



Towards a HEP-workload GEN GPU benchmark: MadGraph5_aMC@NLO

Andrea Valassi (CERN IT-SC-RD)



HEPiX Benchmarking WG meeting, 5th November 2020 - <https://indico.cern.ch/event/946409>



Foreword: GEN workloads and GPUs in WLCG

- MC event generators are essential for HEP physics analyses
 - First step in the GEN-SIM-DIGI-RECO chain for simulated data
 - MadGraph5_aMC@NLO (MGaMC) is one of the main generators at the LHC
 - Growing interest (and concerns) about the software aspects of generators
 - *GEN workloads are a sizeable (and growing) fraction of WLCG CPU budgets*
 - Currently ~12% for ATLAS (mainly Sherpa) and ~5% for CMS (mainly MGaMC)
 - Paper for the HL-LHC LHCC review: <https://arxiv.org/abs/2004.13687>
 - See also the recent LHCC talk: <https://doi.org/10.5281/zenodo.4028834>
 - Code modernization and speedup, faster algorithms, port to new architectures...
 - WLCG does not provide/account GPUs yet, but it will one day
 - Example: HPCs (and the experiments already have access to these resources)
 - Only few workloads could make sizable use of GPUs today (e.g. CMS patatrack)
 - Porting generators to GPUs would open up new ways of using GPUs in WLCG
 - *And generators are natural candidates for exploiting parallelism on GPUs (see later...)*
- ⇒ *Generators are a natural fit for preparing a HEP-workload GPU benchmark*
- Bonus: in principle one can run exactly the same algorithm on CPUs and GPUs






MGaMC on GPU – overview

- Why choose MadGraph and not another MC generator for a GPU port?
 - Earlier efforts at KEK in 2010-2013, which were never released for production
 - Touched both matrix element (ME) and MC integration – see summary at HOW2019
 - We are ~not leveraging on this work (based on old version of MadGraph's ME library)
 - Main reason: active involvement of core MadGraph developer (Olivier Mattelaer)
- Project is maintained on <https://github.com/madgraph5/madgraph4gpu>
 - Overall coordination: Stefan Roiser
 - Composite effort: *physicists & engineers* (CERN, Louvain, Argonne, Bangalore)
 - Core CUDA development: AV, OM, SR, Taran Singhania
 - Abstraction layers, profiling, MC integration: David Smith, Laurence Field, Smita Darmora, Taylor Childers, Tyler Burch, Walter Hopkins
 - Just joined: Josh McFayden, Stephan Hageboek
 - Meetings every two weeks
 - Plus more frequent discussions of core CUDA developers
 - The collaboration is largely a spin-off of the activities in the HSF generator WG
 - <https://hepsoftwarefoundation.org/workinggroups/generators.html>

The core of MGaMC is a code-generating code

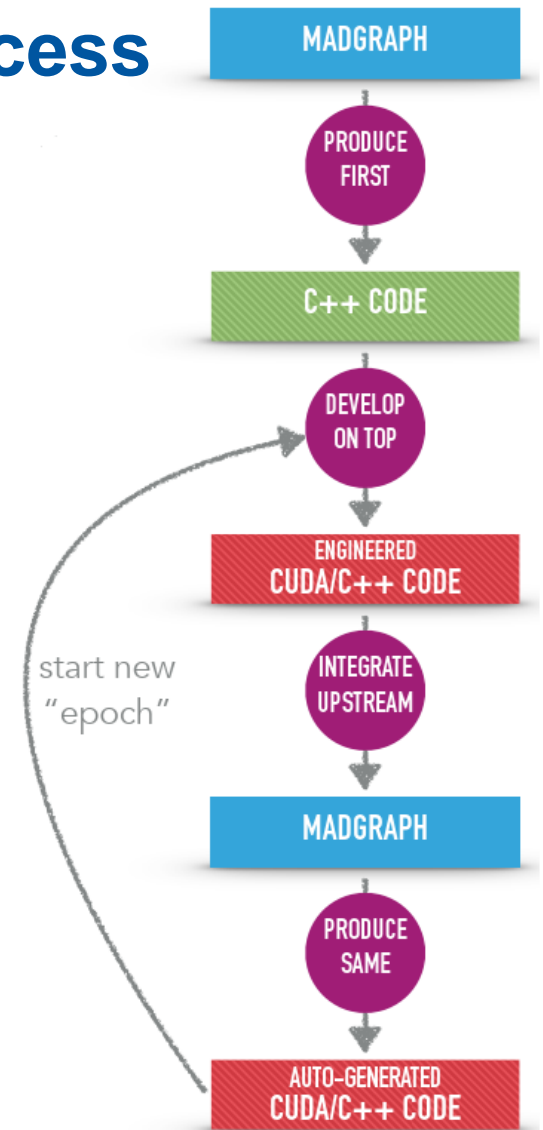
- The heart of a MC generator is the software that calculates the “*matrix element*” (*ME*) for a given process – an element of the “scattering matrix”
 - This is the sum of many contributions, one for each relevant Feynman diagram
 - Simple example: LEP process, e^+e^- to $\mu^+\mu^-$, only two diagrams



- LHC processes (e.g. gg to $t\bar{t}gg$) are much more complex, many more diagrams
 - For every different physics process, MadGraph does two things
 - It identifies which are the relevant Feynman diagrams
 - It generates the code to compute the ME from those Feynman diagrams
- \Rightarrow The core of MadGraph is code-generating code (ALPHA, written in Python)
 - The generated code can be *Fortran* (default!), C++, Python... or *CUDA* (new!)- There are also other (hardcoded) components, the same for all processes

MGaMC on GPU – development process

- MadGraph being a code-generating code complicates the development process
 - We eventually must deliver CUDA-generating code
 - We start from the C++-generating code
- We start from a simple process, $e+e^-$ to $\mu+\mu^-$
 - Few diagrams = Few lines of code to manually port to CUDA, and manually optimize...
 - Then the code-generating code must be adapted
 - And we start another iteration
 - Potentially on a complex process with more diagrams
- Eventually: not only CUDA/Nvidia, also Intel, AMD?
 - Either natively or using abstraction layers: oneAPI, SYCL, HIP/ROCm, Alpaka, Kokkos, OpenCL...
 - Code-generating code would be needed for these too



S. Roiser, [madgraph4gpu meeting](#), Oct 2020

MC generators: simplified computational anatomy

Pseudo-random numbers
Uniform distribution in $[0,1]$
One event i : vector \vec{r}_i (dimension d)
Draw $d \times N_{wgt}$ numbers r (N_{wgt} weighted events)

Phase space sampling
For each event i , map \vec{r}_i to physical phase space $\vec{x}_i = H(\vec{r}_i)$
The resulting \vec{x}_i are distributed according to a known p.d.f. $g(\vec{x})$
Compute the value of $g(\vec{x}_i)$

