



# UNCOALESCED GLOBAL ACCESSES SAMPLE

v2022.3.0 | August 2022



# TABLE OF CONTENTS

- Chapter 1. Introduction..... 1
- Chapter 2. Application..... 2
- Chapter 3. Configuration..... 3
- Chapter 4. Initial version of the kernel..... 4
- Chapter 5. Updated version of the kernel..... 9
- Chapter 6. Resources..... 11

# Chapter 1.

## INTRODUCTION

This sample profiles a memory-bound CUDA kernel which does a simple computation on an array of double3 data type in global memory using the Nsight Compute profiler. The profiler is used to analyze and identify the memory accesses which are uncoalesced and result in inefficient DRAM accesses.

### **Global memory accesses on a GPU**

Global memory resides in device memory and device memory is accessed via 32, 64, or 128-byte memory transactions.

When a warp executes an instruction that accesses global memory, it coalesces the memory accesses of the threads within the warp into one or more of these memory transactions depending on the size of the data accessed by each thread and the distribution of the memory addresses across the threads. If global memory accesses of the threads within a warp cannot be combined into the same memory transaction then we refer to these as uncoalesced global memory accesses. In general, the more transactions are necessary, the more unused bytes are transferred in addition to the bytes accessed by the threads, reducing the instruction throughput accordingly. For example, if a 32-byte memory transaction is generated for each thread's 4-byte access, throughput is divided by 8.

## Chapter 2. APPLICATION

The sample CUDA application adds a floating point constant to an input array of 1,048,576 (1024\*1024) double3 elements in global memory and generates an output array of double3 in global memory of the same size. double3 is a 24-byte built-in vector type which is a structure containing 3 double precision floating point values:

```
struct
{
    double x, y, z;
};
```

The uncoalescedGlobalAccesses sample is available with Nsight Compute under <nsight-compute-install-directory>/extras/samples/uncoalescedGlobalAccesses.

# Chapter 3.

## CONFIGURATION

The profiling results included in this document were collected on the following configuration:

- ▶ Target system: Linux (x86\_64) with a NVIDIA RTX A2000 (Ampere GA106) GPU
- ▶ Nsight Compute version: 2022.1.1

The Nsight Compute UI screen shots in the document are taken by opening the profiling reports on a Windows 10 system.

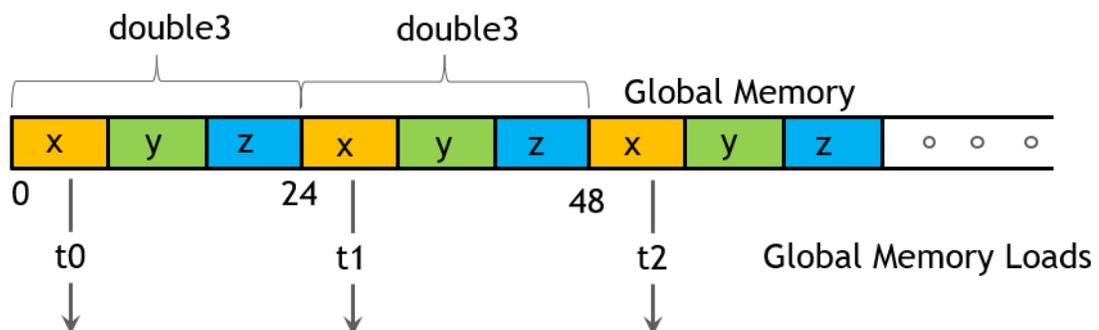
# Chapter 4.

## INITIAL VERSION OF THE KERNEL

The initial version of the sample code provides a naive implementation for the kernel which adds a floating point constant to an input array of double3.

```
__global__ void addConstDouble3(int numElements, double3 *d_in, double k,
double3 *d_out)
{
    int index = blockIdx.x * blockDim.x + threadIdx.x;
    if (index < numElements)
    {
        double3 a = d_in[index];
        a.x += k;
        a.y += k;
        a.z += k;
        d_out[index] = a;
    }
}
```

The instruction `a = d_in[index]` in the kernel code results in each thread in a warp accessing global memory 24-bytes apart. In the first step all threads request a load for `d_in[index].x` as shown in the following diagram. In the second step a load for `d_in[index].y` and in the third step a load for `d_in[index].z` is made by all threads.



