

## Lecture 5: Libraries and tools

Prof Wes Armour  
[wes.armour@eng.ox.ac.uk](mailto:wes.armour@eng.ox.ac.uk)

Prof Mike Giles  
[mike.giles@maths.ox.ac.uk](mailto:mike.giles@maths.ox.ac.uk)

Oxford e-Research Centre  
 Department of Engineering Science

1

1

## Learning outcomes

In this fifth lecture we will learn about GPU libraries and tools for GPU programming.

You will learn about:

- NVIDIA GPU libraries and their usefulness in scientific computing.
- Third party libraries.
- Directives based approaches to GPU computation.
- Tools for GPU programming.

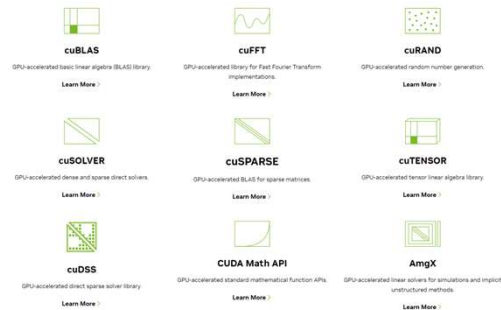
2

2

## Software overview

NVIDIA provides a **rich ecosystem of software tools** that allow you to easily utilise GPUs in your project.

During this lecture we will focus on a range of **libraries and software tools that will make your life easier** when either writing CUDA code or when utilising GPUs in your projects.



<https://developer.nvidia.com/gpu-accelerated-libraries>

3

3

## Dependencies – Advantages / Disadvantages

**Some advantages:**

- **Simple to use** – you don't need to write your own complex code to perform a specific task.
- **Well maintained** – always benefit from the latest optimisations and improvements.
- **Easier(?)** to move from CPU code to GPU code (for example see cuFFTW).

**Some disadvantages:**

- **Reduces portability** - can make installing your code on another system harder (the user also needs to have the dependency installed).
- **Reliance on a third party** - If the dependency isn't maintained it could break your code as other things (e.g. compiler) are updated.
- **Inheritance** - If your code is very dependent on it and the developers stop supporting it – you become the owner (not a great position to be in).

4

4

## CUDA math library

The CUDA Math Library contains all of the typical mathematical functions that you will need for your projects. *It is very similar to Intel's MKL library.*

- various **exponential** and **log** functions
- **trigonometric** functions and their inverses
- **hyperbolic** functions and their inverses
- **error functions** and their inverses
- **Bessel** and **Gamma** functions
- **vector norms** and reciprocals (esp. for graphics)
- To use - "#include math.h"

The library supports standard int, float and double types, but in recent years has also added support for fp8, fp16 and bfloat16.

Typecasting and SIMD intrinsics are also included in this library.



### CUDA Math Library

GPU-accelerated standard mathematical function library

CUDA Math Library | NVIDIA Developer

5

5

## cuBLAS Library

The cuBLAS (CUDA Basic Linear Algebra Subprograms) library provides CUDA accelerated standard BLAS APIs (for 152 different routines) for dense matrices.

- includes **matrix-vector** and **matrix-matrix** product.
- it is possible to call cuBLAS routines from user kernels (via a device API).
- some support for a single routine call to do a "batch" of smaller matrix-matrix multiplications.
- also support for using CUDA streams to do a large number of small tasks concurrently.
- has support for **multi-GPU operation** (cuBLASxt or cuBLASmg).
- has **mixed / low precision** implementations.



### cuBLAS

GPU-accelerated basic linear algebra (BLAS) library

6

6

## cuBLAS Library

To use cuBLAS in your codes, a set of routines are called from your host code. These come in two forms, helper routines and compute routines.

#### Helper routines for:

- memory allocation
- data copying from CPU to GPU, and vice versa
- error reporting



### cuBLAS

GPU-accelerated basic linear algebra (BLAS) library

#### Compute routines for:

- e.g. **matrix-matrix** and **matrix-vector** product
- **Warning!** Some calls are asynchronous, i.e. the call starts the operation but the host code then continues before it has completed!!

cuBLAS | NVIDIA Developer

7

7

## cuBLAS Library

cuBLAS is one of three libraries that utilize **'tensor cores'**. Tensor cores differ from standard processing cores in that they are designed to perform very specific operations, and these operations are executed on mixed-precision data.

