



GPU Profiling with Celeritas

Peter Heywood, Research Software Engineer

The University of Sheffield

2023-06-22

Increase Science Throughput

- Ever-increasing demand for increased simulation throughput

1. Buy more / “better” hardware

2. Improve Software

- Improve implementations
- Improve algorithms (i.e. work efficiency)

Must understand software performance to improve performance

Profile

Celeritas

Celeritas is a new Monte Carlo transport code designed for high-performance simulation of high-energy physics detectors.

The Celeritas project implements HEP detector physics on GPU accelerator hardware with the ultimate goal of supporting the massive computational requirements of the HL-LHC upgrade.

- github.com/celeritas-project/celeritas
- NVIDIA GPUs via [CUDA](#)
- AMD GPUs via [HIP](#)
- Ben Morgan - “Detector Simulations in Particle Physics”

Profiling Tools

- CPU-only profilers
 - `gprof`, `perf`, `Kcachegrind`, `VTune`, ...
- NVIDIA Profiling tools
 - `Nsight Systems`
 - `NVIDIA Nsight Compute`
 - `nvprof`
- AMD Profiling tools
 - `roctracer`
 - `rocsys`
 - `rocprofv2`

Hardware

- Development machine:
 - NVIDIA Titan V (SM 70, 250W)
 - NVIDIA Titan RTX (SM 75, 280W)
 - 16x fewer FP64 units
 - Intel i7-6850K
- HPC:
 - NVIDIA H100 PCI-e (SM 90, 350W)
 - AMD EPYC 7413



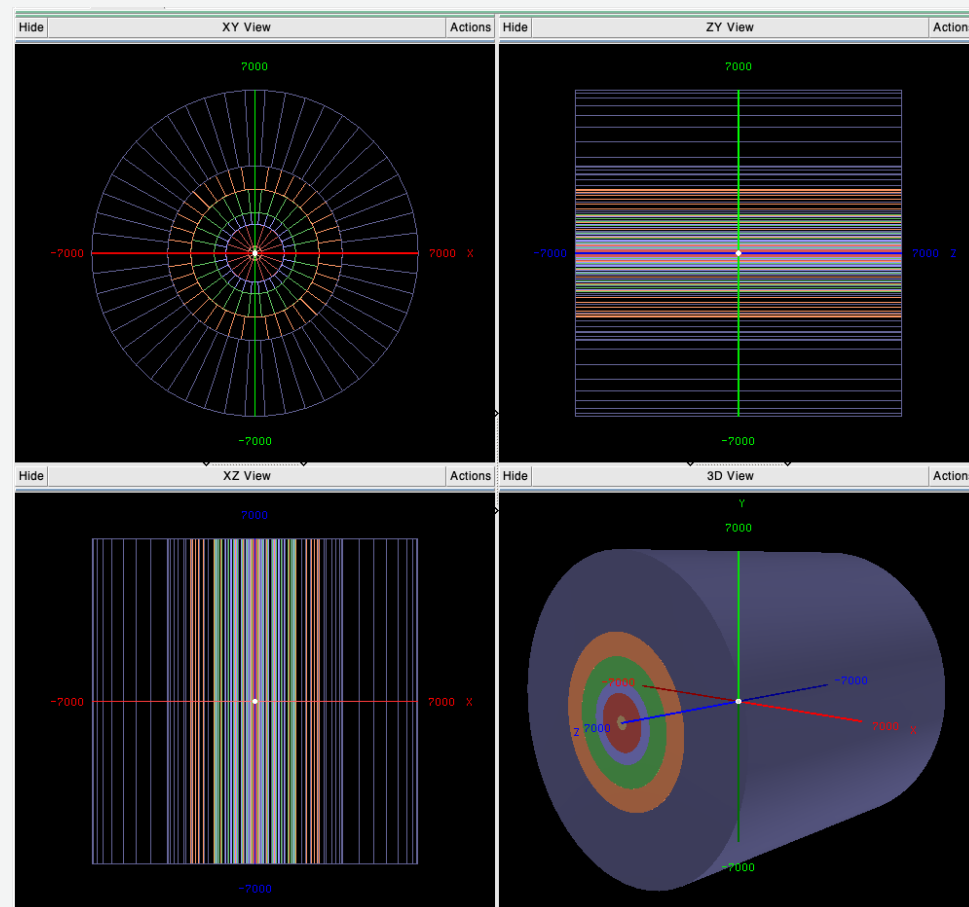
Titan Xp & Titan V GPUs

Inputs / Configuration

- Inputs should ideally be:
 - Representative of real-world use
 - Large enough to fully utilise hardware
 - Small enough to generate usable profile data
- Optimised build
 - `-DCMAKE_BUILD_TYPE=Release, -O3`
 - `-DCMAKE_BUILD_TYPE=RelWithDebInfo, -O2 -g`
- Celeritas `c8db3fce, v0.3.0`
- MPI disabled via `export CELER_DISABLE_PARALLEL=1`

Profile Scenario: Simple CMS

- `celer-sim` test case
- Simple geometry
- Short running
- `simple-cms.gdm1`
- `gamma-3evt-15prim.hepmc3`
- `ctest -R app/celer-sim:device`



github.com/celeritas-project/benchmarks/geant4-validation-app/testem3_evd.png

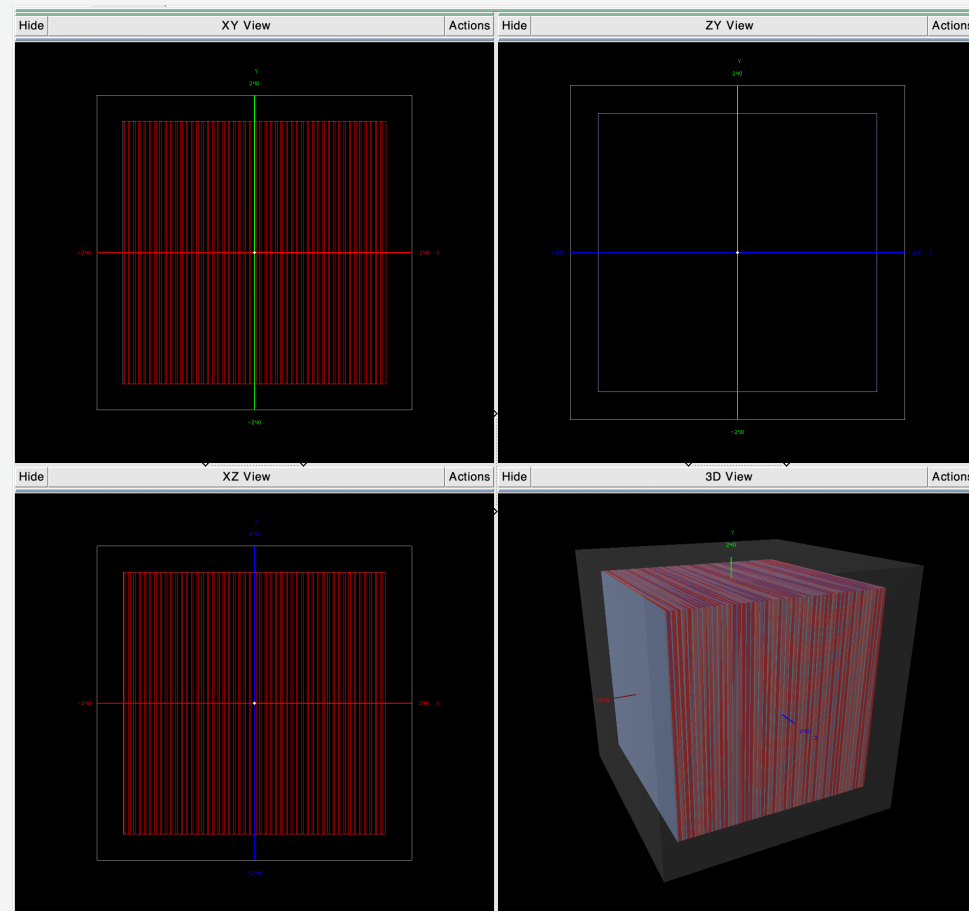
Profile Scenario: TestEm3

- `celer-g4`
- `testem3-flat.gdml`
- `testem3.1k.hepmc3`

```
/control/verbose 0
/tracking/verbose 0
/run/verbose 0
/event/verbose 0
```

```
/celer/outputFile testem3-1k.out.json
/celer/maxNumTracks 524288
/celer/maxNumEvents 2048
/celer/maxInitializers 4194304
/celer/secondaryStackFactor 3
```

```
/celerg4/geometryFile
/celeritas/test/celeritas/data/testem3-
flat.gdml
/celerg4/eventFile
/benchmarks/testem3.1k.hepmc3
```



github.com/celeritas-project/benchmarks/geant4-validation-app/simple_cms_evd.png

Nsight Systems

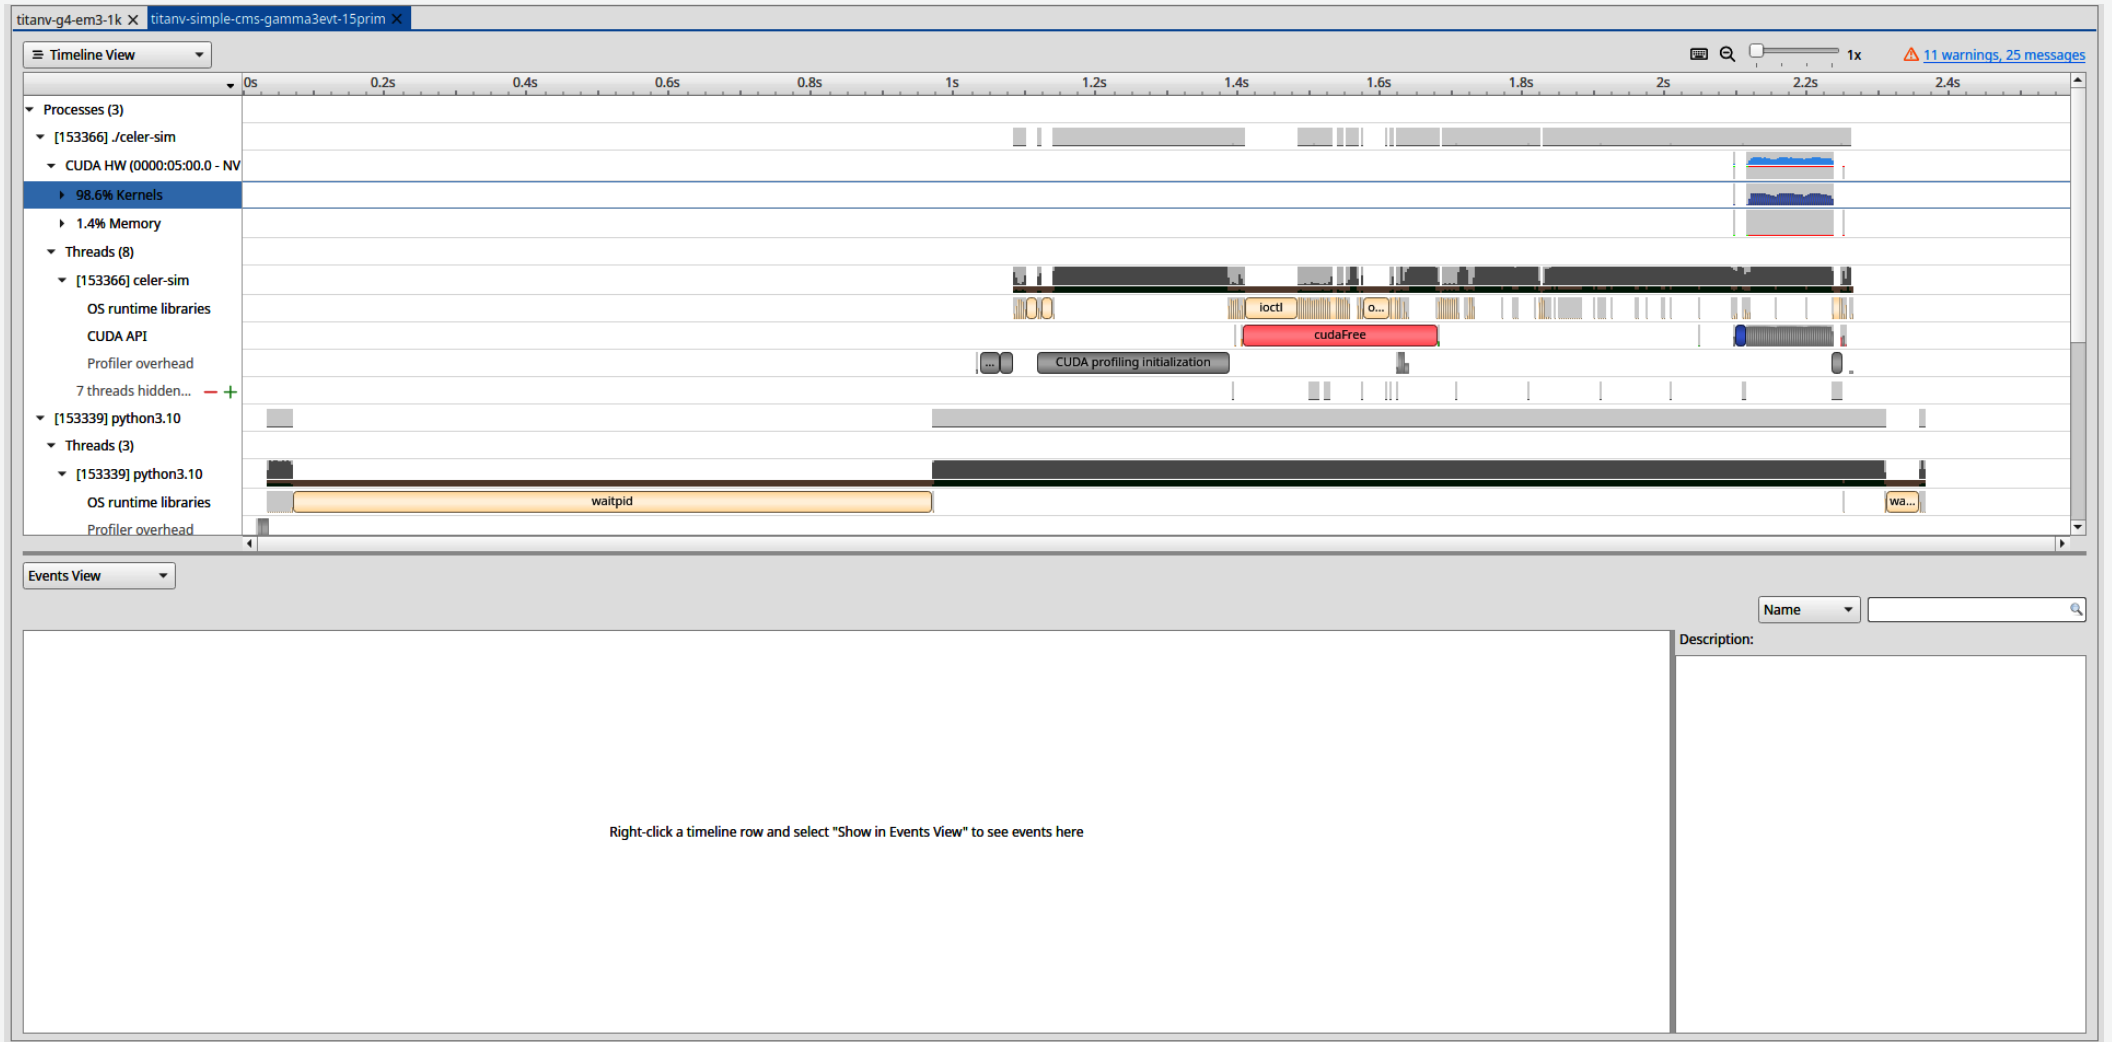
Nsight Systems

- System-wide performance analysis
- CPU + GPU
- Visualise a timeline of events
- CUDA API information, kernel block sizes etc
- Pascal GPUs or newer (SM 60+)

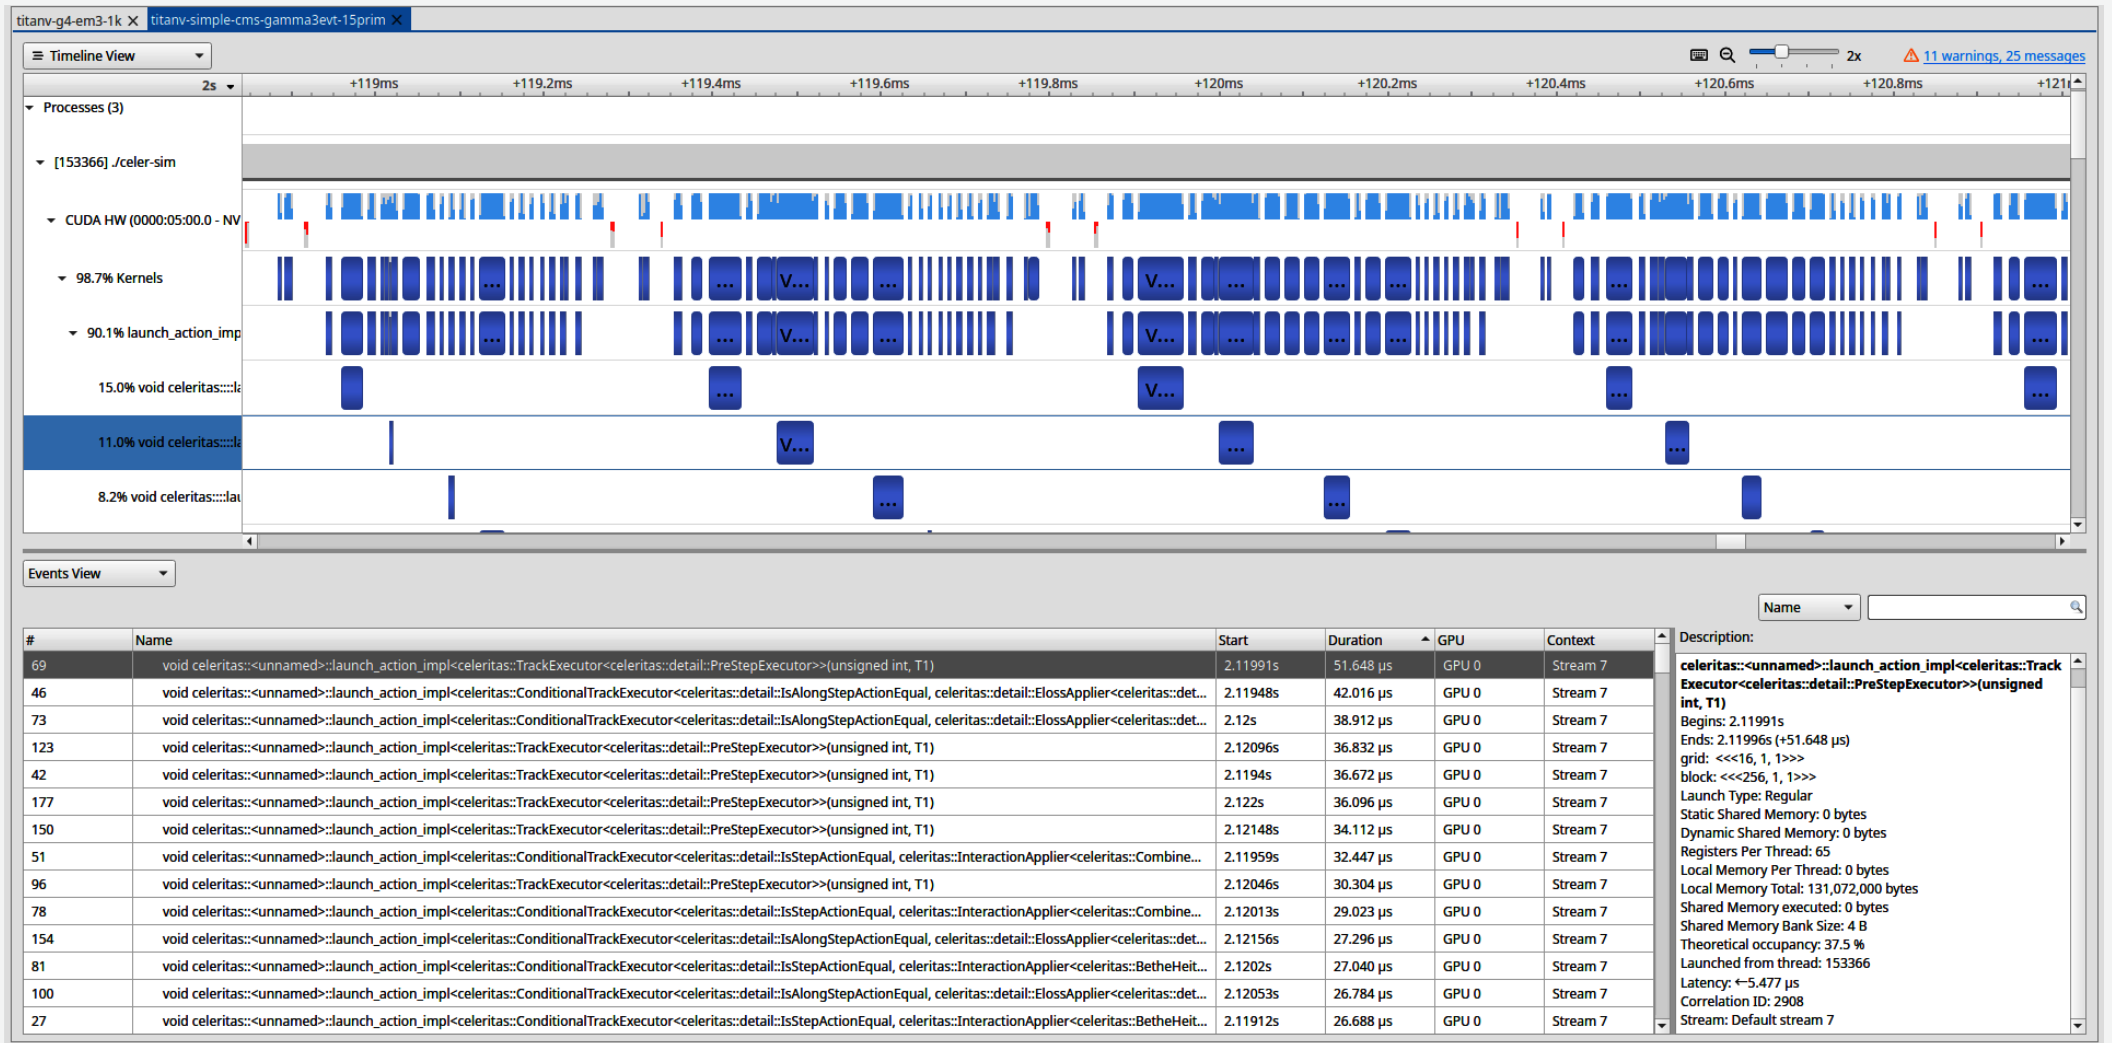
```
1 nsys profile -o timeline ./celer-g4 input.mac  
2 nsys-ui timeline.nsys-rep
```

