




Article

GPU Computing with Python: Performance, Energy Efficiency and Usability [†]

Håvard H. Holm ^{1,2,*} , André R. Brodtkorb ^{3,4}  and Martin L. Sætra ^{4,5} ¹ Mathematics and Cybernetics, SINTEF Digital, P.O. Box 124 Blindern, NO-0314 Oslo, Norway² Department of Mathematical Sciences, Norwegian University of Science and Technology, NO-7491 Trondheim, Norway³ Research and Development Department, Norwegian Meteorological Institute, P.O. Box 43 Blindern, NO-0313 Oslo, Norway; andreb@met.no⁴ Department of Computer Science, Oslo Metropolitan University, P.O. Box 4 St. Olavs plass, NO-0130 Oslo, Norway; martinls@met.no⁵ Information Technology Department, Norwegian Meteorological Institute, P.O. Box 43 Blindern, NO-0313 Oslo, Norway

* Correspondence: havard.heitlo.holm@sintef.no

[†] This paper is an extended version of our paper published in International Conference on Parallel Computing (ParCo2019).

Received: 6 December 2019; Accepted: 6 January 2020; Published: 9 January 2020



Abstract: In this work, we examine the performance, energy efficiency, and usability when using Python for developing high-performance computing codes running on the graphics processing unit (GPU). We investigate the portability of performance and energy efficiency between Compute Unified Device Architecture (CUDA) and Open Compute Language (OpenCL); between GPU generations; and between low-end, mid-range, and high-end GPUs. Our findings showed that the impact of using Python is negligible for our applications, and furthermore, CUDA and OpenCL applications tuned to an equivalent level can in many cases obtain the same computational performance. Our experiments showed that performance in general varies more between different GPUs than between using CUDA and OpenCL. We also show that tuning for performance is a good way of tuning for energy efficiency, but that specific tuning is needed to obtain optimal energy efficiency.

Keywords: GPU computing; CUDA; OpenCL; high-performance computing; shallow-water simulation; power efficiency

1. Introduction

General-purpose computing using the graphical processing unit (GPU), known as GPU computing, was introduced in the early 2000s and has since become a popular concept. The first examples were acceleration of simple algorithms such as matrix–matrix multiplication by rephrasing the algorithm as operations on graphical primitives (see, e.g., [1]). This was cumbersome, and there existed no development tools for general-purpose computing. However, many algorithms were implemented on the GPU as proof-of-concepts, showing large speedups over the central processing unit (CPU) [2]. Today, the development environment for GPU computing has evolved tremendously and is both mature and stable: Advanced debuggers and profilers are available, making debugging, profile-driven development, and performance optimization easier than ever.

The GPU has traditionally been accessed using compiled languages such as C/C++ or Fortran for the CPU code, and a specialized programming language for the GPU. The rationale is often that performance is paramount, and that compiled languages are therefore required. However, for many GPU codes, most of the time is spent in the numerical code running on the GPU. In these

cases, we can possibly use a higher-level language such as Python for the program flow without significantly affecting the performance. The benefit is that using higher-level languages might increase productivity [3]. PyCUDA and PyOpenCL [4] are two Python packages that offer access to CUDA and OpenCL from Python and have become mature and popular packages since their initial release nearly ten years ago.

Herein, we compare the performance, energy efficiency, and usability of PyCUDA and PyOpenCL for important algorithmic primitives found in many high-performance computing (HPC) applications [5]: a memory-bound numerical simulation code and a computationally bound benchmark code. The memory-bound code performs simulation of the shallow-water equations using explicit stencils [6,7] and represents a class of problems that are particularly well suited for GPU computing [8–12]. The computationally bound code computes the Mandelbrot set, and we use it to explore limitations of using Python for GPU computing.

We show that accessing the GPU from Python is as efficient as from C/C++ in many cases, demonstrate how profile-driven development in Python is essential for increasing performance for GPU code (up to 5 times), and show that energy efficiency increases proportionally with performance tuning. Finally, we investigate the portability of the improvements and power efficiency both between CUDA and OpenCL and between different GPUs. Our findings are summarized in tables that justify that using Python can be preferable to C++ and that using CUDA can be preferable to using OpenCL. Our observations should be directly transferable to other similar architectures and problems.

2. Related Work

There are several high-level programming languages and libraries that offer access to the GPU for certain sets of problems and algorithms. OpenACC [13] (open accelerators) is one example which is pragma-based and offers a set of directives to execute code in parallel on the GPU. However, such high-level abstractions are typically only efficient for certain classes of problems and are often unsuitable for nontrivial parallelization or data movement. CUDA [14] and OpenCL [15] are two programming languages that offer full access to the GPU hardware, including the whole memory subsystem. This is an especially important point, since memory movement is a key bottleneck in many numerical algorithms [5] and therefore has a significant impact on energy consumption.

The performance of GPUs has been reported extensively [16], and several authors have shown that GPUs are efficient in terms of energy-to-solution. Huang et al. [17] demonstrated early on that GPUs could not only speed up computational performance but also increase power efficiency dramatically using CUDA. Qi et al. [18] show how OpenCL on a mobile GPU can increase performance of the discrete Fourier transform by 1.4 times and decrease energy use by 37%. Dong et al. [19] analyzed the energy efficiency of GPU BLAST which simulates compressible hydrodynamics using finite elements with CUDA and reported a 2.5 times speedup and a 42% increase in energy efficiency. Klöh [20] reported that there is a wide spread in terms of energy efficiency and performance when comparing 3D wave propagation and full waveform inversion on two different architectures. They compared an Intel (Santa Clara, CA, USA) Xeon coupled with an ARM-based Nvidia (Santa Clara, CA, USA) Jetson TX2 GPU module and found that the Xeon platform performs best in terms of computational speed, whilst the Jetson platform is most energy-efficient. Memeti et al. [21] compared the programming productivity, performance, and energy use of CUDA, OpenACC, OpenCL, and OpenMP for programming a system consisting of a CPU and GPU or a CPU and an Intel Xeon Phi coprocessor. They reported that CUDA, OpenCL, and OpenMP have a similar performance and energy consumption in one benchmark and that OpenCL performs better than OpenACC for another benchmark. In terms of productivity, the actual person writing the code is important, but OpenACC and OpenMP require less effort than CUDA and OpenCL, and CUDA can require significantly less effort than OpenCL.

Previous studies have also shown that CUDA and OpenCL can compete in terms of performance as long as the comparison is fair [22–25], and there have also been proposed automatic source to source compilers that report similar results [26,27]. This paper is an extension of Holm et al. [28], and we do

not know of any other established literature that thoroughly compares the performance, usability, and energy efficiency of CUDA and OpenCL when accessed from Python.

3. GPU Computing in Python

In this work, we focus on using Python to access the GPU through CUDA and OpenCL. These two GPU programming models are conceptually very similar, and both offer the same kind of parallelism primitives. The main idea is that the computational domain is partitioned into equally sized subdomains that are executed independently and in parallel. Even though the programming models are similar, their terminology differs slightly, and in this paper, we use that of CUDA. A full review is outside the scope of this work but can be found in [29,30]. The following sections give an overview of important parts of CUDA and OpenCL and discuss their respective Python wrappers. We give a short introduction to using them from C++ and Python and compare the benefits and drawbacks of both approaches.

3.1. CUDA

CUDA [14] was first released in 2007 and is available on all Nvidia GPUs as Nvidia's proprietary GPU computing platform. It includes third-party libraries and integrations, the directive-based OpenACC [13] compiler, and the CUDA C/C++ programming language. Today, five of the ten fastest supercomputers (including number one) use Nvidia GPUs, as well as nine out of the ten most energy-efficient [31].

CUDA is implemented in the Nvidia device driver, but the compiler (`nvcc`) and libraries are packaged in the CUDA toolkit and software development toolkit (SDK) (available at <https://developer.nvidia.com/cuda-zone>). The toolkit and SDK contain a plethora of examples and libraries. In addition, the toolkit contains Nvidia Nsight, which is an extension for Microsoft Visual Studio (Microsoft, Redmond, WA, USA) and Eclipse (Eclipse Foundation, Inc., Ottawa, Ontario, Canada) (for Linux) for interactive GPU debugging and profiling. Nsight offers code highlighting, a unified CPU and GPU trace of the application, and automatic identification of GPU bottlenecks. The Nvidia Visual Profiler is a standalone cross-platform application for profiling of CUDA programs, and CUDA versions for debugging (`cuda-gdb`) and memory checking (`cuda-memcheck`) also exist.

3.2. OpenCL

OpenCL [15] is a free and open heterogeneous computing platform that was initiated by Apple in 2009, and today, the OpenCL standard is maintained and developed by the Khronos group. Whilst CUDA was made specifically for Nvidia GPUs, OpenCL can run on a number of heterogeneous computing architectures, including GPUs, CPUs, field-programmable gate arrays (FPGAs), and digital signal processors (DSPs). The OpenCL application programming interface (API) is defined in a common C/C++ header, and a runtime library redirects OpenCL calls to the appropriate device driver by using an installable client driver (ICD), specific to each architecture. The actual OpenCL is therefore implemented in each vendor's device driver. Contrary to CUDA, there is no common toolkit, but there are several third-party libraries (for a list of OpenCL resources, see <https://www.khronos.org/opencl/resources>).

Profiling an OpenCL application can be challenging, and the available tools vary depending on your operating system and hardware vendor (an extensive list of OpenCL debugging and profiling tools can be found at https://www.khronos.org/opengl/wiki/Debugging_Tools). It is possible to get timing information on kernel and memory transfer operations by enabling event profiling information and adding counters explicitly in your source code. This requires extra work and makes the code more complex. Visual Studio can measure the amount of run time spent on the GPU, and CodeXL [32] can be used to get more information on AMD (Santa Clara, CA, USA) GPUs. CodeXL is a successor to gDebugger which offers features similar to those found in Nsight in addition to power profiling and is available both as a standalone cross-platform application and as a Visual Studio extension.

