SixTrack as Benchmark WL

Olga Datskova, Riccardo de Maria (presented on 31/01/20)

WIP: SimpleTrack Docker Images

Testing images constructed for nvidia/intel OpenCL runtimes: repo

```
FROM nvidia/opencl:runtime-centos7

ENV PARTICLES 20000
ENV TURNS 15
ENV DEVICE "0.0"

RUN yum install -y python3 python3-devel python3-pip git gcc && pip3 install pyopencl \
    && git clone https://github.com/rdemaria/simpletrack \
    && cd simpletrack && pip3 install -e .

ENTRYPOINT [ "/bin/bash", "-c", "cd /simpletrack/examples/lhc &&
                  python3 benchmark_opencl.py -p $PARTICLES -t $TURNS -d $DEVICE" ]
```

~$ docker run **–gpus all** -e PARTICLES=1000 -e TURNS=10

gitlab-registry.cern.ch/odatskov/simpletrack-docker:nvidia-master

~$ docker run **--device /dev/dri:/dev/dri**

gitlab-registry.cern.ch/odatskov/simpletrack-docker:intel-master

Tested with Docker version 19.03.5

OpenCL Memory Model

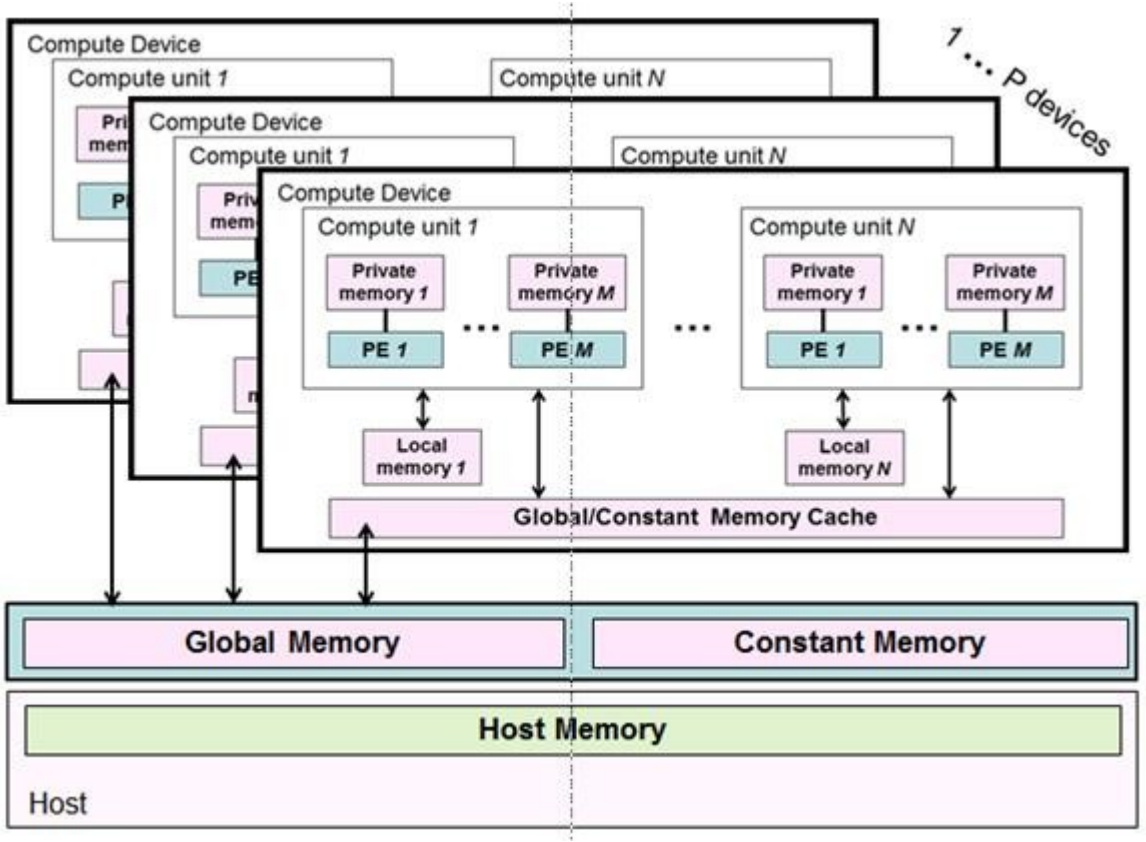

"The named address spaces exposed in an OpenCL Platform" (source)

