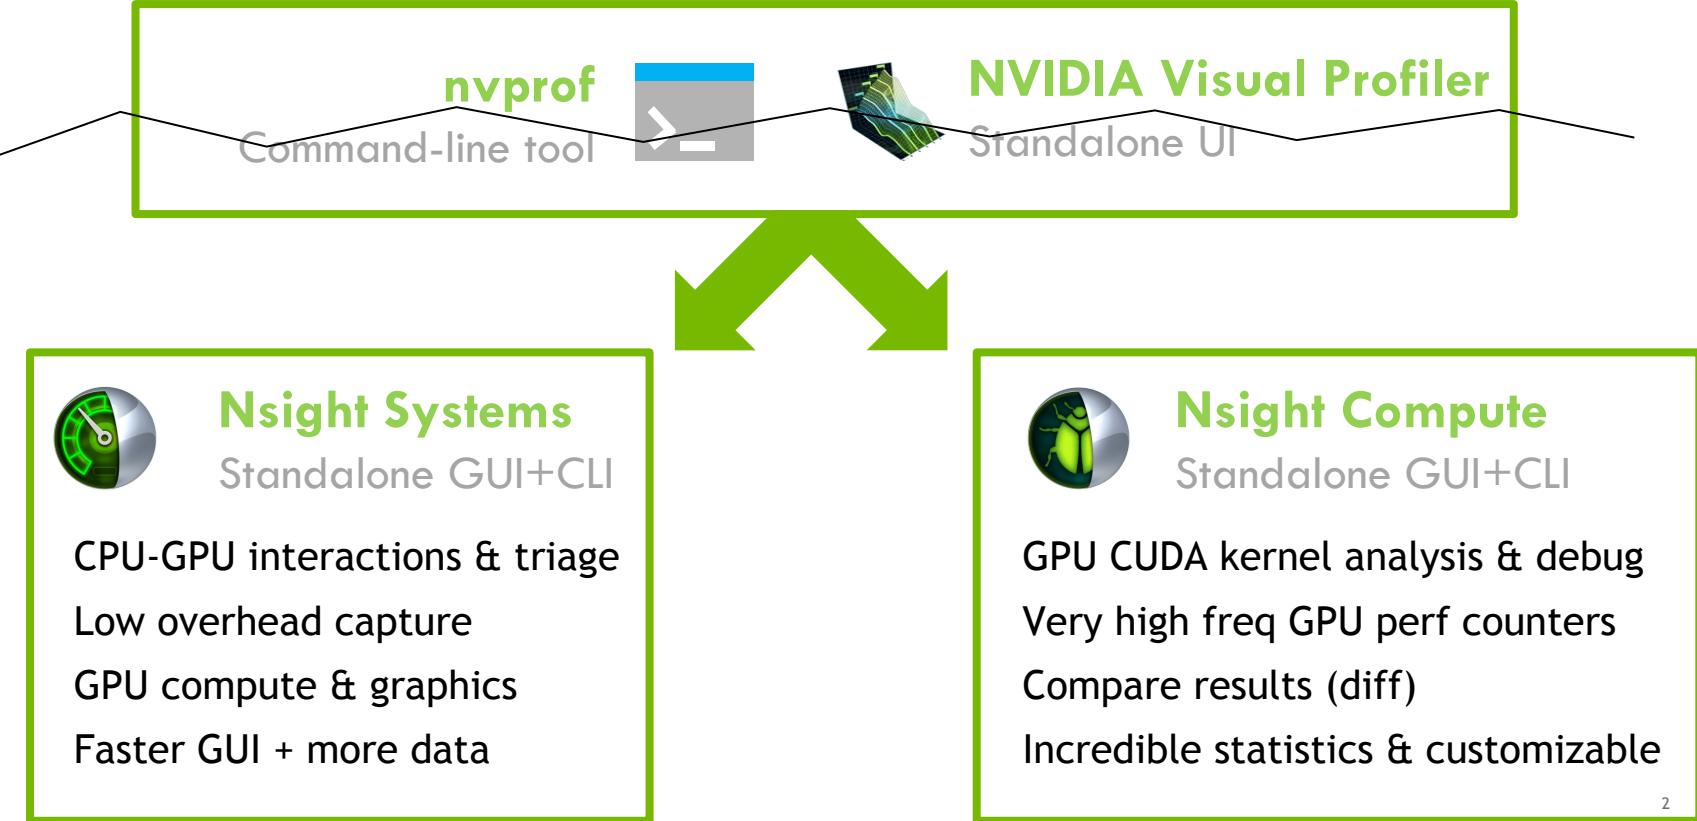




# Introduction to Nsight Systems for Summit

Holly Wilper, [hwilper@nvidia.com](mailto:hwilper@nvidia.com), March 9, 2020

# Legacy Transition



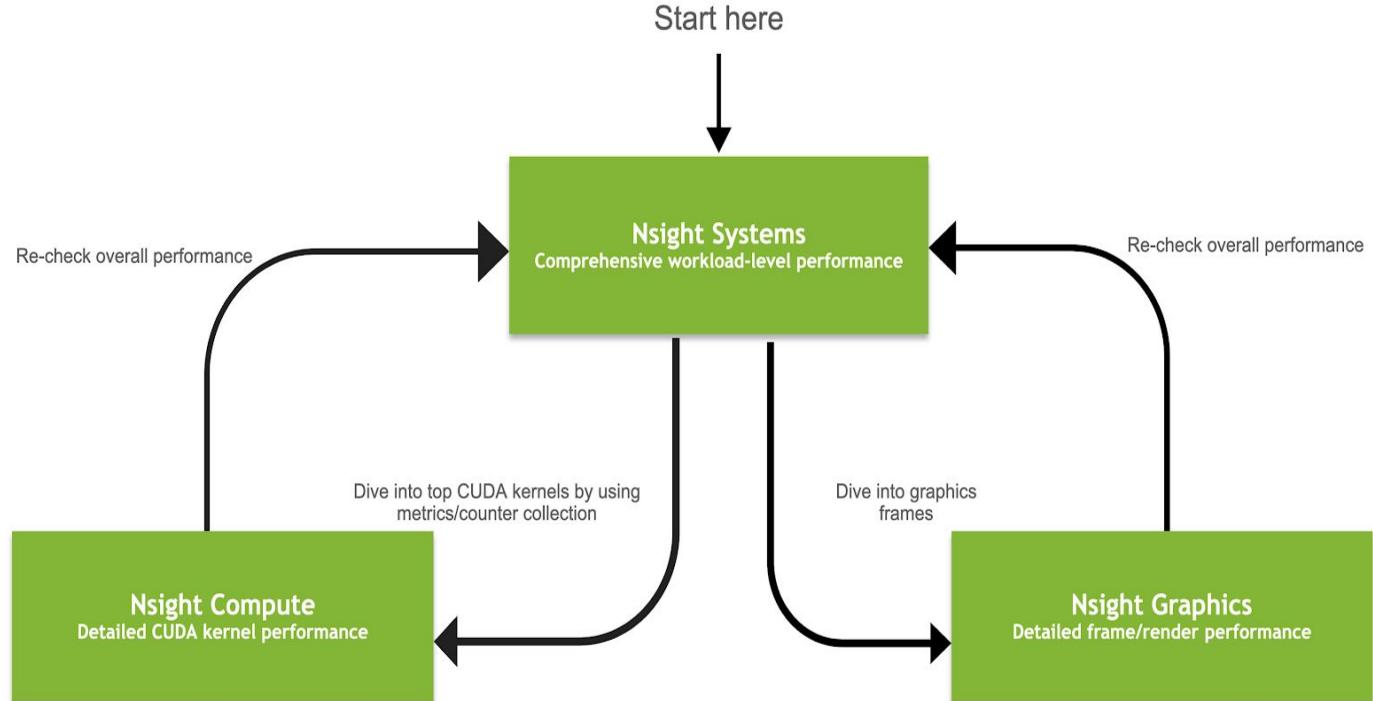
# Nsight Product Family

## Workflow

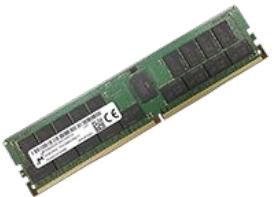
**Nsight Systems** -  
Analyze application  
algorithm system-wide

**Nsight Compute** -  
Debug/optimize CUDA  
kernel

**Nsight Graphics** -  
Debug/optimize graphics  
workloads



# Tuning an Orchestra of Tasks





# Simulation Lattice Microbes



# Pro Visualization & Games

# Overview

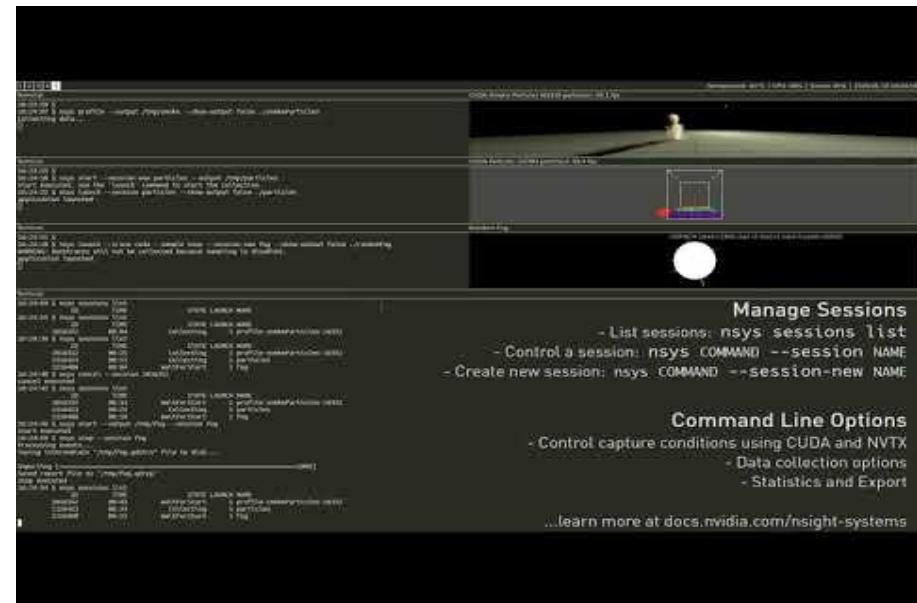
- **System-wide application algorithm tuning with multi-process tree support**
- **Locate optimization opportunities**
  - **Visualize millions of events on a very fast GUI timeline**
  - **See gaps of unused CPU and GPU time**
- **Balance your workload across multiple CPUs and GPUs**
  - **CPU algorithms, utilization, and thread state**
  - **GPU streams, kernels, memory transfers, etc**
- **Multi-platform:**
  - **Host GUI - Linux, Windows, Mac**
  - **Target Arch - x86-64, IBM Power, ARM server, Tegra**

# Timeline Features on Power Arch

- **Compute**
  - CUDA 10+ API & GPU workload ranges & mem transfers with correlation
  - OpenACC
  - OpenMP (next release)
- **OS**
  - Thread state and CPU utilization
  - MPI (OpenMPI & MPICH)
- **NVTX User Annotations API**

# Other Key Features

- Thread call-stack periodic sampling
  - Backtraces via frame pointers (or Dwarf unwind in next release)
  - Hot functions
- Command Line Interface (CLI)
  - No host PC required to record
  - Works in containers & VMs
  - Usable w/ access limitations
  - Scriptable / interactive mode
  - Multiple sessions
  - Multiple reports per launch