Nsight Systems: Simple CMS

Timeline: Simple CMS (Titan V)



Timeline: Simple CMS (Titan V)



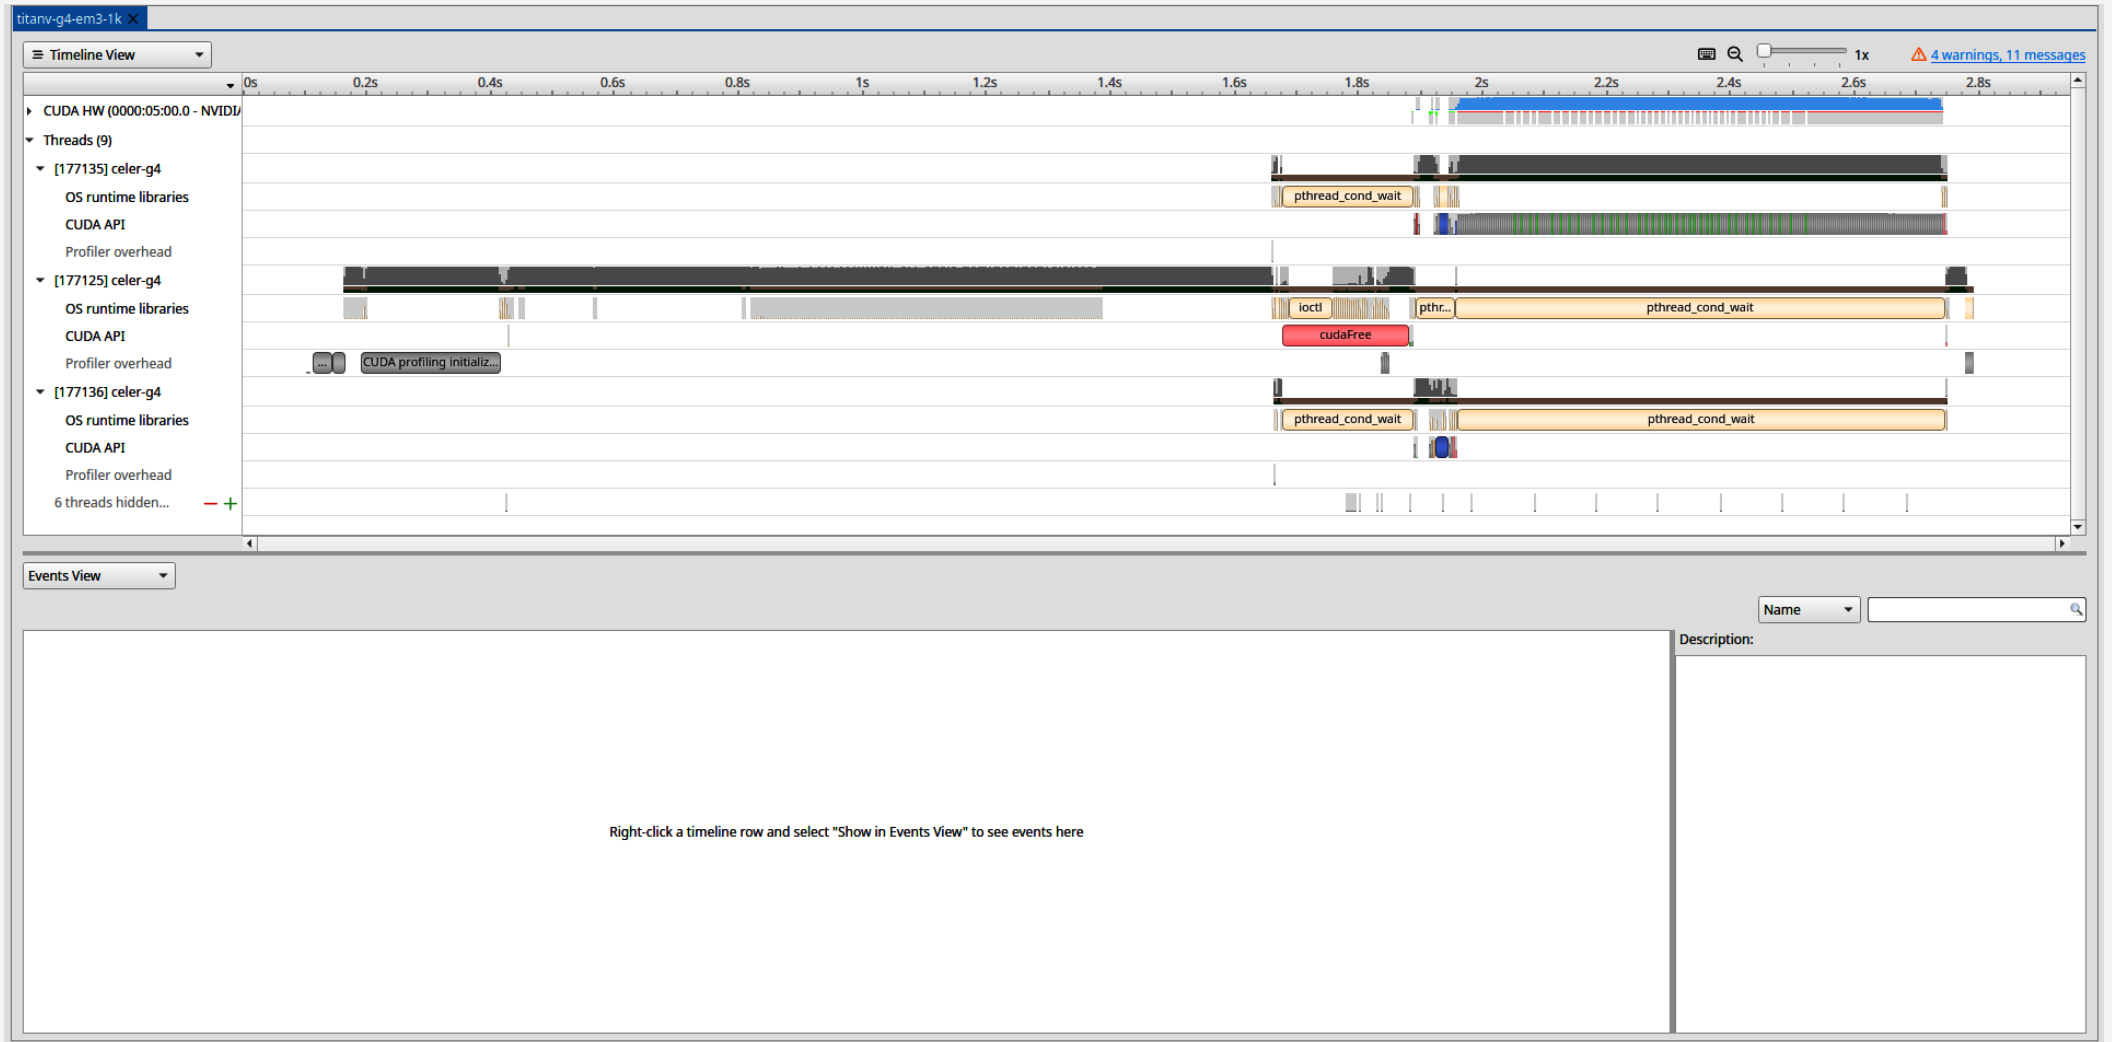
Nsys Table: Simple CMS (Titan V)

#	Name	Start	Duration
599	void celeritas::<unnamed>::launch_action_impl<celeritas::ConditionalTrackExecutor<celeritas::detail::IsStepActionEqual, celer...	2.12798s	88.960 µs
154	void celeritas::<unnamed>::launch_action_impl<celeritas::TrackExecutor<celeritas::detail::PreStepExecutor>>(unsigned int, T1)	2.11991s	51.648 µs
129	void celeritas::<unnamed>::launch_action_impl<celeritas::ConditionalTrackExecutor<celeritas::detail::IsAlongStepActionEqual...	2.11948s	42.016 µs
158	void celeritas::<unnamed>::launch_action_impl<celeritas::ConditionalTrackExecutor<celeritas::detail::IsAlongStepActionEqual...	2.12s	38.912 µs
647	void celeritas::<unnamed>::launch_action_impl<celeritas::TrackExecutor<celeritas::detail::PreStepExecutor>>(unsigned int, T1)	2.12891s	38.688 µs
4851	void celeritas::<unnamed>::launch_action_impl<celeritas::TrackExecutor<celeritas::detail::PreStepExecutor>>(unsigned int, T1)	2.19895s	38.592 µs
4417	void celeritas::<unnamed>::launch_action_impl<celeritas::TrackExecutor<celeritas::detail::PreStepExecutor>>(unsigned int, T1)	2.19257s	38.176 µs
2387	void celeritas::<unnamed>::launch_action_impl<celeritas::TrackExecutor<celeritas::detail::PreStepExecutor>>(unsigned int, T1)	2.15739s	37.856 µs
5170	void celeritas::<unnamed>::launch_action_impl<celeritas::TrackExecutor<celeritas::detail::PreStepExecutor>>(unsigned int, T1)	2.20439s	37.600 µs
5199	void celeritas::<unnamed>::launch_action_impl<celeritas::TrackExecutor<celeritas::detail::PreStepExecutor>>(unsigned int, T1)	2.20491s	37.248 µs
2706	void celeritas::<unnamed>::launch_action_impl<celeritas::TrackExecutor<celeritas::detail::PreStepExecutor>>(unsigned int, T1)	2.16298s	36.992 µs
212	void celeritas::<unnamed>::launch_action_impl<celeritas::TrackExecutor<celeritas::detail::PreStepExecutor>>(unsigned int, T1)	2.12096s	36.832 µs

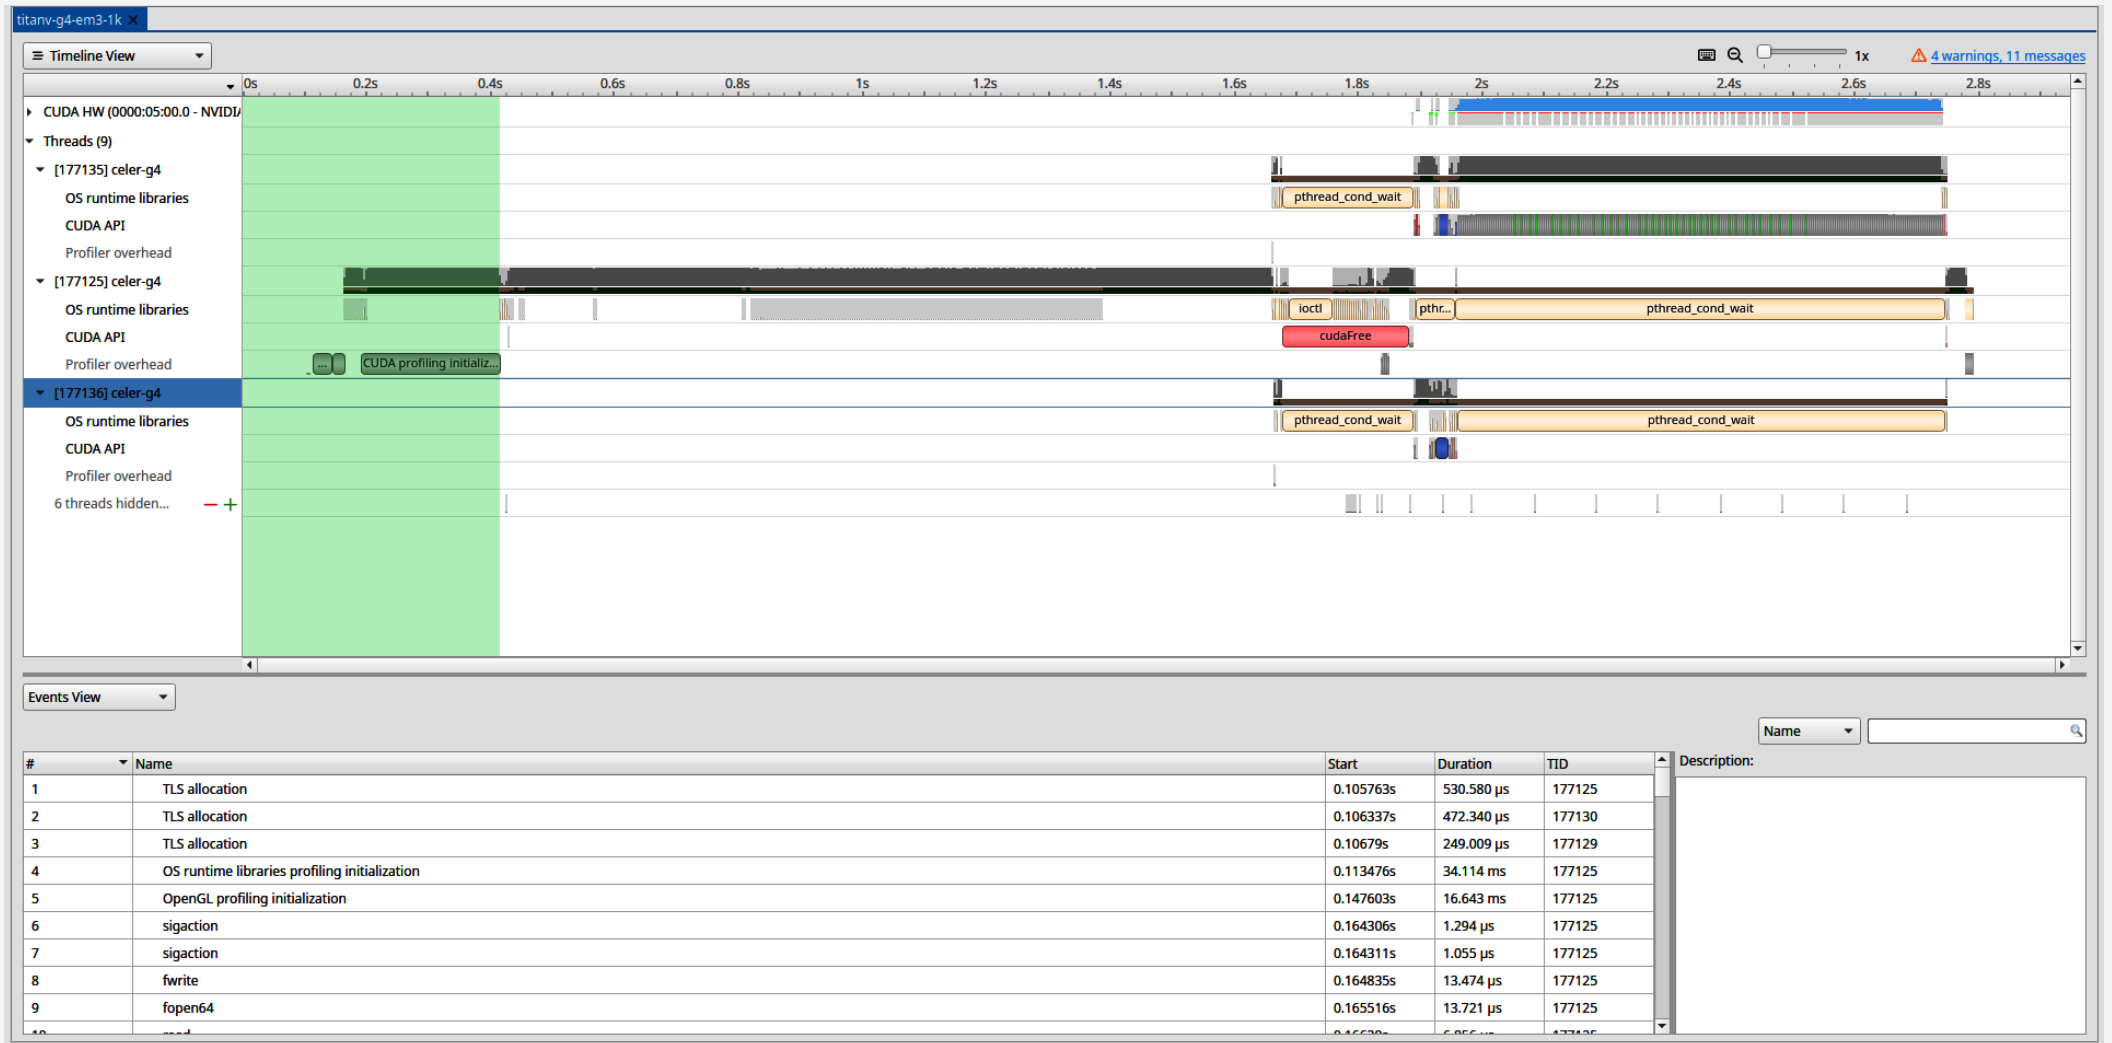
- Longest kernel: 88us
- Launch latency: 5.2us
- Threads: 16 * 256

