

DEUTSCHES ELEKTRONEN-SYNCHROTRON
Ein Forschungszentrum der Helmholtz-Gemeinschaft



DESY 21-058
arXiv:2105.01796
May 2021

A GPU-Based Kalman Filter for Track Fitting

X. Ai, N. Styles

Deutsches Elektronen-Synchrotron DESY, Hamburg

G. Mania

Deutsches Elektronen-Synchrotron DESY, Hamburg

and

Fachbereich Informatik, Universität Hamburg

H. M. Gray

University of California and Lawrence Berkeley National Laboratory, Berkeley, USA

M. Kuhn

Fakultät für Informatik, Otto von Guericke Universität Magdeburg

ISSN 0418-9833

NOTKESTRASSE 85 – 22607 HAMBURG

DESY behält sich alle Rechte für den Fall der Schutzrechtserteilung und für die wirtschaftliche Verwertung der in diesem Bericht enthaltenen Informationen vor.

DESY reserves all rights for commercial use of information included in this report, especially in case of filing application for or grant of patents.

To be sure that your reports and preprints are promptly included in the
HEP literature database
send them to (if possible by air mail):

DESY Zentralbibliothek Notkestraße 85 22607 Hamburg Germany	DESY Bibliothek Platanenallee 6 15738 Zeuthen Germany
---	---

A GPU-based Kalman Filter for Track Fitting

Xiaocong Ai  · Georgiana Mania  · Heather M. Gray  ·
Michael Kuhn  · Nicholas Styles 

Received: date / Accepted: date

Abstract Computing centres, including those used to process High-Energy Physics data and simulations, are increasingly providing significant fractions of their computing resources through hardware architectures other than x86 CPUs, with GPUs being a common alternative. GPUs can provide excellent computational performance at a good price point for tasks that can be suitably parallelized. Charged particle (track) reconstruction is a computationally expensive component of HEP data reconstruction, and thus needs to use available resources in an efficient way. In this paper, an implementation of Kalman filter-based track fitting using CUDA and running on GPUs is presented. This utilizes the ACTS (A Common Tracking Software) toolkit; an open source and experiment-independent toolkit for track reconstruction. The implementation details and parallelization approach are described, along with the specific challenges for such an implementation. Detailed performance benchmarking results are discussed, which show encouraging performance gains over a CPU-based implementation for representative configurations. Finally, a perspective on

the challenges and future directions for these studies is outlined. These include more complex and realistic scenarios which can be studied, and anticipated developments to software frameworks and standards which may open up possibilities for greater flexibility and improved performance.

Keywords particle tracking · track fitting · parallelization · GPU · CUDA · OpenMP

1 Introduction

The reconstruction of the trajectories of charged particles for High-Energy Physics (HEP) experiments is a very computationally demanding task. The most commonly used techniques are adaptive methods based on the Kalman filter [15,25], which account for the trajectories of charged particles in magnetic fields and the energy loss of charged particles in the detector material. See Ref. [40] for a review. As the execution time of such algorithms explodes combinatorially with the number of charged particles, the advent of the upgrade to the Large Hadron Collider (LHC), the High-Luminosity LHC (HL-LHC), portends an even greater challenge, with events containing up to 10,000 tracks.

For many years, HEP has been relying on Moore's Law [27], the observation that the number of transistors on an integrated circuit doubles approximately every two years. As the circuits have begun to approach intrinsic limits in terms of density and power, Moore's Law has begun to slow, further complicating potential performance improvements [39]. In addition, other computing architectures have become increasingly powerful and hence popular, such as graphical processing units (GPUs) and field programmable gate arrays (FPGAs). Therefore there has been a shift towards achieving speed

X. Ai
Deutsches Elektronen-Synchrotron, Hamburg, Germany
E-mail: xiaocong.ai@desy.de

G. Mania
Deutsches Elektronen-Synchrotron, Hamburg, Germany and
University of Hamburg, Informatics Department, Hamburg,
Germany

H.M. Gray
University of California, Berkeley CA, USA and Physics Division,
Lawrence Berkeley National Laboratory

M. Kuhn
Otto von Guericke University Magdeburg, Faculty of Computer Science,
Magdeburg, Germany

N. Styles
Deutsches Elektronen-Synchrotron, Hamburg, Germany

improvements by adding additional cores, particularly at high-performance computing centers. These multi-core systems require highly parallel code in order to be fully exploited, requiring additional knowledge from software developers. Moreover, much of the existing code for high-energy physics experiments is not well-suited to such architectures and hence requires significant development and adaptation to be able to exploit them.

We present a proof-of-concept of a Kalman filter algorithm on GPUs. Detailed studies of the physics and technical performance are presented for three different GPU architectures and compared to performance on CPUs. In particular, this is used to identify and discuss the key challenges in this implementation and to highlight future directions towards the development of an even more performant algorithm. Related studies, demonstrating the potential for trigger applications, have been performed by other groups [9, 38].

2 The Kalman Filter and ACTS

Track reconstruction is typically a multi-stage procedure, wherein candidates can be rejected at each stage. This approach allows high reconstruction efficiency and purity to be achieved in the final output collection, while reducing the overhead from processing unwanted candidates further than necessary. It starts from measurements (deposited energy in sensitive elements of the detector) and combines them in various configurations (including appropriate calibrations at various stages) to form plausible candidate trajectories. Accurate estimations of the parameters which define the mathematical form used to describe the trajectory are then made.

After any required pre-processing of the raw measurements, a typical first step is *Seeding*, in which small sets of compatible measurements are grouped using simple criteria and an initial trajectory estimate made. Seeds passing requirements can then be used as the basis for *Track Finding*, in which additional compatible measurements are added to the trajectory through the detector. Once the full set of measurements for the trajectory is obtained, a *Track Fitting* step can be performed, in order to precisely estimate the parameters and their covariance.

A commonly used approach and important tool in many track reconstruction applications is the Kalman filter [15, 25]. Developed in the late 1950s, the initial application of the Kalman filter procedure was in ballistics [32], where it allowed telemetry data for the heading and acceleration of the projectile to be combined with information on its location. The generalized procedure, in which measurements are combined with predictions based on an underlying model, results in state estimates

more precise than either measurements or predictions alone, and has since been very widely used in many fields.

Within track reconstruction, a typical Kalman filter *step* would proceed as follows (see Figure 1):

1. An initial estimate of the track state (i.e. helix parameters) at a given position is taken as the starting point
2. This track state is propagated according to the track model on to the next *Measurement Surface* (i.e. the reference plane of a sensitive detector), providing a prediction of track state on this surface
3. The prediction is combined with the measurement at this surface, if present, either through a weighted average or the so-called *Gain Matrix* formalism forming a new track state which is used to update the initial estimate
4. This new estimate is then used for further Kalman s.pdf, up to the end of the trajectory

The Kalman procedure has the property that the next state estimate can be determined from the one preceding it. While this is a useful property in many cases, as it requires no ‘history’ to be stored, it has the consequence that only the final state contains the full information about all the s.pdf preceding it, and therefore the best possible precision. In order to allow the prior track states to benefit from this information (e.g. to allow a χ^2 quality metric to be defined based on measurement residuals), an additional stage is needed. This *Smoother* stage can be performed using one of two approaches: either using essentially the same procedure as the forward Kalman filter but in reverse direction as illustrated in Figure 2 or using the Rauch-Tung-Striebel (RTS) smoother [37] formalism with the stored Jacobians between states calculated during the forward Kalman filter s.pdf. The latter approach does not require a second propagation of the track parameters and is therefore expected to have better timing performance.

A Common Tracking Software (ACTS) [30, 10, 28, 11] has its origins in the track reconstruction algorithms used by the ATLAS experiment [13]. ACTS provides a toolkit of algorithms for track reconstruction within a generic, framework- and experiment-independent software package. It also includes a fast simulation package. The ACTS code is designed to be inherently thread-safe to support parallel code execution and the data structures are vectorised. The implementation has been designed to be fully agnostic to detection technologies, detector design, and software frameworks so that it can be used by a range of experiments. The Eigen library [29] is used for algebra operations. In addition, ACTS is designed to be an R&D platform for the development of

