



VIENNA
SCIENTIFIC
CLUSTER



DEEP
LEARNING
INSTITUTE

Accelerated Computing with OpenACC and Deep Learning

27 – 29 April 2021

VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

MODULE TWO: PROFILING

Dr. Volker Weinberg | LRZ | 27.04.2021

MODULE OVERVIEW

Topics to be covered

- Compiling and profiling sequential code
- Explanation of multicore programming
- Compiling and profiling multicore code

COMPILING SEQUENTIAL CODE

NVIDIA'S HPC COMPILERS (AKA PGI)

NVIDIA Compiler Names (PGI names still work)

- `nvc` - The command to compile C code (formerly known as 'pgcc')
- `nvc++` - The command to compile C++ code (formerly known as 'pgc++')
- `nvfortran` - The command to compile Fortran code (formerly known As `pgfortran/pgf90/pgf95/pgf77`)
- The `-fast` flag instructs the compiler to optimize the code to the best of its abilities

```
$ nvc -fast main.c  
$ nvc++ -fast main.cpp  
$ pgfortran -fast main.F90
```

```
$ pgcc -fast main.c  
$ pgc++ -fast main.cpp  
$ pgfortran -fast main.F90
```

NVIDIA'S HPC COMPILERS (AKA PGI)

-Minfo flag

- The Minfo flag will instruct the compiler to print feedback about the compiled code
- -Minfo=accel will give us information about what parts of the code were accelerated via OpenACC
- -Minfo=opt will give information about all code optimizations
- -Minfo=all will give all code feedback, whether positive or negative

```
$ pgcc -fast -Minfo=all main.c  
$ pgc++ -fast -Minfo=all main.cpp  
$ pgfortran -fast -Minfo=all main.f90
```