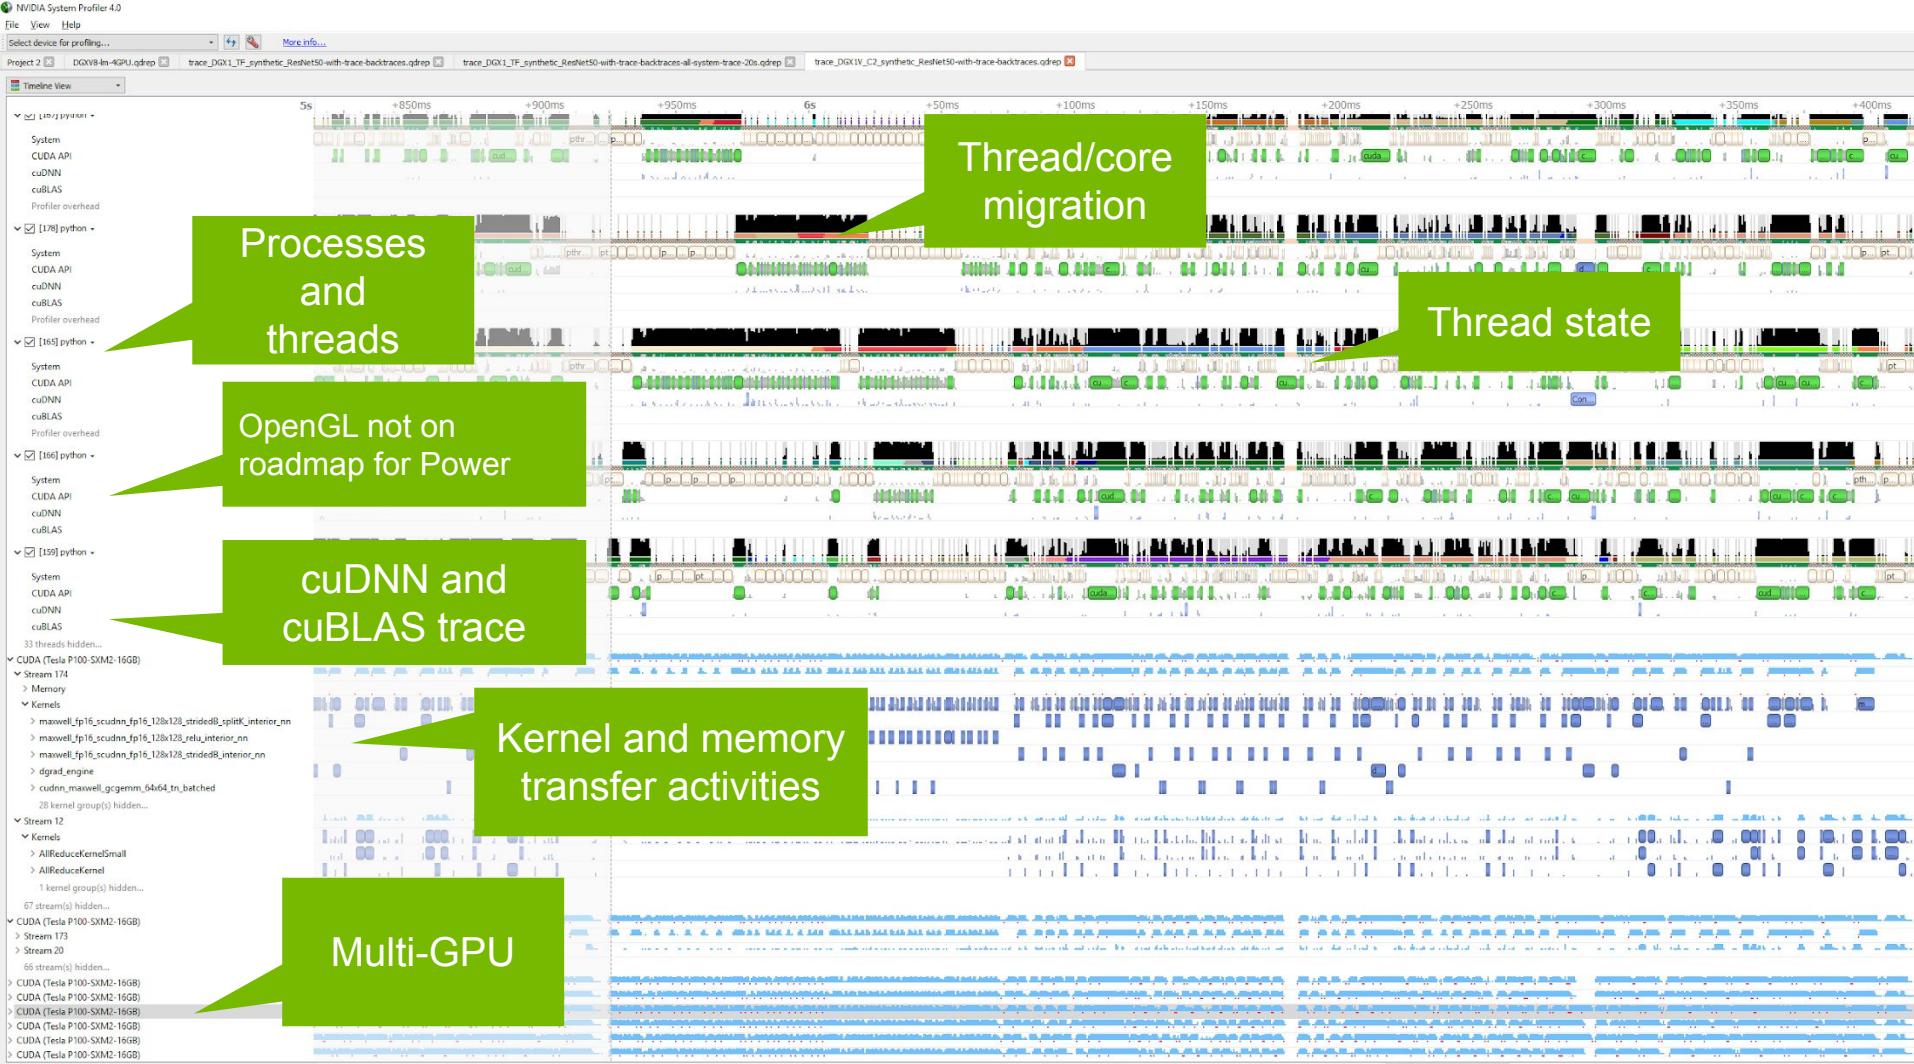


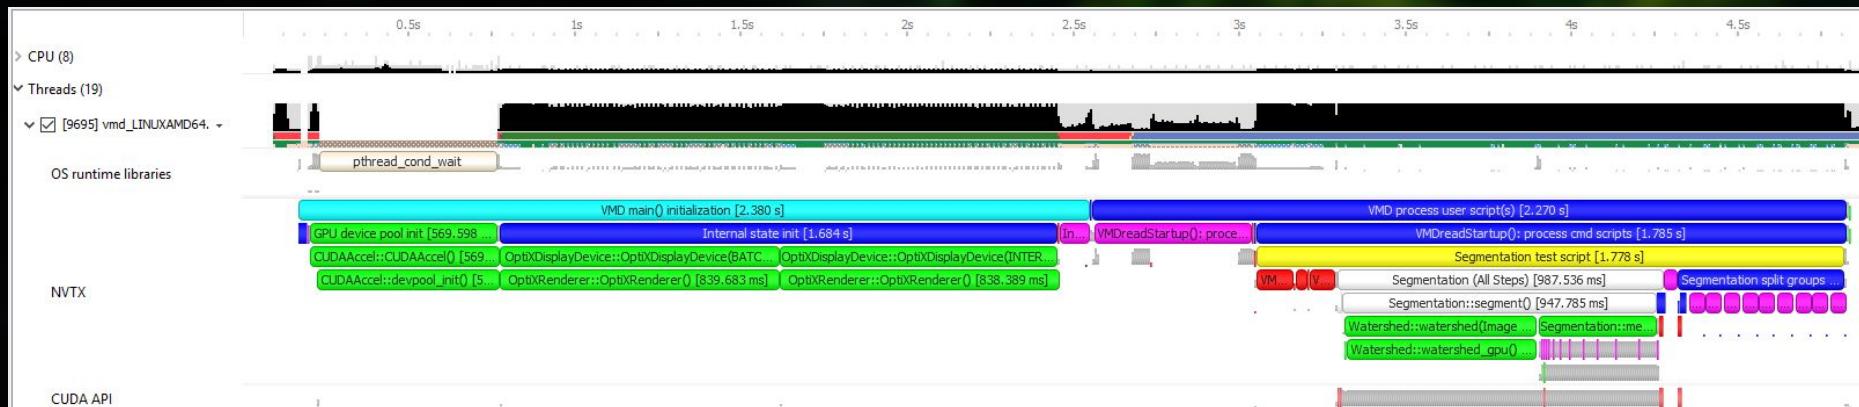


# Demo



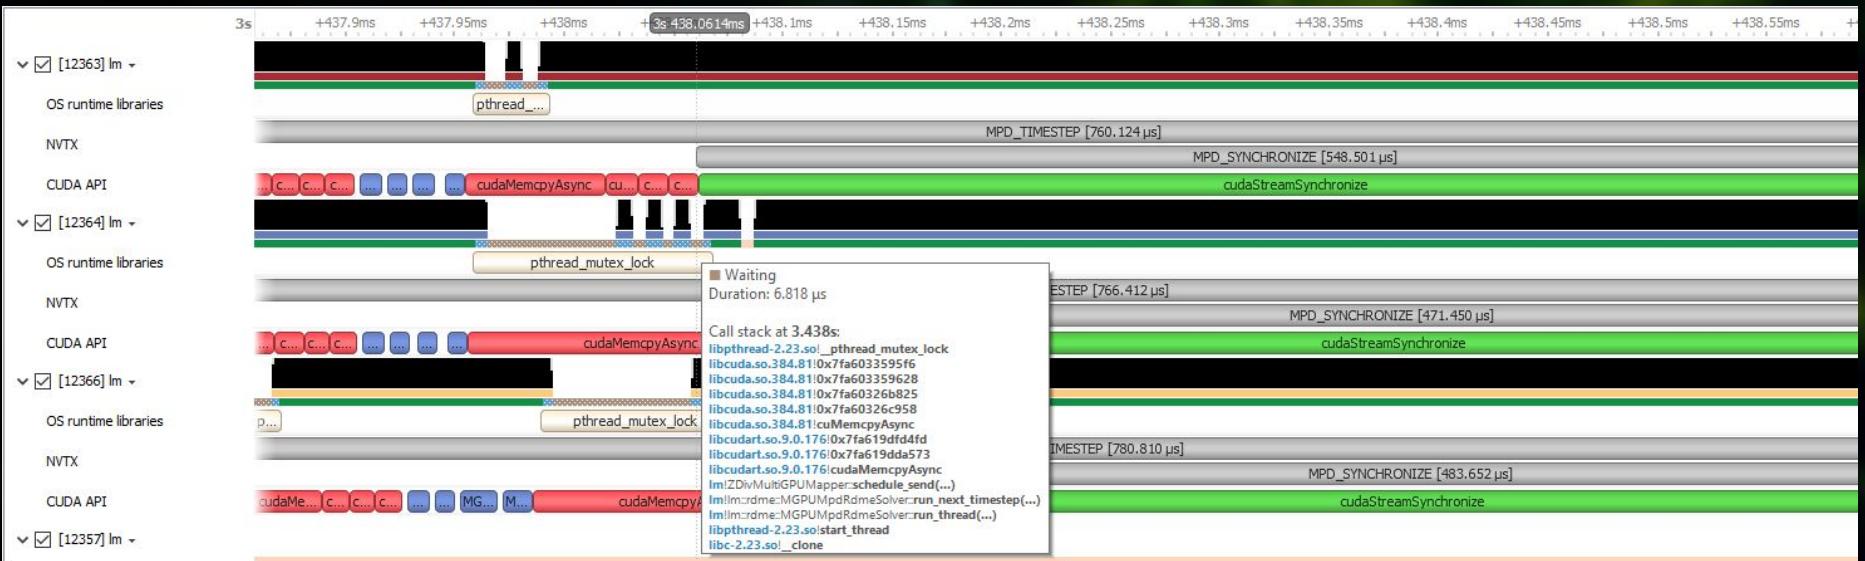
# Feature Highlights





## User annotations APIs for CPU & GPU

Example: Visual Molecular Dynamics (VMD) algorithms visualized with NVTX on CPU



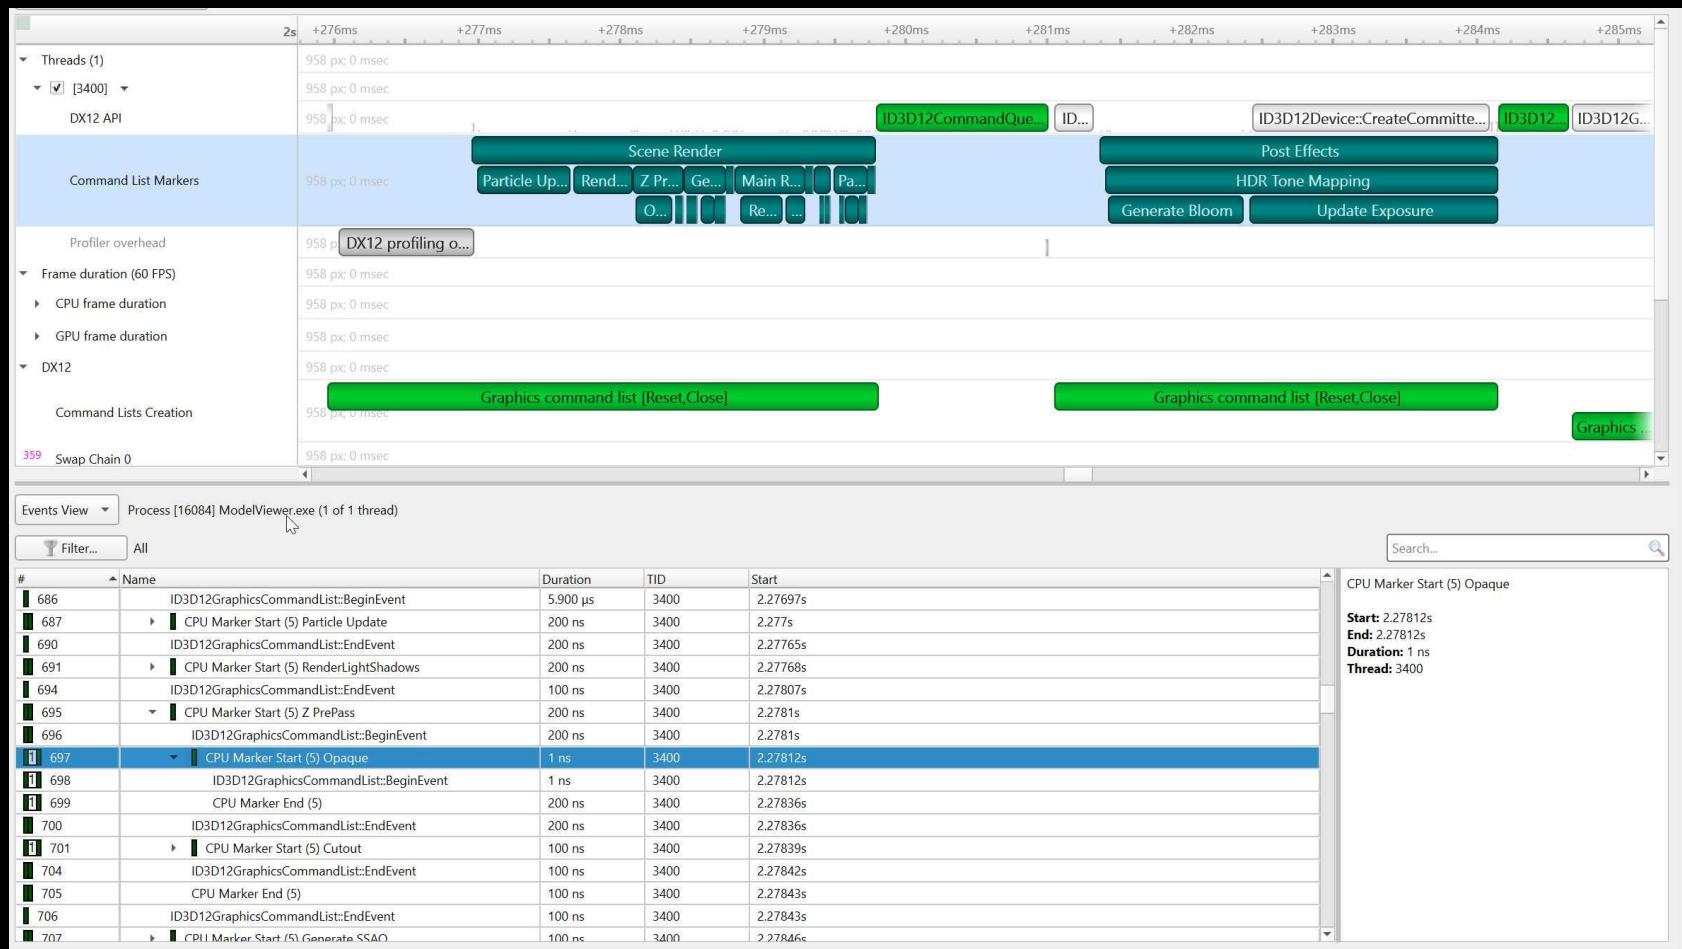
OS runtime trace (OSRT)  
(Coming to Power in future release)