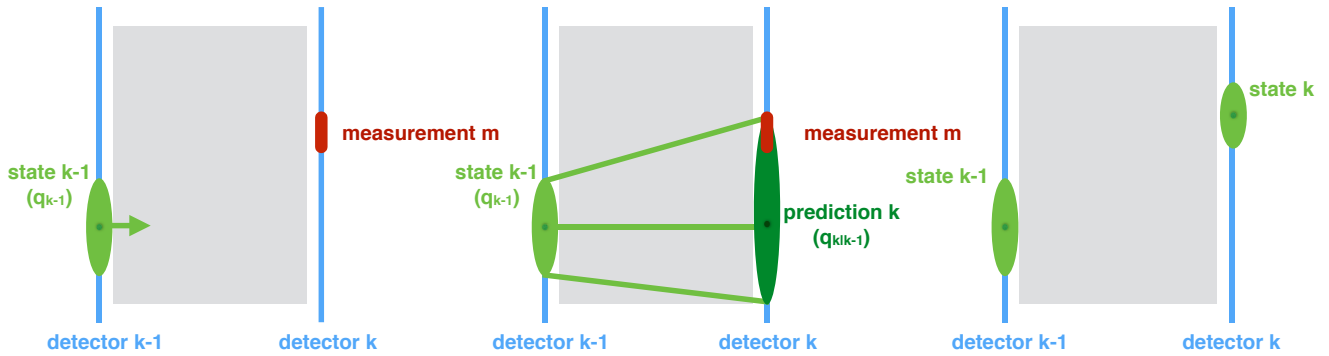


Fig. 1 Illustration of the different s.pdf of the Kalman filter using a simplified detector model consisting of only two layers. (Left) The track state at the $(k - 1)$ th surface is indicated with the light green ellipse, and a measurement on the k th surface is indicated in red. (Center) A prediction is made for the track state on the k th surface and the size of the prediction is indicated with the dark green ellipse. (Right) The track state on the k th surface is updated by including the measurement on the k th surface

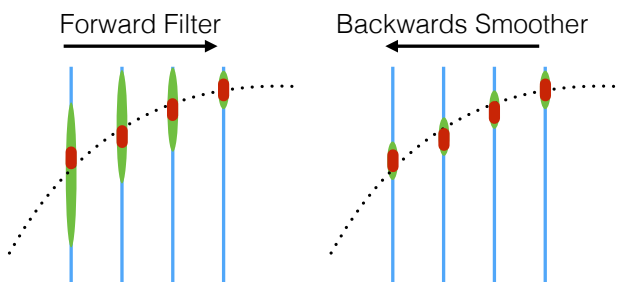


Fig. 2 Illustration of a forward filter and a backwards smoother on a simplified four layer detector geometry. The red points indicate the measurements and their uncertainties on each layer. The green points indicate the predictions. The predictions from the forward filter (left) are obtained when the filter is run from left to right. The predictions from the backwards filter are obtained during a second pass of the filter when it is run from right to left

new algorithms and the porting of existing algorithms to new hardware platforms. See Ref.[30,10,28,11] for further details.

While various representations of trajectories are possible, in this paper we will focus on helical trajectories of charged particles in a solenoidal magnetic field described using the following parameters:

- Two parameters loc_0 and loc_1 describing the spatial coordinates represented in the local frame of the measurement plane. In the special case of describing the parameters at a perigee surface¹, these become

¹ A surface defined at the point of closest approach to a reference point, for example the nominal interaction point in a particle collider

the transverse and longitudinal impact parameters d_0 and z_0 respectively.

- The polar and azimuthal angles of the particle momentum vector direction, ϕ and θ , at that point.
- A curvature parameter, expressed as the ratio of charge to momentum $\frac{q}{p}$.
- The time t .

Both the backwards-propagation and the RTS Kalman smoothing approaches are available within ACTS. The latter approach is used for the performance studies in this paper.

3 Parallelization and Offloading Techniques

There is a wide range of tools and frameworks available that can improve the runtime performance of scientific code via parallelization and offloading. Two of the most widely used frameworks are Open Multi-Processing (OpenMP) [19] and Compute Unified Device Architecture (CUDA) [6]. While the former traditionally allows parallelization on CPU systems via multiple threads, the latter is used to offload parts of the code to massively-parallel Nvidia GPUs.

3.1 OpenMP

OpenMP is a compiler-based high-level approach for thread parallelization on shared memory architectures. One of its outstanding features is that it is very easy to use and does not require knowledge of threading and operating system internals [18]. It is available for C, C++ and Fortran, some of the most-widely used programming languages for scientific computing. OpenMP

achieves its simplicity by being integrated into the compiler, which facilitates the parallelization of applications. Compiler support is widely available, which allows it to be used on personal computers as well as supercomputers. Applications are annotated with so-called *pragmas*, and turned into parallel code by the compiler; for instance, a simple `#pragma omp parallel for` instructs the compiler to parallelize the subsequent *for* loop. OpenMP takes care of thread management and scheduling as well as data decomposition, which allows developers to focus on the problem they want to solve.

One of OpenMP's drawbacks has been its focus on CPU-based parallelism. However, it has recently been extended with improved offloading functionality that allows the compiler to offload certain parts of an application to accelerators such as GPUs and FPGAs. Consequently, OpenMP now can target both CPUs and GPUs, which offers better portability than vendor-specific approaches such as CUDA [20].

3.2 CUDA

CUDA is a parallel computing platform and application programming interface introduced by Nvidia for their line of GPUs [36]. It allows highly-parallel GPUs to be used for general-purpose computations such as those common in high-energy physics. It is available for a range of programming languages, including C, C++, and Fortran. Wrappers are also available for additional programming languages, such as Python, R, Julia and many others.

GPUs are very specialized processing units and feature a high number of compute cores, which can be leveraged for scientific computations. Programs are offloaded to the GPUs in the form of so-called *compute kernels*, that is, single functions and their associated data. A kernel executes in parallel across a set of threads, which can use per-thread registers. Moreover, threads are aggregated into so-called warps that are executed concurrently. Several of these warps can be grouped into a thread block, which has access to a fast region of shared memory that all threads within the block can access. Finally, thread blocks can be combined into grids by the programmer. Thread blocks in a grid can only share data via global memory. Details of the CUDA programming model can be found in Ref. [8].

CUDA extends existing languages and requires dedicated compilers (`nvcc` for C/C++ and `nvfortran` for Fortran). While it allows for the optimal use of Nvidia GPUs, it is not portable and cannot be used for GPUs produced by other vendors. However, there have been attempts to provide abstraction layers or conversion tools for other approaches to be able to run CUDA

code via OpenCL [22, 14]. There are also a variety of libraries that automatically offload compute-intensive operations to the GPUs. Examples include libraries such as cuBLAS for linear algebra and cuFFT for fast-fourier transforms [24].

Competing approaches including OpenACC [31] and OpenCL are available but have not been not as widely adopted so far. For the parallel GPU implementation presented in this paper, we have chosen CUDA because it is the de-facto standard for GPU-accelerated code and is widely supported. Our attempts to develop a GPU-accelerated solution using OpenMP were not successful so far due to offloading support in OpenMP still being in an early stage of development.

4 GPU Implementation

In this section, implementation details of the Kalman filter used to run track fitting on both CPUs and GPUs are presented. The code can be found in Ref. [12].

4.1 Parallelization Strategy

As discussed in Section 2, track fitting is the step in the track reconstruction chain that precisely estimates the reconstructed track parameters and the associated covariance matrices. If track fitting is performed sequentially in a single event, the execution time will increase almost linearly with increasing track multiplicity. However, the dependence of the track fitting execution time on the number of tracks weakens if the track fitting can be parallelized. The implementation of a track-level parallel strategy is straightforward, since the track fitting for each reconstructed track is completely independent. In addition, the algorithm can be parallelized *within* a track fit, i.e. intra-track parallelization. Possible gains come from the matrix operations, e.g. the transportation of the track parameters in a magnetic field and the Kalman filter update and smoothing, which are computationally expensive and have to be repeated for all the propagation s.pdf and measurements by a total of up to $\mathcal{O}(10)$ times per track. However, in practise only very limited intra-track parallelization for those matrix operations can be achieved by using multiple threads, because the sizes of the matrices in one operation are usually relatively small. For example, the largest size matrix operated on in a single track fit in ACTS is the covariance matrix of the track parameters represented in the global coordinate system, which is of size 8×8 .

This paper discusses both parallelization strategies for track fitting on GPUs:

1. Track-level parallelization: Track fitting for different tracks is executed in parallel using different CUDA threads (or blocks if further intra-track parallelization is used).
2. Intra-track parallelization: The matrix operations involved in a single track fit are parallelized as much as possible using multiple threads within a single CUDA block. In this case, the block shared memory is used for the objects relevant with one track fit.

The transportation of the track parameters and their associated covariance matrices in a magnetic field requires a numerical solution to the equation of particle motion. The adaptive Runge-Kutta-Nyström [35] method is used to transport the track parameters in ACTS. When extrapolating the track parameters from one measurement point to the next, the covariance of the track parameters is updated with the transport Jacobian between the measurement points. Because the track parameters are represented in the local coordinate frame of the detector, the transform Jacobians between local and global track parameters at the two measurement points have to be applied. If the fitting is performed using one CUDA block per track, the matrix multiplication for the covariance transport can be parallelized using multiple threads.

4.2 CUDA Considerations and Limitations

Various CUDA programming requirements have consequences for the problem-specific factors that shape the parallelization strategies, and thus have an impact on the final implementation. The most challenging ones are detailed next.

4.2.1 Polymorphism

Virtual functions cannot be called inside a CUDA kernel unless the objects are constructed there. The Curiously Recurring Template Pattern (CRTP) is a C++ design pattern that emulates the behaviour of dynamic polymorphism through having a base class which is a template specialization for the derived class itself.

