Performance Monitoring & Queries on Intel GPUs

Lionel Landwerlin

27 September 2018

Hardware overview
i915 interface
Userspace tools

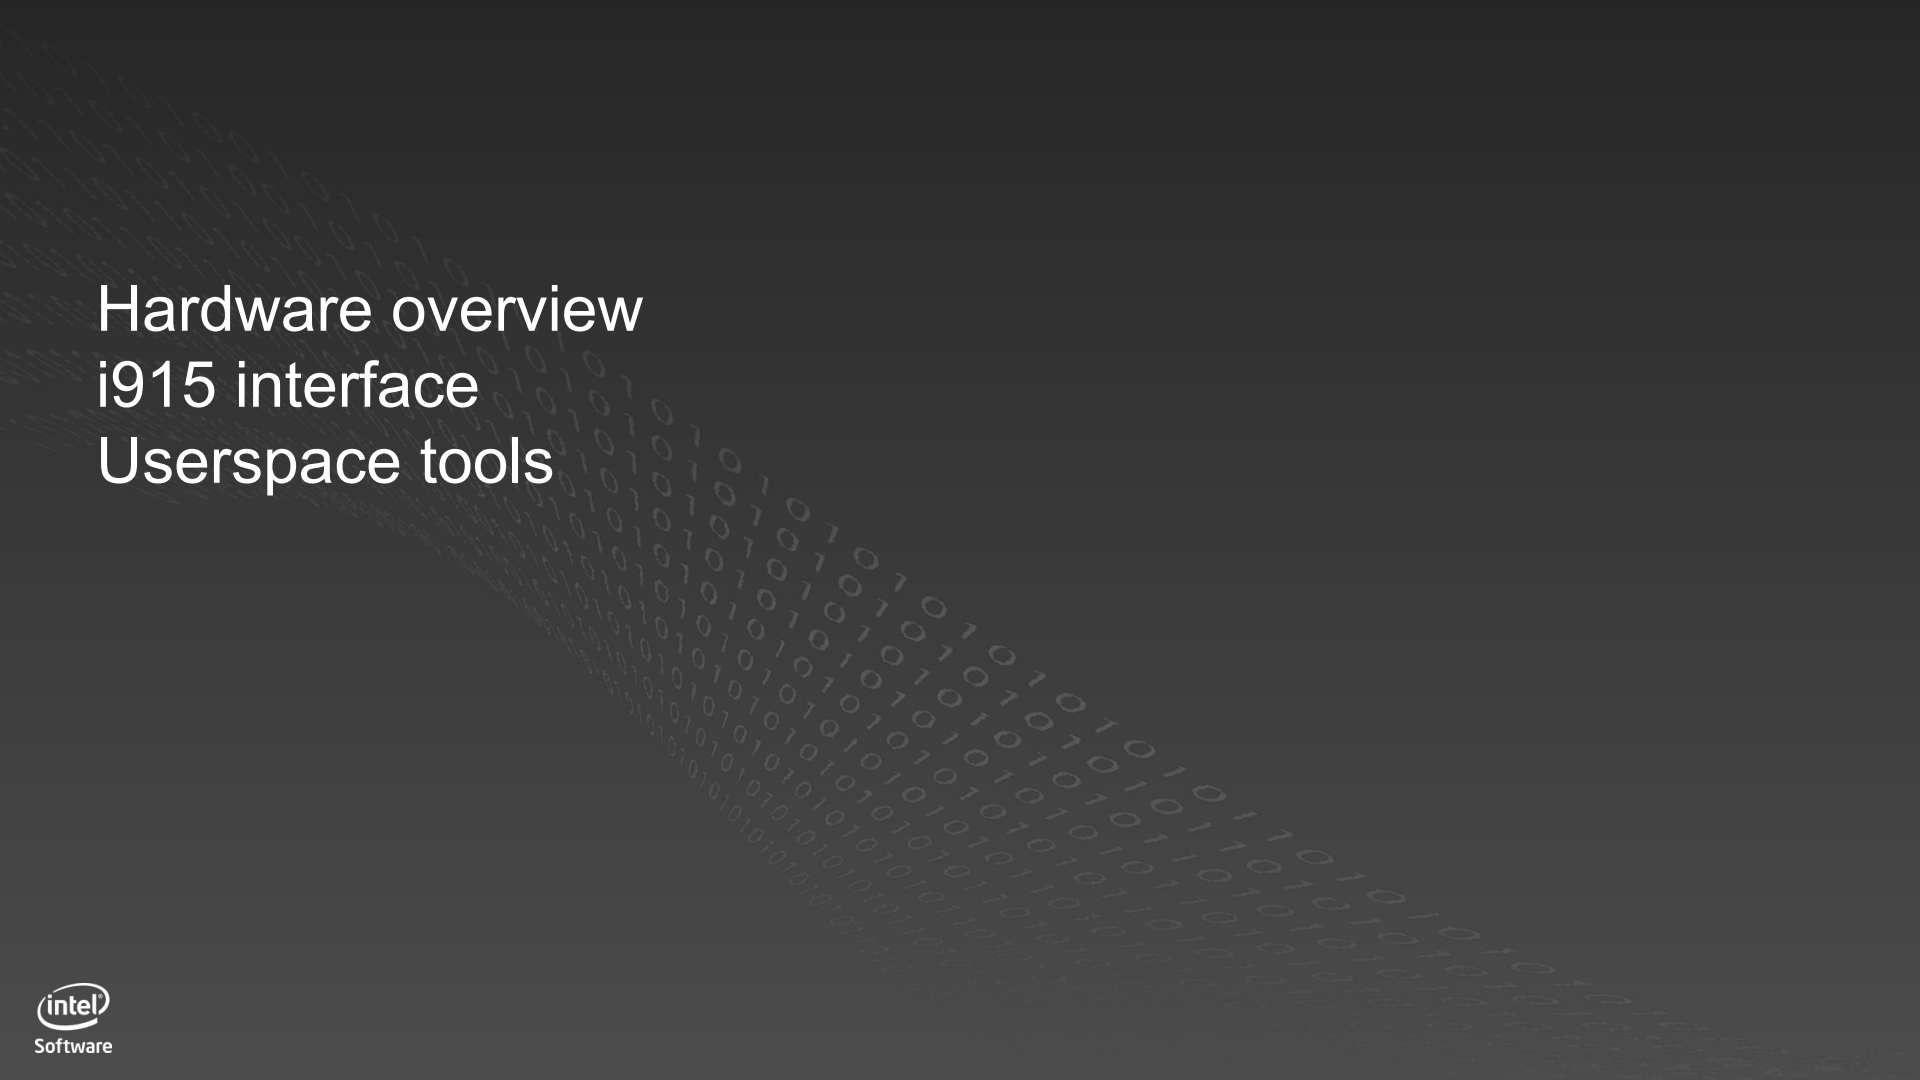

Hardware overview

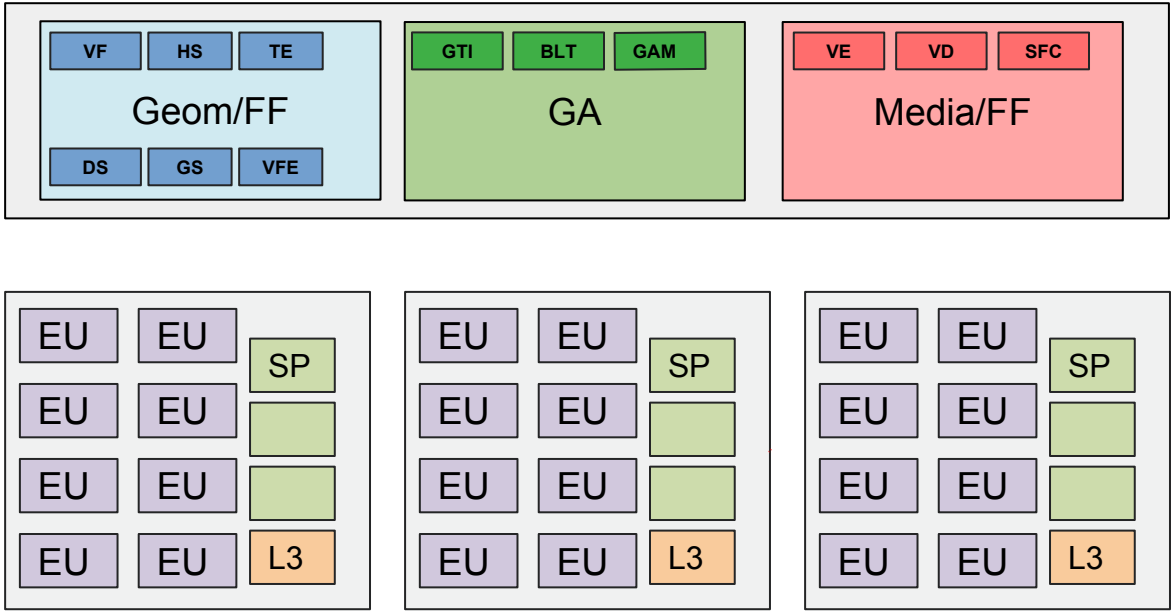

https://01.org/sites/default/files/documentation/intel-gfx-prm-osrc-kbl-vol04-configurations.pdf

Hardware overview

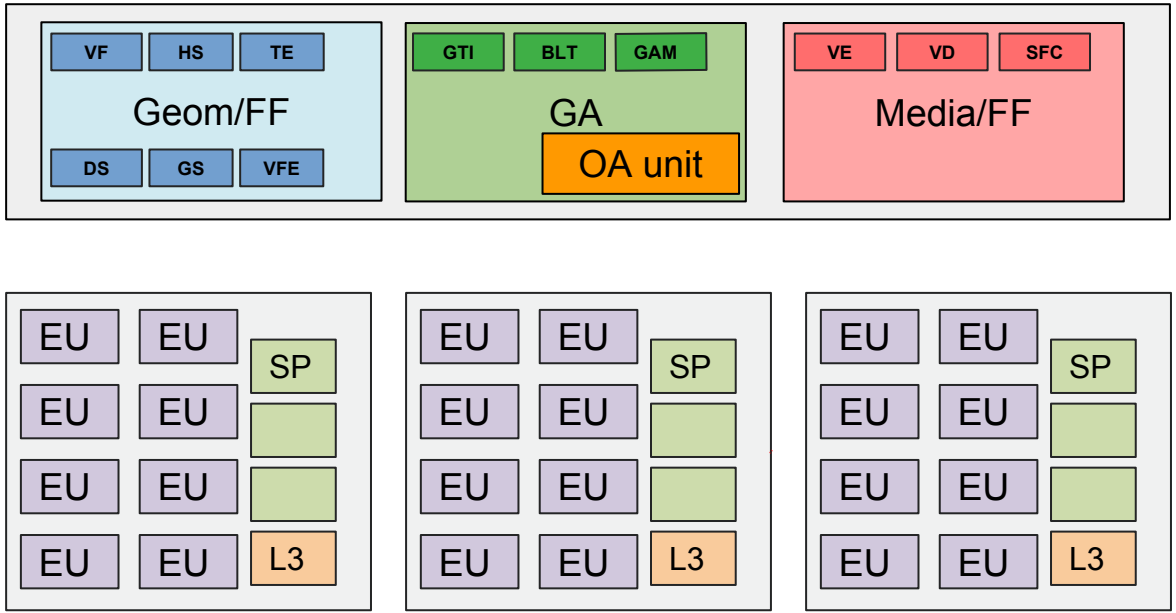

https://01.org/sites/default/files/documentation/intel-gfx-prm-osrc-kbl-vol04-configurations.pdf

Hardware overview

OA unit :

- Writes snapshots of multiple registers to memory on :
  - context switch
  - programmed timer
  - frequency changes
  - request from command streamer (only on 3D engine)

- Snapshots written to :
  - OA buffer (circular buffer up to 16Mb)
  - application address space

https://01.org/sites/default/files/documentation/intel-gfx-prm-osrc-kbl-vol14-observability.pdf