By reducing the precision of your matrix/vector operations in cuBLAS, you can achieve significant acceleration.

<https://developer.nvidia.com/cublas>



8

8

## cuBLAS Library

Since the introduction of Ampere, there has been a 'TF32' variable that serves as a compromise between FP32 and BFLOAT16. It allows for a balance between lower precision and FP32 matrix-matrix operations.

Useful for:

- AI Training
- Liner solvers

<https://developer.nvidia.com/blog/accelerating-ai-training-with-tf32-tensor-cores/#:~:text=TF32%20mode%20is%20the%20default,any%20changes%20to%20mode%20scripts>

## cuBLAS Example

Here is an example of using cuBLAS to perform the SAXPY operation.

SAXPY stands for *Single-Precision A-X Plus Y*

$$y = \alpha x + y$$

Note the use of cuBLAS\_v2.h, cuBLAS.h contains legacy interfaces so you should pick which best fits with your needs.

```

#include <stdio.h>
#include <stdlib.h>
#include <cuda_runtime.h>
#include <cuda.h>
#include <cuBLAS_v2.h>

#define N 8

int main(void) {
    float *x, *y;
    cudaMallocManaged(&x, N * sizeof(float));
    cudaMallocManaged(&y, N * sizeof(float));

    // Initialize x and y arrays here
    for (int i = 0; i < N; i++) {
        x[i] = -1.0f;
        y[i] = (float)(i + 1);
    }

    cublasHandle_t handle;
    cublasCreate(&handle);

    const float alpha = 2.0f;
    cublasSaxpy(handle, N, &alpha, x, 1, y, 1);

    // Synchronize device and check for errors
    cudaDeviceSynchronize();

    // Check the result
    float maxError = 0.0f;
    for (int i = 0; i < N; i++) {
        maxError = fmax(maxError, fabs(y[i] - (float)(i + 1)));
        printf("y[%d]: %f\n", i, y[i]);
    }
    printf("Max error: %f\n", maxError);

    cudaFree(x);
    cudaFree(y);

    cublasDestroy(handle);

    return 0;
}
    
```

<https://docs.nvidia.com/cuda/archive/11.4.4/cublas/index.html>

## cuBLAS Example

Compile with:

```
$ nvcc cuBLAS_example.cu -o test -lcublas
```

Results are as we would expect.

We have initialised  $x$  to be -1.0, we have multiplied it by the scalar  $\alpha$  which we set equal to 2.0.

We then add the output to  $y$ , which was initialised from 0(+3) to 7(+3).

```

y[0]: 1.000000
y[1]: 2.000000
y[2]: 3.000000
y[3]: 4.000000
y[4]: 5.000000
y[5]: 6.000000
y[6]: 7.000000
y[7]: 8.000000
Max error: 0.000000
    
```

<https://developer.nvidia.com/blog/six-ways-saxpy/>

## cuTENSOR

Tensor cores, originally introduced on Volta in 2017, provided hardware-enabled acceleration for matrix-matrix multiplies.

$D = A * B + C$

- These cores originally performed a 4x4 matrix multiply-accumulate (think of it as a FMA for matrices) using the `wmma::instruction`.
- Matrices **A** and **B** would be lower precision and the accumulators would be the same or higher precision.
- With Ampere and Hopper, some of these restrictions have been relaxed.

<https://arxiv.org/pdf/2206.02874.pdf>  
<https://developer.nvidia.com/blog/programming-tensor-cores-cuda-9/>  
<https://developer.nvidia.com/blog/nvidia-automatic-mixed-precision-tensorflow/>  
<https://developer.download.nvidia.com/video/gputechconf/gtc/2019/presentation/s9593-cutorser-high-performance-tensor-operations-in-cuda-v2.pdf>

## cuTENSOR

Current tensor cores, such as those in Hopper, are able to use various precisions and exploit sparsity to achieve further acceleration.

- FP64 inputs with FP32 compute (DMMA).
- FP32 inputs with FP16, BF16, or TF32 compute.
- Complex-times-real operations.
- Conjugate support (without the need for transposition).
- Support for tensors with up to 64-dimensions.