Filter ...

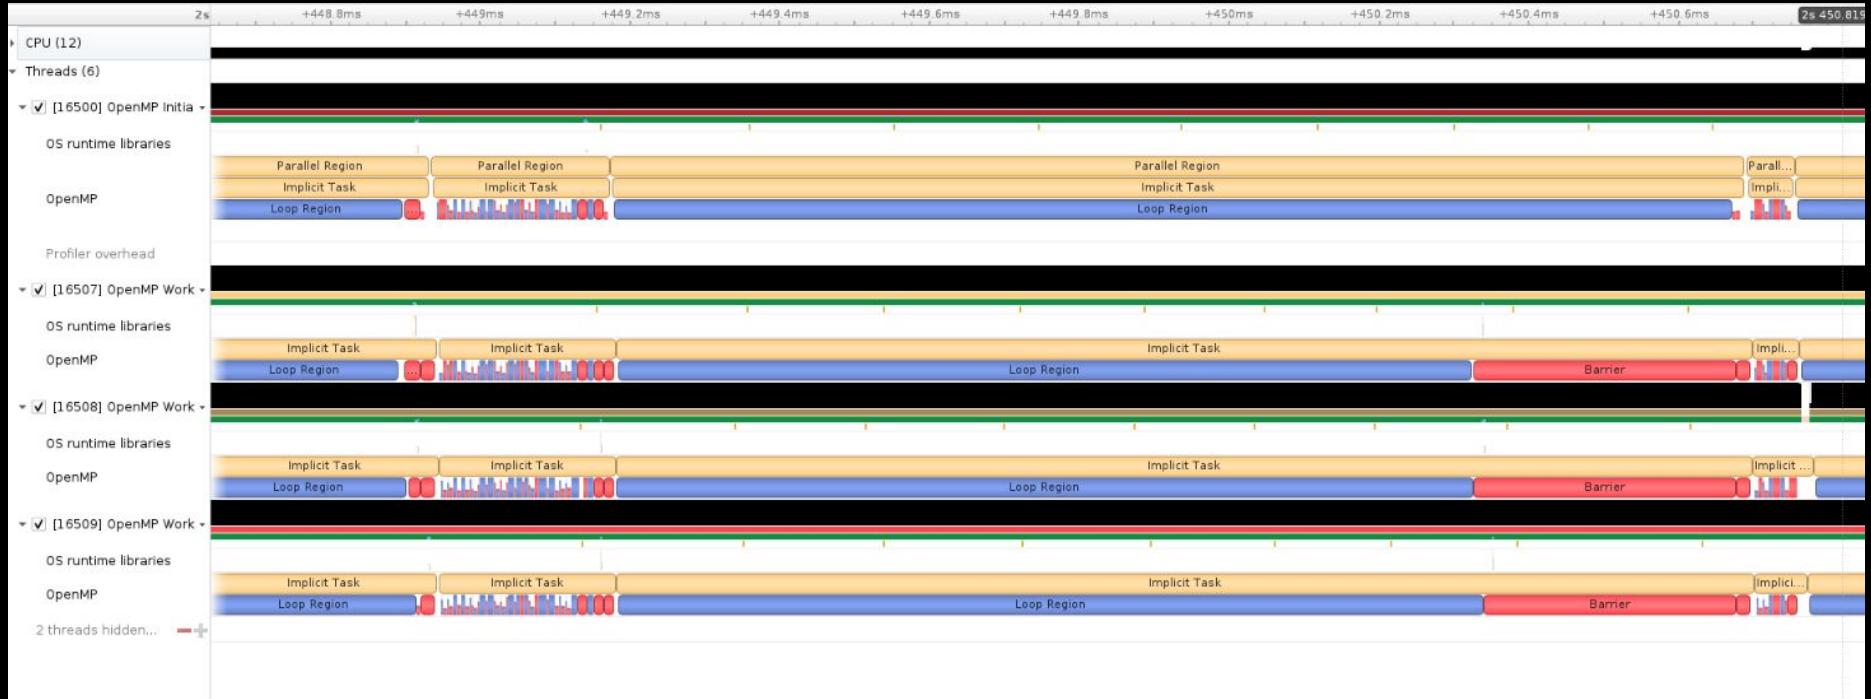
99.82% (23,260 samples) of data is shown due to applied filters.

Symbol Name	Self, %	Module Name
▼ VolumetricData::compute_volume_gradient()	20.14	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ VolumetricData::compute_volume_gradient()	20.14	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ BaseMolecule::add_volume_data(char const*, double const*, double const*, double const*, double const*, int, int, int, float*)	18.30	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ VMDApp::molecule_add_volumetric(int, char const*, double const*, double const*, double const*, double const*, int, int, int, float*)	18.30	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ obj_segmentation(void*, Tcl_Interp*, int, Tcl_Obj* const*)	18.30	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
[Max depth]	18.30	[Max depth]
▼ BaseMolecule::add_volume_data(char const*, float const*, float const*, float const*, float const*, int, int, int, float*, float*, float*)	1.84	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ MolFilePlugin::read_volumetric(Molecule*, int, int const*)	1.84	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ VMDApp::molecule_load(int, char const*, char const*, FileSpec const*)	1.84	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ text_cmd_mol(void*, Tcl_Interp*, int, char const**)	1.84	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ TclInvokeStringCommand	1.84	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ TclEvalObjvInternal	1.84	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ TclExecuteByteCode	1.84	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ TclCompEvalObj	1.84	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ TclEvalObjEx	1.84	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ Tcl_RecordAndEvalObj	1.84	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ TclTextInterp::evalFile(char const*)	1.84	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ VMDApp::logfile_read(char const*)	1.84	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
▼ VMDreadStartup(VMDApp*)	1.84	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11
[Max depth]	1.84	[Max depth]
> 0x7f10ca7022d6	5.13	/usr/lib64/libcuda.so.390.25
> obj_segmentation(void*, Tcl_Interp*, int, Tcl_Obj* const*)	3.44	/home/johns/vmd/src/gtcbuilds/vmd_LINUXAMD64.11

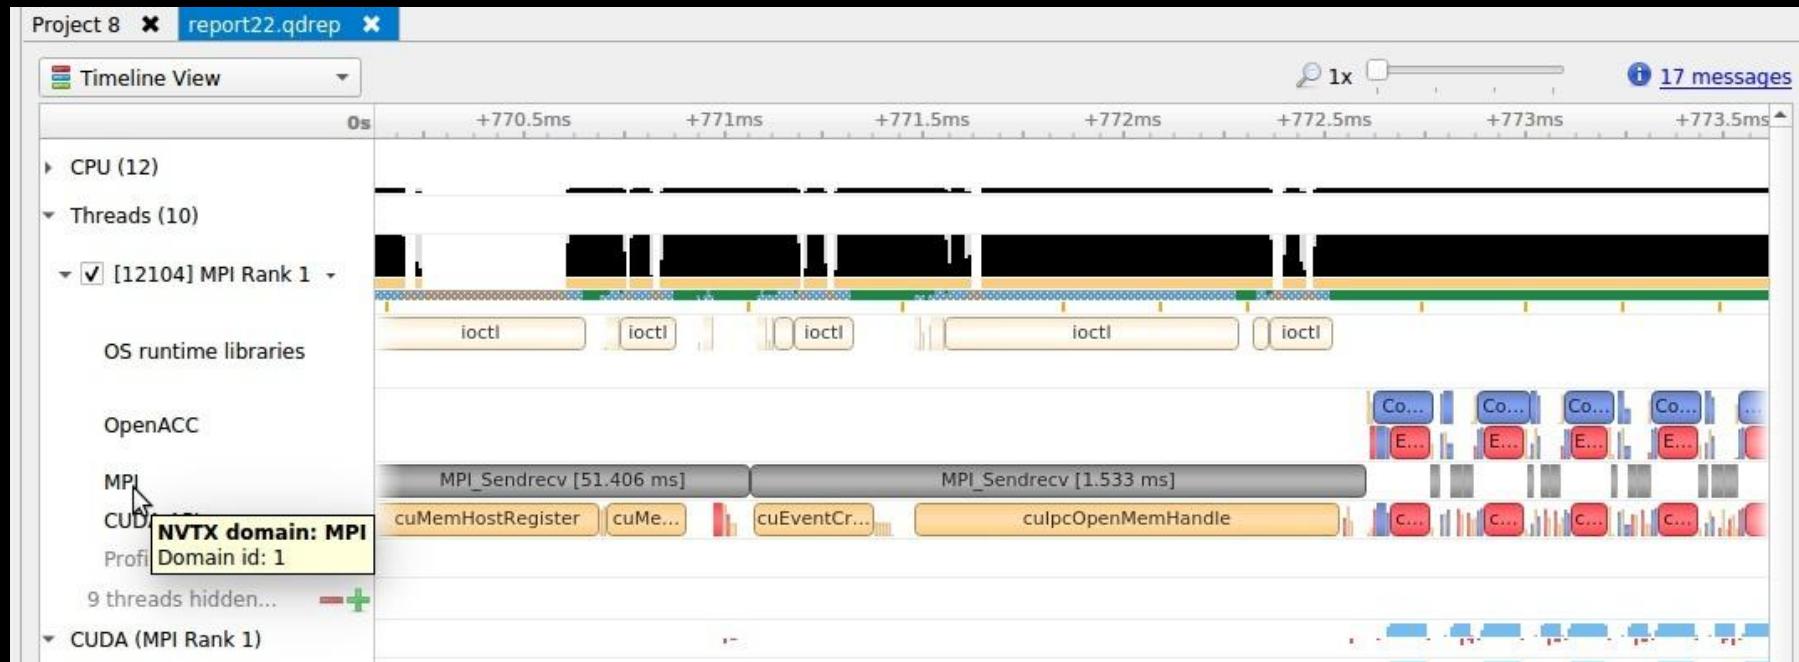
Function table shows statistics from periodic call-stack backtraces