Nsight Systems: TestEm3

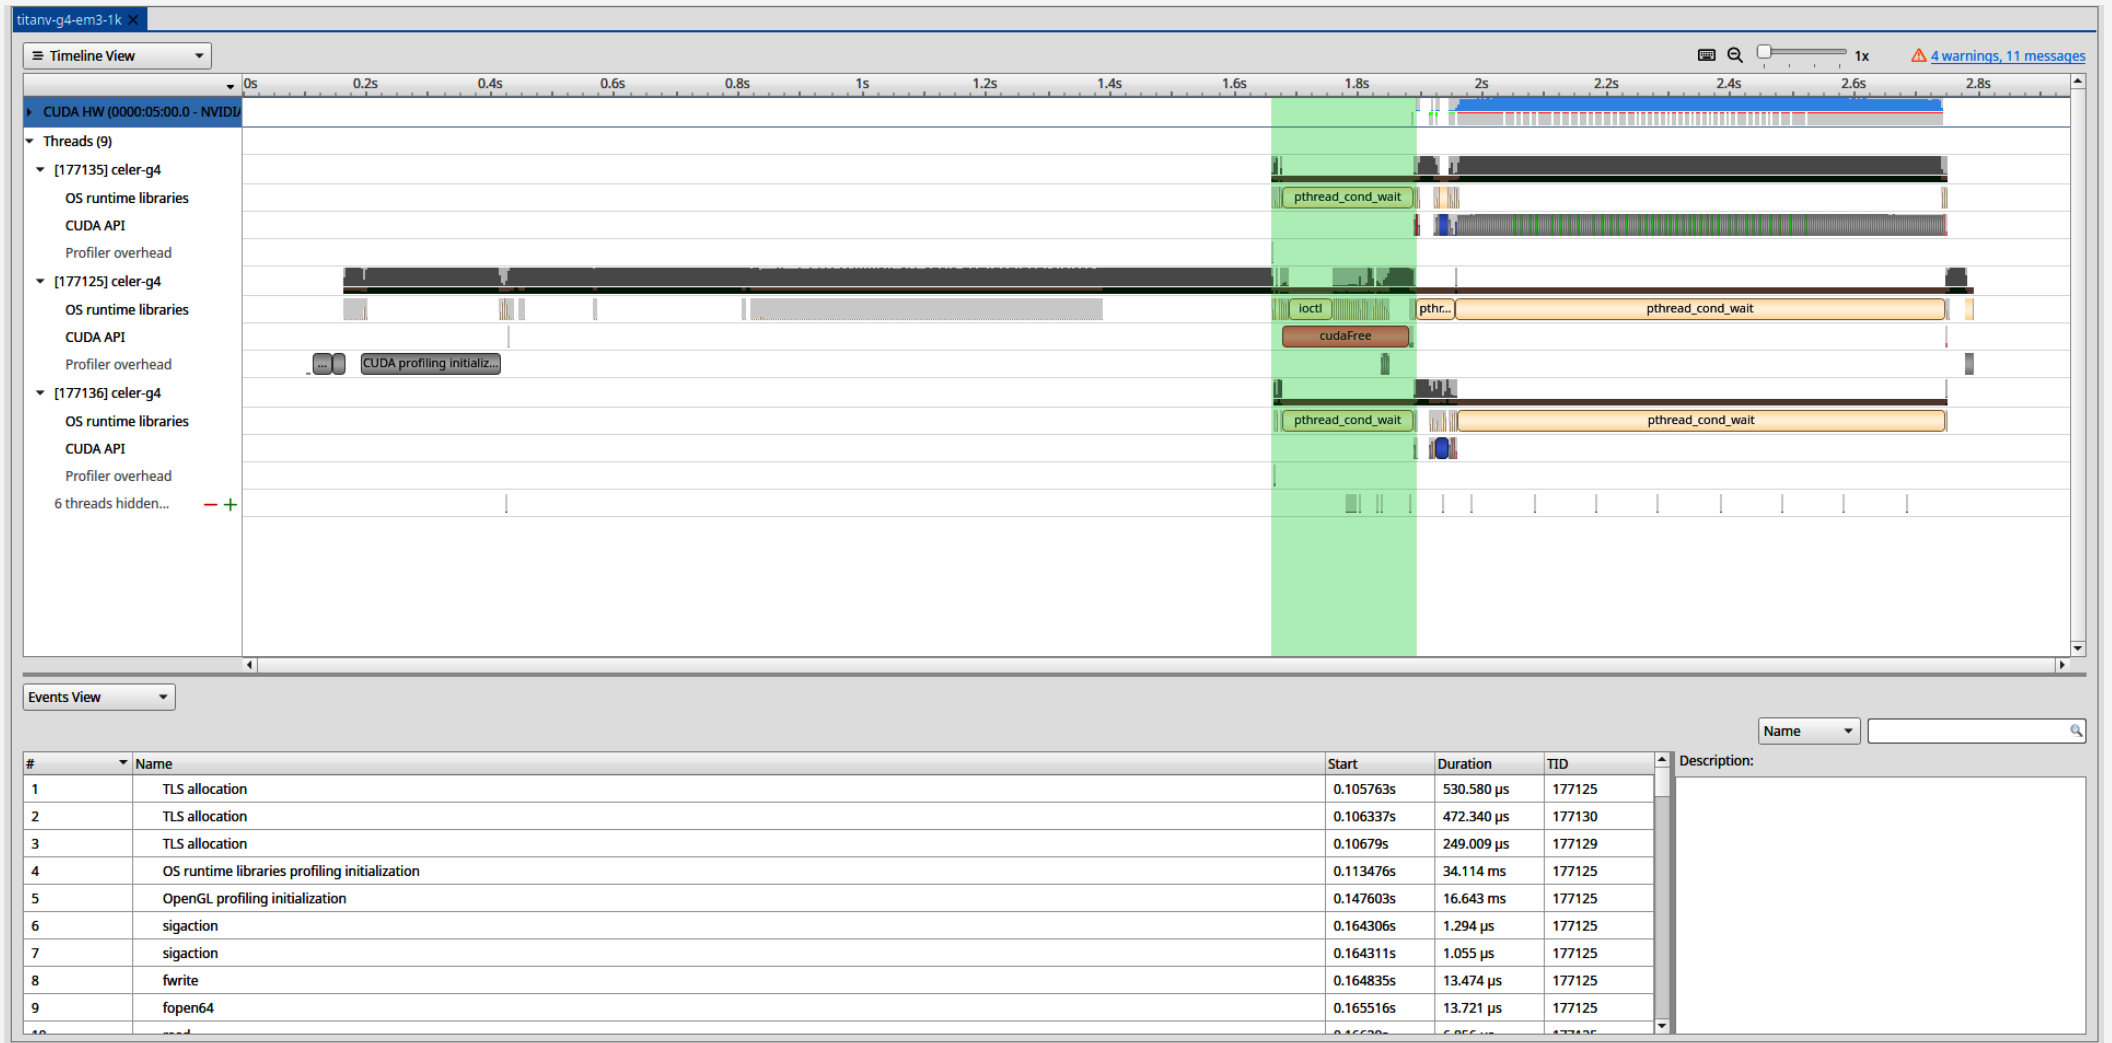
Timeline: TestEm3 (Titan V)



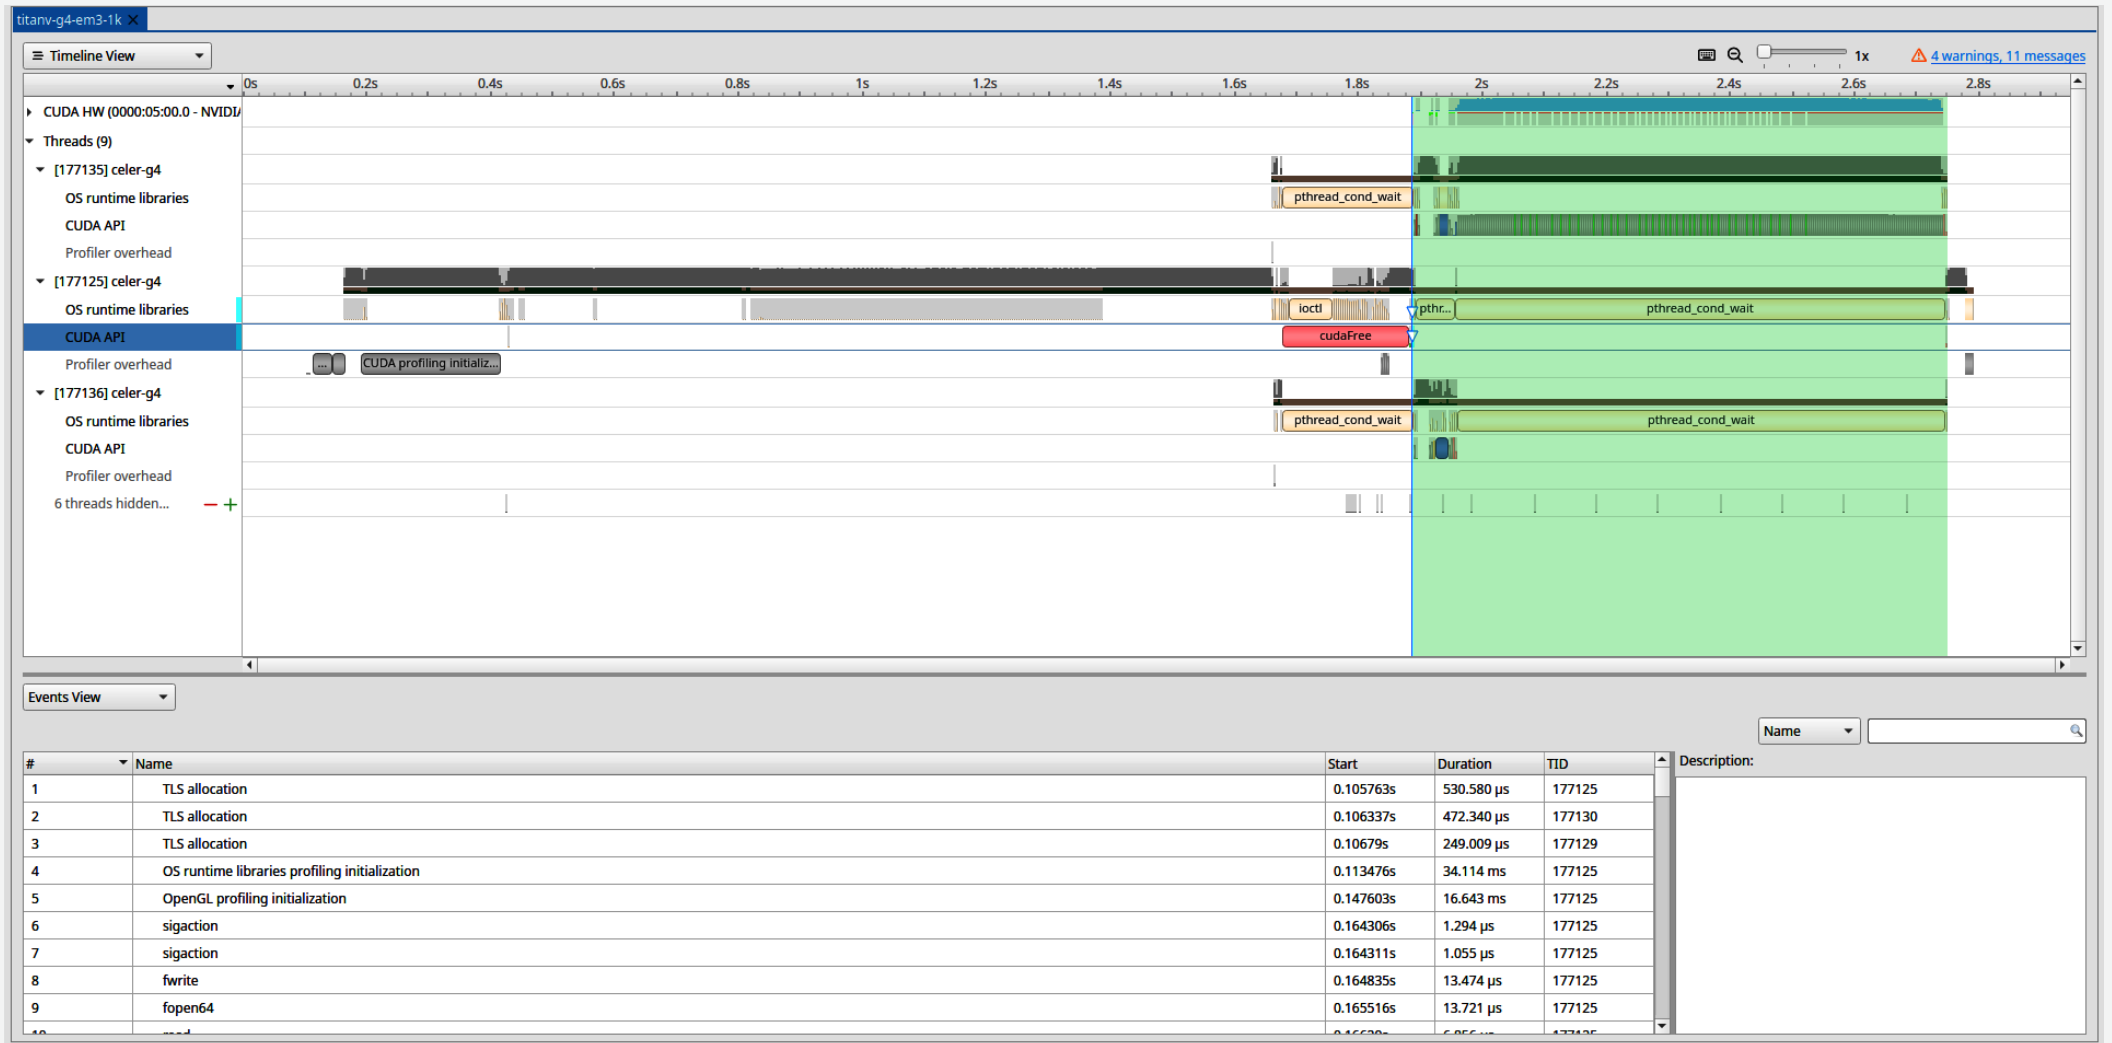
Timeline: TestEm3 (Titan V) Profiling Overheads