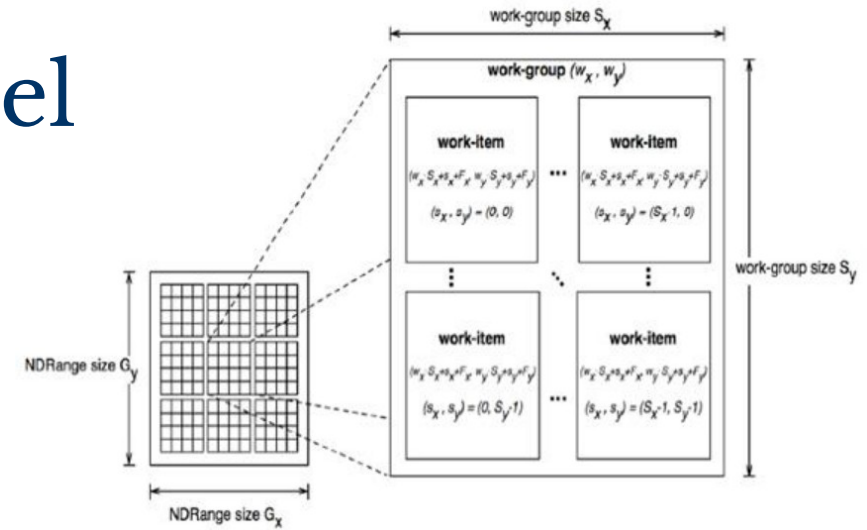

Top: Thread blocks (CUDA)/Workgroups (OpenCL)
Bottom: Execution states and transitions

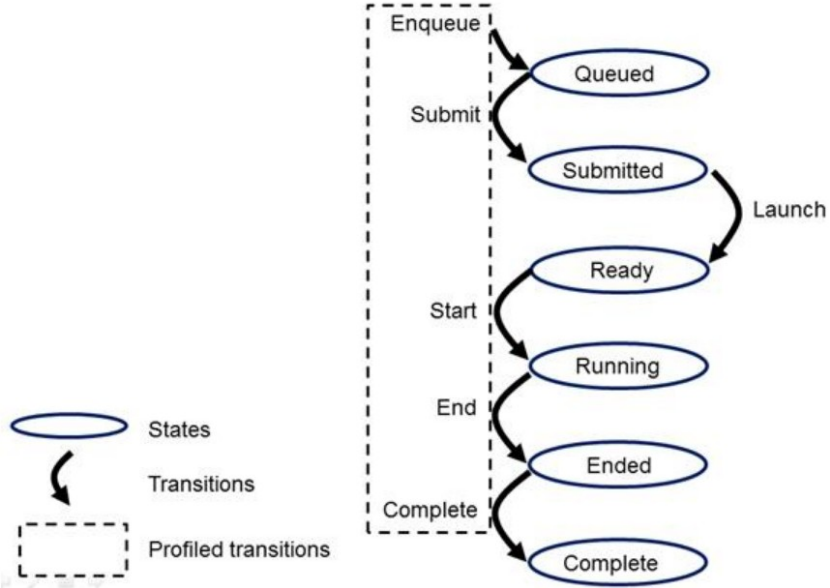

NVidia Tesla V100

Image source: Volta architecture whitepaper (source)

