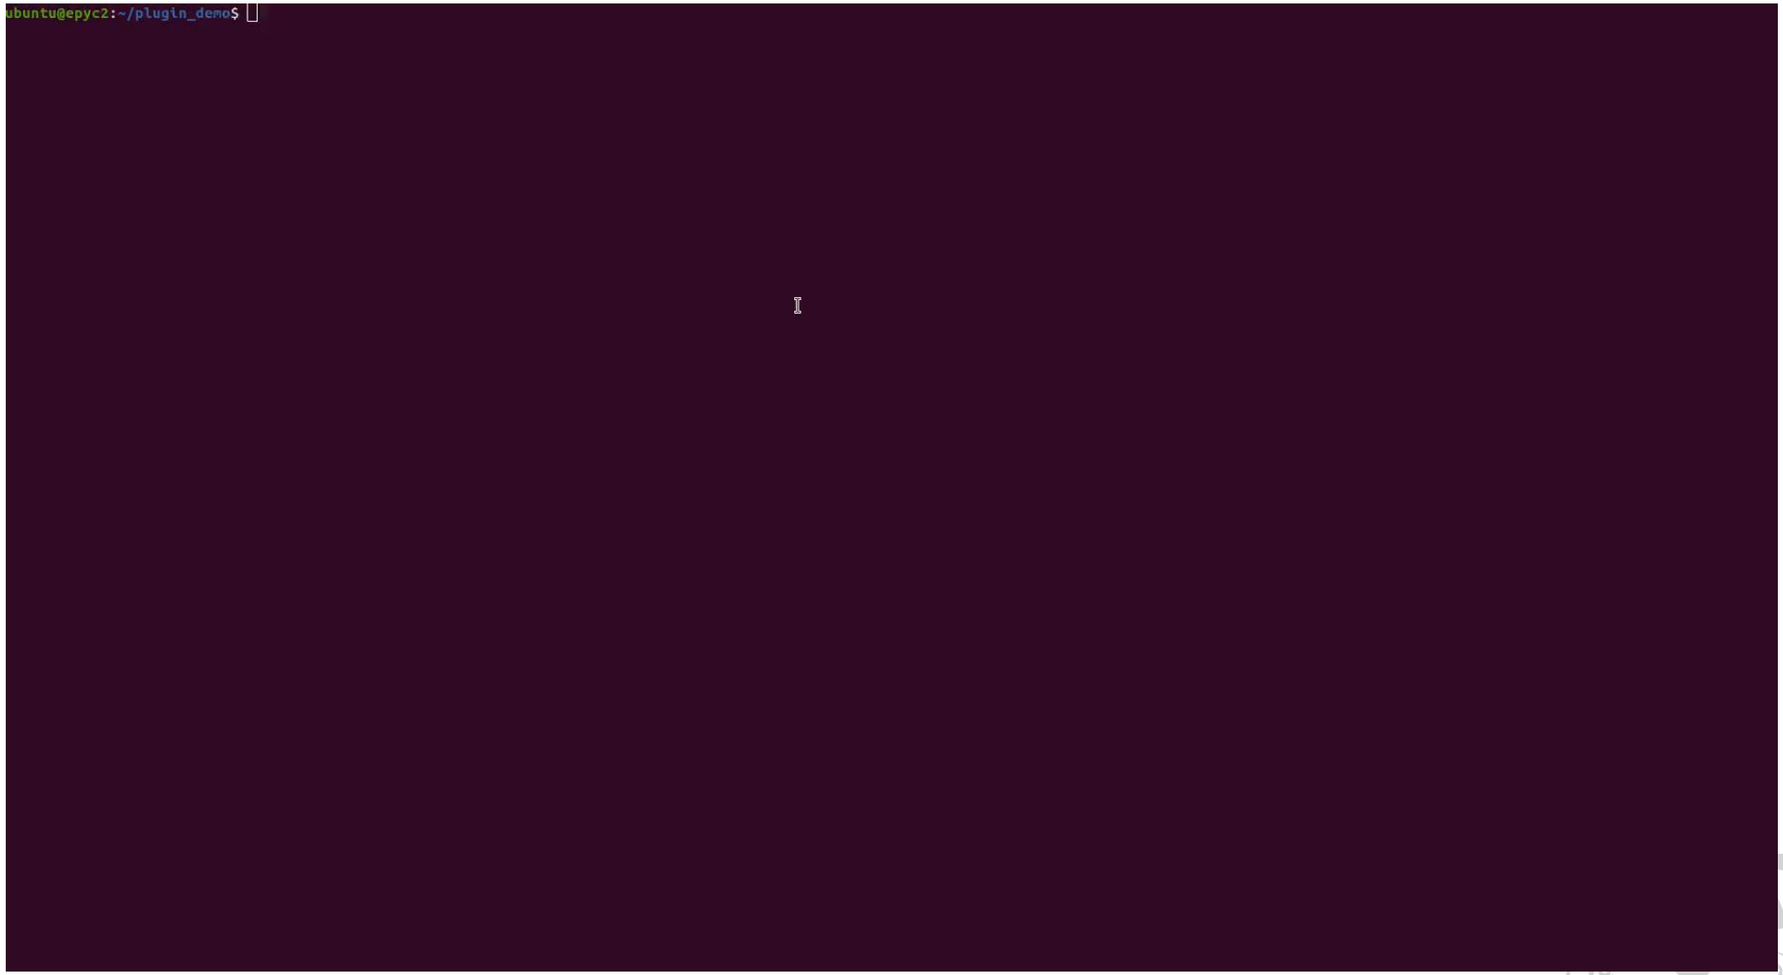

OTF2 support in Trace Compass : Background

OTF2 :

- Binary trace format
- C API
- MPI, OpenMP, Pthreads events

OTF2 support in Trace Compass : Background

Supporting a new trace type in Trace Compass :

1) Parse the trace
2) Convert events into the Trace Compass internal event format
3) Read the events and provide analysis

OTF2 support in Trace Compass : OTF2 to CTF converter

Parse and read the trace in Trace Compass:

- Convert it in CTF format with barectf
- Use the Trace Compass CTF parser

OTF2 support in Trace Compass : OTF2 to CTF converter

- First major version of the converter has been implemented
- Convert every event defined in the OTF2 format
- No features associated to snapshots and markers (future work)

