# **Nsight Systems – Introduction**



## **Typical Optimization Workflow**





### **NVIDIA Performance Analysis Tools**





THONS



### **PROFILING GPU APPLICATION**

Focusing System Operation



### **PROFILING GPU APPLICATION**

### Focusing GPU Computing



#### NVIDIA (Visual) Profiler / Nsight Compute



#### NVIDIA Supports them with cuDNN, cuBLAS, and so on

# System-Wide Application Tuning

#### Maximize your GPU Investment

Locate optimization opportunities

- Visualize millions of events on a timeline
- See gaps of unused CPU and GPU time

Balance your workload across multiple CPUs and GPUs

- CPU utilization and thread state
- GPU streams, kernels, memory transfers, etc.
- NIC performance metrics & tracing of network libraries

Multi-platform support

OpenACC

- Linux, Windows and Mac OS X (host-only)
- x86-64, Power9, ARM server, Tegra (Linux & QNX)



#### From the Macro to the Micro

#### Putting it all together











0-

0 -

0-

0-



More Science, Less Programming



This material is released by NVIDIA Corporation under the Creative Commons Attribution 4.0 International (CC BY 4.0)

135

138

BAR.SYNC 0x0

ISETP.GT.AND PD, PT, R11, BX3f, PT

ISETP.GT.AND PL, PT, R11, 0x1f, PT

BSSY B1, ex7f87d326fc5e



#### File <u>View Tools</u> <u>H</u>elp



# Command Line Interface (CLI) Statistics and Export to SQLite, JSON, etc.



## **Command Line Interface**

The Nsight Systems CLI provides several different commands

- Basic profiling session nsys profile ./app
- Interactive sessions (scriptable) nsys start|launch|stop|cancel nsys session list

nsys status|shutdown

• Statistics and export

#### nsys stats|export

(export to sqlite, hdf, text, json, info)







# CLI Profiling - Some Useful Switches

API tracing

- -t, --trace=cuda,nvtx,osrt,opengl
- (cublas,cusparse,cudnn,mpi,oshmem,ucx,openacc,openmp,vulkan,none)
- Overwrite existing report
- -f, --force-overwrite=[true|false]
- Summary statistics (profile output on command line)
- --stats=[true|false]
- Report file name
  - -o, --output=report# (patterns for hostname, PID and environment variables)



# CLI Profiling - Some Useful Switches

- Callstack sampling
  - -s, --sample=[cpu|none]
  - --sampling-period=number of CPU Instructions Retired events
  - -b, --backtrace=[lbr|fp|dwarf|none]
  - --samples-per-backtrace={1..12}

(The number of CPU IP samples collected for every CPU IP sample backtrace collected.)

- Set the paranoid level: "sudo sh -c 'echo 1 >/proc/sys/kernel/perf\_event\_paranoid'
- CUDA memory usage

OpenAC

- --cuda-memory-usage=[true|**false**]
- (Use nsys profile --help for a list of available options.)

# CLI Profiling - MPI Programs

#### Single Node

nsys profile [nsys\_args] mpirun [mpirun\_args] your\_executable

■ ⇒ This creates one report file.

#### Multiple Nodes

mpirun [mpirun\_args] nsys profile [nsys\_args] your\_executable

Set output report name with -o report\_name\_%q{OMPI\_COMM\_WORLD\_RANK}
 (for OpenMPI, PMI\_RANK for MPICH and SLURM\_PROCID for Slurm)

.⇒ This creates one report file	per MPI rank.	<pre>#!/bin/bash # OMPI_COMM_WORLD_LOCAL_RANK for node local rank</pre>
Profile only	specific ranks. ⇔	<pre>if [ \$OMPI_COMM_WORLD_RANK -eq 0 ]; then     nsys profile -t mpi "\$@" else</pre>
	This motorial i	"\$@" fi

## **CLI Profiling - Additional Output**

 WARNING: The command line includes a target application therefore the CPU contextswitch scope has been set to process-tree.
 Collecting data...

Temporary data is written to **/tmp/nvidia/nsight\_systems** by default. Set **TMPDIR** to specify another location.

#### • ... APP OUTPUT .

```
Processing events...
```

Saving temporary "/tmp/nsys-report-2b96-9038-d0bd-2600.qdstrm" file to disk...

```
Creating final output files...
```

Saved report file to "/tmp/nsys-report-2b96-9038-d0bd-2600.nsys-rep"

Report file moved to "/final/destination/report.nsys-rep"