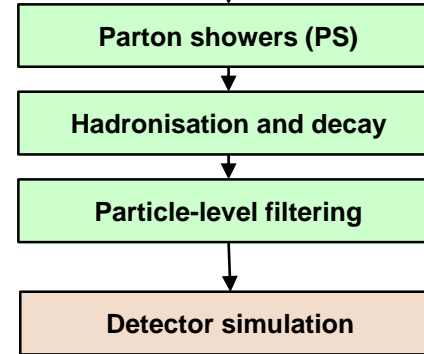
Not shown explicitly here:
parton distribution functions
(initial state parton flavors
and momenta are not fixed)

Phase space sampling optimization
Compute w_{max} over the phase space
Optimize phase space sampling $H(\vec{r}_i)$:
EITHER minimize the variance of I
OR maximize the unweighting efficiency

Matrix element (ME) calculation
For each event i , compute the differential cross-section $f(\vec{x}_i)$
Compute the weight $w_i = f(\vec{x}_i)/g(\vec{x}_i)$

Monte Carlo integration
Average of weights $I = \frac{1}{N} \sum w_i$
→ **Output: I (estimator of $\int x dx$)**

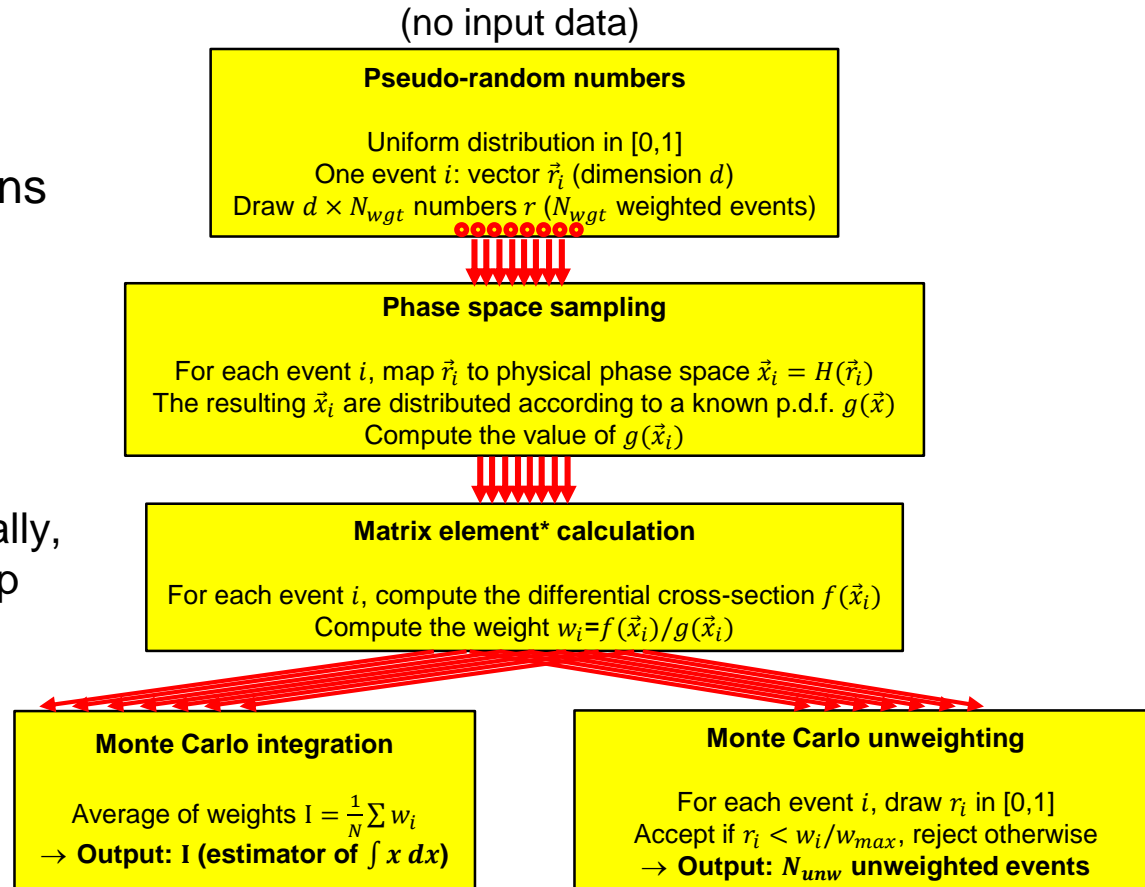
Monte Carlo unweighting
For each event i , draw r_i in $[0,1]$
Accept if $r_i < w_i/w_{max}$, reject otherwise
→ **Output: N_{unw} unweighted events**



