

GPU Instrumentation Group Outbrief

Scalable Tools Workshop

June 22th, 2023

State of instrumentation tooling

- Intel : GT-Pin
- Nvidia : NVBit/SASSI
- Dyninst support for AMD GPUs
 - Extract/inject into fatbin with provided tools
- Compile-time instrumentation

Instrumentation is up to the user

Dyninst can (will) instrument, it's up to the user to decide what to do

- Memory management is a major issue
 - Where data is stored
 - How data is sent to the CPU
- e.g. branch divergence analysis

Challenges for instrumentation

- Registers, LDS are scarce resources
- Vector ops & masking pose a challenge for liveness analysis
- Dynamic instrumentation not planned for now

AMD compiler team is skeptical about how feasible instrumentation is for some kernels : LLNL's Quicksilver, some MIOpen (AMD's ML lib) kernels

Few things can be leveraged :

- ISA provides interesting info (EXEC mask, ...)
 - Some are more cryptic (S_SENDMSG?)
- Facilities to 'ring the doorbell' for the host or other devices
- ATT Tracing?

Need more interaction with vendors