



CUDA Samples

Reference Manual

Table of Contents

Chapter 1. Release Notes.....	1
1.1. CUDA 11.6.....	1
1.2. CUDA 11.5.....	1
1.3. CUDA 11.4 Update 1.....	1
1.4. CUDA 11.4.....	2
1.5. CUDA 11.3.....	2
1.6. CUDA 11.2.....	2
1.7. CUDA 11.1.....	2
1.8. CUDA 11.0.....	3
1.9. CUDA 10.2.....	4
1.10. CUDA 10.1 Update 2.....	4
1.11. CUDA 10.1 Update 1.....	4
1.12. CUDA 10.1.....	4
1.13. CUDA 10.0.....	5
1.14. CUDA 9.2.....	5
1.15. CUDA 9.0.....	5
1.16. CUDA 8.0.....	6
1.17. CUDA 7.5.....	7
1.18. CUDA 7.0.....	7
1.19. CUDA 6.5.....	8
1.20. CUDA 6.0.....	9
1.21. CUDA 5.5.....	9
1.22. CUDA 5.0.....	10
1.23. CUDA 4.2.....	11
1.24. CUDA 4.1.....	11
Chapter 2. Getting Started.....	12
2.1. Getting CUDA Samples.....	12
Windows.....	12
Linux.....	12
2.2. Building Samples.....	12
Windows.....	12
Linux.....	13
2.3. CUDA Cross-Platform Samples.....	13
2.4. Using CUDA Samples to Create Your Own CUDA Projects.....	14
2.4.1. Creating CUDA Projects for Windows.....	14

2.4.2. Creating CUDA Projects for Linux.....	14
Chapter 3. Samples Reference.....	16
3.1. Introduction Reference.....	16
asyncAPI.....	16
c++11_cuda - C++11 CUDA.....	17
clock - Clock.....	17
clock_nvrtc - Clock libNVRTC.....	18
concurrentKernels - Concurrent Kernels.....	18
cplIntegration - C++ Integration.....	18
cppOverload.....	19
cudaOpenMP.....	19
fp16ScalarProduct - FP16 Scalar Product.....	19
matrixMul - Matrix Multiplication (CUDA Runtime API Version).....	20
matrixMul_nvrtc - Matrix Multiplication with libNVRTC.....	20
matrixMulDrv - Matrix Multiplication (CUDA Driver API Version).....	21
matrixMulDynlinkJIT - Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version).....	21
mergeSort - Merge Sort.....	22
simpleAssert.....	22
simpleAssert_nvrtc - simpleAssert with libNVRTC.....	23
simpleAtomicIntrinsics - Simple Atomic Intrinsics.....	23
simpleAtomicIntrinsics_nvrtc - Simple Atomic Intrinsics with libNVRTC.....	23
simpleAttributes.....	24
simpleAWBarrier - Simple Arrive Wait Barrier.....	24
simpleCallback - Simple CUDA Callbacks.....	25
simpleCooperativeGroups - Simple Cooperative Groups.....	25
simpleCubemapTexture - Simple Cubemap Texture.....	25
simpleCUDA2GL - CUDA and OpenGL Interop of Images.....	26
simpleDrvRuntime - Simple Driver-Runtime Interaction.....	26
simpleHyperQ.....	27
simpleIPC.....	27
simpleLayeredTexture - Simple Layered Texture.....	28
simpleMPI.....	28
simpleMultiCopy - Simple Multi Copy and Compute.....	28
simpleMultiGPU - Simple Multi-GPU.....	29
simpleOccupancy.....	29
simpleP2P - Simple Peer-to-Peer Transfers with Multi-GPU.....	30
simplePitchLinearTexture - Pitch Linear Texture.....	30

simplePrintf.....	31
simpleSeparateCompilation - Simple Static GPU Device Library.....	31
simpleStreams.....	31
simpleSurfaceWrite - Simple Surface Write.....	32
simpleTemplates - Simple Templates.....	32
simpleTemplates_nvrtc - Simple Templates with libNVRTC.....	32
simpleTexture - Simple Texture.....	33
simpleTexture3D - Simple Texture 3D.....	33
simpleTextureDrv - Simple Texture (Driver Version).....	34
simpleVoteIntrinsics - Simple Vote Intrinsics.....	34
simpleVoteIntrinsics_nvrtc - Simple Vote Intrinsics with libNVRTC.....	34
simpleZeroCopy.....	35
systemWideAtomics - System wide Atomics.....	35
template - Template.....	36
UnifiedMemoryStreams - Unified Memory Streams.....	36
vectorAdd - Vector Addition.....	36
vectorAdd_nvrtc - Vector Addition with libNVRTC.....	37
vectorAddDrv - Vector Addition Driver API.....	37
vectorAddMMAP - Vector Addition cuMemMap.....	38
3.2. Utilities Reference.....	38
bandwidthTest - Bandwidth Test.....	38
deviceQuery - Device Query.....	39
deviceQueryDrv - Device Query Driver API.....	39
topologyQuery - Topology Query.....	39
3.3. Concepts and Techniques Reference.....	39
boxFilter - Box Filter.....	39
convolutionSeparable - CUDA Separable Convolution.....	40
convolutionTexture - Texture-based Separable Convolution.....	40
cuHook - CUDA Interception Library.....	41
dct8x8 - DCT8x8.....	41
EGLStream_CUDA_CrossGPU.....	42
EGLStream_CUDA_Interop - EGLStream CUDA Interop.....	42
EGLSync_CUDA_Interop - EGLSync CUDA Event Interop.....	43
eigenvalues - Eigenvalues.....	44
FunctionPointers - Function Pointers.....	44
histogram - CUDA Histogram.....	44
imageDenoising - Image denoising.....	45
inlinePTX - Using Inline PTX.....	45

inlinePTX_nvrtc - Using Inline PTX with libNVRTC.....	46
interval - Interval Computing.....	46
MC_EstimatePiInlineP - Monte Carlo Estimation of Pi (inline PRNG).....	47
MC_EstimatePiInlineQ - Monte Carlo Estimation of Pi (inline QRNG).....	47
MC_EstimatePiP - Monte Carlo Estimation of Pi (batch PRNG).....	48
MC_EstimatePiQ - Monte Carlo Estimation of Pi (batch QRNG).....	48
MC_SingleAsianOptionP - Monte Carlo Single Asian Option.....	49
particles - Particles.....	49
radixSortThrust - CUDA Radix Sort (Thrust Library).....	50
reduction - CUDA Parallel Reduction.....	50
reductionMultiBlockCG - Reduction using MultiBlock Cooperative Groups.....	51
scalarProd - Scalar Product.....	51
scan - CUDA Parallel Prefix Sum (Scan).....	51
segmentationTreeThrust - CUDA Segmentation Tree Thrust Library.....	52
shfl_scan - CUDA Parallel Prefix Sum with Shuffle Intrinsics (SHFL_Scan).....	52
sortingNetworks - CUDA Sorting Networks.....	52
streamOrderedAllocation - stream Ordered Allocation.....	53
streamOrderedAllocationIPC - stream Ordered Allocation IPC Pools.....	53
streamOrderedAllocationP2P - stream Ordered Allocation Peer-to-Peer access.....	54
threadFenceReduction.....	54
threadMigration - CUDA Context Thread Management.....	55
3.4. CUDA Features Reference.....	55
bf16TensorCoreGemm - bfloat16 Tensor Core GEMM.....	55
binaryPartitionCG - Binary Partition Cooperative Groups.....	56
bindlessTexture - Bindless Texture.....	56
cdpAdvancedQuicksort - Advanced Quicksort (CUDA Dynamic Parallelism).....	57
cdpBezierTessellation - Bezier Line Tessellation (CUDA Dynamic Parallelism).....	57
cdpQuadtree - Quad Tree (CUDA Dynamic Parallelism).....	58
cdpSimplePrint - Simple Print (CUDA Dynamic Parallelism).....	58
cdpSimpleQuicksort - Simple Quicksort (CUDA Dynamic Parallelism).....	59
cudaCompressibleMemory - CUDA Compressible Memory.....	59
cudaTensorCoreGemm - CUDA Tensor Core GEMM.....	60
dmmaTensorCoreGemm - Double Precision Tensor Core GEMM.....	60
globalToShmemAsyncCopy - Global Memory to Shared Memory Async Copy.....	61
graphMemoryFootprint - Graph Memory Footprint.....	61
graphMemoryNodes - Graph Memory Nodes.....	62
immaTensorCoreGemm - Tensor Core GEMM Integer MMA.....	62
jacobiCudaGraphs - Jacobi CUDA Graphs.....	63

memMapIPCDrv - Memmap IPC Driver API.....	63
newdelete - NewDelete.....	64
ptxjit - PTX Just-in-Time compilation.....	64
simpleCudaGraphs - Simple CUDA Graphs.....	64
StreamPriorities - Stream Priorities.....	65
tf32TensorCoreGemm - tf32 Tensor Core GEMM.....	65
warpAggregatedAtomicsCG - Warp Aggregated Atomics using Cooperative Groups.....	66
3.5. CUDA Libraries Reference.....	66
batchCUBLAS.....	66
batchedLabelMarkersAndLabelCompressionNPP - Batched Label Markers And Label Compression NPP.....	67
- Box Filter with NPP.....	67
cannyEdgeDetectorNPP - Canny Edge Detector NPP.....	68
conjugateGradient - ConjugateGradient.....	68
conjugateGradientCudaGraphs - Conjugate Gradient using Cuda Graphs.....	69
conjugateGradientMultiBlockCG - conjugateGradient using MultiBlock Cooperative Groups.....	69
conjugateGradientMultiDeviceCG - conjugateGradient using MultiDevice Cooperative Groups.....	70
conjugateGradientPrecond - Preconditioned Conjugate Gradient.....	71
conjugateGradientUM - ConjugateGradientUM.....	71
cudaNvSci - CUDA NvSciBuf/NvSciSync Interop.....	72
- NvMedia CUDA Interop.....	72
- cuDLA Error Reporting.....	73
- cuDLA Hybrid Mode.....	73
- cuDLA Standalone Mode.....	74
- cuSolverDn Linear Solver.....	74
cuSolverRf - cuSolverRf Refactorization.....	75
- cuSolverSp Linear Solver.....	75
cuSolverSp_LowlevelCholesky - cuSolverSp LowlevelCholesky Solver.....	75
cuSolverSp_LowlevelQR - cuSolverSp Lowlevel QR Solver.....	76
FilterBorderControlNPP - Filter Border Control NPP.....	76
freedImageInteropNPP - FreeImage and NPP Interopability.....	77
histEqualizationNPP - Histogram Equalization with NPP.....	77
lineOfSight - Line of Sight.....	78
matrixMulCUBLAS - Matrix Multiplication (CUBLAS).....	78
- MersenneTwisterGP11213.....	79
- NVJPEG simple.....	79
- NVJPEG Encoder.....	80

oceanFFT - CUDA FFT Ocean Simulation.....	80
randomFog - Random Fog.....	80
- Simple CUBLAS.....	81
simpleCUBLAS_LU - Simple CUBLAS LU.....	81
simpleCUBLASXT - Simple CUBLAS XT.....	82
simpleCUFFT - Simple CUFFT.....	82
simpleCUFFT_2d_MGPU - SimpleCUFFT_2d_MGPU.....	83
simpleCUFFT_callback - Simple CUFFT Callbacks.....	83
simpleCUFFT_MGPU - Simple CUFFT_MGPU.....	84
watershedSegmentationNPP - Watershed Segmentation NPP.....	84
3.6. Domain Specific Reference.....	85
bicubicTexture - Bicubic B-spline Interpolation.....	85
bilateralFilter - Bilateral Filter.....	85
binomialOptions - Binomial Option Pricing.....	86
binomialOptions_nvrtc - Binomial Option Pricing with libNVRTC.....	86
BlackScholes - Black-Scholes Option Pricing.....	87
BlackScholes_nvrtc - Black-Scholes Option Pricing with libNVRTC.....	87
convolutionFFT2D - FFT-Based 2D Convolution.....	87
dwtHaar1D - 1D Discrete Haar Wavelet Decomposition.....	88
dxtc - DirectX Texture Compressor (DXTC).....	88
fastWalshTransform - Fast Walsh Transform.....	88
FDTD3d - CUDA C 3D FDTD.....	89
fluidsD3D9 - Fluids (Direct3D Version).....	89
fluidsGL - Fluids (OpenGL Version).....	90
fluidsGLES - Fluids (OpenGLES Version).....	90
HSOpticalFlow - Optical Flow.....	91
Mandelbrot.....	91
marchingCubes - Marching Cubes Isosurfaces.....	92
MonteCarloMultiGPU - Monte Carlo Option Pricing with Multi-GPU support.....	92
nbody - CUDA N-Body Simulation.....	93
nbody_opengles - CUDA N-Body Simulation with GLES.....	94
nbody_screen - CUDA N-Body Simulation on Screen.....	94
- NV12toBGRandResize.....	95
p2pBandwidthLatencyTest - Peer-to-Peer Bandwidth Latency Test with Multi-GPUs.....	95
postProcessGL - Post-Process in OpenGL.....	96
quasirandomGenerator - Niederreiter Quasirandom Sequence Generator.....	96
quasirandomGenerator_nvrtc - Niederreiter Quasirandom Sequence Generator with libNVRTC.....	97

recursiveGaussian - Recursive Gaussian Filter.....	97
simpleD3D10 - Simple Direct3D10 (Vertex Array).....	98
simpleD3D10RenderTarget - Simple Direct3D10 Render Target.....	98
simpleD3D10Texture - Simple D3D10 Texture.....	99
- Simple D3D11.....	99
simpleD3D11Texture - Simple D3D11 Texture.....	100
simpleD3D12 - Simple D3D12 CUDA Interop.....	101
simpleD3D9 - Simple Direct3D9 (Vertex Arrays).....	101
simpleD3D9Texture - Simple D3D9 Texture.....	102
- Simple OpenGL.....	102
simpleGLES - Simple OpenGL.....	103
simpleGLES_EGLOutput - Simple OpenGL EGLOutput.....	103
simpleGLES_screen - Simple OpenGL on Screen.....	104
- Vulkan CUDA Interop Sinewave.....	104
- Vulkan CUDA Interop PI Approximation.....	105
SLID3D10Texture - SLI D3D10 Texture.....	106
smokeParticles - Smoke Particles.....	106
SobelFilter - Sobel Filter.....	107
SobolQRNG - Sobol Quasirandom Number Generator.....	107
stereoDisparity - Stereo Disparity Computation (SAD SIMD Intrinsics).....	107
VFlockingD3D10.....	108
volumeFiltering - Volumetric Filtering with 3D Textures and Surface Writes.....	108
volumeRender - Volume Rendering with 3D Textures.....	109
- Vulkan Image - CUDA Interop.....	110
3.7. Performance Reference.....	110
alignedTypes - Aligned Types.....	110
transpose - Matrix Transpose.....	111
UnifiedMemoryPerf - Unified and other CUDA Memories Performance.....	111
Chapter 4. Dependencies.....	113
Third-Party Dependencies.....	113
FreeImage.....	113
Message Passing Interface.....	113
Only 64-Bit.....	113
DirectX.....	114
DirectX 12.....	114
OpenGL.....	114
OpenGL ES.....	114
Vulkan.....	114

OpenMP.....	114
Screen.....	114
X11.....	115
EGL.....	115
EGLOutput.....	115
EGLSync.....	115
NVSCI.....	115
NvMedia.....	115
CUDA Features.....	115
CUFFT Callback Routines.....	115
CUDA Dynamic Paralellism.....	116
Multi-block Cooperative Groups.....	116
Multi-Device Cooperative Groups.....	116
CUBLAS.....	116
CUDA Interprocess Communication.....	116
CUFFT.....	116
CURAND.....	116
CUSPARSE.....	116
CUSOLVER.....	116
NPP.....	117
NVJPEG.....	117
NVRTC.....	117
Stream Priorities.....	117
Unified Virtual Memory.....	117
16-bit Floating Point.....	117
C++11 CUDA.....	117
Chapter 5. Key Concepts and Associated Samples.....	118
Basic Key Concepts.....	118
Advanced Key Concepts.....	126
Chapter 6. CUDA API and Associated Samples.....	133
CUDA Driver API Samples.....	133
CUDA Runtime API Samples.....	140
Chapter 7. Frequently Asked Questions.....	168

List of Tables

Table 1. Basic Key Concepts and Associated Samples	118
Table 2. Advanced Key Concepts and Associated Samples	126
Table 3. CUDA Driver API and Associated Samples	133
Table 4. CUDA Runtime API and Associated Samples	140

Chapter 1. Release Notes

This section describes the release notes for the CUDA Samples only. For the release notes for the whole CUDA Toolkit, please see [CUDA Toolkit Release Notes](#).

1.1. CUDA 11.6

- ▶ All CUDA samples are now only available on [GitHub repository](#). They are no longer available via CUDA toolkit.
- ▶ Added new folder structure for samples.
- ▶ Added Visual Studio 2022 support to all the samples.

1.2. CUDA 11.5

- ▶ All CUDA samples are now available on [GitHub repository](#).
- ▶ Added `4_CUDA_Libraries/cuDLAHybridMode`. Demonstrate usage of cuDLA in hybrid mode. (available only on [GitHub repository](#))
- ▶ Added `4_CUDA_Libraries/cuDLAStandaloneMode`. Demonstrate usage of cuDLA in standalone mode. (available only on [GitHub repository](#))
- ▶ Added `4_CUDA_Libraries/cuDLAErrorReporting`. Demonstrate DLA error detection via CUDA. (available only on [GitHub repository](#))
- ▶ Added `3_CUDA_Features/graphMemoryNodes`. Demonstrates memory allocations and frees within CUDA graphs using Graph APIs and Stream Capture APIs. (available only on [GitHub repository](#))
- ▶ Added `3_CUDA_Features/graphMemoryFootprint`. Demonstrates how graph memory nodes re-use virtual addresses and physical memory. (available only on [GitHub repository](#))

1.3. CUDA 11.4 Update 1

- ▶ Added support for VS Code on linux platform.

1.4. CUDA 11.4

- ▶ Added `7_CUDA Libraries/simpleCUBLAS_LU`. Demonstrates batched matrix LU decomposition using cuBLAS API `cusblas<t>getrfBatched()`.
- ▶ Updated `2_Graphics/simpleVulkan`, `2_Graphics/simpleVulkanMMAP` and `3_Imaging/vulkanImageCUDA`. Demonstrates use of SPIR-V shaders.
- ▶ Removed `7_CUDA Libraries/boundSegmentsNPP`.

1.5. CUDA 11.3

- ▶ Added `0_Simple/streamOrderedAllocationIPC`. Demonstrates IPC pools of stream ordered memory allocated using `cudaMallocAsync` and `cudaMemPool` family of APIs.
- ▶ Updated `2_Graphics/simpleVulkan`. Demonstrates use of timeline semaphore.
- ▶ Updated `0_Simple/globalToShmemAsyncCopy` with a partitioned cuda pipeline producer-consumer GEMM kernel.
- ▶ Updated multiple samples to use pinned memory using `cudaMallocHost()`.

1.6. CUDA 11.2

- ▶ FreedImage is no longer distributed with the CUDA Samples. On Windows, see the [Dependencies](#) section for more details on how to set up FreedImage. On Linux, it is recommended to install FreedImage with your distribution's package manager.

1.7. CUDA 11.1

- ▶ Added `2_Graphics/simpleVulkanMMAP`. Demonstrates Vulkan CUDA Interop via `cuMemMap` APIs where CUDA buffer is imported in vulkan.
- ▶ Added `7_CUDA Libraries/watershedSegmentationNPP`. Demonstrates how to use the NPP watershed segmentation function.
- ▶ Added `7_CUDA Libraries/batchedLabelMarkersAndLabelCompressionNPP`. Demonstrates how to use the NPP label markers generation and label compression functions based on a Union Find (UF) algorithm including both single image and batched image versions.
- ▶ Deprecated Visual Studio 2015 support for all Windows supported samples.
- ▶ Dropped Visual Studio 2012, 2013 support from all the Windows supported samples.

1.8. CUDA 11.0

- ▶ Added `0_Simple/globalToShmemAsyncCopy`. Demonstrates asynchronous copy of data from global to shared memory using cuda pipeline. Also demonstrates arrive-wait barrier for synchronization.
- ▶ Added `0_Simple/simpleAttributes`. Demonstrates the stream attributes that affect L2 locality.
- ▶ Added `0_Simple/dmmaTensorCoreGemm`. Demonstrates double precision GEMM computation using the WMMA API for double precision employing the Tensor Cores. Also makes use of asynchronous copy from global to shared memory using cuda pipeline which leads to further performance gain.
- ▶ Added `0_Simple/bf16TensorCoreGemm`. Demonstrates `__nv_bfloat16` (e8m7) GEMM computation using the WMMA API for `__nv_bfloat16` employing the Tensor Cores. Also makes use of asynchronous copy from global to shared memory using cuda pipeline which leads to further performance gain.
- ▶ Added `0_Simple/tf32TensorCoreGemm`. Demonstrates `tf32` (e8m10) GEMM computation using the WMMA API for `tf32` employing the Tensor Cores. Also makes use of asynchronous copy from global to shared memory using cuda pipeline which leads to further performance gain.
- ▶ Added `0_Simple/simpleAWBarrier`. Demonstrates the arrive wait barriers.
- ▶ Added warp aggregated atomic multi bucket increments kernel using `labeled_partition` cooperative groups in `6_Advanced/warpAggregatedAtomicsCG` which can be used on compute capability 7.0 and above GPU architectures.
- ▶ Added `0_Simple/binaryPartitionCG`. Demonstrates `binary_partition` cooperative groups creation and usage in divergent path.
- ▶ Added `6_Advanced/cudaCompressibleMemory`. Demonstrates compressible memory allocation using `cuMemMap` API.
- ▶ Removed `7_CUDA Libraries/nvgraph_Pagerank`, `7_CUDA Libraries/nvgraph_SemiRingSpMV`, `7_CUDA Libraries/nvgraph_SpectralClustering`, `7_CUDA Libraries/nvgraph_SSSP` as the NVGRAPH library is dropped from CUDA Toolkit 11.0.
- ▶ Added two new reduction kernels in `6_Advanced/reduction` one which demonstrates `reduce_add_sync` intrinsic supported on compute capability 8.0 and another which uses `cooperative_groups::reduce` function which does `thread_block_tile` level reduction introduced from CUDA 11.0
- ▶ Added windows support to `6_Advanced/c++11_cuda`.

1.9. CUDA 10.2

- ▶ Added `6_Advanced/jacobiCudaGraphs`. Demonstrates Instantiated CUDA Graph Update usage.
- ▶ Added `0_Simple/memMapIPCDrv`. Demonstrates Inter Process Communication using `cuMemMap` APIs with one process per GPU for computation.
- ▶ Added `0_Simple/vectorAddMMAP`. Demonstrates how `cuMemMap` API allows the user to specify the physical properties of their memory while retaining the contiguous nature of their access, thus not requiring a change in their program structure.
- ▶ Added `0_Simple/simpleDrvRuntime`. Demonstrates how CUDA Driver and Runtime APIs can work together to load cuda fatbinary of vector add kernel.
- ▶ Added `0_Simple/cudaNvSci`. Demonstrates CUDA-NvSciBuf/NvSciSync Interop.

1.10. CUDA 10.1 Update 2

- ▶ Added `3_Imaging/vulkanImageCUDA`. Demonstrates how to perform Vulkan Image-CUDA Interop.
- ▶ Added `7_CUDALibraries/nvJPEG_encoder`. Demonstrates encoding of jpeg images using NVJPEG Library.
- ▶ Added Windows support to `7_CUDALibraries/nvJPEG`.
- ▶ Removed DirectX SDK (June 2010 or newer) installation requirement, all the DirectX-CUDA samples now use DirectX from Windows SDK shipped with Microsoft Visual Studio 2012 or higher

1.11. CUDA 10.1 Update 1

- ▶ Added `3_Imaging/NV12toBGRandResize`. Demonstrates how to convert and resize NV12 frames to BGR planars frames using CUDA in batch.
- ▶ Added Visual Studio 2019 support to all the samples.

1.12. CUDA 10.1

- ▶ Added `0_Simple/immaTensorCoreGemm`. Demonstrates integer GEMM computation using the Warp Matrix Multiply and Accumulate (WMMA) API for integers employing the Tensor Cores.
- ▶ Added `2_Graphics/simpleD3D12`. Demonstrates Direct3D12 interoperability with CUDA.

- ▶ Added `7_CUDA Libraries/nvJPEG`. Demonstrates single and batched decoding of jpeg images using NVJPEG Library.
- ▶ Added `7_CUDA Libraries/conjugateGradientCudaGraphs`. Demonstrates conjugate gradient solver on GPU using CUBLAS/CUSPARSE library calls captured and called using CUDA Graph APIs.
- ▶ Updated `0_Simple/simpleIPC` to work on Windows OS as well with TCC enabled GPUs.

1.13. CUDA 10.0

- ▶ Added `1_Uilities/UnifiedMemoryPerf`. Demonstrates the performance comparison of Unified Memory and other types of memory like zero copy buffers, pageable, pagelocked memory on a single GPU.
- ▶ Added `2_Graphics/simpleVulkan`. Demonstrates the Vulkan-CUDA Interop. CUDA imports the Vulkan vertex buffer and operates on it to create sinewave, and synchronizes with Vulkan through vulkan semaphores imported by CUDA.
- ▶ Added `0_Simple/simpleCudaGraphs`. Demonstrates how to use CUDA Graphs through Graphs APIs and Stream Capture APIs.
- ▶ Removed `3_Imaging/cudaDecodeGL`, `3_Imaging/cudaDecodeD3D9` as the cuvid library is dropped from CUDA Toolkit 10.0.
- ▶ Removed `6_Advanced/cdpLUdecomposition`, `7_CUDA Libraries/simpleDevLibCUBLAS` as the CUBLAS Device library is dropped from CUDA Toolkit 10.0.

1.14. CUDA 9.2

- ▶ Added `7_CUDA Libraries/boundSegmentsNPP`. Demonstrates `nppiLabelMarkers` to generate connected region segment labels.
- ▶ Added `6_Advanced/conjugateGradientMultiDeviceCG`. Demonstrates a conjugate gradient solver on multiple GPUs using Multi Device Cooperative Groups, also uses Unified Memory optimized using prefetching and usage hints.
- ▶ Updated `0_Simple/fp16ScalarProduct` to use fp16 native operators for half2 and other fp16 features, it also compare results of using native vs intrinsics fp16 operations.

1.15. CUDA 9.0

- ▶ Added `7_CUDA Libraries/nvgraph_SpectralClustering`. Demonstrates Spectral Clustering using NVGRAPH Library.
- ▶ Added `6_Advanced/warpAggregatedAtomicsCG`. Demonstrates warp aggregated atomics using Cooperative Groups.

- ▶ Added `6_Advanced/reductionMultiBlockCG`. Demonstrates single pass reduction using Multi Block Cooperative Groups.
- ▶ Added `6_Advanced/conjugateGradientMultiBlockCG`. Demonstrates a conjugate gradient solver on GPU using Multi Block Cooperative Groups.
- ▶ Added Cooperative Groups(CG) support to several samples notable ones to name are `6_Advanced/cdpQuadtree`, `6_Advanced/cdpAdvancedQuicksort`, `6_Advanced/threadFenceReduction`, `3_Imaging/dxtc`, `4_Finance/MonteCarloMultiGPU`, `0_Simple/matrixMul_nvrtc`.
- ▶ Added `0_Simple/simpleCooperativeGroups`. Illustrates basic usage of Cooperative Groups within the thread block.
- ▶ Added `0_Simple/cudaTensorCoreGemm`. Demonstrates a GEMM computation using the Warp Matrix Multiply and Accumulate (WMMA) API introduced in CUDA 9, as well as the new Tensor Cores introduced in the Volta chip family.
- ▶ Updated `0_Simple/simpleVoteIntrinsics` to use newly added `*_sync` equivalent of the vote intrinsics `_any`, `_all`.
- ▶ Updated `6_Advanced/shfl_scan` to use newly added `*_sync` equivalent of the shfl intrinsics.

1.16. CUDA 8.0

- ▶ Added `7_CUDA Libraries/FilterBorderControlNPP`. Demonstrates how any border version of an NPP filtering function can be used in the most common mode (with border control enabled), can be used to duplicate the results of the equivalent non-border version of the NPP function, and can be used to enable and disable border control on various source image edges depending on what portion of the source image is being used as input.
- ▶ Added `7_CUDA Libraries/cannyEdgeDetectorNPP`. Demonstrates the recommended parameters to use with the `nppiFilterCannyBorder_8u_C1R` Canny Edge Detection image filter function. This function expects a single channel 8-bit grayscale input image. You can generate a grayscale image from a color image by first calling `nppiColorToGray()` or `nppiRGBToGray()`. The Canny Edge Detection function combines and improves on the techniques required to produce an edge detection image using multiple steps.
- ▶ Added `7_CUDA Libraries/cuSolverSp_LowlevelCholesky`. Demonstrates Cholesky factorization using `cuSolverSP`'s low level APIs.
- ▶ Added `7_CUDA Libraries/cuSolverSp_LowlevelQR`. Demonstrates QR factorization using `cuSolverSP`'s low level APIs.
- ▶ Added `7_CUDA Libraries/BiCGStab`. Demonstrates Bi-Conjugate Gradient Stabilized (BiCGStab) iterative method for nonsymmetric and symmetric positive definite linear systems using CUSPARSE and CUBLAS
- ▶ Added `7_CUDA Libraries/nvgraph_Pagerank`. Demonstrates Page Rank computation using `nvGRAPH` Library.

- ▶ Added `7_CUDA Libraries/nvgraph_SemiRingSpMV`. Demonstrates Semi-Ring SpMV using nvGRAPH Library.
- ▶ Added `7_CUDA Libraries/nvgraph_SSSP`. Demonstrates Single Source Shortest Path(SSSP) computation using nvGRAPH Library.
- ▶ Added `7_CUDA Libraries/simpleCUBLASXT`. Demonstrates simple example to use CUBLAS-XT library.
- ▶ Added `6_Advanced/c++11_cuda`. Demonstrates C++11 feature support in CUDA.
- ▶ Added `1_Uutilities/topologyQuery`. Demonstrates how to query the topology of a system with multiple GPU.
- ▶ Added `0_Simple/fp16ScalarProduct`. Demonstrates scalar product calculation of two vectors of FP16 numbers.
- ▶ Added `0_Simple/systemWideAtomics`. Demonstrates system wide atomic instructions on migratable memory.
- ▶ Removed `0_Simple/template_runtime`. Its purpose is served by `0_Simple/template`.

1.17. CUDA 7.5

- ▶ Added `7_CUDA Libraries/cuSolverDn_LinearSolver`. Demonstrates how to use the CUSOLVER library for performing dense matrix factorization using cuSolverDN's LU, QR and Cholesky factorization functions.
- ▶ Added `7_CUDA Libraries/cuSolverRf`. Demonstrates how to use cuSolverRF, a sparse re-factorization package of the CUSOLVER library.
- ▶ Added `7_CUDA Libraries/cuSolverSp_LinearSolver`. Demonstrates how to use cuSolverSP which provides sparse set of routines for sparse matrix factorization.
- ▶ The `2_Graphics/simpleD3D9`, `2_Graphics/simpleD3D9Texture`, `3_Imaging/cudaDecodeD3D9`, and `5_Simulations/fluidsD3D9` samples have been modified to use the Direct3D 9Ex API instead of the Direct3D 9 API.
- ▶ The `7_CUDA Libraries/graphcutNPP` and `7_CUDA Libraries/imageSegmentationNPP` samples have been removed. These samples used the NPP graphcut APIs, which have been deprecated in CUDA 7.5.

1.18. CUDA 7.0

- ▶ Removed support for Windows 32-bit builds.
- ▶ The Makefile `x86_64=1` and `ARMv7=1` options have been deprecated. Please use `TARGET_ARCH` to set the targeted build architecture instead.
- ▶ The Makefile `GCC` option has been deprecated. Please use `HOST_COMPILER` to set the host compiler instead.

- ▶ The CUDA Samples are no longer shipped as prebuilt binaries on Windows. Please use VS Solution files provided to build respective executable.
- ▶ Added `0_Simple/clock_nvrtc`. Demonstrates how to compile clock function kernel at runtime using libNVRTC to measure the performance of kernel accurately.
- ▶ Added `0_Simple/inlinePTX_nvrtc`. Demonstrates compilation of CUDA kernel having PTX embedded at runtime using libNVRTC.
- ▶ Added `0_Simple/matrixMul_nvrtc`. Demonstrates compilation of matrix multiplication CUDA kernel at runtime using libNVRTC.
- ▶ Added `0_Simple/simpleAssert_nvrtc`. Demonstrates compilation of CUDA kernel having `assert()` at runtime using libNVRTC.
- ▶ Added `0_Simple/simpleAtomicIntrinsics_nvrtc`. Demonstrates compilation of CUDA kernel performing atomic operations at runtime using libNVRTC.
- ▶ Added `0_Simple/simpleTemplates_nvrtc`. Demonstrates compilation of templated dynamically allocated shared memory arrays CUDA kernel at runtime using libNVRTC.
- ▶ Added `0_Simple/simpleVoteIntrinsics_nvrtc`. Demonstrates compilation of CUDA kernel which uses vote intrinsics at runtime using libNVRTC.
- ▶ Added `0_Simple/vectorAdd_nvrtc`. Demonstrates compilation of CUDA kernel performing vector addition at runtime using libNVRTC.
- ▶ Added `4_Finance/binomialOptions_nvrtc`. Demonstrates runtime compilation using libNVRTC of CUDA kernel which evaluates fair call price for a given set of European options under binomial model.
- ▶ Added `4_Finance/BlackScholes_nvrtc`. Demonstrates runtime compilation using libNVRTC of CUDA kernel which evaluates fair call and put prices for a given set of European options by Black-Scholes formula.
- ▶ Added `4_Finance/quasirandomGenerator_nvrtc`. Demonstrates runtime compilation using libNVRTC of CUDA kernel which implements Niederreiter Quasirandom Sequence Generator and Inverse Cumulative Normal Distribution functions for the generation of Standard Normal Distributions.

1.19. CUDA 6.5

- ▶ Added `7_CUDA Libraries/cuHook`. Demonstrates how to build and use an intercept library with CUDA.
- ▶ Added `7_CUDA Libraries/simpleCUFFT_callback`. Demonstrates how to compute a 1D-convolution of a signal with a filter using a user-supplied CUFFT callback routine, rather than a separate kernel call.
- ▶ Added `7_CUDA Libraries/simpleCUFFT_MGPU`. Demonstrates how to compute a 1D-convolution of a signal with a filter by transforming both into frequency domain, multiplying them together, and transforming the signal back to time domain on Multiple GPUs.

- ▶ Added `7_CUDA Libraries/simpleCUFFT_2d_MGPU`. Demonstrates how to compute a 2D-convolution of a signal with a filter by transforming both into frequency domain, multiplying them together, and transforming the signal back to time domain on Multiple GPUs.
- ▶ Removed `3_Imaging/cudaEncode`. Support for the CUDA Video Encoder (NVENC) has been removed.
- ▶ Removed `4_Finance/ExcelCUDA2007`. The topic will be covered in a blog post at [Parallel Forall](#).
- ▶ Removed `4_Finance/ExcelCUDA2010`. The topic will be covered in a blog post at [Parallel Forall](#).
- ▶ The `4_Finance/binomialOptions` sample is now restricted to running on GPUs with SM architecture 2.0 or greater.
- ▶ The `4_Finance/quasirandomGenerator` sample is now restricted to running on GPUs with SM architecture 2.0 or greater.
- ▶ The `7_CUDA Libraries/boxFilterNPP` sample now demonstrates how to use the static NPP libraries on Linux and Mac.
- ▶ The `7_CUDA Libraries/conjugateGradient` sample now demonstrates how to use the static CUBLAS and CUSPARSE libraries on Linux and Mac.
- ▶ The `7_CUDA Libraries/MersenneTwisterGP11213` sample now demonstrates how to use the static CURAND library on Linux and Mac.