Figure 4. Volta GV100 Full GPU with 84 SM Units

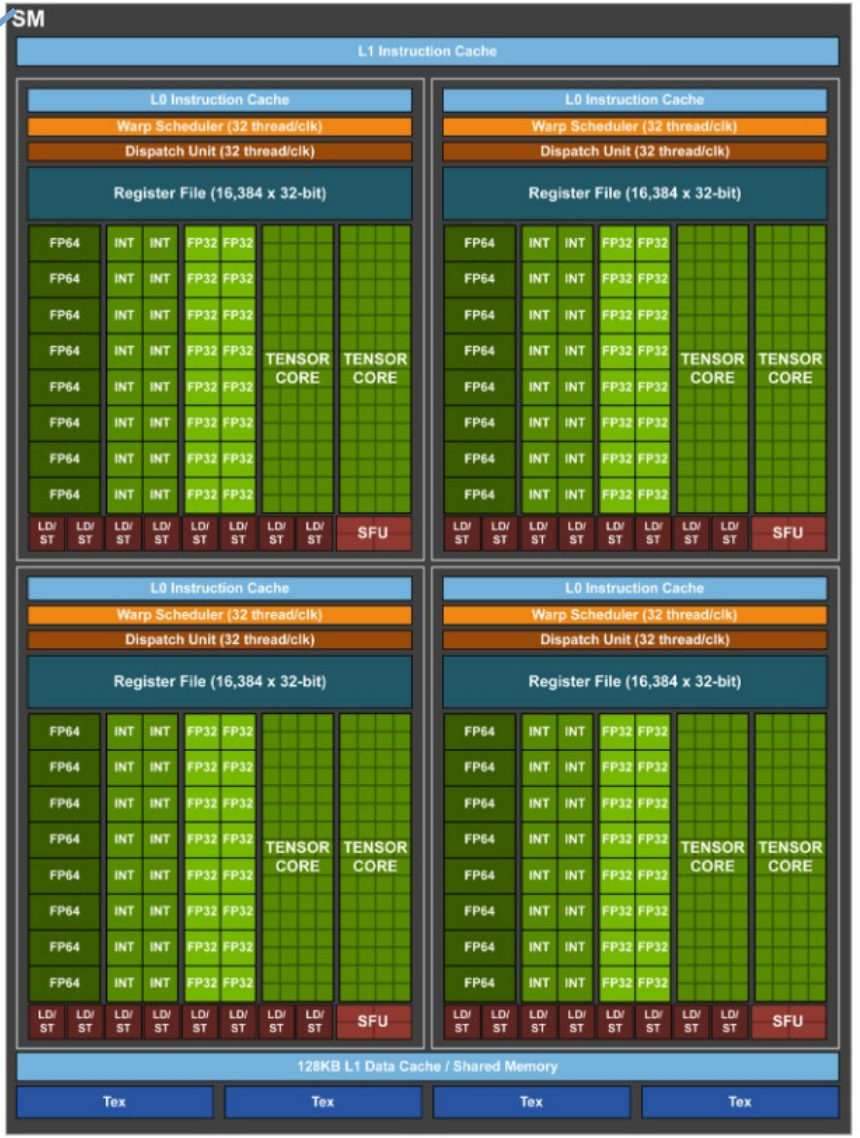

Figure 5. Volta GV100 Streaming Multiprocessor (SM)

OpenCL Event Profiling

- Current SimpleTrack benchmark output:

  29230 **particles*turns/seconds**, 20000 particles, 15 turns, device '0.0', Intel(R) ...

  (number of particles * number of turns) / (host wall time = track + transfer)

- Profiling on kernel and buffer transfer events:

```
self.queue = pyopencl.CommandQueue(self.ctx,
                properties=pyopencl.command_queue_properties.PROFILING_ENABLE))
…

event = self.program.track (…)
print(f"Track: {(event.profile.end - event.profile.start)*1e-9} s")
…

event = pyopencl.enqueue_copy(…)
event.wait()
print(f"Transfer: {(event.profile.end - event.profile.start)*1e-9} s")
```