<https://developer.nvidia.com/cutensor>

<https://docs.nvidia.com/cuda/cutensor/index.html>

<https://developer.nvidia.com/blog/nvidia-hopper-architecture-in-depth/>

<https://github.com/NVIDIA/CUDALibrarySamples/blob/master/cuTENSOR/reduction.cu>

## cuTENSOR

A useful description of the low level tensor operations can be found here:

<https://docs.nvidia.com/cuda/ampere-tuning-guide/index.html#improved-tensor-core-operations>  
(note this is for Ampere)

Here you can find code for a double precision GEMM computation using the Double precision Warp Matrix Multiply and Accumulate (DMMA) here:

[https://github.com/NVIDIA/cuda-samples/tree/master/Samples/3\\_CUDA\\_Features/dmmaTensorCoreGemm](https://github.com/NVIDIA/cuda-samples/tree/master/Samples/3_CUDA_Features/dmmaTensorCoreGemm)

**Note the differences between generations**

Generation	GPU Architecture	Input Matrix format	Output Accumulator format	Matrix Instruction Size (Blocks)
Volta (19-bit precision)	NVIDIA Volta Architecture	FP64	FP32 / FP32	8x8x4
	NVIDIA Turing Architecture	FP16	FP16 / FP32	8x8x4 / 16x8x8 / 16x8x16
	NVIDIA Ampere Architecture	FP16 / BF16x16	FP16 / FP32 (BF16x16 only supports FP32 as accumulator)	16x8x8 / 16x8x16
Hopper (19-bit precision)	NVIDIA Volta Architecture	N/A	N/A	N/A
	NVIDIA Turing Architecture	N/A	N/A	N/A
	NVIDIA Ampere Architecture	TF32 (19-bit)	FP32	16x8x4
Hopper (Integer DMMA)	NVIDIA Volta Architecture	N/A	N/A	N/A
	NVIDIA Turing Architecture	unsigned char/signed char (8-bit precision)	int32	8x8x16
	NVIDIA Ampere Architecture	unsigned char/signed char (8-bit precision)	int32	8x8x16 / 16x8x16 / 16x8x32
Hopper (Integer sub-byte DMMA)	NVIDIA Volta Architecture	N/A	N/A	N/A
	NVIDIA Turing Architecture	unsigned int/signed int (16-bit precision)	int32	8x8x32
	NVIDIA Ampere Architecture	unsigned int/signed int (16-bit precision)	int32	8x8x32 / 16x8x32 / 16x8x64
Hopper (Binary DMMA)	NVIDIA Volta Architecture	N/A	N/A	N/A
	NVIDIA Turing Architecture	single bit	int32	8x8x128
	NVIDIA Ampere Architecture	single bit	int32	8x8x128 / 16x8x128 / 16x8x256
Hopper (8-bit precision)	NVIDIA Volta Architecture	N/A	N/A	N/A
	NVIDIA Turing Architecture	N/A	N/A	N/A
	NVIDIA Ampere Architecture	FP64	FP64	8x8x4

## cuFFT Library

The cuFFT library is a GPU accelerated library that provides Fast Fourier Transforms (FFTs).

- It provides **1D, 2D and 3D FFTs**.
- It encompasses almost all of the variations found in FFTW and other CPU libraries.
- It **includes the cuFFTW library**, a porting tool, to enable users of FFTW to start using GPUs with minimal effort.
- **Provides some device level functionality** (if this is something of interest, you could ask Karel – he has already produced shared memory device level FFTs for some of our projects).

**cuFFT**  
GPU-accelerated library for Fast Fourier Transforms

## cuFFT Library

cuFFT is used exactly like cuBLAS - it has a set of routines called by host code:

Helper routines include “plan” construction.

Compute routines perform 1D, 2D, 3D FFTs:

- `cufftExecC2C()` - complex-to-complex.
- `cufftExecR2C()` - real-to-complex.
- `cufftExecC2R()` - complex-to-real.

(double precision routines have different function calls, e.g. `cufftExec22Z()`)

It supports doing a “batch” of independent transforms, e.g. applying 1D transform to a 3D dataset.

The simpleCUFFT example in SDK is a good starting point.

<https://docs.nvidia.com/cuda/cufft/index.html#introduction>

**cuFFT**  
GPU-accelerated library for Fast Fourier Transforms

## cuFFT Example

Sample time = 1 millisecond

Number of samples (NX) = 1000

So signal length = 1 second

Create a sinusoid with 10 cycles over the duration of the signal.