NVIDIA NSIGHT FAMILY

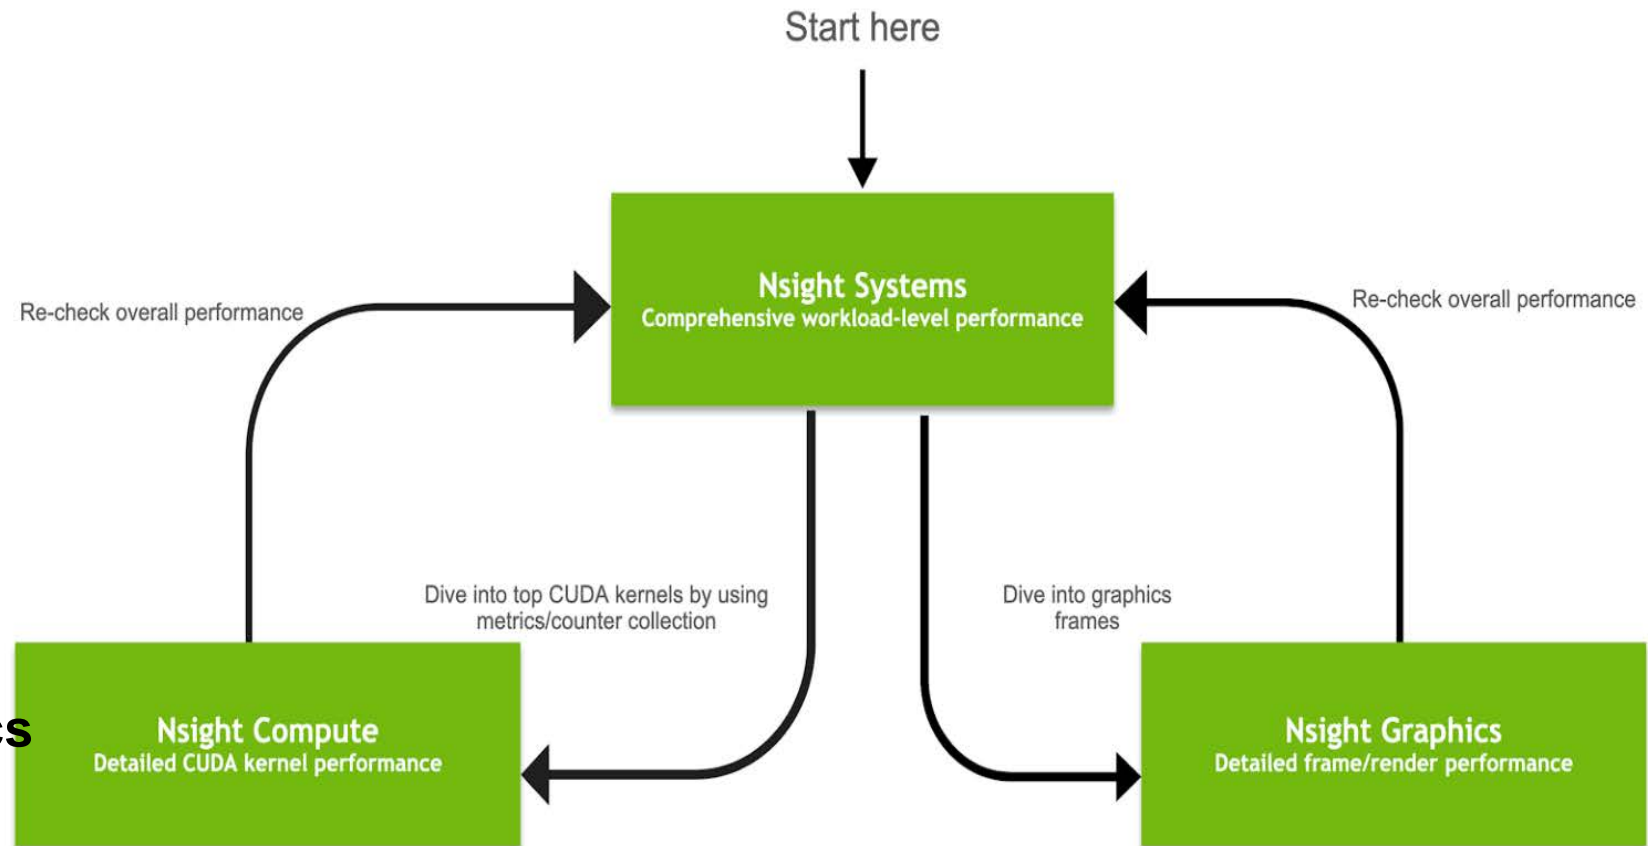
Nsight Product Family

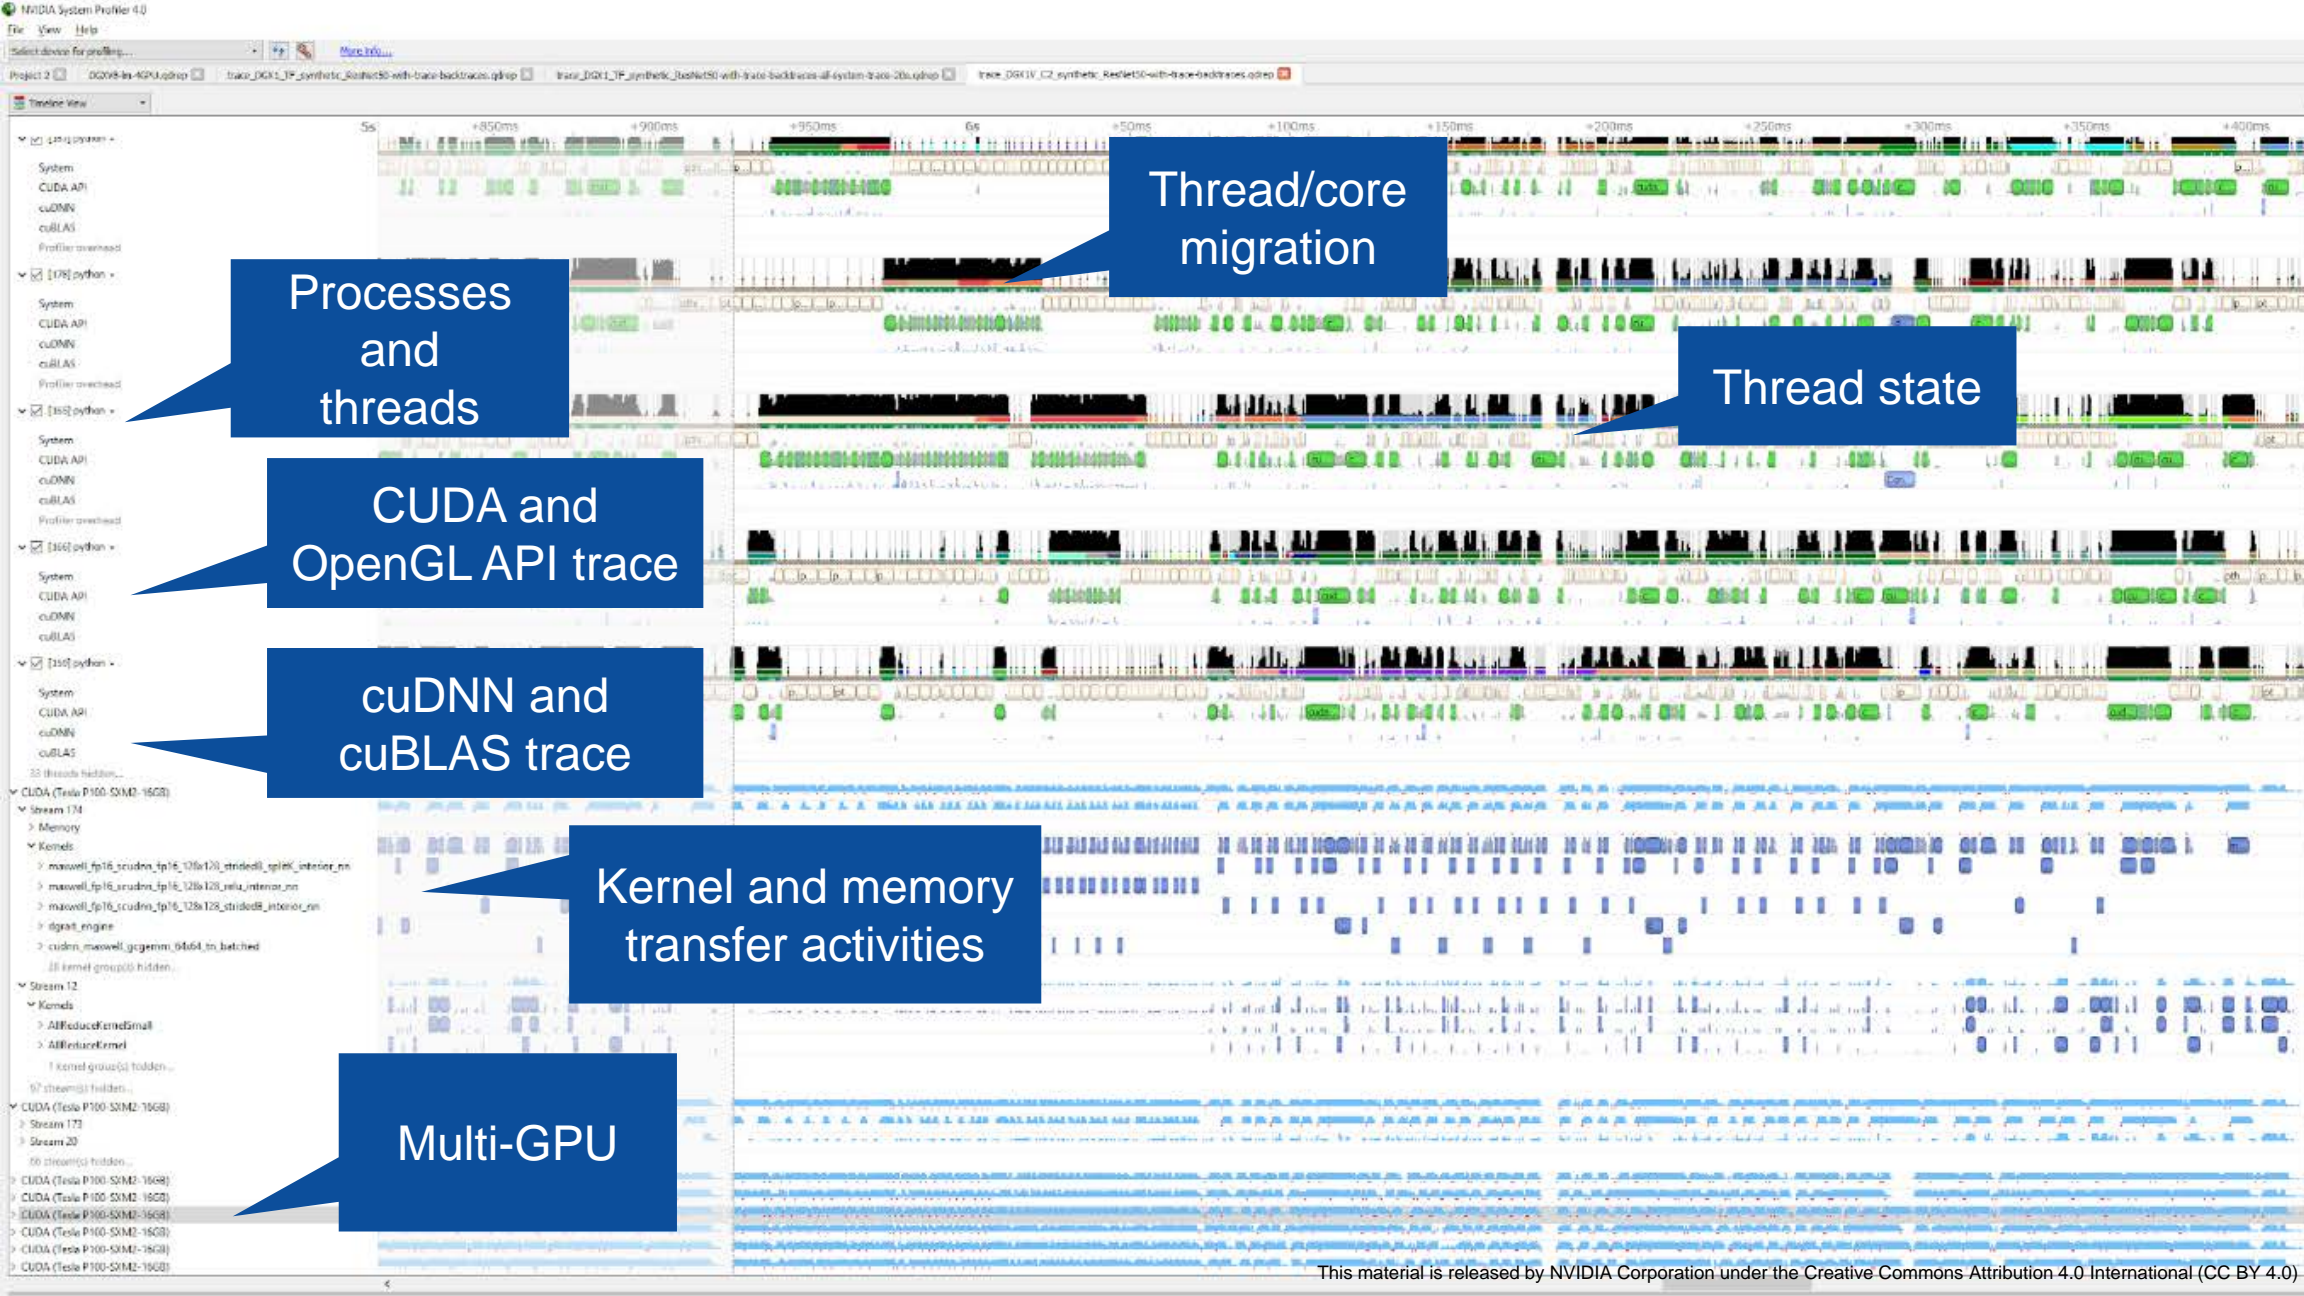
Workflow

Nsight Systems -
Analyze application
algorithm system-wide

Nsight Compute -
Debug/optimize CUDA
kernel

Nsight Graphics -
Debug/optimize graphics
workloads





Processes and threads

Thread/core migration

Thread state

CUDA and OpenGL API trace

cuDNN and cuBLAS trace


Kernel and memory transfer activities

Multi-GPU

PROFILING SEQUENTIAL CODE

OPENACC DEVELOPMENT CYCLE

- **Analyze** your code to determine most likely places needing parallelization or optimization.
- **Parallelize** your code by starting with the most time consuming parts, check for correctness and then analyze it again.
- **Optimize** your code to improve observed speed-up from parallelization.



PROFILING SEQUENTIAL CODE

Step 1: Run Your Code

Record the time it takes for your sequential program to run.

Note the final results to verify correctness later.

Always run a problem that is representative of your real jobs.

Terminal Window

```
$ gcc -fast jacobi.c laplace2d.c
$ ./a.out
  0, 0.250000
 100, 0.002397
 200, 0.001204
 300, 0.000804
 400, 0.000603
 500, 0.000483
 600, 0.000403
 700, 0.000345
 800, 0.000302
 900, 0.000269
total: 39.432648 s
```

PROFILING WITH NSIGHT SYSTEM AND NVTX

PROFILING SEQUENTIAL CODE

Using Command Line Interface (CLI)

NVIDIA Nsight Systems CLI provides

- Simple interface to collect data
- Can be copied to any system and analysed later
- Profiles both serial and parallel code
- For more info enter `nsys --help` on the terminal

To profile a serial application with NVIDIA Nsight Systems, we use NVIDIA Tools Extension (NVTX) API functions in addition to collecting backtraces while sampling.

PROFILING SEQUENTIAL CODE

NVIDIA Tools Extension API (NVTX) library

What is it?

- A C-based Application Programming Interface (API) for annotating events
- Can be easily integrated to the application
- Can be used with NVIDIA Nsight Systems

Why?

- Allows manual instrumentation of the application
- Allows additional information for profiling (e.g: tracing of CPU events and time ranges)

How?