While it is possible to use Visual Profiler for OpenCL, this requires the use of the command-line profiling functionality in the Nvidia driver, which needs to be enabled through environment variables and a configuration file. After running the program with the profiling functionality in effect, the profiling data can be imported into Visual Profiler. Intel Code Builder (part of Intel SDK for OpenCL Applications [33]) and Intel Vtune Amplifier [34] can also be used for OpenCL debugging and profiling, but these tools only support Intel CPUs and Intel Xeon Phi processors.

One disadvantage of OpenCL is that there are large differences between the OpenCL implementations from different vendors, and good performance is likely to rely on vendor-specific extensions. One example is that OpenCL 2.2 is required to use C++ templates in the GPU code, but vendors such as Nvidia only support OpenCL version 1.2. It should also be mentioned that OpenCL has been replaced in favor of Metal [35] by Apple (Cupertino, CA, USA) in their most recent versions of Mac OS X.

3.3. GPU Computing from Python

Researchers spend a large portion of their time writing computer programs [36], and compiled languages such as C/C++ and Fortran have been the de facto standard within scientific computing for decades. These languages are well established, well documented, and give access to a plethora of native and third-party libraries. C++ is the standard way of accessing CUDA and OpenCL today, but developing code with these languages is time-consuming and requires great care. Using higher-level languages such as Python can significantly increase development productivity [3,37]. However, it should be mentioned that the OpenCL and CUDA kernels themselves are not necessarily made simpler or shorter by using Python: The productivity gain comes instead from Python's less verbose code style for the CPU part of the code. This influences every part of the host code, from the boilerplate initialization code and data preprocessing to CUDA/OpenCL API calls, post-processing, and visualization of results.

Since Python was first released in 1994, it has gained momentum within scientific computing. Python is easy to learn, powerful, and emphasizes readability of code. All together, this allows for efficient prototyping of code, but unfortunately often at the expense of performance. Since Python is an interpreted language, it does not have the speed of C/C++ and Fortran. For example, a loop in C/C++ or Fortran will execute at great speed, whilst this construction is relatively slow in Python. This is acceptable for many application areas, but for computationally intensive codes, it becomes a showstopper.

A popular approach to combining the speed of compiled code with Python is to have the program flow written in Python and the performance critical inner loop in a compiled language. The performance critical code can then be called from Python, using, e.g., SWIG (Simplified Wrapper and Interface Generator), Numba [38] or Cython [39]. The following example illustrates this by computing $C = A^T A$ with the dense general matrix multiplication (DGEMM) subroutine from the Basic Linear Algebra Subprograms (BLAS) library supplied with SciPy:

```
1 import numpy as np
2 import scipy.linalg.blas as blas
3 N = 2048
4 A = np.ones((N,N))
5 C = blas.dgemm(alpha=1.0, a=A.T, b=A)
```

Even though Python can be relatively slow, most of the time in this example is spent in the very efficient BLAS library, giving good overall performance. This approach thus offers the speed of compiled code together with the productivity of a high-level language. We can in a similar way execute GPU code from Python, and our experiments show that this gives a good overall performance. Hence, Python can also be a viable option for production-level code, not only for experimental code and prototypes.

Listing 1: Programming example that shows how to fill an array with the numbers 1 to N using PyCUDA. The corresponding C++ code is much longer.

```
1 import pycuda.autoinit
2 import pycuda.driver as drv
3 import pycuda.compiler as compiler
4 import numpy as np
5
6 module = compiler.SourceModule("""
7 __global__ void fill(float *dest) {
8 int i = blockIdx.x*blockDim.x+threadIdx.x;
9 dest[i] = i+1;
10 } """)
11 fill = module.get_function("fill")
12
13 N = 200
14 cpu_data = np.empty(N, dtype=np.float32)
15 fill(drv.Out(cpu_data), block=(N,1,1), grid=(1,1,1))
16 print(cpu_data)
```

There are several libraries that today offer access to the GPU from Python. One class of libraries are OpenCV [40] (open source computer vision) that offers GPU acceleration of algorithms within a specific field. Such libraries are outside the scope of this work, as we focus on general-purpose GPU computing. In addition to these types of libraries, Numba [38] and CuPy [41] are general-purpose programming environments in Python that offer full access to the GPU. However, the GPU programs in Numba are written as Python functions, and the programmer has to rely on Numba for efficient parallelization of the code. While such a design lowers the bar for developers to write code that executes on the GPU, details that are crucial for obtaining the full potential performance might be lost in the abstraction. Additionally, Numba is missing support for dynamic parallelism and texture memory. CuPy also offers functionality to define the GPU functions in terms of Python code but additionally supports raw kernels written in native CUDA.

PyCUDA and PyOpenCL [4] are Python packages that offer access to CUDA and OpenCL, respectively. Both libraries expose the complete API of the underlying programming models and aim to minimize the performance impact. The GPU kernels, which are crucial for the inner loop performance, are written in native low-level CUDA or OpenCL, and memory transfers and kernel launches are made explicit through Python functions. The result is an environment suitable for rapid prototyping of high-performing GPU code. Listing 1 shows a minimal example for filling an array with numbers 1 to N . The first four lines import PyCUDA and numpy. Lines six to ten hold the CUDA source code (written in CUDA C/C++) and compile it using the `SourceModule` interface. Line eleven retrieves a handle to the kernel so that we can call the GPU function from Python. Line 14 allocates data in Python using numpy, which is then handed over to the GPU kernel in line 15. This manages automatic uploading and downloading of data from the GPU through the `drv.Out` class. The example shows how PyCUDA can be used to run a GPU kernel in less than 20 lines of code. The equivalent C++ code would be far longer.

PyCUDA and PyOpenCL can be installed using popular package managers such as apt-get, pip, and conda. However, the fact that they also require that CUDA and OpenCL be properly installed makes things a bit more complicated. The installation procedure on Windows is especially tricky, and it can be a challenge to install the packages successfully. On Linux, a challenge with OpenCL is that Nvidia only supports version 1.2, whilst PyOpenCL tries to compile with version 2.0 by default when it is installed with pip. This results in runtime crashes when loading the PyOpenCL package. The solution is to manually compile the package, which can be cumbersome. An alternative is to

use the packages in apt-get, but these are often very outdated. For PyCUDA, however, we have experienced fewer difficulties installing on Linux.

PyOpenCL ships with integration for Jupyter Notebooks [42], which enables rapid prototyping in an interactive REPL environment (REPL stands for read-eval-print loop and is a class of interactive development environments where you execute the code as you write it). Using the Jupyter Notebook for prototyping is extremely efficient, and our development cycle has typically been as follows:

1. Prototype and develop code in a Jupyter Notebook;
2. Clean up code in the notebook;
3. Move code from notebook to separate Python modules.

The first stage often entails developing the idea and concept interactively, making test programs, and plotting intermediate and final results. Then, when we have a first implementation, we continue by cleaning up the code and making it suitable for moving into a separate Python module. The tests are added to an appropriate continuous integration environment.

Unfortunately, we have experienced that OpenCL occasionally crashes at random within the Jupyter Notebook environment. After a significant amount of debugging, we discovered that the context was not properly cleaned up, and variables that went out of scope were not properly deallocated, thereby causing crashes for longer Jupyter sessions. Unfortunately, these crashes demanded a reload of the GPU driver, which typically means rebooting the system. Our solution was simply to add explicit invocation of Python's garbage collector after dereferencing the variables pointing to GPU main memory.

After having worked with CUDA in the Jupyter Notebook, we discovered that similar problems were also present here. However, because CUDA is somewhat stateful, it uses a stack of contexts and this stack was not properly cleaned up. Our solution was in this case to write a context manager that performed the correct pushing and popping of the CUDA context stack. After having addressed these shortcomings of PyCUDA and PyOpenCL, we have not experienced any crashes caused by the interplay with the Jupyter Notebook.

3.4. C++ Versus Python

The GPU can be accessed both from C++ and Python, and these two approaches have their own benefits and drawbacks. In this section, we examine how a computationally bound application—computing the Mandelbrot set—can be implemented in the four combinations: C++ and CUDA; C++ and OpenCL; Python and CUDA; and Python and OpenCL.

The Mandelbrot set M consists of all complex points c in which

$$z_{n+1} = z_n^2 + c$$

remains bounded as $n \rightarrow \infty$. In practice, one typically initializes a complex number c and $z_0 = 0$, and iterates until $|z_{n+1}| \geq 2$ or an upper limit of iterations has been reached:

```
//Loop until iterations or until it diverges
while (|z| < 2.0 && n < iterations) {
z = z*z + c;
++n;
}
```

Our application consists of a GPU kernel which first initializes the complex coordinate, c , based on the position of each thread in the computational grid. Each thread then executes a loop as shown above before the output value for the thread is written to main memory. If we allow, e.g., 1000 iterations, a single thread may perform up to 1000 iterations of the above while-loop before it writes a single number to the main GPU memory. However, as Figure 1 shows, the number of actual iterations performed by each thread varies dramatically, even between threads located close to each other. Between CUDA

and OpenCL, the kernel code is close to identical, with only syntactic differences (see Section 4.1). The actual kernel that runs on the GPU is the same for both the C++ and the Python variants, and differences between them are therefore found in the host code only.




Figure 1. Mandelbrot set colored using continuous coloring. Each complex coordinate c is determined as within the set if the expression $z_{n+1} = z_n^2 + c$ does not diverge after a given set of iterations. The number of iterations performed for each thread is shown here using a continuous color scale. Locations close to the boundary perform many iterations, whilst only a few iterations are required far away. It is evident that threads in close proximity may perform a very different number of iterations.

