# A Benchmark Set of Highly-efficient CUDA and OpenCL Kernels and its Dynamic Autotuning with Kernel Tuning Toolkit

Filip Petrovič<sup>a</sup>, David Střelák<sup>a,b</sup>, Jana Hozzová<sup>a</sup>, Jaroslav Ol'ha<sup>a</sup>, Richard Trembecký<sup>a</sup>, Siegfried Benkner<sup>c</sup>, Jiří Filipovič<sup>a,c,\*</sup>

<sup>a</sup>Institute of Computer Science, Masaryk University, Botanická 68a, 60200 Brno, Czech Republic <sup>b</sup>Spanish National Centre for Biotechnology, Spanish National Research Council, Calle Darwin, 3, 28049 Madrid, Spain <sup>c</sup>Faculty of Computer Science, University of Vienna, Währinger Str. 29, Vienna 1090, Austria

# Abstract

such as OpenCL (Open Computing Language) or CUDA (Compute Unified Device Architecture) were designed. A code written in those APIs is functionally portable: it can

Email addresses: fillo@mail.muni.cz (Filip Petrovič), 373911@mail.muni.cz (David Střelák), hozzova@mail.muni.cz

422536@mail.muni.cz (Richard Trembecký)

siegfried.benkner@univie.ac.at (Siegfried Benkner), fila@mail.muni.cz (Jiří Filipovič)

Preprint submitted to Elsevier

input size, structure, or application settings, so a code optimized for some input may run sub-optimally when the input is changed [20, 43].

A costly solution to this problem is to manually optimize code for each utilized device and possibly also for multiple sizes or structures of the input. An alternative solution is a technique called *autotuning*. Autotuning allows optimizing the application's *tuning parameters* (properties influencing the application performance) in order to perform the execution more efficiently. It is a general technique with a broad range of applications, which includes areas such as network protocols, compilers, and database

<sup>\*</sup>Corresponding author

<sup>©2020.</sup> This manuscript version is made available under the CC-BY-NC-ND 4.0 license http://creativecommons.org/licenses/by-ncnd/4.0

<sup>(</sup>Jana Hozzová), 348646@mail.muni.cz (Jaroslav Ol'ha),

systems. We focus on *autotuning of code optimization parameters*, which allows changing the application at the level of its source code: from low-level optimizations such as loop-tiling or unrolling to more aggressive changes such as modification of data layout or even using a different algorithm.

In this paper, we introduce the Kernel Tuning Toolkit (KTT), which focuses on autotuning of codes written in OpenCL or CUDA. With KTT, tuning parameters change the source code in a way defined by a programmer via preprocessor macros. Thus, tuning parameters may affect virtually any property of the source code, making autotuning very powerful. KTT targets expert programmers, as potential code optimizations have to be implemented explicitly, requiring detailed knowledge of hardware architectures.

Autotuning can be performed  $offline^1$ , i.e., before the execution of a tuned code. Offline tuning is easier to implement but does not allow an application to re-tune when its environment changes. Online autotuning allows the application to tune itself during runtime by means of changing some runtime parameters. With dynamic autotuning, the application can even build the space of different variants during runtime, i. e., it is able to compile tuned kernels during the tuning process. Although several code parameters autotuning frameworks for heterogeneous computing have been introduced [3, 35, 36, 47], they are intended to be used in a standalone tuning tool, supporting offline autotuning only. On the other hand, KTT can be integrated into application code and supports also dynamic tuning.

A tighter integration into applications has been recently identified as one of the main challenges in autotuning [6]. Kernel Tuning Toolkit was designed to simplify the integration process. It acts as an intermediate layer between the application and OpenCL or CUDA API. Therefore, the application source code has to be adapted to incorporate KTT calls. However, once integrated, the application can transparently switch between execution and tuning of the kernels. For example, the application can re-tune itself if it is executed on new hardware, or start its execution with already optimized tuning parameters, and automatically start re-tuning during runtime when the input changes.

Using KTT, we have developed a benchmark set comprising ten autotuned codes. We have executed the benchmark set on multiple hardware devices, including GPUs from NVIDIA and AMD, CPU and the Xeon Phi. We prove that our autotuned implementations are efficient enough – they often reach performance close to the theoretical peak of the hardware or at least outperform the baseline (i. e., not autotuned) implementations significantly. We also show that autotuning is required to ensure performance portability of the codes.

The search for efficient tuning configurations may be challenging due to the discrete and non-linear nature of

tuning spaces [6]. Therefore, large tuning spaces are usually impossible to explore during application runtime. However, if tuning spaces a are created rationally (i.e., by an expert programmer), their exploration may be feasible even at runtime. Expert programmers have to understand the effect of tuning parameters and set reasonable boundaries to their values. For example, setting the acceptable sizes of work-groups to multiples of 32, as it is suitable for vectorization on CPUs and Xeon Phis and efficient on GPUs executing work-items in warps. We show in this paper that rationally constructed tuning spaces can be moderately-sized (thousands of configurations or less) and still contain enough good configurations required for performance portability. Such tuning spaces can be searched during application runtime without too high overhead. To prove the applicability of KTT in real applications, we demonstrate dynamic autotuning in a CUDA-accelerated 3D Fourier Reconstruction in Xmipp [43].

The paper makes the following major contributions:

- Development of dynamic autotuning techniques in the Kernel Tuning Toolkit. KTT introduces a highlevel API for kernels and data manipulation, which can be easily used and integrated into applications. It allows switching transparently between autotuning and executing tuned kernels. KTT is open-source<sup>2</sup>, fully documented, and contains many examples of its usage.
- Introduction of a benchmark set of autotuned kernels. We have conducted a benchmark set, including multiple kernels relevant for HPC, spanning across multiple application domains such as image processing, linear algebra, computational chemistry, and differential equations. We demonstrate that autotuning of optimization parameters improves the performance portability of the benchmark set across a range of different heterogeneous architectures significantly. We also show that rationally constructed tuning spaces can be searched fast enough to allow dynamic tuning in many cases.
- Demonstration of dynamic tuning with a real-word application. We show that dynamic autotuning can be used in a real-world application, such as a 3D Fourier Reconstruction. Dynamic autotuning is also demonstrated on an application performing batched matrix multiplication with varying matrix sizes. We experimentally evaluate the speed of tuning space search convergence as well as dynamic tuning overhead on these examples.

The rest of the paper is organized as follows. In Section 2, we introduce related work and compare it with our work. The main design decisions and concepts of KTT are described in Section 3. Section 4 introduces a set of

<sup>&</sup>lt;sup>1</sup>We adopt the nomenclature from [6].

<sup>&</sup>lt;sup>2</sup>https://github.com/Fillo7/KTT

ten autotunable benchmarks and evaluates their efficiency and performance portability. Dynamic autotuning is evaluated in Section 5. We conclude and sketch future work in Section 6.

# 2. Related Work

In this section, we compare our work to state-of-the art methods in autotuning in three areas: tuning targets (which properties are tuned), tuning time (when tuning is performed) and search strategies (how the tuning space is searched and evaluated).

Autotuning covers a broad range of empirically tuned parameters related to application performance, such as compiler parameters, or the runtime environment [19, 18]. Some autotuners do not required to modify the application source code, for example, compiler flags tuners [5] or MPI tuner [32]. Other tuners may change application source code in order to test different code optimization variants. We focus on the autotuners altering the code of applications in the rest of this section as they are directly related to our tuner.

Autotuning is already successfully deployed in some high-performance libraries for conventional CPUs, such as ATLAS [48] (linear algebra) or FFTW [16] (signal processing). Libraries for accelerators are also often improved by autotuning [27, 24, 28, 31]. However, those libraries use autotuners specially designed for them. Here, we are interested in generic autotuners. Frameworks for skeletons or DSLs also use autotuning to search for the best combination of the implementation variants empirically [12, 23, 14, 4]. While they cover a broader range of applications compared to autotuned libraries, they are still restricted to a particular problem domain or a set of skeletons.

Code optimizations autotuners generate multiple functionally-equivalent variants of the application source code. They may select one of the predefined variants of a tuned function [34], or generate and compile implementations according to the values of the tuning parameters. We distinguish between compiler-based tuning, where the space of code transformation is generated automatically [38, 44, 40] and user-defined code optimization parameters autotuning [35, 13, 36, 15]. User-defined code optimization parameters tuning requires expert programmers to identify and implement tuning possibilities in the source code manually (e.g., by using preprocessor macros). Even though this approach may be costly in terms of time and expertise of the programmer, it allows to explore highly diversified variants of the code, which usually cannot be generated automatically by compilers: the programmer can change virtually anything, for example, alter algorithms (e.g., use merge sort instead of quicksort) or change the data layout in the memory (e.g., use a structure of arrays instead of an array of structures).

Our Kernel Tuning Toolkit focuses on tuning of userdefined code optimization parameters. Most similar to our work are CLTune [35], AUMA [13], ATF [37, 36], and Kernel Tuner [47], which are problem domain-agnostic autotuners designed for heterogeneous computing.

Existing benchmark sets for heterogeneous computing, such as Parboil [42], SHOC [9], or Polybench/GPU [22] do not support autotuning of code optimization parameters (only work-group size can be typically changed without substantial rewriting the benchmark). To the best of our knowledge, there is no comprehensive benchmark set for code optimization parameters tuning in heterogeneous computing. In [35, 36], two benchmarks are used to evaluate code optimization parameters tuning: GEMM and 2D convolutions. Those benchmarks are also used in our benchmark set. Three benchmarks are used in [47] (one of them is the GEMM introduced in [35]) and in [13]. In our work, a set of ten benchmarks is introduced.

Some forms of dynamic autotuning are supported by problem-specific autotuners, such as SpMV tuning [20] or generic autotuners, such as Active Harmony [44]. Autotuners may also support online autotuning where usually multiple variants of code are produced in an offline phase and searched during runtime. Online tuning is easier to implement than dynamic tuning (there is no runtime compilation), but it is not practical when the number of possible code variants is high. An examples of an online tuner is SOCRATES [17, 33].

None of the frameworks for code optimizations in heterogeneous computing support dynamic tuning natively [35, 13, 36, 47]. To implement dynamic tuning with those frameworks, the programmer has to add a non-trivial amount of glue code, running the tuner during application runtime to find a better tuning configuration and then exporting this configuration into the application, typically by re-compiling OpenCL or CUDA kernels with the JIT compiler. The OpenTuner [3], another similar tuner, is a more generic and low-level tool: it allows us to tune virtually any property of the application, but a higher amount of user effort has to be invested into the integration of the tuner. OpenTuner could be used for dynamic autotuning with higher effort than [35, 13, 36, 47], since a code responsible for a tuned kernel compilation, execution, and testing has to be provided as well. On the other hand, OpenTuner would allow to use results computed by kernels during tuning, which can increase the performance of the tuned application. To the best of our knowledge, KTT is the first autotuning framework combining universal code optimization parameters tuning with native support of dynamic autotuning for heterogeneous computing.

In this paper, we extend state-of-the-art general-purpose code optimizations autotuners for heterogeneous computing with dynamic tuning. Although the concept of dynamic autotuning is well-known, it requires an architecture that hides the OpenCL or CUDA API in order to switch implementation of kernels. In addition we contribute to the state-of-the-art in autotuning by introducing a benchmark set of autotunable codes, evaluating the efficiency and performance portability of the benchmarks, and assessing how difficult it is to amortize the overheads of dynamic tuning.

The space of tuning parameters can be very difficult to search: it is discrete, non-linear, and non-convex. Although a promising method has been recently published [47], the majority of papers report that random search is often as efficient or even more efficient than more sophisticated search methods [25, 39, 35]. Therefore, it can be difficult to search large tuning spaces containing hundreds of thousands of configurations or more. Extremely large tuning spaces, however, result mainly from compiler-based autotuners such as Lift [40] or naively constructed tuning spaces. The papers [25, 39, 35, 47] focus on the analysis of tuning space search methods. However, there has not been much effort invested into studying the size of tuning spaces using a larger number of benchmarks which maintain good performance portability across a wide range of different hardware architectures. In this paper, we constructed a set of ten benchmarks and show that tuning spaces are often small enough to be searched dynamically while still providing performance portability with a nearpeak performance.