The ACTS Kalman filter is designed to be independent of the detector's tracking geometry, which could contain surfaces of different concrete types for different tracking detectors. To realize this design pattern on accelerators, the surfaces are implemented with CRTP, instantiated outside of the Kalman filter, and fed to the algorithm. CRTP is successfully used to define the surfaces as shown in the code sample in Listing 1.

```
template <typename Derived>
inline const typename Derived::SurfaceBoundsType
```

```
    *Surface ::bounds() const {
    return static_cast<const Derived *>(this)->bounds();
}
```

Listing 1 Function definition in Surface base class using CRTP

4.2.2 Thread Memory Limitations

The amount of memory available for a thread, which includes the stack frame size and the maximum number of registers, is automatically configured by the CUDA runtime environment based on device properties including the total amount of shared memory and the cache sizes, and also on the number of parallel threads per execution block. Because the GPU's performance gain is based on the ability to run thousands of threads in parallel, this limits considerably the amount of memory available per thread compared to the CPU. This memory limitation is a major concern for recursive functions, which have to be reimplemented using an iterative approach.

4.2.3 Limited Support for Linear Algebra Libraries

In the Kalman smoothing, the track parameter covariance matrix needs to be inverted to calculate the gain matrix. While the Eigen-based matrix class has a member function that returns the inverse matrix, this method is not supported in the GPU kernels. Also CUDA 10.0 discontinued the support for invoking cuBLAS functions from within the device kernels through `cublas_device` routine [2]. Since at the moment neither of these two linear algebra libraries provides a solution for our scenario, a customized matrix inverter is implemented whose performance impact is discussed in Section 5.

4.2.4 Precision and Rounding

Heterogeneous resources produce slightly different results due to different approaches to floating point arithmetic and rounding. CPUs typically promote float operands to doubles when possible, perform the operations in double precision and then truncate the result to simple precision. Moreover, the x86 floating point units use extended double precision registers (80-bit), while CUDA limits the register sizes to 32-bit and 64-bit as described by the IEEE standard 754-2008 [4, 7].

To mitigate these effects, the customized matrix inverter is always configured to perform the algebra operations using double precision. Floating point precision for other algebra operations is used for benchmarking the performance, and the comparison between the results with float and double precision is discussed in Section 5.3.

4.3 Data Structure and Transfer

Detector geometry information is a fundamental component for track parameter propagation and the integration of material effects during the track reconstruction. Because track reconstruction with a realistic detector description as used in full detector simulation requires significant computational resources, a simplified detector geometry, the so-called *tracking geometry*, is used during the track reconstruction in ACTS for fast navigation and extrapolation of tracks. The basic geometrical component of the tracking geometry in ACTS is the surface. The surface object carries information about its geometrical orientation, shape and boundary, the material approximated from a full detector geometry, and a unique hierarchical geometry identifier.

Magnetic fields are used for measuring the momentum of charged particles at high-energy physics experiments. When a charged particle passes through a magnetic field, it bends with the degree of bending inversely proportional to its momentum. In this paper, a constant magnetic field is used for the performance studies presented. It should be noted that this represents a significant simplification with respect to an inhomogeneous magnetic field, as found in many experiments, for which position-specific field information may need to be stored and retrieved, i.e. more memory might be required. The description of inhomogeneous or nonparametric magnetic fields however is also possible within ACTS.

The ACTS tracking Event Data Model (EDM) includes classes to describe tracking objects such as measurements and track parameters, which are represented with surfaces. The EDM for the track state is designed for the Kalman filter. It consists of a measurement, a predicted track parameter, a filtered track parameter and a smoothed track parameter, all located on a surface. Figure 3 shows a track state on a surface. During the construction of the detector geometry, a unique geometrical identifier is assigned to each detector surface. By storing the geometrical identifier in the tracking EDM, information about the geometry can be shared between the measurements and track parameters, and between the CPU and GPU.

The Kalman filter, detector geometry and magnetic field are global data shared by different tracks during the track reconstruction, and are therefore stored in the kernel global memory. In addition, four sets of track-specific data are required to execute the track fitting kernel on the GPU:

1. Input measurements of the trajectory
2. Starting track parameters to steer the track parameters propagation

3. Track fitting configurations, e.g. a target surface to extract the fitted track parameters
4. Fitted results including the fitting status, a collection of track states on the trajectory and the fitted track parameters on the target surface

Track-specific data are allocated on the pinned memory on the host, i.e. page-locked memory, and this memory is allocated contiguously for each track. The number of detector surfaces intersected by the particle varies with the kinematics of the particle and the detector layout, i.e. the data loads are different between tracks. Managing this load imbalance would require dedicated memory management and task scheduling strategies. In this paper, the detector surfaces are therefore constructed to be boundless² to guarantee that the same number of surfaces are traversed by the tracks. Each surface object requires approximately 120 bytes of memory. In addition to that, the size of memory allocated and transferred for different track-specific data for a detector with 10 plane surfaces is summarized in Table 1. Considering the relatively large size of the track-specific data, and the fact that they are accessed only once during the track fitting, those data are also stored in the kernel global memory.

5 Performance Evaluation

The software performance is studied using a simple telescope-like detector geometry with 10 planar surfaces

² Boundless surfaces always have an intersection with a track as long as the track is not parallel to the surface.

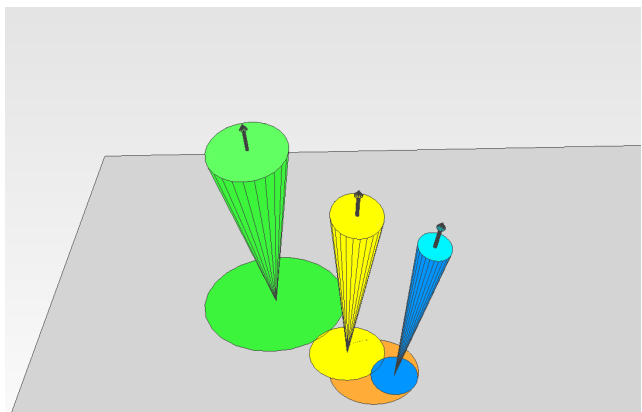


Fig. 3 Demonstration of a track state with a measurement (orange), a predicted track parameter (green), a filtered track parameter (yellow) and a smoothed track parameter (blue) on a plane surface (gray). The covariance matrix of the local coordinates for both the measurement and the fitted track parameters is represented by an ellipse. The momentum direction of the fitted track parameters is represented by an arrow with its covariance matrix represented by a cone oriented in the direction of the arrow

Data type	Size (B)
Input measurements	280
Seeding parameters	168
Fitting configurations	144
Fitting status	1
Fitted states	8,480
Fitted track parameters	216
Total	9,289

Table 1 The size (in bytes) of track-specific memory for a single track

perpendicular to the global x axis and placed equidistantly with 30 mm between two adjacent planes. Realistic HEP track reconstruction applications typically involve a more complicated detector geometry. A constant magnetic field of 2T along the global z axis is used. Samples containing a single muon per event are used for the performance evaluation. Muons are used for this study because they interact only minimally with the detector material and thus high quality track fits are expected. The muons have a transverse momentum uniformly distributed between 1 GeV and 10 GeV with both the azimuthal angle and polar angle fixed to zero.

The Fatras fast simulation engine [23] within the ACTS toolkit is used to generate simulated hits of the single muons on the detector surfaces. Figure 4 illustrates the simulated hits on the detector surfaces for a sample of 10,000 single muons. The dense tracking environment at the HL-LHC is not expected to exceed 10,000 tracks per event.

As the pattern recognition algorithms required to find track candidates from measurements are beyond the scope of this paper, the known trajectories of the simulated particles are used for the fitting in place of track candidates provided by a pattern recognition step. The measurements corresponding to the track candidates are obtained by smearing the positions of the simulated hits with Gaussian distributions to model detector resolution effects. A resolution of 50 μm in both the x and y dimensions of the detector, representative of the resolution of current pixel detectors at the LHC, is used. The initial set of track parameters for the track fit is based on the simulated particle vertex and momentum, smeared by Gaussian noise.

5.1 Hardware and Software Environment

Computing nodes of two supercomputers are used for running the performance tests:

1. Cori Intel Xeon Haswell node, Cori Intel Xeon Phi Knight’s Landing (KNL) node, and Cori GPU node at the National Energy Research Scientific Computing Center (NERSC) [1]

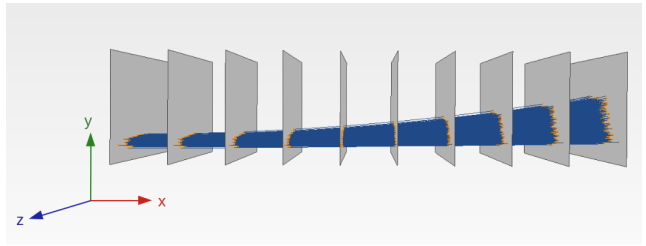


Fig. 4 The detector configuration used to study the performance. It consists of 10 identical planar surfaces (gray planes) perpendicular to the x -axis. The trajectories of 10,000 simulated single muons (blue lines) and the associated simulated hits on the detector surfaces (orange dots) are indicated

2. Intel Xeon Skylake (CL) node and ATLAS-GPU01 node at the National Analysis Facility (NAF) at DESY

Only GPUs from the Nvidia Tesla family are studied. All nodes use the CentOS 7 operating system and all GPUs are using CUDA version 10.2.89. Tables 2 and 3 show the detailed hardware and software configurations of the systems for CPUs and GPUs respectively.