The host code is responsible for allocating output data on the GPU, launching the kernel, and downloading the result from the GPU to the CPU. To benchmark the different variants, we computed the Mandelbrot set for different extents and measured how much time the CPU and GPU uses for different sections of the code. It should be noted that the overhead of running a traditional Python for-loop is significant but can be reduced using, e.g., so-called list comprehensions.

Table 1 shows a summary of our benchmarks. It shows clearly that the kernel launch overhead is larger for Python than for C++, but it becomes negligible when compared to the kernel execution time. Thus, for kernels that last more than a few ms, there will be little performance benefit of using C++, and Python will typically be equally fast. The reason for this is that the CPU in both cases simply will be waiting for the GPU to complete execution for most of the time, and the performance difference will be completely masked. This is in particular shown here in the “Wall time” row, which shows no practical difference between C++ and Python. In fact, for CUDA, the Python variant executes marginally faster. Our explanation to this is that PyCUDA automatically sets the compilation flags for `nvcc` for the specific GPU, yielding more optimized code.

One interesting difference between CUDA and OpenCL from Python is that the compilation time differs. PyCUDA uses `nvcc` to compile and link an executable in a similar fashion to a regular C++ program. PyOpenCL, on the other hand, simply hands over the OpenCL kernel source code to the OpenCL driver, which compiles it on-the-fly. We expect that incorporation of the recent Nvidia runtime compilation library (NVRTC) into PyCUDA will alleviate this shortcoming.

If we compare the other metrics, we see that both the development time and number of lines of code is significantly better for Python. Of particular note is that debugging and visualizing results becomes interactive when using Python and Jupyter Notebooks, thereby increasing the development efficiency dramatically. This is done using standard tools such as Matplotlib [43], making it extremely easy to visualize and explore results. In C++, on the other hand, the only way to reasonably explore the results is through file output (e.g., CSV) and plotting using third-party tools.

Table 1. Performance of Compute Unified Device Architecture (CUDA) and Open Compute Language (OpenCL) for the Mandelbrot application from both Python and C++. The code was run on an Nvidia Tesla M2090 graphical processing unit (GPU) to compute fifty consecutive zooms onto the point $-0.75 + 0.1i$ with a maximum of 5000 iterations. OpenCL 2.0 is available in the host section of the code, but the device section must still use OpenCL 1.2 for the C++ version. The development time is set subjectively by the authors and the lines of code metric contains only the central processing unit (CPU) code related to the actual GPU kernel launch. On Windows, we were unable to get asynchronous execution with OpenCL on several different machines, Python versions, and GPU driver versions. We have not been able to pinpoint the cause of this and suspect that asynchronous execution of OpenCL in Windows is not fully supported. We used page locked memory with CUDA, whilst this was not easily available through OpenCL for the download. The overhead for OpenCL therefore includes the transfer, as it cannot be run concurrently with other operations. It should be noted that OpenCL should support efficient pinned memory transfers, but we were unable to get this to work using several driver versions on several different machines. We therefore report the observed transfer times without pinned memory. The wall time includes all the time the CPU spends from launching the first kernel to having completed downloading the last result from the GPU.

	C++		Python	
	CUDA	OpenCL	CUDA	OpenCL
API version	10.0	1.2 / 2.0	10.0	1.2
Development time	Medium	Medium	Fast	Fast
Approximate lines of code	145	130	100	100
Compilation time	~5 s	~5 s	~5 s	Interactive
Kernel launch overhead	13 μ s	318 μ s	19 μ s	377 μ s
Download overhead	9 μ s	4007 μ s	52 μ s	8872 μ s
Kernel GPU time	480 ms	446 ms	478 ms	444 ms
Download GPU time	4.0 ms	3.9 ms	4.0 ms	8.8 ms
Wall time	24.2 s	22.5 s	24.1 s	22.7 s

4. Porting, Profiling, and Benchmarking Performance and Energy Efficiency

We now describe the process of porting code between PyOpenCL and PyCUDA and optimizing the PyCUDA versions through profile-driven development. We consider edsimulation of the shallow-water equations using three different numerical schemes:

- A linear finite difference scheme;
- A nonlinear finite difference scheme; and
- A high-resolution finite volume scheme.

The schemes are used for simulating real-world ocean currents, and two of them were used operationally in the early days of computational oceanography. All three schemes are essentially stencil operations with an increasing level of complexity, and their details are summarized in Holm et al. [7].

The numerical schemes are algorithmically well suited for the GPU, but little effort has been made to thoroughly optimize the codes' performance on a specific GPU. It is well known in the GPU computing community that performance is not portable between GPUs, neither for OpenCL nor CUDA, and automatically generating good kernel configurations is an active research area (see, e.g., [44–46]). We started by porting the three schemes to CUDA before using the available profiling tools for CUDA to analyze and optimize each scheme. The obtained optimizations were then also back-ported into the original OpenCL code. The profiling and tuning was carried out on a laptop with a dedicated GeForce 840M GPU, representing the low-end part of the GPU performance scale, and on a desktop with a mid-range GeForce GTX780 GPU representing a typical mid-range GPU. We compared the performance of the original and optimized implementations with PyCUDA and PyOpenCL using seven GPUs listed in Table 2, which also includes several high-end server GPUs.

Table 2. A list of the GPUs used in this work. The profile-driven development was carried out on the 840M and GTX780, and the remaining high-end GPUs were used for performance benchmarking. We used the 840M, GTX780, P100, and V100 for benchmarking power efficiency. Note that the performance in giga floating point operations per second (GFLOPS) is for single precision. The K80 GPU consists of two processors on the same card and has a boost feature for temporarily increasing the clock speed to increase performance to 2×4368 GFLOPS.

Model	Class	Architecture (Year)	Memory	GFLOPS	Bandwidth	Power Device
Tesla M2090	Server	Fermi (2011)	6 GiB	1331	178 GB/s	N.A.
Tesla K20	Server	Kepler (2012)	6 GiB	3524	208 GB/s	N.A.
GeForce GTX780	Desktop	Kepler (2013)	3 GiB	3977	288 GB/s	Watt meter
Tesla K80	Server	Kepler (2014)	2×6 GiB	2×2795	2×240 GB/s	N.A.
GeForce 840M	Laptop	Maxwell (2014)	4 GiB	790	16 GB/s	Watt meter
Tesla P100	Server	Pascal (2016)	12 GiB	9523	549 GB/s	nvidia-smi
Tesla V100	Server	Volta (2017)	16 GiB	14,899	900 GB/s	nvidia-smi

4.1. Porting from PyOpenCL to PyCUDA

The porting process requires changing both the kernel code that runs on the GPU and the API calling in the CPU code. The kernels will in most cases run and produce correct results after a simple change of keywords. The CPU API calls, however, are quite different between CUDA and OpenCL and require more attention. This includes handling devices, contexts, and streams, compiling and linking kernels, setting kernel arguments, execution of kernels, and memory transfers between the CPU and GPU (or between GPUs).

A difficulty in the porting process is that it involves a significant amount of changes that all must be completed before it is possible to successfully compile, run, and test the code. The Python interpreter and the CUDA compiler can be helpful in the porting process, as they will indicate the locations where code needs to be altered. A summary of key differences between CUDA and OpenCL, which directly correspond to the steps below, can be found in Table 3. The steps needed to port our simulator from PyOpenCL to PyCUDA should also be applicable for porting other codes:

1. Import PyCUDA instead of PyOpenCL;
2. Change API calls from PyOpenCL to PyCUDA, paying extra attention to context and stream synchronization;
3. Adjust kernel launch parameters. Block sizes for PyCUDA need to be 3D, and global sizes are given in number of blocks instead of total number of threads;
4. Use CUDA indexing in the kernels. Note that `gridDim` needs to be multiplied with `blockDim` to get the CUDA-equivalent of OpenCL `get_global_size()`;
5. Search and replace the remaining keywords in the kernels. Note that GPU functions in OpenCL have no special qualifier and that GPU main memory pointers need no qualifier for function arguments in CUDA.

Even though it might be straightforward to port codes between CUDA and OpenCL, there is a large difference in availability of native and third-party libraries. Built-in and specialized functions and data types (e.g., float3) are also different between CUDA and OpenCL (and between Nvidia and AMD). If your code makes use of libraries or built-in functions and data types, the porting process will be more involved, as substitutes must be found or implemented.

A difference of practical interest is that CUDA uses the (nvcc) compiler for compilation of the GPU code, whilst OpenCL uses the OpenCL driver. The CUDA compiler is a separate program that performs compilation of both CPU and GPU code and links these together into a single file. This process can take a significant amount of time, ranging from seconds to minutes depending on how complex the code is. OpenCL, on the other hand, uses the OpenCL driver for compilation, and this process is very fast in our experience. Both CUDA and OpenCL cache their compilations, so this only applies to the first time a kernel is being used. When developing GPU kernels, however, it becomes

noticeably slower to work with the compilation times of PyCUDA compared to PyOpenCL, with an added compilation overhead of about 5 s.

Table 3. Keywords, functions, and application programming interface (API) calls in CUDA and OpenCL. In many cases, a simple search and replace is sufficient to translate a program.

CUDA	OpenCL
Function qualifiers	
<code>__global__</code>	<code>__kernel</code>
<code>__device__</code>	N/A
Variable qualifiers	
<code>__constant__</code>	<code>__constant</code>
<code>__device__</code>	<code>__global</code>
<code>__shared__</code>	<code>__local</code>
Indexing	
<code>gridDim</code>	<code>get_num_groups()</code>
<code>blockDim</code>	<code>get_local_size()</code>
<code>blockIdx</code>	<code>get_group_id()</code>
<code>threadIdx</code>	<code>get_local_id()</code>
<code>blockIdx*blockDim+threadIdx</code>	<code>get_global_id()</code>
<code>gridDim*blockDim</code>	<code>get_global_size()</code>
Synchronization	
<code>__syncthreads()</code>	<code>barrier()</code>
<code>__threadfence()</code>	N/A
<code>__threadfence_block()</code>	<code>mem_fence()</code>
API calls	
<code>kernel<<<...>>>()</code>	<code>clEnqueueNDRangeKernel()</code>
<code>cudaGetDeviceProperties()</code>	<code>clGetDeviceInfo()</code>

4.2. Profile-Driven Optimization

The shallow-water equations consist of three coupled nonlinear partial differential equations that describe conservation of mass and momentum. In our case, they are formulated using η , the surface deviation from the equilibrium water level, and volume transport hu and hv along the abscissa and ordinate, respectively. Source terms represent varying bathymetry and the Coriolis forces, which take into account that we solve the equations on a rotating sphere. The equations can be used to model gravitational waves, in which the governing motion is horizontal, e.g., the ocean [6,7]. An oceanographic simulation scenario using our simulator is shown in Figure 2.