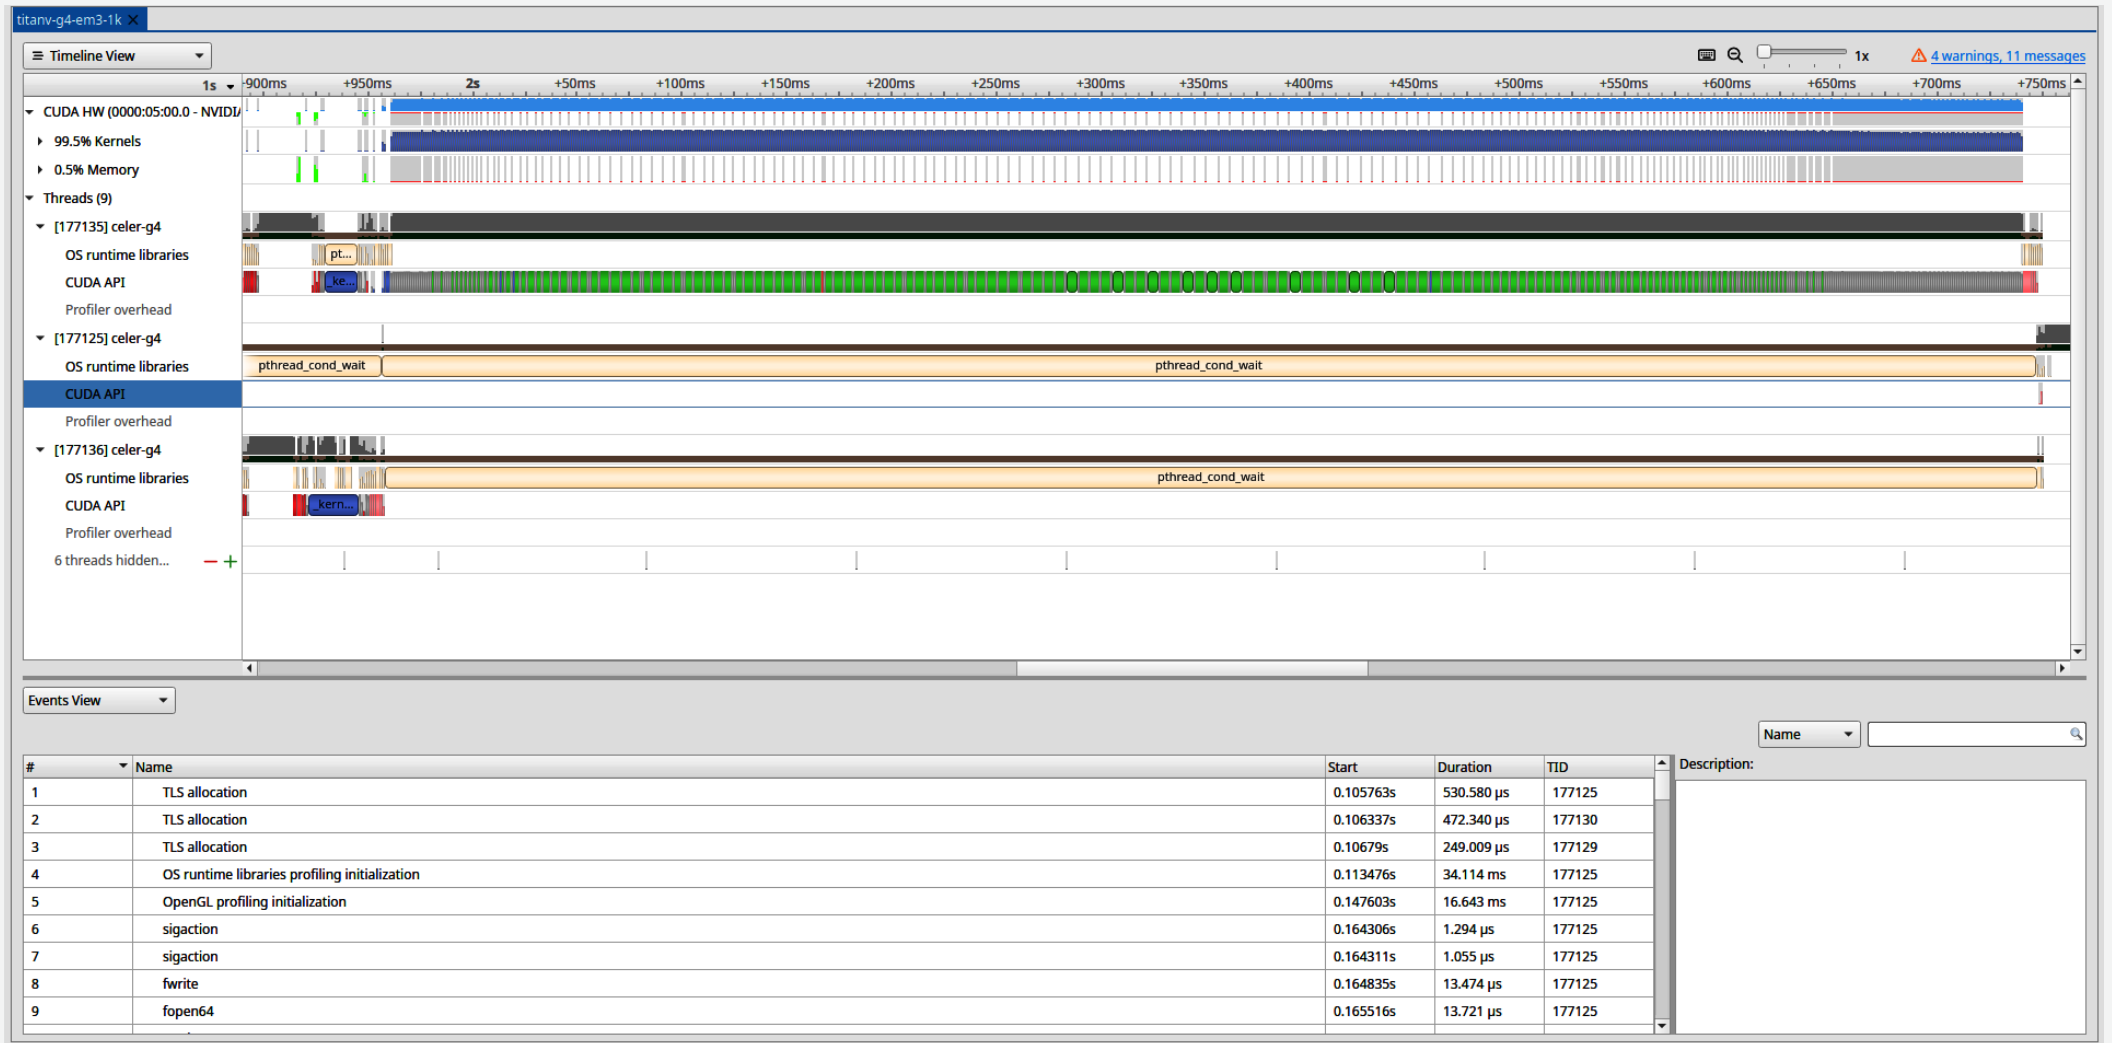
Timeline: TestEm3 (Titan V) GPU Init



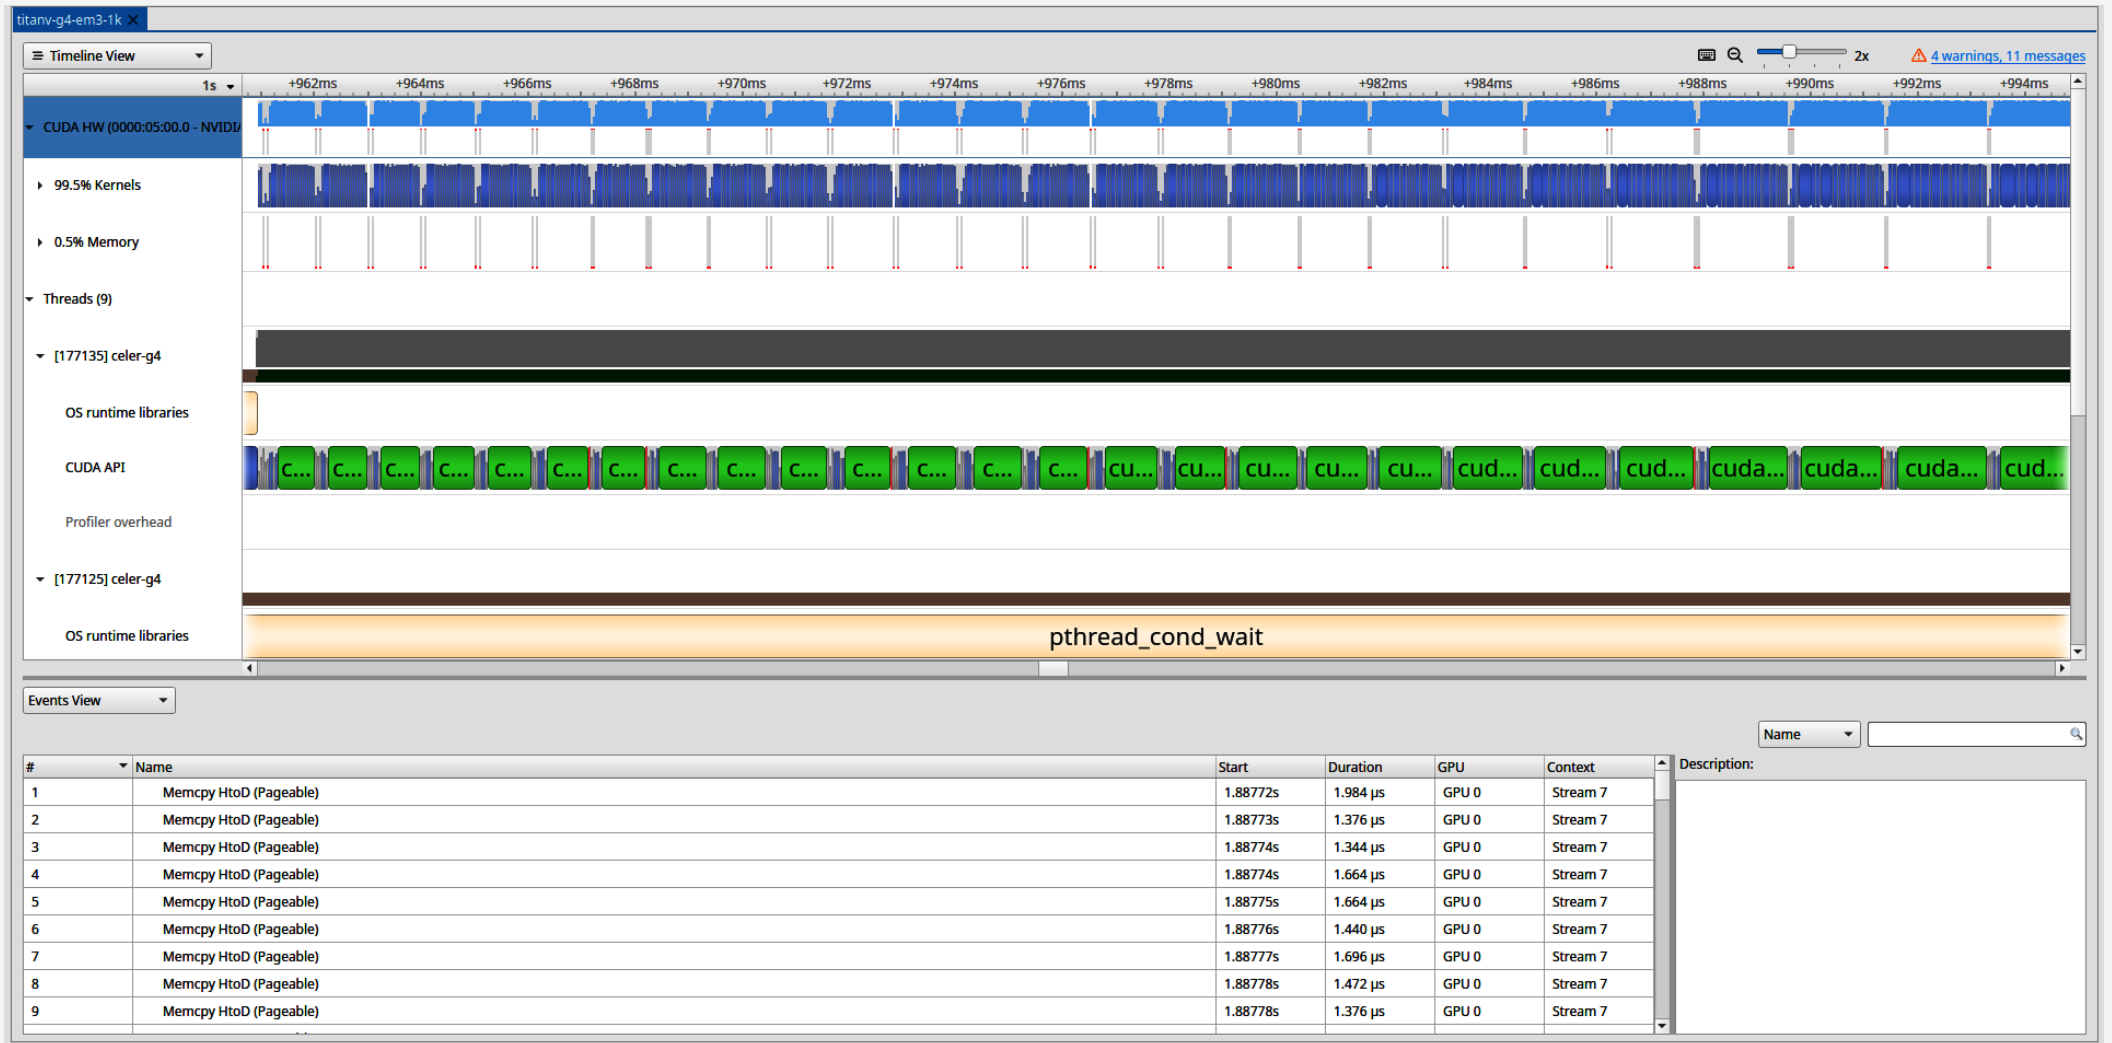
Timeline: TestEm3 (Titan V) GPU region



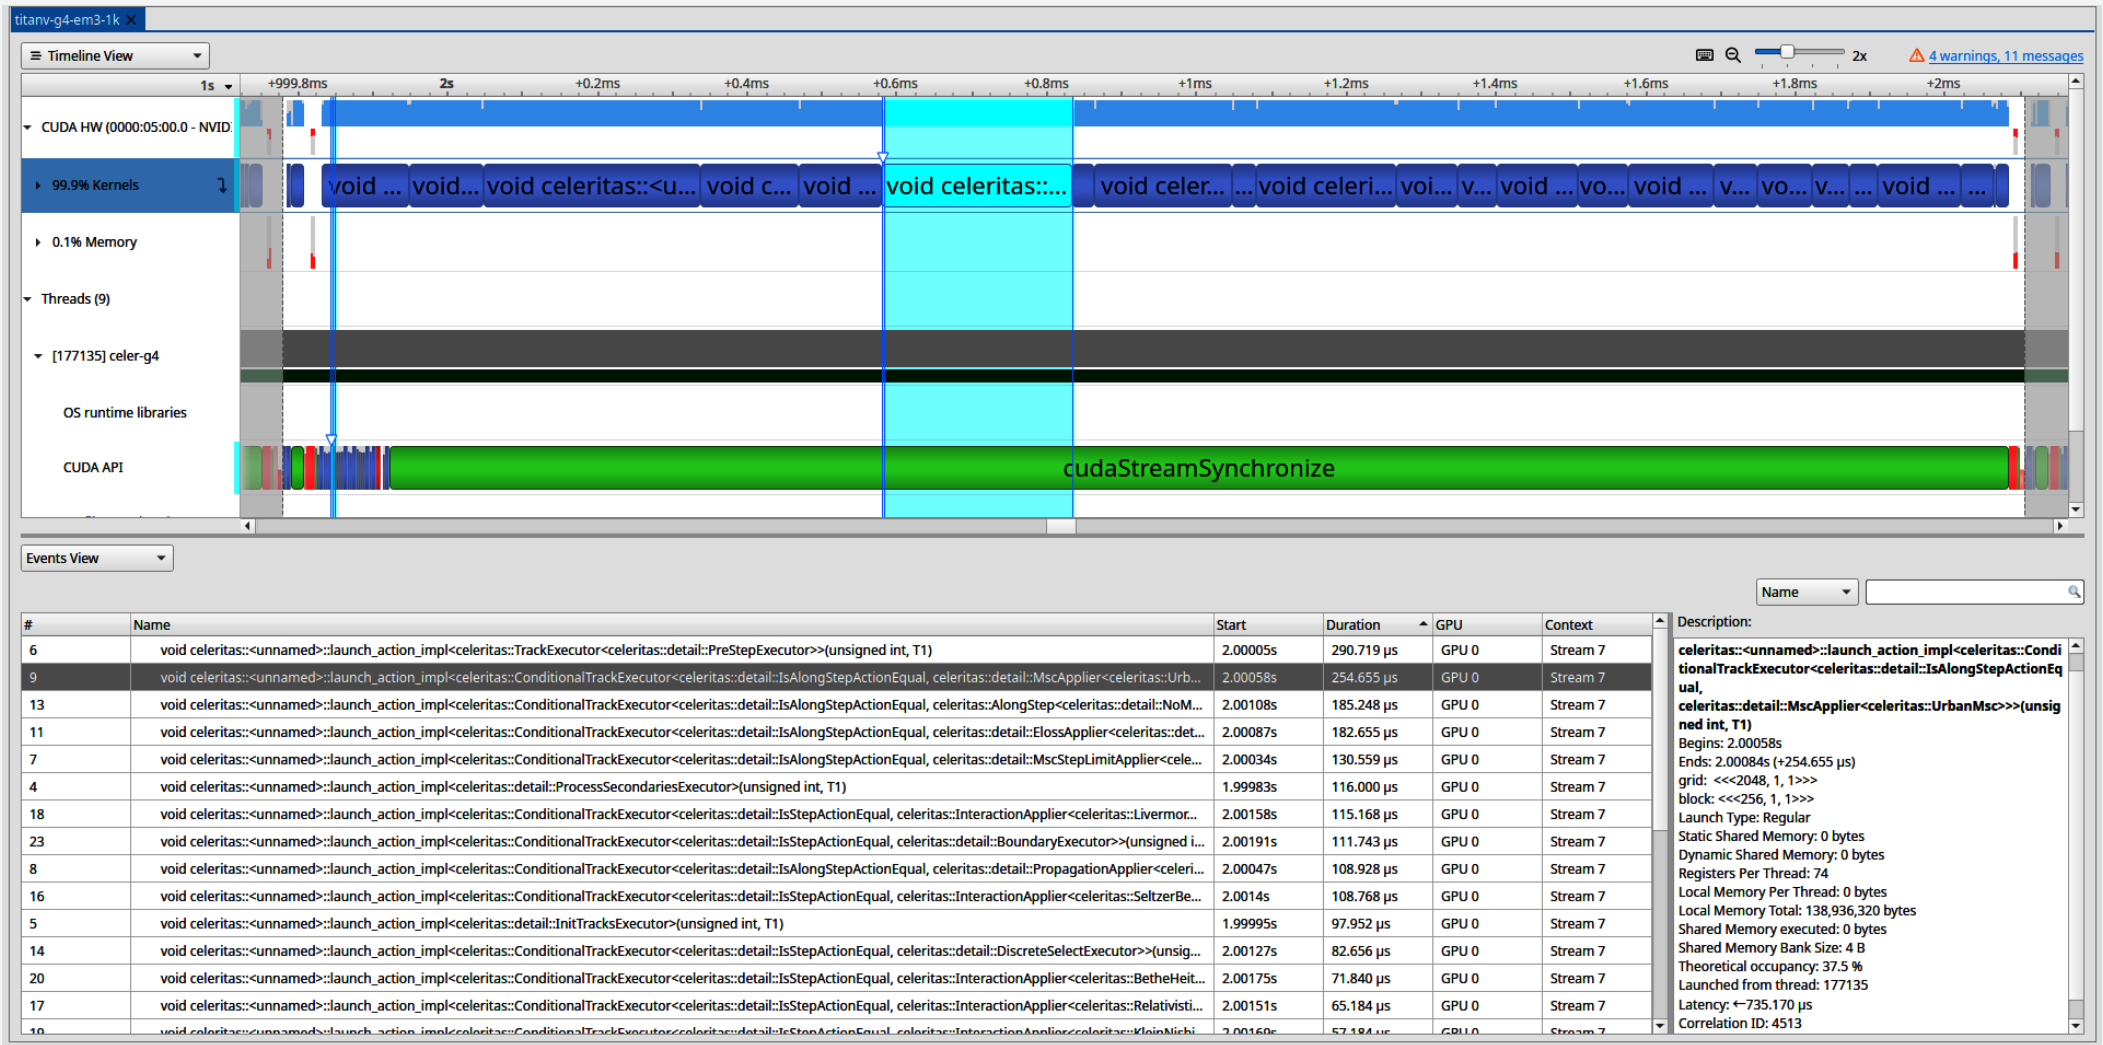
Timeline: TestEm3 (Titan V) GPU region



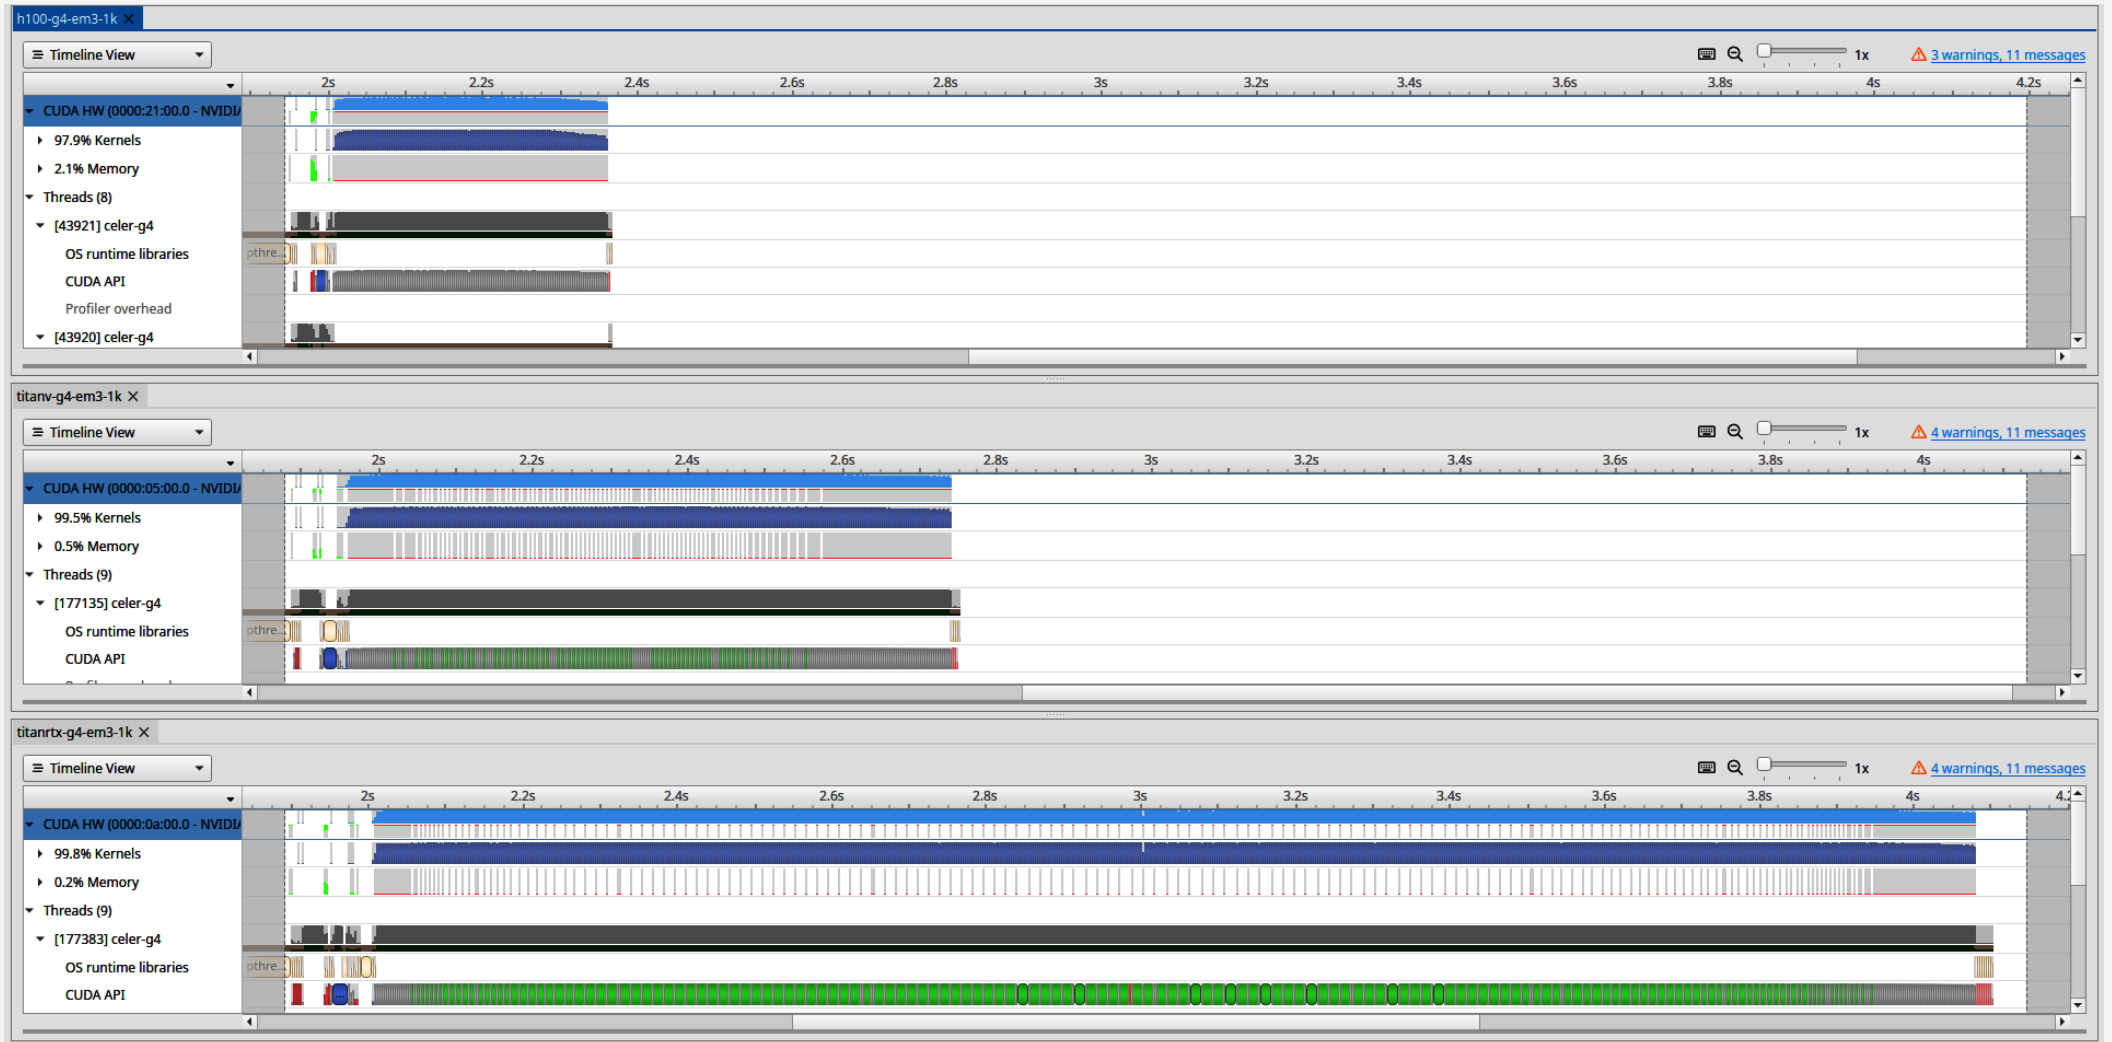
Timeline: TestEm3 (Titan V) GPU region



Timeline: TestEm3 (Titan V) GPU region



Timeline: H100, Titan V, Titan RTX

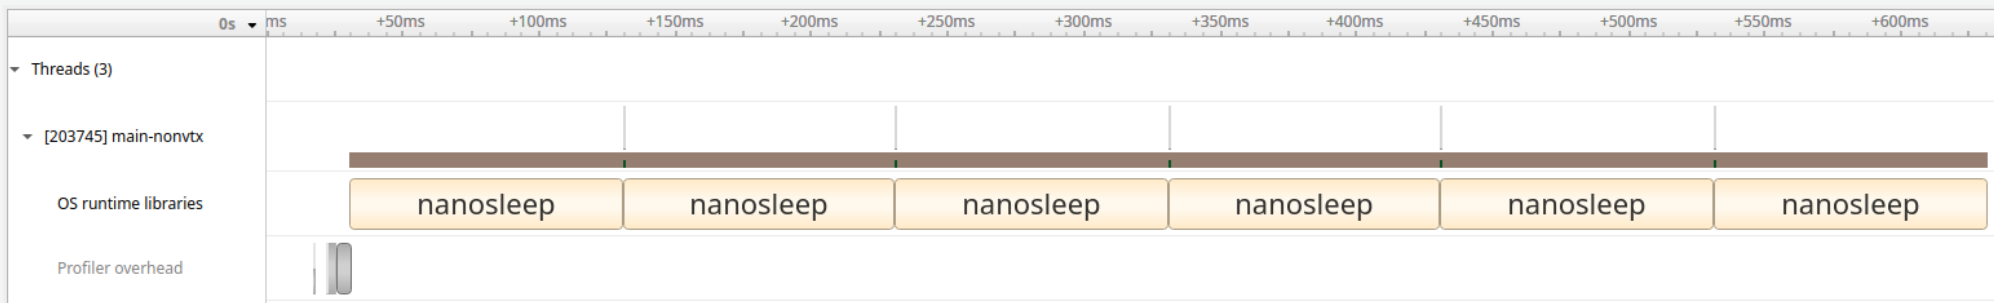


Code Annotation

Code Annotation

- NVIDIA Tools Extension (NVTX)
- AMD ROCTX

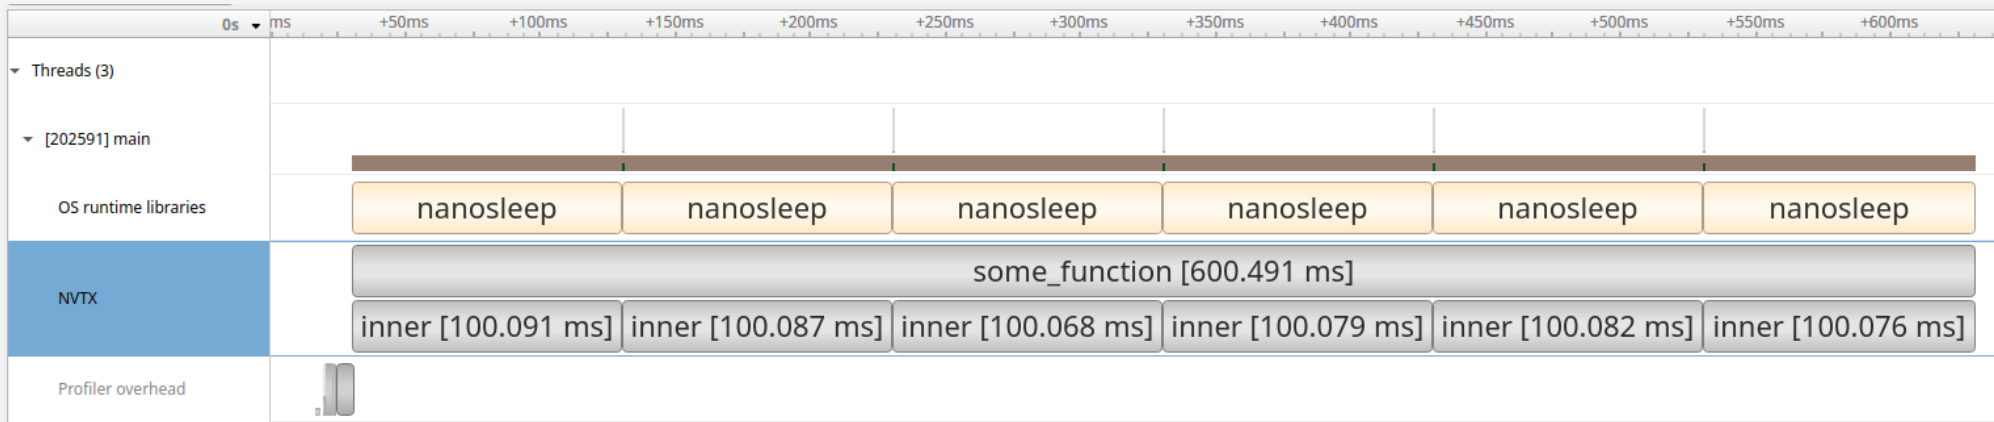
```
1
2 void some_function() {
3
4     for (int i = 0; i < 6; ++i) {
5
6         std::this_thread::sleep_for(std::chrono::milliseconds{100});
7
8     }
9
10 }
```



Code Annotation

- NVIDIA Tools Extension (NVTX)
- AMD ROCTX

```
1 #include <nvtx3/nvToolsExt.h>
2 void some_function() {
3     nvtxRangePush(__FUNCTION__);
4     for (int i = 0; i < 6; ++i) {
5         nvtxRangePush("inner")
6         std::this_thread::sleep_for(std::chrono::milliseconds{100});
7         nvtxRangePop();
8     }
9     nvtxRangePop();
10 }
```



Nsight Compute

Nsight Compute

- Detailed GPU performance metrics
- Compile with `-lineinfo` for line-level profiling
- Capture full metrics via `--set=full`
- Replays GPU kernels many times - significant runtime increase
- Reduce captured kernels via filtering, `-s`, `-c` etc.
- SM 70+ (Volta)

```
1 # All metrics, skip 64 kernels, capture 128.  
2 ncu --set=full -s 64 -c 128 -o metrics celer-g4 input.mac  
3 ncu-ui metrics.ncu-rep
```

- May require `--target-processes`
- Nvidia profiler counters require root or security mitigation disabling since 418.43 (2019-02-22). See [ERR_NVGPUCTRPERM](#).

Nsight Compute: TestEm3 (Titan V)

TestEM3 kernel: Summary

titany-512-g4-em3-1k-s2600-c200.ncu-rep

Page: Summary Result: 119 - 10096 - void celeritas::<unnamed>::launch_action_impl<celeritas::ConditionalTrackExecutor<celeritas::detail::IsAlongStepActionEqual, celeritas::detail::MscApplier<cele... Add Baseline Apply Rules Occupancy Calculator Copy as Image

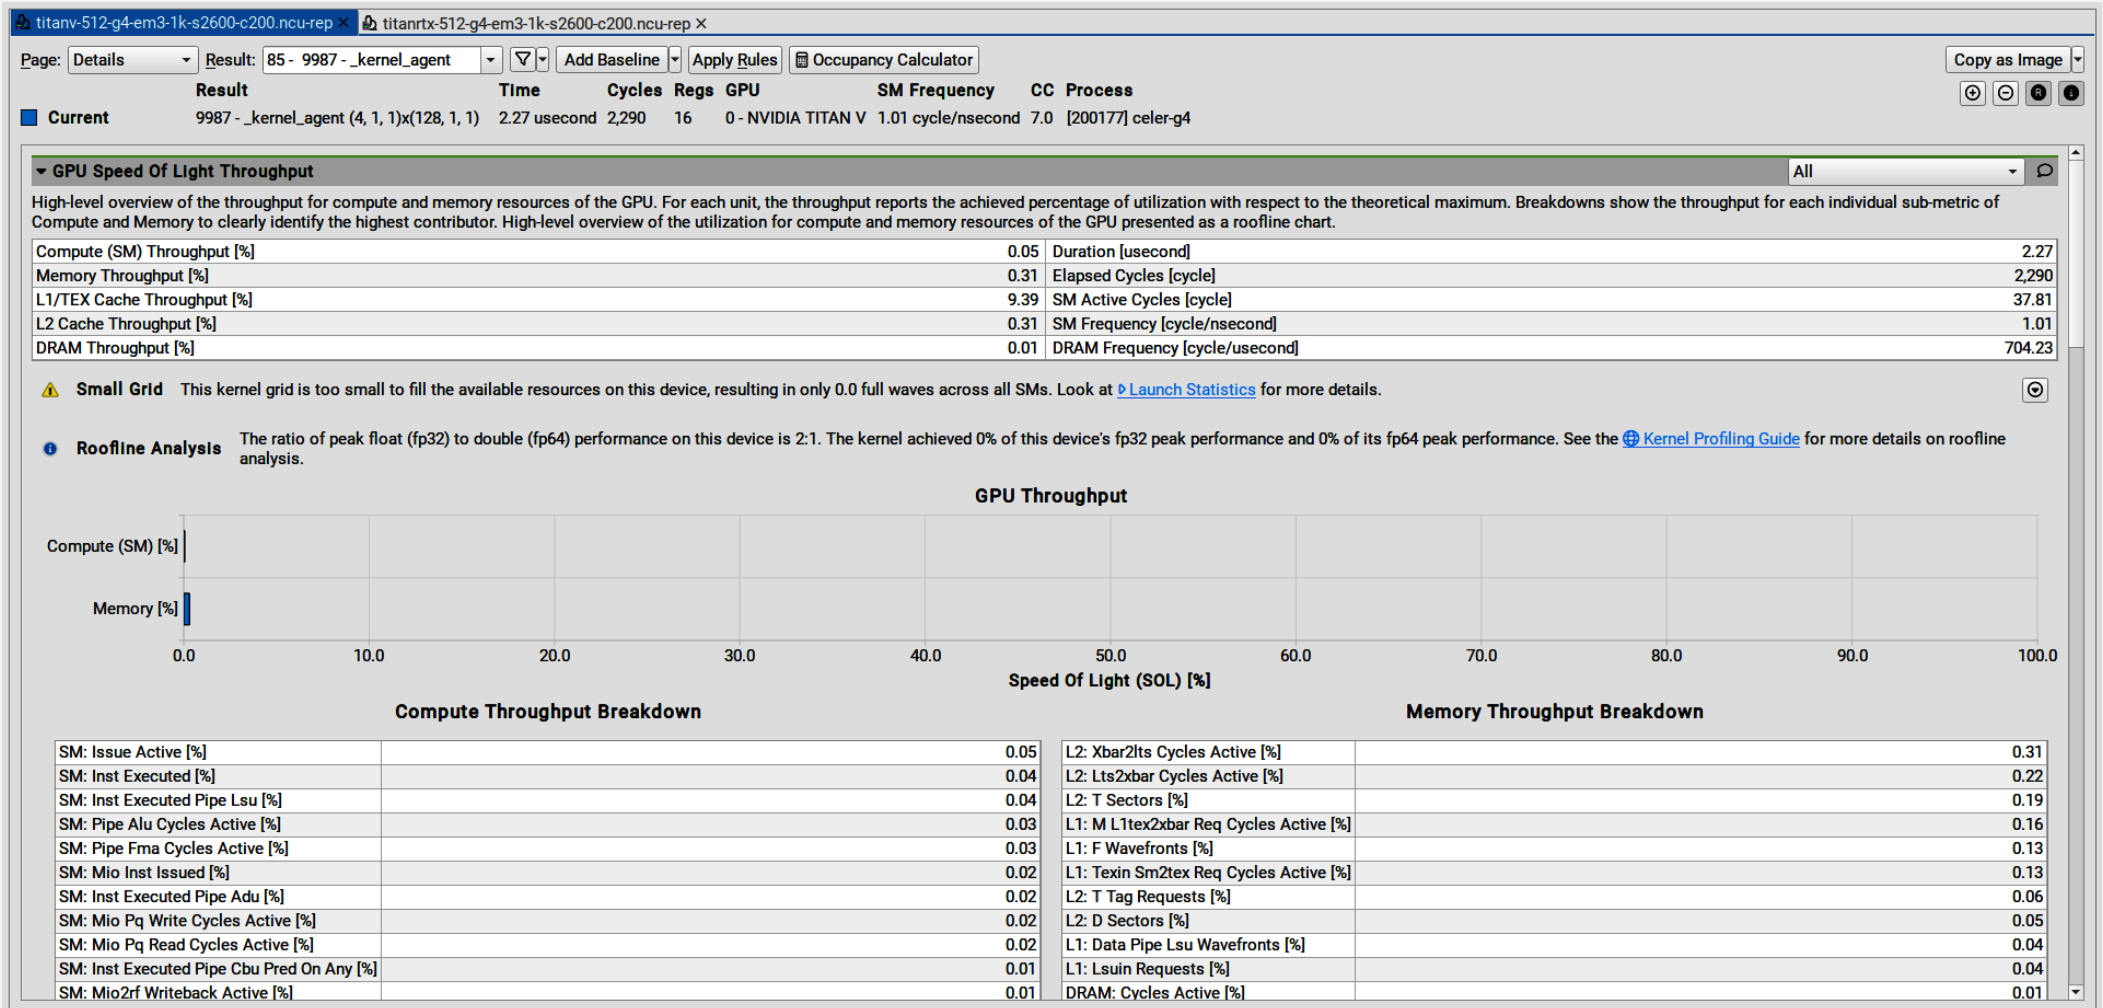
Result **Time** **Cycles** **Regs** **GPU** **SM Frequency** **CC** **Process**

Current 10096 - void celeritas::<unnamed>::launch_action_impl<celeritas::ConditionalTrackExecutor<celeritas::detail::IsAlongStepActionEqual, celeritas::detail::MscApplier<cele... 815.68 usecond 970,111 74 0 - NVIDIA TITAN V 1.19 cycle/nsecond 7.0 [200177] celer-g4

Use the column headers to sort the results in this report. Double-click a result to see detailed metrics.

ID	Issues Detected	Function N:	Demangled Name	Process	Device Name	Grid Size	Block Size	Cycles [cycle]	Duration [msecond]	Compute Throughput	Memory Throughput	# Registers
115	16	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	330, 1, 1	256, 1, 1	1,597,846	1.53	11.32	33.42	
40	16	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	329, 1, 1	256, 1, 1	1,726,823	1.45	10.42	30.59	
140	16	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	331, 1, 1	256, 1, 1	1,603,175	1.42	11.29	33.51	
165	16	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	330, 1, 1	256, 1, 1	1,604,421	1.35	11.15	33.42	
65	16	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	330, 1, 1	256, 1, 1	1,585,620	1.32	11.66	34.13	
190	16	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	331, 1, 1	256, 1, 1	1,462,385	1.32	12.25	36.81	
15	16	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	327, 1, 1	256, 1, 1	1,538,712	1.28	11.50	33.69	
90	16	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	330, 1, 1	256, 1, 1	1,626,690	1.25	11.15	32.81	
119	15	launch_...	void celeritas::<unnamed>::launch_action_1...	[200177] cele...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	970,111	0.82	30.31	33.37	
19	15	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	951,351	0.81	30.91	34.03	