- Import the header only C library `nvToolsExt.h`
- Wrap the code region or a specific function with `nvtxRangePush()` and `nvtxRangePop()`

```

#include <string.h>
#include <stdio.h>
#include <stdlib.h>
#include <omp.h>
#include "laplace2d.h"
#include <nvtx3/nvToolsExt.h>

int main(int argc, char** argv)
{
    const int n = 4096;
    const int m = 4096;
    const int iter_max = 1000;

    const double tol = 1.0e-6;
    double error = 1.0;

    double *restrict A = (double*)malloc(sizeof(double)*n*m);
    double *restrict Anew = (double*)malloc(sizeof(double)*n*m);

    nvtxRangePushA("init");
    initialize(A, Anew, m, n);
    nvtxRangePop();

    printf("Jacobi relaxation Calculation: %d x %d mesh\n", n, m);

    double st = omp_get_wtime();
    int iter = 0;

    nvtxRangePushA("while");
    while ( error > tol && iter < iter_max )
    {
        nvtxRangePushA("calc");
        error = calcNext(A, Anew, m, n);
        nvtxRangePop();

        nvtxRangePushA("swap");
        swap(A, Anew, m, n);
        nvtxRangePop();

        if(iter % 100 == 0) printf("%5d, %0.6f\n", iter, error);

        iter++;
    }
    nvtxRangePop();

    double runtime = omp_get_wtime() - st;

    printf(" total: %f s\n", runtime);

    deallocate(A, Anew);

    return 0;
}

```

jacobi.c
(starting and ending of ranges are highlighted with the same color)

- t Selects the APIs to be traced (nvtx in this example)
- status if true, generates summary of statistics after the collection
- b Selects the backtrace method to use while sampling. The option dwarf uses DWARF's CFI (Call Frame Information).
- force-overwrite if true, overwrites the existing results
- o sets the output (qdrep) filename

```

mozhgank@prmr-dgx-32:~/Code/openacc-training-materials/labs/module4/English/C/solutions/parallel$ nsys profile -t nvtx --stats=true -b dwarf --force-overwrite true -o laplace-seq ./laplace-seq
Collecting data...
Jacobi relaxation calculation: 4096 x 4096 mesh
  0, 0.250000
 100, 0.002397
 200, 0.001204
 300, 0.000804
 400, 0.000603
 500, 0.000483
 600, 0.000403
 700, 0.000345
 800, 0.000302
 900, 0.000269
total: 55.754501 s
Processing events...
Capturing symbol files...
Saving intermediate "/home/mozhgank/Code/openacc-training-materials/labs/module4/English/C/solutions/parallel/laplace-seq.qdstrm" file to disk...
Importing [=====100%]
Saved report file to "/home/mozhgank/Code/openacc-training-materials/labs/module4/English/C/solutions/parallel/laplace-seq.qdrep"
Exporting 70802 events: [=====100%]
Exported successfully to
/home/mozhgank/Code/openacc-training-materials/labs/module4/English/C/solutions/parallel/laplace-seq.sqlite
Generating NVTX Push-Pop Range Statistics...
NVTX Push-Pop Range Statistics (nanoseconds)
Time(%)  Total Time  Instances  Average  Minimum  Maximum  Range
-----
49.9     55754497966    1  55754497966.0  55754497966  55754497966  while
26.5     29577817696   1000  29577817.7   29092956    65008545  calc
23.4     26163892482   1000  26163892.5   25761418    60129514  swap
0.1      137489808     1    137489808.0  137489808   137489808  init

```

NVTX range statistics

"calc" region (calcNext function) takes 26.6%
"swap" region (swap function) takes 23.4% of total execution time


Open laplace-seq.qdrep with Nsight System GUI to view the timeline

PROFILING SEQUENTIAL CODE

Using Nsight Systems

Open the generated report files (*.qdrep) from command line in the Nsight Systems profiler.

File > Open



PROFILING SEQUENTIAL CODE

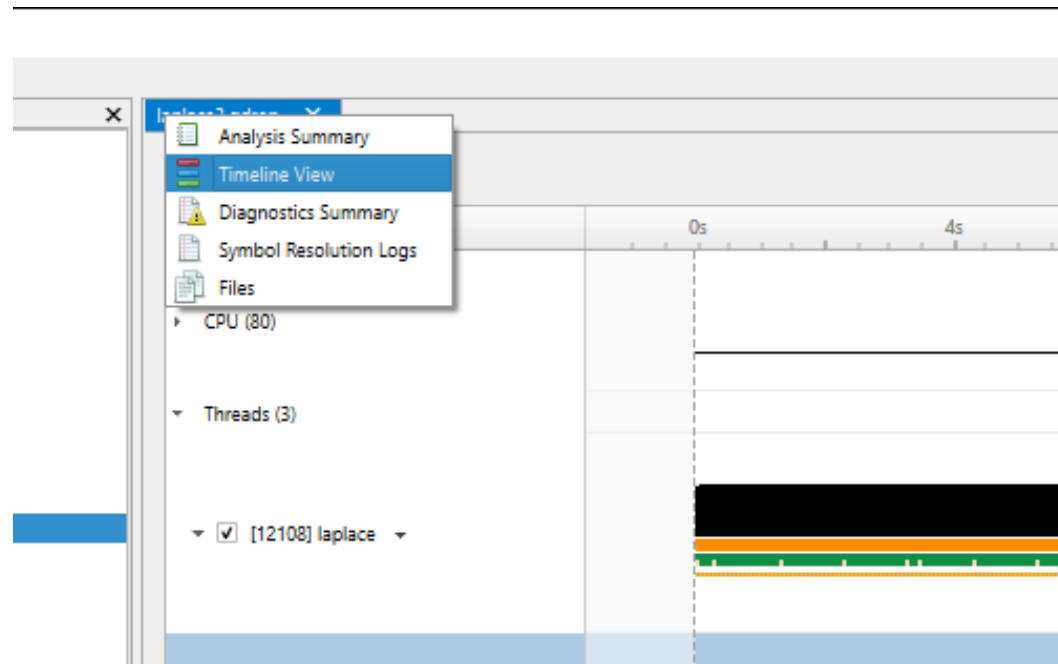
Using Nsight Systems

Navigate through the “view selector”.

“Analysis summary” shows a summary of the profiling session. To review the project configuration used to generate this report, see next slide.

“Timeline View” contains the timeline at the top, and a bottom pane that contains the events view and the function table.

Read more: <https://docs.nvidia.com/nsight-systems>



PROFILING SEQUENTIAL CODE

Using Nsight Systems

Analysis Summary

Profiling session duration: 00:55.623

Total number of threads	3
Number of events collected	70,773
Report size	851.22 KiB
Report capture date	19 March 2020 08:01:16
Host computer	prm-dgx-28
Profiling stop reason	Stopped by user
Imported from	/home/mozhgank/Codes/opensource/training-materials/labs/module2/English/C/replace.c.qdstrm
Import host computer	prm-dgx-28
CLI command used	nsys profile -t mtr -stats=true --force-overwrite true -o replace3_faplace

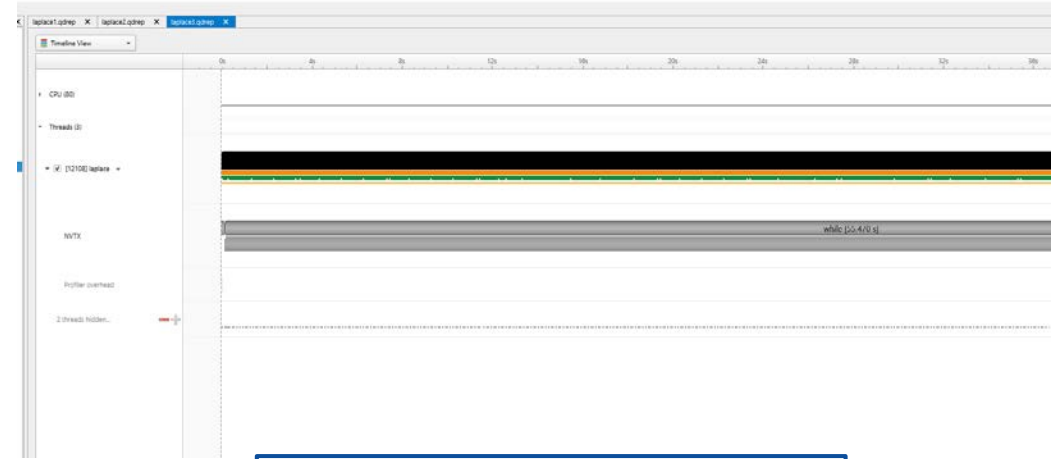
[Show report file in folder](#)

prm-dgx-28 (0:1)

Target

Target name	prm-dgx-28
Platform	Linux
OS	Ubuntu 18.04.3 LTS
Hardware platform	x86_64
Serial number	Local (CLI)
CPU description	Intel(R) Xeon(R) CPU E5-2698 v4 @ 2.20GHz
CUDA driver version	10.2
NVIDIA driver version	440.33.01
ODI rootkit mode	disabled

Analysis Summary



Timeline view
(charts and the hierarchy on the top pane)

