



Updates on Scalability, MPI, DPDK and ROCm in Trace Compass

Arnaud Fiorini

Polytechnique Montréal
Laboratoire DORSAL

Introduction

- MPI:
 - Interface for communication in HPC
 - Instrumented by Score-P and Exa-Tracer through PMPI
- DPDK:
 - Collection of user-space libraries for fast packet processing
 - Supports many processor architectures
- ROCm:
 - Collection of libraries for AMD GPUs
 - Includes profiling and tracing tools

Agenda

- 1 Scalability improvements in Trace Compass
- 2 MPI Exa-Tracer integration
- 3 ROCm integration status
- 4 DPDK instrumentation status



Scalability improvements in Trace Compass

- Improving query speed and sht file size with Partial State System
- Fixing some performance and user experience issues with large traces
- Trace coordinator for analysis pre-processing



MPI Exa-Tracer integration

- Previous implementation based on OTF2 conversion
- Fast prototype thanks to the CallStack plugin
- Refactoring needed to extend the current implementation



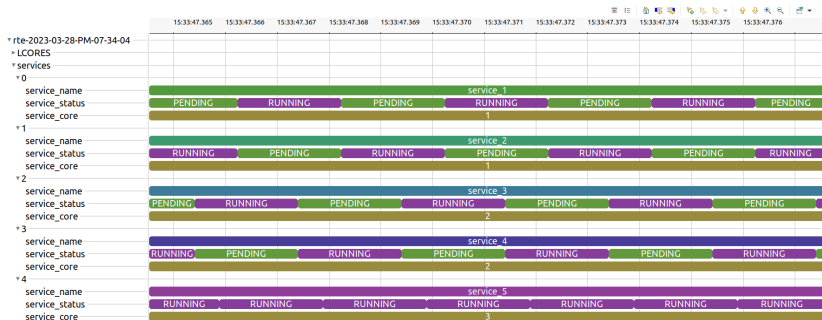
ROCm

- The event format will change soon
- Other sources which output GPU runtime traces can be added
- Will be used as a benchmark for the Partial State System



DPDK

- Work done by Adel is being merged
- Tracepoints are being contributed to DPDK



DPDK

- 46 out of the 98 tracepoints have been implemented
- The Trace Compass integration is not completed
- Feedback and/or use cases would be very appreciated



DPDK – Future Work

- Use cases to generate traces and improving the implementation quality in Trace Compass
- Continue the instrumentation of DPDK



References

- https://doc.dpdk.org/guides/prog_guide/trace_lib.html
- <https://github.com/DPDK/dpdk>
- <https://github.com/ROCM-Developer-Tools/rocprofiler>