### Profile the initial version of the kernel

There are multiple ways to profile kernels with Nsight Compute. For full details see the [Nsight Compute Documentation](#). One example workflow to follow for this sample:

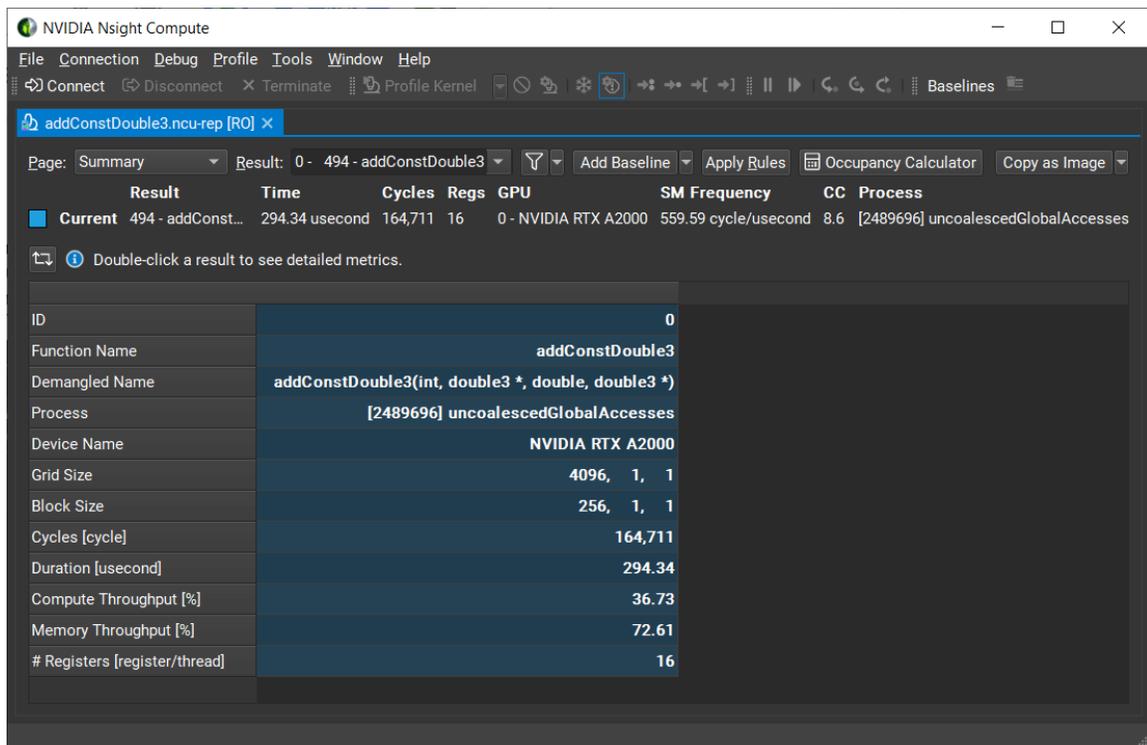
- Refer to the README distributed with the sample on how to build the application

- ▶ Run ncu-ui on the host system
- ▶ Use a local connection if the GPU is on the host system. If the GPU is on a remote system, set up a remote connection to the target system
- ▶ Use the "Profile" activity to profile the sample application
- ▶ Choose the "full" section set
- ▶ Use defaults for all other options

## Summary page

All kernels in the application are profiled and the summary page is displayed. The kernel launch parameters, cycles, duration, compute and memory throughput for each kernel are shown. In this sample we have only one kernel launch.

The duration for this initial version of the kernel is 294 micro seconds and this is used as the baseline for further optimizations.