Sys.	Model Name	S×C×T	Clock Rate (GHz)	Mem. (GB)
1	Intel Xeon E5-2698 v3 (Cori-Haswell)	2×16×2	2.30	128
1	Intel Xeon Phi 7250 (Cori-KNL)	1×68×4	1.40	96
2	Intel Xeon Gold 5115 (NAF-SL)	2×10×2	2.40	376

Table 2 CPU configurations. The Sys. column specifies whether the CPUs are used in the NERSC (1) or NAF (2) system. The S×C×T column represents the number of sockets (S), cores per socket (C) and threads per core (T)

These machines cover a wide range of architectures, with Cori-Haswell representing a standard compute node with two processors and a moderate amount of cores (for a total of 64 threads), while Cori-KNL features a higher core count due to the Xeon Phi’s particular architecture (for a total of 272 threads). Moreover, NAF-P100 and NAF-V100 allow two different GPU generations to be compared, with NAF-P100 being from the older Pascal architecture and NAF-V100 belonging to the newer Volta architecture. The NAF-P100 is connected through a Peripheral Component Interconnect Express (PCIe) serial connector while the NAF-V100 uses the SXM2 connector, a multi-line serial connector that provides both Nvidia NVLink and PCIe connectivity [5]. Cori-V100 is identical to NAF-V100 except for the lower

Sys.	GPU	FP32 Cores	FP64 Cores	Clock Rate (GHz)	Mem. (GB)
1	GV100-SXM2 (Cori-V100)	5,120	2,560	1.53	16
2	GP100-PCIe (NAF-P100)	3,584	1,792	1.48	16
2	GV100-SXM2 (NAF-V100)	5,120	2,560	1.53	32

Table 3 GPU configurations. The Sys. column specifies whether the GPUs are used in the NERSC (1) or NAF (2) system. FP32 and FP64 columns denote the numbers of floating point compute units for single and double precision arithmetic operations, respectively

amount of main memory. Each of the GPUs contains multiple Streaming Multiprocessors (SMs), which are similar to CPUs. However, each GPU is equipped with a large number of SMs, specifically, up to 60 SMs for the P100 and up to 84 SMs for the V100. Each SM contains many CUDA Cores, which execute compute kernels in the form of threads (64 single precision/32 double precision cores per SM for both P100 and V100). Each warp consists of 32 threads, with each warp running on one SM and each SM being able to execute up to 64 warps simultaneously. Due to the large number of total cores, it is important to distribute work across a sufficient number of warps. This allows the dedicated warp schedulers to achieve maximum utilization by keeping the cores busy with instructions. While on P100 all threads share a single program counter (and therefore have to execute the same instruction at the same time), V100 manages execution state per-thread, allowing more independence.

5.2 Tracking Performance

The resolution of the Kalman filter-based track fit is validated by calculating the pulls of the track parameters as follows:

$$v_{pull} = \frac{v_{fit} - v_{truth}}{\sigma_v}, \quad (1)$$

where v_{fit} and σ_v are the value and uncertainty of the fitted track parameter respectively, and v_{truth} is the corresponding parameter for the simulated particle.

The pull distributions of the six fitted perigee track parameters for a simulated sample of 10,000 single muons obtained from Cori-V100 are shown in Figure 5. The pull distributions have means compatible with zero and widths compatible with one, which demonstrates that the track parameters and their uncertainties are estimated correctly by the track fit.

5.3 Computing Performance

The timing performance of the track fitting for different number of tracks on various computing architectures and configurations is measured. Ten tests are run for each measurement. The mean time of the ten tests is taken as the measurement result, and square root of the average of their squared deviations from the mean is taken as the measurement error.

The baseline tests are performed using only track-level parallelization, i.e. without intra-track parallelization using shared memory. CUDA supports user configuration of the runtime properties of the GPU kernels, e.g. grid size and block size, and the launching of multiple kernels with multiple CUDA streams. Furthermore, the number of registers per thread can be controlled via the CUDA `__launch_bounds__()` qualifier when the kernel is defined. Unless explicitly specified, the tests are performed using one CUDA stream with 255 registers per thread as default configuration [5]. Moreover, we choose a grid size of $5,120 \times 1$ to match the number of processing units on the V100 GPU, and a block size of $8 \times 8 \times 1$, which is also the largest size of the matrix dealt with by the Kalman filter in ACTS. Performance with intra-track parallelization and different CUDA configurations are also studied for comparison.

All the tests use single precision arithmetic operations with the exception of (a) the matrix inversion algorithm required by the smoother as detailed in Section 4.2.3, and (b) the explicit scenario which compares the different precisions described in Section 5.3.4.

A singularity container with the executable and the required dependencies used to produce the results presented here is accessible in Ref. [12].

5.3.1 Performance of the Custom Matrix Inversion Algorithm

Because the Eigen-based matrix inversion algorithm used by ACTS cannot be called inside CUDA kernels, a custom algorithm for matrix inversion implemented for this purpose is used. Measurements are performed to compare our custom implementation to the Eigen-based implementation on the CPUs of both systems. As shown in Figure 6, it adds additional time to the fitting when the number of tracks exceeds 100. While the Eigen-based implementation is significantly faster by a constant factor when using only one thread, this effect is much less pronounced when using as many threads. Improving the performance of the custom matrix inversion operations on GPUs to match specialized linear algebra libraries would be expected to further improve the performance.

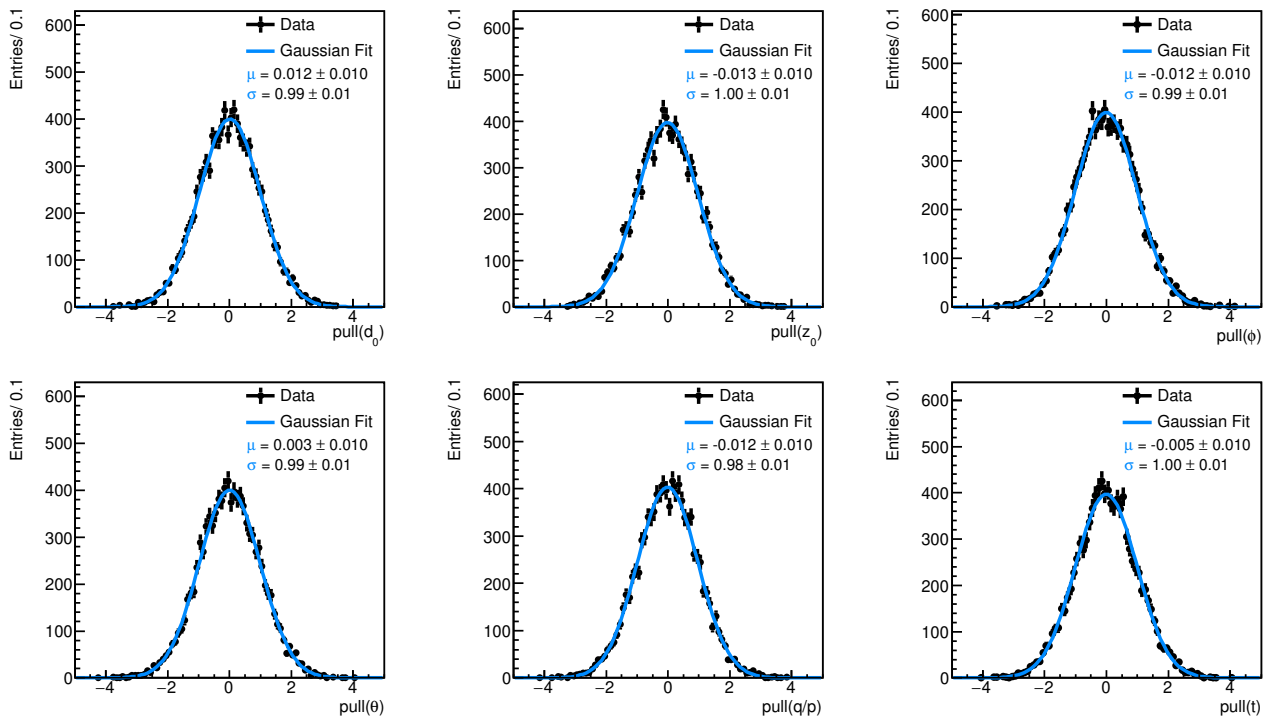


Fig. 5 The distributions of the pull values of the fitted perigee track parameters, $(d_0, z_0, \phi, \theta, \frac{q}{p}, t)$, for a sample of 10,000 single muons obtained from Cori-V100. The black dots are the pull values, and the blue lines are Gaussian fits to the distributions