We started the analysis of the three different numerical schemes using the Nsight extension in Visual Studio and the standalone Visual Profiler application in the exact same manner as if we were profiling from C/C++. Nsight was run on a laptop with a GeForce 840M GPU in Windows, and the Visual Profiler on a desktop with a GTX780 GPU in Linux. Our workflow started by profiling the code, identifying the performance bottleneck, optimizing the bottleneck, and finally profiling to determine if the optimization was successful [47]. To ensure that our optimizations do not introduce errors in the code, we frequently ran integration and regression tests against reference solutions.

An important performance parameter for GPUs is the domain decomposition determined by the block size. CUDA decomposes the work into a grid with equally sized blocks, and all blocks are executed independently. At runtime, the GPU takes the set of blocks and schedules them to the different cores within the GPU. Using a too-small block size will underutilize the GPU, and using a too large block size will similarly exhaust the GPU's resources. Figure 3 shows how the block size has a major impact on performance for three different GPUs and also illustrates that finding the best block

size can be difficult. Because of this, we experimentally obtained the optimal configuration for each scheme before starting profiling, and again after the code had been optimized.

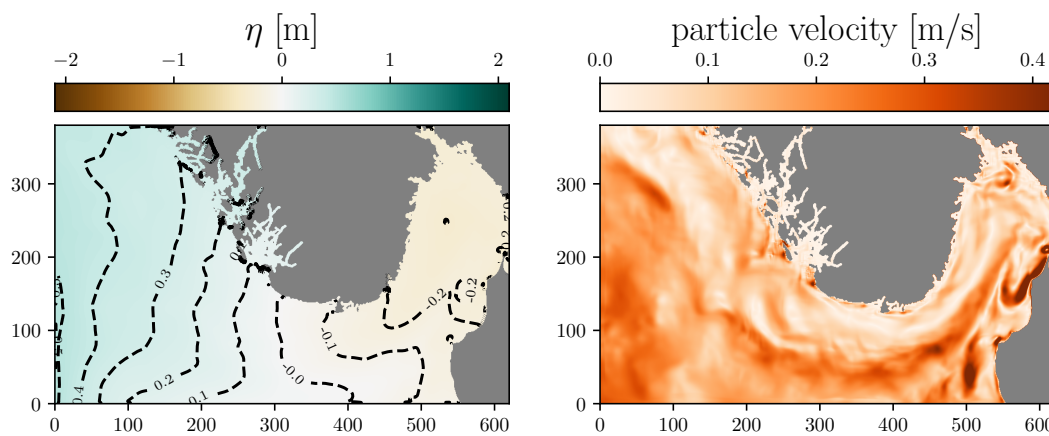


Figure 2. Example of a shallow-water simulation for oceanographic purposes with our code. The left figure shows the sea-surface level, whereas the right figure shows the particle velocity. The simulation covers the North Sea around the southern part of Norway, with Denmark in the lower right corner. The axes shows distances in km, and the grid consists of 1550×950 cells of size $400 \text{ m} \times 400 \text{ m}$. The simulation here is initialized from operational data provided by the Norwegian Meteorological Institute.




Figure 3. Heat map of performance as a function of block width and block height for selected sizes for the high-resolution scheme on three different GPUs. Notice that even though the performance patterns have similarities, the performance on the V100 would be suboptimal if the optimal configuration from the GTX780 is used. The performance increase is 2–5× for all three GPUs from the slowest to the fastest block size.

4.2.1. The High-Resolution Finite Volume Scheme

The first numerical scheme we considered is a high-resolution finite-volume method, which is designed to be well-balanced with respect to steady-state solutions in geostrophic balance [48]. The scheme reads the physical variables, (η, hu, hv) from the GPU main memory, reconstructs an intermediate set of four geostrophic equilibrium variables, calculates inter-cell fluxes, and finally sums up the fluxes and writes the result back to the GPU main memory. Each of the steps are stencil operations building on top of each other, and the kernel relies heavily on the use of shared memory for storing and sharing intermediate results. The integration in time is based on a second-order strong stability preserving Runge-Kutta method, which means that the kernel is called twice for every iteration.

The first Nsight analysis of the high-resolution scheme indicates that the occupancy is very low at only 25%. The occupancy is a measure of how many threads are resident in the cores of the GPU simultaneously and roughly translates to how well the GPU can hide memory latencies. As the GPU

has restricted resources when it comes to the amount of shared memory, number of registers, etc., the occupancy is reduced if each block uses many of these resources. In our case, the limiting factor is the use of shared memory. By reducing the amount of shared memory used by each block, we can expect to get more blocks resident simultaneously on the GPU, which most likely will increase memory throughput and thereby increase performance.

Through six consecutive iterations, we progressively reduced the shared memory by 65% with the following steps (see Figure 4):

1. Recomputing bathymetry in cell intersections instead of storing H_m ;
2. Recomputing face bathymetry instead of storing RH_x and RH_y ;
3. Reusing buffer for physical variables Q for storing the reconstruction variables R ;
4. Recomputing fluxes along the abscissa instead of storing F ;
5. Recomputing fluxes along the ordinate instead of storing G ;
6. Reusing the buffer for derivatives along the abscissa, Q_x , and derivatives along the ordinate, Q_y .

In the third and sixth steps, we managed to reuse other shared memory buffers by reordering the execution flow of the code, and for all other cases, we relied on recomputation. This essentially means that we prefer recomputing results rather than storing and sharing them between threads, thus trading extra computation for less memory storage.

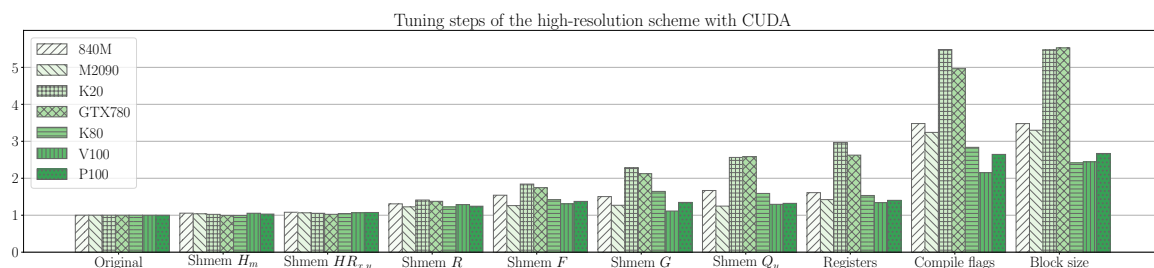


Figure 4. Different optimization stages for the high-resolution kernel showing performance normalized with respect to the original version. Notice that the different optimization strategies impact the performance very differently for different architectures. The platform we profiled on was the 840M GPU, yet the highest performance gain was for the K20 and GTX780 GPUs with over $5\times$ improvement.

The first seven groups in Figure 4 show the impact of the shared memory on performance. The first few buffers we removed were insufficient to increase the occupancy, because it would not free enough space for another resident block per core. However, after stage three, there was space for one extra block, and the occupancy increased. After all six iterations, occupancy increased to 56%, memory bandwidth utilization increased by a factor 1.8, and gigaFLOPS (GFLOPS) (FLOPS stands for floating point operations per second) more than doubled. It should be noted that memory throughput is the important factor here, both since our kernel is memory-bound and because our recomputations artificially increase the number of floating point operations performed by the kernel by 20%.

The next limiting factor was the number of registers (variables) used by the kernel. First, we removed all variables related to debugging and packed four boolean kernel arguments into a single integer using bit operations. Finally, we removed several temporary variables by aggregating these as soon as possible into the final computed flux. The compiler is already very good at optimizing register use, but using the above optimizations reduced the number of registers per thread from 49 to 47, which was sufficient to increase occupancy from 56% to 62%.

Compilation flags can also be used to increase performance with relatively little effort, as can be seen in Figure 4. Here, we used the `--use_fast_math` flag, which enables using fast, albeit less precise, mathematical functions for operations such as exponential and square roots. This gave a dramatic effect and increased memory throughput by a factor of 3.9 relative to the original version. This is not only because the mathematical functions execute faster, but perhaps more importantly because

the fast mathematical functions use less register space, reducing from 47 to 40 registers per thread. This increased occupancy to 69% and left shared memory as the bottleneck once again. In our case, the use of fast mathematical functions is sufficiently accurate, but these compilation flags can ruin the correctness of a program and should be used with care. Other important flags to consider are flags that determine the maximum amount of register and specify cache configurations. After having optimized the kernel, our block size is probably no longer optimal, and we reran the block size optimization.

The performance increased by a factor 3.5 on the 840M, and 5.5 on the GTX780 after these optimization steps. One surprising point is that the optimization steps have a very different effect on different architectures. For example, the K20 and GTX780 increased to over 5 times the performance, whereas the K80, P100, and V100 increased by only half of that. Some steps also actually decreased the performance on one GPU, whilst having a positive effect on others. This was especially noticeable with our attempts at optimizing register use. It should still be noted that none of the optimization steps that reduced shared memory usage led to a slowdown on any of the GPUs, even though these steps increased the computational workload per thread.