1.20. CUDA 6.0

- ▶ New featured samples that support a new CUDA 6.0 feature called UVM-Lite
- ▶ Added `0_Simple/UnifiedMemoryStreams` - new CUDA sample that demonstrates the use of OpenMP and CUDA streams with Unified Memory on a single GPU.
- ▶ Added `1_Uilities/p2pBandwidthTestLatency` - new CUDA sample that demonstrates how measure latency between pairs of GPUs with P2P enabled and P2P disabled.
- ▶ Added `6_Advanced/StreamPriorities` - This sample demonstrates basic use of the new CUDA 6.0 feature stream priorities.
- ▶ Added `7_CUDA Libraries/ConjugateGradientUM` - This sample implements a conjugate gradient solver on GPU using cuBLAS and cuSPARSE library, using Unified Memory.

1.21. CUDA 5.5

- ▶ Linux makefiles have been updated to generate code for the ARMv7 architecture. Only the ARM hard-float floating point ABI is supported. Both native ARMv7 compilation and cross compilation from x86 is supported
- ▶ Performance improvements in CUDA toolkit for Kepler GPUs (SM 3.0 and SM 3.5)

- ▶ Makefiles projects have been updated to properly find search default paths for OpenGL, CUDA, MPI, and OpenMP libraries for all OS Platforms (Mac, Linux x86, Linux ARM).
- ▶ Linux and Mac project Makefiles now invoke NVCC for building and linking projects.
- ▶ Added `0_Simple/cppOverload` - new CUDA sample that demonstrates how to use C++ overloading with CUDA.
- ▶ Added `6_Advanced/cdpBezierTessellation` - new CUDA sample that demonstrates an advanced method of implementing Bezier Line Tessellation using CUDA Dynamic Parallelism. Requires compute capability 3.5 or higher.
- ▶ Added `7_CUDAlibraries/jpegNPP` - new CUDA sample that demonstrates how to use NPP for JPEG compression on the GPU.
- ▶ CUDA Samples now have better integration with Nsight Eclipse IDE.
- ▶ `6_Advanced/ptxjit` sample now includes a new API to demonstrate PTX linking at the driver level.

1.22. CUDA 5.0

- ▶ New directory structure for CUDA samples. Samples are classified accordingly to categories: `0_Simple`, `1_Uutilities`, `2_Graphics`, `3_Imaging`, `4_Finance`, `5_Simulations`, `6_Advanced`, and `7_CUDAlibraries`
- ▶ Added `0_Simple/simpleIPC` - CUDA Runtime API sample is a very basic sample that demonstrates Inter Process Communication with one process per GPU for computation. Requires Compute Capability 2.0 or higher and a Linux Operating System.
- ▶ Added `0_Simple/simpleSeparateCompilation` - demonstrates a CUDA 5.0 feature, the ability to create a GPU device static library and use it within another CUDA kernel. This example demonstrates how to pass in a GPU device function (from the GPU device static library) as a function pointer to be called. Requires Compute Capability 2.0 or higher.
- ▶ Added `2_Graphics/bindlessTexture` - demonstrates use of `cudaSurfaceObject`, `cudaTextureObject`, and MipMap support in CUDA. Requires Compute Capability 3.0 or higher.
- ▶ Added `3_Imaging/stereoDisparity` - demonstrates how to compute a stereo disparity map using SIMD SAD (Sum of Absolute Difference) intrinsics. Requires Compute Capability 2.0 or higher.
- ▶ Added `0_Simple/cdpSimpleQuicksort` - demonstrates a simple quicksort implemented using CUDA Dynamic Parallelism. This sample requires devices with compute capability 3.5 or higher.
- ▶ Added `0_Simple/cdpSimplePrint` - demonstrates simple printf implemented using CUDA Dynamic Parallelism. This sample requires devices with compute capability 3.5 or higher.

- ▶ Added `6_Advanced/cdpLUdecomposition` - demonstrates LU Decomposition implemented using CUDA Dynamic Parallelism. This sample requires devices with compute capability 3.5 or higher.
- ▶ Added `6_Advanced/cdpAdvancedQuicksort` - demonstrates an advanced quicksort implemented using CUDA Dynamic Parallelism. This sample requires devices with compute capability 3.5 or higher.
- ▶ Added `6_Advanced/cdpQuadtree` - demonstrates Quad Trees implemented using CUDA Dynamic Parallelism. This sample requires devices with compute capability 3.5 or higher.
- ▶ Added `7_CUDAlibraries/simpleDevLibCUBLAS` - implements a simple cuBLAS function calls that call GPU device API library running cuBLAS functions. cuBLAS device code functions take advantage of CUDA Dynamic Parallelism and requires compute capability of 3.5 or higher.

1.23. CUDA 4.2

- ▶ Added `segmentationTreeThrust` - demonstrates a method to build image segmentation trees using Thrust. This algorithm is based on Boruvka's MST algorithm.

1.24. CUDA 4.1

- ▶ Added `MersenneTwisterGP11213` - implements Mersenne Twister GP11213, a pseudorandom number generator using the `cuRAND` library.
- ▶ Added `HSOpticalFlow` - When working with image sequences or video it's often useful to have information about objects movement. Optical flow describes apparent motion of objects in image sequence. This sample is a Horn-Schunck method for optical flow written using CUDA.
- ▶ Added `volumeFiltering` - demonstrates basic volume rendering and filtering using 3D textures.
- ▶ Added `simpleCubeMapTexture` - demonstrates how to use `texcubemap` fetch instruction in a CUDA C program.
- ▶ Added `simpleAssert` - demonstrates how to use GPU assert in a CUDA C program.
- ▶ Added `grabcutNPP` - CUDA implementation of Rother et al. GrabCut approach using the 8 neighborhood NPP Graphcut primitive introduced in CUDA 4.1. (C. Rother, V. Kolmogorov, A. Blake. *GrabCut: Interactive Foreground Extraction Using Iterated Graph Cuts*. *ACM Transactions on Graphics (SIGGRAPH'04)*, 2004).

Chapter 2. Getting Started

The CUDA Samples are an educational resource provided to teach CUDA programming concepts. The CUDA Samples are not meant to be used for performance measurements.

For system requirements and installation instructions, please refer to the [Linux Installation Guide](#) and the [Windows Installation Guide](#).

2.1. Getting CUDA Samples

Windows

On Windows, the CUDA Samples are installed using the [CUDA Toolkit Windows Installer](#). By default, the CUDA Samples are installed in:

```
C:\ProgramData\NVIDIA Corporation\CUDA Samples\v11.6\
```

The installation location can be changed at installation time.

Linux

On Linux, to install the CUDA Samples, the CUDA toolkit must first be installed. See the [Linux Installation Guide](#) for more information on how to install the CUDA Toolkit.

Then the CUDA Samples can be installed by running the following command, where <target_path> is the location where to install the samples:

```
$ cuda-install-samples-11.6.sh <target_path>
```

2.2. Building Samples

Windows

The Windows samples are built using the Visual Studio IDE. Solution files (.sln) are provided for each supported version of Visual Studio, using the format:

```
*_vs<version>.sln - for Visual Studio <version>
```

Complete samples solution files exist at:

```
C:\ProgramData\NVIDIA Corporation\CUDA Samples\v11.6\
```

Each individual sample has its own set of solution files at:

```
C:\ProgramData\NVIDIA Corporation\CUDA Samples\v11.6\\
```

To build/examine all the samples at once, the complete solution files should be used. To build/examine a single sample, the individual sample solution files should be used.

Linux

The Linux samples are built using makefiles. To use the makefiles, change the current directory to the sample directory you wish to build, and run `make`:

```
$ cd <sample_dir>
$ make
```

The samples makefiles can take advantage of certain options:

- ▶ **TARGET_ARCH=<arch>** - cross-compile targeting a specific architecture. Allowed architectures are `x86_64`, `armv7l`, `aarch64`, `sbsa`, and `ppc64le`.

By default, `TARGET_ARCH` is set to `HOST_ARCH`. On a `x86_64` machine, not setting `TARGET_ARCH` is the equivalent of setting `TARGET_ARCH=x86_64`.

```
$ make TARGET_ARCH=x86_64
$ make TARGET_ARCH=armv7l
$ make TARGET_ARCH=aarch64
$ make TARGET_ARCH=sbsa
$ make TARGET_ARCH=ppc64le
```

See [here](#) for more details.

- ▶ **dbg=1** - build with debug symbols

```
$ make dbg=1
```

- ▶ **SMS="A B ..."** - override the SM architectures for which the sample will be built, where "A B ..." is a space-delimited list of SM architectures. For example, to generate SASS for SM 35 and SM 50, use `SMS="35 50"`.

```
$ make SMS="35 50"
```

- ▶ **HOST_COMPILER=<host_compiler>** - override the default g++ host compiler. See the [Linux Installation Guide](#) for a list of supported host compilers.

```
$ make HOST_COMPILER=g++
```

2.3. CUDA Cross-Platform Samples

CUDA Samples are now located in <https://github.com/nvidia/cuda-samples>, which includes instructions for obtaining, building, and running the samples.

2.4. Using CUDA Samples to Create Your Own CUDA Projects

2.4.1. Creating CUDA Projects for Windows

Creating a new CUDA Program using the CUDA Samples infrastructure is easy. We have provided a `template` project that you can copy and modify to suit your needs. Just follow these steps:

[<category> refers to one of the following folders: 0_Simple, 1_Utillities, 2_Graphics, 3_Imaging, 4_Finance, 5_Simulations, 6_Advanced, 7_CUDA Libraries.]

1. Copy the content of:

```
C:\ProgramData\NVIDIA Corporation\CUDA Samples\v11.6\<category>\template
```

to a directory of your own:

```
C:\ProgramData\NVIDIA Corporation\CUDA Samples\v11.6\<category>\myproject
```

2. Edit the filenames of the project to suit your needs.
3. Edit the `*.sln`, `*.vcproj` and source files.
Just search and replace all occurrences of `template` with `myproject`.
4. Build the 64-bit, release or debug configurations using:

```
myproject_vs<version>.sln
```

5. Run `myproject.exe` from the release or debug directories located in:

```
C:\ProgramData\NVIDIA Corporation\CUDA Samples\v11.6\bin\win64\[release|debug]
```

6. Now modify the code to perform the computation you require.
See the *CUDA Programming Guide* for details of programming in CUDA.

2.4.2. Creating CUDA Projects for Linux



Note: The default installation folder `<SAMPLES_INSTALL_PATH>` is `NVIDIA_CUDA_11.6_Samples` and `<category>` is one of the following: 0_Simple, 1_Utillities, 2_Graphics, 3_Imaging, 4_Finance, 5_Simulations, 6_Advanced, 7_CUDA Libraries.

Creating a new CUDA Program using the NVIDIA CUDA Samples infrastructure is easy. We have provided a `template` project that you can copy and modify to suit your needs. Just follow these steps:

1. Copy the `template` project:

```
cd <SAMPLES_INSTALL_PATH>/<category>
cp -r template <myproject>
```



```
cd <SAMPLES_INSTALL_PATH>/<category>
```

2. Edit the filenames of the project to suit your needs:

```
mv template.cu myproject.cu  
mv template_cpu.cpp myproject_cpu.cpp
```

3. Edit the Makefile and source files.

Just search and replace all occurrences of `template` with `myproject`.

4. Build the project as (release):

```
make
```

To build the project as (debug), use "make dbg=1":

```
make dbg=1
```

5. Run the program:

```
../../bin/x86_64/linux/release/myproject
```

6. Now modify the code to perform the computation you require.

See the *CUDA Programming Guide* for details of programming in CUDA.

Chapter 3. Samples Reference

This document contains a complete listing of the code samples that are included with the NVIDIA CUDA Toolkit. It describes each code sample, lists the minimum GPU specification, and provides links to the source code and white papers if available.

The code samples are divided into the following categories:

Introduction Reference

Basic CUDA samples for beginners that illustrate key concepts with using CUDA and CUDA runtime APIs.

Utilities Reference

Utility samples that demonstrate how to query device capabilities and measure GPU/CPU bandwidth.

Concepts and Techniques Reference

Samples that demonstrate CUDA related concepts and common problem solving techniques.

CUDA Features Reference

Samples that demonstrate CUDA Features.

CUDA Libraries Reference

Samples that demonstrate how to use CUDA platform libraries (NPP, NVJPEG, NVGRAPH, cuBLAS, cuFFT, cuSPARSE, cuSOLVER and cuRAND).

Domain Specific Reference

Samples that are specific to domain (Graphics, Finance, Image Processing).

Performance Reference

Samples that demonstrate performance optimization.

3.1. Introduction Reference

asyncAPI

This sample illustrates the usage of CUDA events for both GPU timing and overlapping CPU and GPU execution. Events are inserted into a stream of CUDA calls. Since CUDA stream calls are asynchronous, the CPU can perform computations while GPU is executing (including DMA memcopies between the host and device). CPU can query CUDA events to determine whether GPU has completed tasks.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaEventRecord , cudaMallocHost , cudaProfilerStart , cudaEventCreate , cudaEventElapsedTime , cudaDeviceSynchronize , cudaFreeHost , cudaMalloc , cudaEventQuery , cudaProfilerStop , cudaEventDestroy , cudaMemcpyAsync , cudaGetDeviceProperties
Key Concepts	Asynchronous Data Transfers , CUDA Streams and Events
Supported OSes	Linux , Windows

c++11_cuda - C++11 CUDA

This sample demonstrates C++11 feature support in CUDA. It scans a input text file and prints no. of occurrences of x, y, z, w characters.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CPP11
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaMemset , cudaFree , cudaMemcpy
Key Concepts	CPP11 CUDA
Supported OSes	Linux , Windows

clock - Clock

This example shows how to use the clock function to measure the performance of block of threads of a kernel accurately.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaFree , cudaMemcpy
Key Concepts	Performance Strategies
Supported OSes	Linux , Windows

clock_nvrtc - Clock libNVRTC

This example shows how to use the clock function using libNVRTC to measure the performance of block of threads of a kernel accurately.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NVRTC
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuModuleGetFunction , cuMemAlloc , cuLaunchKernel , cuCtxSynchronize , cuMemFree , cuMemcpyDtoH , cuMemcpyHtoD , cudaBlockSize , cudaGridSize
Key Concepts	Performance Strategies , Runtime Compilation
Supported OSes	Linux , Windows

concurrentKernels - Concurrent Kernels

This sample demonstrates the use of CUDA streams for concurrent execution of several kernels on GPU device. It also illustrates how to introduce dependencies between CUDA streams with the new `cudaStreamWaitEvent` function.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaStreamWaitEvent , cudaStreamDestroy , cudaFree , cudaEventRecord , cudaMallocHost , cudaStreamCreate , cudaEventCreate , cudaEventElapsedTime , cudaEventSynchronize , cudaFreeHost , cudaMalloc , cudaEventCreateWithFlags , cudaEventDestroy , cudaMemcpyAsync , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Performance Strategies
Supported OSes	Linux , Windows

cppIntegration - C++ Integration

This example demonstrates how to integrate CUDA into an existing C++ application, i.e. the CUDA entry point on host side is only a function which is called from C++ code and only the file containing this function is compiled with `nvcc`. It also demonstrates that vector types can be used from `cpp`.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaFree , cudaMemcpy
Key Concepts	CPP-CUDA Integration
Supported OSes	Linux , Windows

cppOverload

This sample demonstrates how to use C++ function overloading on the GPU.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaFuncSetCacheConfig , cudaMallocHost , cudaFuncGetAttributes , cudaGetDeviceCount , cudaDeviceSynchronize , cudaFreeHost , cudaMalloc , cudaSetDevice , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	C++ Function Overloading , CUDA Streams and Events
Supported OSes	Linux , Windows

cudaOpenMP

This sample demonstrates how to use OpenMP API to write an application for multiple GPUs.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	OpenMP
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaGetDeviceCount , cudaSetDevice , cudaMalloc , cudaGetLastError , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	CUDA Systems Integration , OpenMP , Multithreading
Supported OSes	Linux , Windows

fp16ScalarProduct - FP16 Scalar Product

Calculates scalar product of two vectors of FP16 numbers.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	FP16
Supported SM Architecture	SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaMallocHost , cudaFreeHost , cudaMalloc , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	CUDA Runtime API
Supported OSes	Linux , Windows

matrixMul - Matrix Multiplication (CUDA Runtime API Version)

This sample implements matrix multiplication and is exactly the same as Chapter 6 of the programming guide. It has been written for clarity of exposition to illustrate various CUDA programming principles, not with the goal of providing the most performant generic kernel for matrix multiplication. To illustrate GPU performance for matrix multiply, this sample also shows how to use the new CUDA 4.0 interface for CUBLAS to demonstrate high-performance performance for matrix multiplication.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaEventRecord , cudaMallocHost , cudaProfilerStart , cudaEventCreate , cudaEventElapsedTime , cudaEventSynchronize , cudaFreeHost , cudaMalloc , cudaProfilerStop , cudaStreamCreateWithFlags , cudaEventDestroy , cudaStreamSynchronize , cudaMemcpyAsync
Key Concepts	CUDA Runtime API , Linear Algebra
Supported OSes	Linux , Windows

matrixMul_nvrtc - Matrix Multiplication with libNVRTC

This sample implements matrix multiplication and is exactly the same as Chapter 6 of the programming guide. It has been written for clarity of exposition to illustrate various CUDA programming principles, not with the goal of providing the most performant generic kernel for matrix multiplication. To illustrate GPU performance for matrix multiply, this sample also

shows how to use the new CUDA 4.0 interface for CUBLAS to demonstrate high-performance performance for matrix multiplication.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NVRTC
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuModuleGetFunction , cuMemAlloc , cuLaunchKernel , cuCtxSynchronize , cuMemFree , cuMemcpyDtoH , cuMemcpyHtoD
Key Concepts	CUDA Runtime API, Linear Algebra , Runtime Compilation
Supported OSes	Linux , Windows

matrixMulDrv - Matrix Multiplication (CUDA Driver API Version)

This sample implements matrix multiplication and uses the new CUDA 4.0 kernel launch Driver API. It has been written for clarity of exposition to illustrate various CUDA programming principles, not with the goal of providing the most performant generic kernel for matrix multiplication. CUBLAS provides high-performance matrix multiplication.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuModuleGetFunction , cuMemcpyHtoD , cuModuleLoadData , cuCtxCreate , cuLaunchKernel , cuDeviceGetName , cuMemAlloc , cuOccupancyMaxPotentialBlockSize , cuDeviceTotalMem , cuMemFree , cuMemcpyDtoH , cuCtxDestroy , cuDeviceGetAttribute
Key Concepts	CUDA Driver API, Matrix Multiply
Supported OSes	Linux , Windows

matrixMulDynlinkJIT - Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version)

This sample revisits matrix multiplication using the CUDA driver API. It demonstrates how to link to CUDA driver at runtime and how to use JIT (just-in-time) compilation from PTX code. It has been written for clarity of exposition to illustrate various CUDA programming principles,

not with the goal of providing the most performant generic kernel for matrix multiplication. CUBLAS provides high-performance matrix multiplication.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuParamSetv , cuMemFree , culnit , cuModuleGetFunction , cuCtxDestroy , cuCtxCreate , cuDeviceGetName , cuCtxSynchronize , cuParamSeti , cuModuleLoadDataEx , cuDeviceGet , cuFuncSetSharedSize , cuMemAlloc , cuDeviceComputeCapability , cuFuncSetBlockShape , cuMemcpyHtoD , cuParamSetSize , cuLaunchGrid , cuDeviceGetCount , cuLaunchKernel , cuMemcpyDtoH
Key Concepts	CUDA Driver API , CUDA Dynamically Linked Library
Supported OSes	Linux , Windows

mergeSort - Merge Sort

This sample implements a merge sort (also known as Batcher's sort), algorithms belonging to the class of sorting networks. While generally subefficient on large sequences compared to algorithms with better asymptotic algorithmic complexity (i.e. merge sort or radix sort), may be the algorithms of choice for sorting batches of short- to mid-sized (key, value) array pairs. Refer to the excellent tutorial by H. W. Lang <http://www.iti.fh-flensburg.de/lang/algorithmen/sortieren/networks/indexen.htm>

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaFree , cudaDeviceSynchronize , cudaMemcpy
Key Concepts	Data-Parallel Algorithms
Supported OSes	Linux , Windows

simpleAssert

This CUDA Runtime API sample is a very basic sample that implements how to use the assert function in the device code. Requires Compute Capability 2.0 .

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDeviceSynchronize , cudaGetErrorString
Key Concepts	Assert
Supported OSes	Linux , Windows

simpleAssert_nvrtc - simpleAssert with libNVRTC

This CUDA Runtime API sample is a very basic sample that implements how to use the assert function in the device code. Requires Compute Capability 2.0 .

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NVRTC
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuModuleGetFunction , cuCtxSynchronize , cuLaunchKernel
Key Concepts	Assert , Runtime Compilation
Supported OSes	Linux , Windows

simpleAtomicIntrinsics - Simple Atomic Intrinsics

A simple demonstration of global memory atomic instructions.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaMallocHost , cudaFreeHost , cudaMalloc , cudaStreamCreateWithFlags , cudaStreamSynchronize , cudaMemcpyAsync
Key Concepts	Atomic Intrinsics
Supported OSes	Linux , Windows

simpleAtomicIntrinsics_nvrtc - Simple Atomic Intrinsics with libNVRTC

A simple demonstration of global memory atomic instructions. This sample makes use of NVRTC for Runtime Compilation.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NVRTC
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuModuleGetFunction , cuMemAlloc , cuLaunchKernel , cuCtxSynchronize , cuMemFree , cuMemcpyDtoH , cuMemcpyHtoD , cudaBlockSize , cudaGridSize
Key Concepts	Atomic Intrinsic s, Runtime Compilation
Supported OSes	Linux , Windows

simpleAttributes

This CUDA Runtime API sample is a very basic example that implements how to use the stream attributes that affect L2 locality. Performance improvement due to use of L2 access policy window can only be noticed on Compute capability 8.0 or higher.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDeviceSetLimit , cudaFree , cudaMallocHost , cudaStreamCreate , cudaFreeHost , cudaMalloc , cudaStreamSetAttribute , cudaMemcpyAsync , cudaStreamSynchronize , cudaGetDeviceProperties
Key Concepts	Attributes usage on stream
Supported OSes	Linux , Windows

simpleAWBarrier - Simple Arrive Wait Barrier

A simple demonstration of arrive wait barriers.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CPP11 , MBCG
Supported SM Architecture	SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaMallocHost , cudaOccupancyMaxActiveBlocksPerMultiprocessor , cudaOccupancyMaxPotentialBlockSize , cudaDeviceGetAttribute , cudaFreeHost , cudaMalloc , cudaStreamCreateWithFlags , cudaLaunchCooperativeKernel , cudaStreamSynchronize , cudaMemcpyAsync

Key Concepts	Arrive Wait Barrier
Supported OSes	Linux , Windows

simpleCallback - Simple CUDA Callbacks

This sample implements multi-threaded heterogeneous computing workloads with the new CPU callbacks for CUDA streams and events introduced with CUDA 5.0.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaStreamDestroy , cudaFree , cudaStreamCreate , cudaHostAlloc , cudaGetDeviceCount , cudaFreeHost , cudaMalloc , cudaSetDevice , cudaMemcpyAsync , cudaStreamAddCallback , cudaGetDeviceProperties
Key Concepts	CUDA Streams , Callback Functions , Multithreading
Supported OSes	Linux , Windows

simpleCooperativeGroups - Simple Cooperative Groups

This sample is a simple code that illustrates basic usage of cooperative groups within the thread block.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDeviceSynchronize , cudaGetErrorString
Key Concepts	Cooperative Groups
Supported OSes	Linux , Windows

simpleCubemapTexture - Simple Cubemap Texture

Simple example that demonstrates how to use a new CUDA 4.1 feature to support cubemap Textures in CUDA C.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaFreeArray , cudaExtent , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaPitchedPtr , cudaCreateTextureObject , cudaMalloc , cudaCreateChannelDesc , cudaPos , cudaMemcpy , cudaGetDeviceProperties

Key Concepts	Texture , Volume Processing
Supported OSes	Linux , Windows

simpleCUDA2GL - CUDA and OpenGL Interop of Images

This sample shows how to copy CUDA image back to OpenGL using the most efficient methods.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaGraphicsUnregisterResource , cudaMemcpyToArray , cudaGraphicsGLRegisterBuffer , cudaHostAlloc , cudaGraphicsResourceGetMappedPointer , cudaProcess , cudaDeviceSynchronize , cudaMalloc , cudaGraphicsSubResourceGetMappedArray , cudaGraphicsGLRegisterImage , cudaGraphicsUnmapResources
Key Concepts	Graphics Interop , Image Processing , Performance Strategies
Supported OSes	Linux , Windows

simpleDrvRuntime - Simple Driver-Runtime Interaction

A simple example which demonstrates how CUDA Driver and Runtime APIs can work together to load cuda fatbinary of vector add kernel and performing vector addition.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuCtxDestroy , cuModuleLoadData , cuCtxCreate , cuLaunchKernel , cuModuleUnload , cunjit , cuModuleGetFunction , cudaFree , cudaMallocHost , cudaFreeHost , cudaMalloc , cudaStreamCreateWithFlags , cudaStreamSynchronize , cudaMemcpyAsync
Key Concepts	CUDA Driver API , CUDA Runtime API , Vector Addition
Supported OSes	Linux , Windows

simpleHyperQ

This sample demonstrates the use of CUDA streams for concurrent execution of several kernels on devices which provide HyperQ (SM 3.5). Devices without HyperQ (SM 2.0 and SM 3.0) will run a maximum of two kernels concurrently.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaStreamDestroy , cudaFree , cudaEventRecord , cudaMallocHost , cudaStreamCreate , cudaEventCreate , cudaEventElapsedTime , cudaEventSynchronize , cudaFreeHost , cudaMalloc , cudaEventDestroy , cudaMemcpy , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	CUDA Systems Integration , Performance Strategies
Supported OSes	Linux , Windows
Whitepaper	HyperQ.pdf

simpleIPC

This CUDA Runtime API sample is a very basic sample that demonstrates Inter Process Communication with one process per GPU for computation. Requires Compute Capability 3.0 or higher and a Linux Operating System, or a Windows Operating System.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	IPC
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDeviceEnablePeerAccess , cudalpcOpenEventHandle , cudaOccupancyMaxActiveBlocksPerMultiprocessor , cudaStreamCreateWithFlags , cudaDeviceCanAccessPeer , cudaMemcpyAsync , cudaStreamDestroy , cudaEventCreate , cudaMalloc , cudaEventDestroy , cudaSetDevice , cudalpcOpenMemHandle , cudaGetDeviceProperties , cudaGetDeviceCount , cudalpcGetEventHandle , cudaGetLastError , cudaStreamSynchronize , cudaStreamWaitEvent , cudaFree , cudalpcCloseMemHandle , cudaEventRecord , cudalpcGetMemHandle , cudaEventSynchronize
Key Concepts	CUDA Systems Integration , Peer to Peer , InterProcess Communication
Supported OSes	Linux , Windows

simpleLayeredTexture - Simple Layered Texture

Simple example that demonstrates how to use a new CUDA 4.0 feature to support layered Textures in CUDA C.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaFreeArray , cudaExtent , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaPitchedPtr , cudaCreateTextureObject , cudaMalloc , cudaCreateChannelDesc , cudaPos , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Texture , Volume Processing
Supported OSes	Linux , Windows

simpleMPI

Simple example demonstrating how to use MPI in combination with CUDA.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	MPI
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaGetLastError , cudaFree , cudaMemcpy
Key Concepts	CUDA Systems Integration , MPI , Multithreading
Supported OSes	Linux , Windows

simpleMultiCopy - Simple Multi Copy and Compute

Supported in GPUs with Compute Capability 1.1, overlapping compute with one memcpy is possible from the host system. For Quadro and Tesla GPUs with Compute Capability 2.0, a second overlapped copy operation in either direction at full speed is possible (PCI-e is symmetric). This sample illustrates the usage of CUDA streams to achieve overlapping of kernel execution with data copies to and from the device.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaStreamDestroy , cudaEventRecord , cudaStreamCreate , cudaHostAlloc , cudaEventCreate , cudaEventElapsedTime , cudaDeviceSynchronize , cudaEventSynchronize , cudaFreeHost , cudaMalloc , cudaEventDestroy , cudaSetDevice , cudaMemcpyAsync , cudaGetDeviceProperties
Key Concepts	CUDA Streams and Events , Asynchronous Data Transfers , Overlap Compute and Copy , GPU Performance
Supported OSes	Linux , Windows

simpleMultiGPU - Simple Multi-GPU

This application demonstrates how to use the new CUDA 4.0 API for CUDA context management and multi-threaded access to run CUDA kernels on multiple-GPUs.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaStreamDestroy , cudaFree , cudaMallocHost , cudaStreamCreate , cudaGetDeviceCount , cudaFreeHost , cudaMalloc , cudaSetDevice , cudaStreamSynchronize , cudaMemcpyAsync
Key Concepts	Asynchronous Data Transfers , CUDA Streams and Events , Multithreading , Multi-GPU
Supported OSes	Linux , Windows

simpleOccupancy

This sample demonstrates the basic usage of the CUDA occupancy calculator and occupancy-based launch configurator APIs by launching a kernel with the launch configurator, and measures the utilization difference against a manually configured launch.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaEventRecord , cudaOccupancyMaxActiveBlocksPerMultiprocessor , cudaEventCreate , cudaOccupancyMaxPotentialBlockSize , cudaEventElapsedTime , cudaDeviceSynchronize , cudaMalloc , cudaMemcpy , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Occupancy Calculator
Supported OSes	Linux , Windows

simpleP2P - Simple Peer-to-Peer Transfers with Multi-GPU

This application demonstrates CUDA APIs that support Peer-To-Peer (P2P) copies, Peer-To-Peer (P2P) addressing, and Unified Virtual Memory Addressing (UVA) between multiple GPUs. In general, P2P is supported between two same GPUs with some exceptions, such as some Tesla and Quadro GPUs.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	only-64-bit
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDeviceEnablePeerAccess , cudaFree , cudaEventRecord , cudaMallocHost , cudaGetDeviceCount , cudaEventElapsedTime , cudaDeviceSynchronize , cudaEventSynchronize , cudaFreeHost , cudaMalloc , cudaEventCreateWithFlags , cudaDeviceCanAccessPeer , cudaEventDestroy , cudaSetDevice , cudaDeviceDisablePeerAccess , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Performance Strategies , Asynchronous Data Transfers , Unified Virtual Address Space , Peer to Peer Data Transfers , Multi-GPU
Supported OSes	Linux , Windows

simplePitchLinearTexture - Pitch Linear Texture

Use of Pitch Linear Textures

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaMemcpyToArray , cudaEventRecord , cudaFreeArray , cudaEventCreate , cudaEventElapsedTime , cudaDestroyTextureObject , cudaEventSynchronize , cudaMallocPitch , cudaCreateTextureObject , cudaEventDestroy , cudaMallocArray
Key Concepts	Texture , Image Processing
Supported OSes	Linux , Windows

simplePrintf

This basic CUDA Runtime API sample demonstrates how to use the printf function in the device code.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDeviceSynchronize , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Debugging
Supported OSes	Linux , Windows

simpleSeparateCompilation - Simple Static GPU Device Library

This sample demonstrates a CUDA 5.0 feature, the ability to create a GPU device static library and use it within another CUDA kernel. This example demonstrates how to pass in a GPU device function (from the GPU device static library) as a function pointer to be called. This sample requires devices with compute capability 2.0 or higher.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaMalloc , cudaGetLastError , cudaMemcpy , cudaMemcpyFromSymbol
Key Concepts	Separate Compilation
Supported OSes	Linux , Windows

simpleStreams

This sample uses CUDA streams to overlap kernel executions with memory copies between the host and a GPU device. This sample uses a new CUDA 4.0 feature that supports pinning of generic host memory. Requires Compute Capability 2.0 or higher.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemcpy , cudaMemcpyAsync , cudaStreamDestroy , cudaMallocHost , cudaHostAlloc , cudaHostRegister , cudaMalloc , cudaEventCreateWithFlags , cudaEventDestroy , cudaSetDeviceFlags , cudaSetDevice , cudaGetDeviceProperties , cudaGetDeviceCount , cudaEventElapsedTime , cudaMemset , cudaFree , cudaEventRecord , cudaStreamCreate , cudaEventSynchronize , cudaFreeHost , cudaHostUnregister

Key Concepts	Asynchronous Data Transfers , CUDA Streams and Events
Supported OSes	Linux , Windows

simpleSurfaceWrite - Simple Surface Write

Simple example that demonstrates the use of 2D surface references (Write-to-Texture)

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaMemcpyToArray , cudaFreeArray , cudaDestroySurfaceObject , cudaCreateSurfaceObject , cudaDestroyTextureObject , cudaDeviceSynchronize , cudaCreateTextureObject , cudaMalloc , cudaMallocArray , cudaCreateChannelDesc , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Texture , Surface Writes , Image Processing
Supported OSes	Linux , Windows

simpleTemplates - Simple Templates

This sample is a templated version of the template project. It also shows how to correctly templatize dynamically allocated shared memory arrays.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaFree , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	C++ Templates
Supported OSes	Linux , Windows

simpleTemplates_nvrtc - Simple Templates with libNVRTC

This sample is a templated version of the template project. It also shows how to correctly templatize dynamically allocated shared memory arrays.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NVRTC
---------------------	-----------------------

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuModuleGetFunction , cuMemAlloc , cuLaunchKernel , cuCtxSynchronize , cuMemFree , cuMemcpyDtoH , cuMemcpyHtoD
Key Concepts	C++ Templates , Runtime Compilation
Supported OSes	Linux , Windows

simpleTexture - Simple Texture

Simple example that demonstrates use of Textures in CUDA.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaMemcpyToArray , cudaFreeArray , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaCreateTextureObject , cudaMalloc , cudaMallocArray , cudaCreateChannelDesc , cudaMemcpy
Key Concepts	CUDA Runtime API , Texture , Image Processing
Supported OSes	Linux , Windows

simpleTexture3D - Simple Texture 3D

Simple example that demonstrates use of 3D Textures in CUDA.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaFreeArray , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaExtent , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaPitchedPtr , cudaCreateTextureObject , cudaMalloc , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaMemcpy
Key Concepts	Graphics Interop , Image Processing , 3D Textures , Surface Writes
Supported OSes	Linux , Windows

simpleTextureDrv - Simple Texture (Driver Version)

Simple example that demonstrates use of Textures in CUDA. This sample uses the new CUDA 4.0 kernel launch Driver API.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuModuleGetFunction , cuTexObjectDestroy , cuModuleLoadData , cuCtxCreate , cuArrayCreate , cuDeviceGetName , cuLaunchKernel , cuMemAlloc , cuCtxSynchronize , cuArrayDestroy , cuTexObjectCreate , cuMemFree , cuMemcpyDtoH , cuCtxDestroy , cuDeviceGetAttribute
Key Concepts	CUDA Driver API , Texture , Image Processing
Supported OSes	Linux , Windows

simpleVoteIntrinsics - Simple Vote Intrinsics

Simple program which demonstrates how to use the Vote (`__any_sync`, `__all_sync`) intrinsic instruction in a CUDA kernel.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaDeviceSynchronize , cudaMalloc , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Vote Intrinsics
Supported OSes	Linux , Windows

simpleVoteIntrinsics_nvrtc - Simple Vote Intrinsics with libNVRTC

Simple program which demonstrates how to use the Vote (any, all) intrinsic instruction in a CUDA kernel with runtime compilation using NVRTC APIs. Requires Compute Capability 2.0 or higher.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NVRTC
---------------------	-----------------------

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuModuleGetFunction , cuMemAlloc , cuLaunchKernel , cuCtxSynchronize , cuMemFree , cuMemcpyDtoH , cuMemcpyHtoD
Key Concepts	Vote Intrinsic , CUDA Driver API , Runtime Compilation
Supported OSes	Linux , Windows

simpleZeroCopy

This sample illustrates how to use Zero MemCopy, kernels can read and write directly to pinned system memory.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaHostAlloc , cudaGetDeviceCount , cudaHostRegister , cudaDeviceSynchronize , cudaFreeHost , cudaHostUnregister , cudaSetDeviceFlags , cudaSetDevice , cudaHostGetDevicePointer , cudaGetDeviceProperties
Key Concepts	Performance Strategies , Pinned System Paged Memory , Vector Addition
Supported OSes	Linux , Windows

systemWideAtomics - System wide Atomics

A simple demonstration of system wide atomic instructions.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	UVM
Supported SM Architecture	SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaDeviceSynchronize , cudaMallocManaged , cudaGetDeviceProperties
Key Concepts	Atomic Intrinsic , Unified Memory
Supported OSes	Linux

template - Template

A trivial template project that can be used as a starting point to create new CUDA projects.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaFree , cudaMemcpy
Key Concepts	Device Memory Allocation
Supported OSes	Linux , Windows

UnifiedMemoryStreams - Unified Memory Streams

This sample demonstrates the use of OpenMP and streams with Unified Memory on a single GPU.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	OpenMP , UVM , CUBLAS
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaStreamDestroy , cudaFree , cudaMallocManaged , cudaStreamCreate , cudaDeviceSynchronize , cudaStreamAttachMemAsync , cudaSetDevice , cudaStreamSynchronize , cudaGetDeviceProperties
Key Concepts	CUDA Systems Integration , OpenMP , CUBLAS , Multithreading , Unified Memory , CUDA Streams and Events
Supported OSes	Linux , Windows

vectorAdd - Vector Addition

This CUDA Runtime API sample is a very basic sample that implements element by element vector addition. It is the same as the sample illustrating Chapter 3 of the programming guide with some additions like error checking.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaMalloc , cudaGetLastError , cudaMemcpy , cudaGetErrorString
Key Concepts	CUDA Runtime API , Vector Addition

Supported OSes [Linux](#), [Windows](#)

vectorAdd_nvrtc - Vector Addition with libNVRTC

This CUDA Driver API sample uses NVRTC for runtime compilation of vector addition kernel. Vector addition kernel demonstrated is the same as the sample illustrating Chapter 3 of the programming guide.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies [NVRTC](#)

Supported SM Architecture [SM 3.5](#), [SM 3.7](#), [SM 5.0](#), [SM 5.2](#), [SM 5.3](#), [SM 6.0](#), [SM 6.1](#), [SM 6.2](#), [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

CUDA API [cuModuleGetFunction](#), [cuMemAlloc](#), [cuLaunchKernel](#), [cuCtxSynchronize](#), [cuMemFree](#), [cuMemcpyDtoH](#), [cuMemcpyHtoD](#), [cudaBlockSize](#), [cudaGridSize](#)

Key Concepts [CUDA Driver API](#), [Vector Addition](#), [Runtime Compilation](#)

Supported OSes [Linux](#), [Windows](#)

vectorAddDrv - Vector Addition Driver API

This Vector Addition sample is a basic sample that is implemented element by element. It is the same as the sample illustrating Chapter 3 of the programming guide with some additions like error checking. This sample also uses the new CUDA 4.0 kernel launch Driver API.