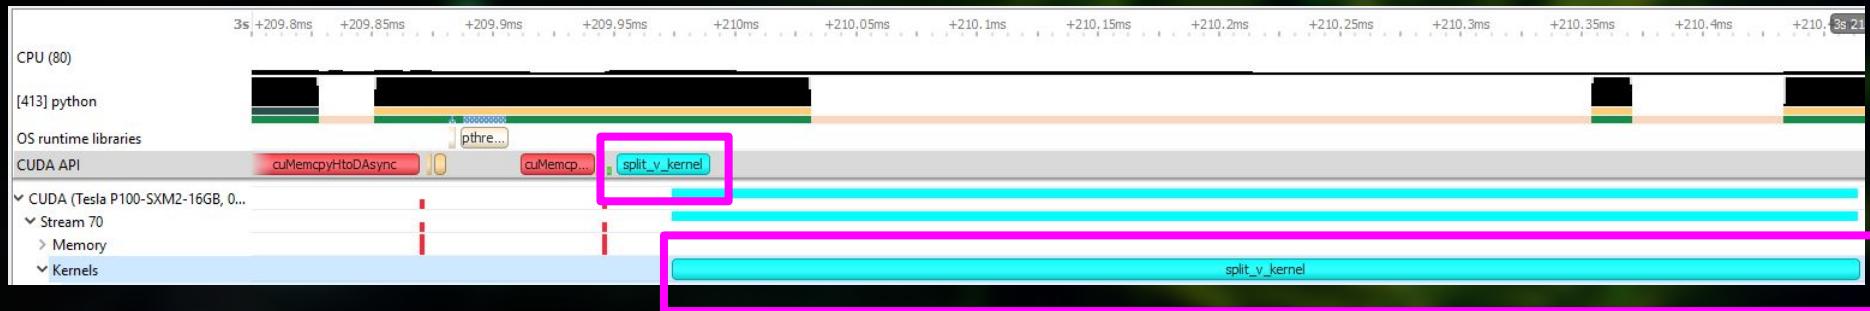
# Event Table



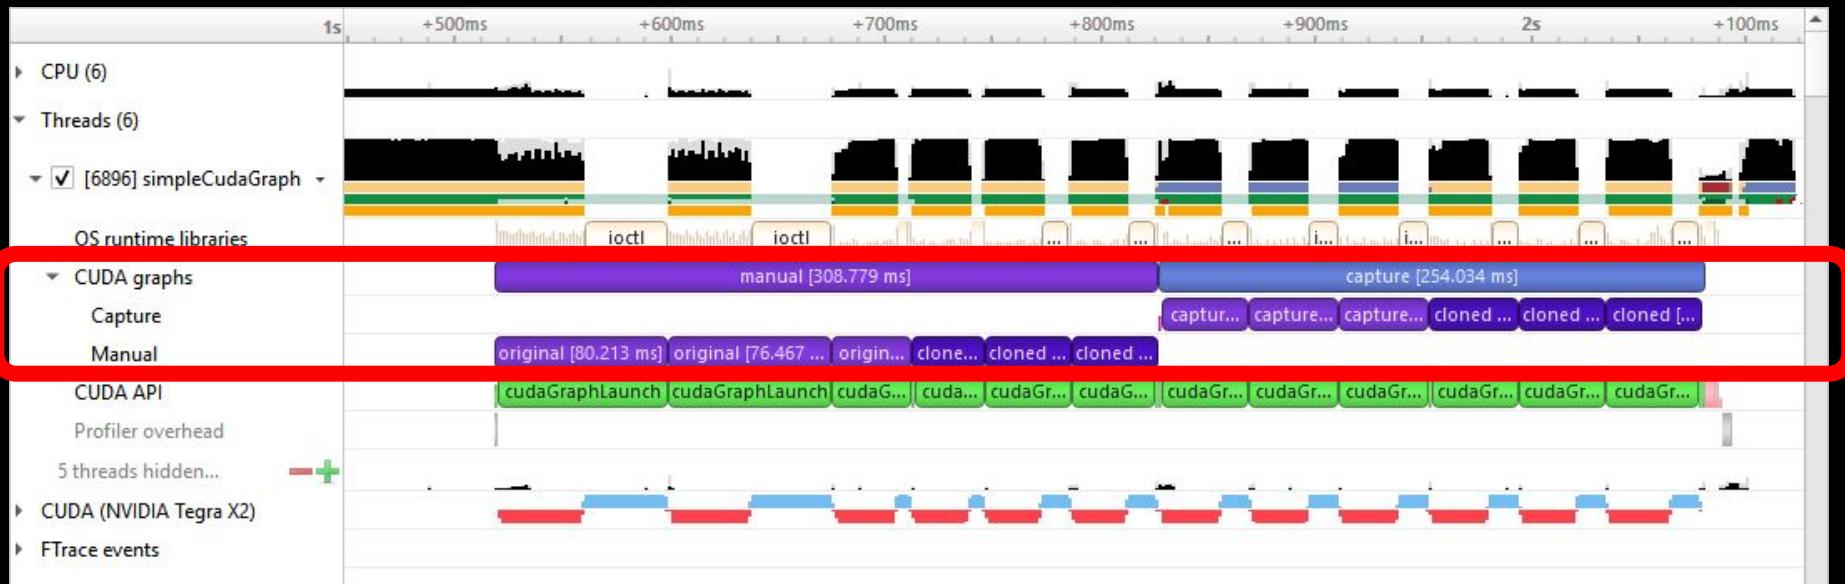
**OpenMP 5.0 trace** (coming soon)  
v4+



## MPI trace



GPU API launch to HW workload correlation



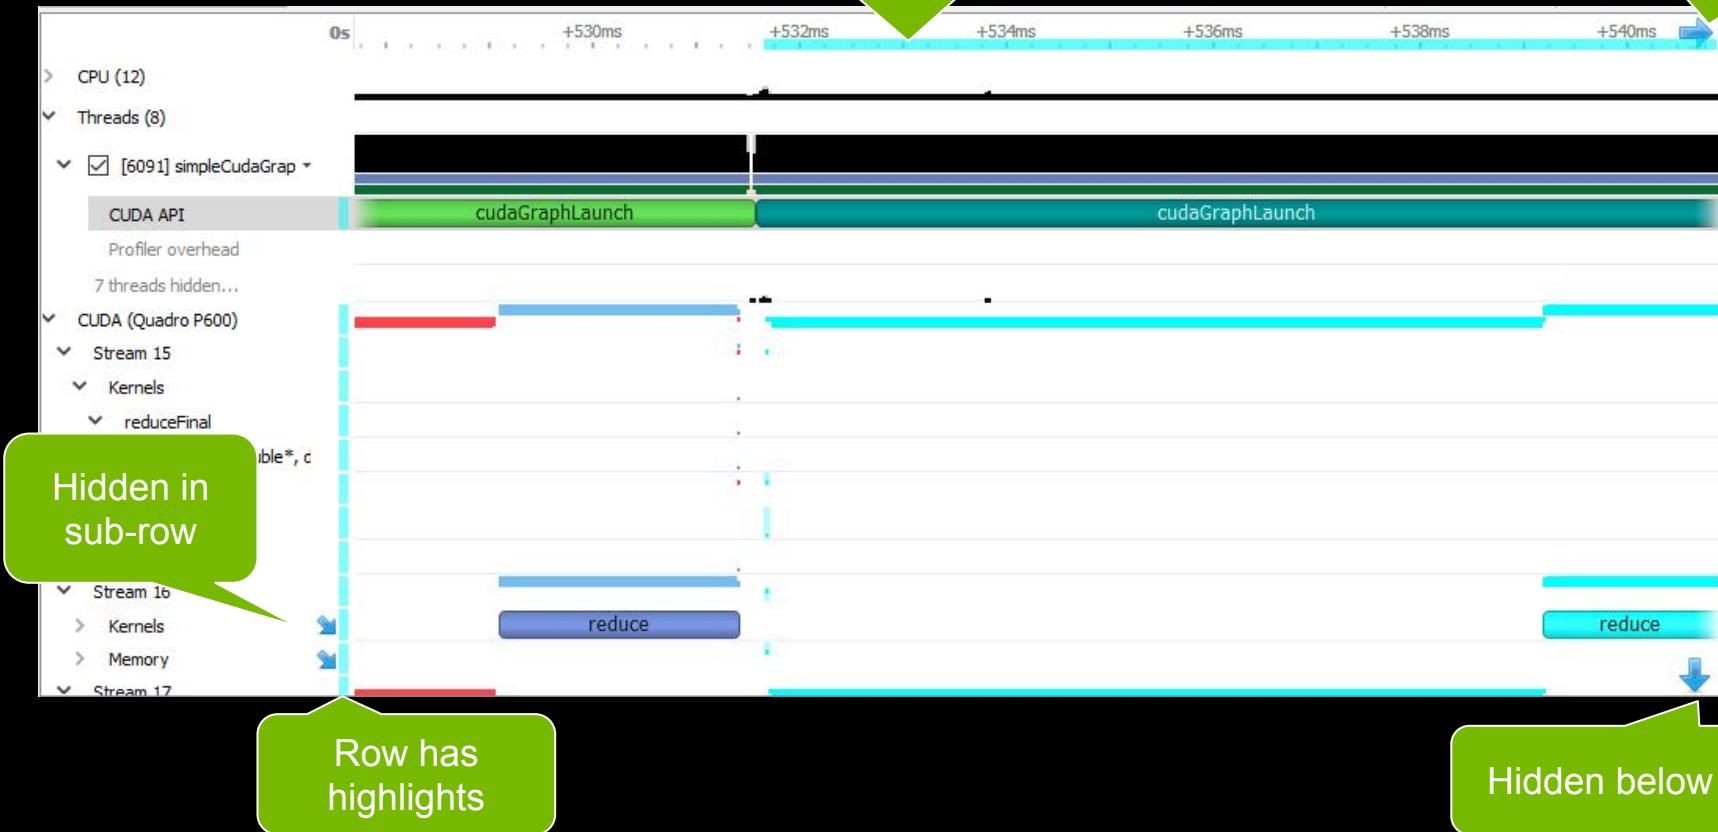
**NVTX Domains - Hoisting & Hierarchies**  
Look like APIs and '/' forms hierarchy



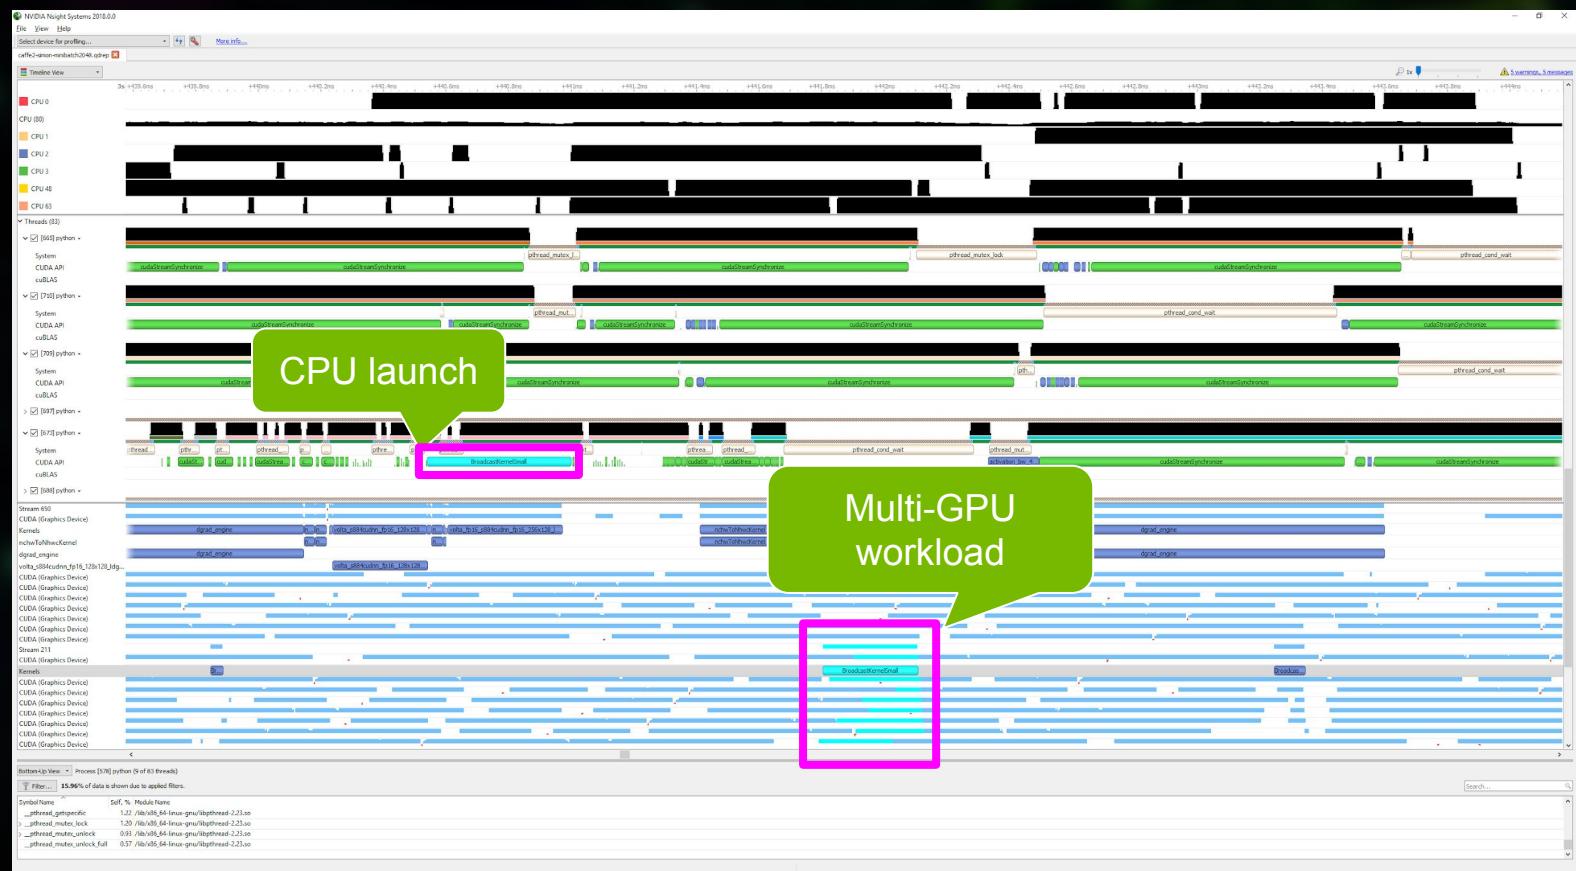
CUDA graph launches show all related GPU ranges

