

PAPER • OPEN ACCESS

Verification of Electromagnetic Physics Models for Parallel Computing Architectures in the GeantV Project

To cite this article: G Amadio *et al* 2017 *J. Phys.: Conf. Ser.* **898** 042019

View the [article online](#) for updates and enhancements.

Related content

- [The GeantV project: preparing the future of simulation](#)
G Amadio, J Apostolakis, M Bandieramonte *et al.*
- [Electromagnetic Physics Models for Parallel Computing Architectures](#)
G Amadio, A Ananya, J Apostolakis *et al.*
- [First experience of vectorizing electromagnetic physics models for detector simulation](#)
G Amadio, J Apostolakis, M Bandieramonte *et al.*



IOP | ebooks™

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

Start exploring the collection - download the first chapter of every title for free.

Verification of Electromagnetic Physics Models for Parallel Computing Architectures in the GeantV Project

G Amadio¹, J Apostolakis², M Bandieramonte^{*2}, S P Behera³, R Brun², P Canal⁴, F Carminati², G Cosmo², L Duhem⁵, D Elvira⁴, G Folger², A Gheata^{2,5}, M Gheata^{2,6}, I Goulas², F Hariri², S Y Jun⁴, D Konstantinov², H Kumawat³, V Ivantchenko², G Lima⁴, T Nikitina², M Novak², W Pokorski², A Ribon², R Seghal³, O Shadura², S Vallecorsa^{2,5} and S Wenzel²

¹ Parallel Computing Center at Sao Paulo State University (UNESP), Sao Paulo, Brazil

² CERN, Route de Meyrin, Meyrin, Switzerland

³ Bhabha Atomic Research Centre (BARC), Mumbai, India

⁴ Fermilab, MS234, P.O. Box 500, Batavia, IL, 60510, USA

⁵ Intel Corporation, Santa Clara, CA, 95052, USA

⁶ Institute of Space Sciences, Bucharest-Magurele, Romania

E-mail: *marilena.bandieramonte@cern.ch

Abstract. An intensive R&D and programming effort is required to accomplish new challenges posed by future experimental high-energy particle physics (HEP) programs. The GeantV project aims to narrow the gap between the performance of the existing HEP detector simulation software and the ideal performance achievable, exploiting latest advances in computing technology. The project has developed a particle detector simulation prototype capable of transporting in parallel particles in complex geometries exploiting instruction level micro-parallelism (SIMD and SIMT), task-level parallelism (multithreading) and high-level parallelism (MPI), leveraging both the multi-core and the many-core opportunities. We present preliminary verification results concerning the electromagnetic (EM) physics models developed for parallel computing architectures within the GeantV project. In order to exploit the potential of vectorization and accelerators and to make the physics model effectively parallelizable, advanced sampling techniques have been implemented and tested. In this paper we introduce a set of automated statistical tests in order to verify the vectorized models by checking their consistency with the corresponding Geant4 models and to validate them against experimental data.

1. Introduction

High-energy particle physics has advanced greatly over recent years and current plans for the future foresee even more ambitious targets and challenges that have to be coped with. The areas involved are diverse: from detector simulation to high-speed data acquisition and storage, from data analytics and machine learning applied to data analysis to network security issues. Amongst all these computing R&D areas, simulation of particle detectors often stands out as the most time consuming part of HEP computing. The GeantV project [1, 2] intends to develop the future generation simulation software aiming at producing detector simulation software



capable of transporting in parallel particles in complex geometries profiting from parallelism at all levels (hardware and software). At the core of GeantV is a scheduling engine orchestrating the progress of the simulation. The project has demonstrated performance gains obtained by propagating multiple tracks from multiple events in parallel, grouping particles according to geometry locality criteria and aiming to do the same for physics locality (i.e. type of particle, energy of the particle, volume of the detector, etc.), increasing the instruction throughput and enhancing data locality in the particle transportation process. The scheduler groups particles into *baskets* that are dispatched to the relevant component (geometry or physics) for vectorized treatment. To improve cache-efficiency and to accomodate SIMD memory access, particles are packed in structures of arrays (SOA).