Machine learning on historical autotuning data can be used to decrease the number of tuning decisions performed during program compilation or execution. In [33], a dynamic selection from a very limited number of code variants is based on a model created from previous tuning runs. In [8], a single tuning parameter can be optimized at compilation time by a neural network trained in multiple trial runs. Contrary to those papers, we focus on multi-dimensional tuning spaces.

#### 3. Architecture of the Kernel Tuning Toolkit

In this section, we introduce the main architectural concepts and the API of the Kernel Tuning Toolkit. We are using the following terminology in the paper. A *tuning parameter* is a variable which affects the code in a user-defined way (e.g., determines loop unroll factor). The *tuning space* is a cross product of all the possible values of all tuning parameters. A *configuration* is a single point in the tuning space (i. e., assignment of concrete values to all tuning parameters), which fully determines one possible implementation variant of the tuned code. The main functionality of KTT is:

- specification of tuning parameters and constraints of tuning space;
- compiling and executing the kernel or *kernel composition* (multiple kernels and host code with shared tuning parameters);
- automatically searching the tuning space;
- managing data transfers automatically (KTT automatically creates and copies data from/to the accelerator);



Figure 1: Schematic view of KTT architecture. The dashed line shows components, which are typically active during dynamic tuning inside the main application loop.

• checking the results of the tuned kernel against a reference implementation computation.

KTT has been designed as a C++ library, which replaces direct access to the OpenCL or CUDA API. By providing a middle layer between the application and the OpenCL or CUDA API, KTT is able to perform autotuning transparently: the kernel execution and tuning can be performed by the same application code. However, in order to allow integration into real-world applications, KTT must support important functionality such as memory management, kernel configuration, execution, and synchronization provided by OpenCL or CUDA. Because KTT forms a middle layer between the application and the CUDA or OpenCL API, it can modify kernel code at runtime, transparently to the application. Moreover, this design allows switching between OpenCL and CUDA easily. When kernel codes for both APIs are provided by the programmer, the KTT is just initialized with the selected API and handles all the communication between the application and OpenCL or CUDA. Because OpenCL and CUDA use a different way to configure the parallelism of the kernel<sup>3</sup>, KTT can automatically translate parallelism configuration for the selected API. The KTT API has been derived from the CLTune project [35], so it is very similar to CLTune when we use it for offline tuning. Additionally, KTT API allows for tuning compositions of multiple kernels, tuning of how kernels are called from host code [15], and novel features for dynamic tuning.

 $<sup>^{3}\</sup>mathrm{OpenCL's}$  NDR ange describes global parallelism, whereas CUDA blocks and threads define different layers of parallelism

The architecture of KTT and its connection to the autotuned application is sketched in Figure 1. The application creates kernels and defines tuning parameters and their acceptable values (with possible constraints passed as lambda functions), and passes them to KTT, where the tuning space is built. Then, it connects input and output buffers to the kernels and starts the tuning process. KTT uses a searcher to search the tuning space and to select a configuration to be executed. In current implementation, random search, simulated annealing and Markov-chain Monte Carlo searchers are available. Then, it compiles the kernel(s) according to the selected configuration, executes and benchmarks it. If dynamic tuning is active, the results of the tuned kernel(s) can be immediately used by the application. The results can be validated against a reference implementation by KTT. The execution of kernel(s) is benchmarked and the performance results are stored in KTT, allowing the searcher to navigate the search process and the application to query for, e.g., the fastest configuration.

#### 3.1. Kernel tuning

The simplest scenario is tuning of a single kernel. In this case, the following steps have to be done in a tunable code:

- initialize the tuner;
- create handlers for kernel arguments;
- create the kernel handler;
- assign input/output arguments to the kernel;
- define tuning parameters, their acceptable values, and constraints;
- start tuning.

The tuner executes and benchmarks different tuning configurations and searches for the one, which results in the shortest kernel runtime.

In many real-world applications, some tuning parameters are shared between multiple kernels (e.g., the memory layout of some intermediate data). KTT framework allows sharing tuning parameters among kernels by using kernel compositions. Moreover, a portion of a computation can be performed on the host (e.g., a tuning parameter may determine how many times a kernel is executed or if a host code performs some pre-computation). KTT uses the *tuning manipulator* when tuning parameters influence the host code. The tuning manipulator class enables users to customize a portion of the framework's code that is responsible for kernel execution and buffer management, and optionally can perform some part of the computation directly in the C++ host code. The tuning manipulator must implement a method *launchComputation*, which can execute multiple kernels, perform computations in C++,

and transfer data between host and device. Tuning manipulators and kernel compositions allow to use tuning parameters, which cannot be implemented when kernels are tuned separately: for example, it is possible to change a format of the intermediate data exchanged between multiple kernels.

# 3.2. Offline and Dynamic tuning

KTT supports different types of autotuning depending on the time when tuning is performed and on the level of integration:

- Offline autotuning is performed prior to the execution of an application, usually by an extra utility. Offline tuning does not require integration of the autotuner into the application. The tuning utility can search for tuning parameters of the computationally most demanding application kernels and then exports values of those parameters to the build system. The disadvantage is that the tuning process cannot be easily repeated *inside* the application, i. e., during application runtime.
- Dynamic autotuning is performed during application runtime. When tuning parameters change application source code, it must be modified according to the actual values of the tuning parameters and recompiled. The application can execute autotuning at any time, e.g., when it is executed on a new hardware device or when a performance-relevant characteristic of the processed data changes<sup>4</sup>. Dynamic tuning can be performed in a blocking manner (the tuner tests several tuning configurations and selects the best one; the results of kernel executions are not passed to the application during tuning) or non-blocking manner (the result of each tested kernel variant is immediately used by the application). With blocking autotuning, KTT automatically replicates input and output arrays, so there is no side effect caused by kernel results on the application. Nonblocking tuning is more suitable for interactive applications or complex parallel workloads with many dependent tasks, where a slow response of some component may be critical to the overall performance.

The Kernel Tuning Toolkit can be integrated into application code so that the application code manages memory objects and executes kernels via the KTT API instead of directly using OpenCL or CUDA. In such a case, the application decides if KTT changes values of the tuning parameters and recompiles kernels (tuning mode) or if KTT just executes the kernels (running mode). Furthermore, the kernels' result can be used by the application even during the tuning process (a non-blocking tuning described above), which improves application performance,

 $<sup>^{4}</sup>$ With the current version of KTT, the decision about when to tune is always made by the application.

especially when the tuning overhead is low (i. e. kernels<sub>21</sub> runtime dominates compilation runtime). 22<sub>23</sub>

#### 3.2.1. Code Example

Let us assume we have two kernels, foo(a) and bar(b). The kernel foo produces a 2D array, which is used as an input for kernel bar: b = foo(a); c = bar(b);. Let us further assume a tuning parameter B\_TRANS, which determines if b is stored transposed. Clearly, the value of B\_TRANS must be the same for both foo and bar, so the kernels must be tuned together. Thus, we create a kernel composition with a tuning manipulator calling both kernels. The tuning manipulator is shown in Listing 1. The class inherited from tuning manipulator must override the method launchComputation, which is responsible for executing the two kernels via KTT in our example, but it could also implement computation in C++ or call KTT functions for data movement or synchronization.

class TunableFoobar : public ktt::TuningManipulator {
 public:

```
TunableFoobar(ktt::KernelId foo, ...):
    // assign kernels and input/output
    // to internal structures
{}
void launchComputation(const ktt::KernelId)
    override {
        // tuning parameters can be queried here
        runKernel(foo);
        runKernel(bar);
    }
private:
    ktt::KernelId foo;
....
};
```

Listing 1: Tuning manipulator

The code setting up KTT is sketched in Listing 2. It  $\frac{1}{2}$  initializes the tuner at line 2, creates kernels (lines 5-8),  $\frac{3}{4}$  their arguments (lines 11-12), and constructs a composition of the kernels (lines 16-22). The composition is cre- $\frac{6}{7}$  ated with a tuning manipulator implemented in a class 8 **TunableFoobar** (see Listing 1). The kernels are created  $\frac{9}{10}$  with an initial configuration of NDRange and work-group11 size (lines 3-4), but this configuration can be altered in  $\frac{12}{13}$  two ways: by defining the relation of NDRange/group size and some tuning parameter (using a pre-defined or lambda function), or directly in launchComputation method by any user code.

```
// Initialize tuner and kernels foo, bar
ktt::Tuner tuner(platformIndex, deviceIndex);
const ktt::DimensionVector ndRange(inputSize);
const ktt::DimensionVector workGroup(128);
ktt::KernelId foo = tuner.addKernelFromFile(
    kernelFile, "foo", ndRange, workGroup);
ktt::KernelId bar = tuner.addKernelFromFile(
    kernelFile, "bar", ndRange, workGroup);
// Creation of kernel arguments a, b, c
ktt::ArgumentId a = tuner->addArgumentVector(srcA,
    ktt::ArgumentAccessType::ReadOnly);
...
// Creation of composition and setting of arguments
ktt::KernelId compositionId = tuner.addComposition(
    "foobar", std::vector<tut:KernelId>{foo, bar, a, b, c</tu>
```

```
std :: make_unique<TunableFoobar>(foo , bar, a, b, c));
tuner.setCompositionKernelArguments(compositionId,
foo, std::vector<size_t>{a, b});
```

```
tuner.setCompositionKernelArguments(compositionId,
    bar, std::vector<size_t>{b, c});
```

```
// Addition of tuning variables
```

24

25

tuner.addParameter(compositionId, "B\_TRANS", {0, 1}); Listing 2: Tuning initialization

After the setup, we can perform kernel tuning. Here, we demonstrate non-blocking dynamic autotuning, which is performed in the main application loop as sketched in Listing 3. In our simple example, we use the variable tuningOn to specify whether dynamic tuning is performed (it can be set by some user-defined function to *true* for a fixed number of iterations, or until some predefined performance is reached). The execution of a composition calling foo and bar can be achieved by two methods: runKernel or tuneKernelByStep. The runKernel executes the composition and stores result in variable c. The execution is performed with a tuning configuration defined by the programmer; usually, the fastest configuration is used. The second method, tuneKernelByStep, also performs the computation and stores results in c, but with a new values of the tuning parameters (selected by KTT using the selected search method). If the tuning space has already been explored, the method tuneKernelByStep executes the configuration, which results in the fastest computation (so it behaves like runKernel executed with the best configuration). If the application is exploring only a subset of the tuning space, it can query the fastest known configuration via the getBestComputationResult method. The rest of the application does not need to be aware whether tuning is performed: the result c is obtained in any case.

```
while(application_run) {
    ..
    if (tuningOn)
        tuner.tuneKernelByStep(compositionId, {c});
    else {
        ktt::ComputationResult best =
           tuner->getBestComputationResult(compositionId);
        tuner.runKernel(compositionId,
            best.getConfiguration(), {c});
    }
    // c is computed here
    ...
}
```

Listing 3: Main loop performing computation

#### 3.3. Independent Queues and Non-blocking Calls

Accelerated codes often employ task-level parallelism to overlap computation on a host, computation on a device, and data movements between the host and the device. Moreover, simultaneous kernel execution may improve the performance of independent kernels when some kernels do not fully utilize the device. Task-level parallelism is realized via non-blocking kernel calls, asynchronous copy and also via multiple queues (OpenCL) or streams (CUDA).