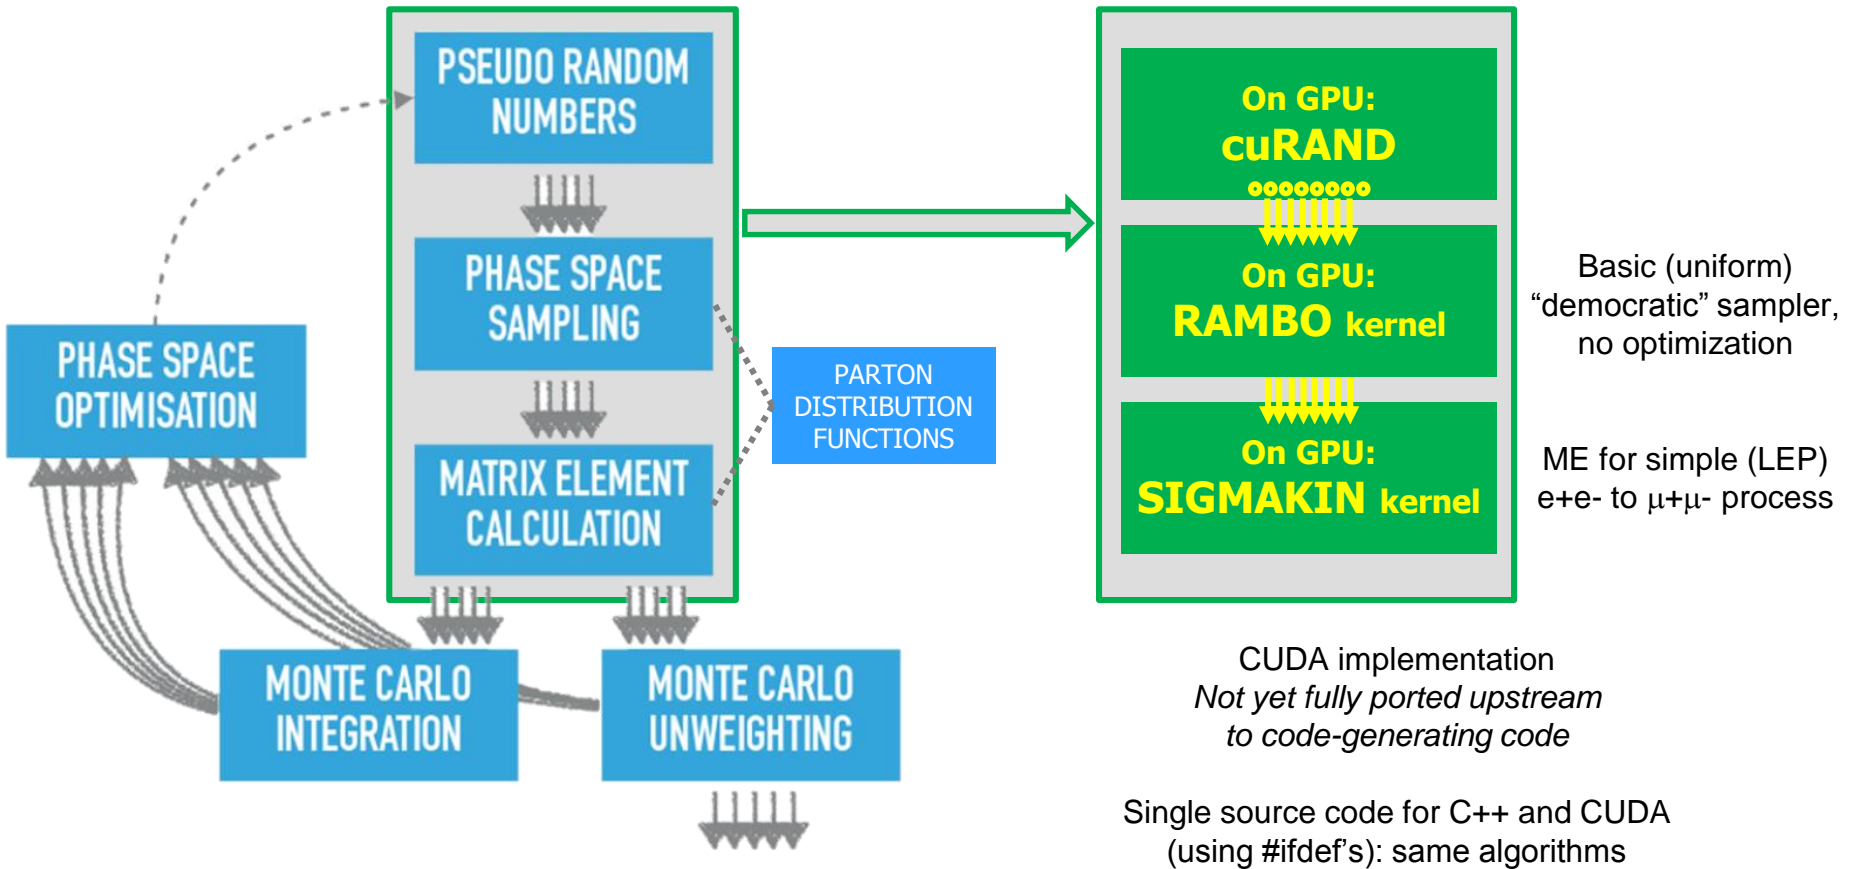
For LHC processes, *the ME calculation is the most computationally-intensive part* of the whole application workflow (but this is not true in a simpler e+e- to $\mu+\mu^-$)

MGaMC on GPU: event-level data-parallel approach

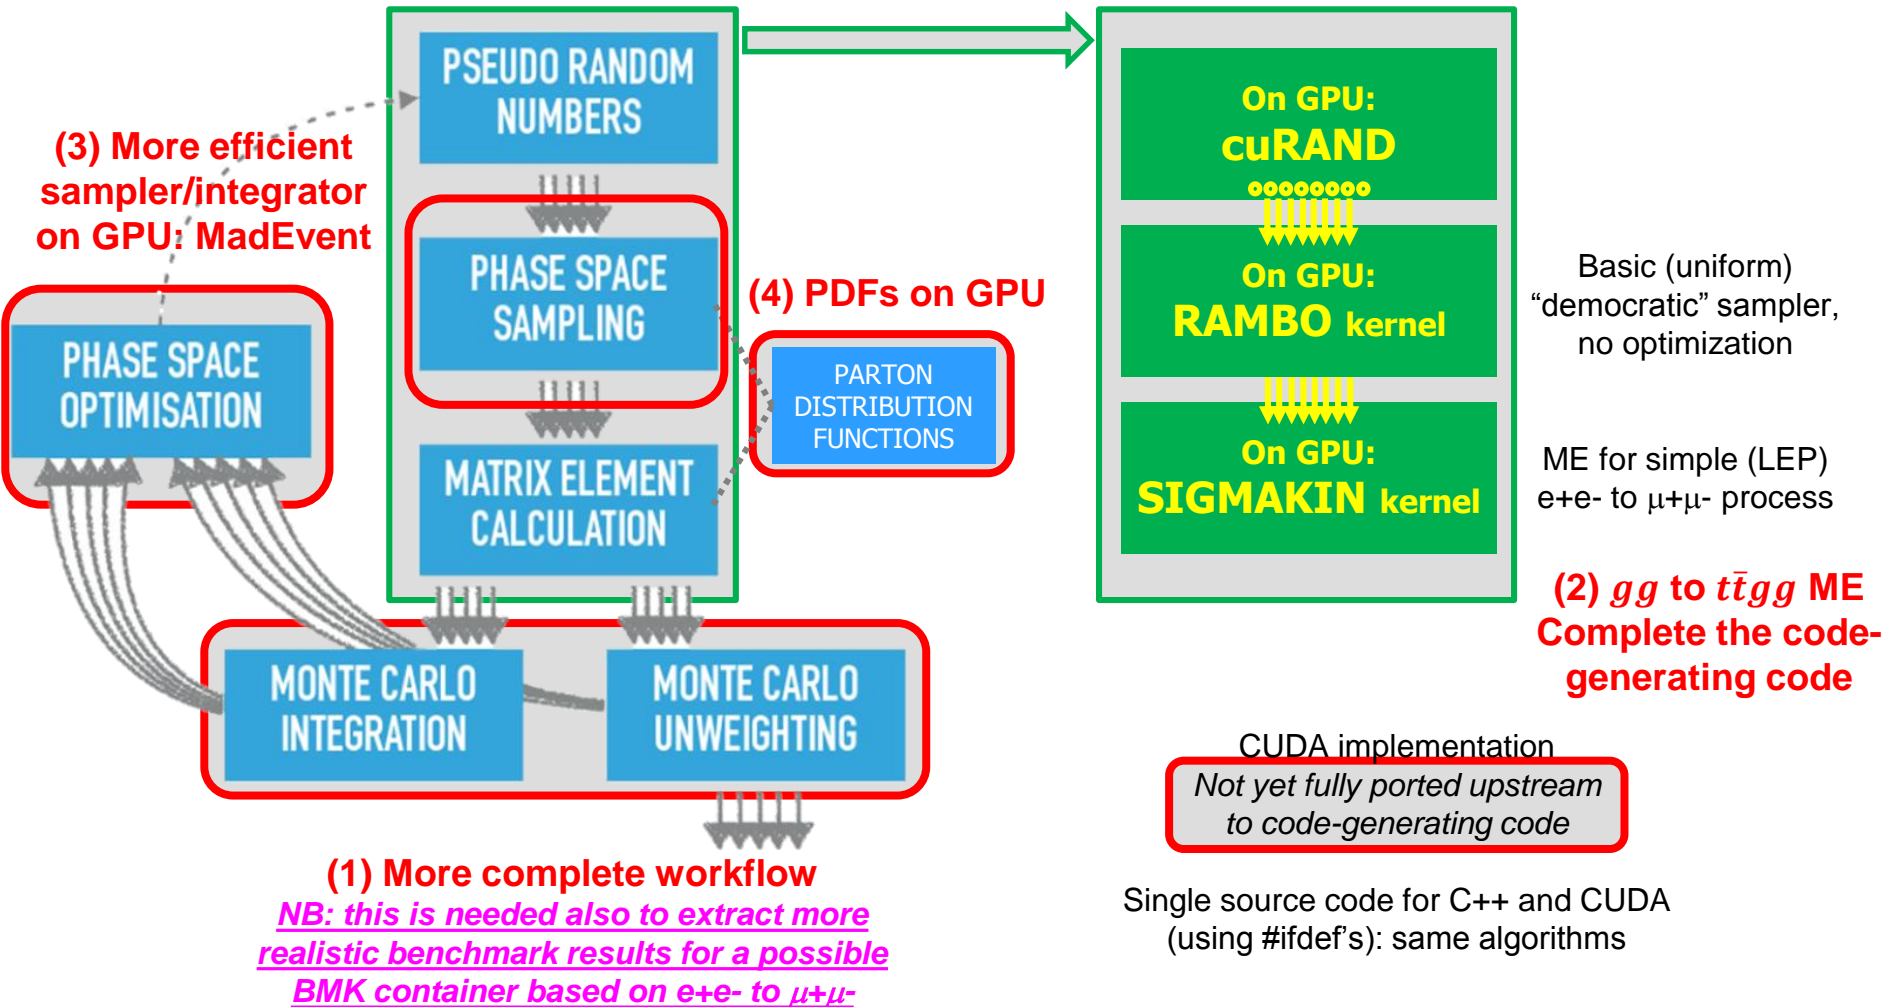
- One main reason why generators are excellent candidates for a GPU port:
 - *The compute-intensive ME calculation is exactly the same function for all events*
 - This reduces the risk of “thread divergence” on GPUs (unlike detector simulation)
- We use an event-level data-parallel approach
 - Execute the same operations on multiple data in parallel
 - *On GPUs: SPMD (Single Program Multiple Data)*
 - On CPUs: SIMD (Single Instruction Multiple Data)?
 - Aka vectorization: eventually, we will also try to speed up the C++ code this way