Code portability is assured by the introduction of software insulation layers between the platform independent code and the hardware specific code, called backends [3, 4]. Apart from the core scheduler, two are the main components of the GeantV project: the geometry and the physics libraries (VecGeom[5, 6] and VecPhys[7, 8]). They have to be designed to process multiple particles in a vectorized manner in order to obtain performance by maximising data and instruction locality. Vectorization of geometry algorithms and of physics models requires the algorithm to be recast to minimise conditional branches and to maximise instruction throughput. To exploit the potential of vectorization and accelerators in the physics models, appropriate sampling techniques have been implemented and tested. Some of these techniques introduce intervals and discrete tables. We identify artefacts that are introduced by discrete sampling techniques and determine the energy range in which these methods provide acceptable approximation. In this paper we introduce a suite of automated statistical tests created to verify the vectorized models by checking their consistency with the corresponding Geant4 models and to validate them against experimental data. We present preliminary verification results concerning the electromagnetic (EM) physics models developed for parallel computing architectures within the GeantV project, focusing in particular, as a case study, on the verification of the Klein-Nishina model for Compton process implemented in GeantV, against the corresponding model extracted from the Geant4 implementation.

2. A physics library for parallel architectures

In the context of a typical HEP simulation, a considerable fraction (30%-50%) of the time is spent in executing algorithms that sample cross sections and implement the physical processes. This makes it necessary to develop a Vectorized Physics library (VecPhys) which, following the parallel approach to the transport of particles proposed in GeantV, is capable of processing several tracks at a time to profit from the instruction level parallelism. The problem is anything but trivial, because during a simulation there are many different types of particles in flight, with different energies and characteristics, and they can undergo various processes and models that need to be sampled. To profit from vectorization gains it is therefore necessary to use a basketizing mechanism able to gather all the particles that undergo the same physics process using the same physics model.

Another crucial phase is the final state sampling that consists, once the specific model for the selected physics process is chosen, of sampling the final state of the particles involved in the interaction (the secondary and the potentially surviving primary). Most of the time the inverse function of the corresponding cumulative distribution function is not analytically calculable, therefore a "one-shot" sampling is not applicable. One of the sampling techniques commonly used in Geant4 and other Monte Carlo codes is the composition-rejection that allows to sample variables according to the probability distribution function (pdf) of the physics model.

However a feature of the rejection algorithm is that different tracks can require different trials to find an acceptable value, which clashes with the need of our vector hardware programming model, which maintains performance only when all lanes of a vector work in lockstep.

A variety of sampling methods has been and is being developed, in order to cope with the needs of different distributions required by EM processes [8]; one of these involves the well known alias sampling technique [9]. Due to the complexity of the code and the revision to try to optimise it, improving both vectorization and accuracy, it is particularly important to have a robust verification suite which will ensure that physics accuracy & performance are maintained to the level required.

Since in a typical HEP event 30-40% of the time is spent on Electromagnetic (EM) processes and most of the secondary particles produced by primary particle interactions with matter are electrons or photons, the development of the VecPhys library has focused on the EM physics with e^- , e^+ and γ as starting point. The goal is to write EM physics models that are able to deal with multiple tracks leading to performance gains while obtaining accurate results. The R&D activity has been evolving following two separated but complementary paths. The first one is focused on the vectorization of existing Geant4 EM physics models, via the exploration, implementation and testing of alternative sampling techniques and on the verification and validation of the new vectorized physics models. The second one is focused on a new implementation of the EM physics models to tackle the new challenges posed to the HEP community by new accelerators projects like, high-luminosity LHC [10] and the Future Circular Collider (FCC)[11]. In this paper we focus on verification/validation created for the first approach.

Table 1. Example of the p-value table automatically generated running the statistical validation suite for a preselected number of interactions of $N_i = 10^5$.

| Input Energy | Validation Quantity | χ -squared test p-value |
|--------------|-----------------------------|------------------------------|
| 0.01MeV | E_{γ}^{out} | 0.31306 |
| | $\cos\theta_{\gamma}^{out}$ | 0.280804 |
| | $\cos\theta_{e^-}^{out}$ | 0.141958 |
| 0.1MeV | E_{γ}^{out} | 0.857909 |
| | $\cos\theta_{\gamma}^{out}$ | 0.560302 |
| | $\cos\theta_{e^-}^{out}$ | $1.5667e^{-15}$ |
| ... | ... | ... |
| ... | ... | ... |
| ... | ... | ... |