In order to reach high performance when integrated into an application, KTT must support this functionality for the tuned kernels. Thus, it is possible to use queues (when using CUDA, KTT queues are implemented as CUDA streams) and non-blocking calls with KTT. However, during the tuning of the kernel, concurrent kernel execution

1

1

2

or non-blocking execution may bias benchmarking (e.g., with concurrent kernel execution, the host code can execute another kernel at the device where the tuned kernel is running, so the measured runtime of the tuned kernel increases). The bias in benchmarking could result in a wrong selection of the best tuning parameter values. Therefore, there are two types of task-level parallelism implemented in KTT:

- *intra-manipulator parallelism* allows simultaneous kernel execution and overlapping computations and memory copy inside a launchComputation method of a tuning manipulator;
- global parallelism also allows simultaneous kernel execution and non-blocking kernel calls at the level of the application code, so the host code can call runKernel in non-blocking mode, allowing to overlap execution of multiple manipulators, or host and device computation.

During the tuning process, global parallelism is not allowed, so only one tuning configuration is executed at a time. Therefore, benchmarking is not biased by executing another code on a computing device or in a CPU thread where KTT is running. However, tuning manipulators may still use intra-manipulator parallelism, so it is still possible to, e.g., execute multiple independent kernels in parallel, or overlap kernel execution with the data copy or the CPU code.

When the tuning process ends, KTT also allows the global parallelism so that kernels or composition calls can be overlapped with another device or host code. Note that the result of the kernel or composition is downloaded to the host memory by default, which enforces synchronization. However, the user can create persistent arguments, which are not copied to the host by KTT unless the application explicitly calls the proper KTT copy method.

# 3.4. Limitations

Recall that KTT forms an intermediate layer between a tuned application and the OpenCL or CUDA API. Therefore, it has to implement the interface to operate those APIs. The current implementation of KTT does not support all the features of CUDA and OpenCL. Due to the lack of OpenCL 2.0 implementation for NVIDIA GPUs, the OpenCL support is limited to OpenCL 1.2 with KTT. Also, some features of CUDA are not supported: texture, surface, and constant memory and cooperative grids. We believe that there is no fundamental problem to support those features in a future version of KTT.

The new features of CUDA and OpenCL, which require changes in kernel code only, do not require any explicit support in KTT, as KTT methods replace only the host API (for example, new warp-level synchronization or warpmatrix operations executed on new CUDA tensor cores can be used with KTT without any explicit support). In its current implementation, a single instance of KTT works with a single computing device. To use multiple devices (e.g., in multi-GPU machine), the programmer has to create multiple instances of KTT and partition the tuning space manually. It also implies that there is no explicit support for tuning which device is to be used for which particular kernel.

### 4. Autotuning Benchmarks

In this section, we introduce a set of ten tunable benchmarks. Each benchmark contains a C++ code, which prepares data and performs tuning with KTT, and OpenCL or CUDA code of tunable kernels. We briefly introduce their implementation and evaluate the benefits of autotuning by measuring their efficiency and assessing their performance portability. All benchmarks have been tuned for and evaluated on seven different hardware devices as listed in Table 1.

| Device                  | Architecture   | SP perf. | BW  |
|-------------------------|----------------|----------|-----|
| $2 \times$ Xeon E5-2650 | Sandy Bridge   | 512      | 102 |
| Xeon Phi 5110P          | Knights Corner | 2,022    | 320 |
| Tesla K20               | Kepler         | 3,524    | 208 |
| GeForce GTX 750         | Maxwell        | 1,044    | 80  |
| GeForce GTX 1070        | Pascal         | 5,783    | 256 |
| Radeon RX Vega 56       | GCN 5          | 8,286    | 410 |
| GeForce RTX 2080Ti      | Turing         | 11,750   | 616 |

Table 1: Devices used in our benchmarks. Arithmetic performance (SP perf.) is measured in single-precision GFlops, memory bandwidth (BW) is measured in GB/s.

#### 4.1. Tuning Parameters

With tuning of code optimization parameters, the tuning parameters can encode virtually any change of the source code. While many benchmarks contain tuning parameters performing the same type of optimization, their implementation may differ from case to case. In this section, we describe the common optimizations parameters implemented in most of the benchmarks.

#### 4.1.1. Work-group Size

On GPUs, the size of work-group allows balancing the amount of reachable parallelism (i. e., amount of workitems which can run simultaneously) and allocated resources (e. g., private and local memory consumption). In general, smaller work-groups (to some extent) allow to allocate of more resources and reduce local barrier overhead. On the other hand, small work-groups may decrease memory locality when some type of memory blocking is used. Very small work-groups may also decrease reachable parallelism due to creation of under-populated warps or due to the limited amount of work-groups which can be placed on GPU simultaneously. On CPUs, work-items are processed in a vectorized loop and thus the work-group size mainly influences the amount of consumed registers and memory locality.

The optimization of work-group size (or block size in CUDA) is a common optimization method, which may be easily implemented without re-compilation of the kernel code. However, most of the integer arithmetic required for array indexing uses the work-group size. Consequently, when the work-group size is encoded by a tuning parameter, indexing arithmetics can be optimized during compilation.

# 4.1.2. Work-item Coarsening

Work-item coarsening (or thread coarsening in CUDA) is a well-known technique [45, 41], optimizing the amount of work per work-item. On GPUs, adding more work per work-item improves private memory locality and instructionlevel parallelism. On the other hand, it also increases the number of used registers, so that the reachable parallelism can be reduced. Work-item coarsening is similar to the loop unrolling on CPUs, as each work-item (i. e., iteration of the generated vectorized loop) performs more computations.

#### 4.1.3. Caching in Local Memory

Local memory (called shared memory in CUDA) is GPU-specific hardware, which allows work-items from the same work-group to share data. It is often used as an explicit cache, where data loaded from global memory are further processed (or where data are collected before they are moved to the global memory). Local memory is faster than global memory and usually also faster than global memory cache. On the other hand, explicit caching may be challenging with more complex memory access patterns. Therefore, it may or may not be efficient to cache data in local memory.

On CPUs, there is no special hardware for local memory – data allocated in the local memory are placed in a buffer in the global memory. Therefore, there is no reason to use it for improving the speed of the code, but it can be still used to share data between work-items.

#### 4.1.4. Caching in Private Memory

Private memory (or registers in CUDA) is the fastest memory available for both GPUs and CPUs. Explicit caching in private memory speeds-up access to the data. However, it may also lead to registers spilling on both GPU and CPU architectures.

# 4.1.5. Tile Size

Memory tiling is a common technique to improve spatial or temporal locality. It is usable for direct global memory access (a tile is stored in the cache by the hardware), or explicit caching in local or private memory. The tile size may or may not be equal to the work-group size (e.g., work-items can process multiple data elements, so the tile size is an integer multiple of work-group size). Bigger tiles ensure better cache locality as long as cache capacity is not exceeded. However, with explicit caching on GPUs, bigger tiles can reduce reachable parallelism by increasing resources consumption.

### 4.1.6. Loop Unrolling

Loop unrolling is a general technique, which allows increasing instruction-level parallelism, reducing branching and simplifying array indexing by common subexpression elimination. It increases the performance of loops if there are enough registers available.

#### 4.1.7. Padding Local Memory

GPU local memory consists of multiple banks (usually 32), which should be accessed in parallel to reach the highest performance. If different data from the same bank are read, a bank conflict occurs and the access into this bank is serialized, resulting in performance degradation. Padding arrays in local memory can prevent bank conflicts in some situations. For example, parallel read of a column of a  $32 \times 32$  matrix in local memory results in a 32-way bank conflict. However, when the matrix is stored as  $33 \times 32$  array, there is no conflict in accessing columns.

### 4.1.8. Explicit Vectorization

The code performed by work-items can be written in a vectorized form. Such a case is similar to loop unrolling with slightly modified effect. With GPUs, it is easier for the compiler to generate faster vector instructions for memory access (both global and local). With CPUs, the OpenCL compiler by default performs de-vectorization and vectorization, but it can be hinted to directly translate vectorized code into vector instructions, which can help if implicit vectorization is not efficient enough. On the other hand, explicit vectorization often increases register usage on GPUs. It may also increase the amount of workload per work-group, which increases registers pressure in case local barriers are called within the kernel.

# 4.2. Benchmark Set Implementation

Here, we introduce the implementation of the benchmark set used in this paper. As the development of autotuning benchmarks is quite a time consuming task (the tuning parameters have to be identified in the code, and their effect has to be implemented), we have composed a benchmark set from already available kernels, kernels developed by our group in several projects, and kernels developed as autotuned variants of previously available nonautotuned kernels. The benchmarks set covers important computational problems spanning across multiple application domains: image processing (3D Fourier Reconstruction and 2D Convolution), linear algebra (BiCG, GEMM, GEMM Batched, Matrix transpose, and Reduction), computational chemistry (Direct Coulomb Summation) and differential equation solvers (N-body and Hotspot). Most of the benchmarks use tuning parameters for performing the optimizations introduced in Section 4.1. Table 2 shows which optimizations are implemented by which particular benchmark. Benchmarks which have been published previously are described briefly here, whereas the unpublished benchmarks are introduced in greater detail. Multiple benchmarks also implement special optimizations not listed in the table – in such case, the optimizations are mentioned in the benchmark description in this section.

The benchmark set is publicly available. Except for 3D Fourier Reconstruction, all benchmarks are bundled with the Kernel Tuning Toolkit as examples of its usage<sup>5</sup>. The autotuned version of 3D Fourier Reconstruction is currently not integrated into the production version of Xmipp, but it can be downloaded from Github<sup>6</sup>.

# 4.2.1. BiCG

BiCG is a kernel used in the biconjugate gradient method. It computes

$$q = Ap \qquad \qquad s = A^T r \tag{1}$$

where A is a matrix and p, q, r, s are vectors. We have adopted the implementation from PolyBench/GPU [22] and implemented kernel fusion and cache tiling similarly to our previous work [14]. In addition to the parameters listed in Table 2, we have created tuning parameters changing the following properties of the code:

- whether BiCG is computed by the fused kernel (loading matrix A only once), or by two separate kernels computing Ap and  $A^Tr$ ;
- the amount of work per work-group (it can iterate over multiple tiles, improving memory locality of output vectors);
- how the reduction of resulting vectors is performed (can be reduced in local memory or global memory);
- how reduction is implemented (using atomic operations, or finishing reduction in a separate kernel).

The implementation uses the tuning manipulator, as tuning parameters change the execution of kernels (e.g., when atomics are not used, an extra kernel is needed to finish computation of vectors q, s).

# 4.2.2. 2D Convolution

The 2D convolution example using  $7 \times 7$  filter is adopted<sup>7</sup> from the CLTune project [35]. The special tuning parameters determine the way of handling shared boundaries of tiles.

# 4.2.3. Direct Coulomb Summation

The direct Coulomb summation precomputes the 3D spatial grid of electric charge around a molecule, used, e.g., in molecular docking [21]. We have introduced the autotuned implementation in [15]. Here, we evaluate a 3D version of the published algorithm. The algorithm tunes, besides those mentioned in Table 2, the following parameters:

- whether input atoms are stored in global or in constant memory;
- whether input atoms are stored as a structure of arrays or as an array of structures.

#### 4.2.4. GEMM

The generalized matrix-matrix multiply (GEMM) is a standard part of BLAS [10]. Its performance is critical for many applications. We have adopted an example from the CLTune project [35] with a complex tuning space containing 241,600 configurations. The large tuning space is mainly caused by applying optimizations listed in Table 2 in multiple dimensions. Moreover, tuning parameters are provided for switching between continuous and strided access to the input matrices.

# 4.2.5. GEMM Batched

Regular BLAS implementations are optimized for large data vectors and matrices. However, some applications, such as deep learning [1], multifrontal solvers for sparse linear systems [11] or Finite Elements Method [29] require executing many instances of BLAS routines operating on very small matrices. Therefore, batched operations (i. e., grouping many BLAS calls that process small matrices together into a single call) are being developed to exploit contemporary highly-parallel hardware.