Hardware overview

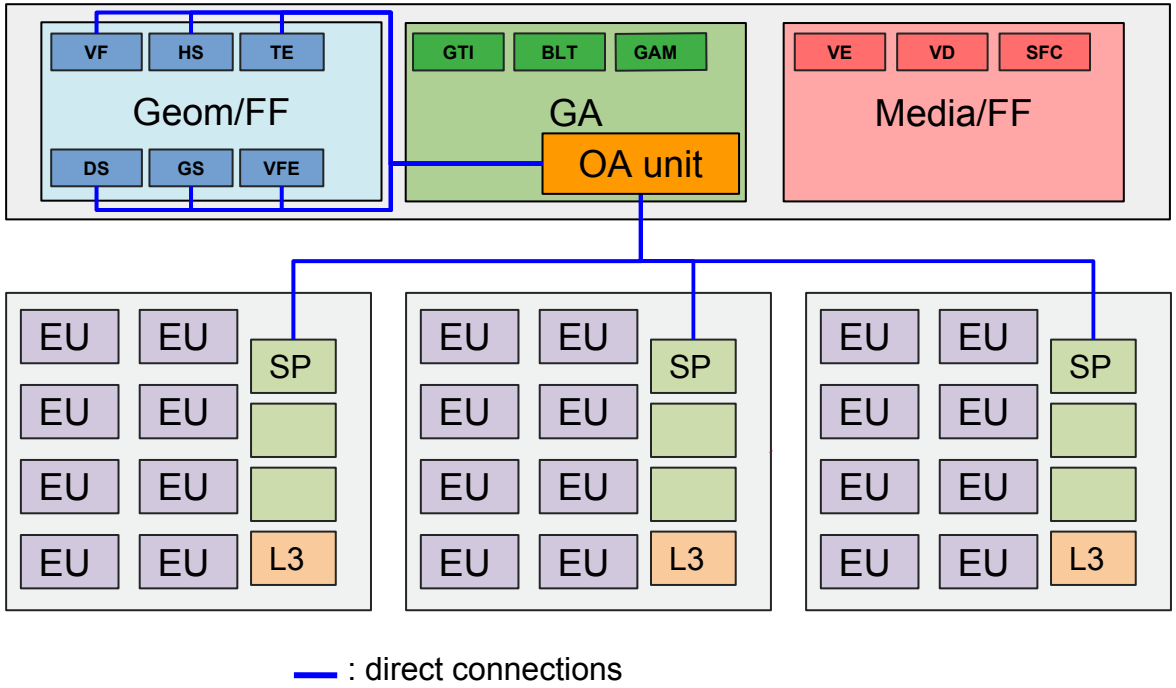

https://01.org/sites/default/files/documentation/intel-gfx-prm-osrc-kbl-vol04-configurations.pdf

Hardware overview

- Direct connections examples :
  - Vertex Shader Threads Dispatched
  - Hull Shader Threads Dispatched
  - Pixel Shader Threads Dispatched
  - 2x2s Rasterized Pixels
  - 2x2s Killed in PS (discard in fragment shader)
  - 2x2s Written To Render Target
  - Blended 2x2s Written to Render Target
  - 2x2s Requested from Sampler
  - Sampler L1 Cache Misses
  - Flexible EU counters
  - …

Mostly 3D counters

https://01.org/sites/default/files/documentation/intel-gfx-prm-osrc-kbl-vol14-observability.pdf

Introduction

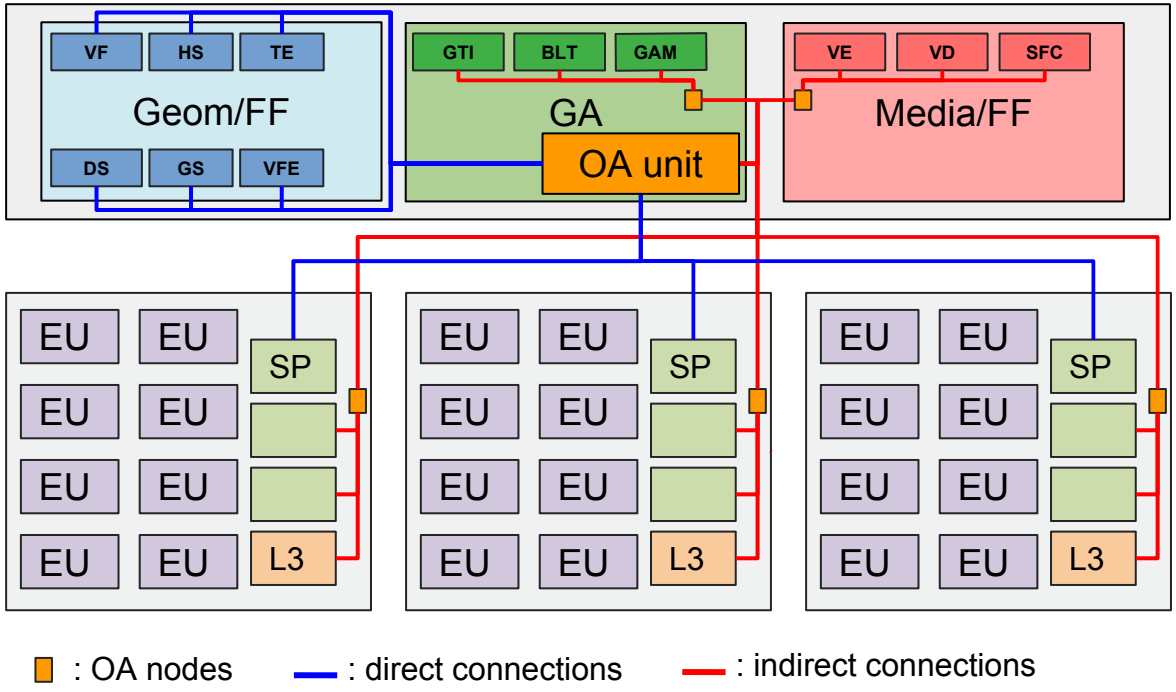

https://01.org/sites/default/files/documentation/intel-gfx-prm-osrc-kbl-vol04-configurations.pdf

Hardware overview

- Indirect connections examples :
  - GTI Depth Throughput
  - Sampler 0/1 Busy
  - L3 Cache Misses
  - Early Depth Bottleneck
  - Hi-Depth Cache Misses
  - Multisampling Color Cache misses
  - Stencil Cache misses
  - …
- HW programming needed to get specific information

https://01.org/sites/default/files/documentation/intel-gfx-prm-osrc-kbl-vol14-observability.pdf

OA reports

256 bytes (Broadwell and above)

- Headers : timestamp + context ID + reason
- A counters : 32 (40 bits) + 4 (32 bits)
  - Mostly 3D counters
- B counters : 8 (32 bits)
- C counters : 8 (32 bits)

https://01.org/sites/default/files/documentation/intel-gfx-prm-osrc-kbl-vol14-observability.pdf

i915 Interface

Exclusive access to the OA unit because of B/C counters programming.