## Highlights in ruler

Hidden to right

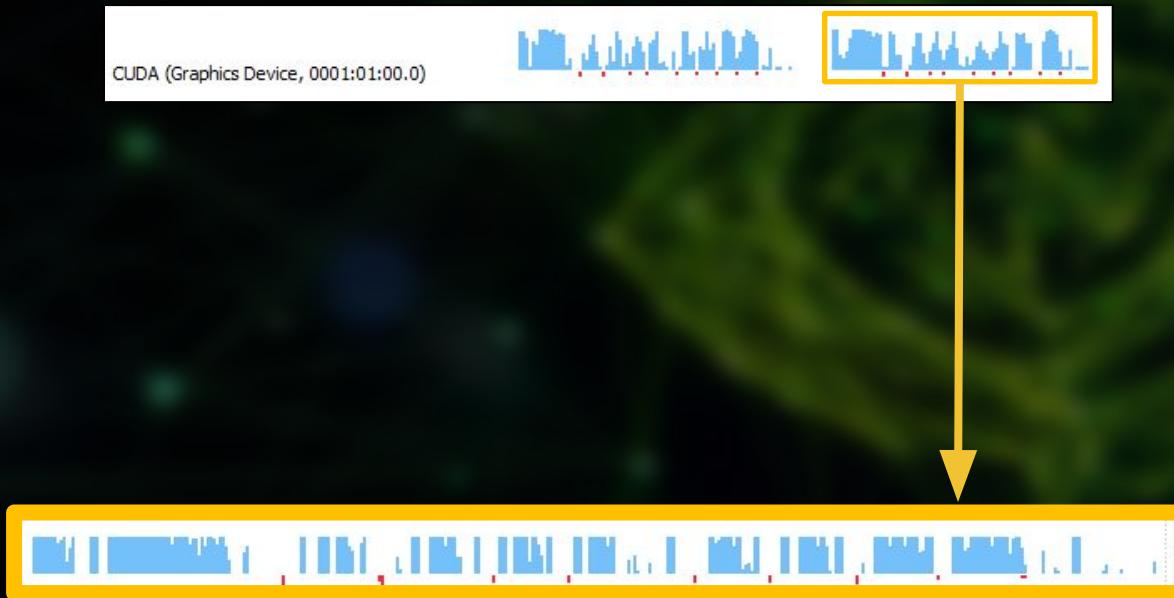


CPU-GPU correlation & location assistance

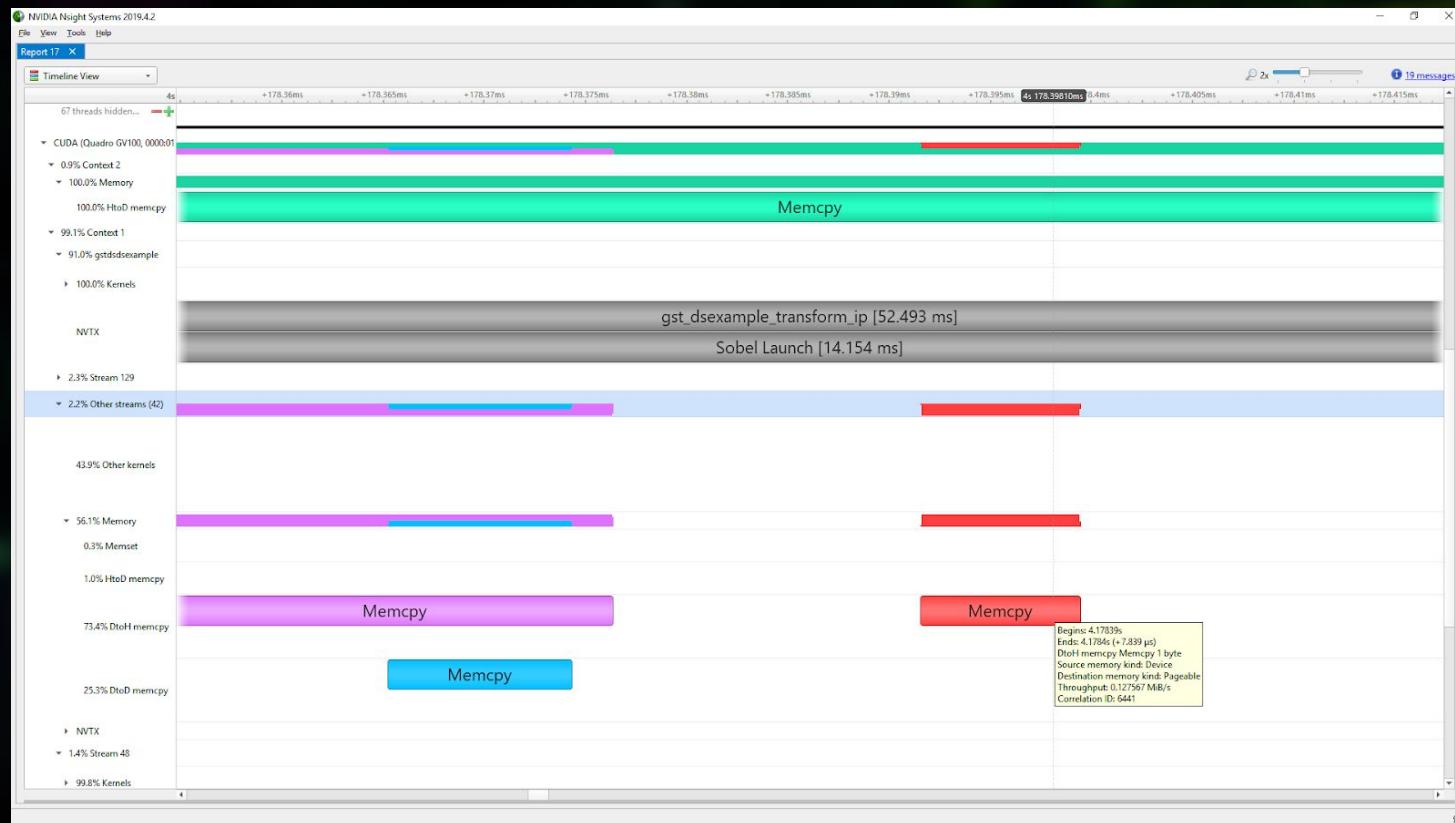


**cudaLaunchCooperativeKernelMultiDevice**  
 From Caffe2 Resnet50 within a container on a DGX-2

## GPU utilization based on percentage time coverage



Zooming in reveals gaps where there were valleys



CUDA memory transfer color palette  
show direction and pageable memory hazards

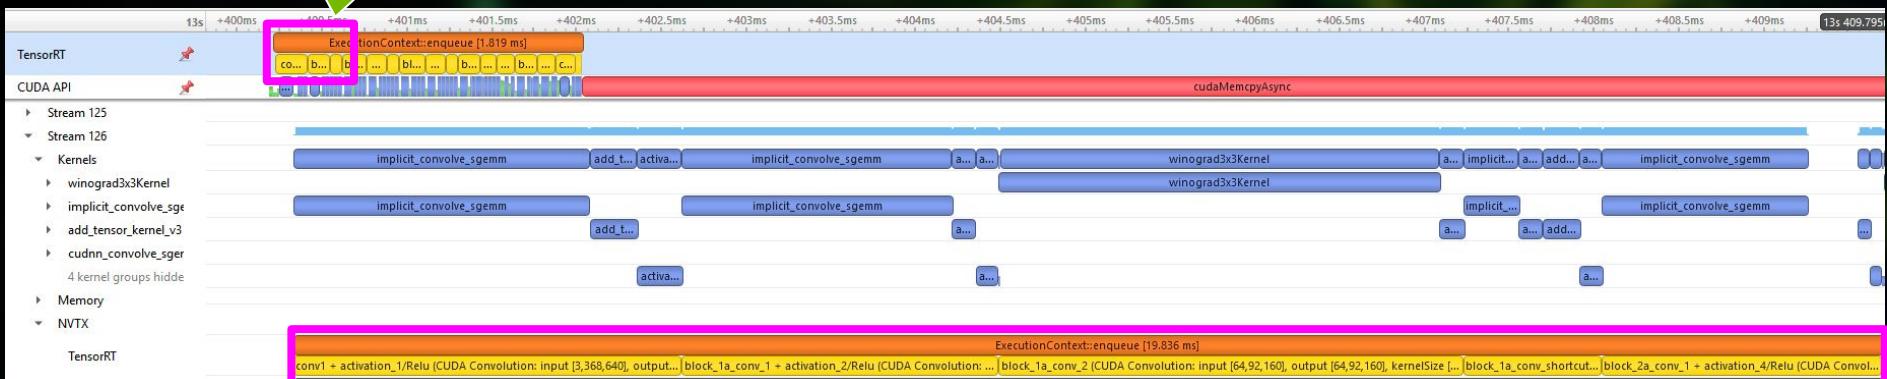


## CUDA unified virtual memory (UVM) transfers



TensorRT trace  
From DeepStream

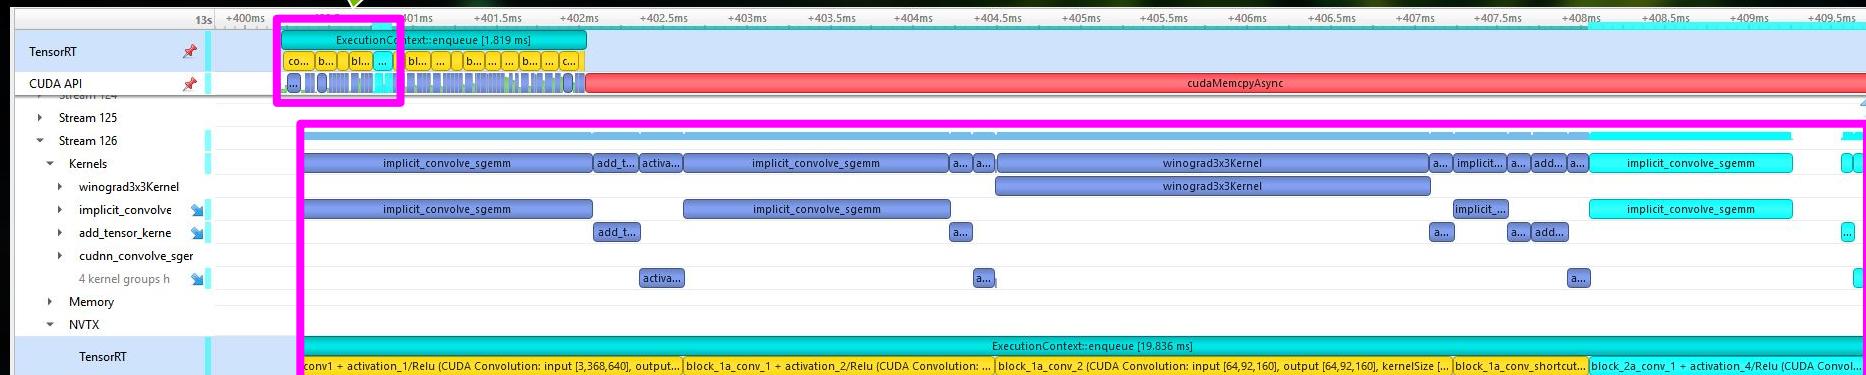
CPU launch



GPU workload

NVTX ranges projected onto the GPU  
From DeepStream

CPU  
launches



GPU workloads

NVTX deep highlighting and correlation  
From DeepStream



TensorFlow Resnet50 DNN nodes  
as NVTX ranges projected onto the GPU

# CLI statistics and export (SQLite & JSON)

Time(%)	Time (ns)	Calls	Avg (ns)	Min (ns)	Max (ns)	Name
25.3	96449693309	492166	195969.8	2590	11910256	cudaStreamSynchronize
17.3	65985078126	1337536	49333.3	6121	207018385	cudaLaunchKernel
15.3	58387659442	303066	192656.6	7710	126095339	cudaMemcpyPeerAsync
13.8	52570248446	88	597389186.9	79402	1109295234	cudaHostAlloc
7.8	29692609029	174420	170236.3	5770	133461628	cudaMemcpyAsync
5.0	19201650230	4717	4070733.6	8389	657446238	cudaMalloc
3.2	12177533056	306181	39772.3	901	124721678	cudaEventRecord
3.1	11678553105	108375	107760.6	5751	10799852	cudaMemcpy2DAsync
2.8	10525939137	120	87716159.5	2445	2684308531	cudaStreamCreateWithFlags
2.2	8474469667	78859	107463.6	6956	97743689	cudaMemsetAsync
1.5	5643462857	10	564346285.7	10283	2824251257	cudaStreamCreate
1.2	4409251983	133120	33122.4	10932	3830401	cudaLaunchCooperativeKernel
0.8	2932703617	441	6650121.6	669	727451161	cudaFree
0.6	2202242513	165248	13326.9	3346	27369081	cudaEventSynchronize
0.1	361857328	145933	2479.6	611	3467557	cudaStreamWaitEvent
0.1	210636728	19	11086143.6	88975	76110537	cudaMallocHost
0.0	16891690	864	19550.6	706	13229050	cudaEventCreateWithFlags
0.0	6833490	1256	5440.7	3696	20444	cudaEventQuery
0.0	1597827	32	49932.1	2931	366719	cudaStreamCreateWithPriority
0.0	1064579	38	28015.2	15625	72513	cudaMemcpy
0.0	1023364	184	5561.8	2589	28453	cudaBindTexture
0.0	376511	184	2046.3	1004	20137	cudaUnbindTexture
0.0	319084	81	3939.3	952	75350	cudaEventCreate
0.0	275150	138	1993.8	949	10096	cudaEventDestroy