Status: where are we now?



Status: which are we the (main) ingredients that we are missing for LHC processes?



Benchmarking based on (our current) $e+e- \rightarrow \mu+\mu-$?

PROFILING

$e+e- \rightarrow \mu+\mu-$
4096 blocks / grid
256 threads / block
10 iterations



$e+e- \rightarrow \mu+\mu$, 4096 blocks / grid, 256 threads / block, 1 iteration



Random Numbers
Copy Weights Device \rightarrow Host
Phase Space Sampling

Copy 4-Momenta Device \rightarrow Host

Matrix Element Calculation
Copy MEs Device \rightarrow Host

S. Roiser, [madgraph4gpu meeting](#), Oct 2020

Most time is currently spent in data copy. Conclusion: $e+e- \rightarrow \mu+\mu$ calculations are too simple

- Benchmarking GPUs today based on MG is possible, but VERY preliminary!
 - 1. Even for $ee\mu\mu$, we need at least a realistic integration/unweighting workflow
 - 2. In any case, it would not be representative of more complex LHC processes

Hardware benchmarking, Software benchmarking

- In the HEPiX BMK WG we mainly deal with *hardware benchmarking*
 - Run the same frozen software on different CPUs/GPUs and compare them
 - Goals: accounting/pledging and procurement of compute resources (à la HS06)
 - Reproducible applications, pre-built and containerized as docker images
- In the madgraph4gpu effort we do a whole lot of *software benchmarking*
 - *Change the software (or change build options) and compare old/new versions*
 - We do more, but we *also* compare software performance on different GPUs
 - Software optimizations depend on hardware (professional/FP64 vs consumer/FP32)
 - Software languages and abstraction layers depend on hardware (Nvidia, AMD, Intel)
- There is a lot in common, but also some important differences
 - Main similarity: we both need reproducible application workloads
 - Main difference: prebuilt library/executables vs source code to rebuild
 - It makes sense to keep these two efforts well synchronized

How do we benchmark our own code today? WIP!!!

```
./gcheck.exe -p 16384 32 12
```

```
*****
```

```
NumIterations      = 12
NumThreadsPerBlock = 32
NumBlocksPerGrid   = 16384
```

CUDA/GPU
Full V100

```
-----
TotalEventsComputed = 6291456
RamboEventsPerSec   = 8.130406e+07 sec^-1
MatrixElemEventsPerSec = 6.703134e+08 sec^-1
```

```
*****
```

```
MeanMatrixElemValue = 1.372152e-02 GeV^0
```

```
*****
```

```
00 CudaFree : 0.928572 sec
0a Proclnit : 0.000625 sec
0b MemAlloc : 0.062333 sec
0c GenCreat : 0.010851 sec
1a GenSeed  : 0.000015 sec
1b GenRnGen : 0.007457 sec
2a Rambolni : 0.000108 sec
2b RamboFin : 0.000049 sec
2c CpDTHwgt : 0.006603 sec
2d CpDTHmom : 0.070621 sec
3a SGoodHel : 0.001733 sec
3b SigmaKin : 0.000081 sec
3c CpDTHmes : 0.009305 sec
4a DumpLoop : 0.022506 sec
8a CompStat : 0.031386 sec
9a GenDestr : 0.000064 sec
9b MemFree  : 0.013876 sec
9c CudReset : 0.027250 sec
9d DumpScrn : 0.000217 sec
9e DumpJson : 0.000003 sec
TOTAL      : 1.193654 sec
TOTAL(123) : 0.095971 sec
TOTAL(23)  : 0.088500 sec
TOTAL(3)   : 0.011118 sec
```

```
*****
```

RamboPerSec
No longer sure this
makes sense... bug?