OTF2 support in Trace Compass : MPI views in Trace Compass

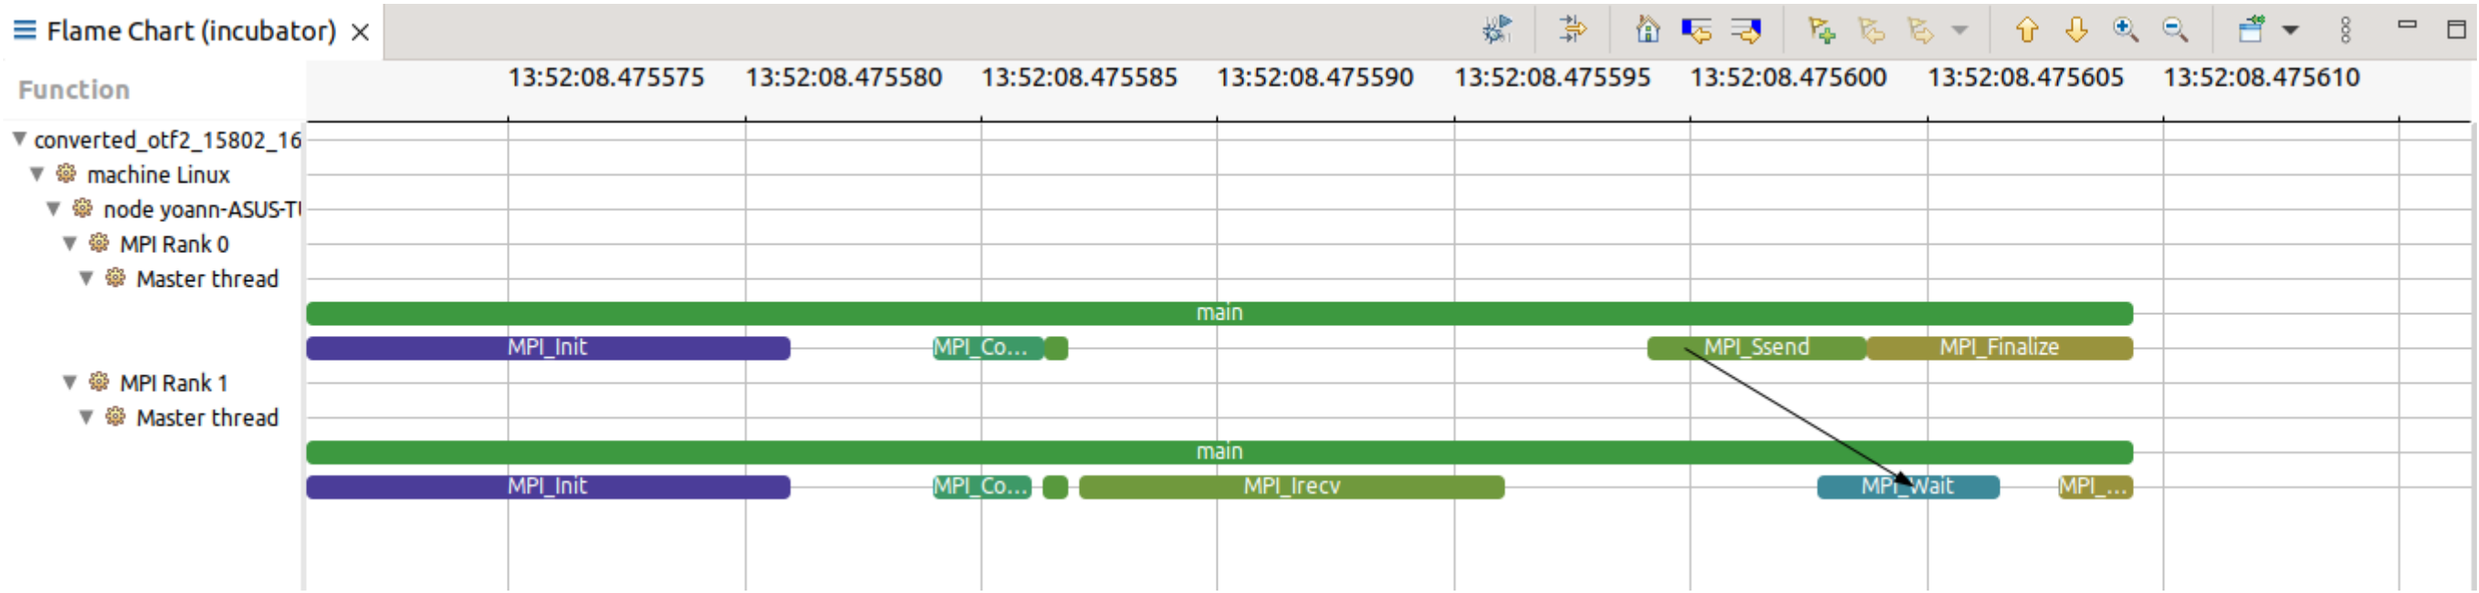

Callstack analysis : Flame Chart view

OTF2 support in Trace Compass : MPI views in Trace Compass

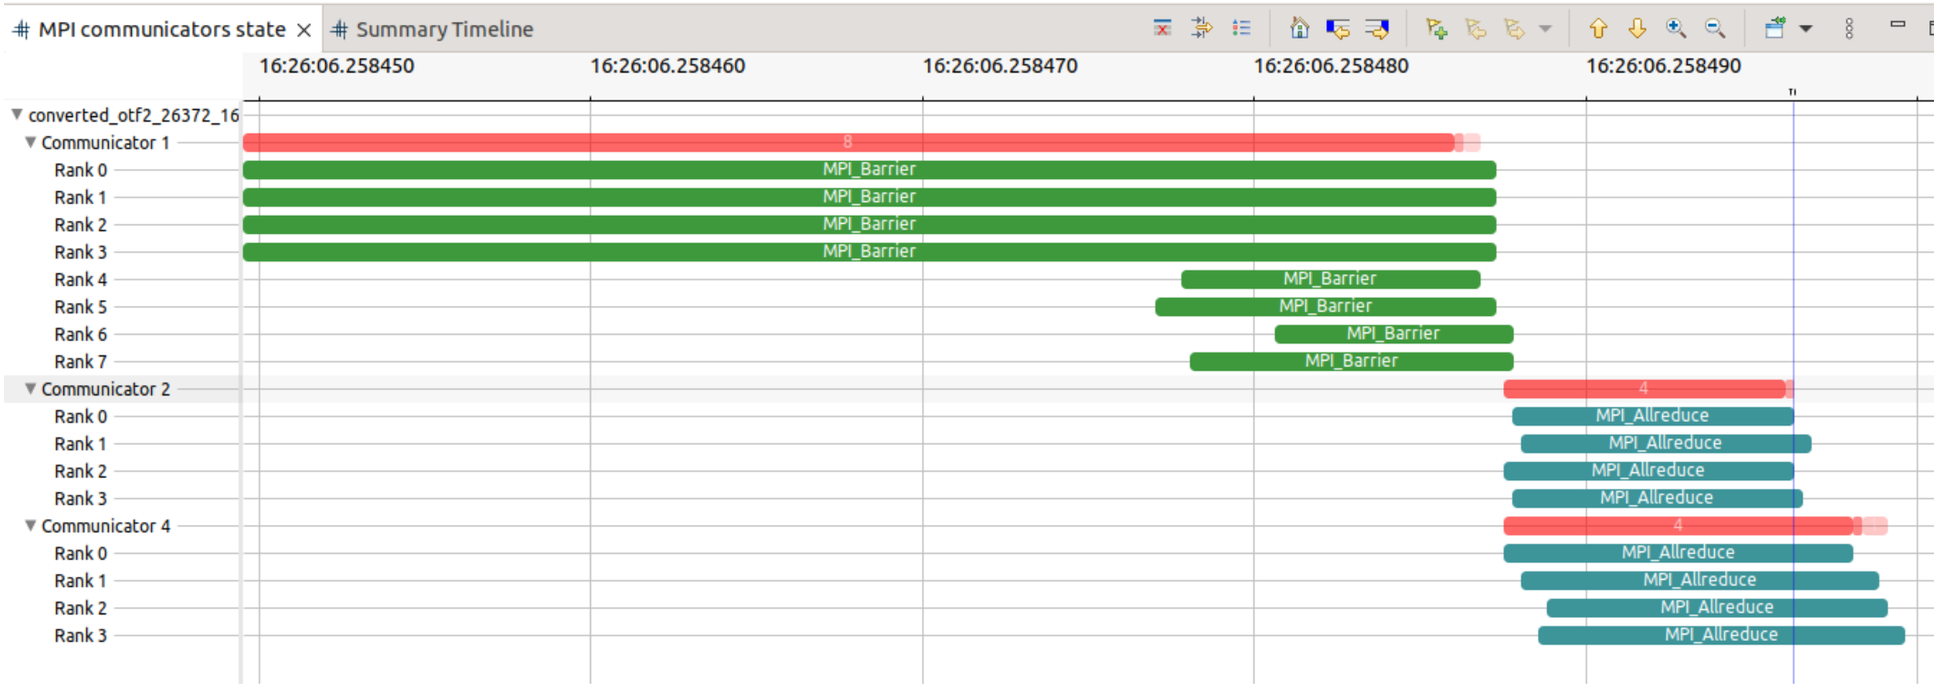

Communicators analysis