194	15	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	967,609	0.81	30.39	33.44	
169	15	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	966,180	0.80	30.44	33.51	
44	15	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	970,301	0.80	30.32	33.28	
94	15	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	978,235	0.80	30.06	33.03	
144	15	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	964,551	0.79	30.49	33.54	
69	15	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	954,240	0.79	30.81	33.88	
166	13	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	718,456	0.61	29.89	51.21	
91	13	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	716,973	0.60	29.95	51.79	
116	13	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	713,324	0.60	30.09	51.77	
16	13	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	709,817	0.60	30.25	52.12	
191	13	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	716,988	0.60	29.95	51.86	
41	13	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	712,247	0.60	30.15	51.92	
66	13	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	712,081	0.60	30.16	51.96	
141	13	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	717,363	0.60	29.93	51.52	
48	14	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	608,635	0.51	16.85	52.23	
23	14	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	609,539	0.51	16.74	52.18	
73	14	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	603,288	0.51	16.92	52.80	
98	14	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	601,340	0.51	16.96	53.00	
148	14	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	604,666	0.51	16.78	52.58	
123	14	launch_a...	void celeritas::<unnamed>::launch_action_impl<c...	[200177] celer...	NVIDIA TITAN V	2048, 1, 1	256, 1, 1	604,396	0.51	16.57	52.49	

TestEM3 2685th kernel: Speed of Light



TestEM3 2719th kernel: Speed of Light

titany-512-g4-em3-1k-s2600-c200.ncu-rep

Page: Details Result: 119 - 10096 - void celeritas::< > Add Baseline Apply Rules Occupancy Calculator Copy as Image

Result Time Cycles Regs GPU SM Frequency CC Process
Current 10096 - void celeritas::<unnamed>::launch_action_impl<celeritas::ConditionalTrackExecutor<celeritas::detail::IsAlongStepActionEqual, celeritas::detail::... 815.68 usecond 970,111 74 0 - NVIDIA TITAN V 1.19 cycle/nsecond 7.0 [200177] celer-g4

GPU Speed Of Light Throughput

All

High-level overview of the throughput for compute and memory resources of the GPU. For each unit, the throughput reports the achieved percentage of utilization with respect to the theoretical maximum. Breakdowns show the throughput for each individual sub-metric of Compute and Memory to clearly identify the highest contributor. High-level overview of the utilization for compute and memory resources of the GPU presented as a roofline chart.

Metric	Value	Metric	Value
Compute (SM) Throughput [%]	30.31	Duration [usecond]	815.68
Memory Throughput [%]	33.37	Elapsed Cycles [cycle]	970,111
L1/TEX Cache Throughput [%]	30.50	SM Active Cycles [cycle]	938,333.76
L2 Cache Throughput [%]	26.49	SM Frequency [cycle/nsecond]	1.19
DRAM Throughput [%]	33.37	DRAM Frequency [cycle/usecond]	843.18

Latency Issue This kernel exhibits low compute throughput and memory bandwidth utilization relative to the peak performance of this device. Achieved compute throughput and/or memory bandwidth below 60.0% of peak typically indicate latency issues. Look at [Scheduler Statistics](#) and [Warp State Statistics](#) for potential reasons.

Roofline Analysis The ratio of peak float (fp32) to double (fp64) performance on this device is 2:1. The kernel achieved close to 0% of this device's fp32 peak performance and 6% of its fp64 peak performance. See the [Kernel Profiling Guide](#) for more details on roofline analysis.