It has been shown that autotuning enables reaching near-peak performance for batched GEMM using very small matrices (up to  $32 \times 32$  elements) [30]. The implementation of batched GEMM has to be altered for different sizes of matrices [30]. We have implemented the batched GEMM kernel from scratch. It is optimized for very small matrices similarly to [30], but also for highly rectangular small matrices. Note that for small matrices, GEMM is memory-bound (it does not expose high flop-to-word ratio). Therefore, optimization strategies are different than for GEMM optimized for larger matrices, resulting in a significantly smaller tuning space. Since our GEMM Batched benchmark is optimized for small matrices only, for bigger matrices, the original GEMM benchmark should be used.

Our implementation uses highly-configurable parallelism. For output matrix of size  $m \times n$ , a work-group of size  $m \times y \times z$  is created, where y, z are tuning parameters. Parameter y defines work-item coarsening: it determines the number of work-items in the y-dimension which process one instance of matrix multiplication and hence the number of elements processed by each work-item. Parameter z

<sup>&</sup>lt;sup>5</sup>https://github.com/Fillo7/KTT/examples

<sup>&</sup>lt;sup>6</sup>https://github.com/I2PC/scipion/tree/jd\_

reconstructFourier\_KTT

 $<sup>^7 \</sup>rm Our$  code uses the same kernel and tuning space, but the application is modified to use KTT API.

| Benchmark        | WG size      | coarsening   | LM caching   | PM caching   | Tile size    | unrolling    | LM padding   | vectorization |
|------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|---------------|
| BiCG             | $\checkmark$ | $\checkmark$ | $\checkmark$ |              | $\checkmark$ | $\checkmark$ |              |               |
| 2D Convolution   | $\checkmark$  |
| Coulomb 3D       | $\checkmark$ | $\checkmark$ |              |              |              | $\checkmark$ |              | $\checkmark$  |
| GEMM             | $\checkmark$ | $\checkmark$ | $\checkmark$ |              | $\checkmark$ | $\checkmark$ |              | $\checkmark$  |
| GEMM Batched     | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |              | $\checkmark$ | $\checkmark$ | $\checkmark$  |
| Hotspot          | $\checkmark$ | $\checkmark$ | $\checkmark$ |              | $\checkmark$ | $\checkmark$ |              |               |
| Matrix Transpose | $\checkmark$ | $\checkmark$ | $\checkmark$ |              | $\checkmark$ |              | $\checkmark$ | $\checkmark$  |
| N-body           | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |              | $\checkmark$ |              | $\checkmark$  |
| Reduction        | $\checkmark$ | $\checkmark$ |              |              |              |              |              | $\checkmark$  |
| Fourier          | $\checkmark$ | $\checkmark$ | $\checkmark$ |              | $\checkmark$ |              |              |               |

Table 2: Common optimizations tuned by benchmarks. Tile size is marked when it can be configured differently than work-group size. The abbreviations used in the names of the columns are as follows: "WG" is work-group, "LM" is local memory, "PM" is private memory.

determines the number of matrix multiplication instances computed by a work-group. Caching in local memory is also implemented for the output matrix: it can be written into global memory directly, or multiple matrices can be arranged in local memory and written together (improves memory coalescing).

#### 4.2.6. Hotspot

The Hotspot kernel, used for calculating a heat distribution on a 2D surface, is based on a kernel from the Rodinia benchmark suite [7]. It implements a 2D finite differences method, which can exploit temporal locality (as it is executed iteratively). We have implemented tuning parameters listed in Table 2, and parameter allowing to tune the number of steps performed in a kernel call (balances temporal locality against redundant computation).

#### 4.2.7. Matrix Transpose

We have implemented autotuning for a tiled matrix transposition sample from NVIDIA CUDA SDK 10.0. The tuning parameters additional to those defined in Table 2 are as follows:

- transposition of work-items (work-items in a warp may read rows and store columns, or read columns and store rows);
- explicit prefetching into the cache.

# 4.2.8. N-body

The computation of gravitational forces between n bodies in space is based on the code sample from NVIDIA CUDA SDK 9.0. It computes a gravitational force between all pairs of bodies, and thus is a very compute-intensive benchmark. We have added the tuning parameters allowing tuning of how input bodies are stored (array of structure or structure of arrays) and also the optimizations defined in Table 2.

# 4.2.9. Reduction

The reduction benchmark computes the sum of all elements in an input vector. We have used the autotuned implementation from our previous work [15]. There are two special optimizations affected by tuning parameters and not listed in Table 2:

- whether the reduction is performed with at most one global barrier only by a fixed number of work-items, or iteratively by multiple kernels scaling with the size of the reduced vector;
- whether, with the fixed number of work-items, the final reduction is performed by extra kernel invocation, or by utilizing atomic operations.

# 4.2.10. 3D Fourier Reconstruction

One of the computationally demanding steps in the image reconstruction pipeline in cryo-electron microscopy is a 3D Fourier reconstruction [2]: the process when 2D samples of arbitrary orientation are inserted into the 3D volume. We have used the autotuned implementation introduced in our previous work [43]. This implementation can be tuned for specific hardware and also for specific samples resolution. In contrast to other benchmarks, 3D Fourier Reconstruction is implemented in CUDA and therefore can be evaluated on NVIDIA GPUs only.

The tuning space allows several optimizations not listed in Table 2:

- atomic writing into output volume (allows to process multiple 2D samples in parallel);
- precomputation or on-the-fly computation of interpolation weights;
- how are work-items mapped to data inside tiles of input 2D samples (optimizing cache locality).

# 4.2.11. Summary

Our benchmarks use a variety of tuning parameters, some of them common for multiple benchmarks, some of them specific for a given computational problem. The size and dimensionality of tuning spaces are summarized in Table 3. Note that the number of tuning parameters can be higher than the number of tuned optimizations described in this section, because some optimizations are

| Benchmark    | dimensions | configurations |
|--------------|------------|----------------|
| BiCG         | 11         | 5,122          |
| Convolution  | 10         | $5,\!248$      |
| Coulomb 3D   | 8          | 1,260          |
| GEMM         | 15         | 241,600        |
| GEMM batched | 11         | 424            |
| Hotspot      | 6          | 480            |
| Transpose    | 9          | 10,752         |
| N-body       | 8          | 9,408          |
| Reduction    | 5          | 175            |
| Fourier      | 6          | 360            |

Table 3: A list of the benchmarks and the size and dimensionality (i.e., the number of tuning parameters) of their tuning spaces.

| Benchmark    | bound.                | ops.      | note                        |
|--------------|-----------------------|-----------|-----------------------------|
| BiCG         | mem                   | $4a^2$    | a: width and height         |
|              |                       |           | of the input matrix         |
| Coulomb 3D   | $\operatorname{comp}$ | $  6ak^3$ | a: number of atoms,         |
|              |                       |           | k: number of grid           |
|              |                       |           | points per dimension        |
| GEMM         | $\operatorname{comp}$ | $2a^3$    | a: width and height         |
|              |                       |           | of all matrices             |
| GEMM batched | mem                   | $12na^2$  | a: width and height         |
|              |                       |           | of all matrices, $n$ :      |
|              |                       |           | number of matrices          |
| Hotspot      | mem                   | $4ia^2$   | a: width and height         |
|              |                       |           | of the input matrix,        |
|              |                       |           | <i>i</i> : number of itera- |
|              |                       |           | tions                       |
| Transpose    | mem                   | $8a^2$    | a: width and height         |
|              |                       | _         | of all matrices             |
| N-body       | comp                  | $  20n^2$ | n: number of bodies         |
| Reduction    | mem                   | 4n        | n: size of input vec-       |
|              |                       |           | tor                         |

Table 4: Number of operations performed by different benchmarks. The column "bound." distinguishes between memory-bound codes (operations in "ops." column refer to transferred bytes) and computebound codes (operations in "ops." column refer to flops).

implemented by multiple tuning parameters (e.g., if optimizations are applied to multiple buffers or multiple dimensions independently). Several benchmarks have been executed with a smaller tuning space on Radeon Vega56 because the AMD ROCm driver has been crashing with some tuning configurations (mainly using vectors of size 16 and higher loop unrolling factors). Those benchmarks are Direct Coulomb Summation, GEMM, and N-Body.

The tuning spaces of benchmarks have been defined during their development. We have not performed any *a posterior* adjustment of the tuning spaces based on the experimental evaluation (e.g., removing poorly-performing configurations). Therefore, we are able to evaluate the difficulty of searching tuning space without bias caused by experimental knowledge of well- or poor-performing configurations.

# 4.3. Efficiency of Benchmarks

If we want to study autotuning spaces (especially concerning how hard it is to search them), we should first prove that those spaces allow us to generate a code with high performance. Here, we demonstrate that our benchmarks either reach performance close enough to theoretical boundaries of the hardware or at least outperform the baseline<sup>8</sup> implementation significantly. We do not evaluate 2D Convolution here: it does not perform at peak performance, but it reaches state-of-the-art performance [35], so it can be considered efficient. We also exclude 3D Fourier Reconstruction – it is a memory latency-bound code, making theoretical performance boundaries difficult to evaluate. However, it has been shown that the autotuned implementation of our gather-based 3D Fourier Reconstruction significantly outperforms state-of-the-art scatter-based approach [43].

We define the efficiency of a benchmark as the relative performance of the benchmark with respect to the relevant hardware performance boundaries (memory or arithmetic throughput). More precisely, we use:

$$efficiency = 100 \cdot \max(\frac{\frac{MEMops}{time}}{MEMpeak}, \frac{\frac{ALUops}{time}}{ALUpeak})$$
(2)

where *time* is the runtime of computation, *MEMpeak* and ALUpeak is peak memory and arithmetic throughput of the hardware<sup>9</sup>. The *MEMops* and *ALUops* are the number of memory or arithmetic operations which are *essential* to solve the task. In other words, we count the number of operations required to solve the problem, not the number of operations required to execute the algorithm (such as array indexing, communication or computations duplicated among work-items). For example, BiCG benchmark is a memory-bound code, which essentially needs to read the input matrix A once. Therefore, even if the unfused implementation reads it twice, the number of operations is computed as the size of the input matrix in bytes divided by the runtime of the implementation. The formulas for computing ALUflops or MEMops of the benchmarks are given in Table 4.

For all benchmarks, we have measured the performance with sufficiently large data as to fully utilize the GPUs. For Batched GEMM, small matrices of size  $16 \times 16$  have been used.

The efficiency of tuned implementations is given in Table 5. The performance of Hotspot benchmark is not close to the theoretical peak, so we measure the speedup over Rodinia implementation. The number of steps per kernel invocation is exposed to the user as a parameter in

 $<sup>^{8}\</sup>mathrm{the}$  implementation we used as a basis for our benchmark, e.g., the Rodinia's Hotspot

<sup>&</sup>lt;sup>9</sup>Only half the memory bandwidth has been considered for dual Intel Xeon E5-2650 because OpenCL provides no mechanism to optimize for NUMA in the dual-socket system (pinning memory buffers and work-groups to NUMA nodes is not possible). Therefore the full system bandwidth is not available.

| Benchmark    | 2080Ti        | 1070          | 750           | K20          | Vega56        | E5-2650      | 5110P         |
|--------------|---------------|---------------|---------------|--------------|---------------|--------------|---------------|
| BiCG         | 88.3%         | 84.7%         | 81.7%         | 50.4%        | 75.6%         | 46.0%        | 6.45%         |
| Coulomb 3D   | 91.8%         | 91.4%         | 84.3%         | 43.2%        | 65.3%         | 74.2%        | 22.2%         |
| GEMM         | 79.8%         | 80.6%         | 91.1%         | 51.3%        | 96.3%         | 37.5%        | 19.7%         |
| GEMM batched | 86.8%         | 81.4%         | 90.0%         | 49.6%        | 86.0%         | 27.7%        | 20.9%         |
| Transpose    | 87.1%         | 80.2%         | 86.3%         | 64.2%        | 86.1%         | 62.5%        | 10.0%         |
| N-body       | 89.7%         | 86.6%         | 87.7%         | 40.6%        | 82.2%         | 77.7%        | 29.9%         |
| Reduction    | 68.7%         | 87.5%         | 89.4%         | 64.1%        | 71.6%         | 33.9%        | 10.1%         |
| Hotspot      | $1.35 \times$ | $1.94 \times$ | $2.06 \times$ | $1.4 \times$ | $2.88 \times$ | $1.2 \times$ | $12.8 \times$ |

Table 5: Performance of benchmarks autotuned for various hardware devices. The performance relative to the theoretical peak of devices (see Table 4) is shown for all benchmarks except for Hotspot, which is compared to the baseline Rodinia implementation.