Measurements

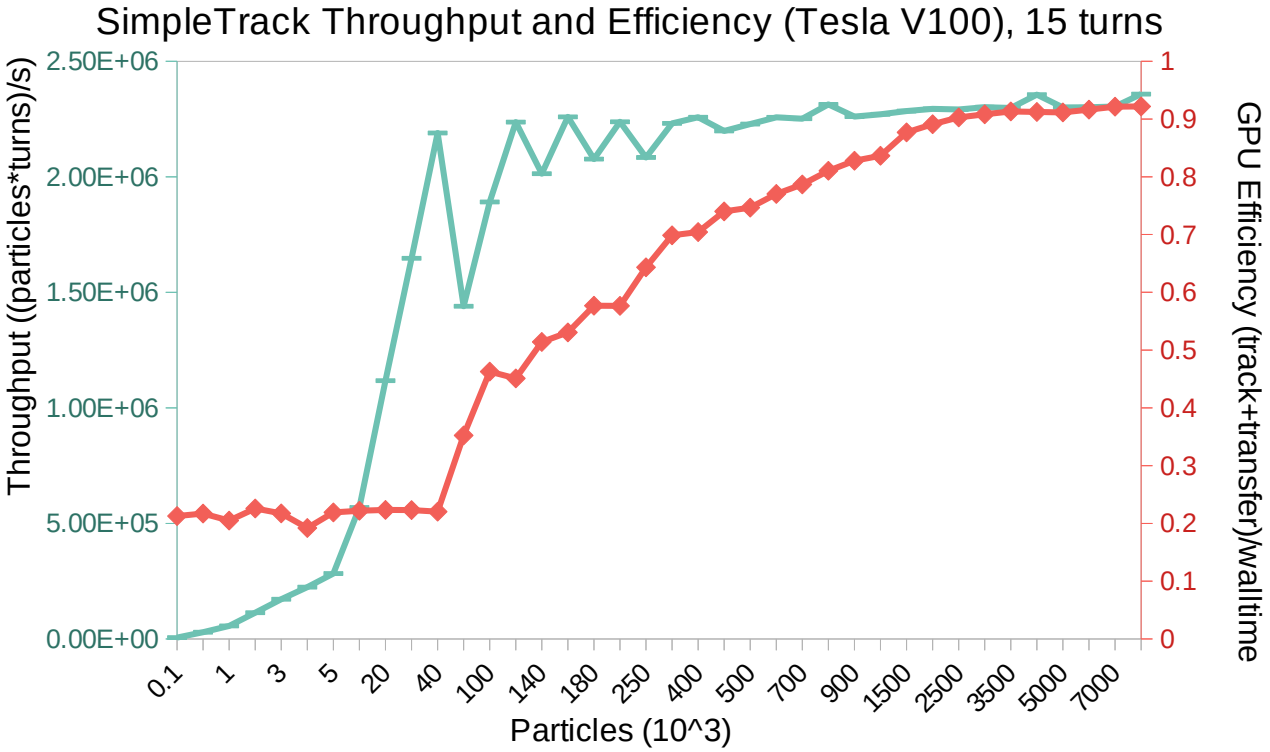

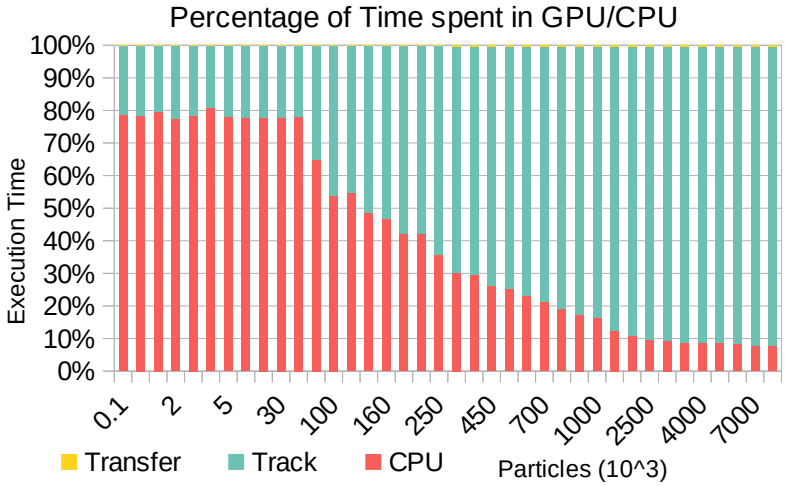

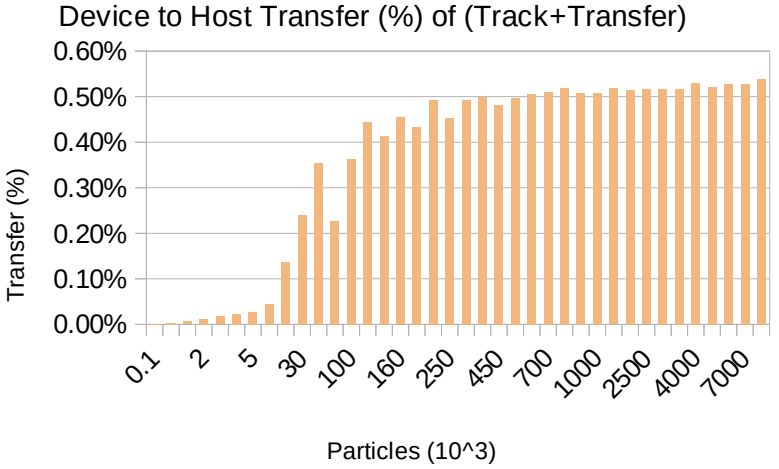

Docker image used: gitlab-registy.cern.ch/odatskov/simpletrack-docker:nvidia-perf