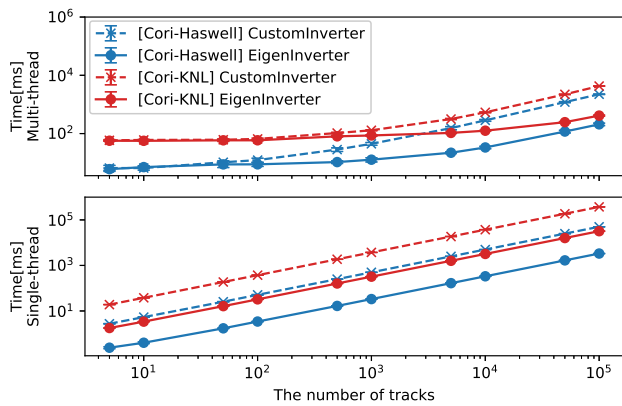


Fig. 6 The fitting time as a function of the number of tracks with different matrix inversion algorithms on Cori-Haswell (dashed blue for custom matrix inversion, and solid blue for Eigen-based matrix inversion) and Cori-KNL (dashed red for custom matrix inversion, and solid red for Eigen-based matrix inversion). The top panel shows the results with 60 and 250 threads on Cori-Haswell and Cori-KNL, respectively, and the bottom panel shows the results with a single thread

5.3.2 Performance of CUDA Code on CPU

The GPU-based track fitting program compiled with `nvcc` can also run on CPUs (although the CUDA driver and CUDA runtime need to be accessible in order to successfully allocate page-locked memory on the host).

In this case, the track-level parallelization is achieved by using OpenMP threads and the host-device model including memory and execution offloading is bypassed. The `nvcc` compiler uses the host compiler (`gcc` in this case) to generate the executable. This approach ensures comparable execution time when the number of tracks is small but it induces a small performance penalty (between 4% and 18%) compared with the standard C++ implementation compiled with OpenMP support when the number of tracks exceeds 1,000, as shown in Figure 7.

Nevertheless, it demonstrates the potential for single-source code targeting heterogeneous hardware resources. This is especially important for large and long-running software projects that might be used on different hardware architectures. However, this still requires the code to be written using CUDA and implies the portability limitations discussed in Section 3.2.

5.3.3 Performance Comparison Between GPU Architectures

Performance of the Kalman filter-based track fitting on the P100 and V100 Nvidia Tesla cards, is compared. Figure 8 shows that fitting time on NAF-P100 is more than a factor of two longer than on Cori-V100, therefore the following tests will focus on the V100 due to its

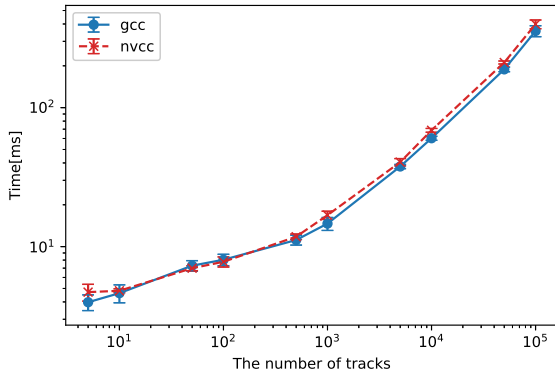


Fig. 7 The fitting time as a function of the number of tracks on NAF-SL using 60 threads with Eigen-based matrix inversion obtained with the `gcc` (solid blue) and `nvcc` (dashed red) compiled executables, respectively

superior performance. These performance differences are expected to a certain extent because the P100 and the V100 come from different hardware generations, with the V100 generally featuring more cores, higher clock rates and an improved interconnect.

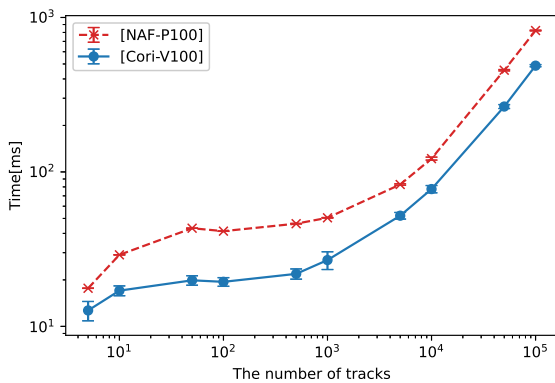


Fig. 8 The fitting time as a function of the number of tracks on NAF-P100 (dashed red) and Cori-V100 (solid blue)

5.3.4 Performance Comparison Between CPU and GPU

Figure 9 compares the fitting time on Cori-V100 with that on Cori-Haswell using different approaches to the matrix inversion. As mentioned previously, only the custom matrix inversion is possible with Cori-V100. When considering the custom matrix inverter, Cori-V100 displays superior performance compared to Cori-Haswell when the number of tracks exceeds 1,000. However, the Eigen-based implementation on CPUs still outperforms

our custom inverter on GPUs, demonstrating that it is important to not only consider the potential benefits from porting the actual code to GPUs but to also take supporting libraries into account.

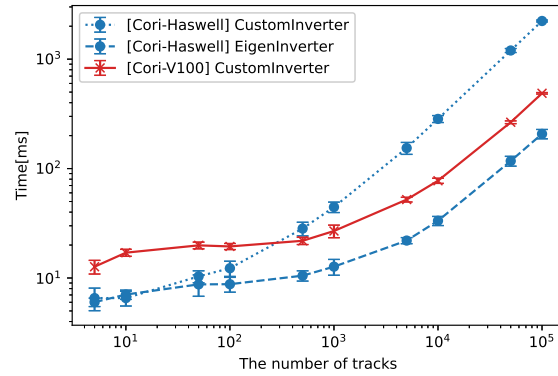


Fig. 9 The fitting time as a function of the number of tracks on Cori-Haswell using 60 threads with Eigen-based matrix inversion (dashed blue) and custom matrix inversion (dotted blue), and on Cori-V100 (solid red)

Our matrix inversion algorithm performs all the operations in double precision regardless of the operands' types. Figure 10 shows that there is little variation in performance between the different operands' types: when running the code on Cori-Haswell, there is virtually no performance difference for more than 1,000 tracks, while double operands are slightly slower on Cori-V100.

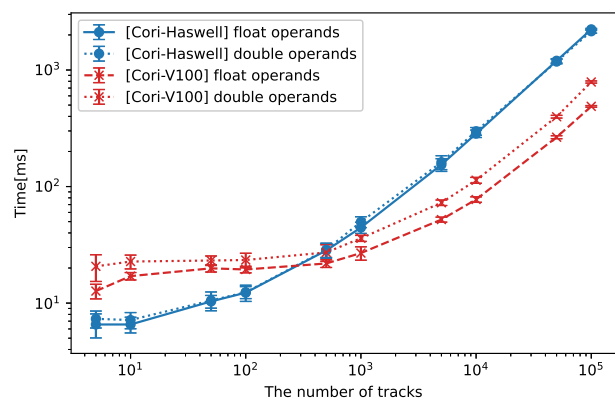


Fig. 10 The fitting time as a function of the number of tracks using `float` and `double` operands on Cori-Haswell (solid blue for float, and dotted blue for double) and Cori-V100 (dashed red for float, and dotted blue for double). The tests are performed using the custom matrix inversion and the ones executed on Cori-Haswell are using 60 OpenMP threads

5.3.5 Performance with Different GPU Configurations

The impact of usage of intra-track parallelization with shared memory, and variations in number of streams per device, grid size and block size as well as the number of fitted tracks are investigated and the most important results are discussed next.

Figure 11 shows that performance gain could be achieved by using intra-track parallelization with shared memory when the number of tracks is below 1,000. However, when the number of tracks exceeds 1,000, using intra-track parallelization results in a performance penalty due to limitation of the available shared memory and resident threads per SM.

Figure 12 shows that no significant performance gain is obtained from using more streams per device.