### **CLI Stats Output**

- CUDA API, kernels and memory operations (by time and by size)
- **OS** runtime

•		NVTX Push-Pop Range Statistics:								
		Time(%)	Total Tim	ne (ns) In	nstances	Average	Minimum	Maximum	Range	
		80,0	16.193.3 3.680.9	76.534 37.672	21.268 8.072	761.396,0 456.013,0	6.823 1.726	2.438.848	MPI:MPI_Waitall MPI:MPI Allreduce	
		0,0	88.5	87.708	4	22.146.927,0	21.714.183	22.579.671	MPI:MPI_Init	
		0,0	24.2	82.234	21.268	1.141,0	872	16.931	MPI:MPI_Isend MPI:MPI_Irecv	
		0,0	8.7	93.048 14.158	4 28	2.198.262,0 136.219,0	2.181.292 2.302	2.215.232	MPI:MPI_Finalize MPI:MPI Barrier	
CUDA Kern	el Statistics:	0,0		66.630	32	8.332,0	832	104.980	MPI:MPI_Reduce	
Time(%)	Total Time (ns)	Instances	Average	Minimum	Maxim	um			Name	
37,0	7.344.350.347	8.568	857.183,0	851.546	5 876.	922 device_te	a_leaf_ppcg_	solve_calc_s	d_new(kernel_info_t, c	
36,0 10,0	7.229.217.310 2.050.961.573	8.568 2.010 1	843.746,0 .020.378,0	839.323 1.005.849	3 899. 9 1.036.	482 device_te 569 device_te	a_leaf_ppcg_ a_leaf_cg_so	solve_update lve_calc_ur(	_r(kernel_info_t, doub kernel_info_t, double,	
9,0	1.879.919.365	2.010	935.283,0	913.754 501_469	4 1.214. 571	808 device_te	a_leaf_cg_so	lve_calc_w(k	ernel_info_t, double*,	
0,0	61.606.046	30 2	.053.534,0	1.176.600	2.890.	541 device_te	a_leaf_calc_	rrn(kernel_i	nfo_t, double const*,	
0,0	58.094.845	12.550	5.088,0	1.88/	15.	184 Vold redu	crion <double< td=""><td>, TREDUCTION</td><td>_ine is the int, double )</td></double<>	, TREDUCTION	_ine is the int, double )	

This material is released by NVIDIA Corporation under the Creative Commons Attribution 4.0 International (CC BY 4.0)

# NVTX NVIDIA Tool Extension Library



# **NVIDIA** Tool Extension Library

NVTX provides means to correlate the profile data with the application code.

```
#include <nvtx3/nvToolsExt.h>
...
nvtxMark("Point in time");
nvtxRangePush("Name of your code region");
// your code goes here
nvtxRangePop();
nvtxRangeId_t rid = nvtxRangeStart("Name of another code region"
// stop might be on another thread
nvtxRangeEnd(rid)
```

Add nvtx to the tracing options of nsys profile.) OpenACC

# NVTX

Domains and Resource Naming

```
nvtxDomainHandle_t yourDomain = nvtxDomainCreateA("My Domain");
nvtxEventAttributes_t evtAttr = {0};
evtAttr.version = NVTX_VERSION;
evtAttr.size = NVTX_EVENT_ATTRIB_STRUCT_SIZE;
evtAttr.messageType = NVTX_MESSAGE_TYPE_ASCII;
eventAttrib.message.ascii = "My Range";
nvtxDomainRangePushEx(yourDomain, &evtAttr);
// your code
```

nvtxDomainRangePop(yourDomain);

```
// #include <sys/syscall.h>
nvtxNameOsThreadA(syscall(SYS gettid), "My Thread");
```



#### **NVTX** Registered Strings

// your code
nvtxDomainRangePop(yourDomain);



#### https://github.com/NVIDIA/NVTX

#### NVTX Fortran Bindings

The NVIDIA HPC SDK Fortran compiler provides NVTX bindings.

- libnvhpcwrapnvtx.[a]so] has to be linked
- Documentation can be found here: <u>https://docs.nvidia.com/hpc-sdk/compilers/fortran-cuda-interfaces/index.html#cfnvtx-runtime</u>

```
use nvtx
...
call nvtxStartRange("YourRange")
! some Fortran code
call nvtxEndRange
```

For other compilers, write your own Fortran NVTX bindings or use existing ones, e.g. <a href="https://raw.githubusercontent.com/maxcuda/NVTX\_example/master/nvtx.f90">https://raw.githubusercontent.com/maxcuda/NVTX\_example/master/nvtx.f90</a>



#### **NVTX** Automatic code instrumentation

The NVIDIA HPC SDK and GNU compilers provide an automatic code instrumentation mechanism using respectively the options -*Minstrument* and *-finstrument-functions* with the \_\_\_cyg\_profile\_\* funcs

C/C++ example with gcc/g++ and option -finstrument-functions :

https://developer.nvidia.com/blog/cuda-pro-tip-generate-custom-application-profile-timelines-nvtx/

 Fortran documentation with nvfortran: <u>https://docs.nvidia.com/hpc-sdk/compilers/fortran-cuda-interfaces/index.html#cfnvtx-auto</u>



## NVTX

"Python developers can either use decorators @nvtx.annotate() or a context manager with nvtx.annotate(..)"

**Get NVTX Python module:** python -m pip install nvtx

Also see this blog or the docs: https://nvtx.readthedocs.io/en/latest/index.html

PyTorch CUDA provides NVTX bindings



# Application



#### **APPLICATION**

**Molecular Simulation** 

#### RDF

The radial distribution function (RDF) denoted in equations by g(r) defines the probability of finding a particle at a distance r from another tagged particle.





# RDF

#### Pseudo Code - C



ATHONS

**OpenACC** 

Across Frames

• Find Distance

• Reduction

# Pseudo Code - Fortran

RDF

#### do iconf=1,nframes if (mod(iconf,1).eq.0) print\*,iconf

do i=1,natoms do j=1,natoms

dx=x(iconf,i)-x(iconf,j)
dy=y(iconf,i)-y(iconf,j)
dz=z(iconf,i)-z(iconf,j)

 $\begin{array}{c} r=dsqrt(dx^{**}2+dy^{**}2+dz^{**}2)\\ if(r<cut)then\\ g(ind)=g(ind)+1.0d0\\ end if\end{array}$ 

enddo

enddo enddo

OpenACC



Across Frames

• Find Distance

Reduction