Events View

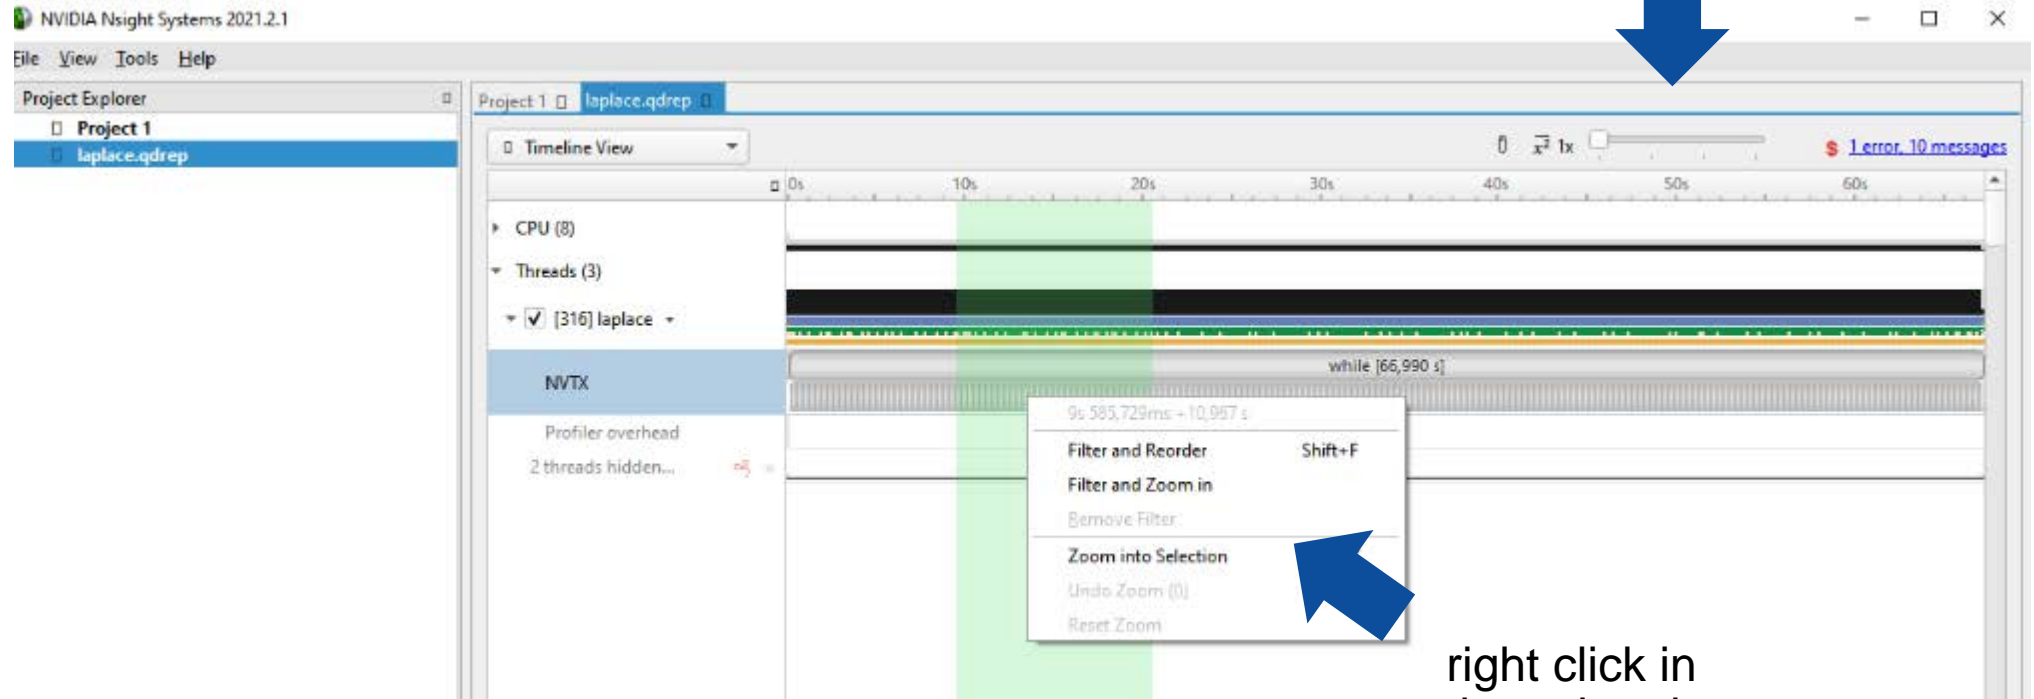
#	Name	Duration	PID	Start
1	+ [1010E] mtr	153.102 ms	12108	0.0095601s
2	+ [1010E] while	15.470 s	12108	0.14700s

Timeline view
(event view and function table on the bottom pane)

PROFILING SEQUENTIAL CODE

Using Nsight Systems

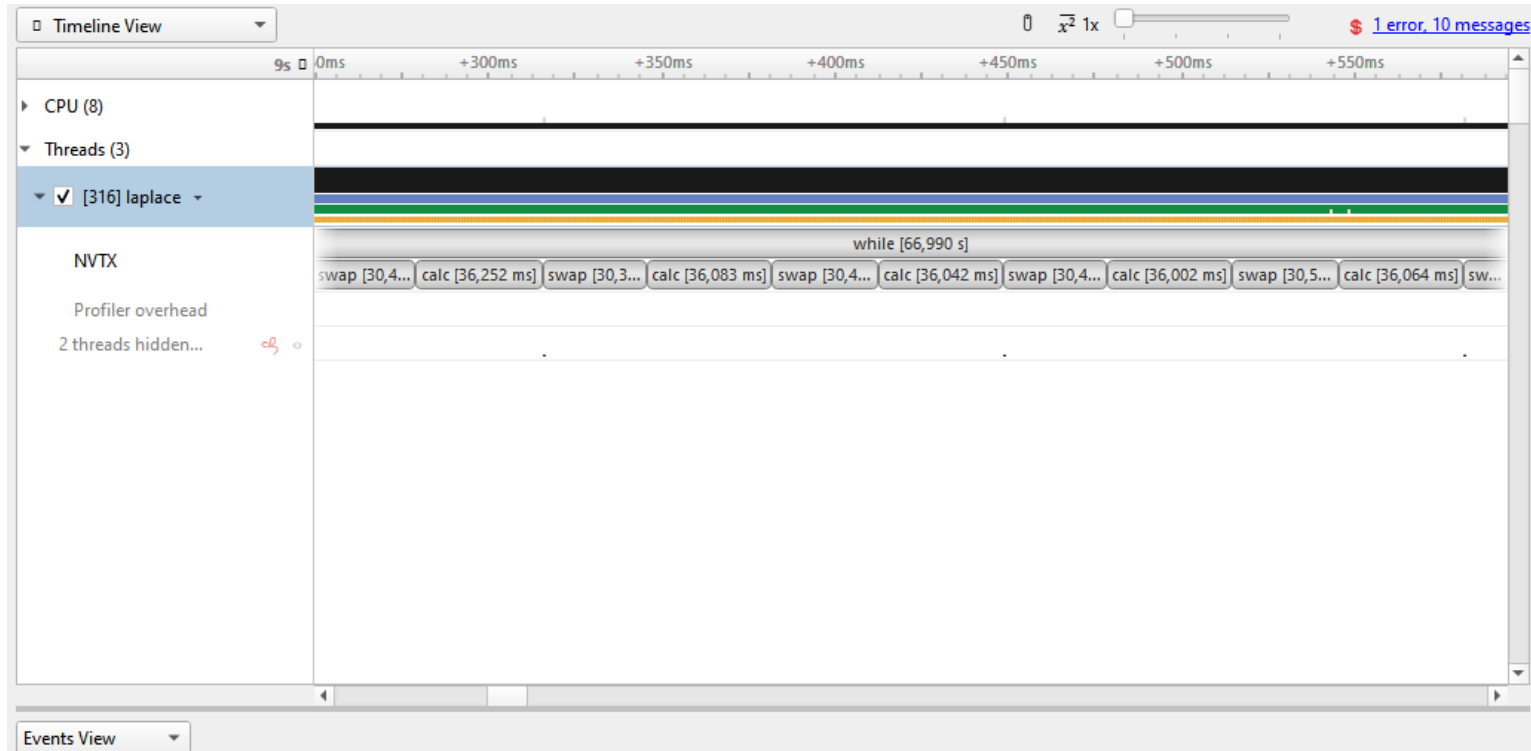
Enlarge view!



right click in
selected region
and Zoom into
selection!

PROFILING SEQUENTIAL CODE

Using Nsight Systems

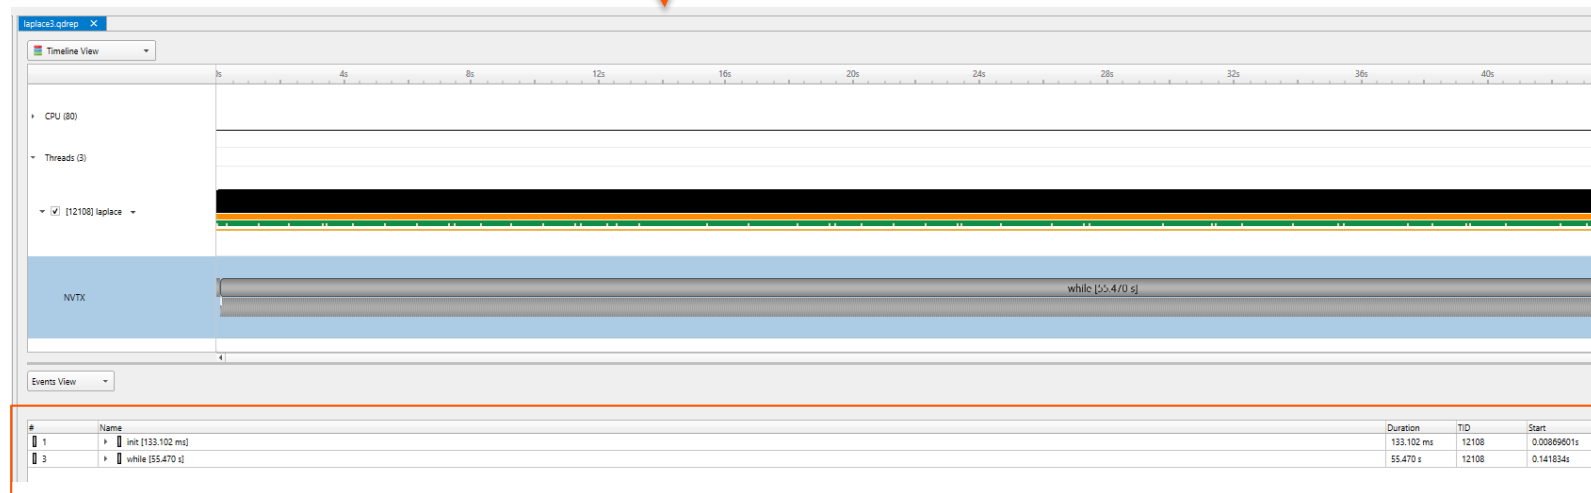
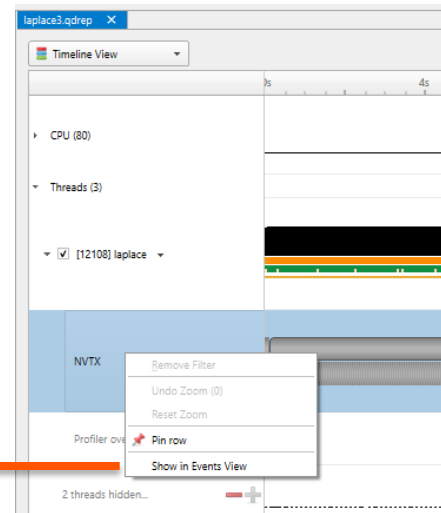