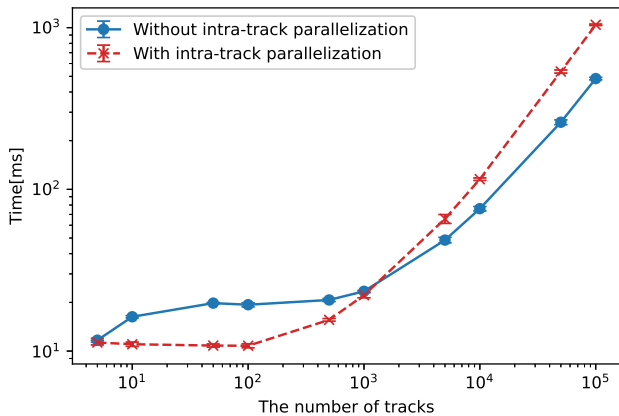


Fig. 11 The fitting time as a function of the number of tracks with linear grid of size $100,000 \times 1$ with (dashed red) or without (solid blue) intra-track parallelization on Cori-V100

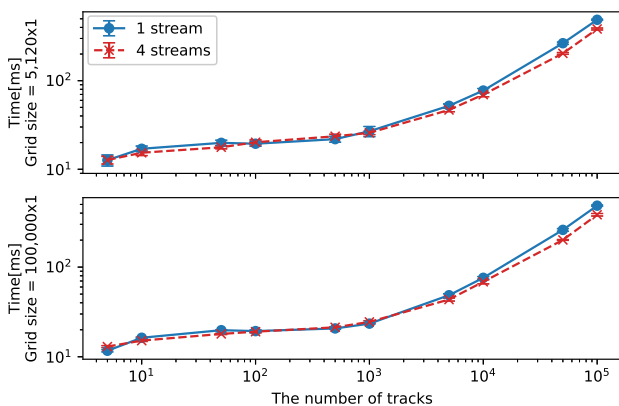


Fig. 12 The fitting time as a function of the number of tracks with linear grids of sizes $5,120 \times 1$ (top) and $100,000 \times 1$ (bottom), with one (solid blue) or four (dashed red) streams per device on Cori-V100

Figure 13 shows the required wall clock time with different grid and block sizes for 10,000 tracks. Note that when there are 1,024 threads per block, the CUDA `__launch_bounds__()` must be specified with the maximum number of threads per block no less than 1,024 to avoid "too many resources requested for launch" errors. This results in a maximum of 64 registers per thread at Cori-V100 with 65,536 32-bit registers per SM. While the performance differences between the two grid sizes are negligible, larger block sizes increase the runtime by up to a factor of 1.5. These results show that it is important to choose block sizes that are appropriate for the underlying hardware, i.e. when the overall workload is not large enough to saturate all the SMs on the GPU, a relatively large block size could lead to further imbalance of the workload distributed to the SMs and hence compromise the performance.

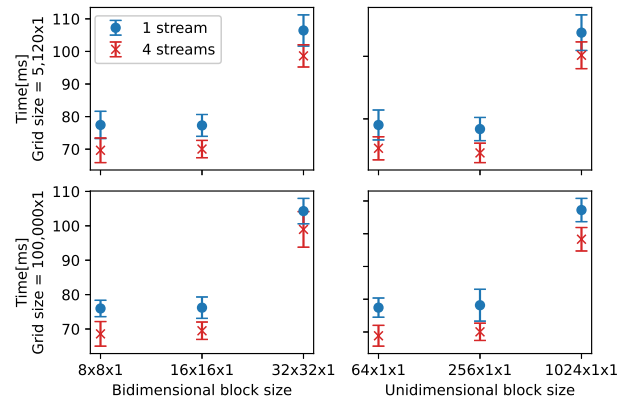


Fig. 13 The fitting time for 10,000 tracks as a function of bidimensional (left) or unidimensional (right) block sizes with linear grids of sizes $5,120 \times 1$ (top) and $100,000 \times 1$ (bottom), with one (blue) or four (red) streams per device on Cori-V100

The number of registers per thread can be reduced with the goal of running more threads per block without exceeding the hardware limits (i.e. the maximum number of blocks per SM). Figure 14 shows how a variation in the number of registers per thread affects the overall performance for a track fitting workload of 10,000 tracks. The performance varies little with the number of registers per thread. This is because the number of resident threads on the SMs are not increased by reducing the number of registers per thread when the overall workload is small. Increasing the block size can enforce at least the same amount of threads as the block size resident on some SMs, but this could lead to unwanted performance compromise due to inefficient GPU utilization. See Section 6 for further discussion.

The different components of the fitting time are analysed using Nsight Systems, one of Nvidia's new perfor-

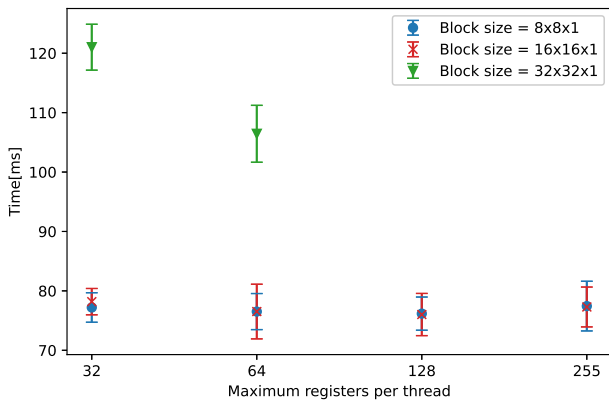


Fig. 14 The fitting time for 10,000 tracks using various number of registers per thread and block sizes on Cori-V100 node. The circle blue, cross red and triangle green represent block sizes of $8 \times 8 \times 1$, $16 \times 16 \times 1$ and $32 \times 32 \times 1$, respectively. When there are 1024 threads per block, a maximum of 64 registers per thread is allowed

mance analysis tools [33]. Figure 15 shows the timeline of those CUDA API activities required for GPU offloading, including the memory allocation on GPU, kernel launching, device synchronization and memory deallocation on the GPU, and the timeline of memory transfer and kernel execution in either one CUDA stream or multiple streams for 10,000 tracks. When using only a single stream, kernel execution accounts for roughly 70% of the total runtime, and memory transfer accounts for roughly 17% with significant impact on the performance. The performance gain from overlapping the data transfer and kernel execution with multiple streams is limited by the memory transfer time. A dedicated CUDA synchronization method might improve this.

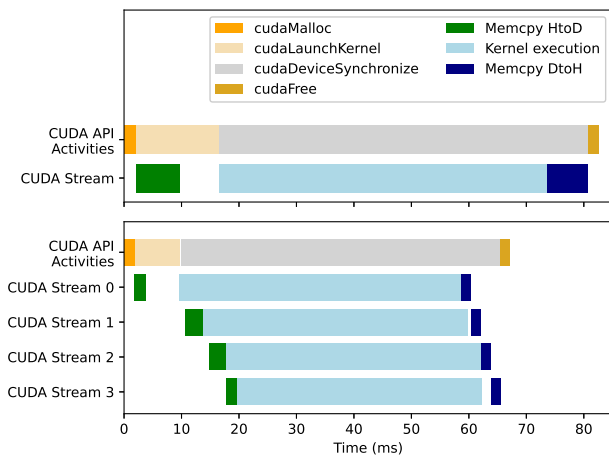


Fig. 15 The timeline of CUDA activities for offloading track fitting for 10,000 tracks to Cori-V100 with one stream (top) and four streams (bottom). The starting times of memory allocation on GPU are taken as the 0 point of the timelines

6 Discussion

The timing performance studies in Section 5 use a single GPU. Potential gains in the timing performance by utilizing multiple GPUs and the GPU occupancy (which could have an impact on the timing performance) are investigated in addition.

While the performance of the GPU-based Kalman filter was tested on one GPU, the implementation can also run on multiple GPUs in parallel. The prototype uses a team of threads on the host, each one fitting the trajectory of a subset of tracks on a different GPU. Despite better problem size scaling, the multi-device solution has a slightly larger overall execution time for smaller numbers of tracks, as shown in Figure 16. Communication via Message Passing Interface (MPI) is required to fully exploit the parallelism by resolving synchronization overhead between the GPUs.

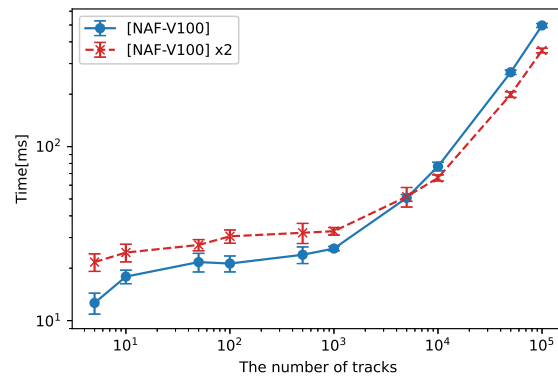


Fig. 16 The fitting time as a function of the number of tracks executed in one stream per device with linear grid of size 5120×1 and block size of $8 \times 8 \times 1$, when using one NAF-V100 (solid blue) and two NAF-V100 in parallel (dashed red)

In addition, the latest versions of the Nvidia HPC Software Development Kit (SDK) provide new tools and libraries designed to maximize performance by optimizing memory transfers and scaling to multiple devices while targeting heterogeneous resources [3]. Additionally, various studies regarding vendor-agnostic offloading approaches show promising results based on standard APIs and/or open-source, non-proprietary solutions [21, 26]. These would be very interesting to explore in future iterations.