Hence frequency of sinusoid is 10 Hz.

**Question: Anything wrong with this code?**

```

#include <cuFFT.h>
#include <cuda.h>
#include <math.h>

#define NX 1000
#define BATCH 1
#define PI 3.14159265358979323846

#define SAMPLE_TIME 0.001 // Sample time in seconds
#define TOTAL_TIME (NX * SAMPLE_TIME) // Total time of the signal

int main()
{
    cufftHandle plan;
    cufftComplex *data;
    cudaMallocManaged(&data, NX * BATCH * sizeof(cufftComplex));

    // Initialize the input data with a sine wave
    for (int i = 0; i < NX; i++) {
        float t = i * SAMPLE_TIME; // Time of the current sample
        data[i].x = sin(PI * 10 * t / TOTAL_TIME); // 10 cycles over the total time
        data[i].y = 0.0f;
    }

    // Create a 1D FFT plan
    cufftPlan1d(&plan, NX, CUFFT_C2C, BATCH);

    // Use the cuFFT plan to transform the signal in place
    cufftExecC2C(plan, data, data, CUFFT_FORWARD);

    // Synchronize the device
    cudaDeviceSynchronize();

    // Print the transformed data (power and frequency)
    for (int i = 0; i < NX; i++) {
        float freq = (i+1) / TOTAL_TIME; // Frequency in Hz
        float power = sqrt(data[i].x * data[i].x + data[i].y * data[i].y); // Power
        printf("Frequency: %f Hz, Power: %f\n", freq, power);
    }

    // Destroy the cuFFT plan
    cufftDestroy(plan);

    // Free device memory
    cudaFree(data);

    return 0;
}
    
```

17

## cuFFT Example

Compile with:

```
$ nvcc cuFFT_example.cu -o test -lcufft
```

Frequency resolution = sampling rate / FFT length

So, given a signal of length 1 second, having 10 cycles in it, we expect a 10Hz response.

Exactly what we see.

```


Frequency: 1.000000 Hz, Power: 0.000004
Frequency: 2.000000 Hz, Power: 0.000006
Frequency: 3.000000 Hz, Power: 0.000001
Frequency: 4.000000 Hz, Power: 0.000002
Frequency: 5.000000 Hz, Power: 0.000001
Frequency: 6.000000 Hz, Power: 0.000003
Frequency: 7.000000 Hz, Power: 0.000006
Frequency: 8.000000 Hz, Power: 0.000004
Frequency: 9.000000 Hz, Power: 0.000004
Frequency: 10.000000 Hz, Power: 500.000000
Frequency: 11.000000 Hz, Power: 0.000004
Frequency: 12.000000 Hz, Power: 0.000005
Frequency: 13.000000 Hz, Power: 0.000013
Frequency: 14.000000 Hz, Power: 0.000005
Frequency: 15.000000 Hz, Power: 0.000010
Frequency: 16.000000 Hz, Power: 0.000002
Frequency: 17.000000 Hz, Power: 0.000001
Frequency: 18.000000 Hz, Power: 0.000006
Frequency: 19.000000 Hz, Power: 0.000002
    
```

18

## cuSPARSE Library

cuSPARSE is a GPU accelerated library that provides various routines to work with sparse matrices.

- Includes **sparse matrix-vector** and **matrix-matrix products**.
- Can be used for iterative solution (but see cuSOLVER for an easy life).
- Also has solution of sparse triangular system
- Note: batched tridiagonal solver is in cuBLAS not cuSPARSE



**cuSPARSE**  
GPU-accelerated BLAS for sparse matrices


19

## cuRAND Library

The cuRAND library is a GPU accelerated library for random number generation.

It has many different algorithms for pseudorandom and quasi-random number generation.

Pseudo: XORWOW, mrg32k3a, Mersenne Twister and Philox 4x32\_10  
Quasi: SOBEL and Scrambled SOBEL



**cuRAND**  
GPU-accelerated random number generation (RNG)

Uniform, Normal, log-Normal and Poisson outputs

This library also includes device level routines for RNG within user kernels.

20

## cuRAND Example

Here is an example of using cuRAND.

This code generates a normal distribution with a mean of zero, standard deviation equal to one.

It prints out the generated numbers to a file.