Rodinia's implementation of Hotspot. To have a fair comparison, we have searched for the best-performing number of steps manually, testing the same values which have been tested by KTT in the autotuned version. As we can see, the performance on GPUs is very good in general. We can reach a performance close to the theoretical peak (75% or more) in most cases for all architectures except Kepler (Tesla K20), which is less efficient than other architectures in all benchmarks. The performance on dual-CPU (Xeon E5-2650) and MIC (Xeon Phi 5110P) is often far from the theoretical peak. The development of OpenCL compiler seems to be not of high priority for CPU-based systems (for example Xeon Phi is not supported in Intel OpenCL from 2015), so this result is not surprising.

Note that the performance of Coulomb 3D and N-body benchmarks has been computed differently for GeForce GTX 2080Ti: the Touring architecture seems to perform transcendental functions in parallel to FP32 instructions. Therefore, we have excluded the reciprocal square root from the computation of overall floating-point operations, using formulas  $5ak^3$  and  $19n^2$  for Coulomb 3D and N-body, respectively (see Table 4). Otherwise, the performance would be overestimated.

#### 4.4. Performance Portability

In this section, we evaluate the performance portability of benchmarks without re-tuning them – i. e., how benchmarks perform if they are executed on a different device than they are tuned for. This evaluation has been performed as follows. We have tuned all benchmarks for all devices d. Then, for each benchmark tuned for device  $d_i$ , we have measured its performance on devices  $d_j, j \neq i$ . The performance is computed as a percentage of the maximal reachable performance for the device:  $100 \frac{perf(d_j)}{perf(d_i)}$ .

Let us compute the performance portability between GeForce GTX 750 and GeForce RTX 2080Ti as an example. When BiCG benchmark is tuned for GeForce GTX 750, it reaches 65 GB/s, when tuned for GeForce RTX 2080Ti, it reaches 544 GB/s. When the code tuned for GeForce RTX 2080Ti is executed on GeForce GTX 750, it reaches 54.6 GB/s, so the performance portability is 84%. When the code tuned for GeForce GTX 750 is executed on GeForce 2080Ti, it reaches performance 381 GB/s, so the performance portability is 70%. Due to vast number of combinations, we compact the results in Table 6 in the following way: we show the average with standard deviation and worst-case performances (i) across GPU architectures, (ii) when GPU code is executed on CPU-derived architecture (CPU and MIC) and (iii) when CPU or MIC code is executed on GPU architecture.

The table clearly demonstrates that the performance is not portable in general. Although the average performance portability is not bad among GPUs, the worst-cases are showing that for some benchmarks, there are combinations of GPUs with very bad performance portability, suggesting the autotuning should be re-executed for different architectures. This is in line with related work, such as [26, 20, 43]. The performance portability is much worse in the case when the benchmarks are tuned for a CPU or MIC and executed on a GPU and vice versa. The poor portability between GPUs and CPU or MIC emphasizes the important role of tuning – although our benchmarks cannot reach peak performance on CPU or MIC, their performance is much higher than in case when the GPU-tuned code is simply executed on a CPU or MIC.

This experiment also reveals a serious limitation of functional portability with OpenCL. OpenCL guarantees the functional portability of the code if it can be executed on a device. When a kernel uses more hardware resources (e.g., number of registers per work-item) than is available on the device, it cannot be executed. It seems that this is the case of finely-tuned kernels, which often use as many resources as possible. When such kernels are executed on a device with a lower amount of resources, they fail. As it is shown in Table 6, this can happen when a code tuned for CPU, MIC, or GPU is executed on a different GPU.

# 5. Dynamic Autotuning

In this section, we experimentally evaluate dynamic autotuning for two applications: batched GEMM and 3D Fourier reconstruction. Moreover, we analytically determine the potential of dynamic autotuning for the rest of the benchmarks.

#### 5.1. Methodology

Recall that with dynamic autotuning, the tuning space is explored at runtime during application execution. There-

|              | GPU→GPU             |       |        | GPU→CPU/MIC         |       |        | CPU/MIC→GPU         |       |        |
|--------------|---------------------|-------|--------|---------------------|-------|--------|---------------------|-------|--------|
| Benchmark    | $avg\pm stdev$      | worst | failed | $avg\pm stdev$      | worst | failed | $avg\pm stdev$      | worst | failed |
| BiCG         | $89.0\%{\pm}12.3\%$ | 57%   | 1      | $44.1\% \pm 17\%$   | 28%   | 0      | $38.8\%{\pm}29.5\%$ | 11%   | 0      |
| Convolution  | $79.4\%{\pm}14.9\%$ | 55%   | 3      | $56.9\%{\pm}18.5\%$ | 33%   | 0      | $10.0\% \pm 3.6\%$  | 6%    | 1      |
| Coulomb 3D   | $95.8\%{\pm}6.5\%$  | 67%   | 0      | $84.8\%{\pm}2.7\%$  | 81%   | 0      | $23.3\%{\pm}16.9\%$ | 3%    | 2      |
| GEMM         | $83.6\%{\pm}16.4\%$ | 31%   | 0      | $18.6\%{\pm}18.5\%$ | 1%    | 0      | $22.3\%{\pm}6.6\%$  | 13%   | 2      |
| GEMM batched | $85.4\%{\pm}17\%$   | 37%   | 0      | $68.2\% \pm 13.2\%$ | 39%   | 0      | $76.7\%{\pm}22.2\%$ | 46%   | 1      |
| Hotspot      | $80.3\%{\pm}17.5\%$ | 46%   | 3      | $70.3\%{\pm}15.6\%$ | 44%   | 0      | $65.1\%{\pm}8.9\%$  | 59%   | 6      |
| Transpose    | $85.0\%{\pm}21.9\%$ | 8%    | 3      | $51.0\%{\pm}27.1\%$ | 11%   | 0      | $34.7\%{\pm}14.7\%$ | 14%   | 0      |
| N-body       | $78.8\%{\pm}24.2\%$ | 2%    | 3      | $45.9\%{\pm}30.1\%$ | 0%    | 0      | $25.7\%{\pm}15.6\%$ | 6%    | 2      |
| Reduction    | $88.4\%{\pm}24\%$   | 12%   | 3      | $53.1\%{\pm}17.4\%$ | 26%   | 0      | $68.3\%{\pm}23.8\%$ | 37%   | 1      |
| Fourier      | $74.5\%{\pm}30\%$   | 31%   | 0      | N/A                 | N/A   | N/A    | N/A                 | N/A   | N/A    |

Table 6: Relative performance of benchmarks ported across GPU architectures, from CPU/MIC to GPU and from GPU to CPU/MIC, without re-tuning. Avg $\pm$ stdev denotes the average and standard deviation of the relative performance, worst shows the worst-case performance, and failed shows the number of cases when some configuration cannot be executed on a device. 3D Fourier Reconstruction has been executed on samples of 128 × 128 pixels on NVIDIA GPUs except for K20.

fore, the implementation variants are compiled and benchmarked during application run-time, resulting in four sources of overhead:

- compilation of OpenCL or CUDA kernels (each explored tuning configuration needs to be compiled by the JIT compiler);
- execution of slower kernels (slower kernels prolong tuning time even in non-blocking autotuning when their results are used for computation);
- enforced global synchronization between tuning runs (during autotuning, execution of the tuning manipulators is not overlapped, see Section 3.3);
- testing kernel output (this step is optional).

These overheads are relevant during the tuning phase only (i. e., when new configurations are searched). However, when a sufficient number of tuned kernel invocations is performed after the tuning, the overhead becomes negligible. Here, we want to know how long the dynamicallytuned code has to run to amortize the tuning overhead under a certain value. Or, alternatively, if dynamic autotuning reaches better performance than a code that has been offline-tuned for a different device or input data.

Note that the overhead of the KTT API (mainly the execution of manipulator in function runKernel) is negligible during the execution of the tuned code.

# 5.2. Batched GEMM

In the previous section, we have introduced an autotuned kernel for batched multiplication of very small matrices. This kernel is tuned for fixed sizes of matrices, e.g., we have used matrices of size  $16 \times 16$  for experiments in Section 4. However, the space of the possible matrix sizes is large. The GEMM kernel performs  $C = A \cdot B$ , where A is an  $i \times j$  matrix, B a  $k \times i$  matrix, and C a  $k \times j$ matrix. Considering small matrices of sizes up to 32 in each dimension i, j, k, we get 32,768 combinations of the sizes. Consider an application or library, which does not know the sizes of multiplied matrices before it is executed. It would be impractical to offline tune the application or library for all possible sizes, so dynamic tuning, performed at runtime once the matrix size is fixed, is of high practical value.

#### 5.2.1. Implementation

We have prepared an experiment, which simulates a real application changing the matrix size from time to time. Our testing application<sup>10</sup> executed the tunable implementation of batched GEMM introduced in Section 4.2.5, but it periodically changes the size of matrices and performs dynamic tuning. More precisely, the application computes batched GEMM in a loop and randomly changes sizes  $i, j, k \in [2, 32]$  every 30 seconds. The application does not save the results of dynamic autotuning, so every time the new sizes are used, the autotuning starts from scratch. The batch size has been selected so that matrices occupy approximately 900 MB of memory (so enough parallelism is also exploited with very small matrices). When the size of matrices is changed, the dynamic tuning using random search starts and is performed until (i) a configuration reaching 75% of the peak memory bandwidth is found, or (ii) 20 configurations have been explored. The first rule allows the application to stop tuning when a configuration resulting in a sufficient performance is reached (we call such a configuration as well-performing configuration). The purpose of the second rule is to stop tuning when a configuration performing close to the theoretical peak cannot be easily found (or does not exist at all). After the tuning is stopped, the computation with the fastest tuning configuration continues until the sizes of matrices are changed again. With an application changing matrices sizes less often, the tuning time could be prolonged. We have measured the performance without tuning overhead (showing whether efficient kernels can be found under limited tuning budget) and with tuning overhead (showing

<sup>&</sup>lt;sup>10</sup>https://github.com/Fillo7/KTT/blob/master/examples/ gemm\_batch/demo.cpp

the real performance of the dynamically tunned application). The time required for initialization and copying of newly created matrices was not benchmarked.

5.2.2. Evaluation

| Device   | Maximum              | Restricted | Incl. overhead |
|----------|----------------------|------------|----------------|
| E5-2650  | $24.5\mathrm{GB/s}$  | 88.6%      | 82.9%          |
| 5110P    | $22.9\mathrm{GB/s}$  | 82.1%      | 72.1%          |
| K20      | $91.2\mathrm{GB/s}$  | 92.7%      | 61.3%          |
| GTX 750  | $63.0\mathrm{GB/s}$  | 91.4%      | 87.8%          |
| GTX 1070 | $205.7\mathrm{GB/s}$ | 97.2%      | 94.3%          |
| Vega 56  | $308.5\mathrm{GB/s}$ | 86.2%      | 74.4%          |
| 2080Ti   | $523.4\mathrm{GB/s}$ | 92.6%      | 85.3%          |

Table 7: Dynamically tuned Batched GEMM on different computational devices. The second column (Maximum) shows the average performance of the fastest configurations (average for all tested matrix sizes). The third column (Restricted) shows averaged relative performance (relative to the maximum) of configurations reachable with dynamic tuning under limited tuning budget (at most 20 configurations explored). Finally, the fourth column (Incl. overhead) shows the averaged relative performance of dynamically tuned code, including tuning overhead.