GPU Throughput

Metric	Value
Compute (SM) [%]	30.31
Memory [%]	33.37

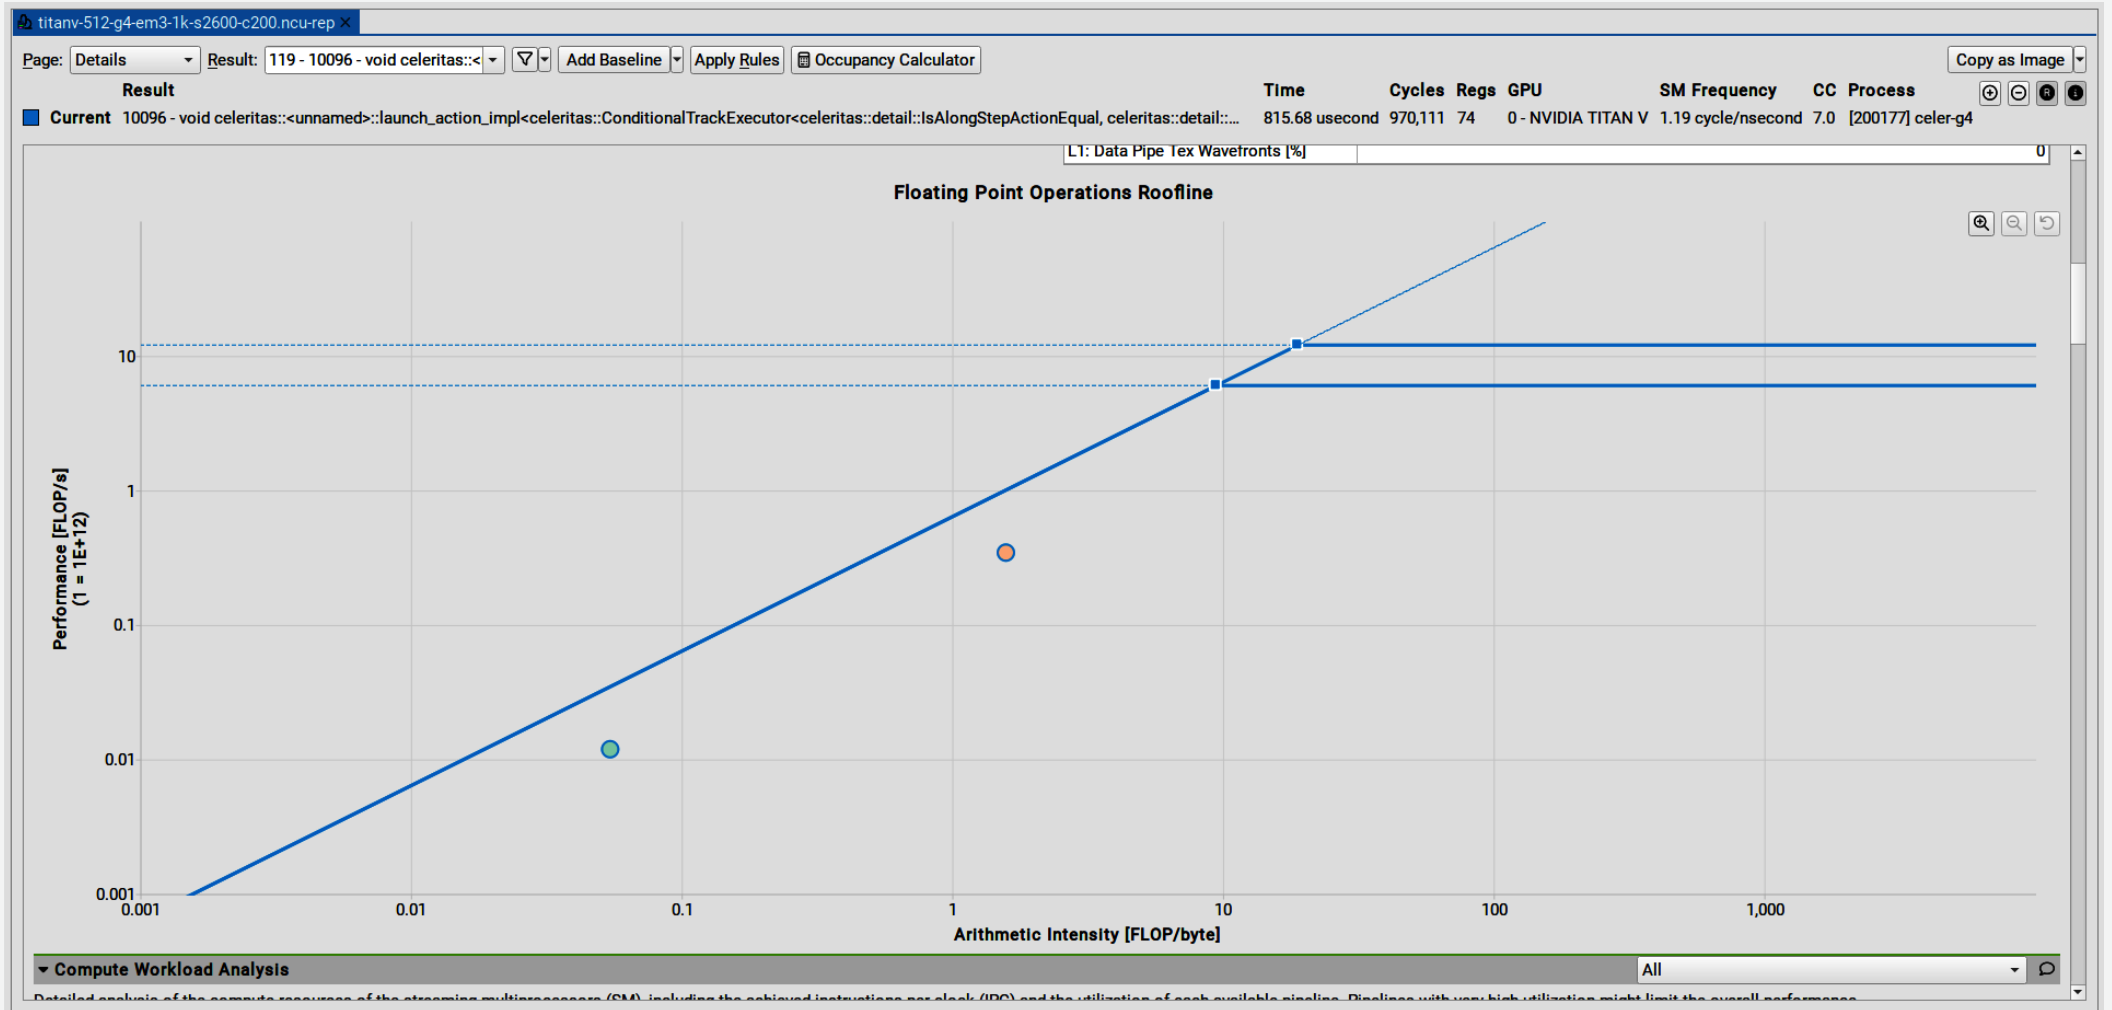
Compute Throughput Breakdown

SM: Pipe Shared Cycles Active [%]	30.31
SM: Pipe Fp64 Cycles Active [%]	30.31
SM: Issue Active [%]	22.05
SM: Inst Executed [%]	22.01
SM: Inst Executed Pipe Cbu Pred On Any [%]	12.59
SM: Pipe Alu Cycles Active [%]	9.59
SM: Pipe Fma Cycles Active [%]	7.47
SM: Inst Executed Pipe Lsu [%]	6.78
SM: Mio2rf Writeback Active [%]	4.83
SM: Inst Executed Pipe Xu [%]	4.26

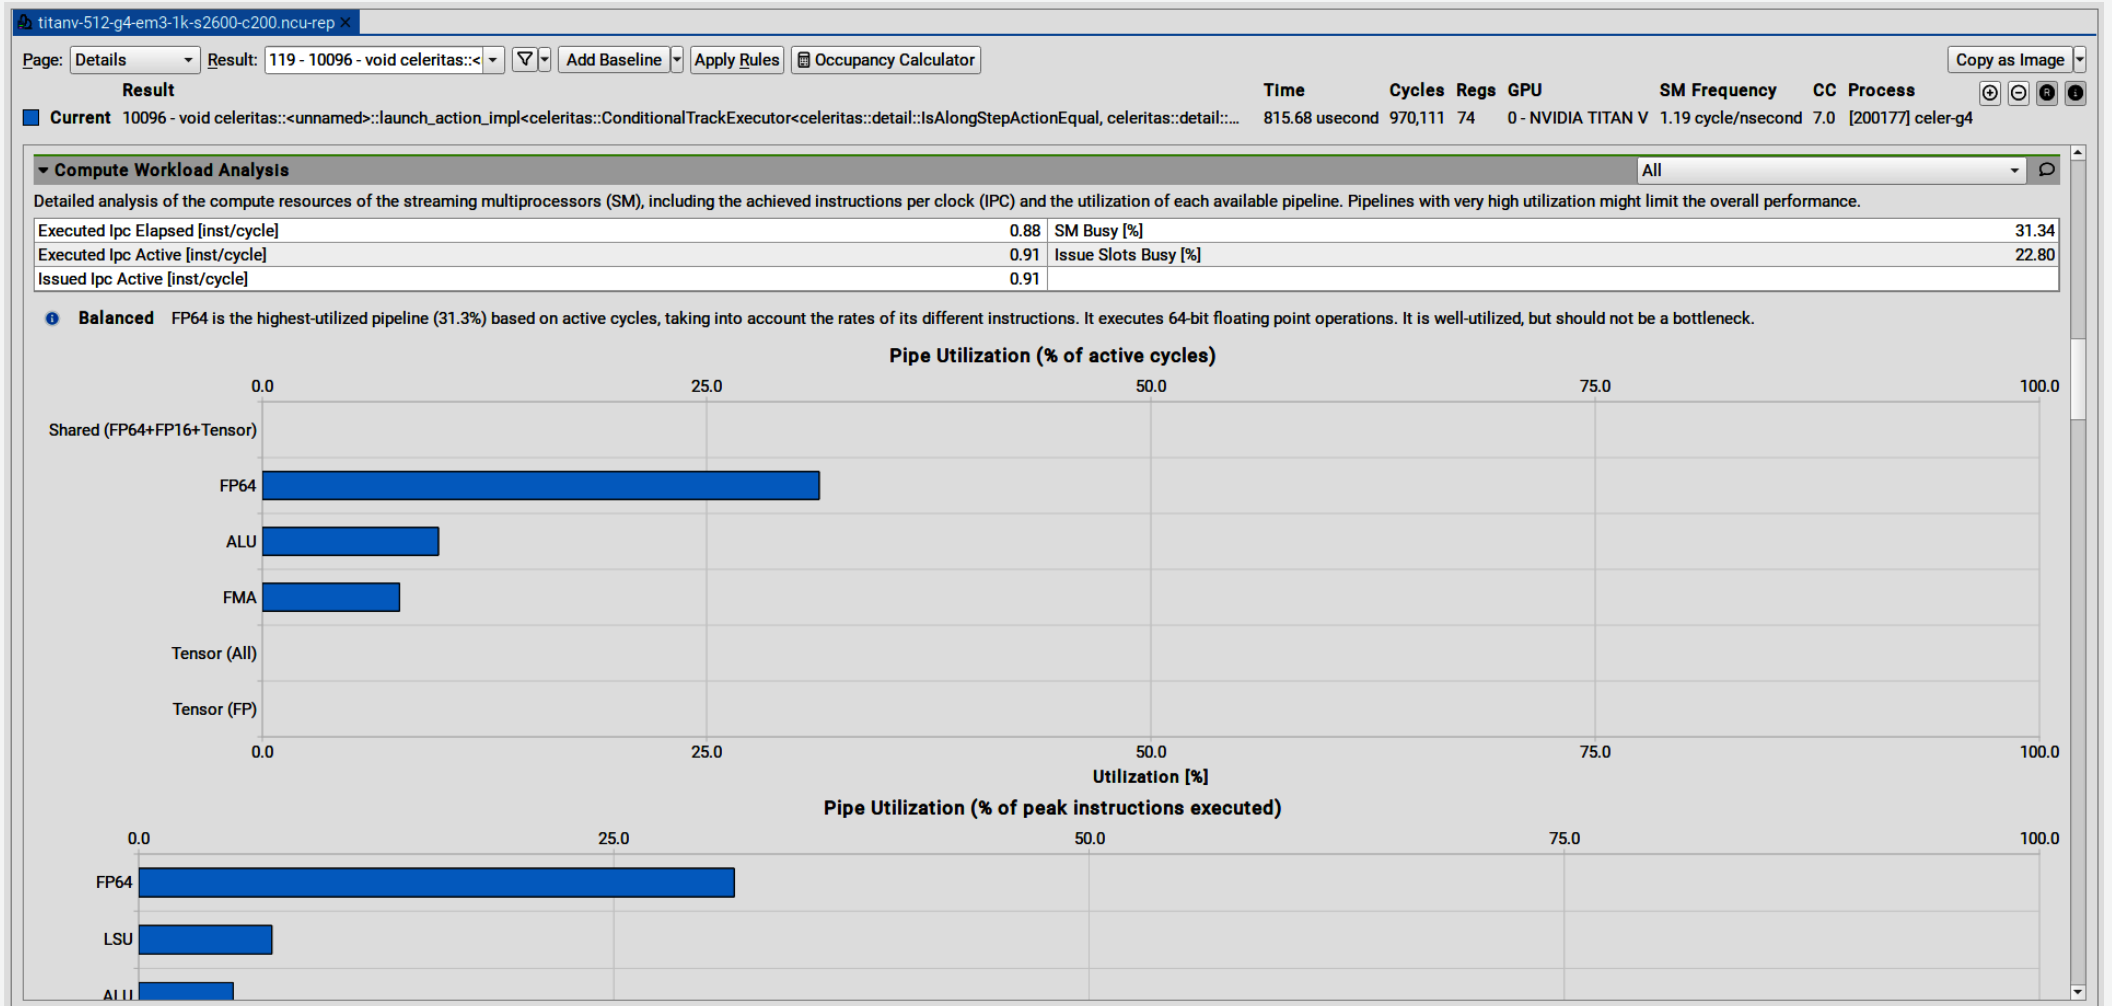
Memory Throughput Breakdown

DRAM: Cycles Active [%]	33.37
L2: Xbar2lts Cycles Active [%]	26.49
DRAM: Dram Sectors [%]	24.76
L2: T Sectors [%]	21.03
L1: M L1tex2xbar Req Cycles Active [%]	15.25
L1: Data Pipe Lsu Wavefronts [%]	15.16
L2: T Tag Requests [%]	13.15
L1: Lsu Writeback Active [%]	11.35
L2: D Sectors Fill Device [%]	9.61
L2: D Sectors [%]	8.34

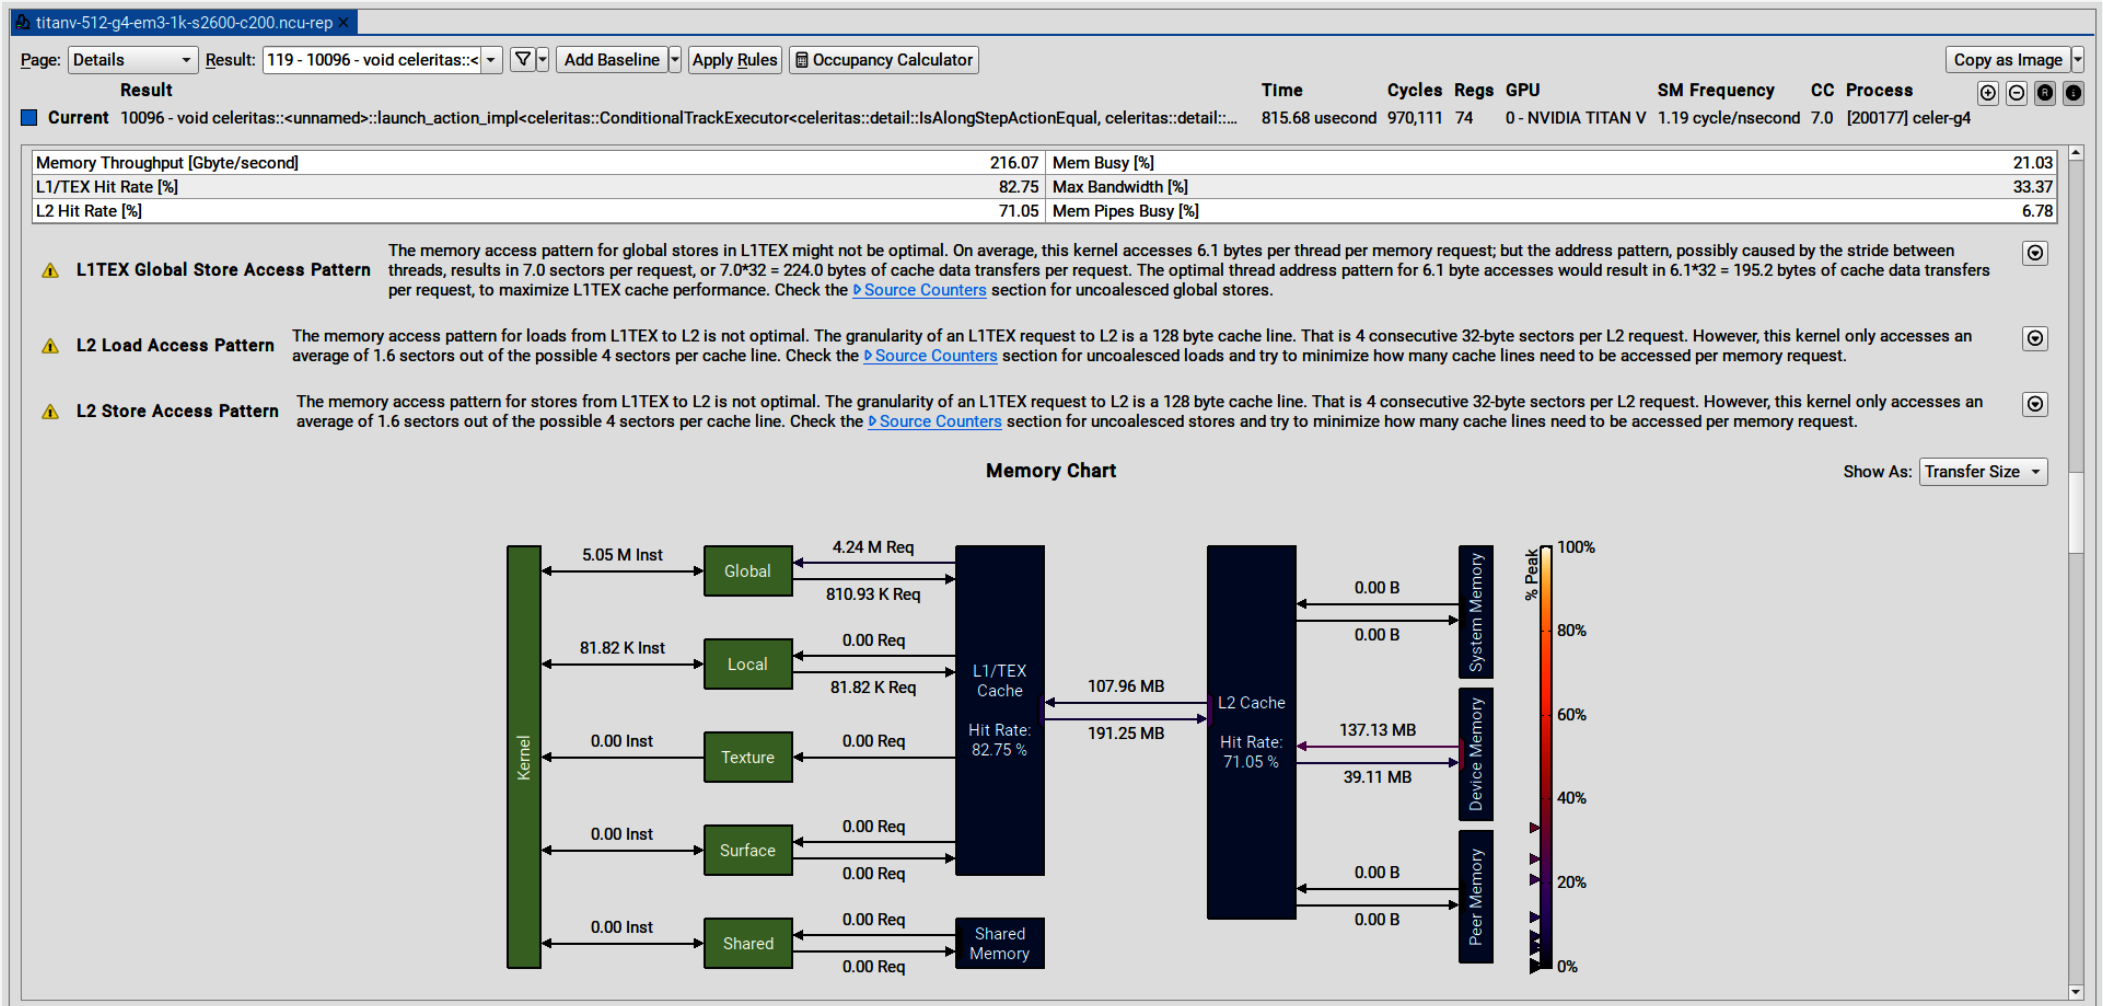
TestEM3 2719th kernel: Roofline



TestEM3 2719th kernel: Compute



TestEM3 2719th kernel: Memory



TestEM3 2719th kernel: Warp State

titany-512-g4-em3-1k-s2600-c200.ncu-rep

Page: Details Result: 119 - 10096 - void celeritas::< > Add Baseline Apply Rules Occupancy Calculator Copy as Image

Result **Time** **Cycles** **Regs** **GPU** **SM Frequency** **CC** **Process**

Current 10096 - void celeritas::<unnamed>::launch_action_impl<celeritas::ConditionalTrackExecutor<celeritas::detail::IsAlongStepActionEqual, celeritas::detail::... 815.68 usecond 970,111 74 0 - NVIDIA TITAN V 1.19 cycle/nsecond 7.0 [200177] celer-g4

Warp State Statistics

Analysis of the states in which all warps spent cycles during the kernel execution. The warp states describe a warp's readiness or inability to issue its next instruction. The warp cycles per instruction define the latency between two consecutive instructions. The higher the value, the more warp parallelism is required to hide this latency. For each warp state, the chart shows the average number of cycles spent in that state per issued instruction. Stalls are not always impacting the overall performance nor are they completely avoidable. Only focus on stall reasons if the schedulers fail to issue every cycle. When executing a kernel with mixed library and user code, these metrics show the combined values.

Warp Cycles Per Issued Instruction [cycle]	22.81	Avg. Active Threads Per Warp	8.18
Warp Cycles Per Executed Instruction [cycle]	22.85	Avg. Not Predicated Off Threads Per Warp	7.61

no_instruction On average, each warp of this kernel spends 10.7 cycles being stalled due to not having the next instruction fetched yet. This represents about 46.8% of the total average of 22.8 cycles between issuing two instructions. A high number of warps not having an instruction fetched is typical for very short kernels with less than one full wave of work in the grid. Excessively jumping across large blocks of assembly code can also lead to more warps stalled for this reason, if this causes misses in the instruction cache.