To compile:

```
nvcc cuRAND_example.cu -o test -lcublas
```

```
#include <stdio.h>
#include <curand.h>

#define N 1000000

int main(int argc, char *argv[]) {
    curandGenerator_t gen;
    float *data;

    /* Allocate Unified Memory - accessible from CPU or GPU */
    cudaMallocManaged(&data, N*sizeof(float));

    /* Create pseudo-random number generator */
    curandCreateGenerator(&gen, CURAND_RNG_PSEUDO_DEFAULT);

    /* Set seed */
    curandSetPseudoRandomGeneratorSeed(gen, 123456789);

    /* Generate N floats on device */
    curandGenerateNormal(gen, data, N, 0.0f, 1.0f);

    /* Wait for GPU to finish before accessing on host */
    cudaDeviceSynchronize();

    /* Open a file for writing */
    FILE *file = fopen("output.txt", "w");
    if (file == NULL) {
        printf("Error opening file!\n");
        return 1;
    }

    /* Write the data to the file */
    for(int i = 0; i < N; i++) {
        fprintf(file, "%1.0f\n", data[i]);
    }

    /* Close the file */
    fclose(file);

    /* Cleanup */
    curandDestroyGenerator(gen);
    cudaFree(data);

    return 0;
}
```

21


## cuRAND Example

Using matplotlib:

```
import matplotlib.pyplot as plt
# Read the data from the file
with open("output.txt", "r") as f:
    data = [float(line) for line in f]

# Create a histogram
plt.hist(data, bins=50, alpha=0.5, color='g')

# Show the plot
plt.show()
```



22

## cuSOLVER

cuSOLVER brings together cuSPARSE and cuBLAS.

Has solvers for dense and sparse systems.

Key LAPACK dense solvers, 3 – 6x faster than MKL.

Sparse direct solvers, 2–14x faster than CPU equivalents.



**cuSOLVER**

GPU-accelerated dense and sparse direct solvers

23

23

## Other notable libraries

CUB (CUDA Unbound): <https://nvlabs.github.io/cub/>

- Provides a collection of basic building blocks at three levels: device, thread block, warp.
- Functions include sort, scan and reduction.
- Thrust uses CUB for CUDA versions of key algorithms.

<http://on-demand.gputechconf.com/gtc/2014/presentations/S4566-cub-collective-software-primitives.pdf>

AmgX (originally named NVAMG): <http://developer.nvidia.com/amgx>

- Library for algebraic multigrid.
- Well suited for implicit unstructured methods.



**AmgX**

GPU-accelerated linear solvers for simulations and implicit unstructured methods

24

24

## Other notable libraries


### cuDNN

- Library for Deep Neural Networks
- Some parts developed by Jeremy Appleyard (NVIDIA) working in Oxford



### nvGraph

- Page Rank, Single Source Shortest Path, Single Source Widest Path
- <https://developer.nvidia.com/nvgraph>



### NPP (NVIDIA Performance Primitives)

- Library for imaging and video processing
- Includes functions for filtering, JPEG decoding, etc.



### CUDA Video Decoder API...


25

25

## NCCL

The NVIDIA Collective Communication Library (NCCL) is a low level library that allows GPUs to communicate within a node and across nodes. It provides functionality such as:

- all-gather,
- all-reduce,
- broadcast,
- point-to-point.



Optimised for PCIe and NVLink (within a node) and over NVIDIA Mellanox Network between nodes.

For ML people – use torch.distributed

<https://developer.nvidia.com/nccl>  
<https://pytorch.org/docs/stable/distributed.html>

26

26

## MAGMA



MAGMA (Matrix Algebra on GPU and Multicore Architectures) has been available for a few years (See nice SC17 handout: <http://www.icl.utk.edu/files/print/2017/magma-sc17.pdf> )

- LAPACK for GPUs – higher level numerical linear algebra, layered on top of cuBLAS.
- It supports interfaces to current LA packages so porting from a previous LAPACK code is easy.
- MAGMA allows applications to exploit heterogeneous systems consisting of multicore CPUs and multi-GPUs.

<https://icl.utk.edu/magma/>  
<https://developer.nvidia.com/magma>

27

27

## ArrayFire

Originally a commercial software (from Accelereyes), but is now open source.

- Supports both CUDA, OpenCL and OneAPI execution.
- C, C++ and Python interfaces.
- Supports NVIDIA and AMD GPUs/APUs, Intel processors and mobile devices from ARM, Qualcomm...
- Wide range of functionality including linear algebra, image and signal processing, random number generation, sorting...
- Actively developed.