We have run the experiments for 3,000 seconds (i.e. 100 changes of matrix sizes) with all devices used in this paper. The matrix sizes have been selected randomly, but the same sizes are used for all devices. We have also performed offline tuning with exhaustive search for those sizes to obtain performance of the fastest configuration at each device. The performance with and without tuning overhead is computed as the relative performance of the fastest configuration found by the offline tuning. The results averaged over all matrix sizes are shown in Table 7. The code executing on dual Xeon E5-2650, Xeon Phi 5110P and Tesla K20, is compiled on Xeon E5-2650, which has quite poor single-core performance and therefore requires a longer time for compilation. The prolonged compilation time does not limit performance on CPU and MIC significantly, because average kernels' runtime is high and therefore, the compilation does not induce significant overhead. On the other hand, the compilation overhead is quite noticeable with Tesla K20. The batched GEMM kernel performs in general very well (i.e., it is close to the theoretical peak) on GPUs except for Tesla K20. Its performance with overhead is also quite close to the peak kernel performance (85% or better) in case of GeForce GTX 750, GeForce GTX 1070 and GeForce RTX 2080Ti, using Core i7-8700 for compilation. A bigger gap between the performance of the fastest kernels discovered during the tuning and performance with overhead can be seen with Radeon RTX Vega 56 (running in a system with Ryzen 7 1700). The main reason is the higher number of poorly performing configurations and, therefore, more tuning steps required to find a well-performing configuration (i.e., the configuration within 75% of peak memory bandwidth, which leads to finalization of the tuning).

For better illustration of tuning performance, the first 300 s of the benchmark execution is shown for well-per-



Figure 2: Performance of dynamically tuned batched GEMM on GeForce GTX1070 + Core i7-8700. The sizes of matrices are changed every 30 s. Performance of actually executed kernels is depicted as dots, whereas lines show performance including overhead. The maximal performance reachable via offline tuning with exhaustive search is shown as horizontal red lines.



Figure 3: Performance of dynamically tuned batched GEMM on Tesla K20 + Xeon E5-2650. The sizes of matrices are changed every 30 s. Performance of actually executed kernels is depicted as dots, whereas lines show performance including overhead. The maximal performance reachable via offline tuning with exhaustive search is shown as horizontal red lines.

forming GeForce GTX 1070 in Figure 2 and for Tesla K20, which suffers from tuning overhead, in Figure 3. Naturally, the performance including tuning overhead drops, when a new matrix size is used and increases in time as tuning overhead is amortized. It can be seen that GeForce GTX 1070, coupled with the modern processor Core i7-8700, is capable of amortizing tuning overhead in a very short time - after 30 seconds of execution, performance with tuning overhead is close to peak kernel performance. Even if multiple configurations are searched before a well-performing one is found (see performance between 270 and 300 seconds for GeForce GTX 1070 in Figure 2), the performance with tuning overhead is close to the performance of the best kernel found during the autotuning. On the other hand, Tesla K20 cannot reach high performance for many matrix sizes. This also means that more configurations

|        | 2080Ti | 1070 | 750  | 680  |
|--------|--------|------|------|------|
| 2080Ti | 100%   | 99%  | 31%  | 49%  |
| 1070   | 99%    | 100% | 31%  | 50%  |
| 750    | 43%    | 67%  | 100% | 94%  |
| 680    | 60%    | 72%  | 71%  | 100% |

Table 8: Performance portability of 3D Fourier reconstruction with  $128 \times 128$  samples. The rows represent GPUs used for offline tuning; the columns represent GPUs used for execution. The percentage shows how performance differs compared to the code using the best combination of tuning parameters (for example, the code tuned for GeForce GTX 1070 and executed on GeForce GTX 750 runs at only 31% of the speed of the code both tuned and executed on GeForce GTX 750).

|         | 128x128 | 91x91 | 64x64 | 50x50 | 32x32 |
|---------|---------|-------|-------|-------|-------|
| 128x128 | 100%    | 100%  | 77%   | 70%   | 32%   |
| 91x91   | 100%    | 100%  | 76%   | 68%   | 33%   |
| 64x64   | 94%     | 94%   | 100%  | 91%   | 67%   |
| 50x50   | 79%     | 78%   | 98%   | 100%  | 86%   |
| 32x32   | 65%     | 67%   | 80%   | 92%   | 100%  |

Table 9: Performance portability on GeForce GTX1070. The rows represent samples resolution used for offline tuning, the columns represent samples resolution used for execution. The percentage shows relative performance compared to the code autotuned for the used resolution.

are searched during the autotuning process. The system with K20 uses an older Xeon E5-2650, which also prolongs kernel compilation time. Therefore, the overhead of the tuning process is significant and would need more kernel invocations to amortize. Tesla K20 is also not able to reach high performance under limited tuning budget (see the difference between kernel performance and performance reachable by offline tuning between 210 s and 240 s, or between 270 s and 300 s for example).

#### 5.3. 3D Fourier reconstruction in Xmipp

We have introduced the CUDA-based GPU acceleration of the 3D Fourier reconstruction in [43], where KTT has been integrated into the 3D Fourier reconstruction for offline tuning and the values of the tuning parameters for various hardware have been manually exported into the production code. The implementation requires autotuning to maintain performance portability across GPUs (see Table 6). Because we were unable to install Xmipp on a system with Tesla K20, we have run the benchmark also on GeForce GTX 680 to get more comprehensive results.

Detailed performance portability across hardware devices is in Table 8. The size of the samples inserted into a 3D domain influence the selection of optimal tuning parameters. The tuning has to be repeated for samples of different sizes. Otherwise, suboptimal performance is obtained, as can be seen in Table 9.

### 5.3.1. Implementation

The pseudocode of the reconstruction is shown in Algorithm 1. The output volumes G (Fourier transform of

the volume) and W (weights for the 3D voxels) are initialized at the beginning of the computation. In the loop body (lines 4-6), the samples are added into the 3D volume. More precisely, the samples are packed into buffers of a predefined size, and their Fourier transform is computed on a CPU (line 4), copied into GPU memory (line 5) and then tuned GPU kernel is executed to insert the samples into volumes G, W (line 6). The whole algorithm is discussed in detail in [43].

|          | ALGORITHM 1: 3D reconstruction                                  |
|----------|---|
|          | Input: s  |
|          | Output: $G, W$  |
| 1        | zero-initialize output volumes $G, W$ in GPU memory;            |
| 2        | initialize buffer of image's Fourier Transform $s_f$ in GPU     |
|          | memory;   |
| 3        | foreach $s \in S$ do  |
| 4        | $s_f \leftarrow FFT(s)$ on CPU;                                 |
| <b>5</b> | upload $s_f$ to GPU;  |
| 6        | insert projections of $s_f$ into $G, W$ ;                       |
| 7        | end   |
| 8        | download $G, W$ to CPU memory;                                  |
| 9        | apply weights $W$ to $G$ and perform inverse transform of $G$ ; |
|          |   |

We use the non-blocking dynamic autotuning, which performs both tuning and computation at the same time. Therefore, all input and output data have to be managed by KTT during the whole program execution. In the loop in line 3, the data are prepared on CPU, then a KTT method tuneKernelByStep, which launches one step of dynamic tuning, is executed. The method selects a new combination of the tuning parameters and executes a tuning manipulator. The tuning manipulator implements lines 5 and 6 of the algorithm. It first copies buffer  $s_f$  into GPU memory and then executes a kernel, which inserts samples from  $s_f$  into volumes G, W. The CPU code is multithreaded, allowing to overlap computation of FFTs with kernel execution. The manipulator uses CUDA streams, so when tuning is finished (and therefore global parallelism is allowed, see Section 3.3) copying buffers may be executed in parallel with kernel execution and even multiple kernels may be executed in parallel (especially when processing small samples). There is a tuning parameter changing whether atomic writes to output volume in global memory are allowed (see Section 4.2.10). Depending on the tuning parameter value, the tuning manipulator method executes kernel iteratively for each projection (atomic writes are disabled), or just once (processing all samples in buffer  $s_f$ in one kernel call).

# 5.3.2. Evaluation

We have designed an experiment demonstrating the usability of dynamic autotuning with the 3D Fourier reconstruction. We have used a real-world setup, performing reconstruction of the Brome Mosaic Virus [46] (EMPIAR entry 10010), processing 1,826,160 samples in resolution  $156 \times 156$ . GPU kernels are processing 1500 samples at

|         | best runtime | tuning 50      | tuning full |
|---------|--------------|----------------|-------------|
| 2080 Ti | 1 m 40 s     | $88\%\pm3\%$   | 54%         |
| 1070    | 5m49s        | $96\%\pm2\%$   | 79%         |
| 750     | 16m59s       | $92\% \pm 4\%$ | 72%         |
| 680     | 15m12s       | $94\%\pm2\%$   | 75%         |

Table 10: The relative performance of dynamically-tuned 3D Fourier reconstruction. The best runtime is measured with  $\bar{0}r\bar{a}culum$ , i.e., the fastest kernel is selected immediately, and no tuning is performed. The relative performance of tuning with searching 50 configurations and with searching the entire tuning space is measured as a percentage of the best runtime. Results for "tuning 50" are shown as an average and standard deviation, whereas other results are shown as an average only (their performance is very stable across multiple executions).

once [43]; therefore, about 1280 kernels are executed to solve the reconstruction (the actual number can be slightly higher due to a small percentage of void samples). All experiments with different GPUs have been performed on a desktop machine with Intel Core i7-8700.

In our experiment, the tuning is performed at the beginning of the computation, when both used hardware and sample size are known. The performance of the dynamically tuned code is compared to the performance of code with ōrāculum (i.e., when the optimal tuning configuration obtained by the offline tuning using exhaustive search is known at the beginning of the computation). We have measured dynamically tuned code in two settings. First, we let KTT perform 50 search steps with random search and then continue with the fastest kernel explored. Second, we perform the exhaustive search and continue with the optimal configuration. As the random search was used, the experiment has been repeated 100 times. Results of this experiment are shown in Table 10. As we can see, the performance penalty of dynamic tuning is smaller than the performance penalty we get for a code that was tuned offline for a different hardware device (see Table 8) or different input size (see Table 9). The performance obtained with dynamic tuning ranges between 88% and 96% of the performance of code with oraculum when 50 configurations are searched, whereas the code mistuned for different GPU can perform at 31% of  $\bar{o}r\bar{a}$  culum in the worst case (see Table 8) and the code tuned for different input size can perform at 32% of  $\bar{o}r\bar{a}$  culum in the worst case (see Table 9).

We further analyze the overhead of dynamic autotuning. Obviously, the more executions of the kernel (in our case, the more samples used to reconstruct the 3D volume), the less overhead of dynamic autotuning. Therefore, for more complicated reconstructions, the performance of dynamically tuned code is closer to the code using the  $\bar{o}r\bar{a}$ culum, whereas trivial reconstruction may suffer from dynamic tuning overhead. Adding more work per kernel (in our case using larger samples) decreases relative overhead of the compilation, but not the overhead caused by the execution of slower kernels and synchronization.

In our experiment, the JIT compiler runs for 45.5 seconds

when the full tuning space is searched. It introduces significant overhead in the experiment with GeForce RTX 2080Ti, as the GPU is very fast – the whole reconstruction with ōrāculum is finished in 1 minute 40 seconds. With all GPUs, some slowdown is caused by executing slow kernels. The performance of average kernel is at 45% of the fastest kernel for RTX 2080Ti, 69% for GeForce GTX 1070, 46% for GeForce GTX 750 and 52% for GeForce GTX 680. The good average performance on GeForce GTX 1070 improves the high relative speed of dynamic tuning with 50 explored configurations.