OTF2 support in Trace Compass : MPI views in Trace Compass

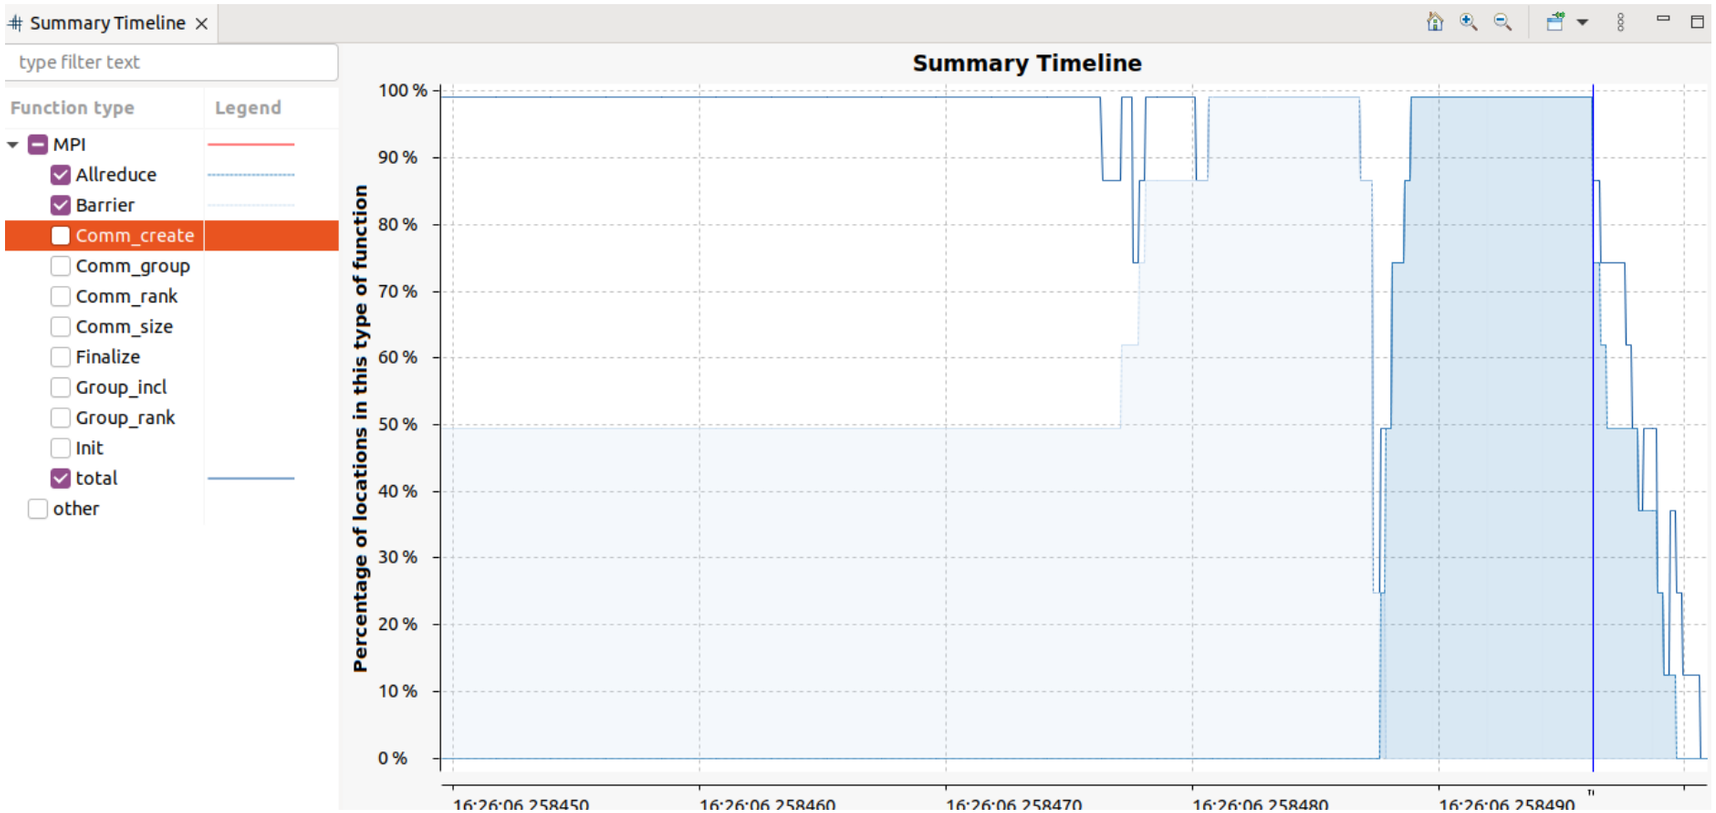

Summary timeline analysis

Demo : Tracing a distributed program running on AMD GPUs

Thank you for your attention!

Any questions?

Annex : Links to the tools

CTF plugin : https://github.com/dorsal-lab/rocprofiler_ctf_plugin

rocprof interfaces :

- ROCProfiler : https://github.com/dorsal-lab/rocprofiler
- ROCTracer : https://github.com/dorsal-lab/roctracer

OTF2 to CTF converter : https://github.com/dorsal-lab/OTF2-to-CTF-converter