



Heterogeneous workload support in HEPscore

David Southwick – IT-DI-OPL

05/11/2020

First look at workload orchestration beyond x86

First attempt made using 'Simpletrack' [1] – a minimal OpenCL python module

- Most work already complete by @odatskov over the last year
- Dockerfiles prepared for various OpenCL compatible hardware (ROCM, Intel, nvidia, pocl) – easily extendable to any HW that supports OpenCL!
- Minimal changes to Simpletrack *summary.json* required to satisfy norms established by CPU workloads, docker image rename
- Open questions remaining on JSON format (metadata, hashes, etc) & calling arguments...

[1] github.com/rdemaria/simpletrack

HEPscore considerations

Things encountered so far:

- **Need a way to modify arguments to containment service (podman/docker/singularity)**
 - Different binds for each service `--nv -B /etc/OpenCL | --gpus all`
 - Switches to use host drivers? Or do we package everything?
 - Unclear if workload images can check correct launch conditions (themselves)
- **(DONE) - Custom registry definition *per workload*: Simpletrack lives outside of main workloads project (for now)**

Overall, HEPscore (v1.0rc9) very flexible - already supports nearly all features required!

Running Example

Using BMK-577 branch & default singularity execution:

```
$ hep-score -f hepscore-gpu.yaml out
HEPscore20 Benchmark
Config Hash:
1afd157b2a1841317b01b75123050013b10d7075ff2dc03d648763a3dbcfab92
System:      Linux bmk-gpu-01.cern.ch 3.10.0-1127.19.1.el7.x86_64 #1 SMP Tue Aug
25 17:23:54 UTC 2020 x86_64
Container Execution: singularity
Registry:    docker://gitlab-registry.cern.ch/hep-benchmarks/hep-workloads
Output:     out/HEPscore_05Nov2020_164917
Date:       Thu Nov 5 16:49:17 2020
```

```
2020-11-05 16:49:17,239 - INFO - Executing 3 runs of lhc-simpletrack-bmk
2020-11-05 16:49:17,239 - INFO - Overriding registry for this container: docker://gitlab-
registry.cern.ch/hep-benchmarks/hep-workloads-gpu/simpletrack
2020-11-05 16:49:17,246 - INFO - Starting run0
2020-11-05 16:50:06,347 - INFO - Starting run1
2020-11-05 16:50:57,749 - INFO - Starting run2

2020-11-05 16:51:46,645 - INFO - Final result: 138114.0
```

Simpletrack output

```
$ singularity run --nv -C -B /tmp/out:/results docker://gitlab-
registry.cern.ch/hep-benchmarks/hep-workloads-gp
u/simpletrack/lhc-simpletrack-bmk:qa-simpletrack
```

```
Base working directory : /results/lhc-simpletrack-p20000-t15-
1604585686_4604
```

Running benchmark for 0.0 device

```
{
  "copies": 1,
  "threads_per_copy": 15,
  "events_per_thread": 20000,
  "wl-scores": {
    "simpletrack": 128699
  },
  "app": {
    "device_id": "0.0",
    "name": "Tesla T4",
    "platform": "NVIDIA CUDA",
    "containment": "singularity"
  }
}
```

Finished running simpletrack

Closing thoughts

Simpletrack from the perspective of HEP is only ~50 lines of bash.

- Possible to simply *python import* Simpletrack directly & reuse already existing functions (as hosts must already have proper drivers installed & enabled...)


Heterogeneous images demand more knowledge from end-user

- What hardware are you running on & is it being properly exposed to the containers
- Opportunity for “sanity checks” inside the workloads, or perhaps based on system scan in HEP-bmk-suite

HEPscore expectations unclear for workload summary JSON – Structure, keys, metadata

- Currently with maximum debug flags users only get “[Error: score not reported for one or more runs](#)”
- Opportunity here to provide feedback of what/where/how went wrong, and tips to fix
- CI test or documentation on JSON structure & standards

Progress tracked under [JIRA epic BMK-575](#)



THANK YOU

home.cern

HEPscore resulting JSON

```
{
  "benchmarks": {
    "lhc-simpletrack-bmk": {
      "version": "qa-simpletrack",
      "registry": "docker://gitlab-registry.cern.ch/hep-benchmarks/hep-workloads-gpu/simpletrack",
      "args": {
        "-p": 30000,
        "-t": 200,
        "-b": "all"
      },
      "run0": {
        "start_at": "Thu Nov 5 16:49:17 2020",
        "end_at": "Thu Nov 5 16:50:06 2020",
        "duration": 49,
        "report": {
          "wl-scores": {
            "simpletrack": 138001
          }
        }
      },
      "run1": {
        "start_at": "Thu Nov 5 16:50:06 2020",
        "end_at": "Thu Nov 5 16:50:57 2020",
        "duration": 51,
        "report": {
          "wl-scores": {
            "simpletrack": 138411
          }
        }
      },
      "run2": {
        "start_at": "Thu Nov 5 16:50:57 2020",
        "end_at": "Thu Nov 5 16:51:46 2020",
        "duration": 49,
        "report": {
          "wl-scores": {
            "simpletrack": 138114
          }
        }
      }
    }
  },
}
```

```
"app": {
  "device_id": "0.0",
  "name": "Tesla T4",
  "platform": "NVIDIA CUDA",
  "containment": "singularity"
},
"run_info": {
  "copies": 1,
  "threads_per_copy": 200,
  "events_per_thread": 30000
}
},
"app_info": {
  "name": "HEPscore20",
  "reference_machine": "CPU Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz",
  "registry": "docker://gitlab-registry.cern.ch/hep-benchmarks/hep-workloads",
  "hash": "1afd157b2a1841317b01b75123050013b10d7075ff2dc03d648763a3dbcfab92",
  "hepscore_ver": "1.0.0.0rc9.dev11"
},
"settings": {
  "method": "geometric_mean",
  "repetitions": 3,
  "scaling": 1,
  "container_exec": "singularity",
  "replay": false
},
"environment": {
  "system": "Linux bmk-gpu-01.cern.ch 3.10.0-1127.19.1.el7.x86_64 #1 SMP Tue Aug 25 17:23:54 UTC 2020 x86_64",
  "date": "Thu Nov 5 16:49:17 2020",
  "singularity_version": "singularity version 3.6.4-1.el7"
},
"wl-scores": {
  "lhc-simpletrack-bmk": {
    "simpletrack": 138114,
    "simpletrack_ref": 1
  }
},
"score": 138114,
"status": "success",
"score_per_core": 34528.5
}
```

Singularity definition

```
Bootstrap: docker
From: nvidia/cuda:10.1-base-centos7

%labels
  Author "David Southwick <David.Southwick@cern.ch>"
  Version 0.1
  Description "simpletrack test image"

%files
  lhc-simpletrack.sh /opt

%post
  DEPS="python3-devel python3-pip git clinfo gcc jq"

  yum install -y -q epel-release
  yum install -y $DEPS
  yum -y clean all --enablerepo='*'
  python3 -m pip install --upgrade pip pyopenc1
  git clone https://github.com/rdemaria/simpletrack.git /opt/simpletrack
  python3 -m pip install -e /opt/simpletrack
  chmod +x /opt/lhc-simpletrack.sh
  mv /opt/lhc-simpletrack.sh /opt/simpletrack/examples/lhc/lhc-simpletrack.sh

%runscript
  cd /opt/simpletrack/examples/lhc/
  exec ./lhc-simpletrack.sh "$@"
```