<https://arrayfire.com/>  
<https://github.com/arrayfire/arrayfire>

28

28

## Thrust

Thrust is a high-level C++ template library with an interface based on the C++ Standard Template Library (STL).


Thrust has a very different philosophy to other libraries - *users write standard C++ code (no CUDA) but get the benefits of GPU acceleration.*

Thrust relies on C++ object-oriented programming – certain objects exist on the GPU, and operations involving them are implicitly performed on the GPU.

It has lots of built-in functions for operations like sort and scan.

It also simplifies memory management and data movement.

<https://thrust.github.io/>



29

29

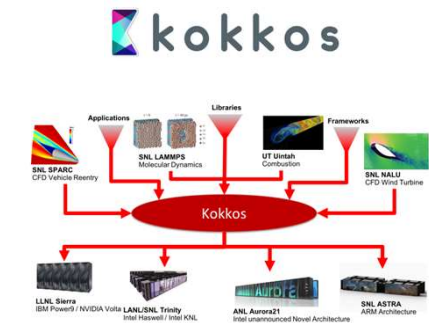
## Kokkos

Kokkos is another high-level C++ template library, similar to Thrust.

It has been developed in the US DoE Labs, so there is considerable investment in both capabilities and on-going software maintenance.

Could be worth investigating if you are considering using Thrust in your projects.

For more information see  
<https://kokkos.github.io/kokkos-core-wiki/>  
<https://kokkos.org/about/>



30

30


## A final word on libraries

NVIDIA maintains webpages with links to a variety of CUDA libraries:

[www.developer.nvidia.com/gpu-accelerated-libraries](http://www.developer.nvidia.com/gpu-accelerated-libraries)

and other tools:

[www.developer.nvidia.com/tools-ecosystem](http://www.developer.nvidia.com/tools-ecosystem)



[https://en.wikipedia.org/wiki/Duke\\_Humfrey%27s\\_Library](https://en.wikipedia.org/wiki/Duke_Humfrey%27s_Library)

31

31

## A note on directive based approaches and other languages

32

32




## OpenMP

**OpenMP 5.0** is a directive based approach to parallelisation.

- It uses a fork-join model.
- Can be used in C / C++ and FORTRAN codes.
- It supports both CPU and GPU hardware.

Is now becoming the industry standard for in node CPU parallelisation.



By Wikipedia user A1 - w:en:File:Fork\_join.svg, CC BY 3.0,  
<https://commons.wikimedia.org/w/index.php?curid=32004077>

33

33

## SYCL


SYCL ("sickle") - C++ Single-source Heterogeneous Programming for OpenCL.

From KHRONOS Group (responsible for OpenCL).

**Provides an abstraction layer that builds on OpenCL.**

It enables **code for heterogeneous processors** to be written in a "single-source" style using completely standard C++.

Supported by Intel, NVIDIA and AMD.



34

34

## Other Languages

**FORTRAN:** CUDA FORTRAN compiler with natural FORTRAN equivalent to CUDA C.

**MATLAB:** can call kernels directly, or use OOP like Thrust to define MATLAB objects which live on the GPU  
<https://uk.mathworks.com/help/parallel-computing/run-matlab-functions-on-a-gpu.html>

**Mathematica:** similar to MATLAB?

**Python:** CuPy (compatible with NumPy – acceleration for array computations), Numba and CUDA python  
<http://mathematician.de/software/pycuda>  
<https://store.continuum.io/cshop/accelerate/>  
<https://developer.nvidia.com/cuda-python>  
<https://developer.nvidia.com/how-to-cuda-python>  
<https://nvidia.github.io/cuda-python/overview.html>

35

35

## Other useful things...



<https://developer.nvidia.com/tools-ecosystem>  
<https://developer.nvidia.com/language-solutions>  
<https://developer.nvidia.com/hpc-sdk>  
<https://developer.nvidia.com/hpc-compilers>

36

36

## Which library should I use for my problem?

37

37

## The seven dwarfs

Phil Colella a senior researcher at Lawrence Berkeley National Laboratory, talked about "7 dwarfs" of numerical computation in 2004.