2 ways to use the i915 API :

- Query mode :
  - Have snapshots filtered by context ID
  - Use in addition to the MI_REPORT_PERF_COUNT instruction
- Monitoring mode :
  - All snapshots available (privileged access)

i915 Interface

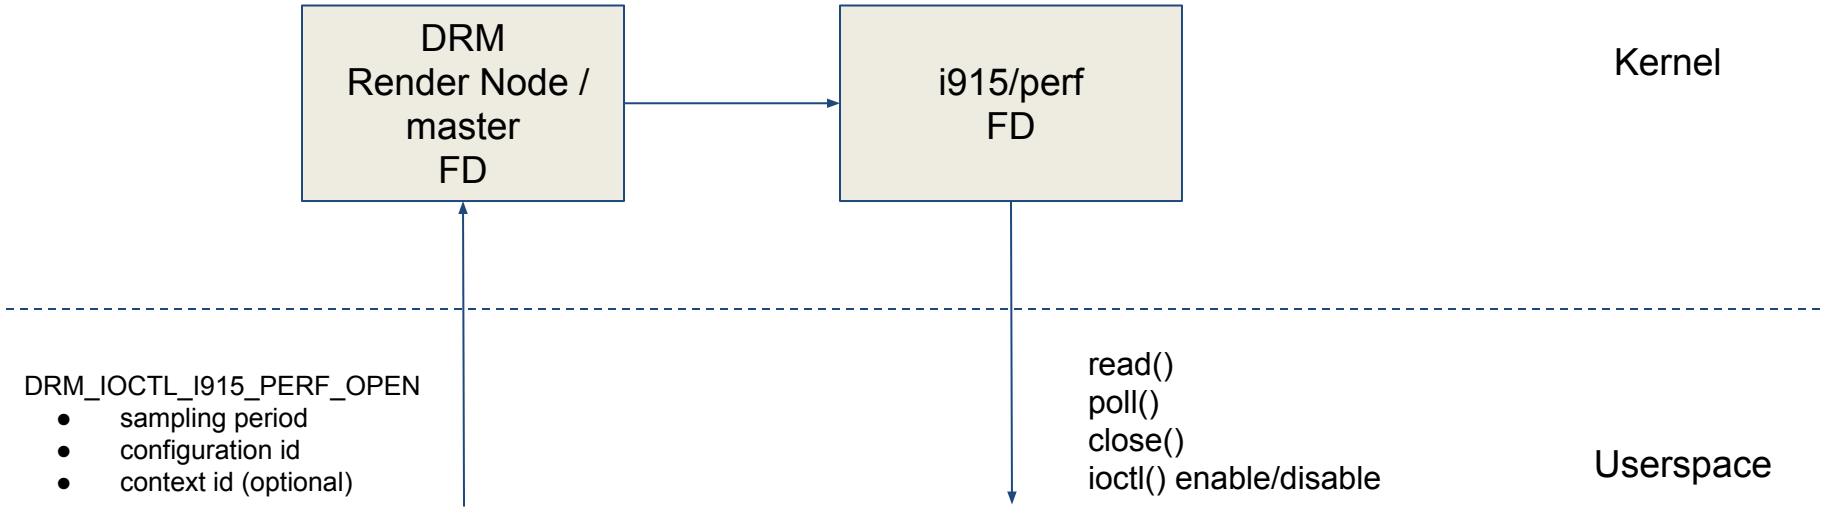

i915 Interface

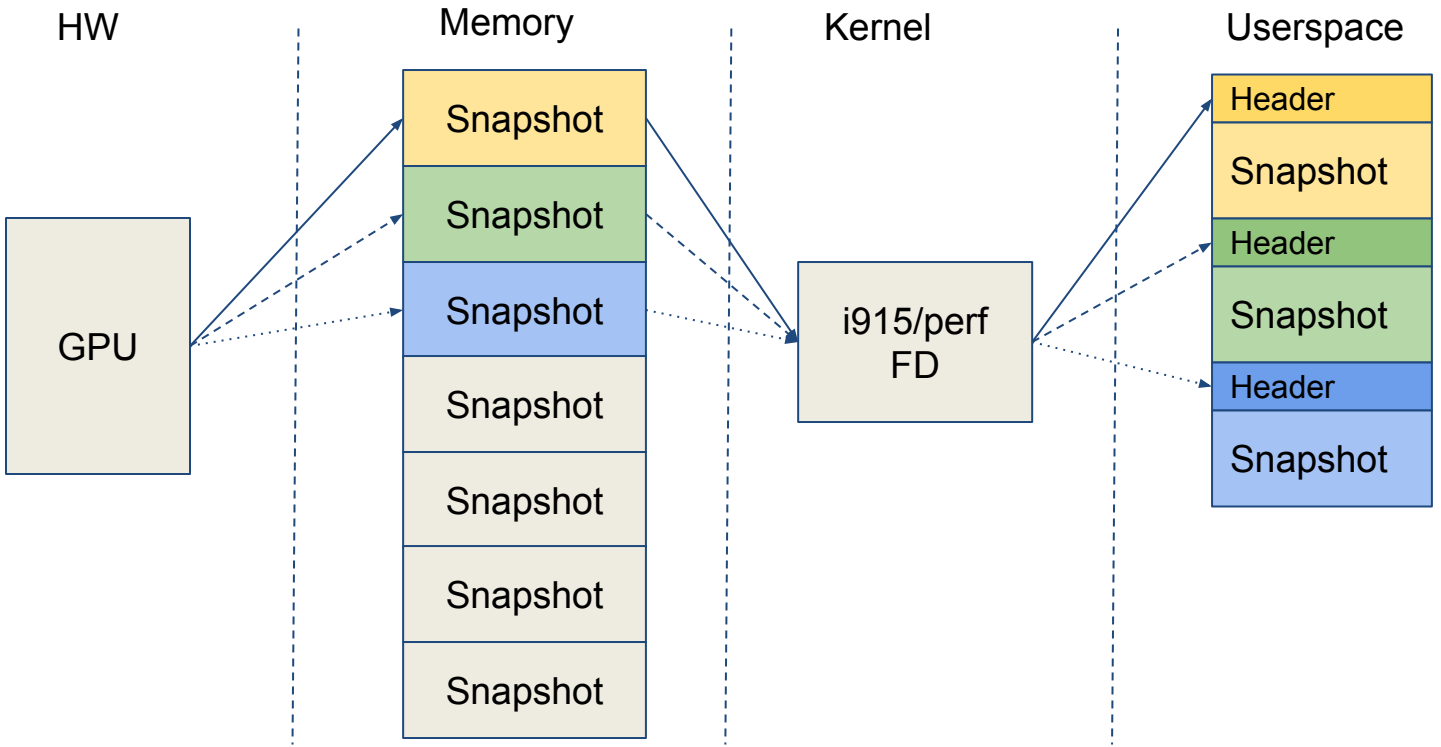

Userspace

- Metrics Discovery (used by Graphics Performance Analyzers / VTUNE)
  - https://github.com/intel/metrics-discovery
- GL_INTEL_performance_query extension
  - https://www.khronos.org/registry/OpenGL/extensions/INTEL/INTEL_performance_query.txt
- GPUTop
  - https://github.com/rib/gputop

OpenGL performance queries

We can’t extract all the performance counters in one pass.