PROFILING SEQUENTIAL CODE

Viewing captured NVTX events and time ranges via Nsight Systems GUI

From the Timeline view, right click on the “NVTX” from the top pane and choose “Show in Events View”.

From the bottom pane, you can now see name of the events captured with the duration.

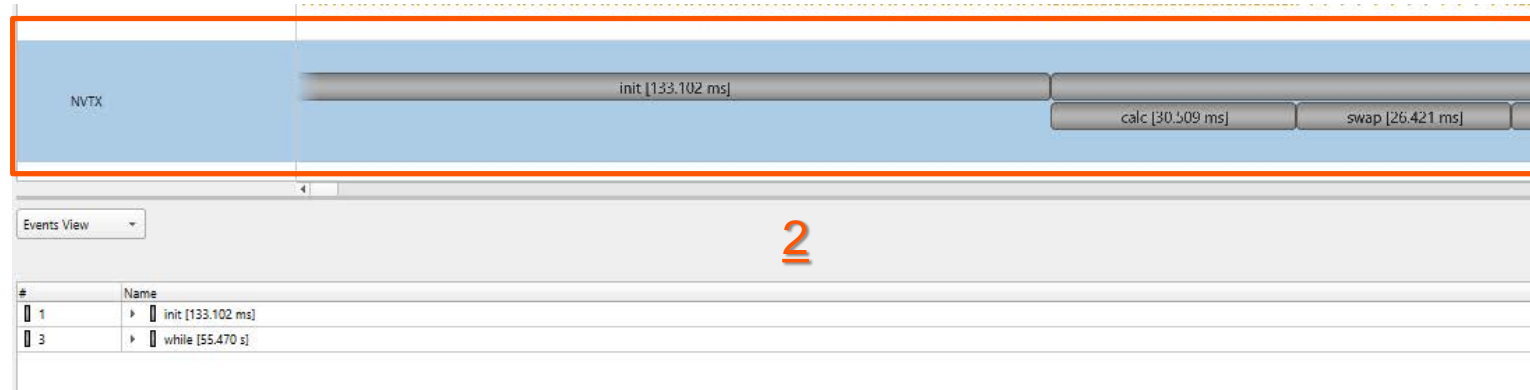
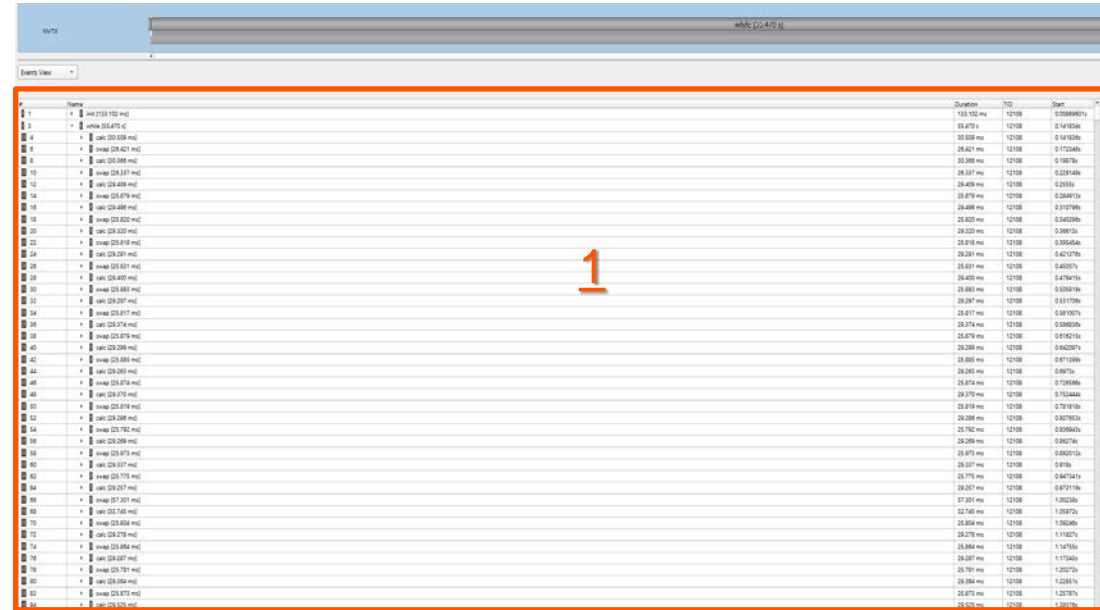


PROFILING SEQUENTIAL CODE

Viewing captured NVTX events and time ranges via Nsight Systems GUI

To have a closer look at events captured by NVTX APIs from GUI, you have two options:

1. Double click on “while” from the bottom pane and expand it.
2. Hover over the top pane and zoom in the chart on the right.

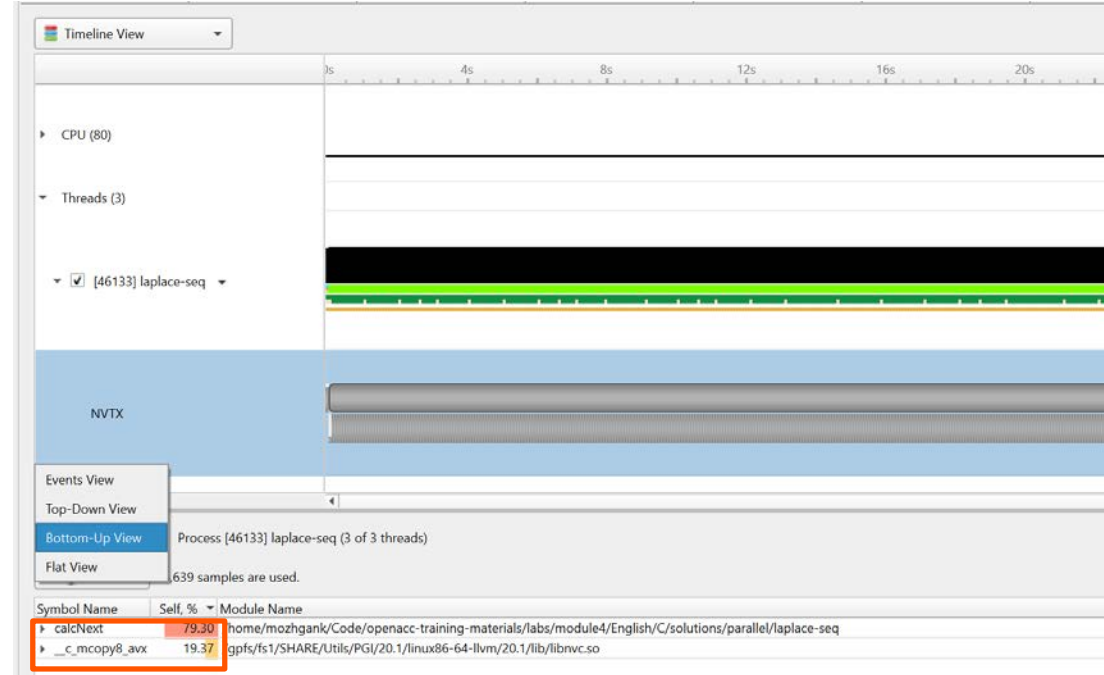


PROFILING SEQUENTIAL CODE

Viewing collected backtraces via Nsight Systems GUI

In the bottom pane, from the function table mode, choose the “Bottom-Up View” trying to find specific bottleneck functions that can be optimized.

In this example, we can see the “calcNext” and “c_mcopy8” (this is actually a compiler optimization that is being applied to our “swap” function) are places that our application is spending more of its time.



PROFILING SEQUENTIAL CODE

Step 2: Profile Your Code

Obtain detailed information about how the code ran.


This can include information such as:

- Total runtime
- Runtime of individual routines
- Hardware counters

Identify the portions of code that took the longest to run. We want to focus on these “hotspots” when parallelizing.