The warp occupancy, defined as the ratio of active warps on an SM to the maximum number of active warps supported by the SM (e.g. 64 warps per SM on the V100), is analyzed using a different Nvidia performance tool: Nsight Compute [33]. Figure 17 shows the dependency of the theoretical warp occupancy and the achieved warp

occupancy on the number of registers per thread and block size for a track fitting workload of 10,000 tracks. The theoretical occupancy is 100% when the number of registers per thread is no larger than 32 and decreases when more registers are required for one thread, hence fewer threads are active. Since the maximum number of thread blocks per SM is 32 on the V100, the block size is not a limiting factor to theoretical occupancy as long as it is no less than 64.

The achieved occupancy is well below the theoretical one, in particular when the block size is small. The reason is that an average of only 119 tracks are distributed to each SM of Cori-V100, providing a total workload of 10,000 tracks. Therefore, only 2 blocks are resident on the SM when the block size is $8 \times 8 \times 1$ (128 threads), and at most 1 resident block when the block size is $16 \times 16 \times 1$ (256 threads) or $32 \times 32 \times 1$ (1024 threads). These correspond to an occupancy of 6.25%, 12.5%, and 50%, respectively. In this case, reducing the number of registers per thread has no impact on the warp occupancy. Event-level parallelization is an effective approach to increase the workload, and hence improve the SM warp occupancy. This can be achieved by an offloading pattern with a fully contained chain of tracking modules that run on GPU requiring minimum data transfer between CPU and GPU [16]. The workload also needs to be accounted for when analyzing the impact of SM warp occupancy on the performance. For instance, better warp occupancy does not necessarily correspond to better timing performance, as shown in Figure 14, for the particular track fitting workload of 10,000 tracks studied here. Further discrepancy between the achieved occupancy and the theoretical one arises from the imbalanced track fitting workload both within blocks and across blocks due to different momentum hence propagation paths between tracks. In this paper, the workload imbalances were already controlled by using a homogeneous detector geometry for all the tracks. For a realistic detector, the concept of tracking regions as presented in Ref. [34,17] can be used for the parallelization of track reconstruction. In the case of track finding, there is additional workload imbalance from the selection of compatible measurements from a pool of non-static number of measurements on a detector surface and possibly splitting the track propagation into multiple branches if there are more than one compatible measurement found. One possible approach to suppress the workload imbalance level would be to group the tracks based on their kinematic properties so that one group of tracks will encounter the same segmented detector region, and assign different groups of tracks to different grids or blocks. See Ref. [34,17] for further discussion.

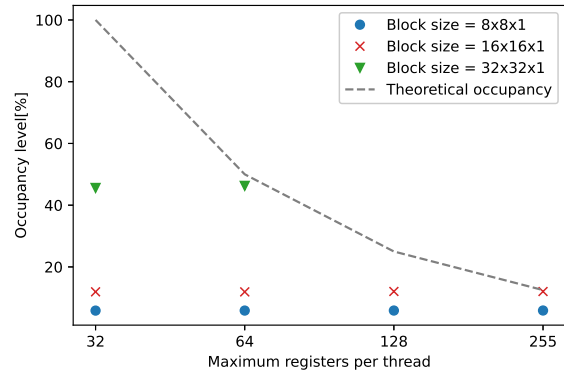


Fig. 17 Warp occupancy levels using various number of registers per thread and block sizes for fitting 10,000 tracks executed in one stream with linear grid of size $5,120 \times 1$ on Cori-V100. The black dashed line is the theoretical warp occupancy, and the circle blue, cross red and triangle green represent the achieved warp occupancy with block sizes of $8 \times 8 \times 1$, $16 \times 16 \times 1$ and $32 \times 32 \times 1$, respectively. When there are 1024 threads per block, a maximum of 64 registers per thread is allowed

7 Conclusion

The reconstruction of charged particle trajectories for current and future high-energy physics experiments is a significant computational challenge. New approaches are needed to cope with the dramatically increased event complexity and rates and with the movement away from x86 architectures. The Kalman filter algorithm is the mainstay of current track reconstruction strategies. We introduced a proof-of-concept implementation of a Kalman filter track fitting algorithm on two different GPU architectures. We have performed studies of its physics and technical performance and compared this to results using CPUs with a particular focus on the limitations observed. Ideas for improvements for future implementations were discussed.

As existing fast matrix inversion algorithms cannot run on GPUs, we developed a custom prototype matrix inversion algorithm, which does not match the CPU performance of highly-optimized state-of-the-art algorithms. When controlling for the matrix inversion algorithm, worse performance for low track multiplicity is obtained with the GPUs compared to the CPUs and the performance is improved by a factor of up to 4.6 with respect to the CPUs for events with more than 1,000 tracks. Significant performance differences are shown between the different GPU architectures.

Parallelization within the track fit was implemented and performance gain was observed with a relatively low track multiplicity. The performance dependence on GPU configurations was also studied. The performance

was largely independent of the grid size and did not change when using multiple kernels. Memory transfer and other overhead can account for up to 30% of the total run time for the track fit.

The typical HL-LHC track multiplicity of 10,000 tracks per event is a relatively small workload for GPUs. For events with 10,000 tracks, performance gain by up to a factor of 1.5 is achieved by using a smaller block size. The small workload is also the main limiting factor in the achieved occupancy on the GPU.

We have demonstrated that the Kalman filter-based track fitting can be successfully implemented and run on GPUs. We have compared different methods for implementation and studied the dependence on the GPU configuration. We have identified limitations of the approach and highlighted areas for future work directions.

Specifically, an evaluation of alternative approaches for GPU offloading, especially those provided by vendor-agnostic interfaces such as OpenMP, can be expected to result in improved performance portability. Moreover, further improvements to the GPU-based matrix inversion algorithm can be expected to bring its performance closer to existing CPU implementations.

Acknowledgements We would like to thank Dr. Attila Krasznahorkay (CERN), Dr. Charles Leggett (Lawrence Berkeley National Laboratory) and Dr. Andreas Salzburger (CERN) for their careful reading of the manuscript and their helpful comments and suggestions. This work used the grid computational resources operated at Deutsches Elektronen-Synchrotron (DESY), Hamburg, Germany and at the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231. We acknowledge the support by the National Science Foundation (NSF) and Data Science in Hamburg - HELMHOLTZ Graduate School for the Structure of Matter (DASHH).

Declarations

Funding This work was funded by the NSF under Cooperative Agreement OAC-1836650, and supported by DASHH with the Grant-No. HIDSS-0002.

Conflict of interest The authors declare that they have no conflict of interest.

Availability of data and material Not applicable. No associated data except for code.

Code availability The code used for this research (including a Singularity container for reproducibility) is available open source [12].

References

1. NERSC Cori System Specification. URL <https://docs.nersc.gov/systems/cori/#system-specification>. Online; accessed 4-February-2021
2. NVIDIA CUDA Toolkit v10.0.130 Release Notes. URL <https://docs.nvidia.com/cuda/archive/10.0/cuda-toolkit-release-notes/index.html#deprecated-features>. Online; accessed 4-February-2021
3. NVIDIA HPC Software Development Kit. URL <https://developer.nvidia.com/hpc-sdk>. Online; accessed 4-February-2021
4. IEEE 754-2008 - IEEE Standard for Floating-Point Arithmetic (2008). URL <https://standards.ieee.org/standard/754-2008.html>
5. NVIDIA Tesla V100 GPU Architecture (2017). URL <https://images.nvidia.com/content/volta-architecture/pdf/volta-architecture-whitepaper.pdf>. Online; accessed 4-February-2021
6. CUDA Toolkit Documentation (2019). URL <https://docs.nvidia.com/cuda/index.html>
7. CUDA Toolkit Documentation - Floating Point and IEEE 754 (2019). URL <https://docs.nvidia.com/cuda/floating-point/index.html>
8. CUDA C++ Programming Guide (2021). URL <https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html>
9. Aaij, R., Albrecht, J., Belous, M., Billoir, P., Boettcher, T., Brea Rodríguez, A., vom Bruch, D., Cámpora Pérez, D.H., Casais Vidal, A., Craik, D.C., Fernandez Declara, P., Funke, L., Gligorov, V.V., Jashal, B., Kazeev, N., Martínez Santos, D., Pisani, F., Pliushchenko, D., Popov, S., Quagliani, R., Rangel, M., Reiss, F., Sánchez Mayordomo, C., Schwemmer, R., Sokoloff, M., Stevens, H., Ustyuzhanin, A., Vilasis Cardona, X., Williams, M.: Allen: A high-level trigger on gpus for lhcb. *Computing and Software for Big Science* 4(1), 7 (2020). DOI 10.1007/s41781-020-00039-7. URL <https://doi.org/10.1007/s41781-020-00039-7>
10. Ai, X.: Acts: A common tracking software. In: Meeting of the Division of Particles and Fields of the American Physical Society (2019)
11. Ai, X.: Tracking with A Common Tracking Software (2020)
12. Ai, X., Mania, G., Gray, H.M., Kuhn, M., Styles, N.: gpukalmanfitter: v2.0 (2021). DOI 10.5281/zenodo.4693389. URL <https://doi.org/10.5281/zenodo.4693389>
13. ATLAS Collaboration: The ATLAS Experiment at the CERN Large Hadron Collider. *JINST* 3, S08003. 437 p (2008). DOI 10.1088/1748-0221/3/08/S08003. URL <https://cds.cern.ch/record/1129811>. Also published by CERN Geneva in 2010
14. Babej, M., Jääskeläinen, P.: HIPCL: tool for porting CUDA applications to advanced opencl platforms through HIP. In: S. McIntosh-Smith (ed.) IWOCCL '20: International Workshop on OpenCL, Munich, Germany, April 27-29, 2020, pp. 18:1–18:3. ACM (2020). DOI 10.1145/3388333.3388641. URL <https://doi.org/10.1145/3388333.3388641>
15. Billoir, P.: Track Fitting With Multiple Scattering: A New Method. *Nucl. Instrum. Meth.* A225, 352–366 (1984). DOI 10.1016/0167-5087(84)90274-6
16. Bocci, A., Kortelainen, M., Innocente, V., Pantaleo, F., Rovere, M.: Heterogeneous reconstruction of tracks and primary vertices with the cms pixel tracker (2020)