*Exactly the same calculation,
exactly the same result*

Throughput adding Rambo – i.e. ME, plus:
- add random-to-momenta mapping
- (CUDA only) add copy momenta to host
- (CUDA only) add copy weight to host
GPU: 6M in 0.09 sec ~ 7E7/sec
CPU: 6M in 18 sec ~ 3E5/sec
GPU is ~x200 a single CPU thread

MEPerSec throughput:
- ME calculation
- (CUDA only) copy MEs to host
GPU: 6M in 0.011 sec ~ 7E8/sec
CPU: 6M in 16 sec ~ 4E5/sec
GPU is ~x1500 a single CPU thread

This is WIP!!! <https://github.com/madgraph5/madgraph4gpu/issues/22>
Device-to-host copies dominate because e+e- to $\mu+\mu-$ MEs are too simple!

```
./check.exe -p 16384 32 12
```

```
*****
```

```
NumIterations      = 12
NumThreadsPerBlock = 32
NumBlocksPerGrid   = 16384
```

C++/CPU
Single thread

```
-----
TotalEventsComputed = 6291456
RamboEventsPerSec   = 3.243099e+06 sec^-1
MatrixElemEventsPerSec = 3.962151e+05 sec^-1
```

```
*****
```

```
MeanMatrixElemValue = 1.372152e-02 GeV^0
```

```
*****
```

```
0a Proclnit : 0.000467 sec
0b MemAlloc : 0.055191 sec
0c GenCreat : 0.000923 sec
1a GenSeed  : 0.000032 sec
1b GenRnGen : 0.321630 sec
2a Rambolni : 0.082014 sec
2b RamboFin : 1.857938 sec
3b SigmaKin : 15.878891 sec
4a DumpLoop : 0.019507 sec
8a CompStat : 0.028952 sec
9a GenDestr : 0.000104 sec
9b MemFree  : 0.001440 sec
9d DumpScrn : 0.000248 sec
9e DumpJson : 0.000002 sec
TOTAL      : 18.247335 sec
TOTAL(123) : 18.140505 sec
TOTAL(23)  : 17.818844 sec
TOTAL(3)   : 15.878891 sec
```

```
*****
```



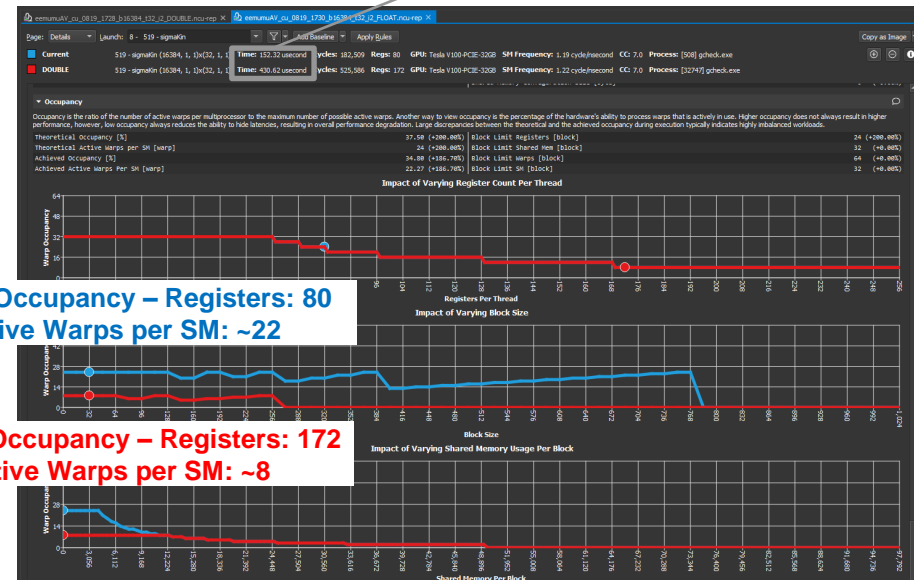
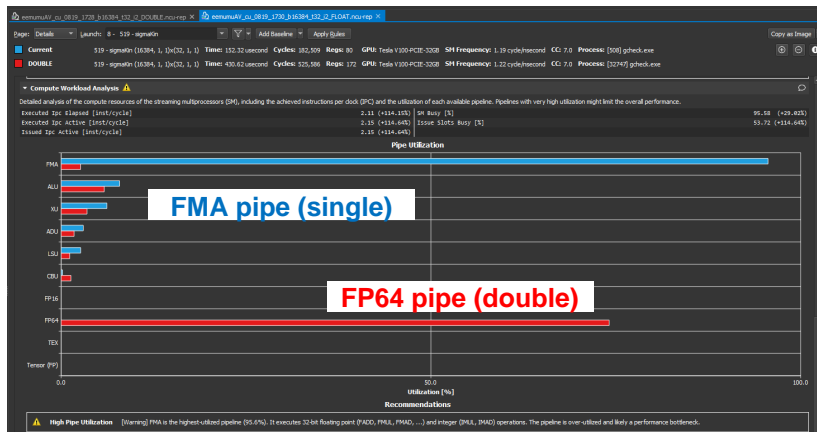
Nvidia GPUs – hackathons and profiling tools

- Some of us attended the <http://gpuhackathons.org> Sheffield event in July
 - Extremely beneficial for us, I highly recommend it for any CUDA developer!
 - Many thanks to our mentors at this specific event! <https://gpuhack.shef.ac.uk>
- This was useful also to understand performance in a benchmarking context
 - Which features of the GPU hardware are relevant (e.g. registers, FP32/64...)
 - Application profiling using two Nvidia tools, [Nsight Systems](#) and [Nsight Compute](#)
- Most plots in these slides are from tools/concepts learnt at the hackathon
 - Some detailed studies in <https://github.com/madgraph5/madgraph4gpu/issues>

One example relevant to GPU benchmarking: float vs. double (consumer vs. professional cards)

- By default MGaMC uses double-precision complex number arithmetic
 - This is required for physics precision... but it is worth checking this again!
- Moving from double to single precision (V100): gain a factor 2.4! ([issue #5](#))
 - Intuitively, being able to use FP32 cores and not only FP64 cores gains a factor 2
 - In addition, single-precision reduces register pressure and increases occupancy

single is 2.4 faster than double (sigmakin = ME kernel)




Outlook: a BMK container for MGaMC on GPU?

- *Creating a container for $e+e^-$ to $\mu+\mu^-$ should be relatively easy and fast*
 - *But first, we need a realistic unweighting workflow and its throughput metric!*
 - This can be interesting to compare GPUs, but is of little relevance to WLCG...
- Creating an LHC-type container needs more development progress first
 - At least a more complete backport of ME calculations to code-generating code
 - A more realistic sampler is also needed for complex processes like gg to $t\bar{t}gg$
 - LHC proton collisions also need a port of PDFs to CUDA
 - But one could also prepare a simpler benchmark before for a gg to $t\bar{t}gg$ process
- *We do not use /cvmfs at all for the moment (no released/installed version)*
 - We would need to build CUDA executables in the CI and embed those in docker
 - There is no input data, this is essentially a pure “GEN” type of workload
- *There is a synergy between the MG and BMK projects, let's stay in touch!*
 - Example: we both need to get hold and test AMD and Intel GPUs at some point!