## Stats/Export - CUDA API summary

Time(%)	Time (ns)	Calls	Avg (ns)	Min (ns)	Max (ns)	Name
13.6	17198585879	66612	258190.5	20864	1327579	nhwc_batch_norm_fwd
11.3	14228434514	20075	708763.9	168352	1786335	nhwc_batch_norm_bwd_add_relu
10.4	13100639465	40158	326227.4	177824	1342749	volta_fp16_s884cudnn_fp16_128x128_ldg8_dgrad_f2f_exp_interior_nhwc_tt_v1
10.1	12688128728	42745	296833.1	168895	760641	volta_fp16_s884cudnn_fp16_128x128_ldg8_relu_f2f_exp_interior_nhwc_tt_v1
9.3	11744819167	41418	283568.0	167680	753631	volta_s884cudnn_fp16_128x128_ldg8_wgrad_idx_exp_interior_nhwc_nt
7.1	8935340081	23848	374678.8	229056	895710	volta_s884cudnn_fp16_64x64_sliced1x4_ldg8_wgrad_idx_exp_interior_nhwc_nt
6.2	7829371510	41408	189078.7	28800	1365372	nhwc_batch_norm_bwd_relu
4.1	5145597181	13820	372329.8	325313	524159	volta_fp16_s884cudnn_fp16_128x128_ldg8_relu_f2f_exp_small_nhwc_tt_v1
3.8	4796840892	12551	382187.9	311649	775264	volta_fp16_s884cudnn_fp16_128x128_ldg8_dgrad_f2f_exp_small_nhwc_tt_v1
2.9	3599110831	482226	7463.5	1760	868095	dctQuantInvJpegKernel
2.4	2968987604	5018	591667.5	124767	1418268	nhwc_batch_norm_bwd
2.0	2550050661	162488	15693.8	1184	1720191	ycbcr_to_format_kernel_roi
1.6	2031335889	3768	539101.9	381407	784829	dgrad_1x1_stride_2x2
1.5	1843345484	1256	1467631.8	1451800	1691492	pooling_bw_max_nhwc_kernel
1.2	1565496722	3781	414043.0	391712	480321	volta_fp16_s884cudnn_fp16_256x64_ldg8_relu_f2f_exp_small_nhwc_tt_v1
1.2	1543111896	3765	409857.1	386239	2808775	volta_fp16_s884cudnn_fp16_256x64_ldg8_dgrad_f2f_exp_small_nhwc_tt_v1
1.1	1373574502	2512	546805.1	434785	760834	dgrad_1d
1.0	1267659803	2514	504240.2	492895	601119	volta_fp16_s884cudnn_fp16_256x128_ldg8_relu_f2f_exp_small_nhwc_tt_v1
1.0	1215018985	58722	20691.0	1152	1325441	MapPlanKernel
0.8	1028518778	5670	179984.3	960	737985	mxnet_generic_kernel
0.7	909701850	65290	13933.2	6400	123808	convertTensor_kernel
0.7	894204599	2536	352604.3	192223	3158098	BatchedSeparableResampleKernel
0.7	868771063	1257	691146.4	670495	821372	first_layer_wgrad_kernel
0.7	859121367	1261	681301.6	675584	793889	first_layer_fwd_kernel
0.7	854624303	1257	679892.0	661086	726018	volta_fp16_s884cudnn_fp16_256x128_ldg8_relu_f2f_exp_interior_nhwc_tt_v1
0.6	719945062	1256	573204.7	555264	648066	dgrad_2d
0.4	479525740	6430	74576.3	2560	1835937	dcAcDecodeKernel
0.4	450010850	1268	354898.1	223551	1985854	BatchedCropMirrorNormalizePermuteKernel
0.3	436136496	1260	346140.1	342111	361025	pooling_fw_kernel_max_nhwc
0.3	364887793	1257	290284.6	164992	395295	volta_fp16_s884cudnn_fp16_256x64_ldg8_dgrad_f2f_exp_interior_nhwc_tt_v1
0.3	354605440	6430	55148.6	2176	1460961	transposeKernel
0.3	336891222	92	3661861.1	892835	26052414	convolve_common_engine_float_NHWC
0.2	302899259	65266	4628.7	928	48576	scalePackedTensor_kernel
0.2	277514711	123110	2254.2	1216	34304	computeOffsetsKernel
0.2	194972504	1261	154617.4	150848	177536	volta_fp16_s884cudnn_fp16_256x64_ldg8_relu_f2f_exp_interior_nhwc_tt_v1
0.1	176725421	46	3841857.0	578114	7942931	wgrad_alg0_engine_NHWC
0.1	160091745	65268	2452.8	1599	20384	computeWgradOffsetsKernel
0.1	116475012	1256	92734.9	91520	242912	pooling_bw_avg_nhwc_kernel
0.1	115380220	2159	53441.5	3008	1234110	destuffKernel
0.1	108445324	57731	1878.5	960	34016	computeBOffsetsKernel
0.1	77051216	23	3350052.9	1496867	7849586	wgrad_alg1_engine_NHWC
0.1	64827363	1256	51614.1	49663	67488	pooling_fw_4d_kernel

## Stats/Export - CUDA kernel summary

Time(%)	Time (ns)	Calls	Avg (ns)	Min (ns)	Max (ns)	Name
86.1	400991604181	118839	33742425.0	2032	29870299636	pthread_cond_wait
6.2	289369285750	7856	36834175.9	2080	108487660	poll
6.0	277985649608	751434	369940.2	1000	2813457310	pthread_mutex_lock
1.5	68744462781	5380	12777781.2	49730	100573603	sem_timedwait
0.1	6271750239	16152	388295.6	1000	28145137	ioctl
0.0	1356083977	988	1372554.6	2152	26485837	mmap
0.0	1049790648	122115	8596.7	1000	8677246	pthread_cond_signal
0.0	651506851	230	2832638.5	230612	27325479	sem_wait
0.0	45465761	6326	7187.1	1002	22511	pthread_cond_broadcast
0.0	40173319	959	41890.8	1814	7576942	munmap
0.0	27428315	92	298133.9	50349	773363	pthread_create
0.0	11388591	3281	3471.1	1000	118420	fflush
0.0	3322283	337	9858.4	5815	44869	fopen
0.0	2404221	259	9282.7	2452	160020	mmap64
0.0	2287088	90	25412.1	1592	212643	mprotect
0.0	2182161	337	6475.3	4275	37453	open64
0.0	2110422	246	8579.0	2585	40699	write
0.0	1137196	539	2109.8	1021	8330	read
0.0	1059939	340	3117.5	2282	14565	fclose
0.0	551730	18	30651.7	19990	123582	backtrace
0.0	121017	105	1152.5	1000	1811	fcntl
0.0	55359	3	18453.0	13893	24165	fopen64
0.0	33996	7	4856.6	1043	25913	fwrite

## Stats/Export - OS Runtime API summary - in 2020.2

Time(%)	Time (ns)	Instances	Avg (ns)	Min (ns)	Max (ns)	Range
4.0	162789746	70	2325567.8	248462	10421312	TensorRT:ExecutionContext::enqueue
0.4	16335803	24	680658.5	94707	2036983	TensorRT:predictions
0.3	10447691	70	149252.7	16094	3604518	TensorRT:conv1 + activation_1/Relu
0.2	10169055	70	145272.2	25936	1169143	TensorRT:conv1 + activation_1/Relu input reformatter 0
0.2	8698016	24	362417.3	9581	2970733	TensorRT:block_3b_conv_2
0.2	8625565	24	359398.5	28257	1242590	TensorRT:average_pooling2d_1
0.2	7733550	46	168120.7	30474	3498224	TensorRT:conv2d_cov/Sigmoid
0.2	6672310	70	95318.7	11271	1151050	TensorRT:block_1a_conv_1 + activation_2/Relu
0.1	5312256	70	75889.4	8316	2133255	TensorRT:block_4a_conv_2
0.1	5197452	24	216560.5	9572	1543570	TensorRT:block_3b_conv_1 + activation_12/Relu
0.1	4922521	70	70321.7	11732	1281643	TensorRT:block_1a_conv_shortcut + add_1 + activation_3/Relu
0.1	4151987	24	172999.5	10592	1625630	TensorRT:block_1b_conv_2
0.1	4121931	70	58884.7	10616	1307464	TensorRT:block_2a_conv_2
0.1	4108970	70	58699.6	9311	834374	TensorRT:block_1a_conv_2
0.1	4027181	24	167799.2	10564	1728250	TensorRT:block_1b_conv_shortcut + add_2 + activation_5/Relu
0.1	4018713	24	167446.4	9636	2431179	TensorRT:block_2b_conv_1 + activation_8/Relu
0.1	3799363	70	54276.6	8595	481048	TensorRT:block_3a_conv_2
0.1	3762886	24	156786.9	13969	1008096	TensorRT:block_4b_conv_shortcut + add_8 + activation_17/Relu
0.1	3376998	24	140708.2	12773	1641372	TensorRT:block_4b_conv_2
0.1	3361499	24	140062.5	11195	1248129	TensorRT:block_2b_conv_2
0.1	3227736	24	134489.0	15949	736287	TensorRT:predictions/Softmax
0.1	3023238	24	125968.2	11972	827474	TensorRT:block_2a_conv_shortcut + add_3 + activation_7/Relu
0.1	2998494	24	124937.2	12327	1217984	TensorRT:block_2a_conv_1 + activation_6/Relu
0.1	2853998	46	62043.4	10971	1983123	TensorRT:block_3a_conv_shortcut + add_3 + activation_7/Relu
0.1	2773491	24	115562.1	12019	809395	TensorRT:block_4a_conv_shortcut + add_7 + activation_15/Relu
0.1	2744711	24	114363.0	10474	892787	TensorRT:block_2b_conv_shortcut + add_4 + activation_9/Relu
0.1	2147200	24	89466.7	12037	645408	TensorRT:block_3a_conv_1 + activation_10/Relu
0.0	1914194	24	79758.1	10599	842947	TensorRT:block_4b_conv_1 + activation_16/Relu
0.0	1732011	24	72167.1	9659	488562	TensorRT:block_3a_conv_shortcut + add_5 + activation_11/Relu
0.0	1675232	24	69801.3	10073	623120	TensorRT:block_4a_conv_1 + activation_14/Relu
0.0	1525283	24	63553.5	10386	540733	TensorRT:block_3b_conv_shortcut + add_6 + activation_13/Relu
0.0	1358781	46	29538.7	10645	241099	TensorRT:block_3a_conv_1 + activation_6/Relu
0.0	1301126	46	28285.3	11929	244583	TensorRT:conv2d_cov

## Stats/Export - NVTX code annotations

Note this includes TensorRT domains

```
$ cat func_table.sql | sqlite3 -column -header ~/Downloads/A_bloom_ALL.sqlite
symbolName      moduleName          flatTimePercentage
-----
__start_thread  /system/lib/libc.so  52.97
__pthread_star  /system/lib/libc.so  52.96
Func at 0x6c1a  /data/app/com.nvidi 51.98
android_main    /data/app/com.nvidi 51.92
art::JValue ar   /system/lib/libart. 46.62
art::interpret  /system/lib/libart. 46.43
art::interpret  /system/lib/libart. 45.97
NvSampleApp::m  /data/app/com.nvidi 45.65
NvSampleApp::m  /data/app/com.nvidi 44.13
bool art::inte   /system/lib/libart. 41.57

symbolName  moduleName          selfTimePercentage
-----
__ioctl     /system/lib/libc.so  9.61
__raw_spin_ [kernel.kallsyms]  2.24
e10_svc_na [kernel.kallsyms]  2.04
__epoll_pw  /system/lib/libc.so  1.68
memcpy      /system/lib/libc.so  1.09
__do_softi  [kernel.kallsyms]  0.95
e10_da     [kernel.kallsyms]  0.91
art::JValu   /system/lib/libart.  0.78
vsnprintf   [kernel.kallsyms]  0.72
kmem_cache  [kernel.kallsyms]  0.66
```