Counters are grouped in query IDs :

- Render Metrics Basic
- Compute Metrics Basic
- Render Metrics for 3D Pipeline Profile
- Memory Reads Distribution
- Memory Writes Distribution
- Compute Metrics Extended
- Compute Metrics L3 Cache
- Metric set HDCAndSF
- Metric set L3_1
- Metric set L3_2
- Metric set L3_3
- Metric set RasterizerAndPixelBackend
- Metric set Sampler
- Metric set TDL_1
- Metric set TDL_2
- Compute Metrics Extra
- Media Vme Pipe
- Gpu Rings Busyness

OpenGL performance queries

GL_INTEL_performance_query :

- List query IDs :
  - `glGetFirstPerfQueryIdINTEL()  / glGetNextPerfQueryIdINTEL()`
- List counters for a given query ID :
  - `glGetPerfCounterInfoINTEL()`
- Query data :
  - `glCreatePerfQueryINTEL() / glBeginPerfQueryINTEL() / glEndPerfQueryINTEL()`
- Get data :
  - `glGetPerfQueryDataINTEL()`

OpenGL performance queries

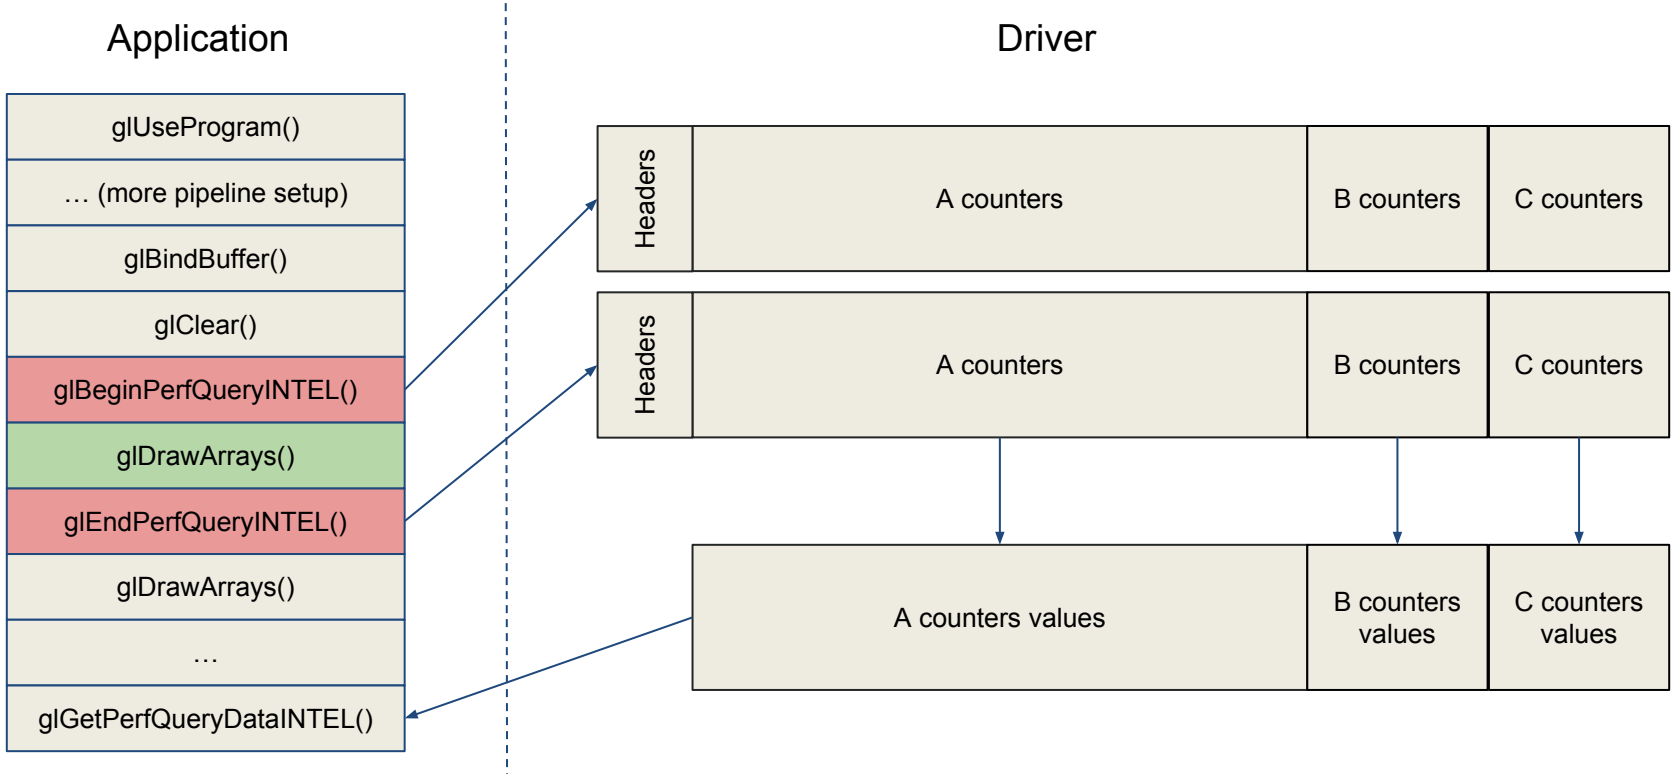

OpenGL performance queries

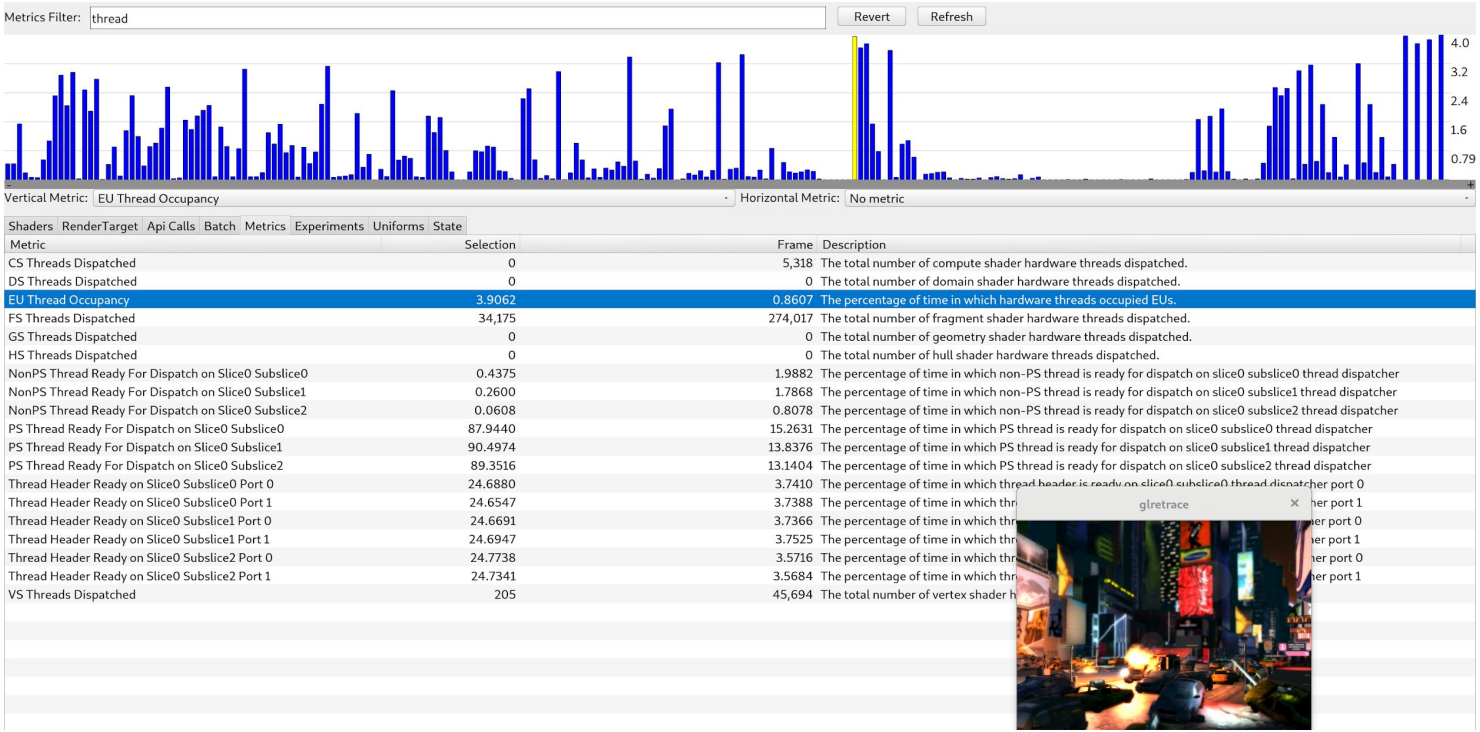

https://github.com/janesma/apitrace

GPUTop

- Client/Server model :
  - Server runs on the target system to monitor
  - Clients connects to the server and process the extracted data

- 2 clients :
  - Command line tool :
    - records accumulated samples in CSV format
    - track an application’s usage
  - User interface :
    - Observe global usage
    - Draw timelines

GPUTop

Server :

```
$ sudo gputop
```

Global monitoring :

```
$ gputop-wrapper -m RenderBasic -c AvgGpuCoreFrequency,RasterizedPixels,Sampler0Busy
```

Application monitoring :

```
$ gputop-wrapper -m RenderBasic -c AvgGpuCoreFrequency,RasterizedPixels,Sampler0Busy -- glxgears
```

Output :

```
 AvgGpuCoreFrequency,  RasterizedPixels, Sampler0Busy
                (Hz),          (pixels),          (%)
           295.3 MHz,    145.6 M pixels,       6.44 %
           295.6 MHz,    119.5 M pixels,       4.84 %
           295.8 MHz,    169.4 M pixels,       7.02 %
           295.6 MHz,    97.31 M pixels,       3.97 %
           295.6 MHz,    120.1 M pixels,       4.87 %
```