Expanded to 13 by a group of UC Berkeley professors in a 2006 report: "A View from Berkeley"  
[www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-183.pdf](http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-183.pdf)



38

38

## The seven dwarfs

These 13 dwarfs define key algorithmic kernels in many scientific computing applications.

They have been **very helpful to focus attention on HPC challenges and development of libraries and problem-solving environments/frameworks.**



39

39

## The seven dwarfs

1. Dense linear algebra
2. Sparse linear algebra
3. Spectral methods
4. N-body methods
5. Structured grids
6. Unstructured grids
7. Monte Carlo



40

40

## 1. Dense Linear Algebra

Many tools available, some from NVIDIA, some third party:

- cuBLAS
- cuSOLVER
- MAGMA
- ArrayFire

CUTLASS, an NVIDIA tool for Fast Linear Algebra in CUDA C++ might also be worth a look if you can't use the above libraries for any reason.

<https://devblogs.nvidia.com/cutlass-linear-algebra-cuda/>



**cuBLAS**

GPU-accelerated basic linear algebra (BLAS) library



**cuSOLVER**

GPU-accelerated dense and sparse direct solvers



**ARRAYFIRE**

41

41


## 2. Sparse Linear Algebra

Iterative solvers

- Some available in Petsc (Portable, Extensible Toolkit for Scientific Computation, for solving PDEs) - <https://petsc.org/release/overview/nutshell/>
- Others can be implemented using sparse matrix-vector multiplication from cuSPARSE (is also now in PETSc).
- NVIDIA has AmgX, an algebraic multigrid library.

Direct solvers

- NVIDIA's cuSOLVER.
- SuperLU project (Gaussian elimination with partial pivoting) <https://portal.nersc.gov/project/sparse/superlu/>
- STRUMPACK (ask Mike) <https://portal.nersc.gov/project/sparse/strumpack//>



**cuSOLVER**

GPU-accelerated dense and sparse direct solvers

42

42

## 3. Spectral methods

cuFFT /cuFFTW

Library provided / maintained by NVIDIA

For those interested in FFTs on GPUs – ask Karel / Jack...



**cuFFT**

GPU-accelerated library for Fast Fourier Transforms

43

43


## 4. N-Body methods

OpenMM:

- <http://openmm.org/>  
open source package to support molecular modelling, developed at Stanford.

Fast multipole methods:

- ExaFMM by Yokota and Barba: <http://www.bu.edu/exafmm/>
- FMM2D by Holm, Engblom, Goude, Holmgren: <http://user.it.uu.se/~stefane/freeware>  
<https://lorenabarba.com/figshare/exafmm-10-years-7-re-writes-the-tortuous-progress-of-computational-research/>
- Software by Takahashi, Cecka, Fong, Darve: <http://onlinelibrary.wiley.com/doi/10.1002/nme.3240/pdf>



<https://docs.nvidia.com/cuda/cuda-samples/index.html#cuda-n-body-simulation>  
[https://developer.download.nvidia.com/compute/cuda/7.1-beta/886\\_website/projects/nbody/doc/nbody\\_gems\\_3\\_ch31.pdf](https://developer.download.nvidia.com/compute/cuda/7.1-beta/886_website/projects/nbody/doc/nbody_gems_3_ch31.pdf)

44

44

## 5. Structured grids

Lots of people have developed one-off applications.


No great need for a library for single block codes (though possible improvements from “tiling”?).

Multi-block codes could benefit from a general-purpose library, mainly for MPI communication.

Oxford OPS project has developed a high-level open-source framework for multi-block codes, using GPUs for code execution and MPI for distributed-memory message-passing.

All implementation details are hidden from “users”, so they don’t have to know about GPU/MPI programming.

For those interested – ask Mike...



Sifflea, Mysid - Drawn by Sifflea, vectorized by Mysid. 45


45

## 6. Unstructured grids

In addition to GPU implementations of specific codes there are projects to create high-level solutions which others can use for their application codes:

- Alonso, Darve and others (Stanford).
- Oxford / Imperial College / Warwick project developed OP2, a general-purpose open-source framework based on a previous framework built on MPI.
- If there’s interest Mike could talk about OP2 and OPS in lecture 8/9.

See <https://op-dsl.github.io/> for both OPS and OP2



L. Zurek, CC BY-SA 3.0 -<http://creativecommons.org/licenses/by-sa/3.0/> - via Wikimedia Commons

46

46

## 7. Monte Carlo methods

- NVIDIA cuRAND library.
- ArrayFire library.
- Some examples in CUDA SDK distribution.
- Nothing else needed except for more output distributions?