We have also measured the overhead of enforced global synchronization. Recall that in such case, the tuning manipulator copies input samples to the GPU and executes GPU kernels without the overlay with another manipulator instance. The overhead is small for 3D Fourier reconstruction: for kernels executed with enforced global synchronization, it is 7% for GeForce RTX 2080Ti and GeForce GTX 1070, and 5% for GeForce GTX 750 and GeForce GTX 680. The global synchronization is enforced when kernels are tuned, but is not needed when tuning is finished. Thus, in our setup, out of the total 1280 kernel executions, only those 50 launched by the tuning manipulator were slowed down.

To conclude, even if the reconstruction program runs in minutes only, dynamic tuning is able to reach better performance than offline tuning in the case offline tuning was performed for different hardware, or different input size.

# 5.4. Dynamic tuning of the benchmark set

The suitability for dynamic tuning for all benchmark can be estimated analytically. We can compare the performance of the best kernel with the average performance of all kernels produced by the tuning space, which allows us to compute the overhead caused by executing slower kernels. We cannot evaluate the relative overhead of kernel compilation, as it depends on application workload (large kernel input prolongs kernel runtime, whereas compilation time remains the same). We also cannot consider the overhead caused by the enforced global synchronization during tuning as it is highly application-dependent if overlapping of manipulators can be leveraged. The performance penalty of enforced synchronization, as well as kernels compilation, is similar for all tuning variants, whereas the performance penalty of slow kernels can be much higher (some tuning variants can produce orders of magnitude slower kernels). Therefore, we consider the overhead of very slow kernels as the most significant one.

In the following we show how to estimate the number n of kernel invocations required in order to amortize the tuning overhead such that the performance of n kernel invocations including the dynamic tuning overhead is a certain fraction of the performance we would have achieved by executing the application using a well-performing kernel n times. We define a well-performing configuration as a configuration producing a kernel with a performance with

which we are satisfied. Note that the well-performing configuration can be easily determined with some benchmarks (e.g., when defined as a percentage of relevant hardware theoretical peak), but it can be also virtually impossible to identify a well-performing configuration until the whole tuning space is searched (e.g., when defined as a configuration reaching a certain fraction of the best configuration performance). In this section, we use a well-performing configuration as a theoretical concept, which is used to determine the number of steps needed to amortize overhead of dynamic tuning.

Let the application with  $\bar{\text{ora}}$ culum be such an application where a well-performing configuration is known at the beginning of its execution (e.g., obtained by previously performed offline tuning). Let the required performance of the dynamically tuned application relative to the performance of the application with  $\bar{\text{ora}}$ culum be rp (so rp = 1.0 means that the dynamically tuned application runs at the same speed as the application that uses some well-performing kernel found during offline tuning). Let the number of tuning steps be s, the average runtime of kernels within the tuning space be  $t_{avg}$  and the runtime of the well-performing kernel be  $t_{well}$ . Then, an average<sup>11</sup> value of rp is computed as:

$$rp = \frac{s \cdot t_{avg} + (n-s) \cdot t_{well}}{n \cdot t_{well}}$$
(3)

The average number of kernel invocations n required to reach relative performance rp can be estimated as:

$$n = \frac{rp \cdot s \cdot \left(\frac{t_{avg}}{t_{well}} - 1\right)}{1 - rp} \tag{4}$$

For example, if the average kernel has runtime  $t_{avg} = 10 ms$ , the well-performing kernel has runtime  $t_{well} = 5 ms$ , we perform s = 100 tuning steps and we want to reach relative performance rp = 0.9 (i. e., dynamic autotuning reach 90% of the performance of an application with  $\bar{o}r\bar{a}$ culum), we need to perform 900 kernel invocations (including those used for tuning).

The real amortization of dynamic autotuning depends on the number of tuning steps required to find a wellperforming kernel. When random search is used, the number of required tuning steps can be computed as follows. Let r be the ratio of well-performing configurations in the tuning space and p be the required probability of finding a well-performing configuration. The number of tuning steps s, which leads to reaching the well-performing configuration with probability p, can be computed as

$$s = \log_{1-r}(1-p) \tag{5}$$

For example, if the ratio of well-performing configurations is r = 0.01, then we need to explore 230 configurations in order to reach a well-performing configuration with probability p = 0.9. Using Equations 4 and 5, we can compute the number of kernel invocations needed to hide overhead caused by executing slow kernels during dynamic tuning. We have set up the following experiment. We define a well-performing configuration as a configuration, which leads to at least 95% of the best configuration performance. Using data gathered from the offline tuning of our benchmark set with exhaustive search, we have computed the number of tuning steps required to find a well-performing configuration with probability 0.9 (using Equation 5). Then, we have computed the number of kernel executions required to decrease dynamic tuning overhead under 10% (i.e. to obtain at least 90% of the performance we would have with  $\bar{o}r\bar{a}culum$ ). The results are given in Table 11.

Table 11 demonstrates that dynamic tuning is feasible even for short program executions (with thousands or tens of thousands of kernel calls) with multiple benchmarks, such as BiCG, Direct Coulomb Summation, Batched GEMM, Hotspot, Transpose, N-body, Reduction and 3D Fourier Reconstruction. Longer execution is needed for 2D Convolution and GEMM benchmarks. This test also shows some interesting differences between the hardware devices used in the test. For example, autotuning of OpenCL code for a CPU is similarly demanding as for GPUs with many benchmarks, whereas it is much harder on the MIC (Xeon Phi) in multiple cases. There are also some benchmarks where different hardware performs highly differently. For example, with Batched GEMM on GeForce GTX 1070, 304 configurations out of 424 are within 95% of the optimum and the average kernel performance is at 95% of the optimum, so it is very easy to find a well-performing kernel and to amortize tuning overhead. With GeForce 750, only 40 configurations produce well-performing kernels, and the average kernel performance is within 62% of the best one, so tuning is harder than for GeForce GTX 1070. With Xeon E5-2650, only three tuning configurations result in a well-performing Batched GEMM kernel, and the average kernel performance is at 51% of the best kernel, so searching a well-performing kernel and amortizing the tuning overhead is significantly harder on the CPU. Interesting differences between GPU and CPU/MIC can be seen in 3D Coulomb Summation benchmark, where tuning for GPUs is harder. Not only is the number of wellperforming kernels different (e.g., 110 for Xeon E5-2650 and 28 for GeForce GTX 1070), but a more significant difference is the average performance – it is much lower for all GPUs (e.g., 5% of the best one on GeForce GTX 1070 vs. 57% on Xeon E5-2650). The poor average speed on GPU is caused by huge register spilling when high unrolling of the inner loop is used. Although it would be easy to remove these slow-performing configurations from the tuning space, we decided to keep the space as it was designed when the benchmark was developed instead of adding a posteriori information for tuning.

<sup>&</sup>lt;sup>11</sup>Here, rp is computed for the average situation with s tuning steps and random search. Obviously, the tuning time may be different from  $s \cdot t_{avg}$  in particular executions.

| Benchmark      | 2080Ti  | 1070    | 750         | K20       | Vega56       | E5-2650    | 5110P      |
|----------------|---------|---------|-------------|-----------|--------------|------------|------------|
| BiCG           | 10,383  | 9,425   | 33.090      | 43,552    | 42,499       | $32,\!338$ | 516,783    |
| 2D Convolution | 265,297 | 98,966  | $197,\!550$ | 165,783   | 99,087       | 7,211      | $3,\!435$  |
| Coulomb 3D     | 17,305  | 16,346  | 4,911       | 5,289     | 117*         | 150        | 631        |
| GEMM           | 20,309  | 151,564 | 764,485     | 205,122   | $18,782^{*}$ | 384,309    | 3,106,384  |
| GEMM batched   | 2       | 2       | 110         | 214       | 440          | 2,341      | 1,214      |
| Hotspot        | 4,314   | 4,467   | 3,309       | $5,\!635$ | 1,489*       | 3,926      | 7,346      |
| Transpose      | 9,398   | 347     | 2,998       | 1,347     | 140,177      | 5,129      | $60,\!688$ |
| N-body         | 7,539   | 33,553  | 2,531       | 20,694    | $554^{*}$    | 2,472      | 1,669,559  |
| Reduction      | 646     | 78      | 40          | 218       | $2,\!198$    | $1,\!650$  | 19,425     |
| 3D Fourier     | 2,239   | 830     | 3,123       | N/A       | N/A          | N/A        | N/A        |
|                |         |         |             |           |              |            |            |

Table 11: The number of kernel invocations required to hide overhead of slow kernels execution. The goal is to find a kernel within 95% of the optimum with 90% probability and decrease tuning overhead under 10% of the runtime. Benchmarks on Radeon RX Vega56 marked with \* are running with smaller tuning space due to ROCm instability.

#### 6. Conclusion and Future Work

In this paper, we have introduced the Kernel Tuning Toolkit – an advanced autoning framework for OpenCL and CUDA applications. Using KTT allows expert programmers to configure applications for offline tuning and dynamic tuning based on arbitrary user-defined code optimization parameters. We have also developed a set of ten benchmarks covering important HPC application areas and demonstrated that autotuning with KTT allows to produce highly efficient implementations (often close to the theoretical peak of the hardware) of these benchmarks for different hardware architectures including CPUs, Xeon Phi co-processors and GPUs. In our experimental evaluations we also demonstrate that autotuning for different architectures is key for performance portability. Moreover, we have shown that rationally-designed tuning spaces are often small enough to be searched during application runtime, making dynamic tuning feasible for a subset of the considered benchmarks. We have demonstrated with two different applications that dynamic tuning outperforms offline tuned implementations quickly if some performancerelevant aspects, such as a the size of data structures, change. Moreover, we have shown that our framework can be integrated into production software, supporting multithreading, overlapping execution of host and device code with memory copies, and utilizing simultaneous kernel execution.

In future work, we would like to focus on the further development and integration of advanced search strategies. We believe that it is possible to accelerate dynamic tuning by gathering properties of the tuning space from previous tuning runs, e.g., determine the relative impact of tuning parameters on performance by analyzing of profiling data. Using more efficient search methods would make dynamic autotuning feasible also for larger tuning spaces with a small number of well-performing configurations. Another line of research will focus on advanced dynamic strategies that can detect when the performance of an application degrades and that can then automatically trigger dynamic re-tuning of the code.

Another planned improvement targets the generation

of tuning spaces. Currently, KTT first generates the whole tuning space and then prunes it based on the constraints given. We plan to speed-up tuning space generation similarly as it has been done in the ATF framework [36].

Furthermore, we plan to investigate the possibilities of connecting KTT with a compiler-based approach. By introducing a DSL for optimizations, the programmer would need to only implement advanced optimizations (such as changing memory layout or the algorithm), whereas simpler optimizations (such as vectorization or loop blocking) would be generated automatically by the compiler.

The vast amount of autotuning results, especially when coupled with profiling counters, can be used by the community to compare behavior and efficiency of different HW architectures, study effects of different code optimizations, or to develop new search strategies. Therefore, we plan to set up a public database containing tuning results with profiling counters and update this database with any new hardware or benchmark available.

Last but not least, KTT has been designed to be independent of the concrete API used for accelerated kernels (e.g., OpenCL or CUDA). It is, therefore, possible to add broader support for APIs, for example, Vulcan support would extend potential applications of KTT towards computer graphics. With non-blocking dynamic tuning, it would be possible to alter shaders at runtime without significant drop of frame rate.

#### Acknowledgements

The work was supported from European Regional Development Fund-Project "CERIT Scientific Cloud" (No. CZ.02.1.01-/0.0/0.0/16\_013/0001802). The project that gave rise to these results received the support of a fellowship from la Caixa Foundation (ID 100010434). The fellowship code is LCF/BQ/DI18-/11660021. This project has received funding from the European Unions Horizon 2020 research and innovation programme under the Marie Skodowska-Curie grant agreement No. 713673. The Spanish Ministry of Economy and Competitiveness through Grants BIO2016-76400-R(AEI/FEDER, UE). Comunidad Autnoma de Madrid through Grant: S2017/BMD-3817.