Supported SM Architecture [SM 3.5](#), [SM 3.7](#), [SM 5.0](#), [SM 5.2](#), [SM 5.3](#), [SM 6.0](#), [SM 6.1](#), [SM 6.2](#), [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

CUDA API [cuModuleGetFunction](#), [cuModuleLoadData](#), [cuCtxCreate](#), [cuLaunchKernel](#), [cuMemAlloc](#), [cuMemcpyDtoH](#), [cuCtxSynchronize](#), [cuMemFree](#), [cuInit](#), [cuCtxDestroy](#), [cuMemcpyHtoD](#)

Key Concepts [CUDA Driver API](#), [Vector Addition](#)

Supported OSes [Linux](#), [Windows](#)

vectorAddMMAP - Vector Addition cuMemMap

This sample replaces the device allocation in the vectorAddDrv sample with cuMemMap-ed allocations. This sample demonstrates that the cuMemMap api allows the user to specify the physical properties of their memory while retaining the contiguous nature of their access, thus not requiring a change in their program structure.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuMemSetAccess , culnit , cuMemAddressReserve , cuModuleGetFunction , cuCtxDestroy , cuCtxCreate , cuMemAddressFree , cuMemGetAllocationGranularity , cuMemUnmap , cuDeviceGetAttribute , cuMemRelease , cuModuleLoadData , cuMemMap , cuMemCreate , cuMemcpyHtoD , cuDeviceCanAccessPeer , cuDeviceGetCount , cuLaunchKernel , cuMemcpyDtoH
Key Concepts	CUDA Driver API , Vector Addition , MMAP
Supported OSes	Linux , Windows

3.2. Utilities Reference

bandwidthTest - Bandwidth Test

This is a simple test program to measure the memcpy bandwidth of the GPU and memcpy bandwidth across PCI-e. This test application is capable of measuring device to device copy bandwidth, host to device copy bandwidth for pageable and page-locked memory, and device to host copy bandwidth for pageable and page-locked memory.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaEventRecord , cudaMallocHost , cudaHostAlloc , cudaEventCreate , cudaGetDeviceCount , cudaEventElapsedTime , cudaDeviceSynchronize , cudaFreeHost , cudaMalloc , cudaEventDestroy , cudaSetDevice , cudaMemcpyAsync , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	CUDA Streams and Events , Performance Strategies
Supported OSes	Linux , Windows

deviceQuery - Device Query

This sample enumerates the properties of the CUDA devices present in the system.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuDeviceGetAttribute , cuSafeCallNoSync , cudaGetDeviceCount , cudaDriverGetVersion , cudaDeviceCanAccessPeer , cudaSetDevice , cudaRuntimeGetVersion , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	CUDA Runtime API , Device Query
Supported OSes	Linux , Windows

deviceQueryDrv - Device Query Driver API

This sample enumerates the properties of the CUDA devices present using CUDA Driver API calls

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuDeviceCanAccessPeer , cuDriverGetVersion , cuDeviceGetCount , cuDeviceGetName , cuDeviceTotalMem , culnit , cuDeviceGetAttribute
Key Concepts	CUDA Driver API , Device Query
Supported OSes	Linux , Windows

topologyQuery - Topology Query

A simple example on how to query the topology of a system with multiple GPU

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaGetDeviceCount , cudaDeviceGetAttribute
Key Concepts	Performance Strategies , Multi-GPU
Supported OSes	Linux , Windows

3.3. Concepts and Techniques Reference

boxFilter - Box Filter

Fast image box filter using CUDA with OpenGL rendering.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaFreeArray , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaCreateTextureObject , cudaMalloc , cudaMallocArray , cudaCreateChannelDesc , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaMemcpy , cudaGetErrorString
Key Concepts	Graphics Interop , Image Processing
Supported OSes	Linux , Windows

convolutionSeparable - CUDA Separable Convolution

This sample implements a separable convolution filter of a 2D signal with a gaussian kernel.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemcpyToSymbol , cudaFree , cudaDeviceSynchronize , cudaMalloc , cudaMemcpy
Key Concepts	Image Processing , Data Parallel Algorithms
Supported OSes	Linux , Windows
Whitepaper	convolutionSeparable.pdf

convolutionTexture - Texture-based Separable Convolution

Texture-based implementation of a separable 2D convolution with a gaussian kernel. Used for performance comparison against convolutionSeparable.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
----------------------------------	---

CUDA API	cudaMemcpyToSymbol , cudaFree , cudaMemcpyToArray , cudaFreeArray , cudaDeviceSynchronize , cudaCreateTextureObject , cudaMalloc , cudaMallocArray , cudaMemcpy
Key Concepts	Image Processing , Texture , Data Parallel Algorithms
Supported OSes	Linux , Windows

cuHook - CUDA Interception Library

This sample demonstrates how to build and use an intercept library with CUDA. The library has to be loaded via LD_PRELOAD, e.g. LD_PRELOAD=<full_path>/libcuhook.so.1 ./cuHook
NOTE: Sample will be waived if the glibc version >= 2.34, as the sample was using these private glibc functions `__libc_dlsym()`, `__libc_dlopen_mode()` which are not exposed in 2.34 version.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuDeviceGetCount , cuCtxCreate , cuMemAlloc , cuHookInfo , cuHookRegisterCallback , cuHook , cuMemFree , cuInit , cuCtxDestroy , cudaFree , cudaDeviceReset
Key Concepts	Debugging
Supported OSes	Linux

dct8x8 - DCT8x8

This sample demonstrates how Discrete Cosine Transform (DCT) for blocks of 8 by 8 pixels can be performed using CUDA: a naive implementation by definition and a more traditional approach used in many libraries. As opposed to implementing DCT in a fragment shader, CUDA allows for an easier and more efficient implementation.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaFreeArray , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaMallocPitch , cudaCreateTextureObject , cudaMallocArray
Key Concepts	Image Processing , Video Compression
Supported OSes	Linux , Windows
Whitepaper	dct8x8.pdf

EGLStream_CUDA_CrossGPU

Demonstrates CUDA and EGL Streams interop, where consumer's EGL Stream is on one GPU and producer's on other and both consumer-producer are different processes.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	EGL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuEGLStreamConsumerConnect , cuMemFree , culnit , cuStreamCreate , cuCtxCreate , cuGraphicsResourceGetMappedEglFrame , cuDeviceGetName , cuCtxSynchronize , cuEGLStreamConsumerAcquireFrame , cuDeviceGet , cuDeviceGetAttribute , cuMemAlloc , cuEGLStreamConsumerReleaseFrame , cuEGLStreamProducerDisconnect , cuEGLStreamProducerConnect , cuEGLStreamConsumerDisconnect , cuMemcpyHtoD , cuEGLStreamProducerReturnFrame , cuCtxPushCurrent , cuCtxPopCurrent , cuEGLStreamProducerPresentFrame cudaDeviceCreateConsumer , cudaFree , cudaConsumerReleaseFrame , cudaDeviceSynchronize , cudaGetValueMismatch , cudaProducerDeinit , cudaProducerPresentFrame , cudaMalloc , cudaProducerInit , cudaProducerReturnFrame , cudaProducerPrepareFrame , cudaConsumerAcquireFrame , cudaMemcpy , cudaGetErrorString , cudaDeviceCreateProducer
Key Concepts	EGLStreams Interop
Supported OSes	Linux

EGLStream_CUDA_Interop - EGLStream CUDA Interop

Demonstrates data exchange between CUDA and EGL Streams.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	EGL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7

CUDA API	cuEGLStreamConsumerConnect , cuArrayDestroy , cuMemFree , cuInit , cuCtxCreate , cuGraphicsResourceGetMappedEglFrame , cuDeviceGetName , cuCtxSynchronize , cuEGLStreamConsumerAcquireFrame , cuDeviceGetAttribute , cuMemcpy , cuMemAlloc , cuEGLStreamConsumerReleaseFrame , cuEGLStreamProducerDisconnect , cuEGLStreamProducerConnect , cuEGLStreamConsumerDisconnect , cuDeviceGetCount , cuEGLStreamProducerReturnFrame , cuCtxPushCurrent , cuCtxPopCurrent , cuMemcpyDtoH , cuEGLStreamProducerPresentFrame cudaDeviceCreateConsumer , cudaConsumerTest , cudaProducerDeinit , cudaProducerInit , cudaProducerReadYUVFrame , cudaProducerTest , cudaProducerReadARGBFrame , cudaDeviceCreateProducer
Key Concepts	EGLStreams Interop
Supported OSes	Linux

EGLSync_CUDA_Interop -EGLSync CUDA Event Interop

Demonstrates interoperability between CUDA Event and EGL Sync/EGL Image using which one can achieve synchronization on GPU itself for GL-EGL-CUDA operations instead of blocking CPU for synchronization.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	EGL , EGLSync , X11 , GLES
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuGraphicsEGLRegisterImage , cuStreamCreate , cuEventCreate , cuCtxCreate , cuGraphicsSubResourceGetMappedArray , cuGraphicsUnregisterResource , cuCtxSynchronize , cuEventCreateFromEGLSync , cuEventDestroy , cuStreamWaitEvent , cuCtxPushCurrent , cuSurfObjectCreate , cuInit , cuEventRecord , cuDeviceGetAttribute cudaFree , cudaDeviceSynchronize , cudaGetValueMismatch , cudaMalloc , cudaMemcpy , cudaGetErrorString
Key Concepts	EGLSync-CUDAEvent Interop , EGLImage-CUDA Interop
Supported OSes	Linux

eigenvalues - Eigenvalues

The computation of all or a subset of all eigenvalues is an important problem in Linear Algebra, statistics, physics, and many other fields. This sample demonstrates a parallel implementation of a bisection algorithm for the computation of all eigenvalues of a tridiagonal symmetric matrix of arbitrary size with CUDA.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaFree , cudaDeviceSynchronize , cudaMemcpy
Key Concepts	Linear Algebra
Supported OSes	Linux , Windows
Whitepaper	eigenvalues.pdf

FunctionPointers - Function Pointers

This sample illustrates how to use function pointers and implements the Sobel Edge Detection filter for 8-bit monochrome images.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemcpyToSymbol , cudaFree , cudaGraphicsMapResources , cudaFreeArray , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaCreateTextureObject , cudaMalloc , cudaMallocArray , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaMemcpy , cudaMemcpyFromSymbol
Key Concepts	Graphics Interop , Image Processing
Supported OSes	Linux , Windows

histogram - CUDA Histogram

This sample demonstrates efficient implementation of 64-bin and 256-bin histogram.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaDeviceSynchronize , cudaMalloc , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Image Processing , Data Parallel Algorithms
Supported OSes	Linux , Windows
Whitepaper	histogram.pdf

imageDenoising - Image denoising

This sample demonstrates two adaptive image denoising techniques: KNN and NLM, based on computation of both geometric and color distance between texels. While both techniques are implemented in the DirectX SDK using shaders, massively speeded up variation of the latter technique, taking advantage of shared memory, is implemented in addition to DirectX counterparts.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaGraphicsUnregisterResource , cudaFreeArray , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaDeviceSynchronize , cudaCreateTextureObject , cudaMalloc , cudaMallocArray , cudaGLRegisterBufferObject , cudaGraphicsUnmapResources , cudaMemcpy
Key Concepts	Image Processing
Supported OSes	Linux , Windows
Whitepaper	imageDenoising.pdf

inlinePTX - Using Inline PTX

A simple test application that demonstrates a new CUDA 4.0 ability to embed PTX in a CUDA kernel.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
----------------------------------	---

CUDA API	cudaFree , cudaMallocHost , cudaGridSize , cudaDeviceSynchronize , cudaBlockSize , cudaFreeHost , cudaMalloc , cudaGetLastError , cudaMemcpy
Key Concepts	Performance Strategies , PTX Assembly , CUDA Driver API
Supported OSes	Linux , Windows

inlinePTX_nvrtc - Using Inline PTX with libNVRTC

A simple test application that demonstrates a new CUDA 4.0 ability to embed PTX in a CUDA kernel.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NVRTC
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuMemAlloc , cuLaunchKernel , cuCtxSynchronize , cuMemFree , cuMemcpyDtoH , cuModuleGetFunction cudaBlockSize , cudaGridSize
Key Concepts	Performance Strategies , PTX Assembly , CUDA Driver API , Runtime Compilation
Supported OSes	Linux , Windows

interval - Interval Computing

Interval arithmetic operators example. Uses various C++ features (templates and recursion). The recursive mode requires Compute SM 2.0 capabilities.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDeviceSetLimit , cudaFree , cudaFuncSetCacheConfig , cudaEventRecord , cudaEventCreate , cudaEventElapsedTime , cudaDeviceSynchronize , cudaSetDevice , cudaMalloc , cudaEventDestroy , cudaGetLastError , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Recursion , Templates
Supported OSes	Linux , Windows

MC_EstimatePiInlineP - Monte Carlo Estimation of Pi (inline PRNG)

This sample uses Monte Carlo simulation for Estimation of Pi (using inline PRNG). This sample also uses the NVIDIA CURAND library.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CURAND
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaFuncGetAttributes , cudaGetDeviceCount , cudaMalloc , cudaSetDevice , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Random Number Generator , Computational Finance , CURAND Library
Supported OSes	Linux , Windows

MC_EstimatePiInlineQ - Monte Carlo Estimation of Pi (inline QRNG)

This sample uses Monte Carlo simulation for Estimation of Pi (using inline QRNG). This sample also uses the NVIDIA CURAND library.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CURAND
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaFuncGetAttributes , cudaGetDeviceCount , cudaMalloc , cudaSetDevice , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Random Number Generator , Computational Finance , CURAND Library
Supported OSes	Linux , Windows

MC_EstimatePiP - Monte Carlo Estimation of Pi (batch PRNG)

This sample uses Monte Carlo simulation for Estimation of Pi (using batch PRNG). This sample also uses the NVIDIA CURAND library.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CURAND
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaFuncGetAttributes , cudaGetDeviceCount , cudaMalloc , cudaSetDevice , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Random Number Generator , Computational Finance , CURAND Library
Supported OSes	Linux , Windows

MC_EstimatePiQ - Monte Carlo Estimation of Pi (batch QRNG)

This sample uses Monte Carlo simulation for Estimation of Pi (using batch QRNG). This sample also uses the NVIDIA CURAND library.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CURAND
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaFuncGetAttributes , cudaGetDeviceCount , cudaMalloc , cudaSetDevice , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Random Number Generator , Computational Finance , CURAND Library
Supported OSes	Linux , Windows

MC_SingleAsianOptionP - Monte Carlo Single Asian Option

This sample uses Monte Carlo to simulate Single Asian Options using the NVIDIA CURAND library.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CURAND
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaFuncGetAttributes , cudaGetDeviceCount , cudaMalloc , cudaSetDevice , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Random Number Generator , Computational Finance , CURAND Library
Supported OSes	Linux , Windows

particles - Particles

This sample uses CUDA to simulate and visualize a large set of particles and their physical interaction. Adding "-particles=<N>" to the command line will allow users to set # of particles for simulation. This example implements a uniform grid data structure using either atomic operations or a fast radix sort from the Thrust library

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemcpyToSymbol , cudaMemset , cudaFree , cudaGraphicsMapResources , cudaGLInit , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaDeviceSynchronize , cudaInit , cudaMalloc , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaMemcpy
Key Concepts	Graphics Interop , Data Parallel Algorithms , Physically-Based Simulation , Performance Strategies

Supported OSes	Linux , Windows
Whitepaper	particles.pdf

radixSortThrust - CUDA Radix Sort (Thrust Library)

This sample demonstrates a very fast and efficient parallel radix sort uses Thrust library. The included RadixSort class can sort either key-value pairs (with float or unsigned integer keys) or keys only.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaEventRecord , cudaEventCreate , cudaEventElapsedTime , cudaEventSynchronize , cudaEventDestroy , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Data-Parallel Algorithms , Performance Strategies
Supported OSes	Linux , Windows
Whitepaper	readme.txt

reduction - CUDA Parallel Reduction

A parallel sum reduction that computes the sum of a large arrays of values. This sample demonstrates several important optimization strategies for Data-Parallel Algorithms like reduction using shared memory, `__shfl_down_sync`, `__reduce_add_sync` and `cooperative_groups reduce`.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	C++11
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaDeviceSynchronize , cudaMalloc , cudaSetDevice , cudaMemcpy , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Data-Parallel Algorithms , Performance Strategies
Supported OSes	Linux , Windows

reductionMultiBlockCG - Reduction using MultiBlock Cooperative Groups

This sample demonstrates single pass reduction using Multi Block Cooperative Groups. This sample requires devices with compute capability 6.0 or higher having compute preemption.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	MBCG , CPP11
Supported SM Architecture	SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaOccupancyMaxActiveBlocksPerMultiprocessor , cudaOccupancyMaxPotentialBlockSize , cudaDeviceSynchronize , cudaSetDevice , cudaMalloc , cudaLaunchCooperativeKernel , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Cooperative Groups , MultiBlock Cooperative Groups
Supported OSes	Linux , Windows

scalarProd - Scalar Product

This sample calculates scalar products of a given set of input vector pairs.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaFree , cudaDeviceSynchronize , cudaMemcpy
Key Concepts	Linear Algebra
Supported OSes	Linux , Windows

scan - CUDA Parallel Prefix Sum (Scan)

This example demonstrates an efficient CUDA implementation of parallel prefix sum, also known as "scan". Given an array of numbers, scan computes a new array in which each element is the sum of all the elements before it in the input array.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaFree , cudaDeviceSynchronize , cudaMemcpy
Key Concepts	Data-Parallel Algorithms , Performance Strategies

Supported OSes [Linux](#), [Windows](#)

segmentationTreeThrust - CUDA Segmentation Tree Thrust Library

This sample demonstrates an approach to the image segmentation trees construction. This method is based on Boruvka's MST algorithm.

Supported SM Architecture [SM 3.5](#), [SM 3.7](#), [SM 5.0](#), [SM 5.2](#), [SM 5.3](#), [SM 6.0](#), [SM 6.1](#), [SM 6.2](#), [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

CUDA API [cudaMemset](#), [cudaEventRecord](#), [cudaEventCreate](#), [cudaEventElapsedTime](#), [cudaEventSynchronize](#), [cudaMemGetInfo](#), [cudaMemcpy](#)

Key Concepts [Data-Parallel Algorithms](#), [Performance Strategies](#)

Supported OSes [Linux](#), [Windows](#)

shfl_scan - CUDA Parallel Prefix Sum with Shuffle Ininsics (SHFL_Scan)

This example demonstrates how to use the shuffle intrinsic `__shfl_up_sync` to perform a scan operation across a thread block.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies [C++11](#)

Supported SM Architecture [SM 3.5](#), [SM 3.7](#), [SM 5.0](#), [SM 5.2](#), [SM 5.3](#), [SM 6.0](#), [SM 6.1](#), [SM 6.2](#), [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

CUDA API [cudaMemset](#), [cudaFree](#), [cudaEventRecord](#), [cudaMallocHost](#), [cudaEventCreate](#), [cudaEventElapsedTime](#), [cudaEventSynchronize](#), [cudaFreeHost](#), [cudaMalloc](#), [cudaMemcpy](#), [cudaGetDeviceProperties](#), [cudaGetDevice](#)

Key Concepts [Data-Parallel Algorithms](#), [Performance Strategies](#)

Supported OSes [Linux](#), [Windows](#)

sortingNetworks - CUDA Sorting Networks

This sample implements bitonic sort and odd-even merge sort (also known as Batcher's sort), algorithms belonging to the class of sorting networks. While generally subefficient, for large

sequences compared to algorithms with better asymptotic algorithmic complexity (i.e. merge sort or radix sort), this may be the preferred algorithms of choice for sorting batches of short-sized to mid-sized (key, value) array pairs. Refer to an excellent tutorial by H. W. Lang <http://www.iti.fh-flensburg.de/lang/algorithmen/sortieren/networks/indexen.htm>

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaFree , cudaDeviceSynchronize , cudaMemcpy
Key Concepts	Data-Parallel Algorithms
Supported OSes	Linux , Windows

streamOrderedAllocation - stream Ordered Allocation

This sample demonstrates stream ordered memory allocation on a GPU using `cudaMallocAsync` and `cudaMemPool` family of APIs.

Supported SM Architecture	SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaStreamDestroy , cudaEventRecord , cudaEventCreate , cudaEventElapsedTime , cudaEventSynchronize , cudaDeviceGetAttribute , cudaMallocAsync , cudaStreamCreateWithFlags , cudaDeviceGetDefaultMemPool , cudaSetDevice , cudaMemPoolSetAttribute , cudaStreamSynchronize , cudaMemcpyAsync , cudaFreeAsync
Key Concepts	Performance Strategies
Supported OSes	Linux , Windows

streamOrderedAllocationIPC - stream Ordered Allocation IPC Pools

This sample demonstrates IPC pools of stream ordered memory allocated using `cudaMallocAsync` and `cudaMemPool` family of APIs.

Supported SM Architecture	SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuDeviceGetAttribute , cuDeviceGet , cudaDeviceEnablePeerAccess , cudaMemPoolImportPointer , cudaOccupancyMaxActiveBlocksPerMultiprocessor , cudaMemPoolDestroy , cudaMallocAsync , cudaStreamCreateWithFlags , cudaDeviceCanAccessPeer , cudaMemcpyAsync , cudaStreamDestroy , cudaSetDevice , cudaGetDeviceProperties , cudaMemPoolSetAccess , cudaGetDeviceCount ,

[cudaDeviceGetAttribute](#), [cudaMemPoolExportPointer](#),
[cudaMemPoolImportFromShareableHandle](#), [cudaMemPoolCreate](#),
[cudaGetLastError](#), [cudaStreamSynchronize](#), [cudaMemPoolGetAccess](#),
[cudaMemPoolExportToShareableHandle](#), [cudaFreeAsync](#)

Key Concepts [Performance Strategies](#)

Supported OSes [Linux](#)

streamOrderedAllocationP2P - stream Ordered Allocation Peer-to-Peer access

This sample demonstrates peer-to-peer access of stream ordered memory allocated using `cudaMallocAsync` and `cudaMemPool` family of APIs.

Supported SM Architecture [SM 6.0](#), [SM 6.1](#), [SM 6.2](#), [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

CUDA API [cudaStreamWaitEvent](#), [cudaStreamDestroy](#), [cudaMemPoolSetAccess](#),
[cudaEventRecord](#), [cudaEventCreate](#), [cudaGetDeviceCount](#),
[cudaMallocAsync](#), [cudaDeviceGetAttribute](#), [cudaStreamCreateWithFlags](#),
[cudaDeviceCanAccessPeer](#), [cudaDeviceGetDefaultMemPool](#), [cudaSetDevice](#),
[cudaStreamSynchronize](#), [cudaMemcpyAsync](#), [cudaFreeAsync](#)

Key Concepts [Performance Strategies](#)

Supported OSes [Linux](#), [Windows](#)

threadFenceReduction

This sample shows how to perform a reduction operation on an array of values using the thread Fence intrinsic to produce a single value in a single kernel (as opposed to two or more kernel calls as shown in the "reduction" CUDA Sample). Single-pass reduction requires global atomic instructions (Compute Capability 2.0 or later) and the `_threadfence()` intrinsic (CUDA 2.2 or later).

Supported SM Architecture [SM 3.5](#), [SM 3.7](#), [SM 5.0](#), [SM 5.2](#), [SM 5.3](#), [SM 6.0](#), [SM 6.1](#), [SM 6.2](#), [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

CUDA API [cudaFree](#), [cudaDeviceSynchronize](#), [cudaMalloc](#), [cudaMemcpy](#),
[cudaGetDeviceProperties](#)

Key Concepts [Cooperative Groups](#), [Data-Parallel Algorithms](#), [Performance Strategies](#)

Supported OSes [Linux](#), [Windows](#)

threadMigration - CUDA Context Thread Management

Simple program illustrating how to the CUDA Context Management API and uses the new CUDA 4.0 parameter passing and CUDA launch API. CUDA contexts can be created separately and attached independently to different threads.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuModuleGetFunction , cuModuleLoadData , cuDeviceGetCount , cuCtxCreate , cuCtxPopCurrent , cuDeviceGetName , cuLaunchKernel , cuMemAlloc , cuMemcpyDtoH , cuModuleUnload , cuCtxPushCurrent , cuDeviceGet , cuMemFree , culnit , cuCtxDestroy , cuDeviceGetAttribute
Key Concepts	CUDA Driver API
Supported OSes	Linux , Windows

3.4. CUDA Features Reference

bf16TensorCoreGemm - bfloat16 Tensor Core GEMM

A CUDA sample demonstrating `__nv_bfloat16` (e8m7) GEMM computation using the Warp Matrix Multiply and Accumulate (WMMA) API introduced with CUDA 11 in Ampere chip family tensor cores for faster matrix operations. This sample also uses async copy provided by cuda pipeline interface for gmem to shmem async loads which improves kernel performance and reduces register pressure.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CPP11
Supported SM Architecture	SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaEventRecord , cudaEventCreate , cudaFuncSetAttribute , cudaEventElapsedTime , cudaEventSynchronize , cudaMalloc , cudaGetLastError , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties

Key Concepts	Matrix Multiply , WMMA , Tensor Cores
Supported OSes	Linux , Windows

binaryPartitionCG - Binary Partition Cooperative Groups

This sample is a simple code that illustrates binary partition cooperative groups and reduce within the thread block.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaMallocHost , cudaOccupancyMaxPotentialBlockSize , cudaMemsetAsync , cudaFreeHost , cudaMalloc , cudaStreamCreateWithFlags , cudaStreamSynchronize , cudaMemcpyAsync
Key Concepts	Cooperative Groups
Supported OSes	Linux , Windows

bindlessTexture - Bindless Texture

This example demonstrates use of `cudaSurfaceObject`, `cudaTextureObject`, and `MipMap` support in CUDA. A GPU with Compute Capability SM 3.0 is required to run the sample.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaGraphicsMapResources , cudaGetMipmappedArrayLevel , cudaGraphicsResourceGetMappedPointer , cudaArrayGetInfo , cudaMemcpy , cudaFreeMipmappedArray , cudaDestroySurfaceObject , cudaPitchedPtr , cudaMalloc , cudaGraphicsUnregisterResource , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaGetLastError , cudaFree , cudaFreeArray , cudaGraphicsGLRegisterBuffer , cudaExtent , cudaCreateSurfaceObject , cudaMallocMipmappedArray , cudaCreateTextureObject , cudaMallocArray , cudaGraphicsUnmapResources
Key Concepts	Graphics Interop , Texture
Supported OSes	Linux , Windows

cdpAdvancedQuicksort - Advanced Quicksort (CUDA Dynamic Parallelism)

This sample demonstrates an advanced quicksort implemented using CUDA Dynamic Parallelism. This sample requires devices with compute capability 3.5 or higher.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CDP
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaEventRecord , cudaEventCreate , cudaEventElapsedTime , cudaDeviceSynchronize , cudaPeekAtLastError , cudaMalloc , cudaStreamCreateWithFlags , cudaGetLastError , cudaMemcpyAsync , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Cooperative Groups , CUDA Dynamic Parallelism
Supported OSes	Linux , Windows

cdpBezierTessellation - Bezier Line Tessellation (CUDA Dynamic Parallelism)

This sample demonstrates bezier tessellation of lines implemented using CUDA Dynamic Parallelism. This sample requires devices with compute capability 3.5 or higher.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CDP
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGetDeviceCount , cudaMalloc , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	CUDA Dynamic Parallelism
Supported OSes	Linux , Windows

cdpQuadtree - Quad Tree (CUDA Dynamic Parallelism)

This sample demonstrates Quad Trees implemented using CUDA Dynamic Parallelism. This sample requires devices with compute capability 3.5 or higher.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CDP
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDeviceSetLimit , cudaFree , cudaMalloc , cudaGetLastError , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Cooperative Groups , CUDA Dynamic Parallelism
Supported OSes	Linux , Windows

cdpSimplePrint - Simple Print (CUDA Dynamic Parallelism)

This sample demonstrates simple printf implemented using CUDA Dynamic Parallelism. This sample requires devices with compute capability 3.5 or higher.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CDP
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDeviceSetLimit , cudaGetLastError , cudaDeviceSynchronize , cudaGetDeviceProperties
Key Concepts	CUDA Dynamic Parallelism
Supported OSes	Linux , Windows

cdpSimpleQuicksort - Simple Quicksort (CUDA Dynamic Parallelism)

This sample demonstrates simple quicksort implemented using CUDA Dynamic Parallelism. This sample requires devices with compute capability 3.5 or higher.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CDP
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDeviceSetLimit , cudaFree , cudaStreamDestroy , cudaDeviceSynchronize , cudaMalloc , cudaStreamCreateWithFlags , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	CUDA Dynamic Parallelism
Supported OSes	Linux , Windows

cudaCompressibleMemory - CUDA Compressible Memory

This sample demonstrates the compressible memory allocation using cuMemMap API.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuMemRelease , cuCtxGetDevice , cuMemGetAllocationPropertiesFromHandle , cuMemSetAccess , cuMemMap , cuMemCreate , cuMemAddressFree , cuMemGetAllocationGranularity , cuMemUnmap , cuMemAddressReserve , cuDeviceGetAttribute , cudaEventRecord , cudaEventCreate , cudaOccupancyMaxPotentialBlockSize , cudaEventElapsedTime , cudaEventSynchronize , cudaMemcpy
Key Concepts	CUDA Driver API , Compressible Memory , MMAP
Supported OSes	Linux , Windows

cudaTensorCoreGemm - CUDA Tensor Core GEMM

CUDA sample demonstrating a GEMM computation using the Warp Matrix Multiply and Accumulate (WMMA) API introduced in CUDA 9. This sample demonstrates the use of the new CUDA WMMA API employing the Tensor Cores introduced in the Volta chip family for faster matrix operations. In addition to that, it demonstrates the use of the new CUDA function attribute `cudaFuncAttributeMaxDynamicSharedMemorySize` that allows the application to reserve an extended amount of shared memory than it is available by default.

Supported SM Architecture	SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaEventRecord , cudaEventCreate , cudaFuncSetAttribute , cudaEventElapsedTime , cudaEventSynchronize , cudaMalloc , cudaGetLastError , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Matrix Multiply , WMMA , Tensor Cores
Supported OSes	Linux , Windows

dmmaTensorCoreGemm - Double Precision Tensor Core GEMM

CUDA sample demonstrates double precision GEMM computation using the Double precision Warp Matrix Multiply and Accumulate (WMMA) API introduced with CUDA 11 in Ampere chip family tensor cores for faster matrix operations. This sample also uses async copy provided by cuda pipeline interface for `gmem` to `shmem` async loads which improves kernel performance and reduces register pressure. Further, this sample also demonstrates how to use cooperative groups async copy interface over a group for performing `gmem` to `shmem` async loads.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CPP11
Supported SM Architecture	SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaEventRecord , cudaEventCreate , cudaFuncSetAttribute , cudaEventElapsedTime , cudaEventSynchronize ,

[cudaMalloc](#), [cudaGetLastError](#), [cudaMemcpy](#), [cudaGetErrorString](#),
[cudaGetDeviceProperties](#)

Key Concepts [Matrix Multiply](#), [WMMA](#), [Tensor Cores](#)

Supported OSes [Linux](#), [Windows](#)

globalToShmemAsyncCopy - Global Memory to Shared Memory Async Copy

This sample implements matrix multiplication which uses asynchronous copy of data from global to shared memory when on compute capability 8.0 or higher. Also demonstrates arrive-wait barrier for synchronization.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies [C++11](#)

Supported SM Architecture [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

CUDA API [cudaFree](#), [cudaEventRecord](#), [cudaMallocHost](#), [cudaEventCreate](#),
[cudaMemsetAsync](#), [cudaEventElapsedTime](#), [cudaEventSynchronize](#),
[cudaDeviceGetAttribute](#), [cudaFreeHost](#), [cudaMalloc](#),
[cudaStreamCreateWithFlags](#), [cudaEventDestroy](#), [cudaStreamSynchronize](#),
[cudaMemcpyAsync](#)

Key Concepts [CUDA Runtime API](#), [Linear Algebra](#), [C++11 CUDA](#)

Supported OSes [Linux](#), [Windows](#)

graphMemoryFootprint - Graph Memory Footprint

This sample demonstrates how graph memory nodes re-use virtual addresses and physical memory.