Lab Code: Laplace Heat Transfer

Total Runtime: 39.43 seconds



PROFILING SEQUENTIAL CODE

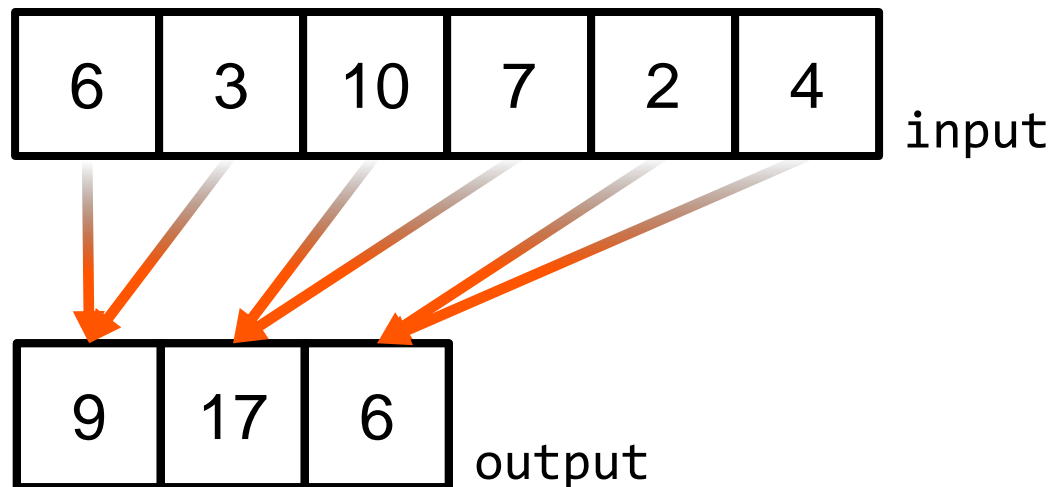
Step 3: Identify Parallelism

Observe the loops contained within the identified hotspots

Are these loops parallelizable?
Can the loop iterations execute independently of each other?
Are the loops multi-dimensional, and does that make them very large?

Loops that are good to parallelize tend to have a lot of iterations to map to parallel hardware.

```
void pairing(int *input, int *output, int N){  
    for(int i = 0; i < N; i++){  
        output[i] = input[i*2] + input[i*2+1];  
    }  
}
```



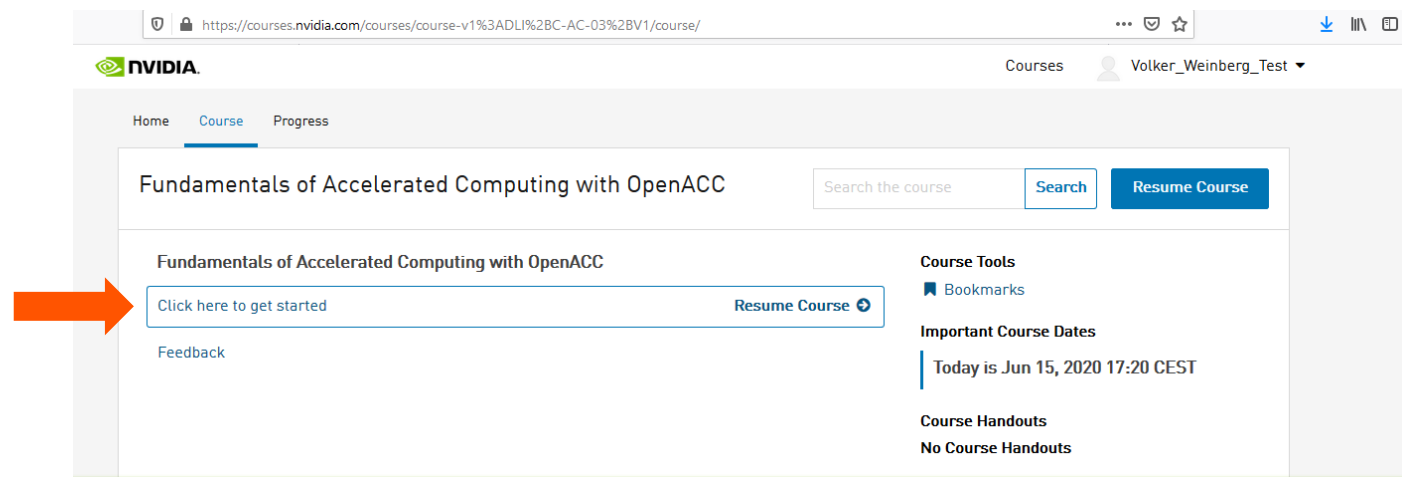
PLEASE START LAB NOW!

TRAINING SETUP

- To get started, follow these steps:
- Create an NVIDIA Developer account at <http://courses.nvidia.com/join> Select "Log in with my NVIDIA Account" and then "Create Account"
- Visit <http://courses.nvidia.com/dli-event> and enter the event code

VSC_OACC_AMBASSADOR_AP21



TRAINING SETUP



The screenshot shows the NVIDIA course page for "Fundamentals of Accelerated Computing with OpenACC". The page includes a search bar, a "Resume Course" button, and a "Click here to get started" button. An orange arrow points to the "Click here to get started" button.

Fundamentals of Accelerated Computing with OpenACC

[Bookmark this page](#)



The footer of the course page features the NVIDIA and Deep Learning Institute logos on the left and a "START" button with a play icon on the right. An orange arrow points to the "START" button.

To get started with this live GPU enabled interactive content please click the "Start" button on the top right of this block.

This will launch a pre-configured GPU workstation, it may take 5-10 minutes.


PROFILING MULTICORE CODE

PROFILING MULTICORE CODE

What is multicore?

- *Multicore* refers to using a CPU with multiple computational cores as our parallel device
- These cores can run independently of each other, but have shared access to memory
- Loop iterations can be spread across CPU threads and can utilize SIMD/vector instructions (SSE, AVX, etc.)
- Parallelizing on a multicore CPU is a good starting place, since data management is unnecessary


CPU



PROFILING MULTICORE CODE

Using a multicore CPU with OpenACC

- OpenACC's generic model involves a combination of a host and a device
- Host generally means a CPU, and the device is some parallel hardware
- When running with a multicore CPU as our device, typically this means that our host/device will be the same
- This also means that their memories will be the same



PROFILING MULTICORE CODE

Compiling code for a specific parallel hardware

- The '-ta' flag will allow us to compile our code for a specific, target parallel hardware
- 'ta' stands for "Target Accelerator," an accelerator being another way to refer to a parallel hardware
- Our OpenACC code can be compiled for many different kinds of parallel hardware without having to change the code

```
$ pgcc -fast -Minfo=accel -ta=multicore laplace2d.c  
calcNext:  
    35, Generating Multicore code  
    36, #pragma acc loop gang
```

PROFILING MULTICORE CODE

Compiling code for a specific parallel hardware

- nsys profile -t nvtx --stats=true --force-overwrite true -o laplace_parallel ./laplace_parallel

NVTX Push-Pop Range Statistics:

Time (%)	Total Time (ns)	Instances	Average	Minimum	Maximum	Range
49.9	24908340742	1	24908340742.0	24908340742	24908340742	while
26.4	13167317033	1000	13167317.0	9986457	52044034	calc
23.4	11711313301	1000	11711313.3	8693117	62627309	swap
0.4	175394843	1	175394843.0	175394843	175394843	init

Report file moved to "/home/openacc/labs/module2/English/C/laplace_parallel.qdrep"

Report file moved to "/home/openacc/labs/module2/English/C/laplace_parallel.sqlite"

PROFILING OPENACC CODE

PARALLEL VS SEQUENTIAL

Compiler feedback

Have a close look at the PGI compiler feedback for both sequential and parallel implementation of the application.

It provides information about how your program was optimized or why a particular optimization was not made.

Note: Adding `-Minfo` flag or `-Minfo=accel` or `-Minfo=all` when compiling, will enable compiler feedback messages, giving details about the parallel code generated.

```
jacobi.c:
main:
    59, Loop not vectorized/parallelized: potential early exits
laplace2d.c:
"laplace2d.c", line 24: warning: function "memset" declared implicitly
    memset(A, 0, n * m * sizeof(double));
    ^
initialize:
    27, Generated vector simd code for the loop
calcNext:
    40, Generated vector simd code for the loop containing reductions
swap:
    56, Memory copy idiom, loop replaced by call to __c_mcopy8
```

Sequential

```
jacobi.c:
main:
    59, Loop not vectorized/parallelized: potential early exits
laplace2d.c:
"laplace2d.c", line 24: warning: function "memset" declared implicitly
    memset(A, 0, n * m * sizeof(double));
    ^
initialize:
    27, Generated vector simd code for the loop
calcNext:
    35, Generating Tesla code
    37, #pragma acc loop gang /* blockIdx.x */
        Generating implicit reduction(max:error)
    40, #pragma acc loop vector(128) /* threadIdx.x */
    35, Generating implicit copyin(A[:]) [if not already present]
        Generating implicit copy(error) [if not already present]
        Generating implicit copyout(Anew[:]) [if not already present]
    40, Loop is parallelizable
swap:
    51, Generating Tesla code
    53, #pragma acc loop gang /* blockIdx.x */
    56, #pragma acc loop vector(128) /* threadIdx.x */
    51, Generating implicit copyout(A[:]) [if not already present]
        Generating implicit copyin(Anew[:]) [if not already present]
    56, Loop is parallelizable
```