#### References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] V. Abrishami, J. R. Bilbao-Castro, J. Vargas, R. Marabini, J. M. Carazo, and C. O. S. Sorzano. A fast iterative convolution weighting approach for gridding-based direct fourier three-dimensional reconstruction with correction for the contrast transfer function. *Ultramicroscopy*, 157:79 – 87, 2015.
- [3] J. Ansel, S. Kamil, K. Veeramachaneni, J. Ragan-Kelley, J. Bosboom, U.-M. O'Reilly, and S. Amarasinghe. OpenTuner: An extensible framework for program autotuning. In *Proceedings* of the 23rd International Conference on Parallel Architectures and Compilation, PACT '14, pages 303–316, 2014.
- [4] E. Bajrovic and S. Benkner. Automatic performance tuning of pipeline patterns for heterogeneous parallel architectures. In 2014 International Conference on Parallel and Distributed Processing, Techniques and Applications, 2014.
- [5] E. Bajrovic, Mijakovic R., J. Dokulil, S. Benkner, and M. Gerndt. Tuning OpenCL applications with the periscope tuning framework. In 2016 49th Hawaii International Conference on System Sciences (HICSS), 2016.
- [6] P. Balaprakash, J. Dongarra, T. Gamblin, M. Hall, J. K. Hollingsworth, B. Norris, and R. Vuduc. Autotuning in highperformance computing applications. *Proceedings of the IEEE*, 106(11):2068–2083, Nov 2018.
- [7] S. Che, M. Boyer, J. Meng, D. Tarjan, Tarjan, J. W. Sheaffer, S. Lee, and K. Skadron. Rodinia: A benchmark suite for heterogeneous computing. In *IEEE International Symposium on Workload Characterization (IISWC)*, 2009.
- [8] C. Cummins, P. Petoumenos, Z. Wang, and H. Leather. Endto-end deep learning of optimization heuristics. In 2017 26th International Conference on Parallel Architectures and Compilation Techniques (PACT), pages 219–232, 2017.
- [9] A. Danalis, G. Marin, C. McCurdy, J. S. Meredith, P. C. Roth, K. Spafford, V. Tipparaju, and J. S. Vetter. The scalable heterogeneous computing (SHOC) benchmark suite. In *Proceedings of* the 3rd Workshop on General-Purpose Computation on Graphics Processing Units, GPGPU-3, pages 63–74. ACM, 2010.
- [10] J. J. Dongarra, J. Du Croz, S. Hammarling, and I. S. Duff. A set of level 3 basic linear algebra subprograms. ACM Trans. Math. Softw., 16(1):1–17, 1990.
- [11] I. S. Duff and J. K. Reid. The multifrontal solution of indefinite sparse symmetric linear. ACM Trans. Math. Softw., 9(3):302– 325, 1983.
- [12] J. Enmyren, U. Dastgeer, and C. W. Kessler. Towards a tunable multi-backend skeleton programming framework for multi-GPU systems. In MCC-3: Swedish Woekshop on Multicore Computing, 2010.
- [13] T. L. Falch and A. C. Elster. Machine learning based autotuning for enhanced OpenCL performance portability. In Proceedings of the 2015 IEEE International Parallel and Distributed Processing Symposium Workshop, 2015.
- [14] J. Filipovič, M. Madzin, J. Fousek, and L. Matyska. Optimizing CUDA code by kernel fusion: application on BLAS. *The Journal of Supercomputing*, 2015.
- [15] J. Filipovič, F. Petrovič, and S. Benkner. Autotuning of OpenCL kernels with global optimizations. In Proceedings of the 1st Workshop on AutotuniNg and aDaptivity AppRoaches for Energy Efficient HPC Systems (ANDARE '17), 2017.

- [16] M. Frigo and S. G. Johnson. The design and implementation of fttw3. *Proceedings of the IEEE*, 93(2):216–231, 2005.
- [17] D. Gadioli, R. Nobre, P. Pinto, E. Vitali, A. H. Ashouri, G. Palermo, J. Cardoso, and C. Silvano. SOCRATES a seamless online compiler and system runtime autotuning framework for energy-aware applications. In 2018 Design, Automation Test in Europe Conference Exhibition (DATE), 2018.
- [18] M. Gerndt, S. Benkner, E. César, C. Navarrete, E. Bajrovic, J. Dokulil, C. Guillén, R. Mijakovic, and A. Sikora. A multiaspect online tuning framework for HPC applications. *Software Quality Journal*, 2017.
- [19] M. Gerndt, E. Cesar, and S. Benkner. Automatic tuning of hpc applications - the periscope tuning framework (ptf). In Automatic Tuning of HPC Applications - The Periscope Tuning Framework (PTF). Shaker Verlag, 2015.
- [20] S. G. D. Gonzalo, S. D. Hammond, C. R. Trott, and W. M. a. Hwu. Revisiting online autotuning for sparse-matrix vector multiplication kernels on next-generation architectures. In 2017 IEEE 19th International Conference on High Performance Computing and Communications; IEEE 15th International Conference on Smart City; IEEE 3rd International Conference on Data Science and Systems (HPCC/SmartCity/DSS), 2017.
- [21] D. S. Goodsell, G. M. Morris, and A. J. Olson. Automated docking of flexible ligands: Applications of autodock. *Journal* of Molecular Recognition, 9(1):1–5, 1996.
- [22] S. Grauer-Gray and L. N. Pouchet. Polybench/gpu 1.0. http://web.cse.ohio-state.edu/~pouchet.2/software/ polybench/GPU/index.html, 2012.
- [23] S. Grauer-Gray, L. Xu, R. Searles, S. Ayalasomayajula, and J. Cavazos. Auto-tuning a high-level language targeted to gpu codes. In 2012 Innovative Parallel Computing (InPar), 2012.
- [24] D. Grewe and A. Lokhmotov. Automatically generating and tuning GPU code for sparse matrix-vector multiplication from a high-level representation. In Fourth Workshop on General Purpose Processing on Graphics Processing Units (GPGPU), 2011.
- [25] T. Kisuki, P. M. W. Knijnenburg, and M. F. P. O'Boyle. Combined selection of tile sizes and unroll factors using iterative compilation. In *Proceedings 2000 International Conference on Parallel Architectures and Compilation Techniques (PACT'00)*, 2000.
- [26] J. Kurzak, S. Tomov, and J. Dongarra. Autotuning GEMM kernels for the Fermi GPU. *IEEE Transactions on Parallel and Distributed Systems*, 23(11):2045–2057, 2012.
- [27] Y. Li, J. Dongarra, and S. Tomov. A note on auto-tuning GEMM for GPUs. In Proceedings of the 9th International Conference on Computational Science: Part I, 2009.
- [28] Y. Li, Y.-Q. Zhang, Y.-Q. Liu, G.-P. Long, and H.-P. Jia. MPFFT: An auto-tuning FFT library for OpenCL GPUs. Journal of Computer Science and Technology, 28(1):90–105, 2013.
- [29] K. Ljungkvist. Matrix-free finite-element operator application on graphics processing units. In Euro-Par 2014: Parallel Processing Workshops, pages 450–461. Springer International Publishing, 2014.
- [30] I. Masliah, A. Abdelfattah, A. Haidar, S. Tomov, M. Baboulin, J. Falcou, and J. Dongarra. High-performance matrix-matrix multiplications of very small matrices. In *Euro-Par 2016: Parallel Processing*, pages 659–671. Springer International Publishing, 2016.
- [31] K. Matsumotoi, N. Nakasato, and S. G. Sedukhin. Performance tuning of matrix multiplication in OpenCL on different GPUs and CPUs. In 2012 SC Companion: High Performance Computing, Networking Storage and Analysis, 2012.
- [32] R. Miceli, G. Civario, A. Sikora, E. César, M. Gerndt, H. Haitof, C. Navarrete, S. Benkner, M. Sandrieser, L. Morin, and F. Bodin. AutoTune: A Plugin-Driven Approach to the Automatic Tuning of Parallel Applications, pages 328–342. Springer, 2013.
- [33] S. Muralidharan, A. Roy, M. Hall, M. Garland, and P. Rai. Architecture-adaptive code variant tuning. SIGARCH Com-

puter Architecture News, 44(2):325–338, 2016.

- [34] S. Muralidharan, M. Shantharam, M. Hall, M. Garland, and B. Catanzaro. Nitro: A framework for adaptive code variant tuning. In *Proceedings of the 2014 IEEE 28th International Parallel and Distributed Processing Symposium*, IPDPS '14, pages 501–512. IEEE Computer Society, 2014.
- [35] C. Nugteren and V. Codreanu. CLTune: A generic auto-tuner for OpenCL kernels. In Proceedings of the IEEE 9th International Symposium on Embedded Multicore/Many-core Systemson-Chip (MCSoC), 2015.
- [36] A. Rasch and S. Gorlatch. ATF: A generic directive-based autotuning framework. Concurrency and Computation: Practice and Experience, 0(0):e4423, 2018.
- [37] A. Rasch, M. Haidl, and S. Gorlatch. ATF: A generic autotuning framework. In 2017 IEEE 19th International Conference on High Performance Computing and Communications; IEEE 15th International Conference on Smart City; IEEE 3rd International Conference on Data Science and Systems (HPC-C/SmartCity/DSS), 2017.
- [38] G. Rudy, M. M. Khan, M. Hall, C. Chen, and J. Chame. A programming language interface to describe transformations and code generation. In *Languages and Compilers for Parallel Computing.* Springer Berlin Heidelberg, 2011.
- [39] K. Seymour, H. You, and J. Dongarra. A comparison of search heuristics for empirical code optimization. In 2008 IEEE International Conference on Cluster Computing, 2008.
- [40] M. Steuwer, T. Remmelg, and C. Dubach. LIFT: A functional data-parallel ir for high-performance GPU code generation. In 2017 IEEE/ACM International Symposium on Code Generation and Optimization (CGO), pages 74–85, 2017.
- [41] J. E. Stone, J. C. Phillips, P. L. Freddolino, D. J. Hardy, L. G. Trabuco, and K. Schulten. Accelerating molecular modeling applications with graphics processors. *Journal of Computational Chemistry*, 28(16), 2007.
- [42] J. A. Stratton, C. Rodrigues, I. J. Sung, N. Obeid, L. W. Chang, N. Anssari, G. D. Liu, and W. M. Hwu. Parboil: A revised benchmark suite for scientificand commercial throughput computing. Technical report, University of Illinois at Urbana-Champaign, 2012.
- [43] D. Střelák, C. O. S. Sorzano, J. M. Carazo, and J. Filipovič. A GPU acceleration of 3D Fourier reconstruction in Cryo-EM. The International Journal of High Performance Computing Applications, 0, 2019.
- [44] A. Tiwari and J. K. Hollingsworth. Online adaptive code generation and tuning. In *IEEE International Parallel Distributed Processing Symposium (IPDPS)*, 2011.
- [45] V. Volkov and J. W. Demmel. Benchmarking GPUs to tune dense linear algebra. In ACM/IEEE conference on Supercomputing (SC), 2008.
- [46] Z. Wang, C. F. Hryc, B. Bammes, P. V. Afonine, J. Jakana, D.-H. Chen, X. Liu, M. L. Baker, C. Kao, S. J. Ludtke, M. F. Schmid, P. D. Adams, and W. Chiu. An atomic model of brome mosaic virus using direct electron detection and real-space optimization. *Nat Commun*, 5:4808, 2014.
- [47] B. van Werkhoven. Kernel tuner: A search-optimizing gpu code auto-tuner. Future Generation Computer Systems, 90:347 – 358, 2019.
- [48] R. C. Whaley and J. J. Dongarra. Automatically tuned linear algebra software. In *Proceedings of the 1998 ACM/IEEE Conference on Supercomputing*, 1998.