Backup slides

MC generators and *HL-LHC software and computing*

- One of the main issues (not the only one!): *HL-LHC computing resource gap*
 - Generator performance must also keep up with higher physics precision



CPU cost of generators as a fraction of WLCG CPU resources: for ATLAS, ballpark of 10%-20% (for CMS, this is lower)

ATLAS considers an overall generator speedup by a factor x2 as an R&D goal for HL-LHC

Side note: the higher the fraction of **negative-weight events from MC generators**, the higher the CPU cost of MC simulation, MC reconstruction and analysis (need more MC events)

- Many other challenges, including:
 - WLCG software workloads on **non-traditional resources (HPCs, GPUs...)**
 - Funding and **careers** (especially at the theory/experiment/computing interface)

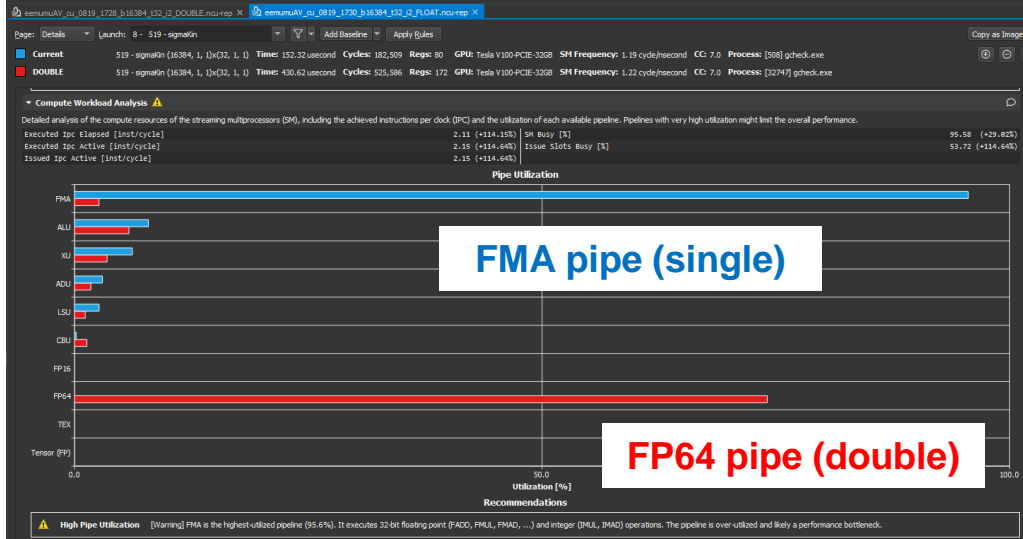


A few performance comparisons

- A few studies from August 2020 (essentially the same code as today)
 - Measurements on Ixbatch: sensitive to external load, reproducible within ~10%
 - Extensive discussion of numbers and Nsight Compute profiler plots on <https://github.com/madgraph5/madgraph4gpu/issues>
 - General approach: change a #define switch to a non-default value in the code
 - Now kept cuComplex and float switches, but hardcoded AOSOA and “local” memory
- From double precision to single precision: gain a factor 2.4! ([issue #5](#))
 - FMA (FP32) used instead of FP64 pipe; fewer registers hence higher occupancy
- Memory layout for 4-momenta ([issue #16](#)): study “requests” vs “sectors”
 - Default: AOSOA with 4 events per array (four 8-byte doubles: 32-byte cache line)
 - AOS ~7% slower: memory not coalesced, #sectors (transactions) factor 4 higher
 - SOA ~2% slower: slightly higher number of registers
 - Side result of this study: improve helicity filtering to reduce #requests ([issue #24](#))
- Memory for wavefunctions in ixxx/oxxx/FV_xx ([issue #7](#)): default is “local”
 - Lose 70% with “global” (become memory-bound!), also lose 35% with “shared”
- From thrust::complex to cuComplex: lose ~5% ([issue #6](#))
 - Require execution of higher number of instructions (something to fix?)

FP: double vs single precision (issue #5)

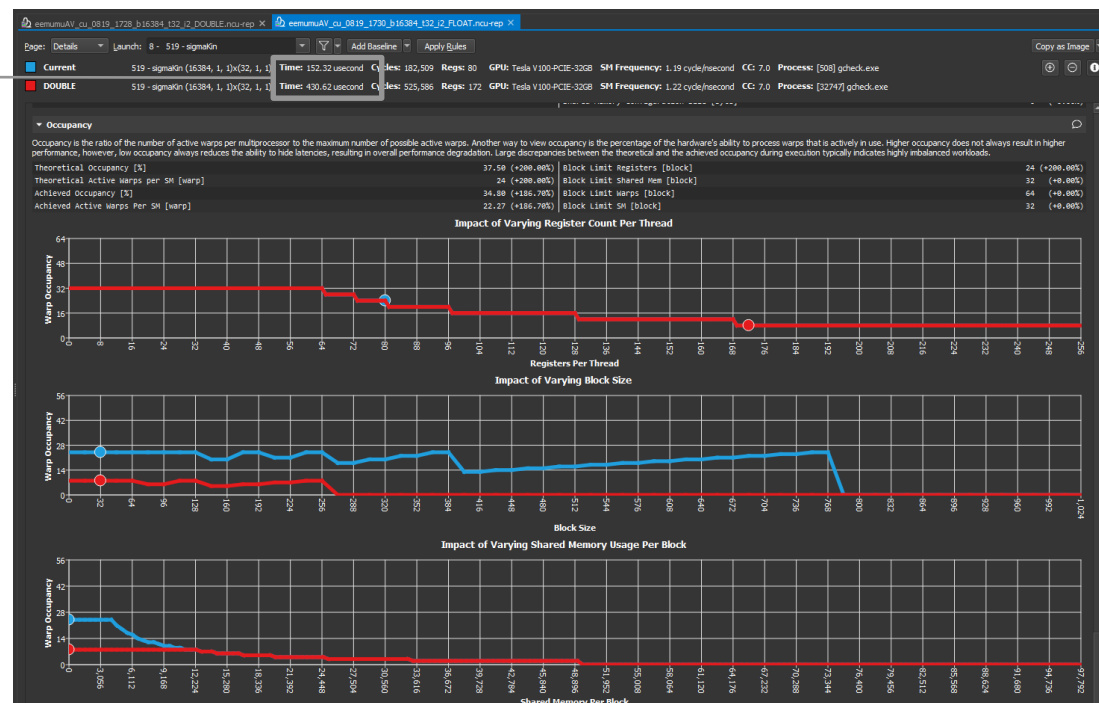
- Two main effects
 - Exploit the FP32 units (unused otherwise!)
 - Fewer registers hence higher warp occupancy



single is 2.4 faster than double (sigmakin = ME kernel)