Parallel

```

#include <math.h>
#include <stdlib.h>

#define OFFSET(x, y, m) (((x)*(m)) + (y))

void initialize(double *restrict A, double *restrict Anew, int m, int n)
{
    memset(A, 0, n * m * sizeof(double));
    memset(Anew, 0, n * m * sizeof(double));

    for(int i = 0; i < m; i++){
        A[i] = 1.0;
        Anew[i] = 1.0;
    }
}

double calcNext(double *restrict A, double *restrict Anew, int m, int n)
{
    double error = 0.0;
    #pragma acc parallel loop reduction(max:error)

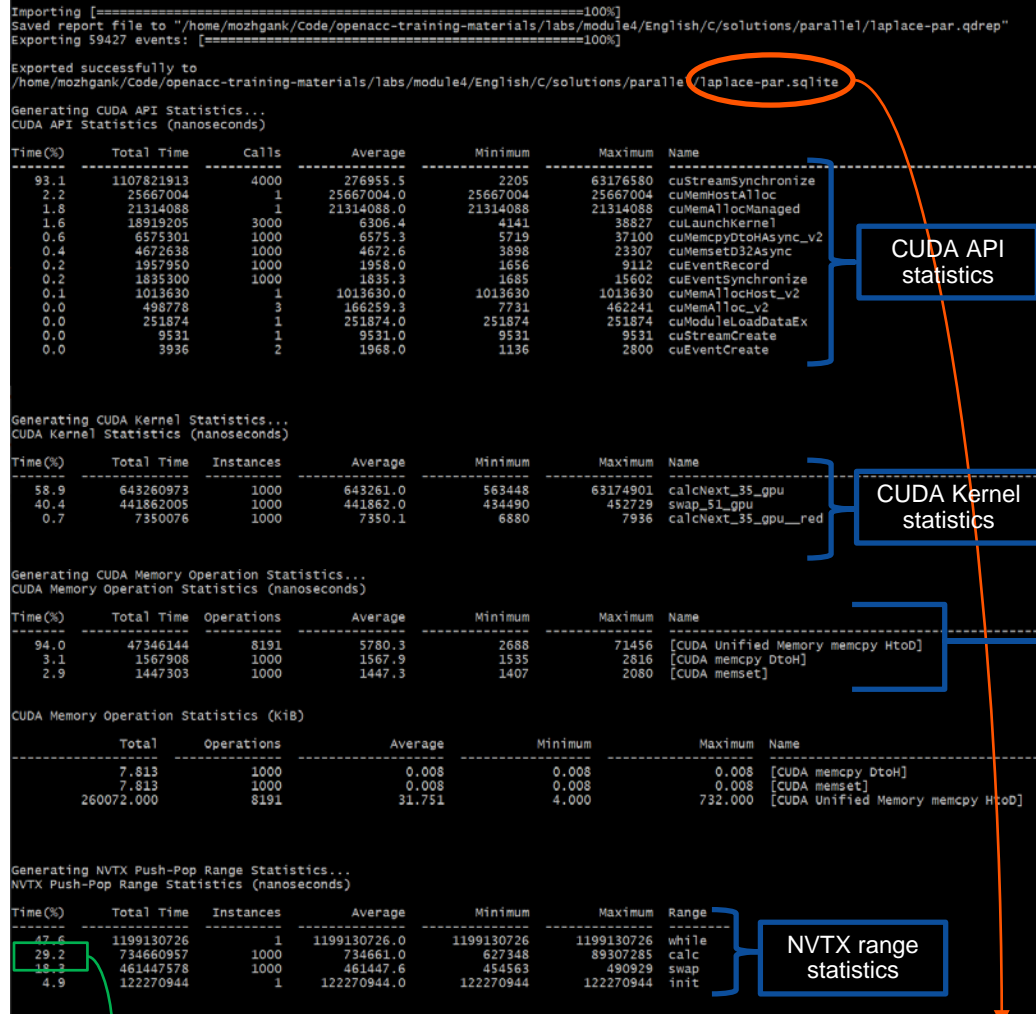
    for( int j = 1; j < n-1; j++)
    {
        #pragma acc loop
        for( int i = 1; i < m-1; i++ )
        {
            Anew[OFFSET(j, i, m)] = 0.25 * ( A[OFFSET(j, i+1, m)] + A[OFFSET(j, i-1, m)]
            + A[OFFSET(j-1, i, m)] + A[OFFSET(j+1, i, m)]);
            error = max( error, fabs(Anew[OFFSET(j, i, m)] - A[OFFSET(j, i, m)]));
        }
    }
    return error;
}

void swap(double *restrict A, double *restrict Anew, int m, int n)
{
    #pragma acc parallel loop
    for( int j = 1; j < n-1; j++)
    {
        #pragma acc loop
        for( int i = 1; i < m-1; i++ )
        {
            A[OFFSET(j, i, m)] = Anew[OFFSET(j, i, m)];
        }
    }
}

void deallocate(double *restrict A, double *restrict Anew)
{
    free(A);
    free(Anew);
}

```

laplace2d.c
 (Parallelised using OpenACC parallel
 directives (pragmas highlighted))



"calc" region (calcNext function) takes 29.2%
 "swap" region (swap function) takes 18.3% of
 total execution time

Open laplace-par.qdrep
 with Nsight System GUI to
 view the timeline

PARALLEL VS SEQUENTIAL SPEEDUP

Viewing captured NVTX events

Have a close look at the captured NVTX events for both serial and parallel implementations.

Time spent in “while” loop has significantly decreased.

Achieved speedup: ~47

```
mozhgank@prn-dgx-32:~/Code/openacc-training-materials/labs/module4/English/C/solutions/parallel$ nsys profile -t nvtx,openacc --stats=true -b dwarf --force-overwrite true -o laplace-par ./laplace-par
WARNING: openACC, cuDNN and cuBLAS rely on CUDA. CUDA tracing has been automatically enabled.
Collecting data...
Jacobi relaxation calculation: 4096 x 4096 mesh
0, 0.250000
100, 0.002397
200, 0.001204
300, 0.000804
400, 0.000603
500, 0.000483
600, 0.000403
700, 0.000345
800, 0.000302
900, 0.000269
total: 1.199133 s
```

Parallel

```
mozhgank@prn-dgx-32:~/Code/openacc-training-materials/labs/module4/English/C/solutions/parallel$ nsys profile -t nvtx --stats=true -b dwarf --force-overwrite true -o laplace-seq ./laplace-seq
Collecting data...
Jacobi relaxation calculation: 4096 x 4096 mesh
0, 0.250000
100, 0.002397
200, 0.001204
300, 0.000804
400, 0.000603
500, 0.000483
600, 0.000403
700, 0.000345
800, 0.000302
900, 0.000269
total: 55.754501 s
```

