



**POLYTECHNIQUE
MONTREAL**

TECHNOLOGICAL
UNIVERSITY

Improvements to the ROCm plugin in Trace Compass

Arnaud Fiorini

Laboratoire DORSAL

Introduction

- Specific hardware, environment
- Complex toolchain
- Traditional tools are hard to use

Agenda

① Demonstrations

- ROCm (ROC-profiler, ROC-tracer, ROCgdb)
- Trace Compass

② Scalability of Trace Analysis

- Trace Compass improvements
- Distributed trace analysis



ROCm

- AMD open source software
- Compute stack for headless system deployments
- Mainly interfaced with OpenCL, OpenMP and HIP
- HIP is a CUDA-like interface to schedule computations on GPUs



GPU Programming model

- Device code, also called **kernel function**

```
1 __global__ void helloworld(char* in, char* out)
2 {
3     int num = hipThreadIdx_x + hipBlockDim_x * hipBlockIdx_x;
4     out[num] = in[num] + 1;
5 }
```



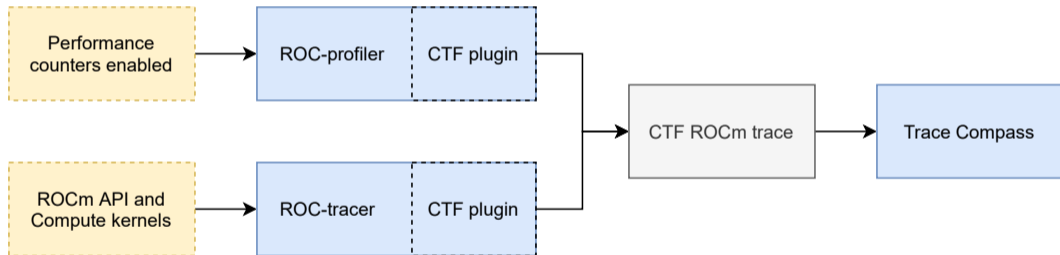
GPU Programming model

- 1 Allocating space on the device
- 2 Copy your memory over to the device
- 3 Perform a **kernel launch** with the kernel function
- 4 Copy back the result to the host
- 5 Free the space on the device

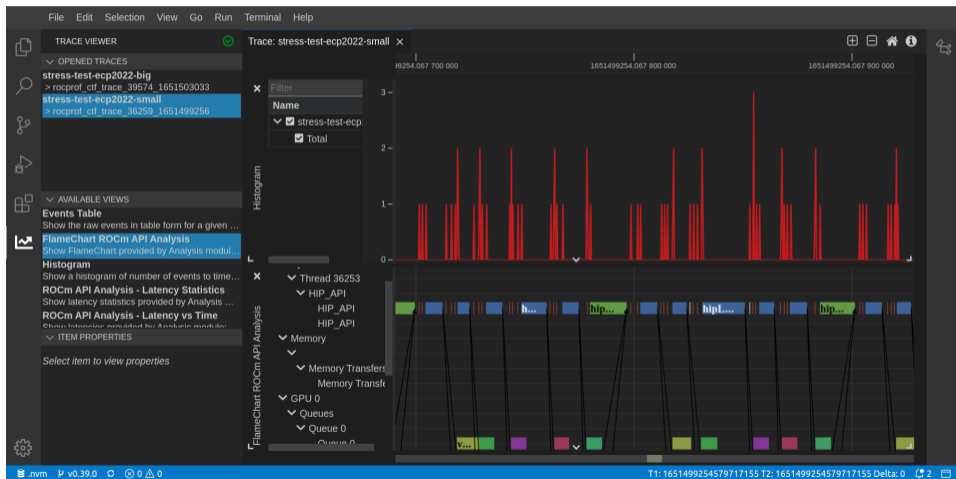
These steps are done by the user with the HIP API (CUDA-like API)



Tracing pipeline



Demo



Scalability

Current scalability improvements on Trace Compass¹:

- Preprocessing to reduce the wait time
- Partial State System
- Distribute the analysis on multiple nodes
- Critical path on disk
- Multiple patches that improve performance

¹Abdellah Rahmani, Quoc-Hao Tran, Geneviève Bastien, Matthew Khouzam



Future Work

- Implement Critical path for GPU and HPC workloads
- Performance metrics analysis
- Overview Analysis for large scale applications



Contributions

- Interfacing with ROCm to handle CPU and GPU traces and performance metrics
- Optimizing Trace Compass scalability to handle large traces
- UX improvements for HPC use cases