The screenshot shows the NVIDIA Nsight Compute application window. The title bar reads "NVIDIA Nsight Compute". The menu bar includes "File", "Connection", "Debug", "Profile", "Tools", "Window", and "Help". The toolbar contains buttons for "Connect", "Disconnect", "Terminate", "Profile Kernel", and "Baselines". The main window displays the "Summary" page for a kernel named "addConstDouble3.ncu-rep [R0]".

At the top of the summary page, there is a navigation bar with "Page: Summary", "Result: 0 - 494 - addConstDouble3", and buttons for "Add Baseline", "Apply Rules", "Occupancy Calculator", and "Copy as Image". Below this is a table with the following columns: "Result", "Time", "Cycles", "Regs", "GPU", "SM Frequency", and "CC Process".

Result	Time	Cycles	Regs	GPU	SM Frequency	CC Process
Current 494 - addConst...	294.34 usecond	164,711	16	0 - NVIDIA RTX A2000	559.59 cycle/usecond	8.6 [2489696] uncoalescedGlobalAccesses

Below the table, there is a note: "Double-click a result to see detailed metrics." This leads to a detailed metrics table:

ID	0
Function Name	addConstDouble3
Demangled Name	addConstDouble3(int, double3 *, double, double3 *)
Process	[2489696] uncoalescedGlobalAccesses
Device Name	NVIDIA RTX A2000
Grid Size	4096, 1, 1
Block Size	256, 1, 1
Cycles [cycle]	164,711
Duration [usecond]	294.34
Compute Throughput [%]	36.73
Memory Throughput [%]	72.61
# Registers [register/thread]	16

## Details page - GPU Speed Of Light Throughput

The details page "GPU Speed Of Light Throughput" section provides a high-level overview of the throughput for compute and memory resources of the GPU used by the kernel.

The screenshot shows the NVIDIA Nsight Compute interface. At the top, the window title is 'NVIDIA Nsight Compute'. Below the menu bar, there are navigation buttons like 'Connect', 'Disconnect', 'Terminate', and 'Profile Kernel'. The main area shows the kernel name 'addConstDouble3.ncu-rep [RO]' and a 'Details' page. A table of performance metrics is displayed:

Result	Time	Cycles	Regs	GPU	SM Frequency	CC	Process
Current	494 - addConstDouble...	294.34 usecond	164,711	16	0 - NVIDIA RTX A2000	559.59 cycle/usecond	8.6 [2489696] uncoalescedGlobalAccesses

Below the table, there is a section titled 'GPU Speed Of Light Throughput' with a table of throughput values:

Metric	Value	Unit
Compute (SM) Throughput [%]	36.73	Duration [usecond]
Memory Throughput [%]	72.61	Elapsed Cycles [cycle]
L1/TEX Cache Throughput [%]	95.53	SM Active Cycles [cycle]
L2 Cache Throughput [%]	56.28	SM Frequency [cycle/usecond]
DRAM Throughput [%]	72.61	DRAM Frequency [cycle/nsecond]

There are three warning messages:

- High Memory Throughput**: Memory is more heavily utilized than Compute. Look at the [Memory Workload Analysis](#) section to identify the DRAM bottleneck. Check memory replay (coalescing) metrics to make sure you're efficiently utilizing the bytes transferred. Also consider whether it is possible to do more work per memory access (kernel fusion) or whether there are values you can (re)compute.
- FP64/32 Utilization**: The ratio of peak float (fp32) to double (fp64) performance on this device is 64:1. The kernel achieved 0% of this device's fp32 peak performance and 18% 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 18% lower than the fp64 pipeline utilization. Check the [Instruction Statistics](#) section to see if using fused instructions can benefit this kernel.

For this kernel it shows a hint for High Memory Throughput and suggests looking at the memory workload analysis section. Click on Memory Workload Analysis.

### Details page - Memory Workload Analysis section

The Memory Workload Analysis shows hints for L1TEX Global store and load access patterns. The description and focus metrics of these performance issues describe how more sectors than necessary are being accessed from memory. A sector is an aligned 32-byte chunk of memory in a cache line or device memory. These additional sector accesses are caused by uncoalesced memory accesses and can negatively impact performance. In this case, for the load or store instructions, each thread is accessing a double (8 bytes) and there are 32 threads in a warp. Therefore, each memory request from a warp should ideally access 256 bytes (8 x 32), which is 8 sectors. However, in this unoptimized version, we see 24 sectors per request. It suggests checking the Source Counters section for uncoalesced global stores and loads. Click on the Source Counters link.