Single Occupancy
Registers: 80
Active Warps per SM: ~22

Double Occupancy
Registers: 172
Active Warps per SM: ~8



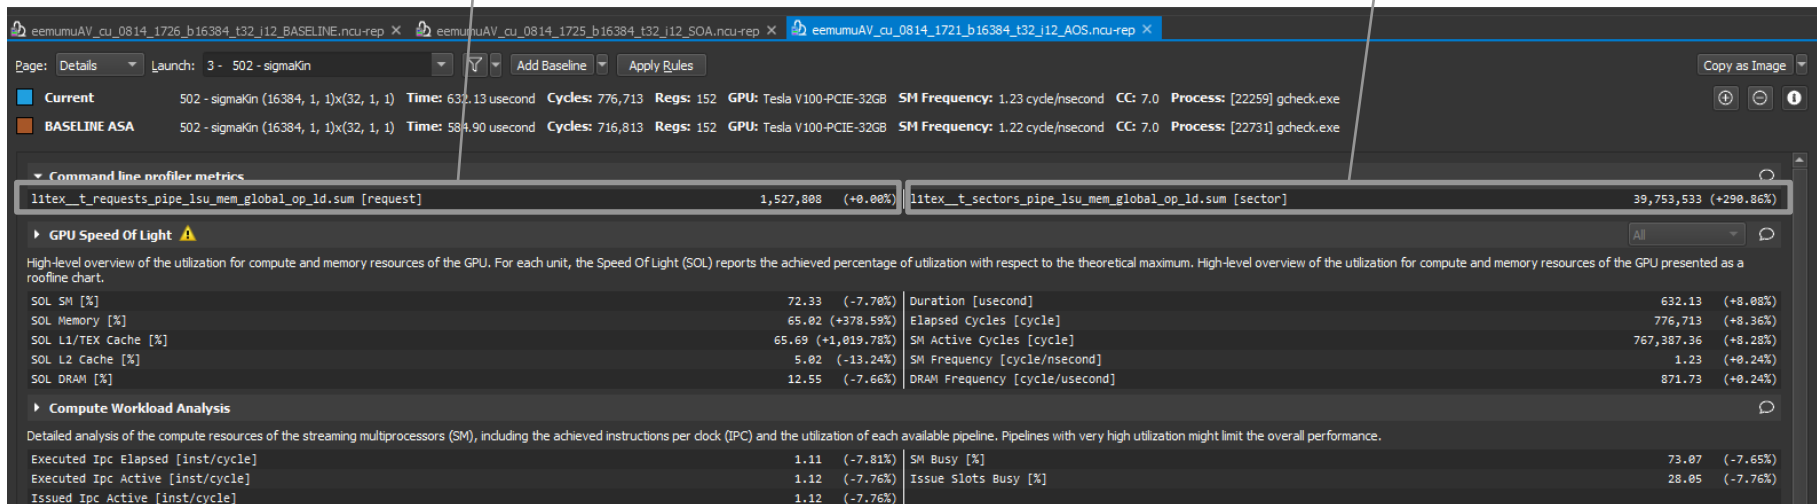
4-momenta memory: AOSOA vs AOS (issue #16)

Number of “requests”:

- It is the same for AOSOA or AOS (the same information needs to be retrieved...)
- Later on, the number show here was dramatically decreased by improving helicity filtering (issue #24)

Number of “sectors” (transactions):

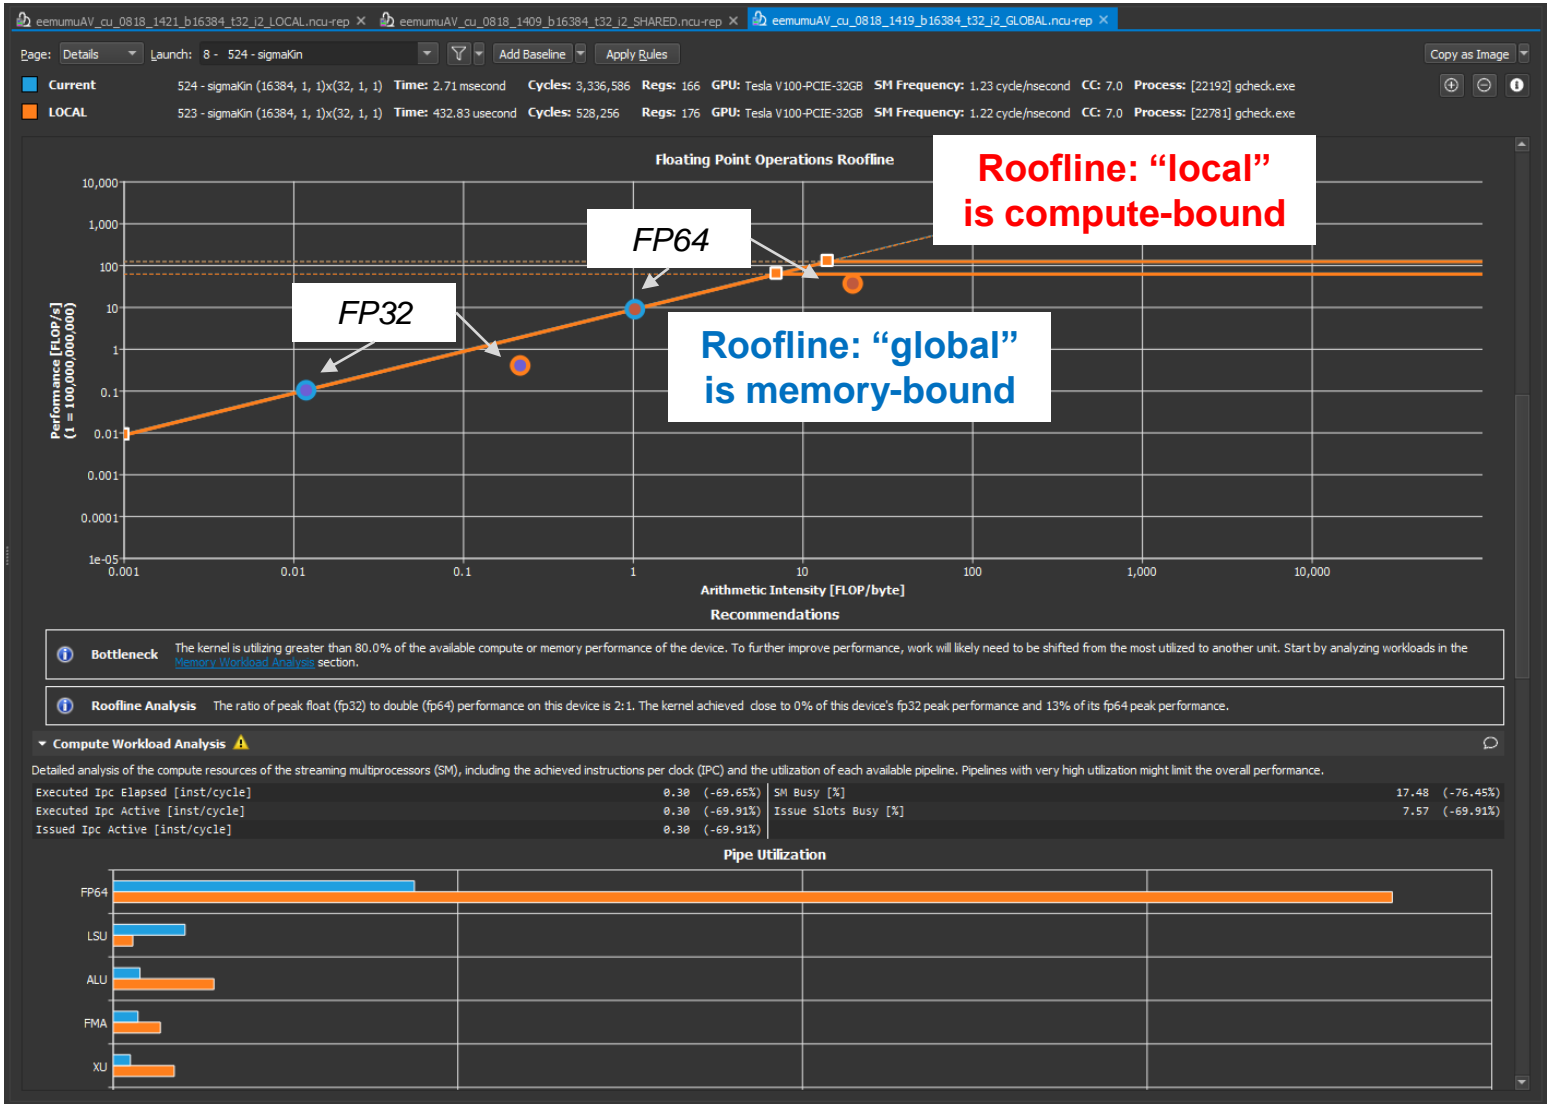
- It is a factor 4 higher for AOS than for AOSOA
- One roundtrip (AOSOA, coalesced) vs four roundtrips (AOS, not coalesced) to retrieve four 8-byte doubles from a 32-byte cache line