17. Cerati, Giuseppe, Elmer, Peter, Gravelle, Brian, Kortelainen, Matti, Krutelyov, Vyacheslav, Lantz, Steven, Masciovecchio, Mario, McDermott, Kevin, Norris, Boyana, Reid, Michael, Reinsvold Hall, Allison, Riley, Daniel, Tadel, Matevz, Wittich, Peter, Wang, Bei, Würthwein, Frank, Yagil, Avraham: Reconstruction of charged particle tracks in realistic detector geometry using a vectorized and parallelized kalman filter algorithm. *EPJ Web Conf.* **245**, 02013 (2020). DOI 10.1051/epjconf/202024502013. URL <https://doi.org/10.1051/epjconf/202024502013>
18. Clark, D.: Openmp: a parallel standard for the masses. *IEEE Concurr.* **6**(1), 10–12 (1998). DOI 10.1109/4434.656771. URL <https://doi.org/10.1109/4434.656771>
19. Dagum, L., Menon, R.: OpenMP: an industry standard API for shared-memory programming. *IEEE Computational Science and Engineering* **5**(1), 46–55 (1998). DOI 10.1109/99.660313
20. Daley, C.S., Ahmed, H., Williams, S., Wright, N.J.: A case study of porting HPGMG from CUDA to openmp target offload. In: K. Milfeld, B.R. de Supinski, L. Koesterke, J. Klinkenberg (eds.) *OpenMP: Portable Multi-Level Parallelism on Modern Systems - 16th International Workshop on OpenMP, IWOMP 2020, Austin, TX, USA, September 22-24, 2020, Proceedings, Lecture Notes in Computer Science*, vol. 12295, pp. 37–51. Springer (2020). DOI 10.1007/978-3-030-58144-2_3. URL https://doi.org/10.1007/978-3-030-58144-2_3
21. Deakin, T., Poenaru, A., Lin, T., McIntosh-Smith, S.: Tracking performance portability on the yellow brick road to exascale. In: 2020 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC), pp. 1–13 (2020). DOI 10.1109/P3HPC51967.2020.00006
22. Du, P., Weber, R., Luszczek, P., Tomov, S., Peterson, G.D., Dongarra, J.J.: From CUDA to opencl: Towards a performance-portable solution for multi-platform GPU programming. *Parallel Comput.* **38**(8), 391–407 (2012). DOI 10.1016/j.parco.2011.10.002. URL <https://doi.org/10.1016/j.parco.2011.10.002>
23. Edmonds, K., Fleischmann, S., Lenz, T., Magass, C., Mechnich, J., Salzburger, A.: The Fast ATLAS Track Simulation (FATRAS). *Tech. Rep. ATL-SOFT-PUB-2008-001. ATL-COM-SOFT-2008-002*, CERN, Geneva (2008). URL <https://cds.cern.ch/record/1091969>
24. Fatica, M.: CUDA toolkit and libraries. In: 2008 IEEE Hot Chips 20 Symposium (HCS), pp. 1–22 (2008). DOI 10.1109/HOTCHIPS.2008.7476520
25. Fruhwirth, R.: Application of Kalman filtering to track and vertex fitting. *Nucl. Instrum. Meth.* **A262**, 444–450 (1987). DOI 10.1016/0168-9002(87)90887-4
26. Gayatri, R., Yang, C., Kurth, T., Deslippe, J.: A case study for performance portability using openmp 4.5. In: S. Chandrasekaran, G. Juckeland, S. Wienke (eds.) *Accelerator Programming Using Directives - 5th International Workshop, WACCPD 2018, Dallas, TX, USA, November 11-17, 2018, Proceedings, Lecture Notes in Computer Science*, vol. 11381, pp. 75–95. Springer (2018). DOI 10.1007/978-3-030-12274-4_4. URL https://doi.org/10.1007/978-3-030-12274-4_4
27. G.E. Moore: Cramming More Components onto Integrated Circuits. *Electronics* **38** **8** (1965). DOI 10.1109/JPROC.1998.658762
28. Gessinger, P., Grasland, H., Gray, H., Kiehn, M., Klimpel, F., Langenberg, R., Salzburger, A., Schlag, B., Zhang, J., Ai, X.: The Acts project: track reconstruction software for HL-LHC and beyond. *EPJ Web Conf.* **245**, 10003 (2020). DOI 10.1051/epjconf/202024510003
29. Guennebaud, G., Jacob, B., et al.: *Eigen v3*. <http://eigen.tuxfamily.org> (2010)
30. Gumpert, C., Salzburger, A., Kiehn, M., Hrdinka, J., Calace, N.: ACTS: from ATLAS software towards a common track reconstruction software. *Tech. Rep. ATL-SOFT-PROC-2017-030. 4*, CERN, Geneva (2017). DOI 10.1088/1742-6596/898/4/042011. URL <https://cds.cern.ch/record/2243297>
31. Herdman, J.A., Gaudin, W.P., Perks, O., Beckingsale, D.A., Mallinson, A.C., Jarvis, S.A.: Achieving portability and performance through openacc. In: S. Chandrasekaran, F.S. Foertter, O.R. Hernandez (eds.) *Proceedings of the First Workshop on Accelerator Programming using Directives, WACCPD '14, New Orleans, Louisiana, USA, November 16-21, 2014*, pp. 19–26. IEEE Computer Society (2014). DOI 10.1109/WACCPD.2014.10. URL <https://doi.org/10.1109/WACCPD.2014.10>
32. Kalman, R.E.: A new approach to linear filtering and prediction problems. *ASME Journal of Basic Engineering* (1960)
33. Knobloch, M., Mohr, B.: Tools for GPU computing - debugging and performance analysis of heterogenous HPC applications. *Supercomput. Front. Innov.* **7**(1), 91–111 (2020). DOI 10.14529/jsfi200105. URL <https://doi.org/10.14529/jsfi200105>
34. Lantz, S., McDermott, K., Reid, M., Riley, D., Wittich, P., Berkman, S., Cerati, G., Kortelainen, M., Hall, A.R., Elmer, P., Wang, B., Giannini, L., Krutelyov, V., Masciovecchio, M., Tadel, M., Würthwein, F., Yagil, A., Gravelle, B., Norris, B.: Speeding up particle track reconstruction using a parallel kalman filter algorithm. *Journal of Instrumentation* **15**(09), P09030–P09030 (2020). DOI 10.1088/1748-0221/15/09/p09030. URL <https://doi.org/10.1088/1748-0221/15/09/p09030>
35. Myrheim, J., Bugge, L.: A fast Runge-Kutta method for fitting tracks in a magnetic field. *Nucl. Instrum. Meth.* **160**(1), 43–48 (1979). DOI 10.1016/0029-554X(79)90163-0
36. Nickolls, J., Buck, I., Garland, M., Skadron, K.: Scalable parallel programming with CUDA. *ACM Queue* **6**(2), 40–53 (2008). DOI 10.1145/1365490.1365500. URL <https://doi.org/10.1145/1365490.1365500>
37. Rauch, H.E., Tung, F., Striebel, C.T.: Maximum likelihood estimates of linear dynamic systems. *AIAA Journal* **3**(8), 1445–1450 (1965). DOI 10.2514/3.3166
38. Rohr, David, Gorbunov, Sergey, Ole Marten, Schmidt, Shahoyan, Ruben: Gpu-based online track reconstruction for the alice tpc in run 3 with continuous read-out. *EPJ Web Conf.* **214**, 01050 (2019). DOI 10.1051/epjconf/201921401050. URL <https://doi.org/10.1051/epjconf/201921401050>
39. Shalf, J.: The future of computing beyond moore’s law. *Philosophical Transactions of the Royal Society A* **378** (2020)
40. Strandlie, A., Frühwirth, R.: Track and vertex reconstruction: From classical to adaptive methods. *Rev. Mod. Phys.* **82**, 1419–1458 (2010). DOI 10.1103/RevModPhys.82.1419. URL <https://link.aps.org/doi/10.1103/RevModPhys.82.1419>