**addConstDouble3.ncu-rep [RO]**

Page: Details Result: 0 - 494 - addConstDouble3 Add Baseline Apply Rules Occupancy Calculator Copy as Image

Result	Time	Cycles	Regs	GPU	SM Frequency	CC	Process
Current 494 - addConstDouble3 (4096, 1, 1)x(25...	294.34 usecond	164,711	16	0 - NVIDIA RTX A2000	559.59 cycle/usecond	8.6	[2489696] uncoalescedGlobalAccesses

**Memory Workload Analysis**

Detailed analysis of the memory resources of the GPU. Memory can become a limiting factor for the overall kernel performance when fully utilizing the involved hardware units (Mem Busy), exhausting the available communication bandwidth between those units (Max Bandwidth), or by reaching the maximum throughput of issuing memory instructions (Mem Pipes Busy). Detailed chart of the memory units. Detailed tables with data for each memory unit.

Memory Throughput [Gbyte/second]	166.44	Mem Busy [%]	56.28
L1/TEX Hit Rate [%]	63.10	Max Bandwidth [%]	72.61
L2 Hit Rate [%]	66.75	Mem Pipes Busy [%]	10.71
L2 Compression Success Rate [%]	0	L2 Compression Ratio	0

**L1TEX Global Load Access Pattern**

The memory access pattern for global loads in L1TEX might not be optimal. On average, this kernel accesses 8.0 bytes per thread per memory request but the address pattern, possibly caused by the stride between threads, results in 24.0 sectors per request, or  $24.0 \times 32 = 768.0$  bytes of cache data transfers per request. The optimal thread address pattern for 8.0 byte accesses would result in  $8.0 \times 32 = 256.0$  bytes of cache data transfers per request, to maximize L1TEX cache performance. Check the [Source Counters](#) section for uncoalesced global loads.

Name	Value	Info
Sectors per L1TEX Request	24	2,359,296 / 98,304 × 8.0

**L1TEX Global Store Access Pattern**

The memory access pattern for global stores in L1TEX might not be optimal. On average, this kernel accesses 8.0 bytes per thread per memory request; but the address pattern, possibly caused by the stride between threads, results in 24.0 sectors per request, or  $24.0 \times 32 = 768.0$  bytes of cache data transfers per request. The optimal thread address pattern for 8.0 byte accesses would result in  $8.0 \times 32 = 256.0$  bytes of cache data transfers per request, to maximize L1TEX cache performance. Check the [Source Counters](#) section for uncoalesced global stores.

Name	Value	Info
Sectors per L1TEX Request	24	2,359,296 / 98,304 × 8.0

## Details page - Source Counters section

The Source Counters section shows a hint for "Uncoalesced Global Accesses". It explains that the metric "L2 Theoretical Sectors Global Excessive" is the indicator for uncoalesced accesses. The table for this metric lists the source lines with the highest value. Click on one of the source lines to view the kernel source at which the bottleneck occurs.

**addConstDouble3.ncu-rep [RO]**

Page: Details Result: 0 - 494 - addConstDouble3 Add Baseline Apply Rules Occupancy Calculator Copy as Image

Result	Time	Cycles	Regs	GPU	SM Frequency	CC	Process
Current 494 - addConstDouble...	294.34 usecond	164,711	16	0 - NVIDIA RTX A2000	559.59 cycle/usecond	8.6	[2489696] uncoalescedGlobalAccesses

**Source Counters**

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]	65,536	Branch Efficiency [%]	0
Branch Instructions Ratio [%]	0.10	Avg. Divergent Branches	0