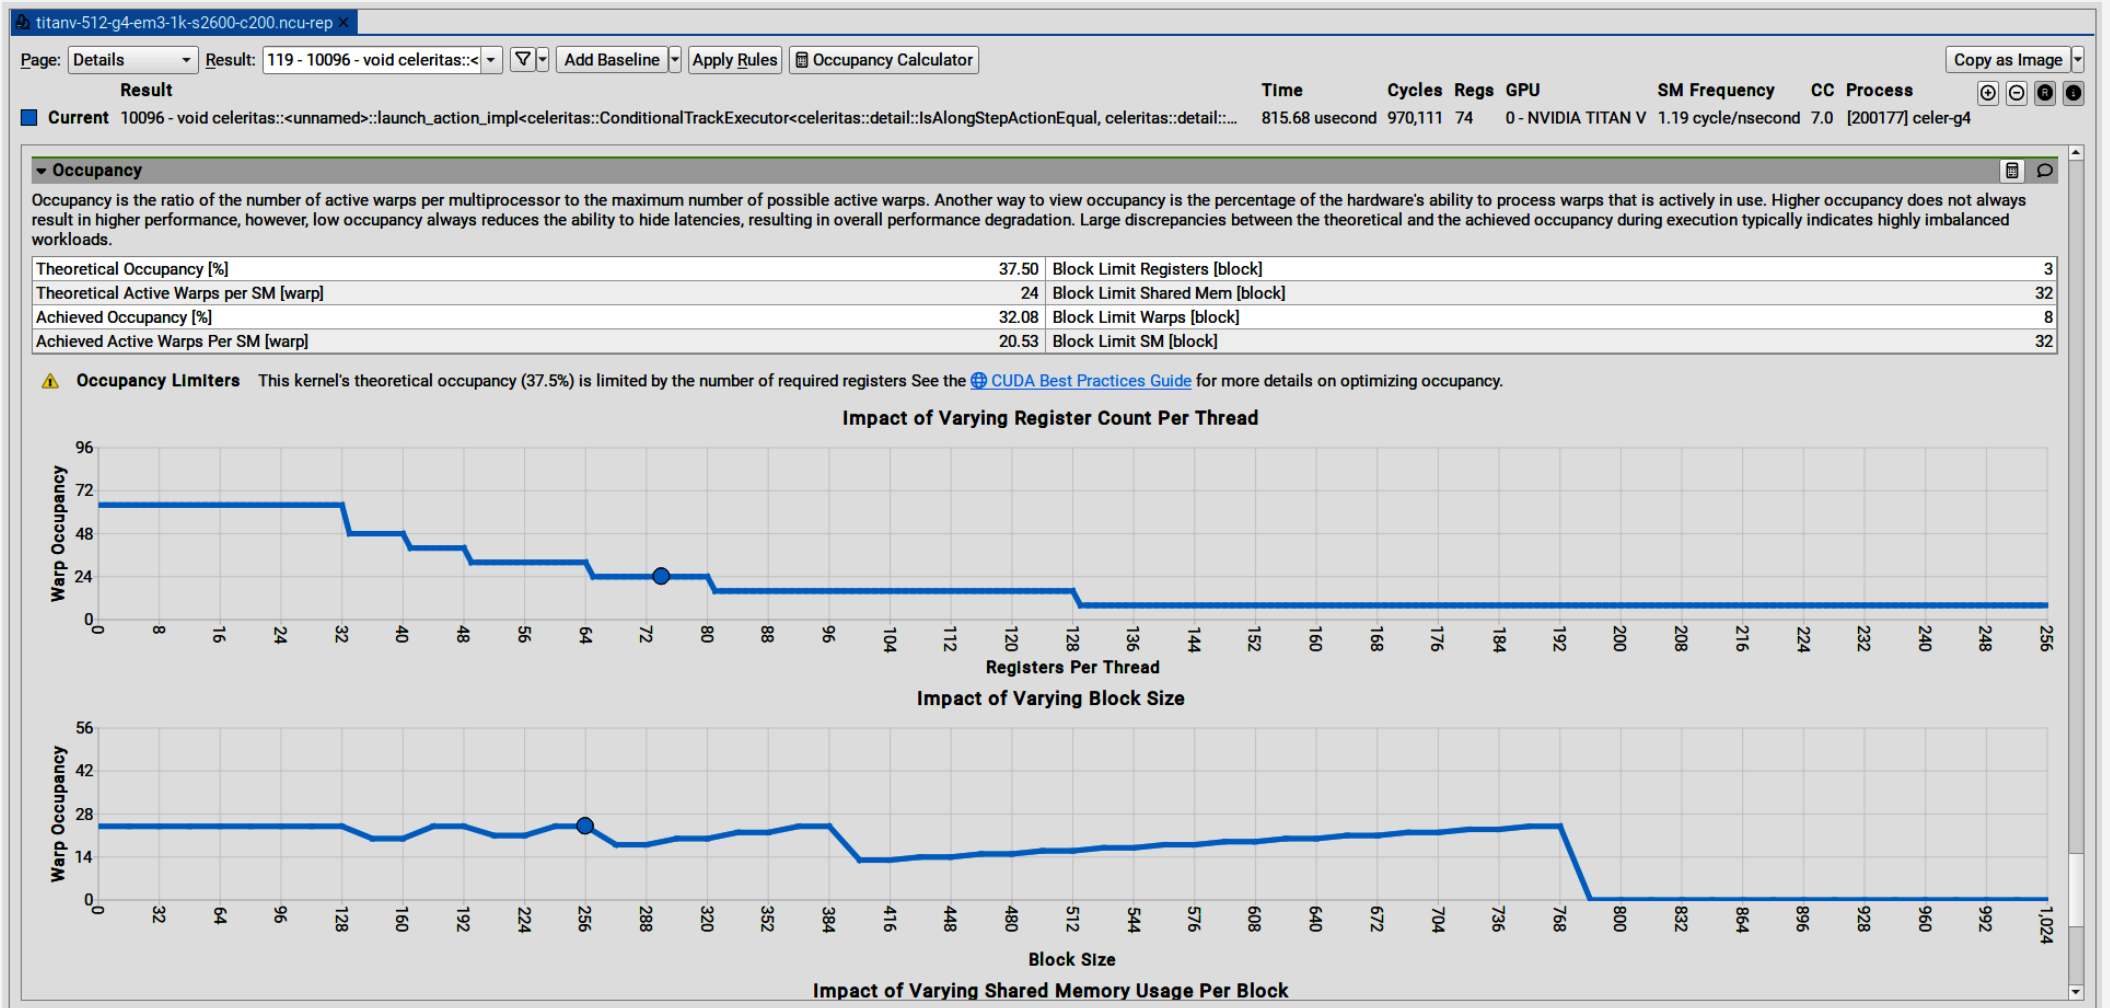
Warp Stall Check the [Source Counters](#) section for the top stall locations in your source based on sampling data. The [Kernel Profiling Guide](#) provides more details on each stall reason.

Thread Divergence Instructions are executed in warps, which are groups of 32 threads. Optimal instruction throughput is achieved if all 32 threads of a warp execute the same instruction. The chosen launch configuration, early thread completion, and divergent flow control can significantly lower the number of active threads in a warp per cycle. This kernel achieves an average of 8.2 threads being active per cycle. This is further reduced to 7.6 threads per warp due to predication. The compiler may use predication to avoid an actual branch. Instead, all instructions are scheduled, but a per-thread condition code or predicate controls which threads execute the instructions. Try to avoid different execution paths within a warp when possible. In addition, ensure your kernel makes use of Independent Thread Scheduling, which allows a warp to reconverge after a data-dependent conditional block by explicitly calling `__syncwarp()`.

Warp State (All Cycles)

Warp State	Approximate Cycles
Stall No Instruction	10.7
Stall Long Scoreboard	6.5
Stall Wait	4.5
Selected	2.5
Stall Math Pipe Throttle	2.0
Stall LG Throttle	2.0
Stall Branch Resolving	1.5

TestEM3 2719th kernel: Occupancy



TestEM3 2719th kernel: Source Counters

titany-512-g4-em3-1k-s2600-c200.ncu-rep

Page: Details Result: 119 - 10096 - void celeritas::<unnamed>::launch_action_impl<celeritas::ConditionalTrackExecutor<celeritas::detail::IsAlongStepActionEqual, celeritas::detail::...>> Add Baseline Apply Rules Occupancy Calculator Copy as Image

Result Time Cycles Regs GPU SM Frequency CC Process
Current 10096 - void celeritas::<unnamed>::launch_action_impl<celeritas::ConditionalTrackExecutor<celeritas::detail::IsAlongStepActionEqual, celeritas::detail::...>> 815.68 usecond 970,111 74 0 - NVIDIA TITAN V 1.19 cycle/nsecond 7.0 [200177] celer-g4

Source Counters All

Source metrics, including branch efficiency and sampled warp stall reasons. Warp Stall Sampling metrics are periodically sampled over the kernel runtime. They indicate when warps were stalled and couldn't be scheduled. See the documentation for a description of all stall reasons. Only focus on stalls if the schedulers fail to issue every cycle.

Branch Instructions [inst]	6,302,731	Branch Efficiency [%]	92.95
Branch Instructions Ratio [%]	0.09	Avg. Divergent Branches	772.08

Uncoalesced Global Accesses This kernel has uncoalesced global accesses resulting in a total of 11555387 excessive sectors (58% of the total 19911571 sectors). Check the L2 Theoretical Sectors Global Excessive table for the primary source locations. The [CUDA Programming Guide](#) had additional information on reducing uncoalesced device memory accesses.

L2 Theoretical Sectors Global Excessive

Location	Value	Value (%)
Collection.hh:452 (0x7f2097898570 in void celeritas::<unnamed>::launch_action_impl<c...	615,468	3
UrbanMsc.hh:202 (0x7f2097896ba0 in void celeritas::<unnamed>::launch_action_impl<c...	256,612	2
SurfaceFunctors.hh:83 (0x7f209789c280 in void celeritas::<unnamed>::launch_action_im...	233,231	2
SurfaceAction.hh:144 (0x7f2097898a00 in void celeritas::<unnamed>::launch_action_imp...	233,231	2
SurfaceAction.hh:144 (0x7f20978986e0 in void celeritas::<unnamed>::launch_action_imp...	233,231	2

Warp Stall Sampling (All Cycles)

Location	Value	Value (%)
MscApplier.hh:53 (0x7f20978954f0 in void celeritas::<unnamed>::launch_action_impl<c...	465	1
UrbanMscHelper.hh:122 (0x7f2097895600 in void celeritas::<unnamed>::launch_action_i...	426	1
Algorithms.hh:106 (0x7f209789e330 in void celeritas::<unnamed>::launch_action_impl<c...	383	1
UrbanMsc.hh:270 (0x7f20978af5d0 in void celeritas::<unnamed>::launch_action_impl<c...	352	1
ArrayUtils.hh:237 (0x7f20978ab520 in void celeritas::<unnamed>::launch_action_impl<c...	318	1

Warp Stall Sampling (Not-issued Cycles)

Location	Value	Value (%)
MscApplier.hh:53 (0x7f20978954f0 in void celeritas::<unnamed>::launch_action_impl<c...	405	1
UrbanMscHelper.hh:122 (0x7f2097895600 in void celeritas::<unnamed>::launch_action_i...	370	1
UrbanMsc.hh:270 (0x7f20978af5d0 in void celeritas::<unnamed>::launch_action_impl<c...	303	1
Algorithms.hh:106 (0x7f209789e330 in void celeritas::<unnamed>::launch_action_impl<c...	275	1
ArrayUtils.hh:237 (0x7f20978ab520 in void celeritas::<unnamed>::launch_action_impl<c...	269	1

Most Instructions Executed

Location	Value	Value (%)
Algorithmsimpl.hh:71 (0x7f209789e3e0 in void celeritas::<unnamed>::launch_action_im...	114,565	0

- Must be compiled with `-lineinfo`

TestEM3 2719th kernel: Source

Page: Source Result: 119 - 10096 - void celeritas::<unnamed>::launch_action_impl<celeritas::ConditionalTrackExecutor<celeritas::detail::IsAlongStepActionEqual, celeritas::detail::MscApplier<cele... 815.68 usecond 970,111 74 0 - NVIDIA TITAN V 1.19 cycle/nsecond 7.0 [200177] celer-g4

View: Source and SASS

Source: TrackExecutor.hh Find...

Navigation: L2 Theoretical Sectors Global Excessive

Source: void celeritas::<unnamed>::launch_action Find...

Navigation: L2 Theoretical Sectors Global Excessive

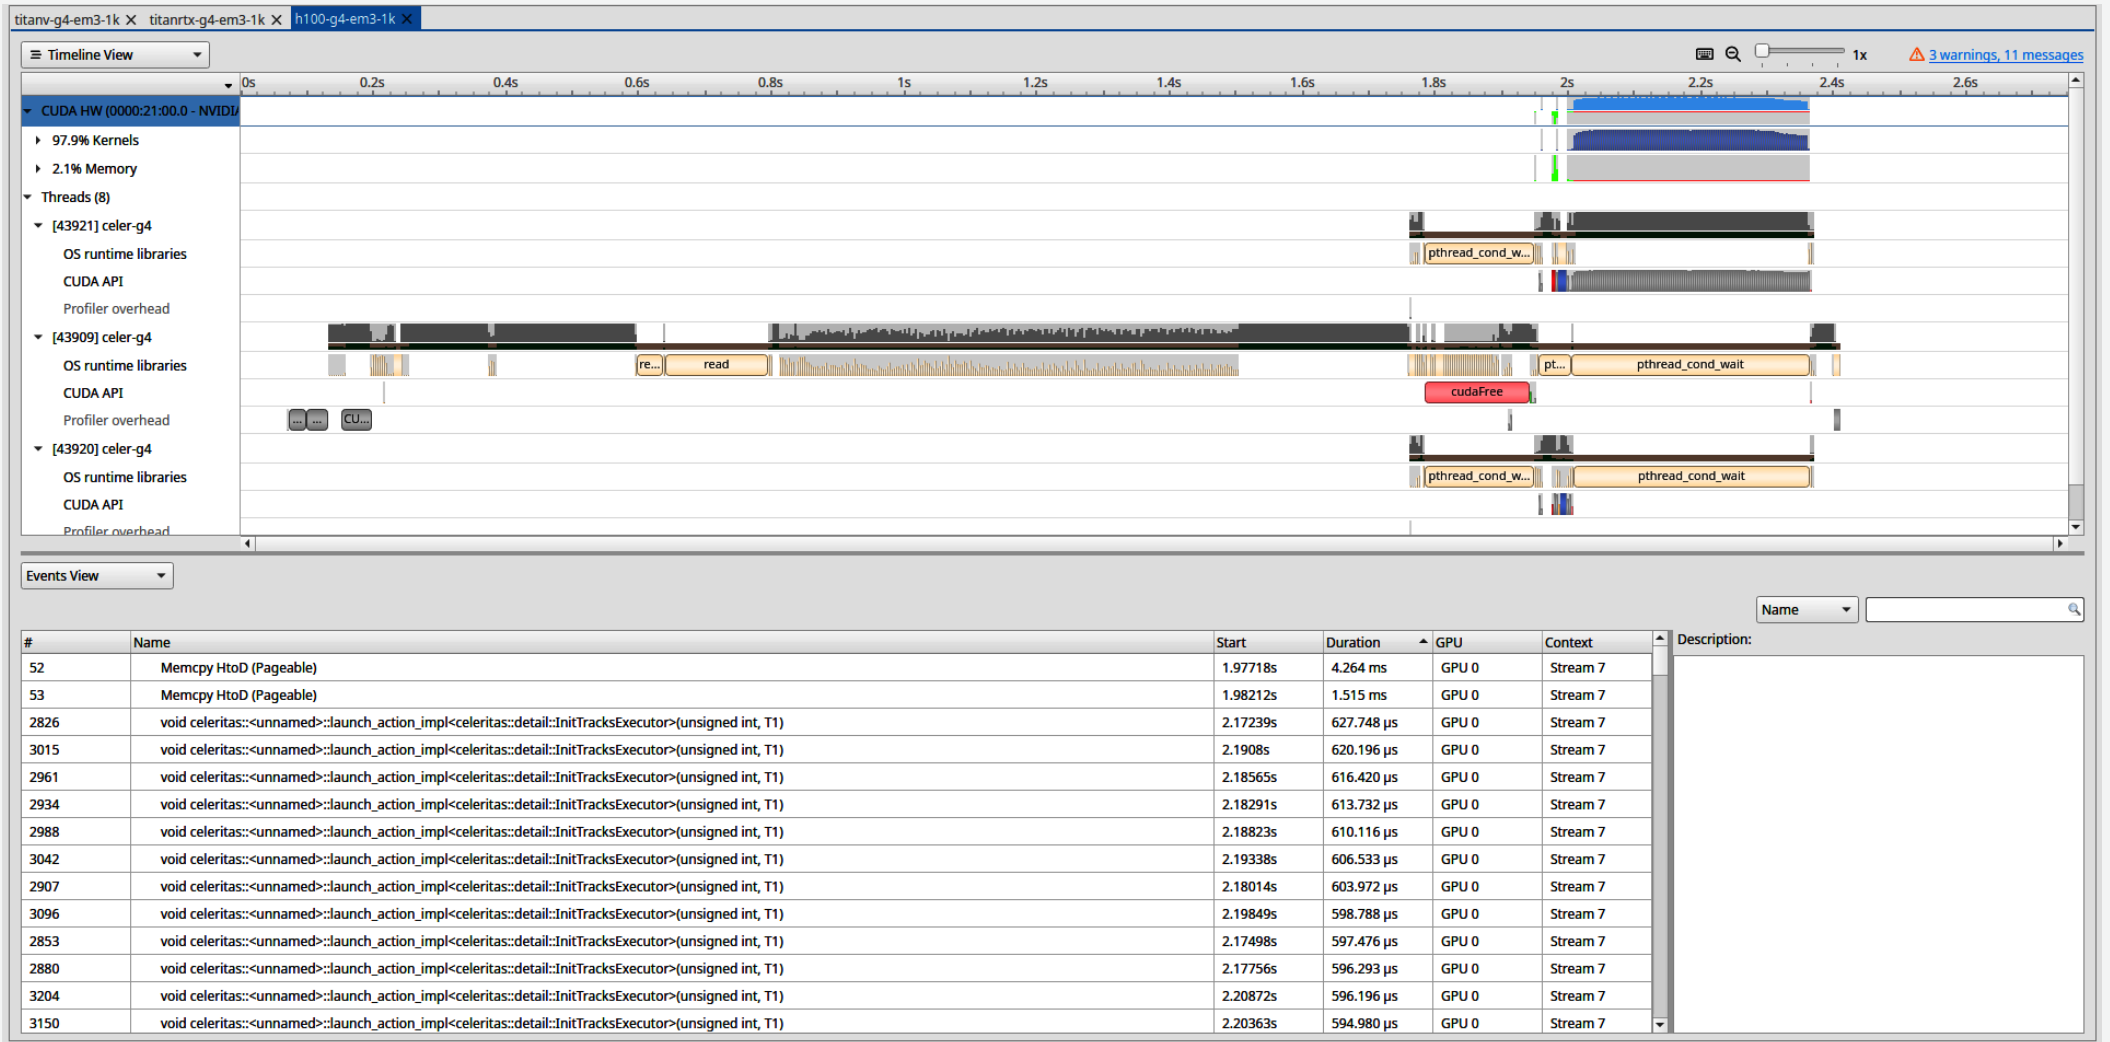
# Source	Address Space	Access Operation	# Address	Source	Address Space	Access Operation	Access Size	Local Sector	Global Excessive
1... , state_{state}			796 00007f20 978984b0	IMAD.MOV.U32 R25, RZ, RZ, 0x7f...					
1... , applies_{celeritas::forward<C>(applies)}			797 00007f20 978984c0	ISETP.NE.AND P0, PT, R37, R0, P...					
1... , execute_impl_{celeritas::forward<Ts>(args) ...}			798 00007f20 978984d0	@!P0 BRA 0x7f209789c630					
1... {			799 00007f20 978984e0	LDG.E.64.SYS R10, [R12+-0x240]	obal	Load	64		
1... }			800 00007f20 978984f0	LDG.E.64.SYS R32, [R14+-0x344]	obal	Load	64		
1... CELER_FUNCTION void operator()(ThreadId thread)			801 00007f20 97898500	LD.E.SYS R9, [R16]	obal	Load	32		
1... {			802 00007f20 97898510	LDG.E.64.SYS R30, [R14+-0x2f4]	obal	Load	64		
1... CELER_EXPECT(thread < state_→size());			803 00007f20 97898520	IMAD.WIDE.U32 R34, R35, 0x18, P...					
1... CoreTrackView const track(*params_, *state_, thread);			804 00007f20 97898530	IMAD.MOV.U32 R11, RZ, RZ, RZ					
1... if (!applies_(track.make_sim_view()))	Global(2)	Load	805 00007f20 97898540	IMAD.WIDE.U32 R36, R37, 0x4, R3...					
1... {			806 00007f20 97898550	IMAD.WIDE.U32 R32, R0, 0x4, R32...					
1... return;			807 00007f20 97898560	IMAD.IADD R35, R35, 0x1, R11					
1... }			808 00007f20 97898570	LD.E.SYS R43, [R36]	obal	Load	32	615,468	
1... return execute_impl_(track);	Global	Load	809 00007f20 97898580	IADD3 R39, R9, R43, RZ					
1... }			810 00007f20 97898590	IADD3 R38, P0, R30, R39, RZ					
1... private:			811 00007f20 978985a0	IMAD.X R39, RZ, RZ, R31, P0					
1... ParamsPtr params_;			812 00007f20 978985b0	LD.E.U8.SYS R38, [R38]	obal	Load	8	34,870	
1... StatePtr state_;			813 00007f20 978985c0	BMOV.32.CLEAR RZ, B2					
1... C applies_;			814 00007f20 978985d0	BSSY B2, 0x7f209789c570					
1... detail::TrackExecutorImpl<Ts...> execute_impl_;			815 00007f20 978985e0	BMOV.32.CLEAR RZ, B1					
1... }			816 00007f20 978985f0	BSSY B1, 0x7f209789c250					
1... }			817 00007f20 97898600	CS2R R10, SRZ					
1... }			818 00007f20 97898610	ISETP.GT.AND P0, PT, R38, 0x4,					
1... }			819 00007f20 97898620	BRB 0x7f209789c310					