Supported SM Architecture [SM 6.0](#), [SM 6.1](#), [SM 6.2](#), [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

CUDA API [cudaGraphAddMemFreeNode](#), [cudaStreamDestroy](#), [cudaFree](#),
[cudaGraphExecDestroy](#), [cudaGraphInstantiate](#), [cudaDeviceGetAttribute](#),
[cudaDriverGetVersion](#), [cudaGraphCreate](#), [cudaGraphAddKernelNode](#),
[cudaGraphAddMemAllocNode](#), [cudaStreamCreateWithFlags](#),
[cudaDeviceGraphMemTrim](#), [cudaStreamSynchronize](#),

	cudaDeviceGetGraphMemAttribute , cudaGraphDestroy , cudaGetDeviceProperties , cudaGraphLaunch
Key Concepts	CUDA Runtime API , Performance Strategies , CUDA Graphs
Supported OSes	Linux , Windows

graphMemoryNodes - Graph Memory Nodes

A demonstration of memory allocations and frees within CUDA graphs using Graph APIs and Stream Capture APIs.

Supported SM Architecture	SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMallocAsync , cudaStreamCreateWithFlags , cudaMemcpy , cudaMemcpyAsync , cudaStreamDestroy , cudaMallocManaged , cudaEventCreate , cudaDriverGetVersion , cudaGraphCreate , cudaGraphAddMemAllocNode , cudaMalloc , cudaEventDestroy , cudaStreamEndCapture , cudaGraphExecDestroy , cudaStreamBeginCapture , cudaDeviceGetAttribute , cudaStreamSynchronize , cudaGraphDestroy , cudaGraphLaunch , cudaGraphAddMemFreeNode , cudaStreamWaitEvent , cudaFree , cudaEventRecord , cudaGraphInstantiate , cudaGraphAddKernelNode , cudaFreeAsync
Key Concepts	CUDA Graphs , Stream Capture
Supported OSes	Linux , Windows

immaTensorCoreGemm - Tensor Core GEMM Integer MMA

CUDA sample demonstrating a integer GEMM computation using the Warp Matrix Multiply and Accumulate (WMMA) API for integer introduced in CUDA 10. This sample demonstrates the use of the CUDA WMMA API employing the Tensor Cores introduced in the Volta chip family for faster matrix operations. In addition to that, it demonstrates the use of the new CUDA function attribute `cudaFuncAttributeMaxDynamicSharedMemorySize` that allows the application to reserve an extended amount of shared memory than it is available by default.

Supported SM Architecture	SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaEventRecord , cudaEventCreate , cudaFuncSetAttribute , cudaEventElapsedTime , cudaEventSynchronize , cudaMalloc , cudaGetLastError , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Matrix Multiply , WMMA , Tensor Cores

Supported OSes [Linux](#), [Windows](#)

jacobiCudaGraphs - Jacobi CUDA Graphs

Demonstrates Instantiated CUDA Graph Update with Jacobi Iterative Method using `cudaGraphExecKernelNodeSetParams()` and `cudaGraphExecUpdate()` approach.

Supported SM Architecture [SM 3.5](#), [SM 3.7](#), [SM 5.0](#), [SM 5.2](#), [SM 5.3](#), [SM 6.0](#), [SM 6.1](#), [SM 6.2](#), [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

CUDA API [cudaGraphAddMemsetNode](#), [cudaStreamCreateWithFlags](#), [cudaMemcpyAsync](#), [cudaMallocHost](#), [cudaPitchedPtr](#), [cudaGraphCreate](#), [cudaMalloc](#), [cudaPos](#), [cudaGraphAddMemcpyNode](#), [cudaStreamEndCapture](#), [cudaGraphExecDestroy](#), [cudaStreamBeginCapture](#), [cudaGraphExecKernelNodeSetParams](#), [cudaStreamSynchronize](#), [cudaGraphLaunch](#), [cudaFree](#), [cudaGraphInstantiate](#), [cudaExtent](#), [cudaMemsetAsync](#), [cudaFreeHost](#), [cudaGraphAddKernelNode](#), [cudaGraphExecUpdate](#)

Key Concepts [CUDA Graphs](#), [Stream Capture](#), [Instantiated CUDA Graph Update](#), [Cooperative Groups](#)

Supported OSes [Linux](#), [Windows](#)

memMapIPCDrv - Memmap IPC Driver API

This CUDA Driver API sample is a very basic sample that demonstrates Inter Process Communication using `cuMemMap` APIs with one process per GPU for computation. Requires Compute Capability 3.0 or higher and a Linux Operating System, or a Windows Operating System

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies [IPC](#)

Supported SM Architecture [SM 3.5](#), [SM 3.7](#), [SM 5.0](#), [SM 5.2](#), [SM 5.3](#), [SM 6.0](#), [SM 6.1](#), [SM 6.2](#), [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

CUDA API [cuCtxSetCurrent](#), [cuMemSetAccess](#), [cuMemcpyDtoHAsync](#), [cuStreamDestroy](#), [cuInit](#), [cuMemAddressReserve](#), [cuCtxDestroy](#), [cuModuleGetFunction](#), [cuModuleLoad](#), [cuStreamCreate](#), [cuCtxCreate](#), [cuMemExportToShareableHandle](#), [cuMemAddressFree](#), [cuMemGetAllocationGranularity](#), [cuModuleLoadDataEx](#), [cuDeviceGet](#), [cuMemUnmap](#), [cuDeviceGetAttribute](#).

	cuMemRelease , cuCtxEnablePeerAccess , cuMemMap , cuMemImportFromShareableHandle , cuMemCreate , cuStreamSynchronize , cuDeviceCanAccessPeer , cuDeviceGetCount , cuLaunchKernel , cuOccupancyMaxActiveBlocksPerMultiprocessor
Key Concepts	CUDA Driver API , cuMemMap IPC , MMAP
Supported OSes	Linux , Windows

newdelete - NewDelete

This sample demonstrates dynamic global memory allocation through device C++ new and delete operators and virtual function declarations available with CUDA 4.0.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDeviceSetLimit , cudaFree , cudaDeviceSynchronize , cudaMalloc , cudaMemcpy
Key Concepts	Device Memory Allocation , C++ Templates
Supported OSes	Linux , Windows

ptxjit - PTX Just-in-Time compilation

This sample uses the Driver API to just-in-time compile (JIT) a Kernel from PTX code. Additionally, this sample demonstrates the seamless interoperability capability of the CUDA Runtime and CUDA Driver API calls. For CUDA 5.5, this sample shows how to use cuLink* functions to link PTX assembly using the CUDA driver at runtime.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuModuleGetFunction , cuLinkAddData , cuModuleLoadData , cuLaunchKernel , cuModuleUnload , cuLinkComplete , cuLinkCreate , cuLinkDestroy , cudaDriverGetVersion , cudaFree , cudaMalloc , cudaMemcpy
Key Concepts	CUDA Driver API
Supported OSes	Linux , Windows

simpleCudaGraphs - Simple CUDA Graphs

A demonstration of CUDA Graphs creation, instantiation and launch using Graphs APIs and Stream Capture APIs.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
----------------------------------	---

CUDA API	cudaGraphAddMemsetNode , cudaGraphsUsingStreamCapture , cudaMemcpyAsync , cudaGraphGetNodes , cudaStreamDestroy , cudaMallocHost , cudaGraphClone , cudaEventCreate , cudaPitchedPtr , cudaGraphCreate , cudaMalloc , cudaPos , cudaGraphAddMemcpyNode , cudaStreamEndCapture , cudaGraphExecDestroy , cudaStreamBeginCapture , cudaGraphAddHostNode , cudaGraphsManual , cudaStreamSynchronize , cudaGraphDestroy , cudaGraphLaunch , cudaStreamWaitEvent , cudaFree , cudaEventRecord , cudaStreamCreate , cudaGraphInstantiate , cudaLaunchHostFunc , cudaExtent , cudaMemsetAsync , cudaFreeHost , cudaGraphAddKernelNode
Key Concepts	CUDA Graphs , Stream Capture
Supported OSes	Linux , Windows

StreamPriorities - Stream Priorities

This sample demonstrates basic use of stream priorities.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	Stream-Priorities
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDeviceGetStreamPriorityRange , cudaEventRecord , cudaEventCreate , cudaEventElapsedTime , cudaEventSynchronize , cudaMalloc , cudaStreamCreateWithPriority , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	CUDA Streams and Events
Supported OSes	Linux

tf32TensorCoreGemm - tf32 Tensor Core GEMM

A CUDA sample demonstrating tf32 (e8m10) GEMM computation using the Warp Matrix Multiply and Accumulate (WMMA) API introduced with CUDA 11 in Ampere chip family tensor cores for faster matrix operations. This sample also uses async copy provided by cuda pipeline interface for gmem to shmem async loads which improves kernel performance and reduces register pressure.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be

installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CPP11
Supported SM Architecture	SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaEventRecord , cudaEventCreate , cudaFuncSetAttribute , cudaEventElapsedTime , cudaEventSynchronize , cudaMalloc , cudaGetLastError , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Matrix Multiply , WMMA , Tensor Cores
Supported OSes	Linux , Windows

warpAggregatedAtomicsCG - Warp Aggregated Atomics using Cooperative Groups

This sample demonstrates how using Cooperative Groups (CG) to perform warp aggregated atomics to single and multiple counters, a useful technique to improve performance when many threads atomically add to a single or multiple counters.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaDeviceGetAttribute , cudaMalloc , cudaMemcpy
Key Concepts	Cooperative Groups , Atomic Intrinsic
Supported OSes	Linux , Windows

3.5. CUDA Libraries Reference

batchCUBLAS

A CUDA Sample that demonstrates how using batched CUBLAS API calls to improve overall performance.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUBLAS
---------------------	------------------------

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuRand , cuEqual , cudaFree , cudaStreamCreate , cudaDeviceSynchronize , cudaMalloc , cudaGetLastError , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Linear Algebra , CUBLAS Library
Supported OSes	Linux , Windows

batchedLabelMarkersAndLabelCompressionNPP - Batched Label Markers And Label Compression NPP

An NPP CUDA Sample that demonstrates how to use the NPP label markers generation and label compression functions based on a Union Find (UF) algorithm including both single image and batched image versions.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NPP
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaMallocHost , cudaMallocPitch , cudaDeviceGetAttribute , cudaFreeHost , cudaDriverGetVersion , cudaMalloc , cudaStreamGetFlags , cudaRuntimeGetVersion , cudaStreamSynchronize , cudaMemcpyAsync , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Performance Strategies , Image Processing , NPP Library , Using NPP Batch Functions
Supported OSes	Linux , Windows

- Box Filter with NPP

A NPP CUDA Sample that demonstrates how to use NPP FilterBox function to perform a Box Filter.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	FreelImage , NPP
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDriverGetVersion , cudaRuntimeGetVersion
Key Concepts	Performance Strategies , Image Processing , NPP Library
Supported OSes	Linux , Windows

cannyEdgeDetectorNPP - Canny Edge Detector NPP

An NPP CUDA Sample that demonstrates the recommended parameters to use with the `nppiFilterCannyBorder_8u_C1R` Canny Edge Detection image filter function. This function expects a single channel 8-bit grayscale input image. You can generate a grayscale image from a color image by first calling `nppiColorToGray()` or `nppiRGBToGray()`. The Canny Edge Detection function combines and improves on the techniques required to produce an edge detection image using multiple steps.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	FreelImage , NPP
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGetDeviceCount , cudaDriverGetVersion , cudaMalloc , cudaDeviceInit , cudaSetDevice , cudaRuntimeGetVersion , cudaGetDeviceProperties
Key Concepts	Performance Strategies , Image Processing , NPP Library
Supported OSes	Linux , Windows

conjugateGradient - ConjugateGradient

This sample implements a conjugate gradient solver on GPU using CUBLAS and CUSPARSE library.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUBLAS , CUSPARSE
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaDeviceSynchronize , cudaMalloc , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Linear Algebra , CUBLAS Library , CUSPARSE Library
Supported OSes	Linux , Windows

conjugateGradientCudaGraphs - Conjugate Gradient using Cuda Graphs

This sample implements a conjugate gradient solver on GPU using CUBLAS and CUSPARSE library calls captured and called using CUDA Graph APIs.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUBLAS , CUSPARSE
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaStreamDestroy , cudaFree , cudaGraphExecDestroy , cudaMallocHost , cudaStreamCreate , cudaGraphInstantiate , cudaOccupancyMaxPotentialBlockSize , cudaStreamBeginCapture , cudaMemsetAsync , cudaStreamEndCapture , cudaFreeHost , cudaMalloc , cudaStreamSynchronize , cudaMemcpyAsync , cudaGraphDestroy , cudaGetDeviceProperties , cudaGraphLaunch
Key Concepts	Linear Algebra , CUBLAS Library , CUSPARSE Library
Supported OSes	Linux , Windows

conjugateGradientMultiBlockCG - conjugateGradient using MultiBlock Cooperative Groups

This sample implements a conjugate gradient solver on GPU using Multi Block Cooperative Groups, also uses Unified Memory.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be

installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	UVM , MBCG
Supported SM Architecture	SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaEventRecord , cudaMallocManaged , cudaOccupancyMaxActiveBlocksPerMultiprocessor , cudaEventCreate , cudaEventElapsedTime , cudaDeviceSynchronize , cudaEventDestroy , cudaLaunchCooperativeKernel , cudaGetDeviceProperties
Key Concepts	Unified Memory , Linear Algebra , Cooperative Groups , MultiBlock Cooperative Groups , CUBLAS Library , CUSPARSE Library
Supported OSes	Linux , Windows

conjugateGradientMultiDeviceCG - conjugateGradient using MultiDevice Cooperative Groups

This sample implements a conjugate gradient solver on multiple GPUs using Multi Device Cooperative Groups, also uses Unified Memory optimized using prefetching and usage hints.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	UVM , MDCG , CPP11
Supported SM Architecture	SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDeviceEnablePeerAccess , cudaMemset , cudaFree , cudaMallocManaged , cudaMemPrefetchAsync , cudaHostAlloc , cudaOccupancyMaxActiveBlocksPerMultiprocessor , cudaStreamCreate , cudaGetDeviceCount , cudaFreeHost , cudaSetDevice , cudaDeviceCanAccessPeer , cudaLaunchCooperativeKernel , cudaStreamSynchronize , cudaMemAdvise , cudaGetDeviceProperties
Key Concepts	Unified Memory , Linear Algebra , Cooperative Groups , MultiDevice Cooperative Groups , CUBLAS Library , CUSPARSE Library
Supported OSes	Linux , Windows

conjugateGradientPrecond - Preconditioned Conjugate Gradient

This sample implements a preconditioned conjugate gradient solver on GPU using CUBLAS and CUSPARSE library.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUBLAS , CUSPARSE
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaMalloc , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Linear Algebra , CUBLAS Library , CUSPARSE Library
Supported OSes	Linux , Windows

conjugateGradientUM - ConjugateGradientUM

This sample implements a conjugate gradient solver on GPU using CUBLAS and CUSPARSE library, using Unified Memory

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	UVM , CUBLAS , CUSPARSE
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaMallocManaged , cudaDeviceSynchronize , cudaMalloc , cudaGetDeviceProperties
Key Concepts	Unified Memory , Linear Algebra , CUBLAS Library , CUSPARSE Library
Supported OSes	Linux , Windows

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NVSCI , NvMedia
Supported SM Architecture	SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuDeviceGetUuid , cudaGetMipmappedArrayLevel , cudaImportNvSciImage , cudaImportExternalSemaphore , cudaStreamCreateWithFlags , cudaExternalMemoryGetMappedMipmappedArray , cudaDestroyExternalMemory , cudaMemcpyAsync , cudaStreamDestroy , cudaSignalExternalSemaphoresAsync , cudaDeviceGetNvSciSyncAttributes , cudaFreeMipmappedArray , cudaMallocHost , cudaDestroySurfaceObject , cudaImportExternalMemory , cudaMalloc , cudaSetDevice , cudaDestroyExternalSemaphore , cudaCreateChannelDesc , cudaStreamSynchronize , cudaFree , cudaFreeArray , cudaCreateSurfaceObject , cudaFreeHost , cudaMallocArray , cudaWaitExternalSemaphoresAsync , cudaImportNvSciSync
Key Concepts	CUDA NvSci Interop , Data Parallel Algorithms , Image Processing
Supported OSes	Linux

- cuDLA Error Reporting

This sample demonstrates how DLA errors can be detected via CUDA.

Supported SM Architecture	SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaStreamDestroy , cudaFree , cudaStreamSynchronize , cudaMemsetAsync , cudaMalloc , cudaStreamCreateWithFlags , cudaSetDevice , cudaGetErrorName , cudaMemcpyAsync
Key Concepts	cuDLA , Data Parallel Algorithms , Image Processing
Supported OSes	Linux

- cuDLA Hybrid Mode

This sample demonstrates cuDLA hybrid mode wherein DLA can be programmed using CUDA.

Supported SM Architecture	SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
----------------------------------	--

CUDA API	cudaStreamDestroy , cudaFree , cudaStreamSynchronize , cudaMemsetAsync , cudaMalloc , cudaStreamCreateWithFlags , cudaSetDevice , cudaGetErrorName , cudaMemcpyAsync
Key Concepts	cuDLA , Data Parallel Algorithms , Image Processing
Supported OSes	Linux

- cuDLA Standalone Mode

This sample demonstrates cuDLA standalone mode wherein DLA can be programmed without using CUDA.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NVSCI
Supported SM Architecture	SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
Key Concepts	cuDLA , Data Parallel Algorithms , Image Processing
Supported OSes	Linux

- cuSolverDn Linear Solver

A CUDA Sample that demonstrates cuSolverDN's LU, QR and Cholesky factorization.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUSOLVER , CUBLAS , CUSPARSE
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuDoubleComplex , cuComplex , cudaMemset , cudaFree , cudaStreamDestroy , cudaStreamCreate , cudaDeviceSynchronize , cudaMalloc , cudaMemcpy
Key Concepts	Linear Algebra , CUSOLVER Library
Supported OSes	Linux , Windows

cuSolverRf - cuSolverRf Refactorization

A CUDA Sample that demonstrates cuSolver's refactorization library - CUSOLVERRF.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUSOLVER , CUBLAS , CUSPARSE
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuDoubleComplex , cuComplex , cuGet_cudaStreamDestroy , cudaFree , cudaStreamCreate , cudaDeviceSynchronize , cudaMalloc , cudaMemcpy
Key Concepts	Linear Algebra , CUSOLVER Library
Supported OSes	Linux , Windows

- cuSolverSp Linear Solver

A CUDA Sample that demonstrates cuSolverSP's LU, QR and Cholesky factorization.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUSOLVER , CUSPARSE
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuDoubleComplex , cuComplex , cudaStreamDestroy , cudaFree , cudaStreamCreate , cudaDeviceSynchronize , cudaMalloc , cudaMemcpyAsync
Key Concepts	Linear Algebra , CUSOLVER Library
Supported OSes	Linux , Windows

cuSolverSp_LowlevelCholesky - cuSolverSp LowlevelCholesky Solver

A CUDA Sample that demonstrates Cholesky factorization using cuSolverSP's low level APIs.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be

installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUSOLVER , CUSPARSE
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuDoubleComplex , cuComplex , cudaStreamDestroy , cudaFree , cudaStreamCreate , cudaMalloc , cudaMemcpy
Key Concepts	Linear Algebra , CUSOLVER Library
Supported OSes	Linux , Windows

cuSolverSp_LowlevelQR - cuSolverSp Lowlevel QR Solver

A CUDA Sample that demonstrates QR factorization using cuSolverSP's low level APIs.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUSOLVER , CUSPARSE
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuDoubleComplex , cuComplex , cuGet , cudaStreamDestroy , cudaFree , cudaStreamCreate , cudaMalloc , cudaMemcpy
Key Concepts	Linear Algebra , CUSOLVER Library
Supported OSes	Linux , Windows

FilterBorderControlNPP - Filter Border Control NPP

This sample demonstrates how any border version of an NPP filtering function can be used in the most common mode, with border control enabled. Mentioned functions can be used to duplicate the results of the equivalent non-border version of the NPP functions. They can be also used for enabling and disabling border control on various source image edges depending on what portion of the source image is being used as input.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be

installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	FreelImage , NPP
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaDeviceReset , cudaGetDeviceCount , cudaDriverGetVersion , cudaDeviceInit , cudaSetDevice , cudaRuntimeGetVersion , cudaGetDeviceProperties
Key Concepts	Performance Strategies , Image Processing , NPP Library
Supported OSes	Linux , Windows

freelImageInteropNPP - FreelImage and NPP Interoperability

A simple CUDA Sample demonstrate how to use FreelImage library with NPP.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	FreelImage , NPP
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaGetDeviceCount , cudaDriverGetVersion , cudaDeviceInit , cudaSetDevice , cudaRuntimeGetVersion , cudaGetDeviceProperties
Key Concepts	Performance Strategies , Image Processing , NPP Library
Supported OSes	Linux , Windows

histEqualizationNPP - Histogram Equalization with NPP

This CUDA Sample demonstrates how to use NPP for histogram equalization for image data.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	FreelImage , NPP
---------------------	--

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGetDeviceCount , cudaDriverGetVersion , cudaMalloc , cudaDeviceInit , cudaSetDevice , cudaRuntimeGetVersion , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Image Processing , Performance Strategies , NPP Library
Supported OSes	Linux , Windows

lineOfSight - Line of Sight

This sample is an implementation of a simple line-of-sight algorithm: Given a height map and a ray originating at some observation point, it computes all the points along the ray that are visible from the observation point. The implementation is based on the Thrust library.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFreeArray , cudaDeviceSynchronize , cudaCreateTextureObject , cudaMallocArray , cudaCreateChannelDesc
Key Concepts	Thrust Library
Supported OSes	Linux , Windows

matrixMulCUBLAS - Matrix Multiplication (CUBLAS)

This sample implements matrix multiplication from Chapter 3 of the programming guide. To illustrate GPU performance for matrix multiply, this sample also shows how to use the new CUDA 4.0 interface for CUBLAS to demonstrate high-performance performance for matrix multiplication.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUBLAS
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaEventRecord , cudaEventCreate , cudaEventElapsedTime , cudaEventSynchronize , cudaMalloc , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	CUDA Runtime API , Performance Strategies , Linear Algebra , CUBLAS

Supported OSes [Linux](#), [Windows](#)

- MersenneTwisterGP11213

This sample demonstrates the Mersenne Twister random number generator GP11213 in cuRAND.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies [CURAND](#)

Supported SM Architecture [SM 3.5](#), [SM 3.7](#), [SM 5.0](#), [SM 5.2](#), [SM 5.3](#), [SM 6.0](#), [SM 6.1](#), [SM 6.2](#), [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

CUDA API [cudaStreamDestroy](#), [cudaFree](#), [cudaMallocHost](#), [cudaFreeHost](#), [cudaMalloc](#), [cudaStreamCreateWithFlags](#), [cudaStreamSynchronize](#), [cudaMemcpyAsync](#)

Key Concepts [CURAND Library](#)

Supported OSes [Linux](#), [Windows](#)

- NVJPEG simple

A CUDA Sample that demonstrates single and batched decoding of jpeg images using NVJPEG Library.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies [NVJPEG](#)

Supported SM Architecture [SM 3.5](#), [SM 3.7](#), [SM 5.0](#), [SM 5.2](#), [SM 5.3](#), [SM 6.0](#), [SM 6.1](#), [SM 6.2](#), [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

CUDA API [cudaStreamDestroy](#), [cudaFree](#), [cudaEventRecord](#), [cudaHostAlloc](#), [cudaEventCreate](#), [cudaEventElapsedTime](#), [cudaEventSynchronize](#), [cudaFreeHost](#), [cudaMalloc](#), [cudaStreamCreateWithFlags](#), [cudaStreamSynchronize](#), [cudaGetDeviceProperties](#)

Key Concepts [Image Decoding](#), [NVJPEG Library](#)

Supported OSes [Linux](#), [Windows](#)

- NVJPEG Encoder

A CUDA Sample that demonstrates single encoding of jpeg images using NVJPEG Library.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NVJPEG
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaEventRecord , cudaEventCreate , cudaEventElapsedTime , cudaDeviceSynchronize , cudaEventSynchronize , cudaMalloc , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Image Encoding , NVJPEG Library
Supported OSes	Linux , Windows

oceanFFT - CUDA FFT Ocean Simulation

This sample simulates an Ocean height field using CUFFT Library and renders the result using OpenGL.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL , CUFFT
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaUpdateHeightmapKernel , cudaGraphicsGLRegisterBuffer , cudaCalculateSlopeKernel , cudaGraphicsResourceGetMappedPointer , cudaMalloc , cudaGenerateSpectrumKernel , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Graphics Interop , Image Processing , CUFFT Library
Supported OSes	Linux , Windows

randomFog - Random Fog

This sample illustrates pseudo- and quasi- random numbers produced by CURAND.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL , CURAND
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaFree , cudaMemcpy , cudaGetErrorString
Key Concepts	3D Graphics , CURAND Library
Supported OSes	Linux , Windows

- Simple CUBLAS

Example of using CUBLAS API interface to perform GEMM operations.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUBLAS
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaFree
Key Concepts	Image Processing , CUBLAS Library
Supported OSes	Linux , Windows

simpleCUBLAS_LU - Simple CUBLAS LU

CUDA sample demonstrating cuBLAS API `cusolverDgtrfBatched()` for lower-upper (LU) decomposition of a matrix.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUBLAS
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaGetErrorEnum , cudaFree , cudaMalloc , cudaMemcpy

Key Concepts	CUBLAS Library , LU decomposition
Supported OSes	Linux , Windows

simpleCUBLASXT - Simple CUBLAS XT

Example of using CUBLAS-XT library which performs GEMM operations over Multiple GPUs.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUBLAS
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGetDeviceCount , cudaGetDeviceProperties
Key Concepts	CUBLAS-XT Library
Supported OSes	Linux , Windows

simpleCUFFT - Simple CUFFT

Example of using CUFFT. In this example, CUFFT is used to compute the 1D-convolution of some signal with some filter by transforming both into frequency domain, multiplying them together, and transforming the signal back to time domain. cuFFT plans are created using simple and advanced API functions.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUFFT
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaFree , cudaMemcpy
Key Concepts	Image Processing , CUFFT Library
Supported OSes	Linux , Windows

simpleCUFFT_2d_MGPU - SimpleCUFFT_2d_MGPU

Example of using CUFFT. In this example, CUFFT is used to compute the 2D-convolution of some signal with some filter by transforming both into frequency domain, multiplying them together, and transforming the signal back to time domain on Multiple GPU.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUFFT
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaXtFree , cudaGetDeviceCount , cudaDeviceSynchronize , cudaMalloc , cudaSetDevice , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Image Processing , CUFFT Library
Supported OSes	Linux , Windows

simpleCUFFT_callback - Simple CUFFT Callbacks

Example of using CUFFT. In this example, CUFFT is used to compute the 1D-convolution of some signal with some filter by transforming both into frequency domain, multiplying them together, and transforming the signal back to time domain. The difference between this example and the Simple CUFFT example is that the multiplication step is done by the CUFFT kernel with a user-supplied CUFFT callback routine, rather than by a separate kernel call.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	callback , CUFFT
Supported SM Architecture	SM 3.5 , SM 5.0 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaMalloc , cudaMemcpy , cudaMemcpyFromSymbol , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Image Processing , CUFFT Library
Supported OSes	Linux

simpleCUFFT_MGPU - Simple CUFFT_MGPU

Example of using CUFFT. In this example, CUFFT is used to compute the 1D-convolution of some signal with some filter by transforming both into frequency domain, multiplying them together, and transforming the signal back to time domain on Multiple GPU.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUFFT
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaGetDeviceCount , cudaDeviceSynchronize , cudaSetDevice , cudaXtFree , cudaGetDeviceProperties
Key Concepts	Image Processing , CUFFT Library
Supported OSes	Linux , Windows

watershedSegmentationNPP - Watershed Segmentation NPP

An NPP CUDA Sample that demonstrates how to use the NPP watershed segmentation function.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NPP
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaDeviceGetAttribute , cudaDriverGetVersion , cudaStreamGetFlags , cudaStreamSynchronize , cudaRuntimeGetVersion , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Performance Strategies , Image Processing , NPP Library
Supported OSes	Linux , Windows

3.6. Domain Specific Reference

bicubicTexture - Bicubic B-spline Interpolation

This sample demonstrates how to efficiently implement a Bicubic B-spline interpolation filter with CUDA texture.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaFreeArray , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaCreateTextureObject , cudaMalloc , cudaMallocArray , cudaCreateChannelDesc , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Graphics Interop , Image Processing
Supported OSes	Linux , Windows

bilateralFilter - Bilateral Filter

Bilateral filter is an edge-preserving non-linear smoothing filter that is implemented with CUDA with OpenGL rendering. It can be used in image recovery and denoising. Each pixel is weight by considering both the spatial distance and color distance between its neighbors. Reference: "C. Tomasi, R. Manduchi, Bilateral Filtering for Gray and Color Images, proceeding of the ICCV, 1998, <http://users.soe.ucsc.edu/~manduchi/Papers/ICCV98.pdf>"

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7

CUDA API	cudaMemcpyToSymbol , cudaFree , cudaGraphicsMapResources , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaMallocPitch , cudaCreateTextureObject , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaRuntimeGetVersion , cudaGetDeviceProperties
Key Concepts	Graphics Interop , Image Processing
Supported OSes	Linux , Windows

binomialOptions - Binomial Option Pricing

This sample evaluates fair call price for a given set of European options under binomial model.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemcpyToSymbol , cudaDeviceSynchronize , cudaMemcpyFromSymbol
Key Concepts	Computational Finance
Supported OSes	Linux , Windows
Whitepaper	binomialOptions.pdf

binomialOptions_nvrtc - Binomial Option Pricing with libNVRTC

This sample evaluates fair call price for a given set of European options under binomial model. This sample makes use of NVRTC for Runtime Compilation.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NVRTC
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuLaunchKernel , cuModuleGetGlobal , cuCtxSynchronize , cuMemcpyDtoH , cuModuleGetFunction , cuMemcpyHtoD , cudaBlockSize , cudaGridSize
Key Concepts	Computational Finance , Runtime Compilation
Supported OSes	Linux , Windows

BlackScholes - Black-Scholes Option Pricing

This sample evaluates fair call and put prices for a given set of European options by Black-Scholes formula.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaFree , cudaDeviceSynchronize , cudaMemcpy
Key Concepts	Computational Finance
Supported OSes	Linux , Windows
Whitepaper	BlackScholes.pdf

BlackScholes_nvrtc - Black-Scholes Option Pricing with libNVRTC

This sample evaluates fair call and put prices for a given set of European options by Black-Scholes formula, compiling the CUDA kernels involved at runtime using NVRTC.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NVRTC
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuModuleGetFunction , cuMemAlloc , cuLaunchKernel , cuCtxSynchronize , cuMemFree , cuMemcpyDtoH , cuMemcpyHtoD , cudaBlockSize , cudaGridSize
Key Concepts	Computational Finance , Runtime Compilation
Supported OSes	Linux , Windows

convolutionFFT2D - FFT-Based 2D Convolution

This sample demonstrates how 2D convolutions with very large kernel sizes can be efficiently implemented using FFT transformations.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be

installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CUFFT
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaCreateTextureObject , cudaMalloc , cudaMemcpy
Key Concepts	Image Processing , CUFFT Library
Supported OSes	Linux , Windows

dwtHaar1D - 1D Discrete Haar Wavelet Decomposition

Discrete Haar wavelet decomposition for 1D signals with a length which is a power of 2.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaFree , cudaMemcpy
Key Concepts	Image Processing , Video Compression
Supported OSes	Linux , Windows

dxtc - DirectX Texture Compressor (DXTC)

High Quality DXT Compression using CUDA. This example shows how to implement an existing computationally-intensive CPU compression algorithm in parallel on the GPU, and obtain an order of magnitude performance improvement.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaDeviceSynchronize , cudaMalloc , cudaMemcpy , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Cooperative Groups , Image Processing , Image Compression
Supported OSes	Linux , Windows
Whitepaper	cuda_dxtc.pdf

fastWalshTransform - Fast Walsh Transform

Naturally(Hadamard)-ordered Fast Walsh Transform for batching vectors of arbitrary eligible lengths that are power of two in size.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaDeviceSynchronize , cudaMalloc , cudaMemcpy
Key Concepts	Linear Algebra , Data-Parallel Algorithms , Video Compression
Supported OSes	Linux , Windows

FDTD3d - CUDA C 3D FDTD

This sample applies a finite differences time domain progression stencil on a 3D surface.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemcpyToSymbol , cudaFree , cudaEventRecord , cudaEventCreate , cudaFuncGetAttributes , cudaGetDeviceCount , cudaEventElapsedTime , cudaDeviceSynchronize , cudaMalloc , cudaEventDestroy , cudaSetDevice , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Performance Strategies
Supported OSes	Linux , Windows

fluidsD3D9 - Fluids (Direct3D Version)

An example of fluid simulation using CUDA and CUFFT, with Direct3D 9 rendering. A Direct3D Capable device is required.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	DirectX
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaGraphicsUnregisterResource , cudaFreeArray , cudaGraphicsResourceGetMappedPointer , cudaDestroyTextureObject , cudaMallocPitch , cudaCreateTextureObject , cudaMalloc , cudaMallocArray , cudaGetLastError , cudaGraphicsUnmapResources , cudaMemcpy , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Graphics Interop , CUFFT Library , Physically-Based Simulation
Supported OSes	Windows

fluidsGL - Fluids (OpenGL Version)

An example of fluid simulation using CUDA and CUFFT, with OpenGL rendering.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL , CUFFT
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaFreeArray , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaDestroyTextureObject , cudaMallocPitch , cudaCreateTextureObject , cudaMalloc , cudaMallocArray , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Graphics Interop , CUFFT Library , Physically-Based Simulation
Supported OSes	Linux , Windows
Whitepaper	fluidsGL.pdf

fluidsGLES - Fluids (OpenGLES Version)

An example of fluid simulation using CUDA and CUFFT, with OpenGLES rendering.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GLES , CUFFT
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaFreeArray , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaDestroyTextureObject , cudaMallocPitch , cudaCreateTextureObject , cudaMalloc , cudaMallocArray , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Graphics Interop , CUFFT Library , Physically-Based Simulation
Supported OSes	Linux

HSOpticalFlow - Optical Flow

Variational optical flow estimation example. Uses textures for image operations. Shows how simple PDE solver can be accelerated with CUDA.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMalloc , cudaMemset , cudaFree , cudaMemcpy
Key Concepts	Image Processing , Data Parallel Algorithms
Supported OSes	Linux , Windows
Whitepaper	OpticalFlow.pdf

Mandelbrot

This sample uses CUDA to compute and display the Mandelbrot or Julia sets interactively. It also illustrates the use of "double single" arithmetic to improve precision when zooming a long way into the pattern. This sample uses double precision. Thanks to Mark Granger of NewTek who submitted this code sample.!