**Uncoalesced Global Accesses**

This kernel has uncoalesced global accesses resulting in a total of 3145728 excessive sectors (67% of the total 4718592 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 (%)
<a href="#">uncoalescedGlobalAccesses.cu:34 (0x7f3d64fb1c20.in...</a>	524,288	17
<a href="#">uncoalescedGlobalAccesses.cu:34 (0x7f3d64fb1c10.in...</a>	524,288	17
<a href="#">uncoalescedGlobalAccesses.cu:34 (0x7f3d64fb1bf0.in...</a>	524,288	17
<a href="#">uncoalescedGlobalAccesses.cu:30 (0x7f3d64fb1bb0.in...</a>	524,288	17
<a href="#">uncoalescedGlobalAccesses.cu:30 (0x7f3d64fb1ba0.in...</a>	524,288	17

## Source page

The CUDA source and SASS(GPU Assembly) for the kernel is shown side by side. When opening the Source page from Source Counters section, the Navigation metric is automatically filled in to match, in this case "L2 Theoretical Sectors Global Excessive". You can see this by the bolding in the column header. The source line at which the bottleneck occurs is highlighted.

It shows uncoalesced global memory load accesses at line #30:

```
double3 a = d_in[index];
```

It shows uncoalesced global memory store accesses at line #34:

```
d_out[index] = a;
```

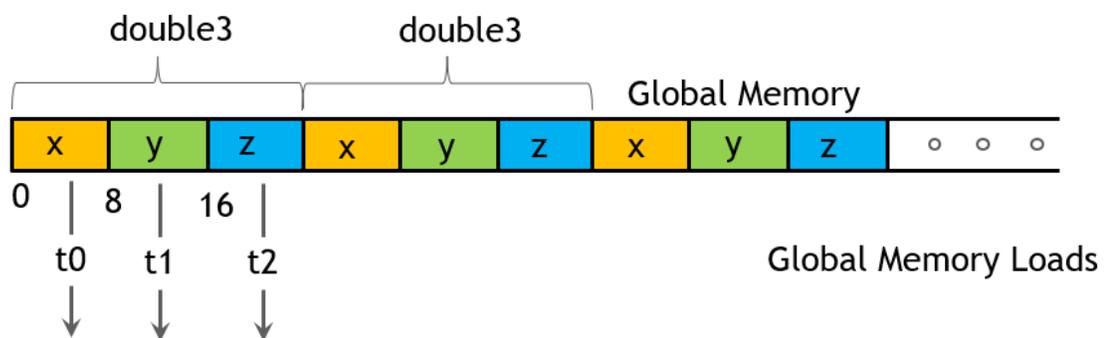
The screenshot displays the NVIDIA Nsight Compute interface. The top panel shows the kernel name 'addConstDouble3.ncu-rep [RO]' and various performance metrics. The 'View' dropdown is set to 'Source and SASS'. The 'Source' view on the left shows the CUDA code with line 30, 'double3 a = d\_in[index];', highlighted. The 'SASS' view on the right shows the assembly code with line 10, 'LD0.E.64 R4, [R2.64];', highlighted. The 'Navigation' dropdown is set to 'L2 Theoretical Sectors Global Excessive'. The 'Theoretical Sectors Global Excessive' column in both views shows a value of 1,572,864.

# Chapter 5.

## UPDATED VERSION OF THE KERNEL

Considering the uncoalesced accesses reported by the profiler we analyze the global load access pattern. Each thread executes 3 reads for the three double values in double3.