GPUTop

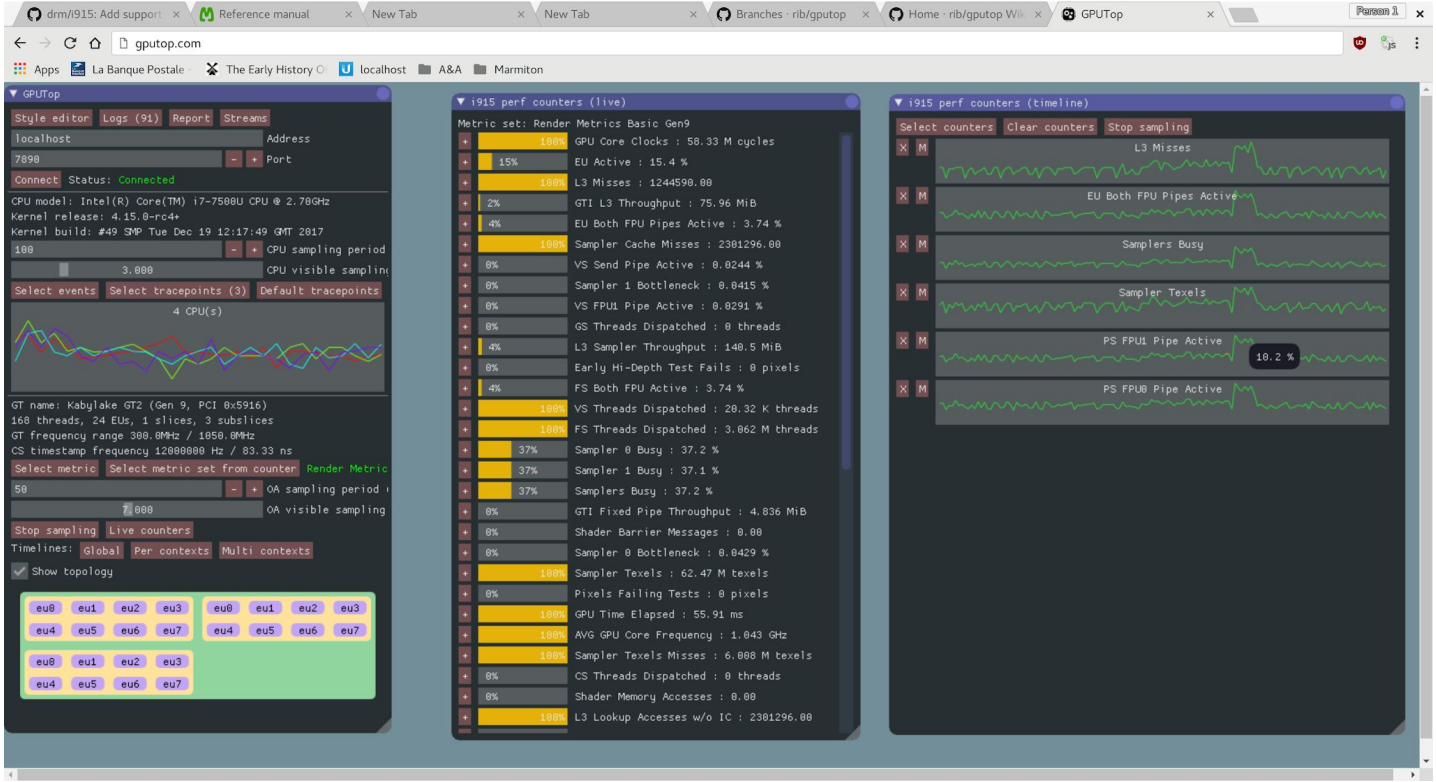
drm/i915: Add support
Reference manual
New Tab
New Tab
Branches · rib/gputop
Home · rib/gputop Wiki
GPUTop
Person 1
gputop.com
Apps
La Banque Postale
The Early History O
localhost
A&A
Marmiton
GPUTop
Style editor
Logs (91)
Report
Streams
localhost
Address
7890
Port
Connect
Status: Connected
CPU model: Intel(R) Core(TM) i7-7500U CPU @ 2.70GHz
Kernel release: 4.15.0-rc4+
Kernel build: #49 SMP Tue Dec 19 12:17:49 GMT 2017
100
CPU sampling period
3.000
CPU visible sampling
Select events
Select tracepoints (3)
Default tracepoints
4 CPU(s)
GT name: Kabylake GT2 (Gen 9, PCI 0x5916)
168 threads, 24 EUs, 1 slices, 3 subslices
GT frequency range 300.0MHz / 1050.0MHz
CS timestamp frequency 12000000 Hz / 83.33 ns
Select metric
Select metric set from counter
Render Metric
50
OA sampling period
7.000
OA visible sampling
Stop sampling
Live counters
Timelines:
Global
Per contexts
Multi contexts
Show topology
eu0 eu1 eu2 eu3 eu4 eu5 eu6 eu7
eu0 eu1 eu2 eu3 eu4 eu5 eu6 eu7
eu0 eu1 eu2 eu3 eu4 eu5 eu6 eu7
i915 perf counters (live)
Metric set: Render Metrics Basic Gen9
100% GPU Core Clocks : 58.33 M cycles
15% EU Active : 15.4 %
100% L3 Misses : 1244590.00
2% GTI L3 Throughput : 75.96 MiB
4% EU Both FPU Pipes Active : 3.74 %
100% Sampler Cache Misses : 2301296.00
0% VS Send Pipe Active : 0.0244 %
0% Sampler 1 Bottleneck : 0.0415 %
0% VS FPU1 Pipe Active : 0.0291 %
0% GS Threads Dispatched : 0 threads
4% L3 Sampler Throughput : 140.5 MiB
0% Early Hi-Depth Test Fails : 0 pixels
4% FS Both FPU Active : 3.74 %
100% VS Threads Dispatched : 20.32 K threads
100% FS Threads Dispatched : 3.062 M threads
37% Sampler 0 Busy : 37.2 %
37% Sampler 1 Busy : 37.1 %
37% Samplers Busy : 37.2 %
0% GTI Fixed Pipe Throughput : 4.836 MiB
0% Shader Barrier Messages : 0.00
0% Sampler 0 Bottleneck : 0.0429 %
100% Sampler Texels : 62.47 M texels
0% Pixels Failing Tests : 0 pixels
100% GPU Time Elapsed : 55.91 ms
100% AVG GPU Core Frequency : 1.043 GHz
100% Sampler Texels Misses : 6.008 M texels
0% CS Threads Dispatched : 0 threads
0% Shader Memory Accesses : 0.00
100% L3 Lookup Accesses w/o IC : 2301296.00
i915 perf counters (timeline)
Select counters
Clear counters
Stop sampling
L3 Misses
EU Both FPU Pipes Active
Samplers Busy
Sampler Texels
PS FPU1 Pipe Active
10.2 %
PS FPU0 Pipe Active