4.2.2. The Nonlinear Finite Difference Scheme

Our second scheme is a nonlinear second-order classical leapfrog scheme, also called centered-in-time centered-in-space. In this scheme, the three physical variables are defined on a staggered grid and are updated using separate kernels. Each kernel reads all three variables from the GPU main memory, performs computations, and writes back the result. The computational work in these three kernels is significantly less than for the high-resolution scheme, making this scheme even more memory-bound. However, this also means that only four shared memory buffers are needed (three for the variables from the current time step and one for the constant depth). The initial profiling shows that the occupancy is 100% for all three kernels.

Due to its simplicity, this type of kernel leaves less headroom for optimization than the aforementioned high-resolution scheme. For example, the compilation flag `--use_fast_math`, which had a dramatic effect for the high-resolution scheme, has only a marginal effect on this kernel because it performs less mathematical operations and uses few registers already. One thing that we can optimize, however, is using concurrent kernel execution, since all three variables can be updated simultaneously. Unfortunately, each kernel already occupies all available resources on the GPU, and therefore, the effect turns out to be marginal, as shown in Figure 5. The figure also shows that the boundary condition kernels took too long and directed our attention to them. These could with relative ease be optimized to take an insignificant amount of time.




Figure 5. Four screenshots from the Nvidia Visual Profiler, illustrating the effect of optimizations for the nonlinear scheme. The top figure shows the original performance, and in the second, we introduce multiple streams so that the three physical variables can be computed independently. In the third figure, we reduce the number of blocks that are needed for the boundary condition kernels, and finally, we compute all three variables within one kernel in the bottom figure. The horizontal time scale is about 7 ms and equal for all four figures and shows that computing all variables in a single kernel gives the highest performance gain.

Because we launched three kernels to update the variables, we read the variables η , hu , and hv multiple times from the GPU main memory for every iteration. By carefully gathering these kernels into one, we can reduce the number of variables read and written from the GPU main memory from 16 to ten. After this optimization, our profiling shows a 36% decrease in memory traffic, compared to the theoretical 37.5%. This directly translates to an equivalent improvement in performance, as shown in the lower panel of Figure 5.

By merging the three kernels into one, the occupancy decreased to 62.5%, as each thread now requires more registers. Several attempts were made to reduce this number, but none of our attempts helped the compiler to do a better job. One way of forcing fewer registers is to set the `--maxrregcount` compilation flag. However, this implies that registers are spilled to local memory (cache on the GPU), and even though the occupancy increases, the actual performance decreases. The final results from optimizing the nonlinear scheme are shown in Figure 6, and the optimizations increase the performance of the CUDA versions by between two to three times on all seven GPUs.

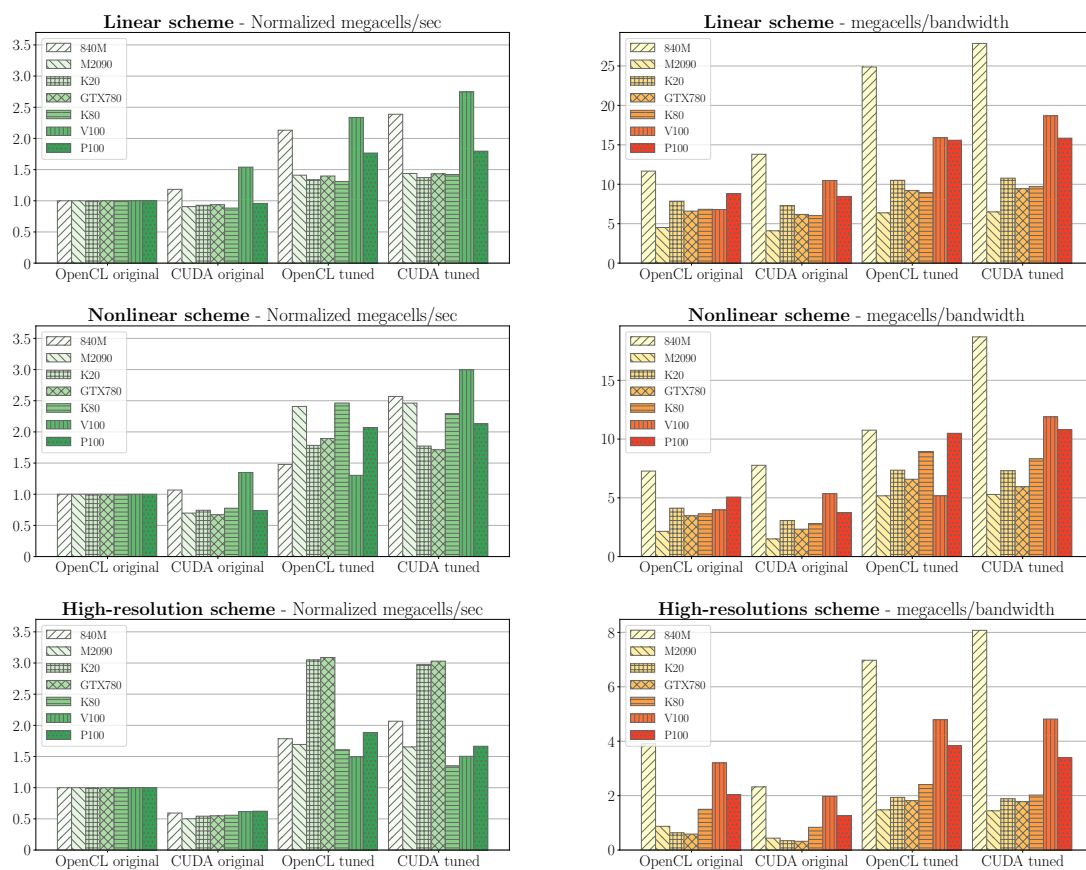


Figure 6. Performance of original, ported, and optimized kernels measured in megacells per second. The left column is normalized with respect to original performance, and the right column with respect to the theoretical GPU bandwidth. Notice that there is relatively little difference between CUDA and OpenCL, whilst there is a significant difference in how effective the tuning is for the different architectures. Furthermore, there is a significant loss of performance when porting from OpenCL to CUDA in our original approach for the high-resolution and nonlinear schemes. From our experience, this relates to how the two languages optimize mathematical expressions with and without the fast math compilation flags. Also notice that the 840M, V100, and P100 GPUs achieved the best performance relative to the GPU bandwidth.

4.2.3. The Linear Finite Difference Scheme

The final scheme solves the linearized shallow-water equations using a forward–backward linear finite-difference scheme [49] and is asymmetric in time. It consists of three simple kernels, in which the most recent results are always used (similarly to a Gauss–Seidel iteration):

$$\begin{aligned} hu^{n+1} &\leftarrow F(\eta^n, hv^n, hu^n), \\ hv^{n+1} &\leftarrow G(\eta^n, hv^n, hu^{n+1}), \\ \eta^{n+1} &\leftarrow H(\eta^n, hv^{n+1}, hu^{n+1}). \end{aligned}$$

In order to reduce the amount of memory read from the GPU main memory, we applied the same strategy as for the nonlinear scheme and carefully combined all three kernels into one. Since the execution order has to comply with the data dependencies, we needed to read extra input data (referred to as ghost cells or computational halo) for all variables at time step n . We then computed hu^{n+1} , hv^{n+1} , and finally η^{n+1} . Our profiling shows that this reduced the amount of memory read and written to GPU main memory by 50%.

The profiling now tells us that memory dependencies are the main bottleneck, and that the GPU has no eligible instructions for about 60% of the cycles. Without a major redesign of the algorithm, there is typically little we can do to improve the kernel further, but for completeness, we added the `--use_fast_math` compilation flag. However, since the kernel is simple and heavily memory-bound, the flag had a negligible impact.

The overall performance increase is more than two times for the laptop 840M GPU and the server V100 GPU, but only around 40% for the others. It should also be noted that this kernel is the one with the smallest negative performance impact when moving from OpenCL to CUDA.

4.3. Backporting Optimizations to OpenCL

To do a fair comparison of the performance of PyCUDA and PyOpenCL, we needed to back-port our optimizations to PyOpenCL. As most of the optimizations are carried out in the CUDA source code, it is a relatively simple matter of copying the optimized code into the original OpenCL kernels and updating the code according to Table 3. In the Python code, the loading of the separate kernels for the linear and nonlinear schemes was replaced by the new merged kernels, and the input arguments were updated to correspond to these optimizations. Active and appropriate use of a version control system and regression testing is crucial in this work.

Porting compiler flags can in general be a greater challenge, as there are no one-to-one overlaps between the compile-time options for the two languages. However, in our case, only the `--use_fast_math` compilation flag was used in the final PyCUDA version, and it has a matching counterpart, `-cl-fast-relaxed-math`, which is passed to PyOpenCL's API when compiling the kernels.

4.4. Comparing Performance

The overall performance gain of our optimization is shown in Figure 6, where all results are given in megacells per second normalized with respect to the original PyOpenCL implementation in the left-hand column and GPU bandwidth (see Table 2) in the right-hand column. The original porting from PyOpenCL to PyCUDA gave a noticeable reduction in performance for the high-resolution scheme on all GPUs. After careful examination, we attributed this to different default compilation flags in PyCUDA and PyOpenCL: In PyCUDA, the fast-math flag was shown to double the performance for the high-resolution scheme, while we found that it gave less than a 5% performance gain with PyOpenCL. Note that the slowdown in the original porting is much less for the linear and nonlinear schemes, as these schemes contain fewer complex mathematical operations, and we instead observe a varying effect on performance of porting the original OpenCL code to CUDA. When examining the

numerical schemes one by one, we see that the optimizations performed for the high-resolution scheme appears to be highly portable when back-ported to PyOpenCL for all GPUs. For the tuned nonlinear scheme, however, we saw that the 840M and V100 GPUs give a significantly higher performance using CUDA than OpenCL. Finally, for the linear scheme, the performance is similar for all GPUs, and only the 840M and V100 GPUs benefit significantly from the optimization effort. In total, we see that certain scheme and GPU combinations result in a significant speedup for CUDA over OpenCL, but we cannot conclude whether this is caused by differences in driver versions or from other factors. We are therefore not able to claim that CUDA performs better than OpenCL in general. When looking at the performance normalized with respect to the specific GPU bandwidth, we see that the 840M laptop GPU offers the highest performance for all schemes, followed by the two most recent high-end GPUs, the V100 and P100.