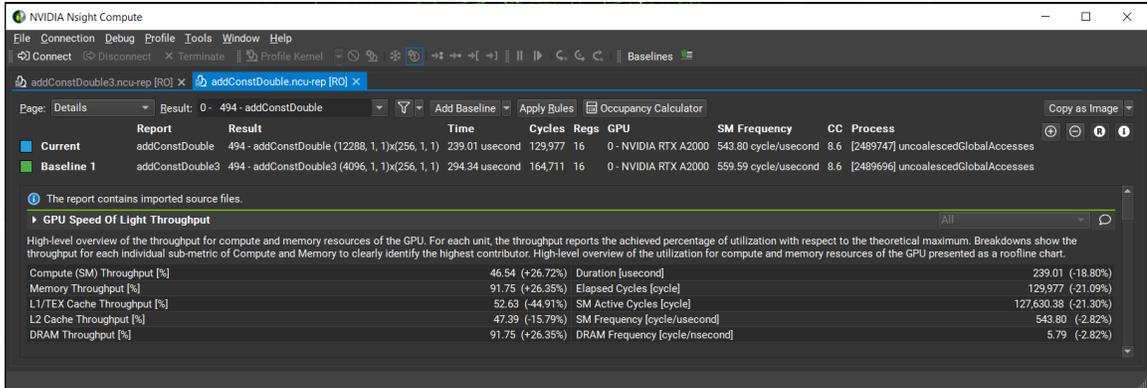
We can treat the double3 array as a double array and each thread can process one double instead of one double3. With this change threads in a warp access consecutive double values and both loads and stores are coalesced.



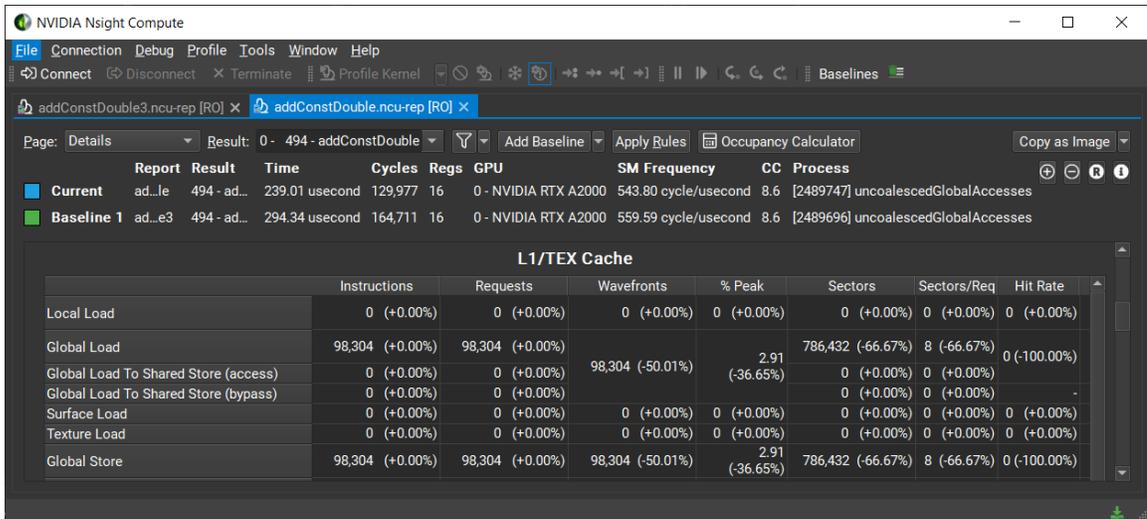
```
__global__ void addConstDouble(int numElements, double *d_in, double k, double *d_out)
{
    int index = blockIdx.x * blockDim.x + threadIdx.x;
    if (index < numElements)
    {
        d_out[index] = d_in[index] + k;
    }
}
```

### Profile the updated kernel

The kernel duration has reduced from 294 microseconds to 239 microseconds. We can set a baseline to the initial version of the kernel and compare the profiling results.



We can confirm that the global memory accesses are coalesced. In the L1/TEX Cache metrics table under the Memory workload analysis section we see that the "Sectors/Req" metric value is 8 for both global loads and global stores.



# Chapter 6.

## RESOURCES

- ▶ GPU Technology Conference 2021 talk S32089: Requests, Wavefronts, Sectors Metrics: Understanding and Optimizing Memory-Bound Kernels with Nsight Compute
- ▶ Nsight Compute Documentation

## **Notice**

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

## **Trademarks**

NVIDIA and the NVIDIA logo are trademarks or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

## **Copyright**

© 2022-2022 NVIDIA Corporation and affiliates. All rights reserved.

This product includes software developed by the Syncro Soft SRL (<http://www.sync.ro/>).