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaGraphicsUnregisterResource , cudaGraphicsGLRegisterBuffer , cudaGLUnregisterBufferObject , cudaGraphicsResourceGetMappedPointer , cudaGLUnmapBufferObject , cudaDeviceSynchronize , cudaGLMapBufferObject , cudaMalloc , cudaGLRegisterBufferObject , cudaGraphicsUnmapResources , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Graphics Interop , Data Parallel Algorithms
Supported OSes	Linux , Windows

marchingCubes - Marching Cubes Isosurfaces

This sample extracts a geometric isosurface from a volume dataset using the marching cubes algorithm. It uses the scan (prefix sum) function from the Thrust library to perform stream compaction.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaGraphicsUnregisterResource , cudaGraphicsGLRegisterBuffer , cudaGLUnregisterBufferObject , cudaGraphicsResourceGetMappedPointer , cudaGLUnmapBufferObject , cudaDestroyTextureObject , cudaGLMapBufferObject , cudaCreateTextureObject , cudaMalloc , cudaCreateChannelDesc , cudaGLRegisterBufferObject , cudaGraphicsUnmapResources , cudaMemcpy
Key Concepts	OpenGL Graphics Interop , Vertex Buffers , 3D Graphics , Physically Based Simulation
Supported OSes	Linux , Windows

MonteCarloMultiGPU - Monte Carlo Option Pricing with Multi-GPU support

This sample evaluates fair call price for a given set of European options using the Monte Carlo approach, taking advantage of all CUDA-capable GPUs installed in the system. This sample use double precision hardware if a GTX 200 class GPU is present. The sample also takes advantage of CUDA 4.0 capability to supporting using a single CPU thread to control multiple GPUs

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	CURAND
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7

CUDA API	cudaMemset , cudaFree , cudaStreamDestroy , cudaEventRecord , cudaMallocHost , cudaStreamCreate , cudaEventCreate , cudaGetDeviceCount , cudaDeviceSynchronize , cudaEventSynchronize , cudaFreeHost , cudaMalloc , cudaEventDestroy , cudaSetDevice , cudaMemcpyAsync , cudaStreamSynchronize , cudaGetDeviceProperties
Key Concepts	Random Number Generator , Computational Finance , CURAND Library
Supported OSes	Linux , Windows
Whitepaper	MonteCarlo.pdf

nbody - CUDA N-Body Simulation

This sample demonstrates efficient all-pairs simulation of a gravitational n-body simulation in CUDA. This sample accompanies the GPU Gems 3 chapter "Fast N-Body Simulation with CUDA". With CUDA 5.5, performance on Tesla K20c has increased to over 1.8TFLOP/s single precision. Double Performance has also improved on all Kepler and Fermi GPU architectures as well. Starting in CUDA 4.0, the nBody sample has been updated to take advantage of new features to easily scale the n-body simulation across multiple GPUs in a single PC. Adding "-numbodies=<bodies>" to the command line will allow users to set # of bodies for simulation. Adding "-numdevices=<N>" to the command line option will cause the sample to use N devices (if available) for simulation. In this mode, the position and velocity data for all bodies are read from system memory using "zero copy" rather than from device memory. For a small number of devices (4 or fewer) and a large enough number of bodies, bandwidth is not a bottleneck so we can achieve strong scaling across these devices.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemcpyToSymbol , cudaGraphicsMapResources , cudaEventRecord , cudaStreamQuery , cudaEventCreate , cudaGraphicsResourceGetMappedPointer , cudaGetDeviceCount , cudaEventElapsedTime , cudaDeviceSynchronize , cudaEventSynchronize , cudaGraphicsResourceSetMapFlags , cudaSetDeviceFlags , cudaEventDestroy , cudaDeviceCanAccessPeer , cudaSetDevice , cudaGraphicsUnmapResources , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Graphics Interop , Data Parallel Algorithms , Physically-Based Simulation
Supported OSes	Linux , Windows
Whitepaper	nbody_gems3_ch31.pdf

nbody_opengles - CUDA N-Body Simulation with GLES

This sample demonstrates efficient all-pairs simulation of a gravitational n-body simulation in CUDA. Unlike the OpenGL nbody sample, there is no user interaction.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GLES
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemcpyToSymbol , cudaGraphicsMapResources , cudaEventRecord , cudaStreamQuery , cudaEventCreate , cudaGraphicsResourceGetMappedPointer , cudaGetDeviceCount , cudaEventElapsedTime , cudaDeviceSynchronize , cudaEventSynchronize , cudaGraphicsResourceSetMapFlags , cudaSetDeviceFlags , cudaEventDestroy , cudaSetDevice , cudaGraphicsUnmapResources , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Graphics Interop , Data Parallel Algorithms , Physically-Based Simulation
Supported OSes	Linux

nbody_screen - CUDA N-Body Simulation on Screen

This sample demonstrates efficient all-pairs simulation of a gravitational n-body simulation in CUDA. Unlike the OpenGL nbody sample, there is no user interaction.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	screen , GLES
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemcpyToSymbol , cudaGraphicsMapResources , cudaEventRecord , cudaStreamQuery , cudaEventCreate , cudaGraphicsResourceGetMappedPointer , cudaGetDeviceCount ,

[cudaEventElapsedTime](#), [cudaDeviceSynchronize](#), [cudaEventSynchronize](#),
[cudaGraphicsResourceSetMapFlags](#), [cudaSetDeviceFlags](#),
[cudaEventDestroy](#), [cudaSetDevice](#), [cudaGraphicsUnmapResources](#),
[cudaGetDeviceProperties](#), [cudaGetDevice](#)

Key Concepts

[Graphics Interop](#), [Data Parallel Algorithms](#), [Physically-Based Simulation](#)

Supported OSes

- NV12toBGRandResize

This code shows two ways to convert and resize NV12 frames to BGR 3 planars frames using CUDA in batch. Way-1, Convert NV12 Input to BGR @ Input Resolution-1, then Resize to Resolution#2. Way-2, resize NV12 Input to Resolution#2 then convert it to BGR Output. NVIDIA HW Decoder, both dGPU and Tegra, normally outputs NV12 pitch format frames. For the inference using TensorRT, the input frame needs to be BGR planar format with possibly different size. So, conversion and resizing from NV12 to BGR planar is usually required for the inference following decoding. This CUDA code provides a reference implementation for conversion and resizing.

Supported SM

[SM 3.5](#), [SM 3.7](#), [SM 5.0](#), [SM 5.2](#), [SM 5.3](#), [SM 6.0](#), [SM 6.1](#), [SM 6.2](#), [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

Architecture**CUDA API**

[cudaStreamDestroy](#), [cudaFree](#), [cudaEventRecord](#), [cudaMallocManaged](#),
[cudaStreamCreate](#), [cudaEventCreate](#), [cudaEventElapsedTime](#),
[cudaDeviceSynchronize](#), [cudaDestroyTextureObject](#), [cudaEventSynchronize](#),
[cudaStreamAttachMemAsync](#), [cudaCreateTextureObject](#), [cudaMalloc](#),
[cudaEventDestroy](#), [cudaMemcpy](#)

Key Concepts

[Graphics Interop](#), [Image Processing](#), [Video Processing](#)

Supported OSes

[Linux](#), [Windows](#)

p2pBandwidthLatencyTest - Peer-to-Peer Bandwidth Latency Test with Multi-GPUs

This application demonstrates the CUDA Peer-To-Peer (P2P) data transfers between pairs of GPUs and computes latency and bandwidth. Tests on GPU pairs using P2P and without P2P are tested.

Supported SM

[SM 3.5](#), [SM 3.7](#), [SM 5.0](#), [SM 5.2](#), [SM 5.3](#), [SM 6.0](#), [SM 6.1](#), [SM 6.2](#), [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

Architecture**CUDA API**

[cudaDeviceEnablePeerAccess](#), [cudaOccupancyMaxPotentialBlockSize](#),
[cudaStreamCreateWithFlags](#), [cudaDeviceCanAccessPeer](#),
[cudaStreamDestroy](#), [cudaHostAlloc](#), [cudaEventCreate](#), [cudaMalloc](#),
[cudaEventDestroy](#), [cudaSetDevice](#), [cudaMemcpyPeerAsync](#),

	cudaGetDeviceProperties , cudaCheckError , cudaGetDeviceCount , cudaEventElapsedTime , cudaGetLastError , cudaDeviceDisablePeerAccess , cudaStreamSynchronize , cudaGetErrorString , cudaStreamWaitEvent , cudaMemset , cudaFree , cudaEventRecord , cudaFreeHost
Key Concepts	Performance Strategies , Asynchronous Data Transfers , Unified Virtual Address Space , Peer to Peer Data Transfers , Multi-GPU
Supported OSes	Linux , Windows

postProcessGL - Post-Process in OpenGL

This sample shows how to post-process an image rendered in OpenGL using CUDA.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaGraphicsUnregisterResource , cudaMemcpyToArray , cudaGraphicsGLRegisterBuffer , cudaHostAlloc , cudaGraphicsResourceGetMappedPointer , cudaProcess , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaCreateTextureObject , cudaGraphicsSubResourceGetMappedArray , cudaGetChannelDesc , cudaMalloc , cudaGraphicsGLRegisterImage , cudaGraphicsUnmapResources
Key Concepts	Graphics Interop , Image Processing
Supported OSes	Linux , Windows

quasirandomGenerator - Niederreiter Quasirandom Sequence Generator

This sample implements Niederreiter Quasirandom Sequence Generator and Inverse Cumulative Normal Distribution functions for the generation of Standard Normal Distributions.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemcpyToSymbol , cudaMemset , cudaFree , cudaDeviceSynchronize , cudaMalloc , cudaMemcpy

Key Concepts	Computational Finance
Supported OSes	Linux , Windows

quasirandomGenerator_nvrtc - Niederreiter Quasirandom Sequence Generator with libNVRTC

This sample implements Niederreiter Quasirandom Sequence Generator and Inverse Cumulative Normal Distribution functions for the generation of Standard Normal Distributions, compiling the CUDA kernels involved at runtime using NVRTC.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	NVRTC
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuMemFree , cuMemcpyDtoH , cuMemAlloc
Key Concepts	Computational Finance , Runtime Compilation
Supported OSes	Linux , Windows

recursiveGaussian - Recursive Gaussian Filter

This sample implements a Gaussian blur using Deriche's recursive method. The advantage of this method is that the execution time is independent of the filter width.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaDeviceSynchronize , cudaMalloc , cudaGraphicsUnregisterResource ,

[cudaGraphicsUnmapResources](#), [cudaMemcpy](#), [cudaGetDeviceProperties](#), [cudaGetDevice](#)

Key Concepts [Graphics Interop](#), [Image Processing](#)

Supported OSes [Linux](#), [Windows](#)

simpleD3D10 - Simple Direct3D10 (Vertex Array)

Simple program which demonstrates interoperability between CUDA and Direct3D10. The program generates a vertex array with CUDA and uses Direct3D10 to render the geometry. A Direct3D Capable device is required.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies [DirectX](#)

Supported SM Architecture [SM 3.5](#), [SM 3.7](#), [SM 5.0](#), [SM 5.2](#), [SM 5.3](#), [SM 6.0](#), [SM 6.1](#), [SM 6.2](#), [SM 7.0](#), [SM 7.2](#), [SM 7.5](#), [SM 8.0](#), [SM 8.6](#), [SM 8.7](#)

CUDA API [cudaGraphicsMapResources](#), [cudaGraphicsUnregisterResource](#), [cudaGraphicsResourceGetMappedPointer](#), [cudaGetDeviceCount](#), [cudaGetLastError](#), [cudaGraphicsUnmapResources](#), [cudaGetErrorString](#), [cudaGetDeviceProperties](#)

Key Concepts [Graphics Interop](#), [3D Graphics](#)

Supported OSes [Windows](#)

simpleD3D10RenderTarget - Simple Direct3D10 Render Target

Simple program which demonstrates interop of rendertargets between Direct3D10 and CUDA. The program uses RenderTarget positions with CUDA and generates a histogram with visualization. A Direct3D10 Capable device is required.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies [DirectX](#)

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaGraphicsMapResources , cudaGraphicsUnregisterResource , cudaGraphicsResourceGetMappedPointer , cudaGetDeviceCount , cudaMalloc , cudaGraphicsSubResourceGetMappedArray , cudaBindTextureToArray , cudaUnbindTexture , cudaGetLastError , cudaGraphicsUnmapResources , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Graphics Interop , Texture
Supported OSes	Windows

simpleD3D10Texture - Simple D3D10 Texture

Simple program which demonstrates how to interoperate CUDA with Direct3D10 Texture. The program creates a number of D3D10 Textures (2D, 3D, and CubeMap) which are generated from CUDA kernels. Direct3D then renders the results on the screen. A Direct3D10 Capable device is required.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	DirectX
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaGraphicsMapResources , cudaGraphicsUnregisterResource , cudaGetDeviceCount , cudaMallocPitch , cudaMalloc , cudaGraphicsSubResourceGetMappedArray , cudaGetLastError , cudaGraphicsUnmapResources , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Graphics Interop , Texture
Supported OSes	Windows

- Simple D3D11

Simple program which demonstrates how to use the CUDA D3D11 External Resource Interoperability APIs to update D3D11 buffers from CUDA and synchronize between D3D11 and CUDA with Keyed Mutexes.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be

installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	DirectX
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaSignalExternalSemaphoresAsync , cudaImportVertexBuffer , cudaAcquireSync , cudaGetDeviceCount , cudaImportExternalSemaphore , cudaReleaseSync , cudaImportExternalMemory , cudaExternalMemoryGetMappedBuffer , cudaImportKeyedMutex , cudaStreamCreateWithFlags , cudaDestroyExternalSemaphore , cudaSetDevice , cudaWaitExternalSemaphoresAsync , cudaDestroyExternalMemory
Key Concepts	Graphics Interop , Image Processing
Supported OSes	Windows

simpleD3D11Texture - Simple D3D11 Texture

Simple program which demonstrates Direct3D11 Texture interoperability with CUDA. The program creates a number of D3D11 Textures (2D, 3D, and CubeMap) which are written to from CUDA kernels. Direct3D then renders the results on the screen. A Direct3D Capable device is required.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	DirectX
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaGraphicsMapResources , cudaGraphicsUnregisterResource , cudaGetDeviceCount , cudaMallocPitch , cudaMalloc , cudaGraphicsSubResourceGetMappedArray , cudaGetLastError , cudaGraphicsUnmapResources , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Graphics Interop , Image Processing
Supported OSes	Windows

simpleD3D12 - Simple D3D12 CUDA Interop

A program which demonstrates Direct3D12 interoperability with CUDA. The program creates a sinewave in DX12 vertex buffer which is created using CUDA kernels. DX12 and CUDA synchronizes using DirectX12 Fences. Direct3D then renders the results on the screen. A DirectX12 Capable NVIDIA GPU is required on Windows10 or higher OS.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	DirectX12
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaSignalExternalSemaphoresAsync , cudaStreamCreate , cudaGetDeviceCount , cudaImportExternalSemaphore , cudaGetDeviceProperties , cudaImportExternalMemory , cudaExternalMemoryGetMappedBuffer , cudaDestroyExternalSemaphore , cudaSetDevice , cudaWaitExternalSemaphoresAsync , cudaStreamSynchronize , cudaDestroyExternalMemory
Key Concepts	Graphics Interop , CUDA DX12 Interop , Image Processing
Supported OSes	Windows

simpleD3D9 - Simple Direct3D9 (Vertex Arrays)

Simple program which demonstrates interoperability between CUDA and Direct3D9. The program generates a vertex array with CUDA and uses Direct3D9 to render the geometry. A Direct3D capable device is required.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	DirectX
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaGraphicsMapResources , cudaGraphicsUnregisterResource , cudaGraphicsResourceGetMappedPointer , cudaGetLastError , cudaGraphicsUnmapResources

Key Concepts	Graphics Interop
Supported OSes	Windows

simpleD3D9Texture - Simple D3D9 Texture

Simple program which demonstrates Direct3D9 Texture interoperability with CUDA. The program creates a number of D3D9 Textures (2D, 3D, and CubeMap) which are written to from CUDA kernels. Direct3D then renders the results on the screen. A Direct3D capable device is required.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	DirectX
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaGraphicsMapResources , cudaGraphicsUnregisterResource , cudaMallocPitch , cudaMalloc , cudaGraphicsSubResourceGetMappedArray , cudaGetLastError , cudaGraphicsUnmapResources
Key Concepts	Graphics Interop , Texture
Supported OSes	Windows

- Simple OpenGL

Simple program which demonstrates interoperability between CUDA and OpenGL. The program modifies vertex positions with CUDA and uses OpenGL to render the geometry.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaDeviceSynchronize , cudaMalloc , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaMemcpy

Key Concepts	Graphics Interop , Vertex Buffers , 3D Graphics
Supported OSes	Linux , Windows

simpleGLES - Simple OpenGL ES

Demonstrates data exchange between CUDA and OpenGL ES (aka Graphics interop). The program modifies vertex positions with CUDA and uses OpenGL ES to render the geometry.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GLES
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaDeviceSynchronize , cudaMalloc , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaMemcpy
Key Concepts	Graphics Interop , Vertex Buffers , 3D Graphics
Supported OSes	Linux

simpleGLES_EGLOutput - Simple OpenGL ES EGLOutput

Demonstrates data exchange between CUDA and OpenGL ES (aka Graphics interop). The program modifies vertex positions with CUDA and uses OpenGL ES to render the geometry, and shows how to render directly to the display using the EGLOutput mechanism and the DRM library. `` NOTE: On Orin platform, execute this command before running sample: \$ sudo modprobe nvidia-drm modeset=1 ``

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	EGLOutput , GLES
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaDeviceSynchronize ,

	cudaMalloc , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaMemcpy
Key Concepts	Graphics Interop , Vertex Buffers , 3D Graphics
Supported OSes	Linux

simpleGL ES_screen - Simple OpenGL ES on Screen

Demonstrates data exchange between CUDA and OpenGL ES (aka Graphics interop). The program modifies vertex positions with CUDA and uses OpenGL ES to render the geometry.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	screen , GLES
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaDeviceSynchronize , cudaMalloc , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaMemcpy
Key Concepts	Graphics Interop , Vertex Buffers , 3D Graphics
Supported OSes	

- Vulkan CUDA Interop Sinewave

This sample demonstrates Vulkan CUDA Interop. CUDA imports the Vulkan vertex buffer and operates on it to create sinewave, and synchronizes with Vulkan through vulkan semaphores imported by CUDA. This sample depends on Vulkan SDK, GLFW3 libraries, for building this sample please refer to "Build_instructions.txt" provided in this sample's directory

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , VULKAN
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7

CUDA API	cudaTimelineSemaphore , cudaSignalExternalSemaphoresAsync , cudaOccupancyMaxActiveBlocksPerMultiprocessor , cudaGetDeviceCount , cudaImportExternalSemaphore , cudaVertMem , cudaImportExternalMemory , cudaDestroyExternalMemory , cudaExternalMemoryGetMappedBuffer , cudaSignalSemaphore , cudaStreamCreateWithFlags , cudaWaitSemaphore , cudaHeightMap , cudaDestroyExternalSemaphore , cudaSetDevice , cudaWaitExternalSemaphoresAsync , cudaStreamSynchronize , cudaGetDeviceProperties
Key Concepts	Graphics Interop , CUDA Vulkan Interop , Data Parallel Algorithms
Supported OSes	Linux , Windows

- Vulkan CUDA Interop PI Approximation

This sample demonstrates Vulkan CUDA Interop via cuMemMap APIs. CUDA exports buffers that Vulkan imports as vertex buffer. CUDA invokes kernels to operate on vertices and synchronizes with Vulkan through vulkan semaphores imported by CUDA. This sample depends on Vulkan SDK, GLFW3 libraries, for building this sample please refer to "Build_instructions.txt" provided in this sample's directory

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , VULKAN
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuMemRelease , cuMemExportToShareableHandle , cuMemSetAccess , cuMemMap , cuMemCreate , cuMemAddressFree , cuMemGetAllocationGranularity , cuMemUnmap , cuMemAddressReserve cudaOccupancyMaxActiveBlocksPerMultiprocessor , cudaImportExternalSemaphore , cudaStreamCreateWithFlags , cudaMemcpyAsync , cudaStreamDestroy , cudaSignalExternalSemaphoresAsync , cudaMallocHost , cudaMalloc , cudaSetDevice , cudaGetDeviceProperties , cudaGetDeviceCount , cudaDeviceGetAttribute , cudaSignalSemaphore , cudaWaitSemaphore , cudaDestroyExternalSemaphore , cudaStreamSynchronize , cudaFree , cudaLaunchHostFunc , cudaMemsetAsync , cudaFreeHost , cudaWaitExternalSemaphoresAsync
Key Concepts	cuMemMap IPC , MMAP , Graphics Interop , CUDA Vulkan Interop , Data Parallel Algorithms
Supported OSes	Linux , Windows

SLID3D10Texture - SLI D3D10 Texture

Simple program which demonstrates SLI with Direct3D10 Texture interoperability with CUDA. The program creates a D3D10 Texture which is written to from a CUDA kernel. Direct3D then renders the results on the screen. A Direct3D Capable device is required.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	DirectX
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cuCtxPushCurrent , cuCtxPopCurrent , cudaMemset , cudaFree , cudaGraphicsMapResources , cudaGraphicsUnregisterResource , cudaGraphicsResourceSetMapFlags , cudaGetDeviceCount , cudaMallocPitch , cudaMalloc , cudaGraphicsSubResourceGetMappedArray , cudaGetLastError , cudaGraphicsUnmapResources , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Performance Strategies , Graphics Interop , Image Processing , 2D Textures
Supported OSes	Windows

smokeParticles - Smoke Particles

Smoke simulation with volumetric shadows using half-angle slicing technique. Uses CUDA for procedural simulation, Thrust Library for sorting algorithms, and OpenGL for graphics rendering.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaCreateTextureObject , cudaExtent , cudaMemcpyToSymbol , cudaPitchedPtr
Key Concepts	Graphics Interop , Data Parallel Algorithms , Physically-Based Simulation
Supported OSes	Linux , Windows
Whitepaper	smokeParticles.pdf

SobelFilter - Sobel Filter

This sample implements the Sobel edge detection filter for 8-bit monochrome images.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaGraphicsMapResources , cudaFreeArray , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaCreateTextureObject , cudaMalloc , cudaMallocArray , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaMemcpy , cudaGetErrorString
Key Concepts	Graphics Interop , Image Processing
Supported OSes	Linux , Windows

SobolQRNG - Sobol Quasirandom Number Generator

This sample implements Sobol Quasirandom Sequence Generator.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaDeviceSynchronize , cudaMalloc , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Computational Finance
Supported OSes	Linux , Windows

stereoDisparity - Stereo Disparity Computation (SAD SIMD Intrinsics)

A CUDA program that demonstrates how to compute a stereo disparity map using SIMD SAD (Sum of Absolute Difference) intrinsics. Requires Compute Capability 2.0 or higher.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaEventRecord , cudaEventCreate , cudaEventElapsedTime , cudaDeviceSynchronize , cudaEventSynchronize , cudaCreateTextureObject , cudaMalloc , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Image Processing , Video Intrinsic s
Supported OSes	Linux , Windows

VFlockingD3D10

The sample models formation of V-shaped flocks by big birds, such as geese and cranes. The algorithms of such flocking are borrowed from the paper "V-like formations in flocks of artificial birds" from Artificial Life, Vol. 14, No. 2, 2008. The sample has CPU- and GPU-based implementations. Press 'g' to toggle between them. The GPU-based simulation works many times faster than the CPU-based one. The printout in the console window reports the simulation time per step. Press 'r' to reset the initial distribution of birds.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	DirectX
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemset , cudaFree , cudaGraphicsMapResources , cudaEventRecord , cudaGraphicsUnregisterResource , cudaEventCreate , cudaGraphicsResourceGetMappedPointer , cudaGetDeviceCount , cudaEventElapsedTime , cudaDeviceSynchronize , cudaEventSynchronize , cudaGraphicsResourceSetMapFlags , cudaMalloc , cudaGetLastError , cudaGraphicsUnmapResources , cudaMemcpy , cudaGetErrorString , cudaGetDeviceProperties
Key Concepts	Graphics Interop , Data Parallel Algorithms , Physically-Based Simulation , Performance Strategies
Supported OSes	Windows

volumeFiltering - Volumetric Filtering with 3D Textures and Surface Writes

This sample demonstrates 3D Volumetric Filtering using 3D Textures and 3D Surface Writes.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaGraphicsMapResources , cudaGraphicsResourceGetMappedPointer , cudaMemcpy , cudaMemcpyToSymbol , cudaDestroySurfaceObject , cudaPitchedPtr , cudaMalloc , cudaGraphicsUnregisterResource , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaMemset , cudaFree , cudaFreeArray , cudaGraphicsGLRegisterBuffer , cudaExtent , cudaCreateSurfaceObject , cudaCreateTextureObject , cudaMallocArray , cudaGraphicsUnmapResources
Key Concepts	Graphics Interop , Image Processing , 3D Textures , Surface Writes
Supported OSes	Linux , Windows

volumeRender - Volume Rendering with 3D Textures

This sample demonstrates basic volume rendering using 3D Textures.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , GL
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaMemcpyToSymbol , cudaMemset , cudaFree , cudaGraphicsMapResources , cudaFreeArray , cudaGraphicsGLRegisterBuffer , cudaGraphicsResourceGetMappedPointer , cudaExtent , cudaDeviceSynchronize , cudaDestroyTextureObject , cudaPitchedPtr , cudaCreateTextureObject , cudaMalloc , cudaProfilerStop , cudaMallocArray , cudaGraphicsUnregisterResource , cudaGraphicsUnmapResources , cudaMemcpy
Key Concepts	Graphics Interop , Image Processing , 3D Textures
Supported OSes	Linux , Windows

- Vulkan Image - CUDA Interop

This sample demonstrates Vulkan Image - CUDA Interop. CUDA imports the Vulkan image buffer, performs box filtering over it, and synchronizes with Vulkan through Vulkan semaphores imported by CUDA. This sample depends on Vulkan SDK, GLFW3 libraries, for building this sample please refer to "Build_instructions.txt" provided in this sample's directory

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	X11 , VULKAN
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaGetMipmappedArrayLevel , cudaImportExternalSemaphore , cudaExternalMemoryGetMappedMipmappedArray , cudaMemcpy , cudaDestroyExternalMemory , cudaSignalExternalSemaphoresAsync , cudaFreeMipmappedArray , cudaVkSemaphoreSignal , cudaVkImportImageMem , cudaDestroySurfaceObject , cudaImportExternalMemory , cudaMalloc , cudaSetDevice , cudaGetDeviceProperties , cudaGetDeviceCount , cudaDestroyTextureObject , cudaUpdateVkImage , cudaDestroyExternalSemaphore , cudaFree , cudaStreamCreate , cudaVkSemaphoreWait , cudaExtent , cudaVkImportSemaphore , cudaCreateSurfaceObject , cudaMallocMipmappedArray , cudaCreateTextureObject , cudaWaitExternalSemaphoresAsync
Key Concepts	Graphics Interop , CUDA Vulkan Interop , Data Parallel Algorithms
Supported OSes	Linux , Windows

3.7. Performance Reference

alignedTypes - Aligned Types

A simple test, showing huge access speed gap between aligned and misaligned structures. It measures per-element copy throughput for aligned and misaligned structures on big chunks of data.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
----------------------------------	---

CUDA API	cudaMemset , cudaFree , cudaDeviceSynchronize , cudaMalloc , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	Performance Strategies
Supported OSes	Linux , Windows

transpose - Matrix Transpose

This sample demonstrates Matrix Transpose. Different performance are shown to achieve high performance.

Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaFree , cudaEventRecord , cudaEventCreate , cudaEventElapsedTime , cudaEventSynchronize , cudaMalloc , cudaEventDestroy , cudaGetLastError , cudaMemcpy , cudaGetDeviceProperties , cudaGetDevice
Key Concepts	Performance Strategies , Linear Algebra
Supported OSes	Linux , Windows
Whitepaper	MatrixTranspose.pdf

UnifiedMemoryPerf - Unified and other CUDA Memories Performance

This sample demonstrates the performance comparison using matrix multiplication kernel of Unified Memory with/without hints and other types of memory like zero copy buffers, pageable, pagelocked memory performing synchronous and Asynchronous transfers on a single GPU.

This sample depends on other applications or libraries to be present on the system to either build or run. If these dependencies are not available on the system, the sample will not be installed. If these dependencies are available, but not installed, the sample will waive itself at build time.

Dependencies	UVM
Supported SM Architecture	SM 3.5 , SM 3.7 , SM 5.0 , SM 5.2 , SM 5.3 , SM 6.0 , SM 6.1 , SM 6.2 , SM 7.0 , SM 7.2 , SM 7.5 , SM 8.0 , SM 8.6 , SM 8.7
CUDA API	cudaStreamDestroy , cudaFree , cudaMallocHost , cudaMallocManaged , cudaMemPrefetchAsync , cudaStreamCreate , cudaStreamAttachMemAsync , cudaFreeHost , cudaMalloc , cudaMemcpyAsync , cudaStreamSynchronize , cudaHostGetDevicePointer , cudaMemcpy , cudaGetDeviceProperties
Key Concepts	CUDA Systems Integration , Unified Memory , CUDA Streams and Events , Pinned System Paged Memory

Supported OSes

[Linux](#), [Windows](#)

Chapter 4. Dependencies

Some CUDA Samples rely on third-party applications and/or libraries, or features provided by the CUDA Toolkit and Driver, to either build or execute. These dependencies are listed below.

If a sample has a dependency that is not available on the system, the sample will not be installed. If a sample has a third-party dependency that is available on the system, but is not installed, the sample will waive itself at build time.

Each sample's dependencies are listed in the [Samples Reference](#) section.

Third-Party Dependencies

These third-party dependencies are required by some CUDA samples. If available, these dependencies are either installed on your system automatically, or are installable via your system's package manager (Linux) or a third-party website.

FreeImage

FreeImage is an open source imaging library. FreeImage can usually be installed on Linux using your distribution's package manager system. FreeImage can also be downloaded from the [FreeImage website](#).

To set up FreeImage on a Windows system, extract the FreeImage DLL distribution into the `7_CUDALibraries/common/` folder such that `7_CUDALibraries/common/FreeImage/Dist/x64/` contains the `.h`, `.dll`, and `.lib` files.

Message Passing Interface

MPI (Message Passing Interface) is an API for communicating data between distributed processes. A MPI compiler can be installed using your Linux distribution's package manager system. It is also available on some online resources, such as [Open MPI](#). On Windows, to build and run MPI-CUDA applications one can install [MS-MPI SDK](#).

Only 64-Bit

Some samples can only be run on a 64-bit operating system.

DirectX

DirectX is a collection of APIs designed to allow development of multimedia applications on Microsoft platforms. For Microsoft platforms, NVIDIA's CUDA Driver supports DirectX. Several CUDA Samples for Windows demonstrates CUDA-DirectX Interoperability, for building such samples one needs to install Microsoft Visual Studio 2012 or higher which provides Microsoft Windows SDK for Windows 8.

DirectX 12

DirectX 12 is a collection of advanced low-level programming APIs which can reduce driver overhead, designed to allow development of multimedia applications on Microsoft platforms starting with Windows 10 OS onwards. For Microsoft platforms, NVIDIA's CUDA Driver supports DirectX. Few CUDA Samples for Windows demonstrates CUDA-DirectX12 Interoperability, for building such samples one needs to install [Windows 10 SDK or higher](#), with VS 2015 or VS 2017.

OpenGL

OpenGL is a graphics library used for 2D and 3D rendering. On systems which support OpenGL, NVIDIA's OpenGL implementation is provided with the CUDA Driver.

OpenGL ES

OpenGL ES is an embedded systems graphics library used for 2D and 3D rendering. On systems which support OpenGL ES, NVIDIA's OpenGL ES implementation is provided with the CUDA Driver.

Vulkan

Vulkan is a low-overhead, cross-platform 3D graphics and compute API. Vulkan targets high-performance realtime 3D graphics applications such as video games and interactive media across all platforms. On systems which support Vulkan, NVIDIA's Vulkan implementation is provided with the CUDA Driver. For building and running Vulkan applications one needs to install the [Vulkan SDK](#).

OpenMP

OpenMP is an API for multiprocessing programming. OpenMP can be installed using your Linux distribution's package manager system. It usually comes preinstalled with GCC. It can also be found at the [OpenMP website](#).

Screen

Screen is a windowing system found on the QNX operating system. Screen is usually found as part of the root filesystem.

X11

X11 is a windowing system commonly found on *-nix style operating systems. X11 can be installed using your Linux distribution's package manager, and comes preinstalled on Mac OS X systems.

EGL

EGL is an interface between Khronos rendering APIs (such as OpenGL, OpenGL ES or OpenVG) and the underlying native platform windowing system.

EGLOutput

EGLOutput is a set of EGL extensions which allow EGL to render directly to the display.

EGLSync

EGLSync is a set of EGL extensions which provides sync objects that are synchronization primitive, representing events whose completion can be tested or waited upon.

NVSCI

NvSci is a set of communication interface libraries out of which CUDA interoperates with NvSciBuf and NvSciSync. NvSciBuf allows applications to allocate and exchange buffers in memory. NvSciSync allows applications to manage synchronization objects which coordinate when sequences of operations begin and end.

NvMedia

NvMedia provides powerful processing of multimedia data for true hardware acceleration across NVIDIA Tegra devices. Applications leverage the NvMedia Application Programming Interface (API) to process the image and video data.

CUDA Features

These CUDA features are needed by some CUDA samples. They are provided by either the CUDA Toolkit or CUDA Driver. Some features may not be available on your system.

CUFFT Callback Routines

CUFFT Callback Routines are user-supplied kernel routines that CUFFT will call when loading or storing data. These callback routines are only available on Linux x86_64 and ppc64le systems.

CUDA Dynamic Paralellism

CDP (CUDA Dynamic Paralellism) allows kernels to be launched from threads running on the GPU. CDP is only available on GPUs with SM architecture of 3.5 or above.

Multi-block Cooperative Groups

Multi Block Cooperative Groups(MBCG) extends Cooperative Groups and the CUDA programming model to express inter-thread-block synchronization. MBCG is available on GPUs with Pascal and higher architecture.

Multi-Device Cooperative Groups

Multi Device Cooperative Groups extends Cooperative Groups and the CUDA programming model enabling thread blocks executing on multiple GPUs to cooperate and synchronize as they execute. This feature is available on GPUs with Pascal and higher architecture.

CUBLAS

CUBLAS (CUDA Basic Linear Algebra Subroutines) is a GPU-accelerated version of the BLAS library.

CUDA Interprocess Communication

IPC (Interprocess Communication) allows processes to share device pointers.

CUFFT

CUFFT (CUDA Fast Fourier Transform) is a GPU-accelerated FFT library.

CURAND

CURAND (CUDA Random Number Generation) is a GPU-accelerated RNG library.

CUSPARSE

CUSPARSE (CUDA Sparse Matrix) provides linear algebra subroutines used for sparse matrix calculations.

CUSOLVER

CUSOLVER library is a high-level package based on the CUBLAS and CUSPARSE libraries. It combines three separate libraries under a single umbrella, each of which can be used independently or in concert with other toolkit libraries. The intent of CUSOLVER is to provide useful LAPACK-like features, such as common matrix factorization and triangular solve routines for dense matrices, a sparse least-squares solver and an eigenvalue solver. In

addition cuSolver provides a new refactorization library useful for solving sequences of matrices with a shared sparsity pattern.

NPP

NPP (NVIDIA Performance Primitives) provides GPU-accelerated image, video, and signal processing functions.