4.5. Measuring Power Consumption

We measure power consumption in two ways. The first method is using the `nvidia-smi` application, which can be used to monitor GPU state parameters such as utilization, temperature, power draw, etc. By programmatically running `nvidia-smi` in the background during benchmark experiments, we can obtain a log containing a high-resolution power draw profile for the runtime of the benchmark. The downside of using `nvidia-smi` is that information about power draw is only supported on recent high-performance GPUs, and we have therefore only benchmarked the power consumption with `nvidia-smi` on the two most recent Tesla GPUs, the P100 and V100. Further, `nvidia-smi` monitors the energy consumption of the GPU only, meaning that we do not have any information about energy consumed by the CPU. For each experiment, `nvidia-smi` is started in the background exactly 3 s before the benchmark and is configured to log the power draw every 20 ms. This background process is stopped again exactly 3 s after the end of the simulation. This approach allows us to measure the energy consumption of the idle GPU both before and after each benchmark, and we ignore the idle sections when computing the mean and total power consumption for each experiment. All results presented here are with the idle load subtracted from the experiment results.

The second method is to measure the total amount of energy used by the entire computer through a watt meter. The use of the watt meter requires physical access to the computer, and we are therefore restricted to doing measurements on the laptop and desktop, containing the GeForce 840M and GeForce GTX780 GPUs, respectively. The watt meter offers no automatic logging or reading but displays the total power used with an accuracy of 1 Wh. To get sufficiently accurate readings, we need to run each benchmark long enough to keep the GPU busy for approximately one hour, after which we read the total and mean power consumption for each experiment. Before and after each benchmark, we also record the background power of the idle system, and the maximum recorded power during the experiment, to monitor whether the operating system is putting any nonrelated background load on the computer. It should also be noted that the battery was removed from the laptop during these experiments. Similarly to the first method, we subtract the idle loads from the result of each experiment, but note that the addition increased load from the CPU will still be included.

4.6. Comparing Energy Efficiency

Figure 7 shows the results from the power efficiency experiments using the watt meter on the laptop (840M) and desktop (GTX780). The top row repeats the results for computational performance also shown in Figure 6 for the relevant GPUs, whereas the second row shows the normalized mean power consumption with respect to the original OpenCL versions. The first thing we notice is that CUDA seems to require less power on the 840M compared to OpenCL for all versions of the three schemes. On the GTX780, however, there are no differences between the two programming models for equivalent versions. In fact, only the tuned high-resolution scheme seems to be different from the others (using about 30% more power), and this behavior can also be seen on the 840M. The power efficiency of the three schemes is shown in the bottom row in Figure 7, and we see that on the 840M,

the tuned CUDA versions are the most power-efficient. This is because CUDA is both more efficient and uses less power on this system. On the GTX780, CUDA and OpenCL have equivalent power efficiency for all tuned schemes.

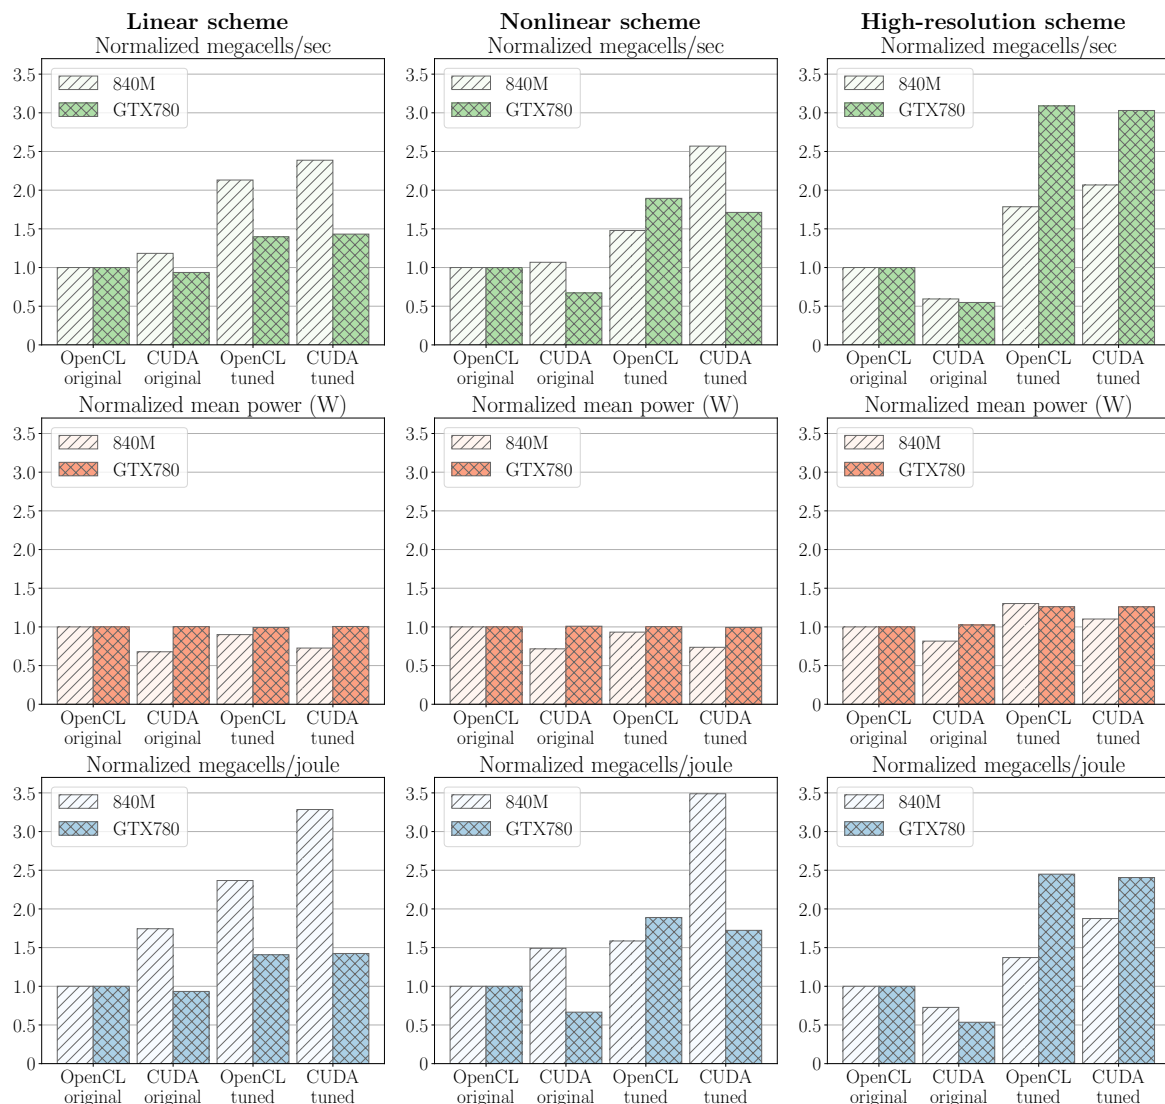


Figure 7. Comparison of original, ported, and optimized codes measured in megacells per second (top row), mean power usage (mid row), and megacells per joule (bottom row), normalized with respect to the original OpenCL implementation, for the laptop (840M) and desktop (GTX780) GPUs. The power is measured through a watt meter and represents the power consumed by the entire computer. Note that the CUDA versions require less power than the OpenCL versions on the 840M, whereas there are no differences between equivalent versions on the GTX780. In terms of power efficiency, CUDA is more efficient than OpenCL on the M840, whereas the GTX780 gives the same power efficiency.

The results for the Tesla P100 and Tesla V100 are shown in Figure 8. The top row shows the computational performance in megacells per second, repeating the results from Figure 6. The second row shows the mean power used by each version of the three schemes. Note here that on the V100, both CUDA versions for the nonlinear scheme use 60%–90% more power than the OpenCL versions, which is the opposite result compared to the 840M results. For both the linear and high-resolution schemes, the results do not differ significantly in favor of either CUDA or OpenCL, but the tuned OpenCL version uses slightly more power for the high-resolution scheme. When we consider the power efficiency in the bottom row, we see that the tuned CUDA versions are the best versions for

all schemes. In particular for the nonlinear scheme, this is mostly due to the large difference in computational efficiency between CUDA and OpenCL for this particular scheme on this particular GPU. For the P100, however, the mean power consumption is almost the same for all versions of each scheme, except for the high-resolution scheme, which has an increased mean power consumption for the tuned versions. The increase in mean power consumption is, however, less than the increase in computational performance, and the high-resolution scheme also has a higher energy efficiency when tuned.



Figure 8. Comparison of original, ported, and optimized codes measured in megacells per second (top row), mean power usage (mid row), and megacells per joule (bottom row), normalized with respect to the original OpenCL implementation, for the Tesla P100 and V100 GPUs. The power is measured through nvidia-smi and represents the power consumed by the GPU only. On the V100, there are only minor differences in mean power consumption between different versions of the linear and high-resolution scheme, but CUDA uses more power than OpenCL for the nonlinear scheme. CUDA is, however, more power-efficient than OpenCL for all three tuned schemes on this GPU. On the P100, the mean power consumption is equal for all versions of the linear and nonlinear schemes, which means that power efficiency becomes equivalent to computational efficiency. For the high-resolution scheme, we see that the mean power consumption increases for the tuned codes, but the computational performance increases more, meaning that the power efficiency is still improved by 30%–40%.