Sequential

#	Name	Duration	TID	Start
1	init [122.271 ms]	122.271 ms	45547	0.84707s
3	while [1.199 s]	1.199 s	45547	0.969399s

Parallel

#	Name	Duration	TID	Start
1	init [137.490 ms]	137.490 ms	46133	0.00975564s
3	while [55.754 s]	55.754 s	46133	0.147279s

Sequential

PROFILING PARALLEL CODE

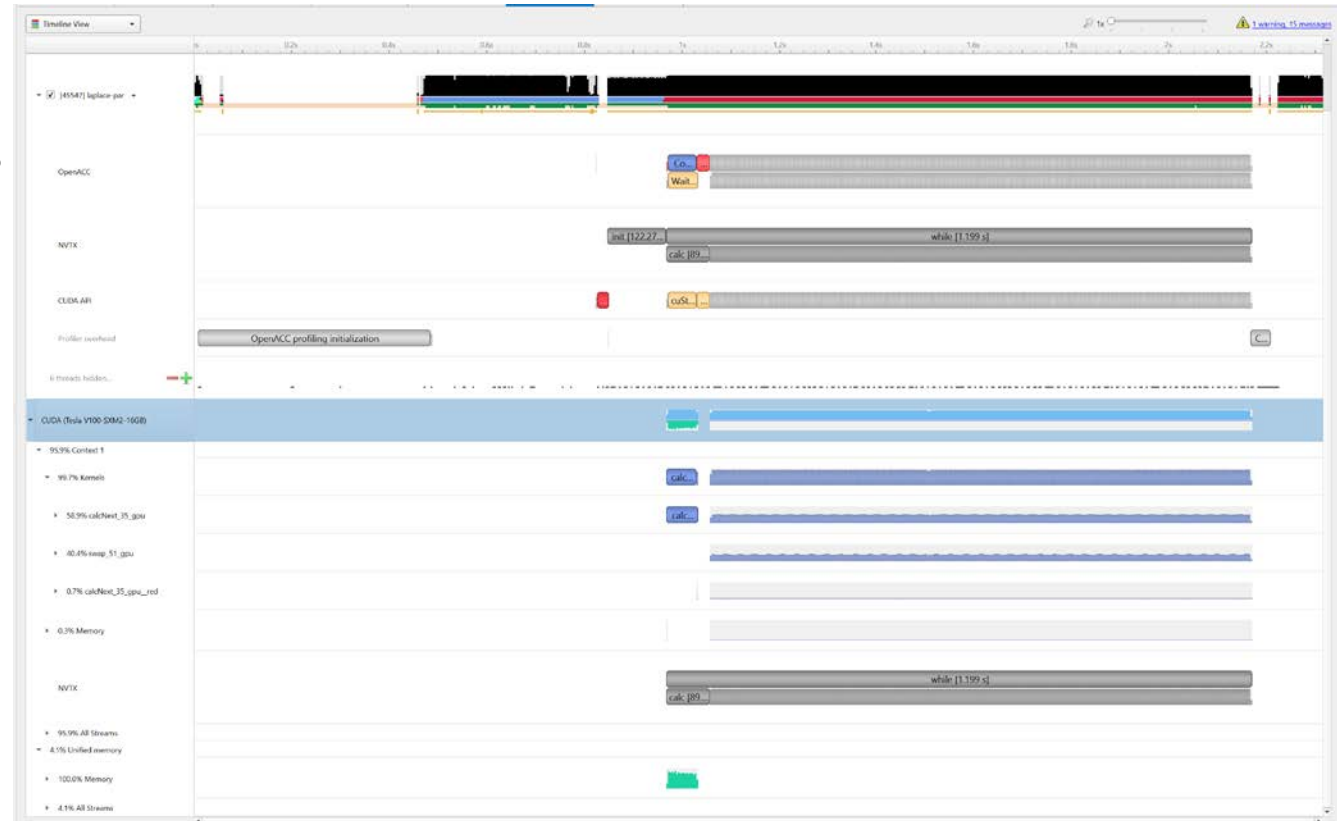
Viewing timeline via Nsight Systems

Contents of the tree-like hierarchy on the left depend on the project settings used to collect this report.

If a certain feature has not been enabled, corresponding rows will not be shown on the timeline.

In this example, we chose to trace NVTX and OpenACC while sampling.

Note: Kernel launches are represented by **blue** and memory transfers are displayed in **green**.



LAB CODE

LAPLACE HEAT TRANSFER

Introduction to lab code - visual


We will observe a simple simulation of heat distributing across a metal plate.

We will apply a consistent heat to the top of the plate.

Then, we will simulate the heat distributing across the plate.

Very Hot

Room Temp



LAPLACE HEAT TRANSFER

Introduction to lab code - technical

The lab simulates a very basic 2-dimensional heat transfer problem. We have two 2-dimensional arrays, **A** and **Anew**.

The arrays represent a 2-dimensional, metal plate. Each element in the array is a **double** value that represents temperature.

We will simulate the distribution of heat until a **minimum change value** is achieved, or until we exceed a **maximum number of iterations**.

A

0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0

Anew

0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0

LAPLACE HEAT TRANSFER

Introduction to lab code - technical

We initialize the top row to be a temperature of 1.0

The **calcNext** function will iterate through all of the inner elements of array A, and update the corresponding elements in Anew

We will take the average of the neighboring cells, and record it in **Anew**.

The **swap** function will copy the contents of Anew to A

A

0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0



Anew

0.0	0.0	0.0	0.0
0.0	0.25	0.25	0.0
0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0

LAPLACE HEAT TRANSFER

Introduction to lab code

A

1.0	1.0	1.0	1.0
0.0	0.25	0.25	0.0
0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0



Anew

1.0	1.0	1.0	1.0
0.0	0.25	0.25	0.0
0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0

The **swap** function will copy the contents of Anew to A

KEY CONCEPTS

In this module we discussed...

- Compiling sequential and parallel code
- CPU profiling for sequential and parallel execution
- Specifics of our Laplace Heat Transfer lab code

LAB GOALS

In this lab you will do the following...

- Build and run the example code using the NVIDIA's HPC compiler
- Use Nsight Systems to understand where the program spends its time

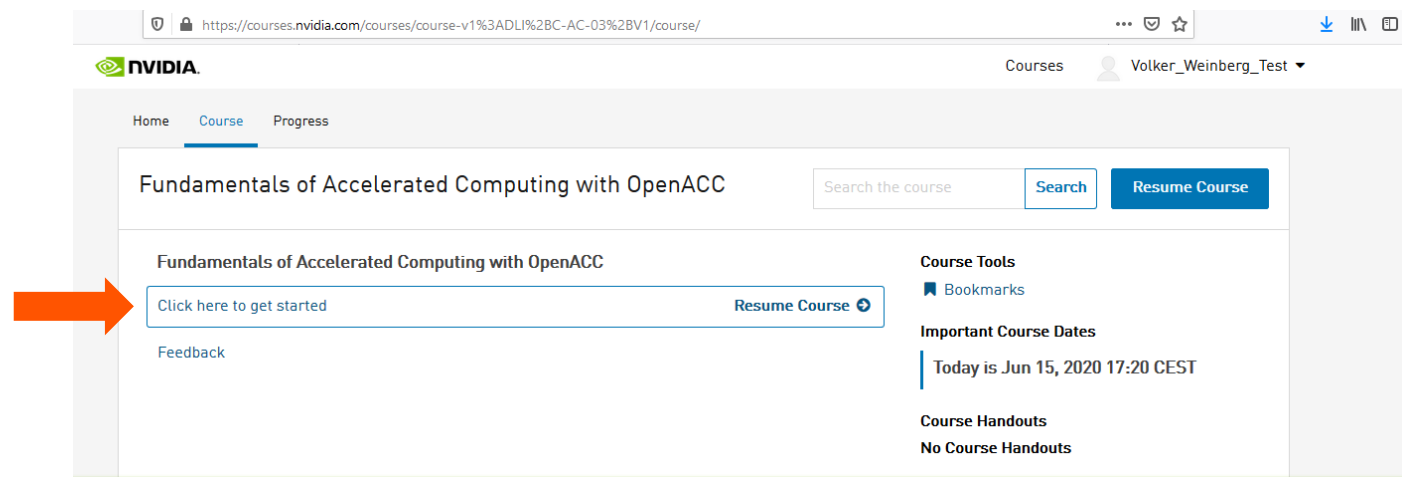
THANK YOU

TRAINING SETUP

- To get started, follow these steps:
- Create an NVIDIA Developer account at <http://courses.nvidia.com/join> Select "Log in with my NVIDIA Account" and then "Create Account"
- Visit <http://courses.nvidia.com/dli-event> and enter the event code

VSC_OACC_AMBASSADOR_AP21

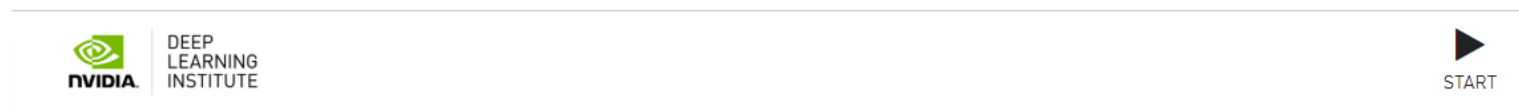
TRAINING SETUP



The screenshot shows the NVIDIA course page for "Fundamentals of Accelerated Computing with OpenACC". The page includes a search bar, a "Resume Course" button, and a "Click here to get started" button. An orange arrow points to the "Click here to get started" button.

Fundamentals of Accelerated Computing with OpenACC

[Bookmark this page](#)



The NVIDIA Deep Learning Institute logo is on the left. On the right, there is a "START" button with a play icon. An orange arrow points to the "START" button.

To get started with this live GPU enabled interactive content please click the "Start" button on the top right of this block.

This will launch a pre-configured GPU workstation, it may take 5-10 minutes.

TRAINING SETUP

Fundamentals of Accelerated Computing with OpenACC

[Bookmark this page](#)



DEEP
LEARNING
INSTITUTE

1 : 56 : 47

REMAINING TIME



LAUNCH TASK



STOP TASK

Please wait 5 - 10 minutes while your interactive GPU enabled environment loads. When the "Launch" button appears, click it to get started.

TRAINING SETUP

Files Running Clusters

Select items to perform actions on them. Upload New ▾ ↻

<input type="checkbox"/> 0 ▾	📁 /	Name ▾	Last Modified	File size
<input type="checkbox"/>	📁 module2		vor 4 Tagen	
<input type="checkbox"/>	📁 module3		vor 4 Tagen	
<input type="checkbox"/>	📁 module4		vor 4 Tagen	
<input type="checkbox"/>	📁 module5		vor 4 Tagen	
<input type="checkbox"/>	📁 module6		vor 4 Tagen	
<input type="checkbox"/>	📄 START HERE.ipynb		vor 4 Tagen	1.36 kB



TRAINING SETUP

Welcome to the OpenACC labs

Please select the appropriate lab below.



- [Module 2](#) - Application Profiling with Nsight Systems Lab - This lab introduces students to application profiling using the Nsight Systems profiler.
- [Module 3](#) - OpenACC Directives Basics - This lab introduces OpenACC directives.
- [Module 4](#) - GPU Programming with OpenACC - This lab introduces GPU programming with OpenACC.
- [Module 5](#) - Data Management with OpenACC - This lab introduces OpenACC data management directives.
- [Module 6](#) - OpenACC Loop Optimizations - This lab introduces students to loop optimizations in OpenACC.

Application Profiling with Nsight Systems Lab

This lab is meant to accompany Module 2 of the OpenACC.org teaching materials. The purpose of this lab is to introduce students to application profiling using the Nsight Systems profiler. Lab instructions and source code is available for C/C++ and Fortran.

Please see the following files to begin the lab:



- [C/C++](#)
- [Fortran](#)

TRAINING SETUP

- NOTE: To be able to see the Nsight System profiler output, please download Nsight System latest version from <https://developer.nvidia.com/nsight-systems>.