## Export - Thread call stack samples



# Running on Summit

```
[skottap@login5.summit skottap]$ module avail
----- /autofs/nccs-svm1_sw/summit/modulefiles/site/linux-rhel7-ppc64le/sp
adios2/2.2.0          fftw/3.3.8      hypre/2.11.1    netcdf-
amgx/2.0.0.130.2-unthreaded hdf5/1.8.18     hypre/2.13.0 (D)  netcdf-
boost/1.59.0          (D)           hdf5/1.10.3   mpip/3.4.1     netcdf/
autoconf/2.69         (D)           gdbm/1.18.1   libtool/2.4.2   openss
automake/1.16.1       (D)           hdf5/1.8.18     m4/1.4.18     (D)  pkgcon
boost/1.59.0          hdf5/1.10.3   metis/5.1.0    python
bzzip2/1.0.6          libdwarf/20180129 ncurses/6.1   readline
cmake/3.14.2          libelf/0.8.13   netlib-lapack/3.8.0 spectr
diffutils/3.7          (D)           libsodium/1.0.15 numactl/2.0.11   sqlite
DefApps               (L)           hwloc/2.0.2   (D)           pi
antlr/2.7.7          icu4c/58.2    libassuan/2.4.5 py
apr-util/1.6.0        libbsd/0.9.1   libevent/2.1.8  py
autoconf/2.69         libfabric/1.7.0  libgcrypt/1.8.1 py
automake/1.16.1       libbevent/2.1.8  libgd/2.2.4    py
bison/3.0.5           libfftw3/3.3.8  libgpg-error/1.27 py
c-biosc/1.12.1        libibverbs/1.1.3  libksba/1.3.5  py
cairo/1.16.0-py3      libltdl/2.4.5   libluban/1.3.0  py
cmake/3.6.1           liblxml/2.9.9.8 libpng/1.6.34   py
cmake/3.9.2           liblxcb/1.13   libpspell/0.15.5 py
cmake/3.11.3          liblxrender/0.9.10 libxml2/2.9.8   py
cmake/3.13.4          liblxsm/1.0-20180531 libxpm/3.5.12   py
cmake/3.14.2          liblxswm/1.0-20190225 libxrandr/0.9.10 py
cmake/3.15.2          liblxswm/1.0-20190225 libxrender/0.9.10 py
cuda/9.1.85           liblxswm/1.0-20190225 libxsm/1.0-20180531 py
cuda/9.2.148          liblxswm/1.0-20190225 libxsm/1.0-20190225 py
cuda/10.1.105         liblxswm/1.0-20190225 libxsm/1.0-20190225 py
cuda/10.1.168          liblxswm/1.0-20190225 libxsm/1.0-20190225 py
cuda/10.1.243          liblxswm/1.0-20190225 libxsm/1.0-20190225 py
curl/7.63.0           liblxswm/1.0-20190225 libxsm/1.0-20190225 py
darshan-runtime/3.1.7-hdf5pre110 liblxswm/1.0-20190225 libxsm/1.0-20190225 py
darshan-runtime/3.1.7-hdf5post110 liblxswm/1.0-20190225 libxsm/1.0-20190225 py
darshan-runtime/3.1.7-hdf5post110 liblxswm/1.0-20190225 libxsm/1.0-20190225 py
darshan-util/3.1.6    liblxswm/1.0-20190225 libxsm/1.0-20190225 py
darshan-util/3.1.7    liblxswm/1.0-20190225 libxsm/1.0-20190225 py
diffutils/3.7          liblxswm/1.0-20190225 libxsm/1.0-20190225 py
emacs/25.1             liblxswm/1.0-20190225 libxsm/1.0-20190225 py
essl/6.1.0-2           liblxswm/1.0-20190225 libxsm/1.0-20190225 py
essl/6.2.0-20190419    liblxswm/1.0-20190225 libxsm/1.0-20190225 py
expat/2.2.5            liblxswm/1.0-20190225 libxsm/1.0-20190225 py
fontconfig/2.12.3      liblxswm/1.0-20190225 libxsm/1.0-20190225 py
freetype/2.9.1          liblxswm/1.0-20190225 libxsm/1.0-20190225 py
gcc/4.8.5              liblxswm/1.0-20190225 libxsm/1.0-20190225 py
gcc/5.4.0              liblxswm/1.0-20190225 libxsm/1.0-20190225 py
gcc/6.4.0              liblxswm/1.0-20190225 libxsm/1.0-20190225 py
gcc/7.4.0              liblxswm/1.0-20190225 libxsm/1.0-20190225 py
gcc/8.1.0              liblxswm/1.0-20190225 libxsm/1.0-20190225 py
gcc/8.1.1              liblxswm/1.0-20190225 libxsm/1.0-20190225 py
gcc/9.1.0              liblxswm/1.0-20190225 libxsm/1.0-20190225 py
gdb/8.0                liblxswm/1.0-20190225 libxsm/1.0-20190225 py
gdb/8.2-py3            liblxswm/1.0-20190225 libxsm/1.0-20190225 py
gdbm/1.18.1            liblxswm/1.0-20190225 libxsm/1.0-20190225 py
gettext/0.19.8.1       liblxswm/1.0-20190225 libxsm/1.0-20190225 py
git-lfs/2.8.0           liblxswm/1.0-20190225 libxsm/1.0-20190225 py
git/2.9.3              liblxswm/1.0-20190225 libxsm/1.0-20190225 py
[skottap@login5.summit skottap]$ module load nsys/2020.1.1.65
[skottap@login5.summit skottap]$ nsys --version
NVIDIA Nsight Systems version 2020.1.1.65-085319d
```

# Load Nsys Module

\$ module avail

\$ module load nsight-systems/2020.1.1.65

# Create and run script

```
#!/bin/bash
#BSUB -P VEN201
#BSUB -W 2:00
#BSUB -nnodes 2
#BSUB -alloc_flags gpumps
#BSUB -o stdoutput.%J
#BSUB -e stderror.%J
cd /gpfs/alpine/world-shared/ven201/skottap/GROMACS_2020_NEW_VERSION/water_boxes/water-cut1.0_GMX50_bare/0048
export OMP_NUM_THREADS=7
jsrun -n 1 -a 6 -c 42 -g 6 -r 1 -l CPU-CPU -d plane:6 -b packed:7
--smpiargs="-disable gpu hooks" nsys profile -o
/gpfs/alpine/world-shared/ven201/skottap/GROMACS_2020_NEW_VERSION/new_%q{OMPI_COMM_WORLD_RANK} -f true --stats=true
/gpfs/alpine/world-shared/ven201/skottap/GROMACS_2020_NEW_VERSION/gromacs-2020/build/bin/gmx_mpi mdrun -ntomp 7 -pme gpu -npme 1 -noconfout -nb gpu -pin off -nsteps 10000
```

```
[skottap@login5.summit skottap]$ cd GROMACS_2020_NEW_VERSION/
[skottap@login5.summit GROMACS_2020_NEW_VERSION]$ vim batch.lsf
[skottap@login5.summit GROMACS_2020_NEW_VERSION]$ bsub batch.lsf
Job <954934> is submitted to default queue <batch>.
[skottap@login5.summit GROMACS_2020_NEW_VERSION]$ bjobs
JOBID    USER      STAT   SLOTS   QUEUE      START_TIME    FINISH_TIME   JOB_NAME
954934  skottap  PEND       -    batch            -              -          Not_Specified
```

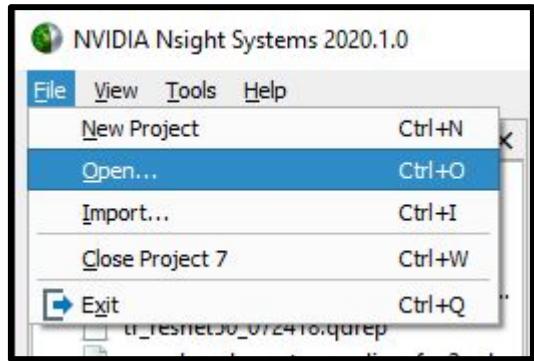
# Copy File back for Viewing

Copy report files back to x86\_64 Linux/Windows/Mac to view in GUI

```
rsync -avz --progress $USER@dtn.ccs.ornl.gov:/path/*.qdrep
```

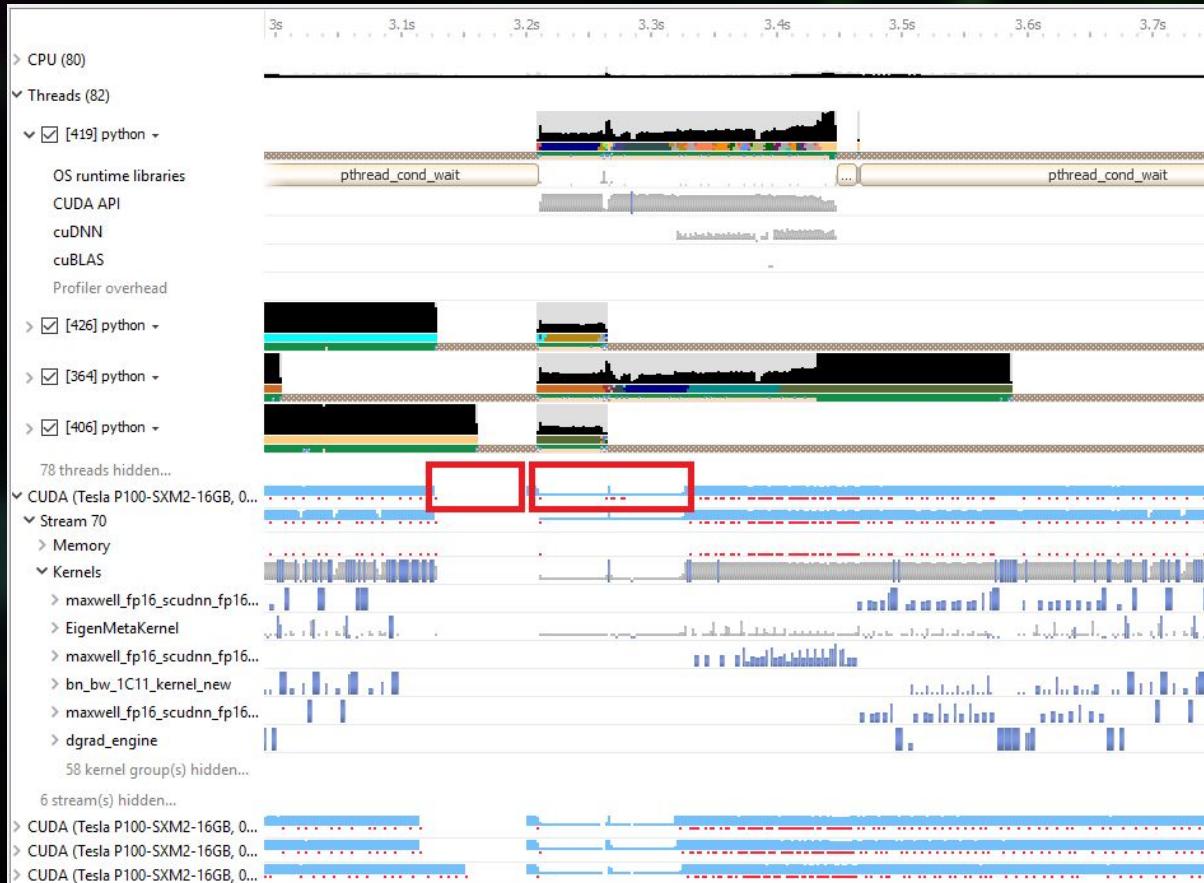
```
skottapalli@SKOTTAPALLI-PC:~$ rsync -avz --progress skottap@dtn.ccs.ornl.gov  
_NEW_VERSION/*.qdrep .  
*****  
NOTICE TO USERS  
  
This is a Federal computer system and is the property of the United States  
Government. It is for authorized use only. Users (authorized or  
unauthorized) have no explicit or implicit expectation of privacy.  
  
Any or all uses of this system and all files on this system may be  
intercepted, monitored, recorded, copied, audited, inspected, and disclosed  
to authorized site, Department of Energy, and law enforcement personnel, as  
well as authorized officials of other agencies, both domestic and foreign.  
By using this system, the user consents to such interception, monitoring,  
recording, copying, auditing, inspection, and disclosure at the discretion  
of authorized site or Department of Energy personnel.  
  
Unauthorized or improper use of this system may result in administrative  
disciplinary action and civil and criminal penalties. By continuing to use  
this system you indicate your awareness of and consent to these terms and  
conditions of use. LOG OFF IMMEDIATELY if you do not agree to the  
conditions stated in this warning.  
*****  
Enter PASSCODE:  
receiving incremental file list  
new_0.qdrep  
    212,738 100% 433.72kB/s  0:00:00 (xfr#1, to-chk=5/6)  
new_1.qdrep  
    213,255 100% 426.76kB/s  0:00:00 (xfr#2, to-chk=4/6)  
new_2.qdrep  
    213,199 100% 233.15kB/s  0:00:00 (xfr#3, to-chk=3/6)  
new_3.qdrep  
    208,793 100% 140.23kB/s  0:00:01 (xfr#4, to-chk=2/6)  
new_4.qdrep  
    206,400 100% 244.32kB/s  0:00:00 (xfr#5, to-chk=1/6)  
new_5.qdrep  
    209,189 100% 175.35kB/s  0:00:01 (xfr#6, to-chk=0/6)  
  
sent 138 bytes received 475,781 bytes 23,215.56 bytes/sec  
total size is 1,263,574 speedup is 2.66  
skottapalli@SKOTTAPALLI-PC:~$
```

# Open File for Visualization

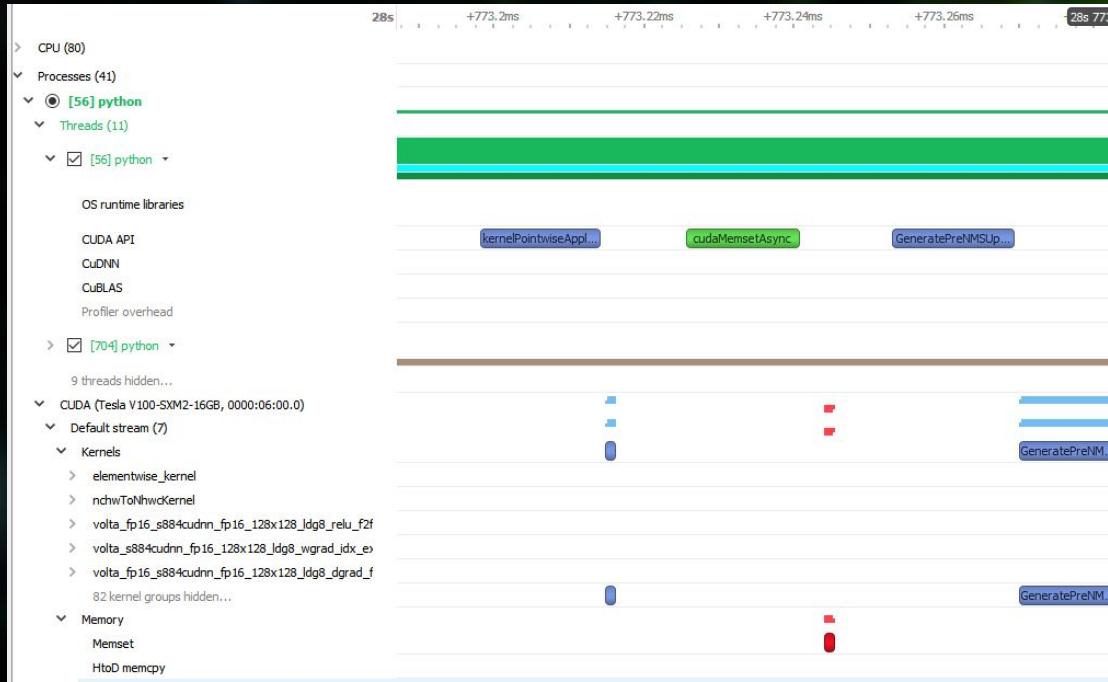


GUI available for Windows, Linux x86\_64, and MacOS. Not available for Linux on Power.