In general, we observe that all experiments show a mean power usage within about 30% of the original OpenCL versions, with the exception of the nonlinear scheme on the V100. On the other side, the computational performance increases up to 5 times (the high-resolution scheme on the GTX780). This shows that the most important factor for improving power efficiency is to increase computational performance.

4.7. Tuning Block Size Configuration for Energy Efficiency

In the previous section, all benchmarks were configured using the optimal block size configuration for computational efficiency. Here, we analyze how sensitive power efficiency is to the block size configuration, and whether there are other optimal configurations if we aim for power efficiency instead of computational performance.

Figure 9 shows the benchmark results for the three different schemes for a wide range of block size configurations. The top row shows computational efficiency in terms of computed megacells per second and confirms how sensitive the performance for all three schemes is to different configurations, similarly to the results from Figure 3. The optimal configurations with respect to computational efficiency are (16,16), (16,4), and (32,8) for the linear, nonlinear, and high-resolution scheme, respectively. The second row shows the mean power consumption for each execution. By comparing the heat map of mean power usage with the computational efficiency, we see that high computational performance does not necessarily translate into high mean power consumption. As a specific example, the most efficient configuration for the linear scheme has a lower mean energy consumption than half of the configurations. The final row of Figure 9 shows energy efficiency with respect to block size configuration, and again, we see that the results do not follow an easily predictable pattern across the different block sizes. Notice how the most energy efficient block sizes (32,24) for the linear and nonlinear scheme, and (32,32) for the high-resolution scheme) can be recognized as configurations that have particularly low mean power consumption and are less than medium fast. None of the most computational efficient configurations are among the very best configurations in terms of energy efficiency. These results show that block tuning for computational efficiency and for energy efficiency are not the same. However, because a performance tuned code has a lower time-to-solution, we can in general claim that performance tuning increases the power efficiency proportionally to performance.




Figure 9. Cont.




Figure 9. Computational performance (top), mean power consumption (middle), and energy efficiency (bottom) for the three numerical schemes for different block size configurations on the Tesla P100 GPU. Notice that all three measures change rapidly and unpredictably between the different configurations, and that there does not seem to be a direct translation between high computational performance and the mean power consumption. The bottom row shows that it is just as important to tune for optimal block size configuration when optimizing for energy efficiency as when optimizing for computational performance, and that these two objectives cannot be expected to be fulfilled simultaneously.

5. Summary

We have presented our experiences from working with CUDA and OpenCL from Python for an extensive period of time (the authors have worked with GPU computing using C++ for over ten years and over four years using Python). Our conclusion is that using Python is as computationally efficient as C++ for our use, and we believe this to be true for many other application areas as well. We also benchmarked three different OpenCL codes, our ported code in CUDA, our optimized CUDA code, and finally, our OpenCL code with optimizations found using the CUDA tools. Our results are shown for seven different GPUs, thus representing many of the GPU architectures in use today. Finally, we looked at the power consumption for all versions of the code and the potential for tuning for energy efficiency.

One of our most important findings is that working with Python is significantly faster than using C++, as we have previously done. We notice that the amount of research we are able to do in the same amount of time increases dramatically, and that the code quality is higher with fewer bugs and crashes. The combination of PyCUDA/PyOpenCL and the Jupyter Notebook led to a very productive development environment after addressing the initial crashes that forced us to reboot often. For our application area, the overhead of using Python becomes negligible with respect to performance.

The original motivation for using OpenCL was to support GPUs and similar architectures from multiple vendors. Our motivation for changing from OpenCL to CUDA was because of the better software ecosystem for CUDA, and we have been very happy with our decision. CUDA appears to be a much more stable and mature development ecosystem with better tools for development, debugging, and profiling for our hardware.

We found it interesting that our initial port from OpenCL to CUDA imposed a performance penalty, due to different default compiler optimizations. Even though some authors have reported OpenCL to be slower than CUDA, we find no conclusive results that support this in general. The performance gain varied much more with the GPU being used than whether we used CUDA or OpenCL. In most of our experiments, tuned versions of CUDA and OpenCL were found to have the same computational performance, with a few examples showing CUDA being faster than OpenCL. Additionally, we found that the performance gain of a single optimization strategy will have vastly different effects on the run time for different GPUs. Even though we profiled and optimized mainly using a laptop GPU, the highest relative performance gain was for a server class and a desktop class GPU (the high-resolution scheme on the GTX780 and K20).

There do not seem to be any clear relationships between the power consumption when comparing different schemes, optimization levels, GPUs, and programming models. When we consider power efficiency, we see that CUDA performs better than OpenCL for all tuned schemes on the Tesla V100 and GeForce 840M GPUs, whereas there are much smaller differences on the GeForce GTX780 and Tesla P100 GPUs. When we examine the impact of performance tuning on power efficiency, there appears to be a systematic and clear relationship: A fast code is a power-efficient code.

In terms of mean power consumption, we found no clear relationships when comparing different schemes, optimization levels, GPUs, or programming models, but it seems to be just as important to consider block size configurations, as we measured a factor two difference between the lowest and highest mean power for different configurations. We also saw that optimizing block size for power efficiency is not necessarily the same as optimizing for computational efficiency.

Author Contributions: The authors have all contributed to all aspects related to the work, including methodology, software, benchmarking, and preparing the manuscript. The authors have contributed close to equally to most parts of the paper, but H.H.H. has had a lead role in the aspect directly related to power efficiency and has had the role of lead and corresponding author. All authors have read and agreed to the published version of the manuscript.

Funding: This work is supported by the Research Council of Norway through grant number 250935 (GPU Ocean). The Tesla K20 computations were performed on resources provided by UNINETT Sigma2—the National Infrastructure for High-Performance Computing and Data Storage in Norway under project number nn9550k.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

GPU	Graphics Processing Unit
HPC	High-Performance Computing
CPU	Central Processing Unit
CUDA	Compute Unified Device Architecture
OpenCL	Open Compute Language
FPGA	Field-Programmable Gate Array
ICD	Installable Client Driver
SDK	Software Development Kit
API	Application Programming Interface
SWIG	Simplified Wrapper and Interface Generator
DSP	Digital Signal Processors
BLAS	Basic Linear Algebra Subprograms

DGEMM Dense General Matrix Multiplication
REPL Read-Eval-Print Loop
FLOPS Floating Point Operations per Second
NVRTC Nvidia Runtime Compilation Library

References

1. Larsen, E.; McAllister, D. Fast matrix multiplies using graphics hardware. In Proceedings of the 2001 ACM/IEEE Conference on Supercomputing, SC'01, Denver, CO, USA, 10–16 November 2001. [CrossRef]
2. Owens, J.; Luebke, D.; Govindaraju, N.; Harris, M.; Krüger, J.; Lefohn, A.; Purcell, T. A Survey of General-Purpose Computation on Graphics Hardware. *Comput. Graph. Forum* **2007**, *26*, 80–113. [CrossRef]
3. Nanz, S.; Furia, C. A Comparative Study of Programming Languages in Rosetta Code. In Proceedings of the IEEE International Conference on Software Engineering, Florence, Italy, 16–24 May 2015; Volume 1, pp. 778–788. [CrossRef]
4. Klöckner, A.; Pinto, N.; Lee, Y.; Catanzaro, B.; Ivanov, P.; Fasih, A. PyCUDA and PyOpenCL: A Scripting-Based Approach to GPU Run-Time Code Generation. *Parallel Comput.* **2012**, *38*, 157–174. [CrossRef]
5. Asanović, K.; Bodik, R.; Catanzaro, B.; Gebis, J.; Husbands, P.; Keutzer, K.; Patterson, D.; Plishker, W.; Shalf, J.; Williams, S.; et al. *The Landscape of Parallel Computing Research: A View from Berkeley*; Technical Report; EECS Department, University of California: Berkeley, CA, USA, 2006.
6. Brodtkorb, A. Simplified Ocean Models on the GPU. Available online: <https://sintef.brage.unit.no/sintef-xmlui/bitstream/handle/11250/2565319/500-Article2bText-1025-1-10-20180815.pdf?sequence=2&isAllowed=y> (accessed on 7 January 2020).
7. Holm, H.; Brodtkorb, A.; Christensen, K.; Broström, G.; Sætra, M. Evaluation of Selected Finite-Difference and Finite-Volume Approaches to Rotational Shallow-Water Flow. *Commun. Comput. Phys.* **2019**, accepted.
8. Hagen, T.; Henriksen, M.; Hjelmervik, J.; Lie, K.A., How to Solve Systems of Conservation Laws Numerically Using the Graphics Processor as a High-Performance Computational Engine. In *Geometric Modelling, Numerical Simulation, and Optimization: Applied Mathematics at SINTEF*; Hasle, G., Lie, K.A., Quak, E., Eds.; Springer: Berlin/Heidelberg, Germany, 2007; pp. 211–264. [CrossRef]
9. Brodtkorb, A.; Sætra, M.; Altinakar, M. Efficient Shallow Water Simulations on GPUs: Implementation, Visualization, Verification, and Validation. *Comput. Fluids* **2012**, *55*, 1–12. [CrossRef]
10. de la Asunción, M.; Mantas, J.; Castro, M. Simulation of one-layer shallow water systems on multicore and CUDA architectures. *J. Supercomput.* **2011**, *58*, 206–214. [CrossRef]
11. Brodtkorb, A.; Sætra, M. Explicit shallow water simulations on GPUs: Guidelines and best practices. In Proceedings of the XIX International Conference on Computational Methods for Water Resources, Urbana-Champaign, IL, USA, 17–21 June 2012; pp. 17–22.
12. Holewinski, J.; Pouchet, L.N.; Sadayappan, P. High-performance Code Generation for Stencil Computations on GPU Architectures. In Proceedings of the 26th ACM International Conference on Supercomputing, Venice, Italy, June 25–29 2012. [CrossRef]
13. OpenACC-Standard.org. The OpenACC Application Programming Interface Version 2.7. Available online: <https://www.openacc.org/sites/default/files/inline-files/APIGuide2.7.pdf> (accessed on 7 January 2020).
14. NVIDIA. NVIDIA CUDA C Programming Guide Version 10.1. Available online: <https://docs.nvidia.com/cuda/archive/10.1/cuda-c-programming-guide/index.html> (accessed on 7 January 2020).
15. Khronos OpenCL Working Group. The OpenCL Specification v. 2.2. Available online: https://www.khronos.org/registry/OpenCL/specs/2.2/pdf/OpenCL_API.pdf (accessed on 7 January 2020).
16. Brodtkorb, A.; Dyken, C.; Hagen, T.; Hjelmervik, J.; Storaasli, O. State-of-the-art in heterogeneous computing. *Sci. Program.* **2010**, *18*, 1–33. [CrossRef]
17. Huang, S.; Xiao, S.; Feng, W.C. On the energy efficiency of graphics processing units for scientific computing. In Proceedings of the 2009 IEEE International Symposium on Parallel & Distributed Processing, Rome, Italy, 23–29 May 2009; pp. 1–8.
18. Qi, Z.; Wen, W.; Meng, W.; Zhang, Y.; Shi, L. An energy efficient OpenCL implementation of a fingerprint verification system on heterogeneous mobile device. In Proceedings of the 2014 IEEE 20th International Conference on Embedded and Real-Time Computing Systems and Applications, Chongqing, China, 20–22 August 2014; pp. 1–8.