NVJPEG

NVJPEG library provides high-performance, GPU accelerated JPEG decoding functionality for image formats commonly used in deep learning and hyperscale multimedia applications.

NVRTC

NVRTC (CUDA RunTime Compilation) is a runtime compilation library for CUDA C++.

Stream Priorities

Stream Priorities allows the creation of streams with specified priorities. Stream Priorities is only available on GPUs with SM architecture of 3.5 or above.

Unified Virtual Memory

UVM (Unified Virtual Memory) enables memory that can be accessed by both the CPU and GPU without explicit copying between the two. UVM is only available on Linux and Windows systems.

16-bit Floating Point

FP16 is a 16-bit floating-point format. One bit is used for the sign, five bits for the exponent, and ten bits for the mantissa. FP16 is only available on specific mobile platforms.

C++11 CUDA

NVCC Support of [C++11 features](#).

Chapter 5. Key Concepts and Associated Samples

The tables below describe the key concepts of the CUDA Toolkit and lists the samples that illustrate how that concept is used.

Basic Key Concepts

Basic Concepts demonstrates how to make use of CUDA features.

Table 1. Basic Key Concepts and Associated Samples

Basic Key Concept	Description	Samples
3D Graphics	<i>3D Rendering</i>	Random Fog, Simple Direct3D10 (Vertex Array), Simple OpenGL, Simple OpenGLES, Simple OpenGLES EGLOutput, Simple OpenGLES on Screen
3D Textures	<i>Volume Textures</i>	Simple Texture 3D
Arrive Wait Barrier		Simple Arrive Wait Barrier
Assert	<i>GPU Assert</i>	simpleAssert, simpleAssert with libNVRTC
Asynchronous Data Transfers	<i>Overlapping I/O and Compute</i>	Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Simple Multi Copy and Compute, Simple Multi-GPU, Simple Peer-to-Peer Transfers with Multi-GPU, asyncAPI, simpleStreams
Atomic Intrinsic	<i>Using atomics with GPU kernels</i>	Simple Atomic Intrinsic, Simple Atomic Intrinsic with libNVRTC, System wide

Basic Key Concept	Description	Samples
		Atomics, Warp Aggregated Atomics using Cooperative Groups
Attributes usage on stream		simpleAttributes
C++ Function Overloading	<i>Use C++ overloading with GPU kernels</i>	cppOverload
C++ Templates	<i>Using Templates with GPU kernels</i>	NewDelete, Simple Templates, Simple Templates with libNVRTC
CPP-CUDA Integration		C++ Integration
CPP11 CUDA	<i>Samples demonstrating how to use C++11 feature support in CUDA.</i>	Global Memory to Shared Memory Async Copy
CUBLAS	<i>CUDA BLAS samples</i>	Matrix Multiplication (CUBLAS), Unified Memory Streams
CUBLAS Library	<i>CUDA BLAS samples</i>	Simple CUBLAS, Simple CUBLAS LU, batchCUBLAS
CUBLAS-XT Library	<i>cuBLAS XT is a library which further accelerates Level 3 BLAS calls by spreading work across multiple GPUs connected to the same motherboard.</i>	Simple CUBLAS XT
CUDA Driver API	<i>Samples that show the CUDA Driver API</i>	CUDA Compressible Memory, Device Query Driver API, Matrix Multiplication (CUDA Driver API Version), Memmap IPC Driver API, Simple Driver-Runtime Interaction, Simple Texture (Driver Version), Simple Vote Intrinsic with libNVRTC, Using Inline PTX, Using Inline PTX with libNVRTC, Vector Addition Driver API, Vector Addition cuMemMap, Vector Addition with libNVRTC
CUDA Dynamic Parallelism	<i>Dynamic Parallelism with GPU Kernels (SM 3.5)</i>	Simple Print (CUDA Dynamic Parallelism)

Basic Key Concept	Description	Samples
CUDA Graphs	<i>Samples demonstrating how to use CUDA Graphs API to create, instantiate and launch CUDA Operations.</i>	Graph Memory Footprint, Graph Memory Nodes, Jacobi CUDA Graphs, Simple CUDA Graphs
CUDA NvSci Interop	<i>Samples showing CUDA and NvSciBuf/NvSciSync Interop.</i>	CUDA NvSciBuf/NvSciSync Interop, NvMedia CUDA Interop
CUDA Runtime API	<i>Samples that use the Runtime API</i>	Device Query, FP16 Scalar Product, Global Memory to Shared Memory Async Copy, Graph Memory Footprint, Matrix Multiplication (CUBLAS), Matrix Multiplication (CUDA Runtime API Version), Matrix Multiplication with libNVRTC, Simple Driver-Runtime Interaction, Simple Texture, Vector Addition
CUDA Streams	<i>Stream API defines a sequence of operations that can be overlapped with I/O</i>	Simple CUDA Callbacks
CUDA Streams and Events	<i>Synchronizing Kernels with Event Timers and Streams</i>	Bandwidth Test, Simple Multi Copy and Compute, Simple Multi-GPU, Unified Memory Streams, Unified and other CUDA Memories Performance, asyncAPI, cppOverload, simpleStreams
CUDA Systems Integration	<i>Samples that integrate with Multi Process (OpenMP, IPC, and MPI)</i>	Unified Memory Streams, Unified and other CUDA Memories Performance, cudaOpenMP, simpleIPC, simpleMPI
CUFFT Library	<i>Samples that use the CUDA FFT accelerated library</i>	Simple CUFFT, Simple CUFFT Callbacks, Simple CUFFT_MGPU, SimpleCUFFT_2d_MGPU
CURAND Library	<i>Samples that use the CUDA random number generator</i>	MersenneTwisterGP11213, Random Fog
CUSOLVER Library	<i>Samples that use the cuSOLVER accelerated library</i>	cuSolverDn Linear Solver , cuSolverRf Refactorization, cuSolverSp Linear Solver , cuSolverSp Lowlevel QR Solver, cuSolverSp LowlevelCholesky Solver
Callback Functions	<i>Creating Callback functions with GPU kernels</i>	Simple CUDA Callbacks

Basic Key Concept	Description	Samples
Compressible Memory		CUDA Compressible Memory
Computational Finance	<i>Finance Algorithms</i>	Black-Scholes Option Pricing, Black-Scholes Option Pricing with libNVRTC
Cooperative Groups	<i>Cooperative Groups is an extension to the CUDA programming model that allows the CUDA program to express the granularity at which different-sized groups of threads are communicating.</i>	Advanced Quicksort (CUDA Dynamic Parallelism), Binary Partition Cooperative Groups, DirectX Texture Compressor (DXTC), Jacobi CUDA Graphs, Quad Tree (CUDA Dynamic Parallelism), Reduction using MultiBlock Cooperative Groups, Simple Cooperative Groups, Warp Aggregated Atomics using Cooperative Groups, conjugateGradient using MultiBlock Cooperative Groups, conjugateGradient using MultiDevice Cooperative Groups, threadFenceReduction
Data Parallel Algorithms	<i>Samples that show good usage of Data Parallel Algorithms</i>	CUDA Separable Convolution, Texture-based Separable Convolution
Debugging	<i>Samples useful for debugging</i>	CUDA Interception Library, simplePrintf
Device Memory Allocation	<i>Samples that show GPU Device side memory allocation</i>	NewDelete, Template
Device Query	<i>Sample showing simple device query of information</i>	Device Query, Device Query Driver API
EGLImage-CUDA Interop	<i>Samples demonstrating how to use EGL Image and CUDA Interop.</i>	EGLSync CUDA Event Interop
EGLStreams Interop	<i>Samples demonstrating how to use EGL Streams and CUDA Interop.</i>	EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU
EGLSync-CUDAEvent Interop	<i>Samples demonstrating interoperability between CUDA Event and EGL Sync for achieving GPU-side synchronization between EGL and CUDA operations without blocking CPU for synchronization.</i>	EGLSync CUDA Event Interop

Basic Key Concept	Description	Samples
GPU Performance	<i>Samples demonstrating high performance and data I/O</i>	Simple Multi Copy and Compute
Graphics Interop	<i>Samples that demonstrate interop between graphics APIs and CUDA</i>	Bicubic B-spline Interpolation, Bilateral Filter, Box Filter, CUDA and OpenGL Interop of Images, NV12toBGRandResize, Simple D3D10 Texture, Simple D3D11, Simple D3D11 Texture, Simple D3D12 CUDA Interop, Simple D3D9 Texture, Simple Direct3D10 (Vertex Array), Simple Direct3D10 Render Target, Simple Direct3D9 (Vertex Arrays), Simple OpenGL, Simple OpenGL ES, Simple OpenGL ES EGLOutput, Simple OpenGL ES on Screen, Simple Texture 3D
Image Decoding	<i>Samples demonstrating image decoding on GPU.</i>	NVJPEG simple
Image Encoding		NVJPEG Encoder
Image Processing	<i>Samples that demonstrate image processing algorithms in CUDA</i>	Batched Label Markers And Label Compression NPP, Bicubic B-spline Interpolation, Bilateral Filter, Box Filter, Box Filter with NPP, CUDA Separable Convolution, CUDA and OpenGL Interop of Images, Canny Edge Detector NPP, Filter Border Control NPP, FreeImage and NPP Interopability, Histogram Equalization with NPP, NV12toBGRandResize, Pitch Linear Texture, Simple CUBLAS, Simple CUFFT, Simple CUFFT Callbacks, Simple CUFFT_MGPU, Simple D3D11, Simple D3D11 Texture, Simple D3D12 CUDA Interop, Simple Surface Write, Simple Texture, Simple Texture (Driver Version), Simple Texture 3D, SimpleCUFFT_2d_MGPU, Texture-based Separable Convolution, Watershed Segmentation NPP

Basic Key Concept	Description	Samples
Instantiated CUDA Graph Update		Jacobi CUDA Graphs
InterProcess Communication	<i>Samples that demonstrate Inter Process Communication between processes</i>	simpleIPC
LU decomposition		Simple CUBLAS LU
Linear Algebra	<i>Samples demonstrating linear algebra with CUDA</i>	Global Memory to Shared Memory Async Copy, Matrix Multiplication (CUBLAS), Matrix Multiplication (CUDA Runtime API Version), Matrix Multiplication with libNVRTC, batchCUBLAS, cuSolverDn Linear Solver , cuSolverRf Refactorization, cuSolverSp Linear Solver , cuSolverSp Lowlevel QR Solver, cuSolverSp LowlevelCholesky Solver
MMAP		CUDA Compressible Memory, Memmap IPC Driver API, Vector Addition cuMemMap, Vulkan CUDA Interop PI Approximation
MPI	<i>Samples demonstrating how to use CUDA with MPI programs</i>	simpleMPI
Matrix Multiply	<i>Samples demonstrating matrix multiply CUDA</i>	CUDA Tensor Core GEMM, Double Precision Tensor Core GEMM, Matrix Multiplication (CUDA Driver API Version), Tensor Core GEMM Integer MMA, bfloat16 Tensor Core GEMM, tf32 Tensor Core GEMM
Multi-GPU	<i>Samples demonstrating how to take advantage of multiple GPUs and CUDA</i>	Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Simple Multi-GPU, Simple Peer-to-Peer Transfers with Multi-GPU, Topology Query
Multithreading	<i>Samples demonstrating how to use multithreading with CUDA</i>	Simple CUDA Callbacks, Simple Multi-GPU, Unified Memory Streams, cudaOpenMP, simpleMPI

Basic Key Concept	Description	Samples
NPP Library	<i>Samples demonstrating how to use NPP (NVIDIA Performance Primitives) for image processing</i>	Batched Label Markers And Label Compression NPP, Box Filter with NPP, Canny Edge Detector NPP, Filter Border Control NPP, FreeImage and NPP Interopability, Histogram Equalization with NPP, Watershed Segmentation NPP
NVJPEG Library	<i>NVJPEG library samples</i>	NVJPEG Encoder, NVJPEG simple
Occupancy Calculator	<i>Samples demonstrating how to use the CUDA Occupancy Calculator</i>	simpleOccupancy
OpenMP	<i>Samples demonstrating how to use OpenMP</i>	Unified Memory Streams, cudaOpenMP
Overlap Compute and Copy	<i>Samples demonstrating how to overlap Compute and Data I/O</i>	Simple Multi Copy and Compute
PTX Assembly	<i>Samples demonstrating how to use PTX code with CUDA</i>	Using Inline PTX, Using Inline PTX with libNVRTC
Peer to Peer	<i>Samples demonstrating how to handle P2P data transfers between multiple GPUs</i>	simpleIPC
Peer to Peer Data Transfers	<i>Samples demonstrating how to handle P2P data transfers between multiple GPUs</i>	Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Simple Peer-to-Peer Transfers with Multi-GPU
Performance Strategies	<i>Samples demonstrating high performance with CUDA</i>	Bandwidth Test, Batched Label Markers And Label Compression NPP, Box Filter with NPP, CUDA and OpenGL Interop of Images, Canny Edge Detector NPP, Clock, Clock libNVRTC, Filter Border Control NPP, FreeImage and NPP Interopability, Graph Memory Footprint, Histogram Equalization with NPP, Matrix Multiplication (CUBLAS), Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Simple Peer-to-Peer Transfers with Multi-GPU, Topology Query, Using Inline PTX, Using Inline PTX with libNVRTC, Watershed Segmentation NPP, simpleZeroCopy, stream Ordered

Basic Key Concept	Description	Samples
		Allocation, stream Ordered Allocation IPC Pools, stream Ordered Allocation Peer-to-Peer access
Pinned System Paged Memory	<i>Samples demonstrating how to properly handle data I/O efficiently between the CPU host and GPU video memory</i>	Unified and other CUDA Memories Performance, simpleZeroCopy
Separate Compilation	<i>Samples demonstrating how to use CUDA library linking</i>	Simple Static GPU Device Library
Stream Capture	<i>Samples demonstrating how to use Stream Capture API to create CUDA Graphs.</i>	Graph Memory Nodes, Jacobi CUDA Graphs, Simple CUDA Graphs
Surface Writes	<i>Samples demonstrating how to use Surface Writes with GPU kernels</i>	Simple Surface Write, Simple Texture 3D
Texture	<i>Samples demonstrating how to use textures GPU kernels</i>	Pitch Linear Texture, Simple Cubemap Texture, Simple D3D10 Texture, Simple D3D9 Texture, Simple Direct3D10 Render Target, Simple Layered Texture, Simple Surface Write, Simple Texture, Simple Texture (Driver Version), Texture-based Separable Convolution
Unified Memory	<i>Samples demonstrating how to use Unified Memory</i>	ConjugateGradientUM, System wide Atomics, Unified Memory Streams, Unified and other CUDA Memories Performance, conjugateGradient using MultiBlock Cooperative Groups, conjugateGradient using MultiDevice Cooperative Groups
Unified Virtual Address Space	<i>Samples demonstrating how to use UVA with CUDA programs</i>	Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Simple Peer-to-Peer Transfers with Multi-GPU
Using NPP Batch Functions		Batched Label Markers And Label Compression NPP
Vector Addition	<i>Samples demonstrating how to use Vector Addition with CUDA programs</i>	Simple Driver-Runtime Interaction, Vector Addition, Vector Addition Driver API, Vector

Basic Key Concept	Description	Samples
		Addition cuMemMap, Vector Addition with libNVRTC, simpleZeroCopy
Vertex Buffers	<i>Samples demonstrating how to use Vertex Buffers with CUDA kernels</i>	Simple OpenGL, Simple OpenGL ES, Simple OpenGL ES EGLOutput, Simple OpenGL ES on Screen
Video Processing		NV12toBGRandResize
Volume Processing	<i>Samples demonstrating how to use 3D Textures for volume rendering</i>	Simple Cubemap Texture, Simple Layered Texture
Vote Intrinsic	<i>Samples demonstrating how to use vote intrinsics with CUDA</i>	Simple Vote Intrinsic, Simple Vote Intrinsic with libNVRTC
cuDLA		cuDLA Error Reporting, cuDLA Hybrid Mode, cuDLA Standalone Mode
cuMemMap IPC		Memmap IPC Driver API, Vulkan CUDA Interop PI Approximation

Advanced Key Concepts

Advanced Concepts demonstrate advanced techniques and algorithms implemented with CUDA.

Table 2. Advanced Key Concepts and Associated Samples

Advanced Key Concept	Description	Samples
2D Textures	<i>Texture Mapping</i>	SLI D3D10 Texture
3D Graphics	<i>3D Rendering</i>	Marching Cubes Isosurfaces
3D Textures	<i>Volume Textures</i>	Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes
C++11 CUDA	<i>Samples demonstrating how to use C++11 feature support in CUDA.</i>	C++11 CUDA
CUBLAS Library	<i>CUDA BLAS samples</i>	Conjugate Gradient using Cuda Graphs, ConjugateGradient, ConjugateGradientUM,

Advanced Key Concept	Description	Samples
		Preconditioned Conjugate Gradient, conjugateGradient using MultiBlock Cooperative Groups, conjugateGradient using MultiDevice Cooperative Groups
CUDA DX12 Interop	<i>Samples showing CUDA and D3D12 interop.</i>	Simple D3D12 CUDA Interop
CUDA Driver API	<i>Samples that show the CUDA Driver API</i>	CUDA Context Thread Management, Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), PTX Just-in-Time compilation
CUDA Dynamic Parallelism	<i>Dynamic Parallelism with GPU Kernels (SM 3.5)</i>	Advanced Quicksort (CUDA Dynamic Parallelism), Bezier Line Tessellation (CUDA Dynamic Parallelism), Quad Tree (CUDA Dynamic Parallelism), Simple Quicksort (CUDA Dynamic Parallelism)
CUDA Dynamically Linked Library	<i>Dynamic loading of the CUDA DLL using CUDA Driver API</i>	Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version)
CUDA Streams and Events	<i>Synchronizing Kernels with Event Timers and Streams</i>	Stream Priorities
CUDA Systems Integration	<i>Samples that integrate with Multi Process (OpenMP, IPC, and MPI)</i>	simpleHyperQ
CUDA Vulkan Interop	<i>Samples demonstrating how to use Vulkan - CUDA Interop.</i>	Vulkan CUDA Interop PI Approximation, Vulkan CUDA Interop Sinewave, Vulkan Image - CUDA Interop
CUFFT Library	<i>Samples that use the CUDA FFT accelerated library</i>	CUDA FFT Ocean Simulation, FFT-Based 2D Convolution, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGL ES Version)
CURAND Library	<i>Samples that use the CUDA random number generator</i>	Monte Carlo Estimation of Pi (batch PRNG), Monte Carlo Estimation of Pi (batch QRNG), Monte Carlo Estimation of Pi (inline PRNG), Monte Carlo Estimation

Advanced Key Concept	Description	Samples
		of Pi (inline QRNG) , Monte Carlo Option Pricing with Multi-GPU support, Monte Carlo Single Asian Option
CUSPARSE Library	<i>Samples that use the cuSPARSE (Sparse Vector Matrix Multiply) functions</i>	Conjugate Gradient using Cuda Graphs, ConjugateGradient, ConjugateGradientUM, Preconditioned Conjugate Gradient, conjugateGradient using MultiBlock Cooperative Groups, conjugateGradient using MultiDevice Cooperative Groups
Computational Finance	<i>Finance Algorithms</i>	Binomial Option Pricing, Binomial Option Pricing with libNVRTC, Monte Carlo Estimation of Pi (batch PRNG), Monte Carlo Estimation of Pi (batch QRNG), Monte Carlo Estimation of Pi (inline PRNG), Monte Carlo Estimation of Pi (inline QRNG) , Monte Carlo Option Pricing with Multi-GPU support, Monte Carlo Single Asian Option, Niederreiter Quasirandom Sequence Generator, Niederreiter Quasirandom Sequence Generator with libNVRTC, Sobol Quasirandom Number Generator
Data Parallel Algorithms	<i>Samples that show good usage of Data Parallel Algorithms</i>	CUDA Histogram, CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA NvSciBuf/NvSciSync Interop, Mandelbrot, NvMedia CUDA Interop, Optical Flow, Particles, Smoke Particles, VFlockingD3D10, Vulkan CUDA Interop PI Approximation, Vulkan CUDA Interop Sinewave, Vulkan Image - CUDA Interop, cuDLA Error Reporting, cuDLA Hybrid Mode, cuDLA Standalone Mode
Data-Parallel Algorithms	<i>Samples that show good usage of Data Parallel Algorithms</i>	CUDA Parallel Prefix Sum (Scan), CUDA Parallel Prefix Sum with Shuffle Intrinsics (SHFL_Scan), CUDA Parallel Reduction, CUDA Radix Sort (Thrust Library), CUDA Segmentation Tree Thrust Library, CUDA

Advanced Key Concept	Description	Samples
		Sorting Networks, Fast Walsh Transform, Merge Sort, threadFenceReduction
Graphics Interop	<i>Samples that demonstrate interop between graphics APIs and CUDA</i>	Bindless Texture, CUDA FFT Ocean Simulation, CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGLES Version), Function Pointers, Mandelbrot, Particles, Post-Process in OpenGL, Recursive Gaussian Filter, SLI D3D10 Texture, Smoke Particles, Sobel Filter, VFlockingD3D10, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes, Vulkan CUDA Interop PI Approximation, Vulkan CUDA Interop Sinewave, Vulkan Image - CUDA Interop
Image Compression	<i>Samples that demonstrate image and video compression</i>	DirectX Texture Compressor (DXTC)
Image Processing	<i>Samples that demonstrate image processing algorithms in CUDA</i>	1D Discrete Haar Wavelet Decomposition, CUDA FFT Ocean Simulation, CUDA Histogram, CUDA NvSciBuf/NvSciSync Interop, DCT8x8, DirectX Texture Compressor (DXTC), FFT-Based 2D Convolution, Function Pointers, Image denoising, NvMedia CUDA Interop, Optical Flow, Post-Process in OpenGL, Recursive Gaussian Filter, SLI D3D10 Texture, Sobel Filter, Stereo Disparity Computation (SAD SIMD Intrinsics), Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes, cuDLA Error Reporting, cuDLA Hybrid Mode, cuDLA Standalone Mode
Linear Algebra	<i>Samples demonstrating linear algebra with CUDA</i>	Conjugate Gradient using Cuda Graphs, ConjugateGradient, ConjugateGradientUM, Eigenvalues, Fast Walsh Transform, Matrix Transpose, Preconditioned

Advanced Key Concept	Description	Samples
		Conjugate Gradient, Scalar Product, conjugateGradient using MultiBlock Cooperative Groups, conjugateGradient using MultiDevice Cooperative Groups
MultiBlock Cooperative Groups	<i>Multi Block Cooperative Groups enables to express inter-thread-block synchronization.</i>	Reduction using MultiBlock Cooperative Groups, conjugateGradient using MultiBlock Cooperative Groups
MultiDevice Cooperative Groups	<i>Multi Device Cooperative Groups enables thread blocks executing on multiple GPUs to cooperate and synchronize as they execute.</i>	conjugateGradient using MultiDevice Cooperative Groups
OpenGL Graphics Interop	<i>Samples demonstrating how to use interoperability CUDA with OpenGL</i>	Marching Cubes Isosurfaces
Performance Strategies	<i>Samples demonstrating high performance with CUDA</i>	Aligned Types, CUDA C 3D FDTD, CUDA Parallel Prefix Sum (Scan), CUDA Parallel Prefix Sum with Shuffle Ininsics (SHFL_Scan), CUDA Parallel Reduction, CUDA Radix Sort (Thrust Library), CUDA Segmentation Tree Thrust Library, Concurrent Kernels, Matrix Transpose, Particles, SLI D3D10 Texture, VFlockingD3D10, simpleHyperQ, threadFenceReduction
Physically Based Simulation	<i>Samples demonstrating high performance collisions and/or physocal interactions</i>	Marching Cubes Isosurfaces
Physically-Based Simulation	<i>Samples demonstrating high performance collisions and/or physocal interactions</i>	CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGLES Version), Particles, Smoke Particles, VFlockingD3D10
Random Number Generator	<i>Samples demonstrating how to use random number generation with CUDA</i>	Monte Carlo Estimation of Pi (batch PRNG), Monte Carlo Estimation of Pi (batch QRNG), Monte Carlo Estimation of Pi (inline PRNG), Monte Carlo Estimation of Pi (inline QRNG) , Monte Carlo Option

Advanced Key Concept	Description	Samples
		Pricing with Multi-GPU support, Monte Carlo Single Asian Option
Recursion	<i>Samples demonstrating recursion on CUDA</i>	Interval Computing
Runtime Compilation	<i>Samples demonstrating how to use NVRTC APIs for runtime compilation of CUDA Kernels</i>	Binomial Option Pricing with libNVRTC, Black-Scholes Option Pricing with libNVRTC, Clock libNVRTC, Matrix Multiplication with libNVRTC, Niederreiter Quasirandom Sequence Generator with libNVRTC, Simple Atomic Intrinsic with libNVRTC, Simple Templates with libNVRTC, Simple Vote Intrinsic with libNVRTC, Using Inline PTX with libNVRTC, Vector Addition with libNVRTC, simpleAssert with libNVRTC
Surface Writes	<i>Samples demonstrating how to use Surface Writes with GPU kernels</i>	Volumetric Filtering with 3D Textures and Surface Writes
Templates	<i>Samples demonstrating how to use templates GPU kernels</i>	Interval Computing
Tensor Cores	<i>Samples demonstrating use of Tensor Cores, introduced in the Volta chip family. Useful for faster matrix operations.</i>	CUDA Tensor Core GEMM, Double Precision Tensor Core GEMM, Tensor Core GEMM Integer MMA, bfloat16 Tensor Core GEMM, tf32 Tensor Core GEMM
Texture	<i>Samples demonstrating how to use textures GPU kernels</i>	Bindless Texture
Thrust Library		Line of Sight
Vertex Buffers	<i>Samples demonstrating how to use Vertex Buffers with CUDA kernels</i>	Marching Cubes Isosurfaces
Video Compression	<i>Samples demonstrating how to use video compression with CUDA</i>	1D Discrete Haar Wavelet Decomposition, DCT8x8, Fast Walsh Transform
Video Intrinsic	<i>Samples demonstrating how to use video intrinsic with CUDA</i>	Stereo Disparity Computation (SAD SIMD Intrinsic)

Advanced Key Concept	Description	Samples
WMMA	<i>Samples demonstrating how to use Warp Matrix Multiply and Accumulate (WMMA) CUDA APIs.</i>	CUDA Tensor Core GEMM, Double Precision Tensor Core GEMM, Tensor Core GEMM Integer MMA, bfloat16 Tensor Core GEMM, tf32 Tensor Core GEMM

Chapter 6. CUDA API and Associated Samples

The tables below list the samples associated with each CUDA API.

CUDA Driver API Samples

The table below lists the samples associated with each CUDA Driver API.

Table 3. CUDA Driver API and Associated Samples

CUDA Driver API	Samples
cuArrayCreate	Simple Texture (Driver Version)
cuArrayDestroy	EGLStream CUDA Interop, Simple Texture (Driver Version)
cuComplex	cuSolverDn Linear Solver , cuSolverRf Refactorization, cuSolverSp Linear Solver , cuSolverSp Lowlevel QR Solver, cuSolverSp LowlevelCholesky Solver
cuCtxCreate	CUDA Context Thread Management, CUDA Interception Library, EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU, EGLSync CUDA Event Interop, Matrix Multiplication (CUDA Driver API Version), Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), Memmap IPC Driver API, Simple Driver- Runtime Interaction, Simple Texture (Driver Version), Vector Addition Driver API, Vector Addition cuMemMap
cuCtxDestroy	CUDA Context Thread Management, CUDA Interception Library, Matrix Multiplication (CUDA Driver API Version), Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), Memmap IPC Driver API, Simple Driver- Runtime Interaction, Simple Texture (Driver Version), Vector Addition Driver API, Vector Addition cuMemMap

CUDA Driver API	Samples
cuCtxEnablePeerAccess	Memmap IPC Driver API
cuCtxGetDevice	CUDA Compressible Memory
cuCtxPopCurrent	CUDA Context Thread Management, EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU, SLI D3D10 Texture
cuCtxPushCurrent	CUDA Context Thread Management, EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU, EGLSync CUDA Event Interop, SLI D3D10 Texture
cuCtxSetCurrent	Memmap IPC Driver API
cuCtxSynchronize	Binomial Option Pricing with libNVRTC, Black-Scholes Option Pricing with libNVRTC, Clock libNVRTC, EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU, EGLSync CUDA Event Interop, Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), Matrix Multiplication with libNVRTC, Simple Atomic Intrinsic with libNVRTC, Simple Templates with libNVRTC, Simple Texture (Driver Version), Simple Vote Intrinsic with libNVRTC, Using Inline PTX with libNVRTC, Vector Addition Driver API, Vector Addition with libNVRTC, simpleAssert with libNVRTC
cuDeviceCanAccessPeer	Device Query Driver API, Memmap IPC Driver API, Vector Addition cuMemMap
cuDeviceComputeCapability	Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version)
cuDeviceGet	CUDA Context Thread Management, EGLStream_CUDA_CrossGPU, Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), Memmap IPC Driver API, stream Ordered Allocation IPC Pools
cuDeviceGetAttribute	CUDA Compressible Memory, CUDA Context Thread Management, Device Query, Device Query Driver API, EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU, EGLSync CUDA Event Interop, Matrix Multiplication (CUDA Driver API Version), Memmap IPC Driver API, Simple Texture (Driver Version), Vector Addition cuMemMap, stream Ordered Allocation IPC Pools
cuDeviceGetCount	CUDA Context Thread Management, CUDA Interception Library, Device Query Driver API, EGLStream CUDA Interop, Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), Memmap IPC Driver API, Vector Addition cuMemMap

CUDA Driver API	Samples
cuDeviceGetName	CUDA Context Thread Management, Device Query Driver API, EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU, Matrix Multiplication (CUDA Driver API Version), Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), Simple Texture (Driver Version)
cuDeviceGetUuid	CUDA NvSciBuf/NvSciSync Interop, NvMedia CUDA Interop
cuDeviceTotalMem	Device Query Driver API, Matrix Multiplication (CUDA Driver API Version)
cuDoubleComplex	cuSolverDn Linear Solver , cuSolverRf Refactorization, cuSolverSp Linear Solver , cuSolverSp Lowlevel QR Solver, cuSolverSp LowlevelCholesky Solver
cuDriverGetVersion	Device Query Driver API
cuEGLStreamConsumerAcquireFrame	EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU
cuEGLStreamConsumerConnect	EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU
cuEGLStreamConsumerDisconnect	EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU
cuEGLStreamConsumerReleaseFrame	EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU
cuEGLStreamProducerConnect	EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU
cuEGLStreamProducerDisconnect	EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU
cuEGLStreamProducerPresentFrame	EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU
cuEGLStreamProducerReturnFrame	EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU
cuEqual	batchCUBLAS
cuEventCreate	EGLSync CUDA Event Interop
cuEventCreateFromEGLSync	EGLSync CUDA Event Interop
cuEventDestroy	EGLSync CUDA Event Interop
cuEventRecord	EGLSync CUDA Event Interop
cuFuncSetBlockShape	Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version)
cuFuncSetSharedSize	Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version)

CUDA Driver API	Samples
cuGet	cuSolverRf Refactorization, cuSolverSp Lowlevel QR Solver
cuGraphicsEGLRegisterImage	EGLSync CUDA Event Interop
cuGraphicsResourceGetMappedEglFrame	EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU
cuGraphicsSubResourceGetMappedArray	EGLSync CUDA Event Interop
cuGraphicsUnregisterResource	EGLSync CUDA Event Interop
cuHook	CUDA Interception Library
cuHookInfo	CUDA Interception Library
cuHookRegisterCallback	CUDA Interception Library
culnit	CUDA Context Thread Management, CUDA Interception Library, Device Query Driver API, EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU, EGLSync CUDA Event Interop, Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), Memmap IPC Driver API, Simple Driver-Runtime Interaction, Vector Addition Driver API, Vector Addition cuMemMap
cuLaunchGrid	Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version)
cuLaunchKernel	Binomial Option Pricing with libNVRTC, Black-Scholes Option Pricing with libNVRTC, CUDA Context Thread Management, Clock libNVRTC, Matrix Multiplication (CUDA Driver API Version), Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), Matrix Multiplication with libNVRTC, Memmap IPC Driver API, PTX Just-in-Time compilation, Simple Atomic Intrinsic with libNVRTC, Simple Driver-Runtime Interaction, Simple Templates with libNVRTC, Simple Texture (Driver Version), Simple Vote Intrinsic with libNVRTC, Using Inline PTX with libNVRTC, Vector Addition Driver API, Vector Addition cuMemMap, Vector Addition with libNVRTC, simpleAssert with libNVRTC
cuLinkAddData	PTX Just-in-Time compilation
cuLinkComplete	PTX Just-in-Time compilation
cuLinkCreate	PTX Just-in-Time compilation
cuLinkDestroy	PTX Just-in-Time compilation

CUDA Driver API	Samples
cuMemAddressFree	CUDA Compressible Memory, Memmap IPC Driver API, Vector Addition cuMemMap, Vulkan CUDA Interop PI Approximation
cuMemAddressReserve	CUDA Compressible Memory, Memmap IPC Driver API, Vector Addition cuMemMap, Vulkan CUDA Interop PI Approximation
cuMemAlloc	Black-Scholes Option Pricing with libNVRTC, CUDA Context Thread Management, CUDA Interception Library, Clock libNVRTC, EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU, Matrix Multiplication (CUDA Driver API Version), Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), Matrix Multiplication with libNVRTC, Niederreiter Quasirandom Sequence Generator with libNVRTC, Simple Atomic Intrinsic with libNVRTC, Simple Templates with libNVRTC, Simple Texture (Driver Version), Simple Vote Intrinsic with libNVRTC, Using Inline PTX with libNVRTC, Vector Addition Driver API, Vector Addition with libNVRTC
cuMemCreate	CUDA Compressible Memory, Memmap IPC Driver API, Vector Addition cuMemMap, Vulkan CUDA Interop PI Approximation
cuMemExportToShareableHandle	Memmap IPC Driver API, Vulkan CUDA Interop PI Approximation
cuMemFree	Black-Scholes Option Pricing with libNVRTC, CUDA Context Thread Management, CUDA Interception Library, Clock libNVRTC, EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU, Matrix Multiplication (CUDA Driver API Version), Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), Matrix Multiplication with libNVRTC, Niederreiter Quasirandom Sequence Generator with libNVRTC, Simple Atomic Intrinsic with libNVRTC, Simple Templates with libNVRTC, Simple Texture (Driver Version), Simple Vote Intrinsic with libNVRTC, Using Inline PTX with libNVRTC, Vector Addition Driver API, Vector Addition with libNVRTC
cuMemGetAllocationGranularity	CUDA Compressible Memory, Memmap IPC Driver API, Vector Addition cuMemMap, Vulkan CUDA Interop PI Approximation
cuMemGetAllocationPropertiesFromHandle	CUDA Compressible Memory

CUDA Driver API	Samples
cuMemImportFromShareableHandle	Memmap IPC Driver API
cuMemMap	CUDA Compressible Memory, Memmap IPC Driver API, Vector Addition cuMemMap, Vulkan CUDA Interop PI Approximation
cuMemRelease	CUDA Compressible Memory, Memmap IPC Driver API, Vector Addition cuMemMap, Vulkan CUDA Interop PI Approximation
cuMemSetAccess	CUDA Compressible Memory, Memmap IPC Driver API, Vector Addition cuMemMap, Vulkan CUDA Interop PI Approximation
cuMemUnmap	CUDA Compressible Memory, Memmap IPC Driver API, Vector Addition cuMemMap, Vulkan CUDA Interop PI Approximation
cuMemcpy	EGLStream CUDA Interop
cuMemcpyDtoH	Binomial Option Pricing with libNVRTC, Black-Scholes Option Pricing with libNVRTC, CUDA Context Thread Management, Clock libNVRTC, EGLStream CUDA Interop, Matrix Multiplication (CUDA Driver API Version), Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), Matrix Multiplication with libNVRTC, Niederreiter Quasirandom Sequence Generator with libNVRTC, Simple Atomic Intrinsic with libNVRTC, Simple Templates with libNVRTC, Simple Texture (Driver Version), Simple Vote Intrinsic with libNVRTC, Using Inline PTX with libNVRTC, Vector Addition Driver API, Vector Addition cuMemMap, Vector Addition with libNVRTC
cuMemcpyDtoHAsync	Memmap IPC Driver API
cuMemcpyHtoD	Binomial Option Pricing with libNVRTC, Black-Scholes Option Pricing with libNVRTC, Clock libNVRTC, EGLStream_CUDA_CrossGPU, Matrix Multiplication (CUDA Driver API Version), Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), Matrix Multiplication with libNVRTC, Simple Atomic Intrinsic with libNVRTC, Simple Templates with libNVRTC, Simple Vote Intrinsic with libNVRTC, Vector Addition Driver API, Vector Addition cuMemMap, Vector Addition with libNVRTC
cuModuleGetFunction	Binomial Option Pricing with libNVRTC, Black-Scholes Option Pricing with libNVRTC, CUDA Context Thread

CUDA Driver API	Samples
	Management, Clock libNVRTC, Matrix Multiplication (CUDA Driver API Version), Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), Matrix Multiplication with libNVRTC, Memmap IPC Driver API, PTX Just-in-Time compilation, Simple Atomic Intrinsic with libNVRTC, Simple Driver-Runtime Interaction, Simple Templates with libNVRTC, Simple Texture (Driver Version), Simple Vote Intrinsic with libNVRTC, Using Inline PTX with libNVRTC, Vector Addition Driver API, Vector Addition cuMemMap, Vector Addition with libNVRTC, simpleAssert with libNVRTC
cuModuleGetGlobal	Binomial Option Pricing with libNVRTC
cuModuleLoad	Memmap IPC Driver API
cuModuleLoadData	CUDA Context Thread Management, Matrix Multiplication (CUDA Driver API Version), PTX Just-in-Time compilation, Simple Driver-Runtime Interaction, Simple Texture (Driver Version), Vector Addition Driver API, Vector Addition cuMemMap
cuModuleLoadDataEx	Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version), Memmap IPC Driver API
cuModuleUnload	CUDA Context Thread Management, PTX Just-in-Time compilation, Simple Driver-Runtime Interaction
cuOccupancyMaxActiveBlocksPerMultiprocessor	Memmap IPC Driver API
cuOccupancyMaxPotentialBlockSize	Matrix Multiplication (CUDA Driver API Version)
cuParamSetSize	Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version)
cuParamSeti	Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version)
cuParamSetv	Matrix Multiplication (CUDA Driver API version with Dynamic Linking Version)
cuRand	batchCUBLAS
cuSafeCallNoSync	Device Query
cuStreamCreate	EGLStream_CUDA_CrossGPU, EGLSync CUDA Event Interop, Memmap IPC Driver API
cuStreamDestroy	Memmap IPC Driver API
cuStreamSynchronize	Memmap IPC Driver API

CUDA Driver API	Samples
cuStreamWaitEvent	EGLSync CUDA Event Interop
cuSurfObjectCreate	EGLSync CUDA Event Interop
cuTexObjectCreate	Simple Texture (Driver Version)
cuTexObjectDestroy	Simple Texture (Driver Version)

CUDA Runtime API Samples

The table below lists the samples associated with each CUDA Runtime API.