# Example problems



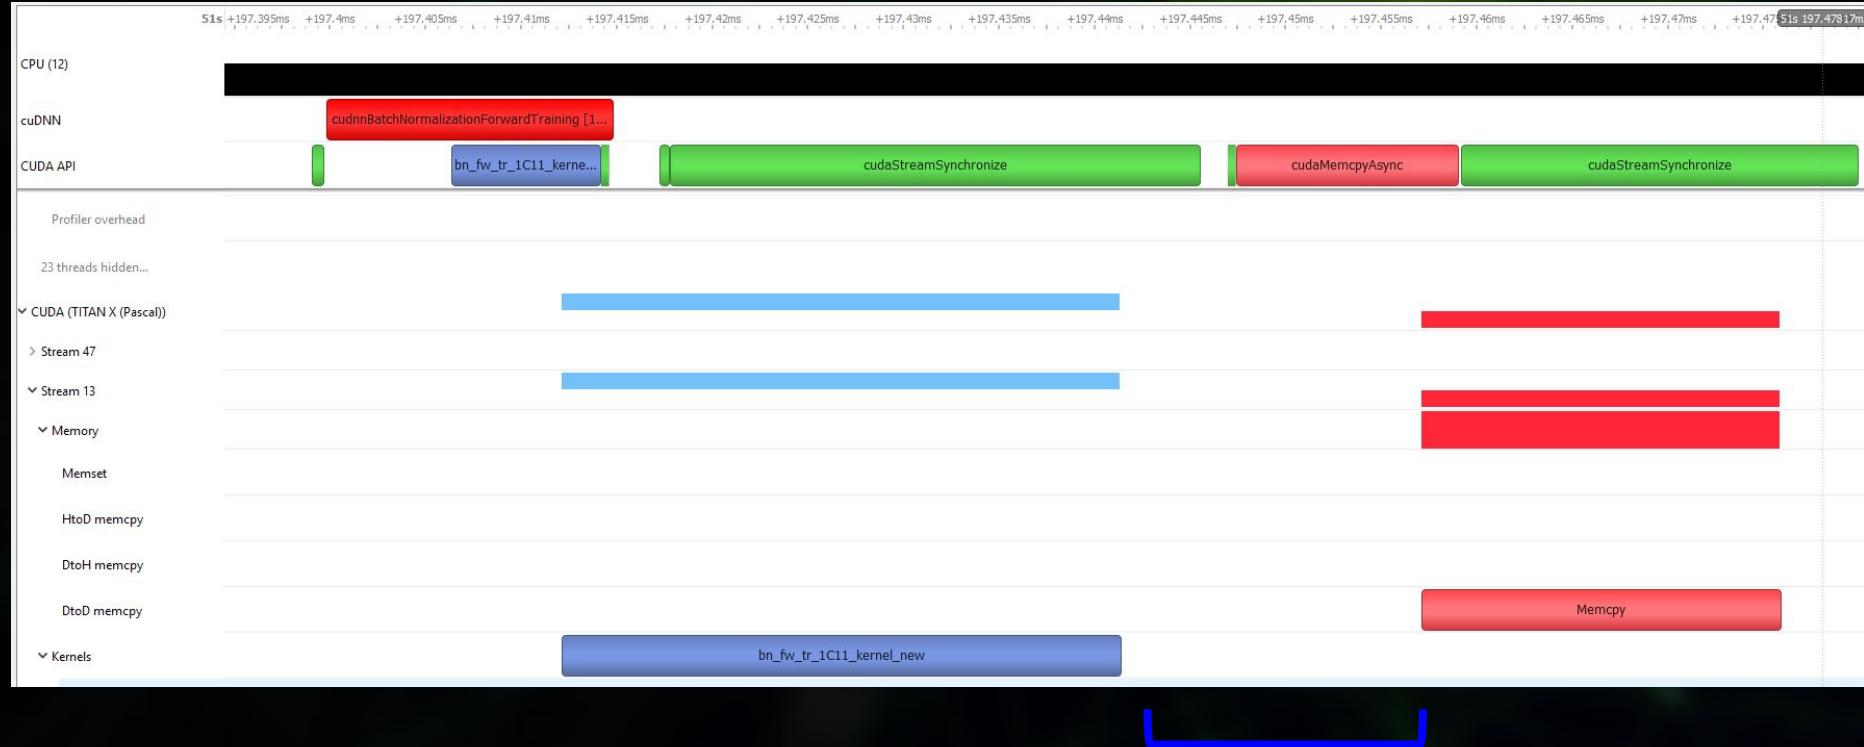
GPU idle and low utilization level of detail



Fusion opportunities  
CPU launch cost + small GPU work size  $\approx$  GPU sparse idle  
This can apply to DNN nodes/layers



**cudaMemcpyAsync behaving synchronous  
Device to host pageable memory  
Mitigate with pinned memory**



Example GPU idle caused by stream synchronization

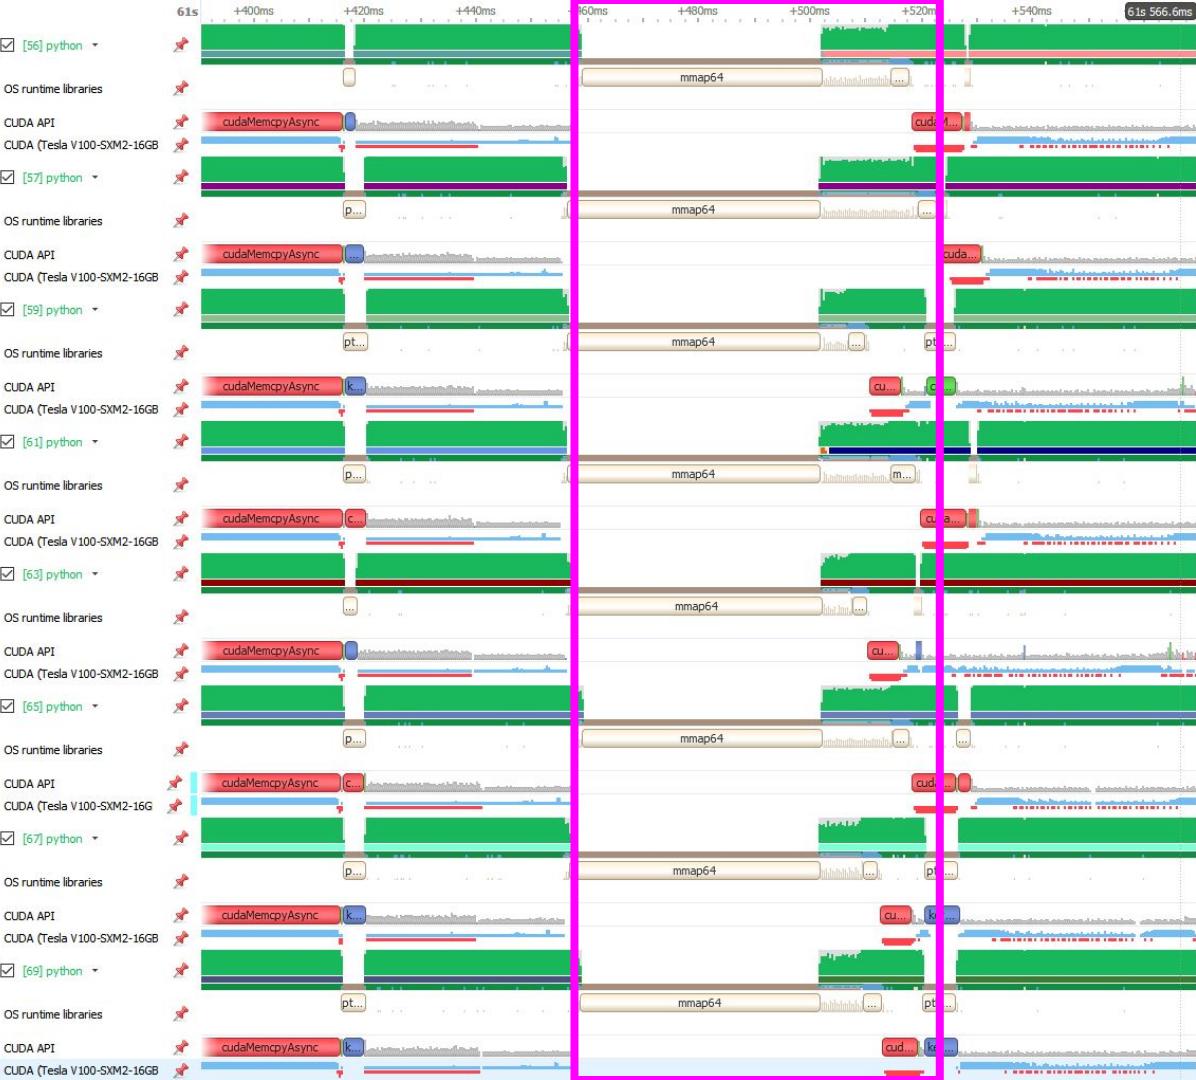
# OS Runtime API Trace

## Example: Mask-RCNN

Map/unmap hiccups

Mitigate by pipelining

- Map 1 batch ahead
- Unmap last batch
- Swap pointers here instead



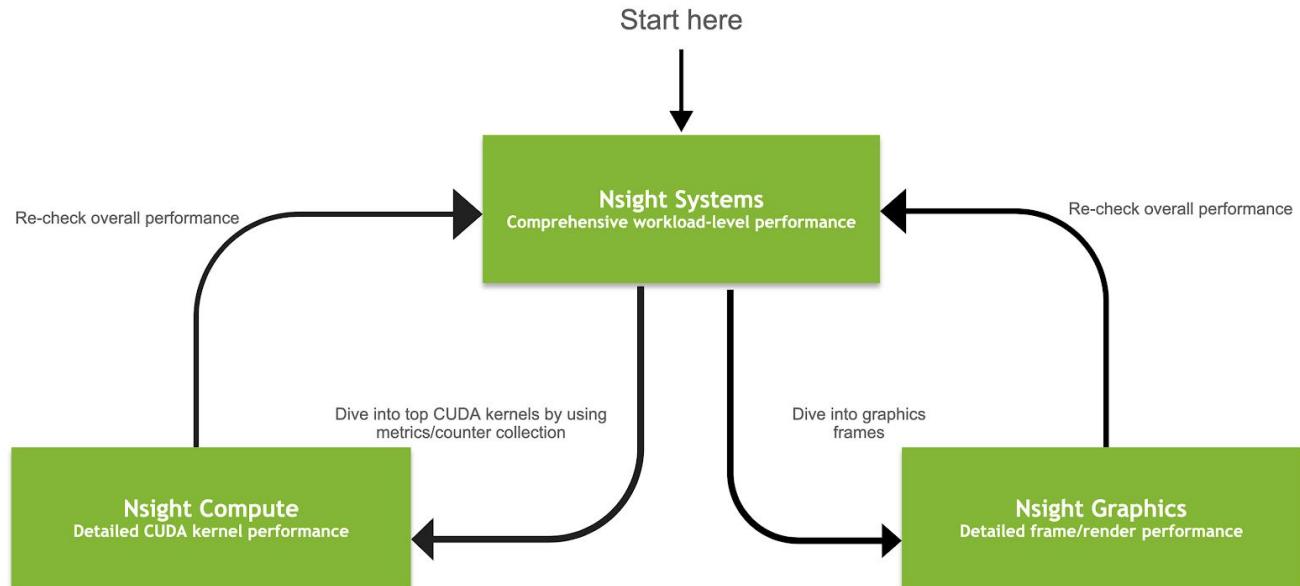
# Nsight Product Family

## Workflow

**Nsight Systems** -  
Analyze application  
algorithm system-wide

**Nsight Compute** -  
Debug/optimize CUDA  
kernel

**Nsight Graphics** -  
Debug/optimize graphics  
workloads





# THANK YOU!

**Download** <https://developer.nvidia.com/nsight-systems>

**NOTE: Website versions newer than CUDA Toolkit**

**Forums** <https://devtalk.nvidia.com>

**Email** [nsight-systems@nvidia.com](mailto:nsight-systems@nvidia.com)



# Backup