19. Dong, T.; Dobrev, V.; Kolev, T.; Rieben, R.; Tomov, S.; Dongarra, J. A step towards energy efficient computing: Redesigning a hydrodynamic application on CPU-GPU. In Proceedings of the 2014 IEEE 28th International Parallel and Distributed Processing Symposium, Phoenix, AZ, USA, 19–23 May 2014; pp. 972–981.
20. Klôh, V.; Yokoyama, D.; Yokoyama, A.; Silva, G.; Ferro, M.; Schulze, B. Performance and Energy Efficiency Evaluation for HPC Applications in Heterogeneous Architectures. In Proceedings of the 2018 Symposium on High Performance Computing Systems (WSCAD), São Paulo, Brazil, 1–3 October 2018.
21. Memeti, S.; Li, L.; Pllana, S.; Kołodziej, J.; Kessler, C. Benchmarking OpenCL, OpenACC, OpenMP, and CUDA: programming productivity, performance, and energy consumption. In Proceedings of the 2017 Workshop on Adaptive Resource Management and Scheduling for Cloud Computing, Washington, DC, USA, 28 July 2017.
22. Du, P.; Weber, R.; Luszczek, P.; Tomov, S.; Peterson, G.; Dongarra, J. From CUDA to OpenCL: Towards a performance-portable solution for multi-platform GPU programming. *Parallel Comput.* **2012**, *38*, 391–407. [[CrossRef](#)]
23. Fang, J.; Varbanescu, A.; Sips, H. A Comprehensive Performance Comparison of CUDA and OpenCL. In Proceedings of the 2011 International Conference on Parallel Processing, Taipei City, Taiwan, 13–16 September 2011; pp. 216–225. [[CrossRef](#)]
24. Gimenes, T.; Pisani, F.; Borin, E. Evaluating the Performance and Cost of Accelerating Seismic Processing with CUDA, OpenCL, OpenACC, and OpenMP. In Proceedings of the 2018 IEEE International Parallel and Distributed Processing Symposium (IPDPS), Vancouver, BC, Canada, 21–25 May 2018; pp. 399–408. [[CrossRef](#)]
25. Karimi, K.; Dickson, N.; Hamze, F. A Performance Comparison of CUDA and OpenCL. *arXiv* **2011**, arXiv:1005.2581.
26. Martinez, G.; Gardner, M.; Feng, W.-c. CU2CL: A CUDA-to-OpenCL Translator for Multi- and Many-Core Architectures. In Proceedings of the 2011 IEEE 17th International Conference on Parallel and Distributed Systems, Tainan, Taiwan, 7–9 December 2011; pp. 300–307. [[CrossRef](#)]
27. Kim, J.; Dao, T.; Jung, J.; Joo, J.; Lee, J. Bridging OpenCL and CUDA: a comparative analysis and translation. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, Austin, TX, USA, 15–20 November 2015; pp. 1–12. [[CrossRef](#)]
28. Holm, H.; Brodtkorb, A.; Sætra, M. Performance and Energy Efficiency of CUDA and OpenCL for GPU Computing using Python. In proceedings of The International Conference on Parallel Computing ParCo2019, Prague, Czech Republic, 10–13 September 2019.
29. Kaeli, D.; Mistry, P.; Schaa, D.; Zhang, D. *Heterogeneous Computing with OpenCL 2.0*; Morgan Kaufmann: Burlington, MA, USA, 2015.
30. Sanders, J.; Kandrot, E. *CUDA by Example: An Introduction to General-Purpose GPU Programming*; Addison-Wesley Professional: Boston, MA, USA, 2010.
31. University of Mannheim; University of Tennessee; NERSC/LBNL. Top 500 Supercomputer Sites. Available online: <http://www.top500.org> (accessed on 7 January 2020).
32. AMD Developer Tools Team. CodeXL Quick Start Guide. Available online: https://github.com/GPUOpen-Tools/CodeXL/releases/download/v2.0/CodeXL_Quick_Start_Guide.pdf (accessed on 7 January 2020).
33. Intel. Intel SDK for OpenCL Applications. Available online: <https://software.intel.com/en-us/opencl-sdk> (accessed on 7 January 2020).
34. Intel. Intel VTune Amplifier. Available online: <https://software.intel.com/en-us/vtune> (accessed on 7 January 2020).
35. Apple Inc. Metal Programming Guide. Available online: <https://developer.apple.com/library/archive/documentation/Miscellaneous/Conceptual/MetalProgrammingGuide/Introduction/Introduction.html> (accessed on 7 January 2020).
36. Wilson, G.; Aruliah, D.; Brown, C.; Hong, N.; Davis, M.; Guy, R.; Haddock, S.; Huff, K.; Mitchell, I.; Plumbley, M.; et al. Best Practices for Scientific Computing. *PLoS Biol.* **2014**, *12*, e1001745. [[CrossRef](#)] [[PubMed](#)]
37. Prechelt, L. *An Empirical Comparison of C, C++, Java, Perl, Python, R, and Tcl for a Search/string-Processing Program*; Technical Report; Karlsruhe Institute of Technology: Karlsruhe, Germany, 2000.
38. Lam, S.; Pitrou, A.; Seibert, S. Numba: A LLVM-based Python JIT Compiler. In Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC, LLVM '15, Austin, TX, USA, 15 November 2015; pp. 7:1–7:6. [[CrossRef](#)]

39. Behnel, S.; Bradshaw, R.; Seljebotn, D. Cython tutorial. In Proceedings of the 8th Python in Science Conference, Pasadena, CA, USA, 18–23 August 2009; Varoquaux, G., van der Walt, S., Millman, J., Eds.; pp. 4–14.
40. Kaehler, A.; Bradski, G. *Learning OpenCV: Computer Vision in C++ with the OpenCV Library*; O'Reilly Media, Inc.: Sebastopol, CA, USA, 2017.
41. Okuta, R.; Unno, Y.; Nishino, D.; Hido, S.; Loomis, C. CuPy: A NumPy-Compatible Library for NVIDIA GPU Calculations. In Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-First Annual Conference on Neural Information Processing Systems (NIPS), Long Beach, CA, USA, 19 May 2017.
42. Kluyver, T.; Ragan-Kelley, B.; Pérez, F.; Granger, B.; Bussonnier, M.; Frederic, J.; Kelley, K.; Hamrick, J.; Grout, J.; Corlay, S.; et al. Jupyter Notebooks—A publishing format for reproducible computational workflows. In *Positioning and Power in Academic Publishing: Players, Agents and Agendas*; IOS Press BV: Amsterdam, The Netherlands, 2016; pp. 87–90. [[CrossRef](#)]
43. Hunter, J. Matplotlib: A 2D graphics environment. *Comput. Sci. Eng.* **2007**, *9*, 90–95. [[CrossRef](#)]
44. Singh, R.; Wood, P.; Gupta, R.; Bagchi, S.; Laguna, I. Snowpack: Efficient Parameter Choice for GPU Kernels via Static Analysis and Statistical Prediction. In Proceedings of the 8th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems, Denver, CO, USA, 13 November 2017; ACM: New York, NY, USA, 2017; pp. 8:1–8:8.
45. Price, J.; McIntosh-Smith, S. Analyzing and Improving Performance Portability of OpenCL Applications via Auto-tuning. In Proceedings of the 5th International Workshop on OpenCL, Toronto, ON, Canada, 16–18 May 2017; ACM: New York, NY, USA, 2017; pp. 14:1–14:4.
46. Falch, T.; Elster, A. Machine learning-based auto-tuning for enhanced performance portability of OpenCL applications. *Concurr. Comput. Pract. Exp.* **2017**, *29*, e4029. [[CrossRef](#)]
47. Brodtkorb, A.; Hagen, T.; Sætra, M. Graphics processing unit (GPU) programming strategies and trends in GPU computing. *J. Parallel Distrib. Comput.* **2013**, *73*, 4–13. [[CrossRef](#)]
48. Chertock, A.; Dudzinski, M.; Kurganov, A.; Lukáčová-Medvidová, M. Well-balanced schemes for the shallow water equations with Coriolis forces. *Numer. Math.* **2017**. [[CrossRef](#)]
49. Sielecki, A. An Energy-Conserving Difference Scheme for the Storm Surge Equations. *Mon. Weather Rev.* **1968**, *96*, 150–156. [[CrossRef](#)]



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).