Table 4. CUDA Runtime API and Associated Samples

CUDA Runtime API	Samples
cudaAcquireSync	Simple D3D11
cudaArrayGetInfo	Bindless Texture
cudaBindTextureToArray	Simple Direct3D10 Render Target
cudaBlockSize	Binomial Option Pricing with libNVRTC, Black-Scholes Option Pricing with libNVRTC, Clock libNVRTC, Simple Atomic Intrinsic with libNVRTC, Using Inline PTX, Using Inline PTX with libNVRTC, Vector Addition with libNVRTC
cudaCalculateSlopeKernel	CUDA FFT Ocean Simulation
cudaCheckError	Peer-to-Peer Bandwidth Latency Test with Multi-GPUs
cudaConsumerAcquireFrame	EGLStream_CUDA_CrossGPU
cudaConsumerReleaseFrame	EGLStream_CUDA_CrossGPU
cudaConsumerTest	EGLStream CUDA Interop
cudaCreateChannelDesc	Bicubic B-spline Interpolation, Box Filter, Line of Sight, Marching Cubes Isosurfaces, NvMedia CUDA Interop, Simple Cubemap Texture, Simple Layered Texture, Simple Surface Write, Simple Texture
cudaCreateSurfaceObject	Bindless Texture, NvMedia CUDA Interop, Simple Surface Write, Volumetric Filtering with 3D Textures and Surface Writes, Vulkan Image - CUDA Interop

CUDA Runtime API	Samples
cudaCreateTextureObject	Bicubic B-spline Interpolation, Bilateral Filter, Bindless Texture, Box Filter, CUDA NvSciBuf/NvSciSync Interop, DCT8x8, FFT-Based 2D Convolution, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGLES Version), Function Pointers, Image denoising, Line of Sight, Marching Cubes Isosurfaces, NV12toBGRandResize, Pitch Linear Texture, Post-Process in OpenGL, Simple Cubemap Texture, Simple Layered Texture, Simple Surface Write, Simple Texture, Simple Texture 3D, Smoke Particles, Sobel Filter, Stereo Disparity Computation (SAD SIMD Intrinsics), Texture-based Separable Convolution, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes, Vulkan Image - CUDA Interop
cudaDestroyExternalMemory	CUDA NvSciBuf/NvSciSync Interop, NvMedia CUDA Interop, Simple D3D11, Simple D3D12 CUDA Interop, Vulkan CUDA Interop Sinewave, Vulkan Image - CUDA Interop
cudaDestroyExternalSemaphore	CUDA NvSciBuf/NvSciSync Interop, NvMedia CUDA Interop, Simple D3D11, Simple D3D12 CUDA Interop, Vulkan CUDA Interop PI Approximation, Vulkan CUDA Interop Sinewave, Vulkan Image - CUDA Interop
cudaDestroySurfaceObject	Bindless Texture, NvMedia CUDA Interop, Simple Surface Write, Volumetric Filtering with 3D Textures and Surface Writes, Vulkan Image - CUDA Interop
cudaDestroyTextureObject	Bicubic B-spline Interpolation, Bilateral Filter, Bindless Texture, Box Filter, CUDA NvSciBuf/NvSciSync Interop, DCT8x8, FFT-Based 2D Convolution, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGLES Version), Function Pointers, Marching Cubes Isosurfaces, NV12toBGRandResize, Pitch Linear Texture, Post-Process in OpenGL, Simple Cubemap Texture, Simple Layered Texture, Simple Surface Write, Simple Texture, Simple Texture 3D, Sobel Filter, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes, Vulkan Image - CUDA Interop
cudaDeviceCanAccessPeer	CUDA N-Body Simulation, Device Query, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Simple Peer-to-Peer Transfers with Multi-GPU, conjugateGradient using MultiDevice Cooperative Groups, simpleIPC, stream Ordered Allocation IPC Pools, stream Ordered Allocation Peer-to-Peer access

CUDA Runtime API	Samples
cudaDeviceCreateConsumer	EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU
cudaDeviceCreateProducer	EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU
cudaDeviceDisablePeerAccess	Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Simple Peer-to-Peer Transfers with Multi-GPU
cudaDeviceEnablePeerAccess	Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Simple Peer-to-Peer Transfers with Multi-GPU, conjugateGradient using MultiDevice Cooperative Groups, simpleIPC, stream Ordered Allocation IPC Pools
cudaDeviceGetAttribute	Batched Label Markers And Label Compression NPP, CUDA NvSciBuf/NvSciSync Interop, Global Memory to Shared Memory Async Copy, Graph Memory Footprint, Graph Memory Nodes, Simple Arrive Wait Barrier, Topology Query, Vulkan CUDA Interop PI Approximation, Warp Aggregated Atomics using Cooperative Groups, Watershed Segmentation NPP, stream Ordered Allocation, stream Ordered Allocation IPC Pools, stream Ordered Allocation Peer-to-Peer access
cudaDeviceGetDefaultMemPool	stream Ordered Allocation, stream Ordered Allocation Peer-to-Peer access
cudaDeviceGetGraphMemAttribute	Graph Memory Footprint
cudaDeviceGetNvSciSyncAttributes	CUDA NvSciBuf/NvSciSync Interop, NvMedia CUDA Interop
cudaDeviceGetStreamPriorityRange	Stream Priorities
cudaDeviceGraphMemTrim	Graph Memory Footprint
cudaDeviceId	CUDA NvSciBuf/NvSciSync Interop
cudaDeviceInit	Canny Edge Detector NPP, Filter Border Control NPP, FreedImage and NPP Interoperability, Histogram Equalization with NPP
cudaDeviceReset	CUDA Interception Library, Filter Border Control NPP
cudaDeviceSetLimit	Interval Computing, NewDelete, Quad Tree (CUDA Dynamic Parallelism), Simple Print (CUDA Dynamic Parallelism), Simple Quicksort (CUDA Dynamic Parallelism), simpleAttributes
cudaDeviceSynchronize	Advanced Quicksort (CUDA Dynamic Parallelism), Aligned Types, Bandwidth Test, Bicubic B-spline Interpolation, Bilateral Filter, Bindless Texture, Binomial Option Pricing,

CUDA Runtime API	Samples
	<p>Black-Scholes Option Pricing, Box Filter, CUDA C 3D FDTD, CUDA Histogram, CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA Parallel Prefix Sum (Scan), CUDA Parallel Reduction, CUDA Separable Convolution, CUDA Sorting Networks, CUDA and OpenGL Interop of Images, ConjugateGradient, ConjugateGradientUM, DCT8x8, DirectX Texture Compressor (DXTC), EGLStream_CUDA_CrossGPU, EGLSync CUDA Event Interop, Eigenvalues, FFT-Based 2D Convolution, Fast Walsh Transform, Function Pointers, Image denoising, Interval Computing, Line of Sight, Mandelbrot, Merge Sort, Monte Carlo Option Pricing with Multi-GPU support, NV12toBGRandResize, NVJPEG Encoder, NewDelete, Niederreiter Quasirandom Sequence Generator, Particles, Post-Process in OpenGL, Recursive Gaussian Filter, Reduction using MultiBlock Cooperative Groups, Scalar Product, Simple CUFFT_MGPU, Simple Cooperative Groups, Simple Cubemap Texture, Simple Layered Texture, Simple Multi Copy and Compute, Simple OpenGL, Simple OpenGLES, Simple OpenGLES EGLOutput, Simple OpenGLES on Screen, Simple Peer-to-Peer Transfers with Multi-GPU, Simple Print (CUDA Dynamic Parallelism), Simple Quicksort (CUDA Dynamic Parallelism), Simple Surface Write, Simple Texture, Simple Texture 3D, Simple Vote Intrinsics, SimpleCUFFT_2d_MGPU, Sobel Filter, Sobol Quasirandom Number Generator, Stereo Disparity Computation (SAD SIMD Intrinsics), System wide Atomics, Texture-based Separable Convolution, Unified Memory Streams, Using Inline PTX, VFlockingD3D10, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes, asyncAPI, batchCUBLAS, conjugateGradient using MultiBlock Cooperative Groups, cppOverload, cuSolverDn Linear Solver , cuSolverRf Refactorization, cuSolverSp Linear Solver , simpleAssert, simpleOccupancy, simplePrintf, simpleZeroCopy, threadFenceReduction</p>
<p>cudaDriverGetVersion</p>	<p>Batched Label Markers And Label Compression NPP, Box Filter with NPP, Canny Edge Detector NPP, Device Query, Filter Border Control NPP, FreeImage and NPP Interopability, Graph Memory Footprint, Graph Memory Nodes, Histogram Equalization with NPP, PTX Just-in-Time compilation, Watershed Segmentation NPP</p>

CUDA Runtime API	Samples
<p>cudaEventCreate</p>	<p>Advanced Quicksort (CUDA Dynamic Parallelism), Bandwidth Test, CUDA C 3D FDTD, CUDA Compressible Memory, CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA Parallel Prefix Sum with Shuffle Intrinsics (SHFL_Scan), CUDA Radix Sort (Thrust Library), CUDA Segmentation Tree Thrust Library, CUDA Tensor Core GEMM, Concurrent Kernels, Double Precision Tensor Core GEMM, Global Memory to Shared Memory Async Copy, Graph Memory Nodes, Interval Computing, Matrix Multiplication (CUBLAS), Matrix Multiplication (CUDA Runtime API Version), Matrix Transpose, Monte Carlo Option Pricing with Multi-GPU support, NV12toBGRandResize, NVJPEG Encoder, NVJPEG simple, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Pitch Linear Texture, Simple CUDA Graphs, Simple Multi Copy and Compute, Stereo Disparity Computation (SAD SIMD Intrinsics), Stream Priorities, Tensor Core GEMM Integer MMA, VFlockingD3D10, asyncAPI, bfloat16 Tensor Core GEMM, conjugateGradient using MultiBlock Cooperative Groups, simpleHyperQ, simpleIPC, simpleOccupancy, stream Ordered Allocation, stream Ordered Allocation Peer-to-Peer access, tf32 Tensor Core GEMM</p>
<p>cudaEventCreateWithFlags</p>	<p>Concurrent Kernels, Simple Peer-to-Peer Transfers with Multi-GPU, simpleStreams</p>
<p>cudaEventDestroy</p>	<p>Bandwidth Test, CUDA C 3D FDTD, CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA Radix Sort (Thrust Library), Concurrent Kernels, Global Memory to Shared Memory Async Copy, Graph Memory Nodes, Interval Computing, Matrix Multiplication (CUDA Runtime API Version), Matrix Transpose, Monte Carlo Option Pricing with Multi-GPU support, NV12toBGRandResize, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Pitch Linear Texture, Simple Multi Copy and Compute, Simple Peer-to-Peer Transfers with Multi-GPU, asyncAPI, conjugateGradient using MultiBlock Cooperative Groups, simpleHyperQ, simpleIPC, simpleStreams</p>
<p>cudaEventElapsedTime</p>	<p>Advanced Quicksort (CUDA Dynamic Parallelism), Bandwidth Test, CUDA C 3D FDTD, CUDA Compressible Memory, CUDA N-Body Simulation, CUDA N-Body</p>

CUDA Runtime API	Samples
	<p>Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA Parallel Prefix Sum with Shuffle Intrinsics (SHFL_Scan), CUDA Radix Sort (Thrust Library), CUDA Segmentation Tree Thrust Library, CUDA Tensor Core GEMM, Concurrent Kernels, Double Precision Tensor Core GEMM, Global Memory to Shared Memory Async Copy, Interval Computing, Matrix Multiplication (CUBLAS), Matrix Multiplication (CUDA Runtime API Version), Matrix Transpose, NV12toBGRandResize, NVJPEG Encoder, NVJPEG simple, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Pitch Linear Texture, Simple Multi Copy and Compute, Simple Peer-to-Peer Transfers with Multi-GPU, Stereo Disparity Computation (SAD SIMD Intrinsics), Stream Priorities, Tensor Core GEMM Integer MMA, VFlockingD3D10, asyncAPI, bfloat16 Tensor Core GEMM, conjugateGradient using MultiBlock Cooperative Groups, simpleHyperQ, simpleOccupancy, simpleStreams, stream Ordered Allocation, tf32 Tensor Core GEMM</p>
<p>cudaEventQuery</p>	<p>asyncAPI</p>
<p>cudaEventRecord</p>	<p>Advanced Quicksort (CUDA Dynamic Parallelism), Bandwidth Test, CUDA C 3D FDTD, CUDA Compressible Memory, CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA Parallel Prefix Sum with Shuffle Intrinsics (SHFL_Scan), CUDA Radix Sort (Thrust Library), CUDA Segmentation Tree Thrust Library, CUDA Tensor Core GEMM, Concurrent Kernels, Double Precision Tensor Core GEMM, Global Memory to Shared Memory Async Copy, Graph Memory Nodes, Interval Computing, Matrix Multiplication (CUBLAS), Matrix Multiplication (CUDA Runtime API Version), Matrix Transpose, Monte Carlo Option Pricing with Multi-GPU support, NV12toBGRandResize, NVJPEG Encoder, NVJPEG simple, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Pitch Linear Texture, Simple CUDA Graphs, Simple Multi Copy and Compute, Simple Peer-to-Peer Transfers with Multi-GPU, Stereo Disparity Computation (SAD SIMD Intrinsics), Stream Priorities, Tensor Core GEMM Integer MMA, VFlockingD3D10, asyncAPI, bfloat16 Tensor Core GEMM, conjugateGradient using MultiBlock Cooperative Groups, simpleHyperQ, simpleIPC, simpleOccupancy, simpleStreams, stream</p>

CUDA Runtime API	Samples
	Ordered Allocation, stream Ordered Allocation Peer-to-Peer access, tf32 Tensor Core GEMM
cudaEventSynchronize	CUDA Compressible Memory, CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA Parallel Prefix Sum with Shuffle Intrinsic (SHFL_Scan), CUDA Radix Sort (Thrust Library), CUDA Segmentation Tree Thrust Library, CUDA Tensor Core GEMM, Concurrent Kernels, Double Precision Tensor Core GEMM, Global Memory to Shared Memory Async Copy, Matrix Multiplication (CUBLAS), Matrix Multiplication (CUDA Runtime API Version), Matrix Transpose, Monte Carlo Option Pricing with Multi-GPU support, NV12toBGRandResize, NVJPEG Encoder, NVJPEG simple, Pitch Linear Texture, Simple Multi Copy and Compute, Simple Peer-to-Peer Transfers with Multi-GPU, Stereo Disparity Computation (SAD SIMD Intrinsic), Stream Priorities, Tensor Core GEMM Integer MMA, VFlockingD3D10, bfloat16 Tensor Core GEMM, simpleHyperQ, simpleIPC, simpleStreams, stream Ordered Allocation, tf32 Tensor Core GEMM
cudaExtent	Bindless Texture, Jacobi CUDA Graphs, Simple CUDA Graphs, Simple Cubemap Texture, Simple Layered Texture, Simple Texture 3D, Smoke Particles, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes, Vulkan Image - CUDA Interop
cudaExternalMemoryGetMappedBuffer	CUDA NvSciBuf/NvSciSync Interop, Simple D3D11, Simple D3D12 CUDA Interop, Vulkan CUDA Interop Sinewave
cudaExternalMemoryGetMappedMipmappedArray	CUDA NvSciBuf/NvSciSync Interop, NvMedia CUDA Interop, Vulkan Image - CUDA Interop
cudaFree	1D Discrete Haar Wavelet Decomposition, Advanced Quicksort (CUDA Dynamic Parallelism), Aligned Types, Bandwidth Test, Batched Label Markers And Label Compression NPP, Bezier Line Tessellation (CUDA Dynamic Parallelism), Bicubic B-spline Interpolation, Bilateral Filter, Binary Partition Cooperative Groups, Bindless Texture, Black-Scholes Option Pricing, Box Filter, C++ Integration, C++11 CUDA, CUDA C 3D FDTD, CUDA FFT Ocean Simulation, CUDA Histogram, CUDA Interception Library, CUDA NvSciBuf/NvSciSync Interop, CUDA Parallel Prefix Sum (Scan), CUDA Parallel Prefix Sum with Shuffle Intrinsic (SHFL_Scan), CUDA Parallel

CUDA Runtime API	Samples
	<p>Reduction, CUDA Separable Convolution, CUDA Sorting Networks, CUDA Tensor Core GEMM, CUDA and OpenGL Interop of Images, Canny Edge Detector NPP, Clock, Concurrent Kernels, Conjugate Gradient using Cuda Graphs, ConjugateGradient, ConjugateGradientUM, DCT8x8, DirectX Texture Compressor (DXTC), Double Precision Tensor Core GEMM, EGLStream_CUDA_CrossGPU, EGLSync CUDA Event Interop, Eigenvalues, FFT-Based 2D Convolution, FP16 Scalar Product, Fast Walsh Transform, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGLES Version), Function Pointers, Global Memory to Shared Memory Async Copy, Graph Memory Footprint, Graph Memory Nodes, Histogram Equalization with NPP, Image denoising, Interval Computing, Jacobi CUDA Graphs, Mandelbrot, Marching Cubes Isosurfaces, Matrix Multiplication (CUBLAS), Matrix Multiplication (CUDA Runtime API Version), Matrix Transpose, Merge Sort, MersenneTwisterGP11213, Monte Carlo Estimation of Pi (batch PRNG), Monte Carlo Estimation of Pi (batch QRNG), Monte Carlo Estimation of Pi (inline PRNG), Monte Carlo Estimation of Pi (inline QRNG) , Monte Carlo Option Pricing with Multi-GPU support, Monte Carlo Single Asian Option, NV12toBGRandResize, NVJPEG Encoder, NVJPEG simple, NewDelete, Niederreiter Quasirandom Sequence Generator, NvMedia CUDA Interop, Optical Flow, PTX Just-in-Time compilation, Particles, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Pitch Linear Texture, Post-Process in OpenGL, Preconditioned Conjugate Gradient, Quad Tree (CUDA Dynamic Parallelism), Random Fog, Recursive Gaussian Filter, Reduction using MultiBlock Cooperative Groups, SLI D3D10 Texture, Scalar Product, Simple Arrive Wait Barrier, Simple Atomic Intrinsic, Simple CUBLAS, Simple CUBLAS LU, Simple CUBLAS XT, Simple CUDA Callbacks, Simple CUDA Graphs, Simple CUFFT, Simple CUFFT Callbacks, Simple Cubemap Texture, Simple D3D10 Texture, Simple D3D11, Simple D3D11 Texture, Simple D3D12 CUDA Interop, Simple D3D9 Texture, Simple Driver- Runtime Interaction, Simple Layered Texture, Simple Multi Copy and Compute, Simple Multi-GPU, Simple OpenGL, Simple OpenGLES, Simple OpenGLES EGLOutput, Simple OpenGLES on Screen, Simple Peer-to-Peer Transfers with Multi-GPU, Simple Quicksort (CUDA Dynamic Parallelism),</p>

CUDA Runtime API	Samples
	<p>Simple Static GPU Device Library, Simple Surface Write, Simple Templates, Simple Texture, Simple Texture 3D, Simple Vote Intrinsic, SimpleCUFFT_2d_MGPU, Sobel Filter, Sobol Quasirandom Number Generator, Stereo Disparity Computation (SAD SIMD Intrinsic), System wide Atomics, Template, Tensor Core GEMM Integer MMA, Texture-based Separable Convolution, Unified Memory Streams, Unified and other CUDA Memories Performance, Using Inline PTX, VFlockingD3D10, Vector Addition, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes, Vulkan CUDA Interop PI Approximation, Vulkan Image - CUDA Interop, Warp Aggregated Atomics using Cooperative Groups, Watershed Segmentation NPP, asyncAPI, batchCUBLAS, bfloat16 Tensor Core GEMM, conjugateGradient using MultiBlock Cooperative Groups, conjugateGradient using MultiDevice Cooperative Groups, cppOverload, cuDLA Error Reporting, cuDLA Hybrid Mode, cuSolverDn Linear Solver , cuSolverRf Refactorization, cuSolverSp Linear Solver , cuSolverSp Lowlevel QR Solver, cuSolverSp LowlevelCholesky Solver, cudaOpenMP, simpleAttributes, simpleHyperQ, simpleIPC, simpleMPI, simpleOccupancy, simpleStreams, tf32 Tensor Core GEMM, threadFenceReduction</p>
<p>cudaFreeArray</p>	<p>Bicubic B-spline Interpolation, Bindless Texture, Box Filter, DCT8x8, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGLES Version), Function Pointers, Image denoising, Line of Sight, NvMedia CUDA Interop, Pitch Linear Texture, Simple Cubemap Texture, Simple Layered Texture, Simple Surface Write, Simple Texture, Simple Texture 3D, Sobel Filter, Texture-based Separable Convolution, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes</p>
<p>cudaFreeAsync</p>	<p>Graph Memory Nodes, stream Ordered Allocation, stream Ordered Allocation IPC Pools, stream Ordered Allocation Peer-to-Peer access</p>
<p>cudaFreeHost</p>	<p>Bandwidth Test, Batched Label Markers And Label Compression NPP, Binary Partition Cooperative Groups, CUDA NvSciBuf/NvSciSync Interop, CUDA Parallel Prefix Sum with Shuffle Intrinsic (SHFL_Scan), Concurrent Kernels, Conjugate Gradient using Cuda Graphs, FP16 Scalar Product, Global Memory to Shared Memory</p>

CUDA Runtime API	Samples
	Async Copy, Jacobi CUDA Graphs, Matrix Multiplication (CUDA Runtime API Version), MersenneTwisterGP11213, Monte Carlo Option Pricing with Multi-GPU support, NVJPEG simple, NvMedia CUDA Interop, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Simple Arrive Wait Barrier, Simple Atomic Intrinsic, Simple CUDA Callbacks, Simple CUDA Graphs, Simple Driver-Runtime Interaction, Simple Multi Copy and Compute, Simple Multi-GPU, Simple Peer-to-Peer Transfers with Multi-GPU, Unified and other CUDA Memories Performance, Using Inline PTX, Vulkan CUDA Interop PI Approximation, asyncAPI, conjugateGradient using MultiDevice Cooperative Groups, cppOverload, simpleAttributes, simpleHyperQ, simpleStreams, simpleZeroCopy
cudaFreeMipmappedArray	Bindless Texture, CUDA NvSciBuf/NvSciSync Interop, NvMedia CUDA Interop, Vulkan Image - CUDA Interop
cudaFuncGetAttributes	CUDA C 3D FDTD, Monte Carlo Estimation of Pi (batch PRNG), Monte Carlo Estimation of Pi (batch QRNG), Monte Carlo Estimation of Pi (inline PRNG), Monte Carlo Estimation of Pi (inline QRNG) , Monte Carlo Single Asian Option, cppOverload
cudaFuncSetAttribute	CUDA Tensor Core GEMM, Double Precision Tensor Core GEMM, Tensor Core GEMM Integer MMA, bfloat16 Tensor Core GEMM, tf32 Tensor Core GEMM
cudaFuncSetCacheConfig	Interval Computing, cppOverload
cudaGLInit	Particles
cudaGLMapBufferObject	Mandelbrot, Marching Cubes Isosurfaces
cudaGLRegisterBufferObject	Image denoising, Mandelbrot, Marching Cubes Isosurfaces
cudaGLUnmapBufferObject	Mandelbrot, Marching Cubes Isosurfaces
cudaGLUnregisterBufferObject	Mandelbrot, Marching Cubes Isosurfaces
cudaGenerateSpectrumKernel	CUDA FFT Ocean Simulation
cudaGetChannelDesc	Post-Process in OpenGL
cudaGetDevice	Batched Label Markers And Label Compression NPP, CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA Parallel Prefix Sum with Shuffle Intrinsic (SHFL_Scan),

CUDA Runtime API	Samples
	<p>CUDA Parallel Reduction, CUDA Radix Sort (Thrust Library), Concurrent Kernels, DirectX Texture Compressor (DXTC), Fluids (Direct3D Version), Matrix Transpose, Recursive Gaussian Filter, Simple CUFFT Callbacks, Sobol Quasirandom Number Generator, Watershed Segmentation NPP, batchCUBLAS, cudaOpenMP, simpleHyperQ, simpleOccupancy, simplePrintf</p>
<p>cudaGetDeviceCount</p>	<p>Bandwidth Test, Bezier Line Tessellation (CUDA Dynamic Parallelism), CUDA C 3D FDTD, CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA NvSciBuf/NvSciSync Interop, Canny Edge Detector NPP, Device Query, Filter Border Control NPP, FreedImage and NPP Interoperability, Histogram Equalization with NPP, Monte Carlo Estimation of Pi (batch PRNG), Monte Carlo Estimation of Pi (batch QRNG), Monte Carlo Estimation of Pi (inline PRNG), Monte Carlo Estimation of Pi (inline QRNG) , Monte Carlo Option Pricing with Multi-GPU support, Monte Carlo Single Asian Option, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, SLI D3D10 Texture, Simple CUBLAS XT, Simple CUDA Callbacks, Simple CUFFT_MGPU, Simple D3D10 Texture, Simple D3D11, Simple D3D11 Texture, Simple D3D12 CUDA Interop, Simple Direct3D10 (Vertex Array), Simple Direct3D10 Render Target, Simple Multi-GPU, Simple Peer-to-Peer Transfers with Multi-GPU, SimpleCUFFT_2d_MGPU, Topology Query, VFlockingD3D10, Vulkan CUDA Interop PI Approximation, Vulkan CUDA Interop Sinewave, Vulkan Image - CUDA Interop, conjugateGradient using MultiDevice Cooperative Groups, cppOverload, cudaOpenMP, simpleIPC, simpleStreams, simpleZeroCopy, stream Ordered Allocation IPC Pools, stream Ordered Allocation Peer-to-Peer access</p>
<p>cudaGetDeviceProperties</p>	<p>Advanced Quicksort (CUDA Dynamic Parallelism), Aligned Types, Bandwidth Test, Batched Label Markers And Label Compression NPP, Bezier Line Tessellation (CUDA Dynamic Parallelism), Bicubic B-spline Interpolation, Bilateral Filter, CUDA C 3D FDTD, CUDA FFT Ocean Simulation, CUDA Histogram, CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA Parallel Prefix Sum with Shuffle Intrinsics (SHFL_Scan), CUDA Parallel Reduction, CUDA Radix Sort (Thrust Library), CUDA Tensor Core GEMM, Canny Edge Detector NPP, Concurrent Kernels,</p>

CUDA Runtime API	Samples
	<p>Conjugate Gradient using Cuda Graphs, ConjugateGradient, ConjugateGradientUM, Device Query, DirectX Texture Compressor (DXTC), Double Precision Tensor Core GEMM, FP16 Scalar Product, Filter Border Control NPP, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGLES Version), FreeImage and NPP Interopability, Graph Memory Footprint, Histogram Equalization with NPP, Interval Computing, Mandelbrot, Matrix Multiplication (CUBLAS), Matrix Transpose, Monte Carlo Estimation of Pi (batch PRNG), Monte Carlo Estimation of Pi (batch QRNG), Monte Carlo Estimation of Pi (inline PRNG), Monte Carlo Estimation of Pi (inline QRNG) , Monte Carlo Option Pricing with Multi-GPU support, Monte Carlo Single Asian Option, NVJPEG Encoder, NVJPEG simple, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Preconditioned Conjugate Gradient, Quad Tree (CUDA Dynamic Parallelism), Recursive Gaussian Filter, Reduction using MultiBlock Cooperative Groups, SLI D3D10 Texture, Simple CUBLAS XT, Simple CUDA Callbacks, Simple CUFFT Callbacks, Simple CUFFT_MGPU, Simple Cubemap Texture, Simple D3D10 Texture, Simple D3D11 Texture, Simple D3D12 CUDA Interop, Simple Direct3D10 (Vertex Array), Simple Direct3D10 Render Target, Simple Layered Texture, Simple Multi Copy and Compute, Simple Peer-to-Peer Transfers with Multi-GPU, Simple Print (CUDA Dynamic Parallelism), Simple Quicksort (CUDA Dynamic Parallelism), Simple Surface Write, Simple Templates, Simple Vote Intrinsic, SimpleCUFFT_2d_MGPU, Sobol Quasirandom Number Generator, Stereo Disparity Computation (SAD SIMD Intrinsic), Stream Priorities, System wide Atomics, Tensor Core GEMM Integer MMA, Unified Memory Streams, Unified and other CUDA Memories Performance, VFlockingD3D10, Vulkan CUDA Interop PI Approximation, Vulkan CUDA Interop Sinewave, Vulkan Image - CUDA Interop, Watershed Segmentation NPP, asyncAPI, batchCUBLAS, bfloat16 Tensor Core GEMM, conjugateGradient using MultiBlock Cooperative Groups, conjugateGradient using MultiDevice Cooperative Groups, cppOverload, cudaOpenMP, simpleAttributes, simpleHyperQ, simpleIPC, simpleOccupancy, simplePrintf, simpleStreams, simpleZeroCopy, stream Ordered Allocation IPC Pools, tf32 Tensor Core GEMM, threadFenceReduction</p>

CUDA Runtime API	Samples
cudaGetErrorEnum	Simple CUBLAS LU
cudaGetErrorName	cuDLA Error Reporting, cuDLA Hybrid Mode
cudaGetErrorString	Advanced Quicksort (CUDA Dynamic Parallelism), Bandwidth Test, Box Filter, CUDA Tensor Core GEMM, Device Query, Double Precision Tensor Core GEMM, EGLStream_CUDA_CrossGPU, EGLSync CUDA Event Interop, Monte Carlo Estimation of Pi (batch PRNG), Monte Carlo Estimation of Pi (batch QRNG), Monte Carlo Estimation of Pi (inline PRNG), Monte Carlo Estimation of Pi (inline QRNG) , Monte Carlo Single Asian Option, NVJPEG Encoder, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Random Fog, SLI D3D10 Texture, Simple Cooperative Groups, Simple D3D10 Texture, Simple D3D11 Texture, Simple Direct3D10 (Vertex Array), Simple Direct3D10 Render Target, Sobel Filter, Sobol Quasirandom Number Generator, Tensor Core GEMM Integer MMA, VFlockingD3D10, Vector Addition, batchCUBLAS, bfloat16 Tensor Core GEMM, cudaOpenMP, simpleAssert, tf32 Tensor Core GEMM
cudaGetLastError	Advanced Quicksort (CUDA Dynamic Parallelism), Bindless Texture, CUDA Tensor Core GEMM, Double Precision Tensor Core GEMM, Fluids (Direct3D Version), Interval Computing, Matrix Transpose, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Quad Tree (CUDA Dynamic Parallelism), SLI D3D10 Texture, Simple D3D10 Texture, Simple D3D11 Texture, Simple D3D9 Texture, Simple Direct3D10 (Vertex Array), Simple Direct3D10 Render Target, Simple Direct3D9 (Vertex Arrays), Simple Print (CUDA Dynamic Parallelism), Simple Static GPU Device Library, Tensor Core GEMM Integer MMA, Using Inline PTX, VFlockingD3D10, Vector Addition, batchCUBLAS, bfloat16 Tensor Core GEMM, cudaOpenMP, simpleIPC, simpleMPI, stream Ordered Allocation IPC Pools, tf32 Tensor Core GEMM
cudaGetMipmappedArrayLevel	Bindless Texture, CUDA NvSciBuf/NvSciSync Interop, NvMedia CUDA Interop, Vulkan Image - CUDA Interop
cudaGetValueMismatch	EGLStream_CUDA_CrossGPU, EGLSync CUDA Event Interop
cudaGraphAddHostNode	Simple CUDA Graphs
cudaGraphAddKernelNode	Graph Memory Footprint, Graph Memory Nodes, Jacobi CUDA Graphs, Simple CUDA Graphs