**cuRAND**

GPU-accelerated random number generation (RNG)

**ARRAYFIRE**

47

47

Useful tools

48

48

## Tools - Debugging

`compute-sanitizer --tool memcheck`  
 A command line tool that detects array out-of-bounds errors, and mis-aligned device memory accesses – very useful because such errors can be tough to track down otherwise.

`compute-sanitizer --tool racecheck`  
 This checks for shared memory race conditions:

- Write-After-Write (WAW): two threads write data to the same memory location but the order is uncertain.
- Read-After-Write (RAW) and Write-After-Read (WAR): one thread writes and another reads, but the order is uncertain.

`compute-sanitizer --tool initcheck`  
 This detects the reading of uninitialized device memory.

`compute-sanitizer --tool synccheck`  
 This detects incorrect use of `__syncthreads()` and related intrinsics.

The screenshot shows the NVIDIA Compute Sanitizer website with sections for Memcheck, Racecheck, Initcheck, and Synccheck. Below the website is a terminal window showing the execution of the tool on a test program:

```

user@machine:~$ compute_sanitizer --tool memcheck ./curand_test
===== COMPUTE-SANITIZER =====
===== ERROR SUMMARY: 0 errors =====
user@machine:~$ compute_sanitizer --tool racecheck ./curand_test
===== COMPUTE-SANITIZER =====
===== RACECHECK SUMMARY: 0 hazards displayed (0 errors, 0 warnings) =====
user@machine:~$ compute_sanitizer --tool initcheck ./curand_test
===== COMPUTE-SANITIZER =====
===== ERROR SUMMARY: 0 errors =====
    
```

## Tools – CUDA-GDB

For those familiar with the GNU debugger – GDB, this is an extension of GDB that allows users to debug both GPU and CPU code.

All existing GDB debugging features are included for debugging host code and then further functionality allows the user to debug device code.

Supports C/C++ and Fortran (that includes CUDA code).



**CUDA-GDB**  
 Delivers a seamless debugging experience that allows you to debug both the CPU and GPU portions of your application simultaneously. Use CUDA-GDB on Linux or MacOS, from the command line, DDD or EMACS.

## Tools - IDEs

Integrated Development Environments (IDE):

Nsight Systems – Unified IDE for Windows/Linux/Mac/Jetson:  
<https://developer.nvidia.com/nsight-systems>

Nsight Visual Studio edition – NVIDIA plugin for Microsoft Visual Studio  
<http://developer.nvidia.com/nvidia-nsight-visual-studio-edition>

Nsight Eclipse plugins  
<https://docs.nvidia.com/cuda/nsight-eclipse-plugins-guide/index.html>

these come with editor, debugger, profiler integration



**NVIDIA® Nsight™**  
 The ultimate development platform for heterogeneous computing. Work with powerful debugging and profiling tools, optimize the performance of your CPU and GPU code. Find out about the Eclipse Edition and the graphics debugging enabled Visual Studio Edition.

## Tools - Profiling

NVIDIA Profiler `ncu` or `ncu-ui` for a graphical interface.

- This is a standalone piece of software for Linux and Windows systems.
- It uses hardware counters to collect a lot of useful information.
- Lots of things can be measured, but the limited number of counters means that, for some larger applications, it runs the application multiple to gather necessary information.
- The `ncu` CLI can be useful if you want to profile on a machine that you don't have a graphical interface to.
- Do `ncu --help` for more info on different options.



**NVIDIA Visual Profiler**  
 This is a cross-platform performance profiling tool that delivers developers vital feedback for optimizing CUDA C/C++ applications. First introduced in 2008, Visual Profiler supports all CUDA capable NVIDIA GPUs shipped since 2006 on Linux, Mac OS X, and Windows.

<https://docs.nvidia.com/nsight-compute/NsightCompute/index.html>  
<https://docs.nvidia.com/cuda/profiler-users-guide/index.html>

## What have we learnt?

In this lecture we've looked at the **wide software ecosystem that now surrounds GPU computing** and how that can be used to make your life as a programmer easier.

We've looked at **directives based approaches** and how these are useful.

Finally we've looked at **tools** that allow us to develop CUDA code in an easy and maintainable way.



53