- Must be compiled with `-lineinfo`

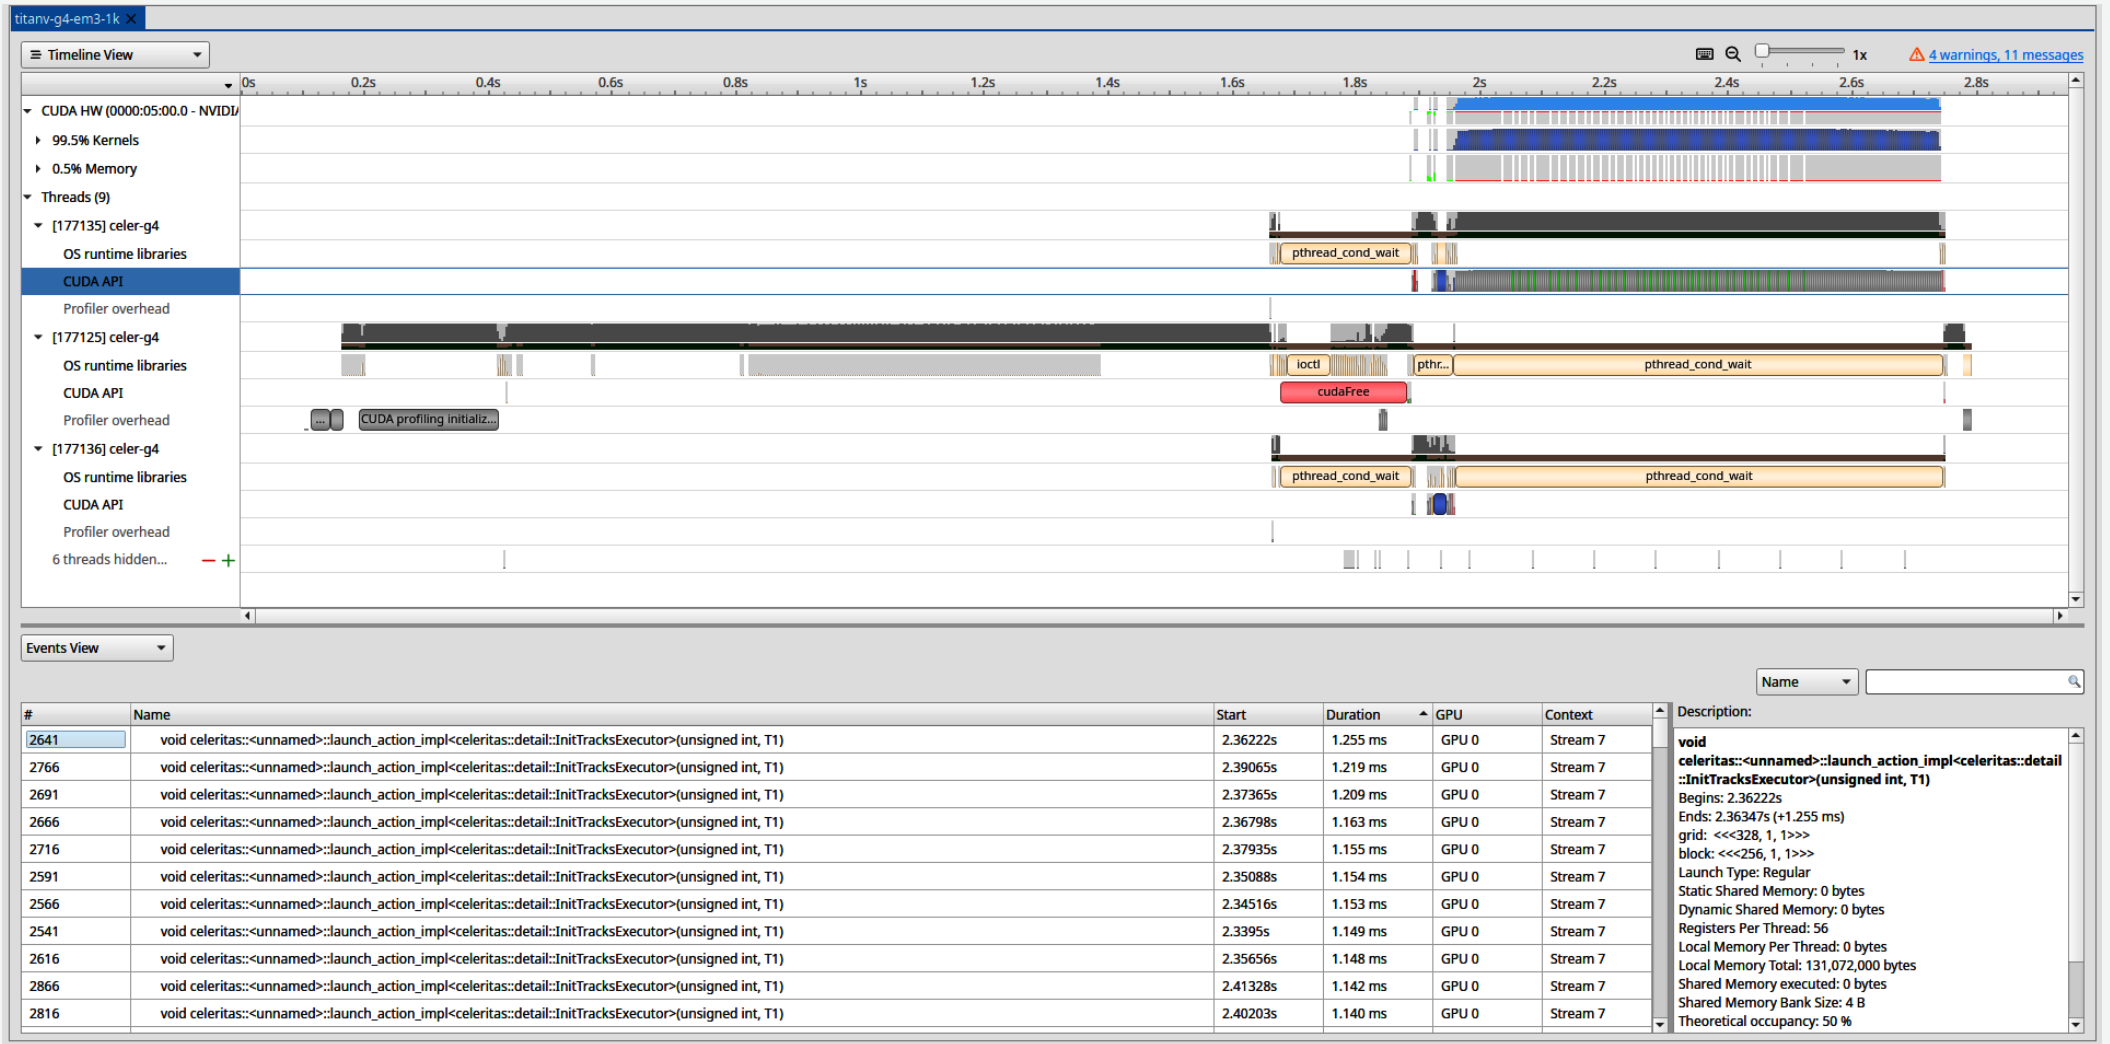
Profile your code

Additional Slides

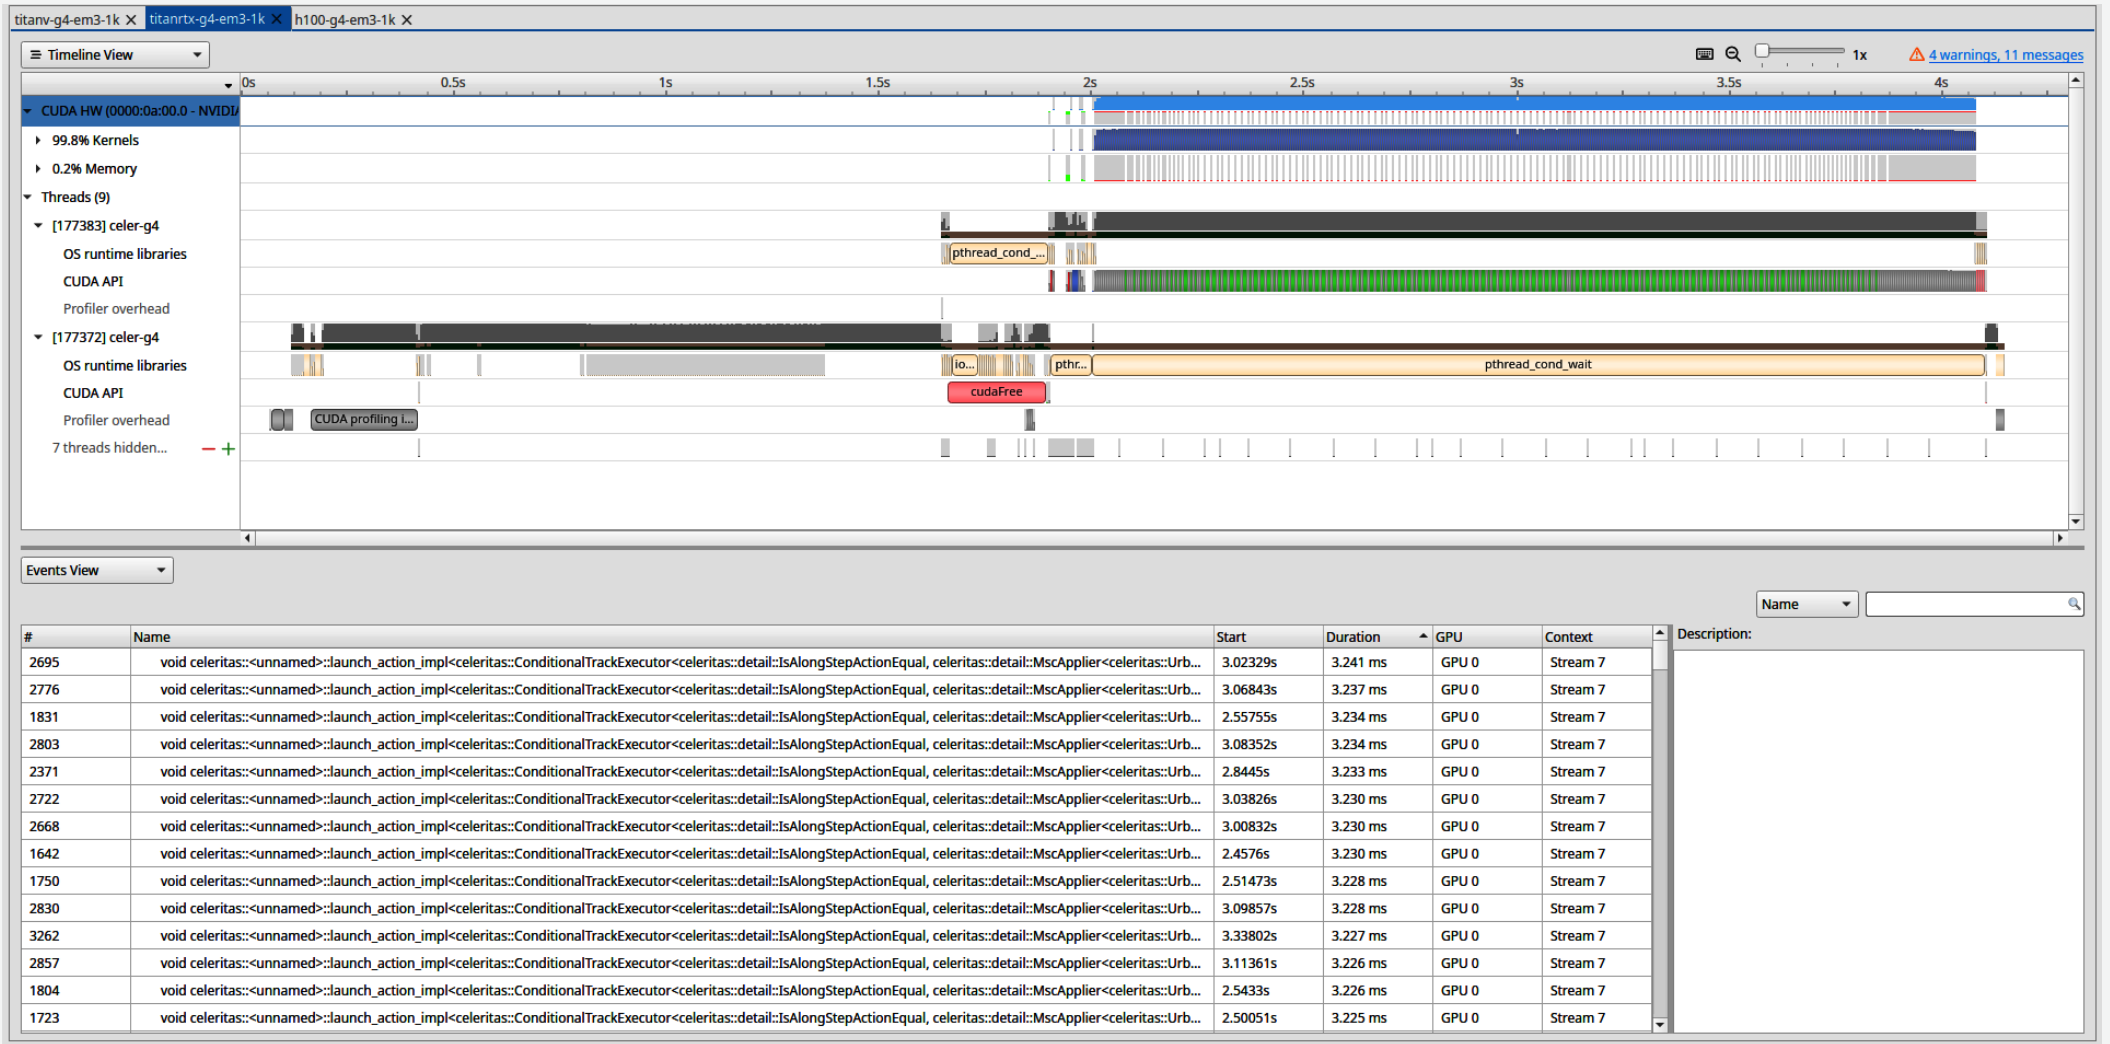
TestEm3 H100



TestEm3 Titan V



TestEm3 Titan RTX



TestEM3 Titan RTX 2719th kernel: speed

titantv-512-g4-em3-1k-s2600-c200.ncu-rep X
titanrtx-512-g4-em3-1k-s2600-c200.ncu-rep X

Page: Details
Result: 119 - 10096 - void celeritas::< >
Add Baseline
Apply Rules
Occupancy Calculator
Copy as Image

Result
Time
Cycles
Regs
GPU
SM Frequency
CC
Process

Current
10096 - void celeritas::<unnamed>::launch_action_impl<celeritas::ConditionalTrackExecutor<celeritas::detail::IsAlongStepActionEqual, celeritas::detail...
4.63 msecond
6,320,717
74
0 - NVIDIA TITAN RTX
1.36 cycle/nsecond
7.5
[200752] celer-g4

GPU Speed Of Light Throughput

High-level overview of the throughput for compute and memory resources of the GPU. For each unit, the throughput reports the achieved percentage of utilization with respect to the theoretical maximum. Breakdowns show the throughput for each individual sub-metric of Compute and Memory to clearly identify the highest contributor. High-level overview of the utilization for compute and memory resources of the GPU presented as a roofline chart.

Compute (SM) Throughput [%]	84.61	Duration [msecond]	4.63
Memory Throughput [%]	10.83	Elapsed Cycles [cycle]	6,320,717
L1/TEX Cache Throughput [%]	5.25	SM Active Cycles [cycle]	6,159,249.22
L2 Cache Throughput [%]	10.83	SM Frequency [cycle/nsecond]	1.36
DRAM Throughput [%]	6.22	DRAM Frequency [cycle/nsecond]	6.50

High Throughput The kernel is utilizing greater than 80.0% of the available compute or memory performance of the device. To further improve performance, work will likely need to be shifted from the most utilized to another unit. Start by analyzing workloads in the [Compute Workload Analysis](#) section.

FP64/32 Utilization The ratio of peak float (fp32) to double (fp64) performance on this device is 32:1. The kernel achieved close to 0% of this device's fp32 peak performance and 15% of its fp64 peak performance. If [Compute Workload Analysis](#) determines that this kernel is fp64 bound, consider using 32-bit precision floating point operations to improve its performance. See the [Kernel Profiling Guide](#) for more details on roofline analysis.

FP64/32 Utilization The achieved fp64 performance is 69% lower than the fp64 pipeline utilization. Check the [Instruction Statistics](#) section to see if using fused instructions can benefit this kernel.

GPU Throughput

Compute Throughput Breakdown		Memory Throughput Breakdown	
SM: Pipe Fp64 Cycles Active [%]	84.61	L2: T Sectors [%]	10.83
SM: Mio Pq Write Cycles Active [%]	22.24	L2: Lts2xbar Cycles Active [%]	9.71
SM: Mio Inst Issued [%]	4.30	DRAM: Cycles Active [%]	6.22
SM: Issue Active [%]	3.77	L2: Xbar2lts Cycles Active [%]	4.49
SM: Inst Executed [%]	3.76	DRAM: Dram Sectors [%]	3.91
SM: Inst Executed Pipe Lsu [%]	2.32	L2: T Tag Requests [%]	3.45
SM: Pipe Alu Cycles Active [%]	1.64	L2: D Sectors [%]	3.24
SM: Inst Executed Pipe Cbu Pred On Any [%]	1.62	L1: Data Pipe Lsu Wavefronts [%]	2.63

Spack Compute Capability

- Spack only accepts a single `cuda_arch` value
- Requires a full dependency rebuild (~90 mins)

Arch	Variant
Volta	<code>variants: +cuda cuda_arch=70 cxxstd=17</code>
Ampere	<code>variants: +cuda cuda_arch=80 cxxstd=17</code>
Hopper	<code>variants: +cuda cuda_arch=90 cxxstd=17</code>

Dockerfile nsys

nvidia/cuda:11.8.0-devel-ubuntu22.04 does not include nsys

Nsys 2022.4.2 (CUDA 11.8.0):

```
1 # Install nsys for profiling. ncu is included
2 RUN if [ "$DOCKERFILE_DISTRO" = "ubuntu" ] ; then \
3     apt-get -yqq update \
4     && apt-get -yqq install --no-install-recommends nsight-systems-2022.4.2 \
5     && rm -rf /var/lib/apt/lists/* ; \
6 fi
```

- Note: nsys and ncu will be removed from the `-ci` containers, which are smaller for bandwidth reasons.

Docker to Apptainer / Singularity

- apptainer/singularity build can convert docker files to apptainer images
- from a registry via `apptainer build img.sif docker://registry/image:tag`
- locally via deameon `apptainer build img.sif docker-deamon:registry/image:tag`
- locally via docker archive files
- https://apptainer.org/docs/user/main/docker_and_oci.html

```
1 # Build the appropriate container
2 cd celeritas/scripts/Docker
3 # Build the cuda Docker container, sm_70. Wait ~90 minutes.
4 ./build.sh cuda
5 # If the image hasn't been pushed to a registry, apptainer requires a local path
6 rm -f docker-temp.tar && docker save $(docker images --format="{{.Repository}}")
7 # Convert to an apptainer container in the working dir
8 apptainer build -F celeritas-dev-jammy-cuda11.sif docker-archive:image.tar
```

Docker to Apptainer / Singularity

- Docker and Apptainer have different defaults when executing images
 - Default directory bindings
 - environment variable mapping
 - in-container user
 - entrypoints
- Likely need to run with various flags to achieve similar behaviour

Docker to Apptainer / Singularity

```
# ephemeral, does not bind home dir by default
docker run --rm -ti --gpus all -v ./src celeritas/dev-jammy-cuda11:2023-06-19
# apptainer, runs as the current user, with the calling users env vars and default
bindings
apptainer run --nv --bind ./:/celeritas-project celeritas-dev-jammy-cuda11-2023-
06-19.sif
# TUoS HPC - this is not perfect
apptainer run --nv --bind ./:/celeritas-project
/mnt/parscratch/users/ac1phey/celeritas-dev-jammy-cuda11-sm90.sif
```

lineinfo

Add `-lineinfo` to

```
mkdir build-lineinfo && cd build-lineinfo
cmake .. -DCMAKE_CUDA_ARCHITECTURES=70 -DCMAKE_BUILD_TYPE=Release -
DCMAKE_CUDA_FLAGS_RELEASE="-O3 -DNDEBUG -lineinfo" -DCELERITAS_DEBUG=OFF
cmake --build . -j `nproc`
```