GPUTop - timelines

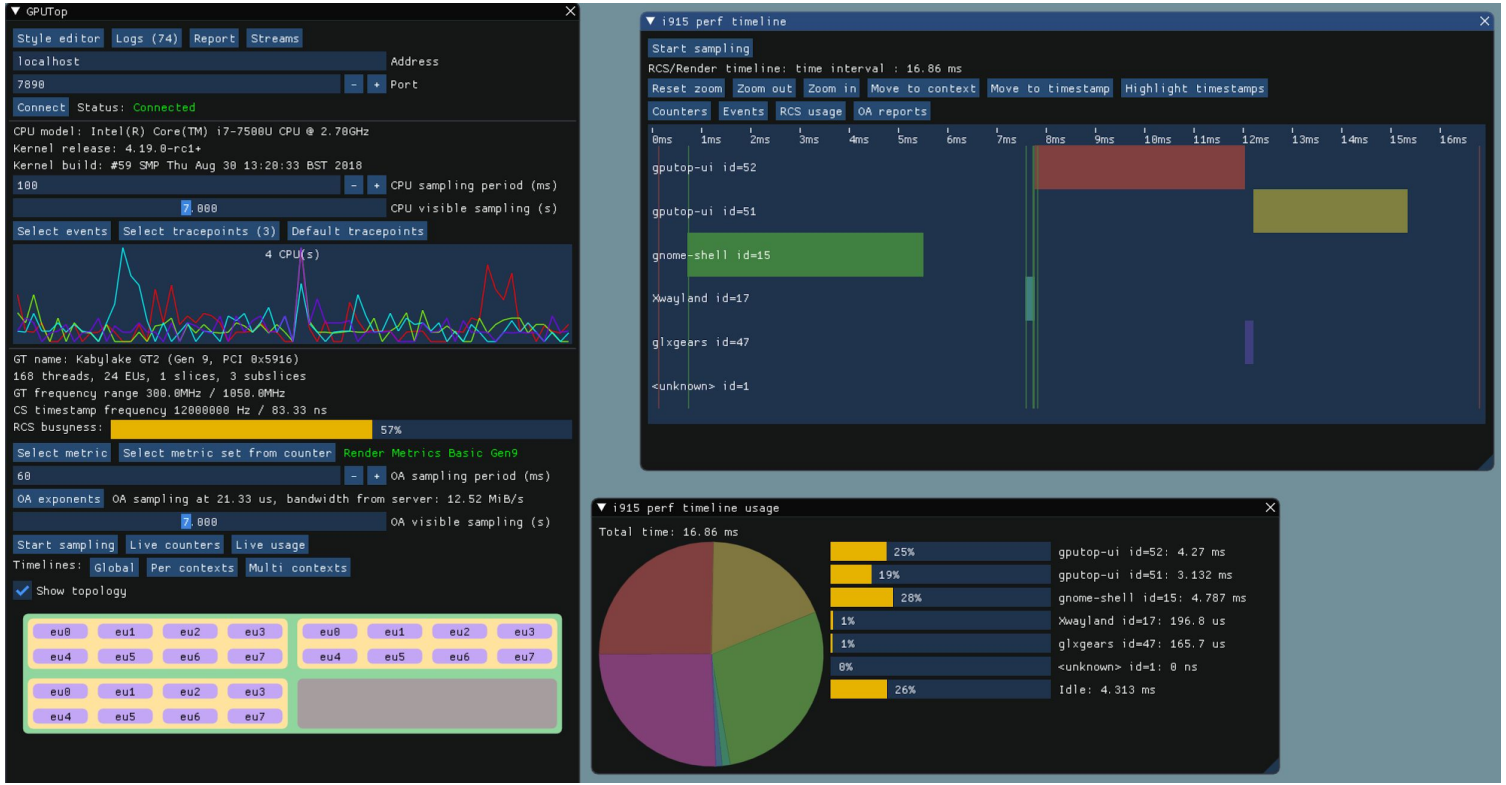

GPUTop - high frequency sampling

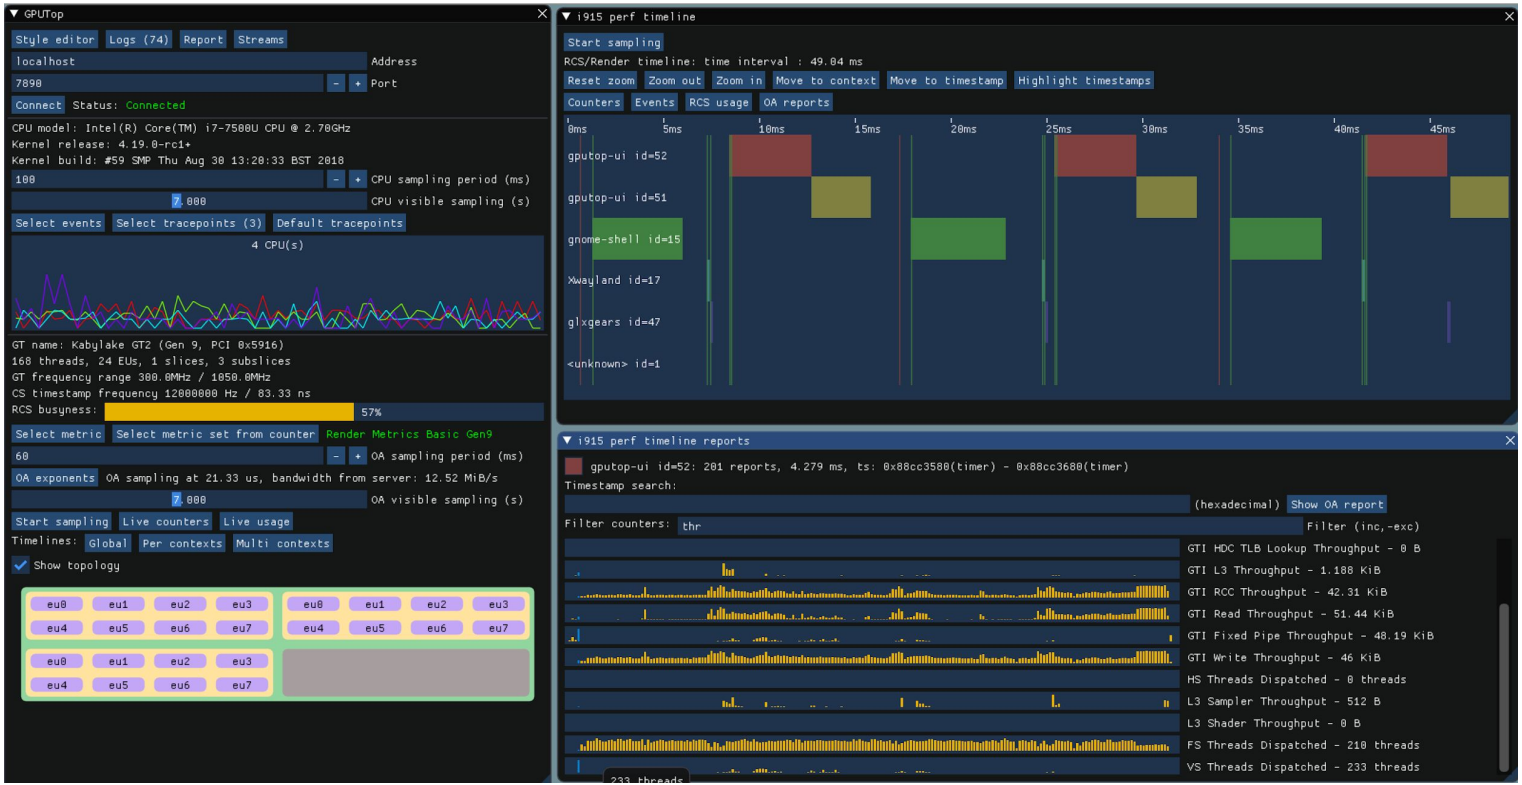

Give performance queries a try :

https://github.com/janesma/apitrace

Give GPUTop a try (kernel 4.14 recommended) :

https://github.com/rib/gputop

http://gputop.com

Questions?