A few useful links:

- <https://developer.nvidia.com/blog/using-nsight-compute-to-inspect-your-kernels/>
- <https://docs.nvidia.com/nsight-compute/2019.5/NsightComputeCli/index.html#nvprof-metric-comparison>
- <https://stackoverflow.com/questions/60535867/what-is-a-transaction-and-a-request-in-the-gld-transactions-per-request-metric>

Wavefunction memory: “local” vs “global” (issue #7)



4. Sampling

NVIDIA Nsight Compute supports periodic sampling of the warp program counter and warp scheduler state on desktop devices of compute capability 6.1 and above.

At a fixed interval of cycles, the sampler in each streaming multiprocessor selects an active warp and outputs the program counter and the warp scheduler state. The tool selects the minimum interval for the device. On small devices, this can be every 32 cycles. On larger chips with more multiprocessors, this may be 2048 cycles. The sampler selects a random active warp. On the same cycle the scheduler may select a different warp to issue.

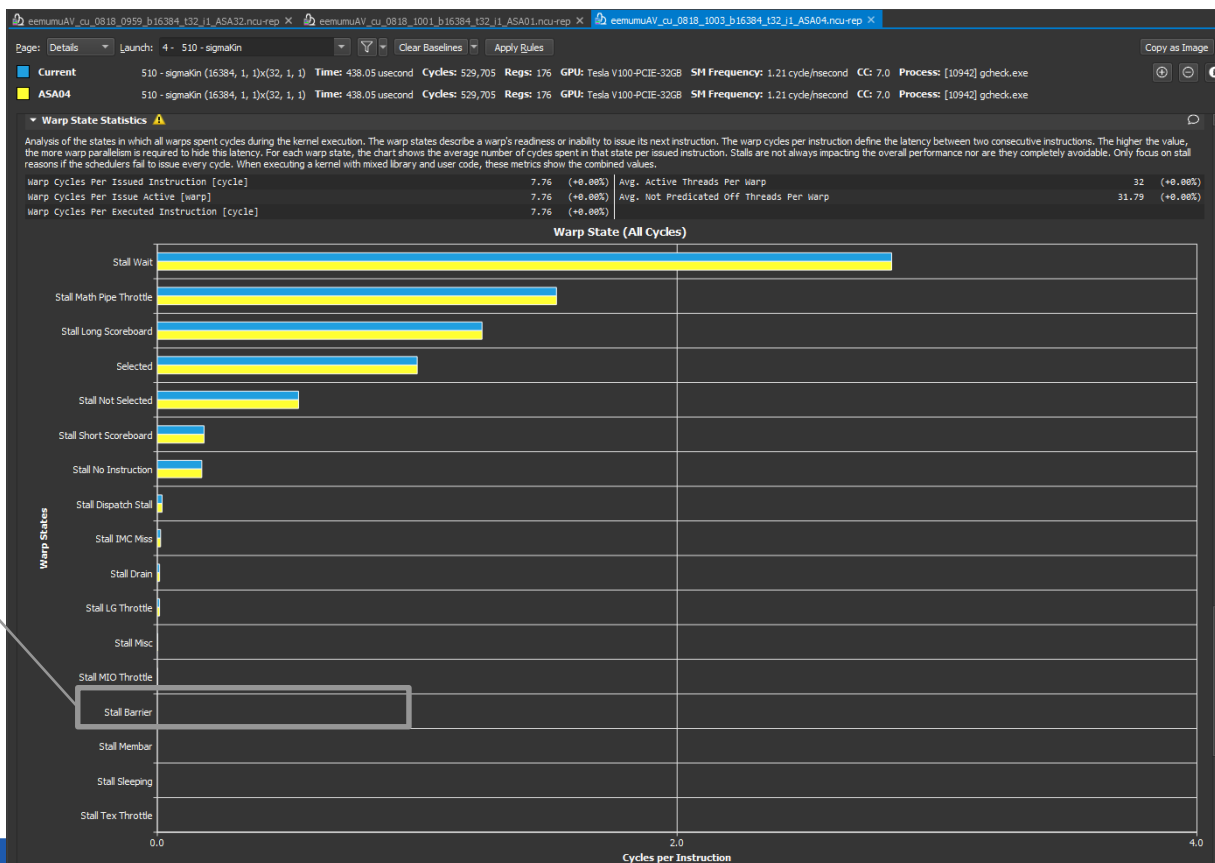
4.1. Warp Scheduler States

Table 2. Warp Scheduler States

State	Hardware Support	Description
Allocation	5.2-6.1	Warp was stalled waiting for a branch to resolve, waiting for all memory operations to retire, or waiting to be allocated to the micro-scheduler.
Barrier	5.2+	Warp was stalled waiting for sibling warps at a CTA barrier. A high number of warps waiting at a barrier is commonly caused by diverging code paths before a barrier. This causes some warps to wait a long time until other warps reach the synchronization point. Whenever possible, try to divide up the work into blocks of uniform workloads. Also, try to identify which barrier instruction causes the most stalls, and optimize the code executed before that synchronization point first.

No evidence for “thread divergence” (issue #25)

<https://docs.nvidia.com/nsight-compute/ProfilingGuide/index.html#statistical-sampler>



- Nsight Compute:
 - Stall Barrier = 0

- Similarly, no indication for thread divergence from Nsight System