CUDA Runtime API	Samples
cudaGraphAddMemAllocNode	Graph Memory Footprint, Graph Memory Nodes
cudaGraphAddMemFreeNode	Graph Memory Footprint, Graph Memory Nodes
cudaGraphAddMemcpyNode	Jacobi CUDA Graphs, Simple CUDA Graphs
cudaGraphAddMemsetNode	Jacobi CUDA Graphs, Simple CUDA Graphs
cudaGraphClone	Simple CUDA Graphs
cudaGraphCreate	Graph Memory Footprint, Graph Memory Nodes, Jacobi CUDA Graphs, Simple CUDA Graphs
cudaGraphDestroy	Conjugate Gradient using Cuda Graphs, Graph Memory Footprint, Graph Memory Nodes, Simple CUDA Graphs
cudaGraphExecDestroy	Conjugate Gradient using Cuda Graphs, Graph Memory Footprint, Graph Memory Nodes, Jacobi CUDA Graphs, Simple CUDA Graphs
cudaGraphExecKernelNodeSetParams	Jacobi CUDA Graphs
cudaGraphExecUpdate	Jacobi CUDA Graphs
cudaGraphGetNodes	Simple CUDA Graphs
cudaGraphInstantiate	Conjugate Gradient using Cuda Graphs, Graph Memory Footprint, Graph Memory Nodes, Jacobi CUDA Graphs, Simple CUDA Graphs
cudaGraphLaunch	Conjugate Gradient using Cuda Graphs, Graph Memory Footprint, Graph Memory Nodes, Jacobi CUDA Graphs, Simple CUDA Graphs
cudaGraphicsGLRegisterBuffer	Bicubic B-spline Interpolation, Bilateral Filter, Bindless Texture, Box Filter, CUDA FFT Ocean Simulation, CUDA and OpenGL Interop of Images, Fluids (OpenGL Version), Fluids (OpenGLES Version), Function Pointers, Image denoising, Mandelbrot, Marching Cubes Isosurfaces, Particles, Post-Process in OpenGL, Recursive Gaussian Filter, Simple OpenGL, Simple OpenGLES, Simple OpenGLES EGLOutput, Simple OpenGLES on Screen, Simple Texture 3D, Sobel Filter, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes
cudaGraphicsGLRegisterImage	CUDA and OpenGL Interop of Images, Post-Process in OpenGL
cudaGraphicsMapResources	Bicubic B-spline Interpolation, Bilateral Filter, Bindless Texture, Box Filter, CUDA FFT Ocean Simulation, CUDA

CUDA Runtime API	Samples
	<p>N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA and OpenGL Interop of Images, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGLES Version), Function Pointers, Image denoising, Mandelbrot, Marching Cubes Isosurfaces, Particles, Post-Process in OpenGL, Recursive Gaussian Filter, SLI D3D10 Texture, Simple D3D10 Texture, Simple D3D11 Texture, Simple D3D9 Texture, Simple Direct3D10 (Vertex Array), Simple Direct3D10 Render Target, Simple Direct3D9 (Vertex Arrays), Simple OpenGL, Simple OpenGLES, Simple OpenGLES EGLOutput, Simple OpenGLES on Screen, Simple Texture 3D, Sobel Filter, VFlockingD3D10, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes</p>
<p>cudaGraphicsResourceGetMappedPointer</p>	<p>Bicubic B-spline Interoplation, Bilateral Filter, Bindless Texture, Box Filter, CUDA FFT Ocean Simulation, CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA and OpenGL Interop of Images, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGLES Version), Function Pointers, Image denoising, Mandelbrot, Marching Cubes Isosurfaces, Particles, Post-Process in OpenGL, Recursive Gaussian Filter, Simple Direct3D10 (Vertex Array), Simple Direct3D10 Render Target, Simple Direct3D9 (Vertex Arrays), Simple OpenGL, Simple OpenGLES, Simple OpenGLES EGLOutput, Simple OpenGLES on Screen, Simple Texture 3D, Sobel Filter, VFlockingD3D10, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes</p>
<p>cudaGraphicsResourceSetMapFlags</p>	<p>CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, SLI D3D10 Texture, VFlockingD3D10</p>
<p>cudaGraphicsSubResourceGetMappedArray</p>	<p>CUDA and OpenGL Interop of Images, Post-Process in OpenGL, SLI D3D10 Texture, Simple D3D10 Texture, Simple D3D11 Texture, Simple D3D9 Texture, Simple Direct3D10 Render Target</p>
<p>cudaGraphicsUnmapResources</p>	<p>Bicubic B-spline Interoplation, Bilateral Filter, Bindless Texture, Box Filter, CUDA FFT Ocean Simulation, CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA and OpenGL Interop of Images, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGLES Version), Function</p>

CUDA Runtime API	Samples
	<p>Pointers, Image denoising, Mandelbrot, Marching Cubes Isosurfaces, Particles, Post-Process in OpenGL, Recursive Gaussian Filter, SLI D3D10 Texture, Simple D3D10 Texture, Simple D3D11 Texture, Simple D3D9 Texture, Simple Direct3D10 (Vertex Array), Simple Direct3D10 Render Target, Simple Direct3D9 (Vertex Arrays), Simple OpenGL, Simple OpenGL ES, Simple OpenGL ES EGLOutput, Simple OpenGL ES on Screen, Simple Texture 3D, Sobel Filter, VFlockingD3D10, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes</p>
<p>cudaGraphicsUnregisterResource</p>	<p>Bicubic B-spline Interpolation, Bilateral Filter, Bindless Texture, Box Filter, CUDA FFT Ocean Simulation, CUDA and OpenGL Interop of Images, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGL ES Version), Function Pointers, Image denoising, Mandelbrot, Marching Cubes Isosurfaces, Particles, Post-Process in OpenGL, Recursive Gaussian Filter, SLI D3D10 Texture, Simple D3D10 Texture, Simple D3D11 Texture, Simple D3D9 Texture, Simple Direct3D10 (Vertex Array), Simple Direct3D10 Render Target, Simple Direct3D9 (Vertex Arrays), Simple OpenGL, Simple OpenGL ES, Simple OpenGL ES EGLOutput, Simple OpenGL ES on Screen, Simple Texture 3D, Sobel Filter, VFlockingD3D10, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes</p>
<p>cudaGraphsManual</p>	<p>Simple CUDA Graphs</p>
<p>cudaGraphsUsingStreamCapture</p>	<p>Simple CUDA Graphs</p>
<p>cudaGridSize</p>	<p>Binomial Option Pricing with libNVRTC, Black-Scholes Option Pricing with libNVRTC, Clock libNVRTC, Simple Atomic Intrinsic with libNVRTC, Using Inline PTX, Using Inline PTX with libNVRTC, Vector Addition with libNVRTC</p>
<p>cudaHeightMap</p>	<p>Vulkan CUDA Interop Sinewave</p>
<p>cudaHostAlloc</p>	<p>Bandwidth Test, CUDA and OpenGL Interop of Images, NVJPEG simple, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Post-Process in OpenGL, Simple CUDA Callbacks, Simple Multi Copy and Compute, conjugateGradient using MultiDevice Cooperative Groups, simpleStreams, simpleZeroCopy</p>
<p>cudaHostGetDevicePointer</p>	<p>Unified and other CUDA Memories Performance, simpleZeroCopy</p>

CUDA Runtime API	Samples
cudaHostRegister	simpleStreams, simpleZeroCopy
cudaHostUnregister	simpleStreams, simpleZeroCopy
cudaImportExternalMemory	CUDA NvSciBuf/NvSciSync Interop, NvMedia CUDA Interop, Simple D3D11, Simple D3D12 CUDA Interop, Vulkan CUDA Interop Sinewave, Vulkan Image - CUDA Interop
cudaImportExternalSemaphore	CUDA NvSciBuf/NvSciSync Interop, NvMedia CUDA Interop, Simple D3D11, Simple D3D12 CUDA Interop, Vulkan CUDA Interop PI Approximation, Vulkan CUDA Interop Sinewave, Vulkan Image - CUDA Interop
cudaImportKeyedMutex	Simple D3D11
cudaImportNvSciImage	CUDA NvSciBuf/NvSciSync Interop, NvMedia CUDA Interop
cudaImportNvSciRawBuf	CUDA NvSciBuf/NvSciSync Interop
cudaImportNvSciSemaphore	CUDA NvSciBuf/NvSciSync Interop
cudaImportNvSciSync	NvMedia CUDA Interop
cudaImportVertexBuffer	Simple D3D11
cudaInit	Particles
cudaIpcCloseMemHandle	simpleIPC
cudaIpcGetEventHandle	simpleIPC
cudaIpcGetMemHandle	simpleIPC
cudaIpcOpenEventHandle	simpleIPC
cudaIpcOpenMemHandle	simpleIPC
cudaLaunchCooperativeKernel	Reduction using MultiBlock Cooperative Groups, Simple Arrive Wait Barrier, conjugateGradient using MultiBlock Cooperative Groups, conjugateGradient using MultiDevice Cooperative Groups
cudaLaunchHostFunc	Simple CUDA Graphs, Vulkan CUDA Interop PI Approximation
cudaMalloc	1D Discrete Haar Wavelet Decomposition, Advanced Quicksort (CUDA Dynamic Parallelism), Aligned Types, Bandwidth Test, Batched Label Markers And Label Compression NPP, Bezier Line Tessellation (CUDA Dynamic Parallelism), Bicubic B-spline Interpolation, Binary Partition

CUDA Runtime API	Samples
	<p>Cooperative Groups, Bindless Texture, Black-Scholes Option Pricing, Box Filter, C++ Integration, C++11 CUDA, CUDA C 3D FDTD, CUDA FFT Ocean Simulation, CUDA Histogram, CUDA Parallel Prefix Sum (Scan), CUDA Parallel Prefix Sum with Shuffle Intrinsic (SHFL_Scan), CUDA Parallel Reduction, CUDA Separable Convolution, CUDA Sorting Networks, CUDA Tensor Core GEMM, CUDA and OpenGL Interop of Images, Canny Edge Detector NPP, Clock, Concurrent Kernels, Conjugate Gradient using Cuda Graphs, ConjugateGradient, ConjugateGradientUM, DirectX Texture Compressor (DXTC), Double Precision Tensor Core GEMM, EGLStream_CUDA_CrossGPU, EGLSync CUDA Event Interop, Eigenvalues, FFT-Based 2D Convolution, FP16 Scalar Product, Fast Walsh Transform, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGLES Version), Function Pointers, Global Memory to Shared Memory Async Copy, Graph Memory Nodes, Histogram Equalization with NPP, Image denoising, Interval Computing, Jacobi CUDA Graphs, Mandelbrot, Marching Cubes Isosurfaces, Matrix Multiplication (CUBLAS), Matrix Multiplication (CUDA Runtime API Version), Matrix Transpose, Merge Sort, MersenneTwisterGP11213, Monte Carlo Estimation of Pi (batch PRNG), Monte Carlo Estimation of Pi (batch QRNG), Monte Carlo Estimation of Pi (inline PRNG), Monte Carlo Estimation of Pi (inline QRNG) , Monte Carlo Option Pricing with Multi-GPU support, Monte Carlo Single Asian Option, NV12toBGRandResize, NVJPEG Encoder, NVJPEG simple, NewDelete, Niederreiter Quasirandom Sequence Generator, NvMedia CUDA Interop, Optical Flow, PTX Just-in-Time compilation, Particles, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Post-Process in OpenGL, Preconditioned Conjugate Gradient, Quad Tree (CUDA Dynamic Parallelism), Random Fog, Recursive Gaussian Filter, Reduction using MultiBlock Cooperative Groups, SLI D3D10 Texture, Scalar Product, Simple Arrive Wait Barrier, Simple Atomic Intrinsic, Simple CUBLAS, Simple CUBLAS LU, Simple CUDA Callbacks, Simple CUDA Graphs, Simple CUFFT, Simple CUFFT Callbacks, Simple Cubemap Texture, Simple D3D10 Texture, Simple D3D11 Texture, Simple D3D9 Texture, Simple Direct3D10 Render Target, Simple Driver- Runtime Interaction, Simple Layered Texture, Simple Multi</p>

CUDA Runtime API	Samples
	<p>Copy and Compute, Simple Multi-GPU, Simple OpenGL, Simple OpenGLES, Simple OpenGLES EGLOutput, Simple OpenGLES on Screen, Simple Peer-to-Peer Transfers with Multi-GPU, Simple Quicksort (CUDA Dynamic Parallelism), Simple Static GPU Device Library, Simple Surface Write, Simple Templates, Simple Texture, Simple Texture 3D, Simple Vote Intrinsic, SimpleCUFFT_2d_MGPU, Sobel Filter, Sobol Quasirandom Number Generator, Stereo Disparity Computation (SAD SIMD Intrinsic), Stream Priorities, Template, Tensor Core GEMM Integer MMA, Texture-based Separable Convolution, Unified and other CUDA Memories Performance, Using Inline PTX, VFlockingD3D10, Vector Addition, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes, Vulkan CUDA Interop PI Approximation, Vulkan Image - CUDA Interop, Warp Aggregated Atomics using Cooperative Groups, asyncAPI, batchCUBLAS, bfloat16 Tensor Core GEMM, cppOverload, cuDLA Error Reporting, cuDLA Hybrid Mode, cuSolverDn Linear Solver , cuSolverRf Refactorization, cuSolverSp Linear Solver , cuSolverSp Lowlevel QR Solver, cuSolverSp LowlevelCholesky Solver, cudaOpenMP, simpleAttributes, simpleHyperQ, simpleIPC, simpleMPI, simpleOccupancy, simpleStreams, tf32 Tensor Core GEMM, threadFenceReduction</p>
<p>cudaMallocArray</p>	<p>Bicubic B-spline Interpolation, Bindless Texture, Box Filter, DCT8x8, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGLES Version), Function Pointers, Image denoising, Line of Sight, NvMedia CUDA Interop, Pitch Linear Texture, Simple Surface Write, Simple Texture, Sobel Filter, Texture-based Separable Convolution, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes</p>
<p>cudaMallocAsync</p>	<p>Graph Memory Nodes, stream Ordered Allocation, stream Ordered Allocation IPC Pools, stream Ordered Allocation Peer-to-Peer access</p>
<p>cudaMallocHost</p>	<p>Bandwidth Test, Batched Label Markers And Label Compression NPP, Binary Partition Cooperative Groups, CUDA NvSciBuf/NvSciSync Interop, CUDA Parallel Prefix Sum with Shuffle Intrinsic (SHFL_Scan), Concurrent Kernels, Conjugate Gradient using Cuda Graphs, FP16 Scalar Product, Global Memory to Shared Memory</p>

CUDA Runtime API	Samples
	<p>Async Copy, Jacobi CUDA Graphs, Matrix Multiplication (CUDA Runtime API Version), MersenneTwisterGP11213, Monte Carlo Option Pricing with Multi-GPU support, NvMedia CUDA Interop, Simple Arrive Wait Barrier, Simple Atomic Intrinsics, Simple CUDA Graphs, Simple Driver- Runtime Interaction, Simple Multi-GPU, Simple Peer-to-Peer Transfers with Multi-GPU, Unified and other CUDA Memories Performance, Using Inline PTX, Vulkan CUDA Interop PI Approximation, asyncAPI, cppOverload, simpleAttributes, simpleHyperQ, simpleStreams</p>
<p>cudaMallocManaged</p>	<p>ConjugateGradientUM, Graph Memory Nodes, NV12toBGRandResize, System wide Atomics, Unified Memory Streams, Unified and other CUDA Memories Performance, conjugateGradient using MultiBlock Cooperative Groups, conjugateGradient using MultiDevice Cooperative Groups</p>
<p>cudaMallocMipmappedArray</p>	<p>Bindless Texture, Vulkan Image - CUDA Interop</p>
<p>cudaMallocPitch</p>	<p>Batched Label Markers And Label Compression NPP, Bilateral Filter, DCT8x8, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGLES Version), Pitch Linear Texture, SLI D3D10 Texture, Simple D3D10 Texture, Simple D3D11 Texture, Simple D3D9 Texture</p>
<p>cudaMemAdvise</p>	<p>conjugateGradient using MultiDevice Cooperative Groups</p>
<p>cudaMemGetInfo</p>	<p>CUDA Segmentation Tree Thrust Library</p>
<p>cudaMemPoolCreate</p>	<p>stream Ordered Allocation IPC Pools</p>
<p>cudaMemPoolDestroy</p>	<p>stream Ordered Allocation IPC Pools</p>
<p>cudaMemPoolExportPointer</p>	<p>stream Ordered Allocation IPC Pools</p>
<p>cudaMemPoolExportToShareableHandle</p>	<p>stream Ordered Allocation IPC Pools</p>
<p>cudaMemPoolGetAccess</p>	<p>stream Ordered Allocation IPC Pools</p>
<p>cudaMemPoolImportFromShareableHandle</p>	<p>stream Ordered Allocation IPC Pools</p>
<p>cudaMemPoolImportPointer</p>	<p>stream Ordered Allocation IPC Pools</p>
<p>cudaMemPoolSetAccess</p>	<p>stream Ordered Allocation IPC Pools, stream Ordered Allocation Peer-to-Peer access</p>
<p>cudaMemPoolSetAttribute</p>	<p>stream Ordered Allocation</p>

CUDA Runtime API	Samples
<p>cudaMemPrefetchAsync</p>	<p>Unified and other CUDA Memories Performance, conjugateGradient using MultiDevice Cooperative Groups</p>
<p>cudaMemcpy</p>	<p>1D Discrete Haar Wavelet Decomposition, Advanced Quicksort (CUDA Dynamic Parallelism), Aligned Types, Bandwidth Test, Bezier Line Tessellation (CUDA Dynamic Parallelism), Bicubic B-spline Interpolation, Bindless Texture, Black-Scholes Option Pricing, Box Filter, C++ Integration, C++11 CUDA, CUDA C 3D FDTD, CUDA Compressible Memory, CUDA FFT Ocean Simulation, CUDA Histogram, CUDA Parallel Prefix Sum (Scan), CUDA Parallel Prefix Sum with Shuffle Intrinsic (SHFL_Scan), CUDA Parallel Reduction, CUDA Segmentation Tree Thrust Library, CUDA Separable Convolution, CUDA Sorting Networks, CUDA Tensor Core GEMM, Clock, ConjugateGradient, DirectX Texture Compressor (DXTC), Double Precision Tensor Core GEMM, EGLStream_CUDA_CrossGPU, EGLSync CUDA Event Interop, Eigenvalues, FFT-Based 2D Convolution, FP16 Scalar Product, Fast Walsh Transform, Fluids (Direct3D Version), Fluids (OpenGL Version), Fluids (OpenGLES Version), Function Pointers, Graph Memory Nodes, Histogram Equalization with NPP, Image denoising, Interval Computing, Mandelbrot, Marching Cubes Isosurfaces, Matrix Multiplication (CUBLAS), Matrix Transpose, Merge Sort, Monte Carlo Estimation of Pi (batch PRNG), Monte Carlo Estimation of Pi (batch QRNG), Monte Carlo Estimation of Pi (inline PRNG), Monte Carlo Estimation of Pi (inline QRNG) , Monte Carlo Single Asian Option, NV12toBGRandResize, NewDelete, Niederreiter Quasirandom Sequence Generator, Optical Flow, PTX Just-in-Time compilation, Particles, Preconditioned Conjugate Gradient, Quad Tree (CUDA Dynamic Parallelism), Random Fog, Recursive Gaussian Filter, Reduction using MultiBlock Cooperative Groups, Scalar Product, Simple CUBLAS LU, Simple CUFFT, Simple CUFFT Callbacks, Simple Cubemap Texture, Simple Direct3D10 Render Target, Simple Layered Texture, Simple OpenGL, Simple OpenGLES, Simple OpenGLES EGLOutput, Simple OpenGLES on Screen, Simple Peer-to-Peer Transfers with Multi-GPU, Simple Quicksort (CUDA Dynamic Parallelism), Simple Static GPU Device Library, Simple Surface Write, Simple Templates, Simple Texture, Simple Texture 3D, Simple Vote Intrinsic, SimpleCUFFT_2d_MGPU, Sobel Filter,</p>

CUDA Runtime API	Samples
	<p>Sobol Quasirandom Number Generator, Stereo Disparity Computation (SAD SIMD Intrinsics), Stream Priorities, Template, Tensor Core GEMM Integer MMA, Texture-based Separable Convolution, Unified and other CUDA Memories Performance, Using Inline PTX, VFlockingD3D10, Vector Addition, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes, Vulkan Image - CUDA Interop, Warp Aggregated Atomics using Cooperative Groups, batchCUBLAS, bfloat16 Tensor Core GEMM, cppOverload, cuSolverDn Linear Solver , cuSolverRf Refactorization, cuSolverSp Lowlevel QR Solver, cuSolverSp LowlevelCholesky Solver, cudaOpenMP, simpleHyperQ, simpleMPI, simpleOccupancy, simpleStreams, tf32 Tensor Core GEMM, threadFenceReduction</p>
<p>cudaMemcpyAsync</p>	<p>Advanced Quicksort (CUDA Dynamic Parallelism), Bandwidth Test, Batched Label Markers And Label Compression NPP, Binary Partition Cooperative Groups, CUDA NvSciBuf/NvSciSync Interop, Concurrent Kernels, Conjugate Gradient using Cuda Graphs, Global Memory to Shared Memory Async Copy, Graph Memory Nodes, Jacobi CUDA Graphs, Matrix Multiplication (CUDA Runtime API Version), MersenneTwisterGP11213, Monte Carlo Option Pricing with Multi-GPU support, NvMedia CUDA Interop, Simple Arrive Wait Barrier, Simple Atomic Intrinsics, Simple CUDA Callbacks, Simple CUDA Graphs, Simple Driver-Runtime Interaction, Simple Multi Copy and Compute, Simple Multi-GPU, Unified and other CUDA Memories Performance, Vulkan CUDA Interop PI Approximation, asyncAPI, cuDLA Error Reporting, cuDLA Hybrid Mode, cuSolverSp Linear Solver , simpleAttributes, simpleIPC, simpleStreams, stream Ordered Allocation, stream Ordered Allocation IPC Pools, stream Ordered Allocation Peer-to-Peer access</p>
<p>cudaMemcpyFromSymbol</p>	<p>Binomial Option Pricing, Function Pointers, Simple CUFFT Callbacks, Simple Static GPU Device Library</p>
<p>cudaMemcpyPeerAsync</p>	<p>Peer-to-Peer Bandwidth Latency Test with Multi-GPUs</p>
<p>cudaMemcpyToArray</p>	<p>CUDA and OpenGL Interop of Images, Pitch Linear Texture, Post-Process in OpenGL, Simple Surface Write, Simple Texture, Texture-based Separable Convolution</p>

CUDA Runtime API	Samples
<p>cudaMemcpyToSymbol</p>	<p>Bilateral Filter, Binomial Option Pricing, CUDA C 3D FDTD, CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA Separable Convolution, Function Pointers, Niederreiter Quasirandom Sequence Generator, Particles, Smoke Particles, Texture-based Separable Convolution, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes</p>
<p>cudaMemset</p>	<p>Advanced Quicksort (CUDA Dynamic Parallelism), Aligned Types, C++11 CUDA, CUDA Parallel Prefix Sum with Shuffle Intrinsics (SHFL_Scan), CUDA Segmentation Tree Thrust Library, CUDA Tensor Core GEMM, Double Precision Tensor Core GEMM, FFT-Based 2D Convolution, Fast Walsh Transform, Monte Carlo Option Pricing with Multi-GPU support, Niederreiter Quasirandom Sequence Generator, Optical Flow, Particles, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Preconditioned Conjugate Gradient, SLI D3D10 Texture, Simple D3D10 Texture, Simple D3D11 Texture, Simple D3D9 Texture, Simple Multi Copy and Compute, Tensor Core GEMM Integer MMA, VFlockingD3D10, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes, Warp Aggregated Atomics using Cooperative Groups, asyncAPI, bfloat16 Tensor Core GEMM, conjugateGradient using MultiDevice Cooperative Groups, cuSolverDn Linear Solver , cudaOpenMP, simpleStreams, tf32 Tensor Core GEMM</p>
<p>cudaMemsetAsync</p>	<p>Binary Partition Cooperative Groups, Conjugate Gradient using Cuda Graphs, Global Memory to Shared Memory Async Copy, Jacobi CUDA Graphs, Simple CUDA Graphs, Vulkan CUDA Interop PI Approximation, cuDLA Error Reporting, cuDLA Hybrid Mode</p>
<p>cudaNvSci</p>	<p>CUDA NvSciBuf/NvSciSync Interop</p>
<p>cudaNvSciApp</p>	<p>CUDA NvSciBuf/NvSciSync Interop</p>
<p>cudaNvSciSignal</p>	<p>CUDA NvSciBuf/NvSciSync Interop</p>
<p>cudaNvSciWait</p>	<p>CUDA NvSciBuf/NvSciSync Interop</p>
<p>cudaOccupancyMaxActiveBlocksPerMultiprocessor</p>	<p>Reduction using MultiBlock Cooperative Groups, Simple Arrive Wait Barrier, Vulkan CUDA Interop PI Approximation,</p>

CUDA Runtime API	Samples
	Vulkan CUDA Interop Sinewave, conjugateGradient using MultiBlock Cooperative Groups, conjugateGradient using MultiDevice Cooperative Groups, simpleIPC, simpleOccupancy, stream Ordered Allocation IPC Pools
cudaOccupancyMaxPotentialBlockSize	Binary Partition Cooperative Groups, CUDA Compressible Memory, Conjugate Gradient using Cuda Graphs, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Reduction using MultiBlock Cooperative Groups, Simple Arrive Wait Barrier, simpleOccupancy
cudaPeekAtLastError	Advanced Quicksort (CUDA Dynamic Parallelism)
cudaPitchedPtr	Bindless Texture, Jacobi CUDA Graphs, Simple CUDA Graphs, Simple Cubemap Texture, Simple Layered Texture, Simple Texture 3D, Smoke Particles, Volume Rendering with 3D Textures, Volumetric Filtering with 3D Textures and Surface Writes
cudaPos	Jacobi CUDA Graphs, Simple CUDA Graphs, Simple Cubemap Texture, Simple Layered Texture
cudaProcess	CUDA and OpenGL Interop of Images, Post-Process in OpenGL
cudaProducerDeinit	EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU
cudaProducerInit	EGLStream CUDA Interop, EGLStream_CUDA_CrossGPU
cudaProducerPrepareFrame	EGLStream_CUDA_CrossGPU
cudaProducerPresentFrame	EGLStream_CUDA_CrossGPU
cudaProducerReadARGBFrame	EGLStream CUDA Interop
cudaProducerReadYUVFrame	EGLStream CUDA Interop
cudaProducerReturnFrame	EGLStream_CUDA_CrossGPU
cudaProducerTest	EGLStream CUDA Interop
cudaProfilerStart	Matrix Multiplication (CUDA Runtime API Version), asyncAPI
cudaProfilerStop	Matrix Multiplication (CUDA Runtime API Version), Volume Rendering with 3D Textures, asyncAPI
cudaReleaseSync	Simple D3D11
cudaRuntimeGetVersion	Batched Label Markers And Label Compression NPP, Bilateral Filter, Box Filter with NPP, Canny Edge Detector

CUDA Runtime API	Samples
	NPP, Device Query, Filter Border Control NPP, FreeImage and NPP Interopability, Histogram Equalization with NPP, Watershed Segmentation NPP
cudaSetDevice	Bandwidth Test, CUDA C 3D FDTD, CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, CUDA NvSciBuf/NvSciSync Interop, CUDA Parallel Reduction, Canny Edge Detector NPP, Device Query, Filter Border Control NPP, FreeImage and NPP Interopability, Histogram Equalization with NPP, Interval Computing, Monte Carlo Estimation of Pi (batch PRNG), Monte Carlo Estimation of Pi (batch QRNG), Monte Carlo Estimation of Pi (inline PRNG), Monte Carlo Estimation of Pi (inline QRNG) , Monte Carlo Option Pricing with Multi-GPU support, Monte Carlo Single Asian Option, NvMedia CUDA Interop, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Reduction using MultiBlock Cooperative Groups, Simple CUDA Callbacks, Simple CUFFT_MGPU, Simple D3D11, Simple D3D12 CUDA Interop, Simple Multi Copy and Compute, Simple Multi-GPU, Simple Peer-to-Peer Transfers with Multi-GPU, SimpleCUFFT_2d_MGPU, Unified Memory Streams, Vulkan CUDA Interop PI Approximation, Vulkan CUDA Interop Sinewave, Vulkan Image - CUDA Interop, conjugateGradient using MultiDevice Cooperative Groups, cppOverload, cuDLA Error Reporting, cuDLA Hybrid Mode, cudaOpenMP, simpleIPC, simpleStreams, simpleZeroCopy, stream Ordered Allocation, stream Ordered Allocation IPC Pools, stream Ordered Allocation Peer-to-Peer access
cudaSetDeviceFlags	CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES, simpleStreams, simpleZeroCopy
cudaSignalExternalSemaphoresAsync	CUDA NvSciBuf/NvSciSync Interop, NvMedia CUDA Interop, Simple D3D11, Simple D3D12 CUDA Interop, Vulkan CUDA Interop PI Approximation, Vulkan CUDA Interop Sinewave, Vulkan Image - CUDA Interop
cudaSignalSemaphore	Vulkan CUDA Interop PI Approximation, Vulkan CUDA Interop Sinewave
cudaStreamAddCallback	Simple CUDA Callbacks
cudaStreamAttachMemAsync	NV12toBGRandResize, Unified Memory Streams, Unified and other CUDA Memories Performance

CUDA Runtime API	Samples
<p>cudaStreamBeginCapture</p>	<p>Conjugate Gradient using Cuda Graphs, Graph Memory Nodes, Jacobi CUDA Graphs, Simple CUDA Graphs</p>
<p>cudaStreamCreate</p>	<p>Concurrent Kernels, Conjugate Gradient using Cuda Graphs, Monte Carlo Option Pricing with Multi-GPU support, NV12toBGRandResize, Simple CUDA Callbacks, Simple CUDA Graphs, Simple D3D12 CUDA Interop, Simple Multi Copy and Compute, Simple Multi-GPU, Unified Memory Streams, Unified and other CUDA Memories Performance, Vulkan Image - CUDA Interop, batchCUBLAS, conjugateGradient using MultiDevice Cooperative Groups, cuSolverDn Linear Solver , cuSolverRf Refactorization, cuSolverSp Linear Solver , cuSolverSp Lowlevel QR Solver, cuSolverSp LowlevelCholesky Solver, simpleAttributes, simpleHyperQ, simpleStreams</p>
<p>cudaStreamCreateWithFlags</p>	<p>Advanced Quicksort (CUDA Dynamic Parallelism), Binary Partition Cooperative Groups, CUDA NvSciBuf/NvSciSync Interop, Global Memory to Shared Memory Async Copy, Graph Memory Footprint, Graph Memory Nodes, Jacobi CUDA Graphs, Matrix Multiplication (CUDA Runtime API Version), MersenneTwisterGP11213, NVJPEG simple, NvMedia CUDA Interop, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Simple Arrive Wait Barrier, Simple Atomic Intrinsic, Simple D3D11, Simple Driver-Runtime Interaction, Simple Quicksort (CUDA Dynamic Parallelism), Vulkan CUDA Interop PI Approximation, Vulkan CUDA Interop Sinewave, cuDLA Error Reporting, cuDLA Hybrid Mode, simpleIPC, stream Ordered Allocation, stream Ordered Allocation IPC Pools, stream Ordered Allocation Peer-to-Peer access</p>
<p>cudaStreamCreateWithPriority</p>	<p>Stream Priorities</p>
<p>cudaStreamDestroy</p>	<p>CUDA NvSciBuf/NvSciSync Interop, Concurrent Kernels, Conjugate Gradient using Cuda Graphs, Graph Memory Footprint, Graph Memory Nodes, MersenneTwisterGP11213, Monte Carlo Option Pricing with Multi-GPU support, NV12toBGRandResize, NVJPEG simple, NvMedia CUDA Interop, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Simple CUDA Callbacks, Simple CUDA Graphs, Simple Multi Copy and Compute, Simple Multi-GPU, Simple Quicksort (CUDA Dynamic Parallelism), Unified Memory Streams, Unified and other CUDA Memories</p>

CUDA Runtime API	Samples
	Performance, Vulkan CUDA Interop PI Approximation, cuDLA Error Reporting, cuDLA Hybrid Mode, cuSolverDn Linear Solver , cuSolverRf Refactorization, cuSolverSp Linear Solver , cuSolverSp Lowlevel QR Solver, cuSolverSp LowlevelCholesky Solver, simpleHyperQ, simpleIPC, simpleStreams, stream Ordered Allocation, stream Ordered Allocation IPC Pools, stream Ordered Allocation Peer-to-Peer access
cudaStreamEndCapture	Conjugate Gradient using Cuda Graphs, Graph Memory Nodes, Jacobi CUDA Graphs, Simple CUDA Graphs
cudaStreamGetFlags	Batched Label Markers And Label Compression NPP, Watershed Segmentation NPP
cudaStreamQuery	CUDA N-Body Simulation, CUDA N-Body Simulation on Screen, CUDA N-Body Simulation with GLES
cudaStreamSetAttribute	simpleAttributes
cudaStreamSynchronize	Batched Label Markers And Label Compression NPP, Binary Partition Cooperative Groups, CUDA NvSciBuf/ NvSciSync Interop, Conjugate Gradient using Cuda Graphs, Global Memory to Shared Memory Async Copy, Graph Memory Footprint, Graph Memory Nodes, Jacobi CUDA Graphs, Matrix Multiplication (CUDA Runtime API Version), MersenneTwisterGP11213, Monte Carlo Option Pricing with Multi-GPU support, NVJPEG simple, NvMedia CUDA Interop, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Simple Arrive Wait Barrier, Simple Atomic Intrinsic, Simple CUDA Graphs, Simple D3D12 CUDA Interop, Simple Driver-Runtime Interaction, Simple Multi-GPU, Unified Memory Streams, Unified and other CUDA Memories Performance, Vulkan CUDA Interop PI Approximation, Vulkan CUDA Interop Sinewave, Watershed Segmentation NPP, conjugateGradient using MultiDevice Cooperative Groups, cuDLA Error Reporting, cuDLA Hybrid Mode, simpleAttributes, simpleIPC, stream Ordered Allocation, stream Ordered Allocation IPC Pools, stream Ordered Allocation Peer-to-Peer access
cudaStreamWaitEvent	Concurrent Kernels, Graph Memory Nodes, Peer-to-Peer Bandwidth Latency Test with Multi-GPUs, Simple CUDA Graphs, simpleIPC, stream Ordered Allocation Peer-to-Peer access

CUDA Runtime API	Samples
cudaTimelineSemaphore	Vulkan CUDA Interop Sinewave
cudaUnbindTexture	Simple Direct3D10 Render Target
cudaUpdateHeightmapKernel	CUDA FFT Ocean Simulation
cudaUpdateVkImage	Vulkan Image - CUDA Interop
cudaVertMem	Vulkan CUDA Interop Sinewave
cudaVkImportImageMem	Vulkan Image - CUDA Interop
cudaVkImportSemaphore	Vulkan Image - CUDA Interop
cudaVkSemaphoreSignal	Vulkan Image - CUDA Interop
cudaVkSemaphoreWait	Vulkan Image - CUDA Interop
cudaWaitExternalSemaphoresAsync	CUDA NvSciBuf/NvSciSync Interop, NvMedia CUDA Interop, Simple D3D11, Simple D3D12 CUDA Interop, Vulkan CUDA Interop PI Approximation, Vulkan CUDA Interop Sinewave, Vulkan Image - CUDA Interop
cudaWaitSemaphore	Vulkan CUDA Interop PI Approximation, Vulkan CUDA Interop Sinewave
cudaXtFree	Simple CUFFT_MGPU, SimpleCUFFT_2d_MGPU

Chapter 7. Frequently Asked Questions

Answers to frequently asked questions about CUDA can be found at <http://developer.nvidia.com/cuda-faq> and in the [CUDA Toolkit Release Notes](#).

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

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

© 2007-2022 NVIDIA Corporation & affiliates. All rights reserved.