3. Statistical verification and validation suite

In the context of high energy physics detector simulation it is critical to verify and validate physics models. For this purpose we have been developing a statistical verification and validation suite to compare, test and validate all the relevant physical quantities of every specific physics process. It consists of different automated regression analysis tests that can be run on the results of a simulation to identify deficiencies of algorithms or errors in implementation. These tests are generic and can be used either to check whether a random sample is compatible with a given theoretical distribution or to assess if two empirical distributions are sampled from the same theoretical distribution. Before running a statistical test it is necessary to define the hypothesis that we want to verify H_0 , known as null hypothesis, and the alternative H_1 , that usually is the hypothesis complementary to H_0 . Then it is necessary to choose a statistic test $t(x)$, which is a function of the sample values $x = \{x_1, ..., x_N\}$, in such a way that the test result highlights the difference between the distributions belonging to H_0 , $f(t|H_0)$, and those belonging to H_1 . Three

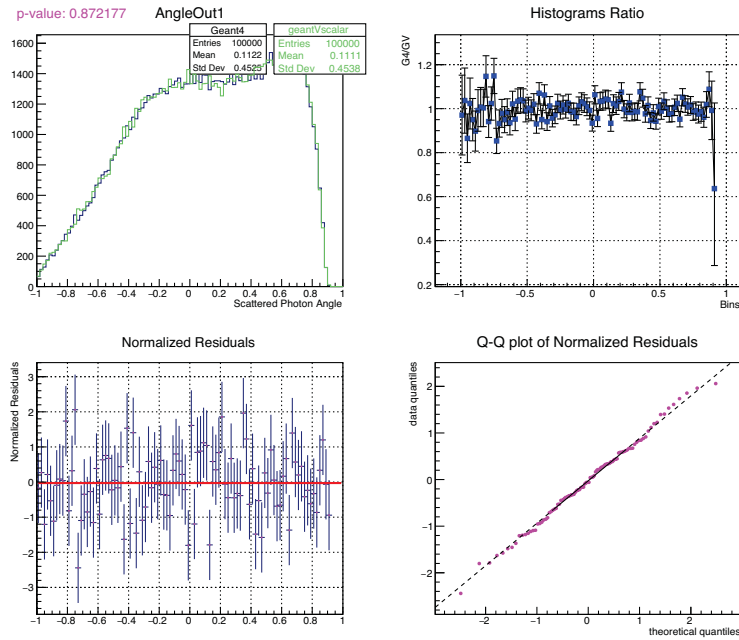


Figure 1. Statistical validation and verification suite output example. From the top right in a clockwise direction: direct comparison between the benchmark results and the obtained one, ratio between the histograms entries, q-q plot of the normalised residuals and normalised residuals from the Pearson χ -squared test.

goodness-of-fit test available in literature have been included so far: the Pearson χ -squared test with analysis of the residuals [12], the Kolmogorov-Smirnov [13] and the Anderson-Darling test [14]. The suite runs the appropriate tests and gives as feedback a binary answer easy to interpret: i.e. the null hypothesis H_0 is accepted or rejected. Furthermore it gives as output a continuous parameter, the p-value $p(t)$, which is a function of the statistical test, to assess the goodness-of-fit in a quantitative way measuring the compatibility of the sample with the null hypothesis. Moreover the tool produces graphical outputs which include analysis of the residuals helping in the identification of potential problems. Before starting the test it is necessary to set some parameters: the input files that contain data to compare, the primary particle input energies, the physical quantities and the name of the physical models to be tested. It is necessary also to specify whether the distributions are binned or unbinned in order to run the appropriate statistical tests. At the end of the test, two output files are produced: one containing the table with the p-values related to all the physical quantities tested for all energy ranges considered, and the other one containing the diagrams with the results.

In Table 1 is shown an example of a first p-value table produced as output running Pearson χ -squared test on the Klein-Nishina model for Compton. For every input energy all the validation quantities are listed with the respective p-values evaluated from the selected test. As it can be seen some of the physics quantities don't pass the test (i.e. $\cos(\theta_e^{out})$ with a p-value = $1.5667e^{-15}$). The graphical output is instead arranged into quadrants, as it is shown in Figure 1, and plots, starting from the top right and proceeding clockwise: direct comparison between the reference sample (i.e. the Geant4 output) and the one obtained from simulation (i.e. the GeantV scalar version output), the ratio between histograms entries, the quantile-quantile plot (q-q plot [15]) of the normalized residuals, and the normalized residuals from the Pearson's χ -squared test.

The analysis of the normalized residuals is very useful because it highlights clearly what

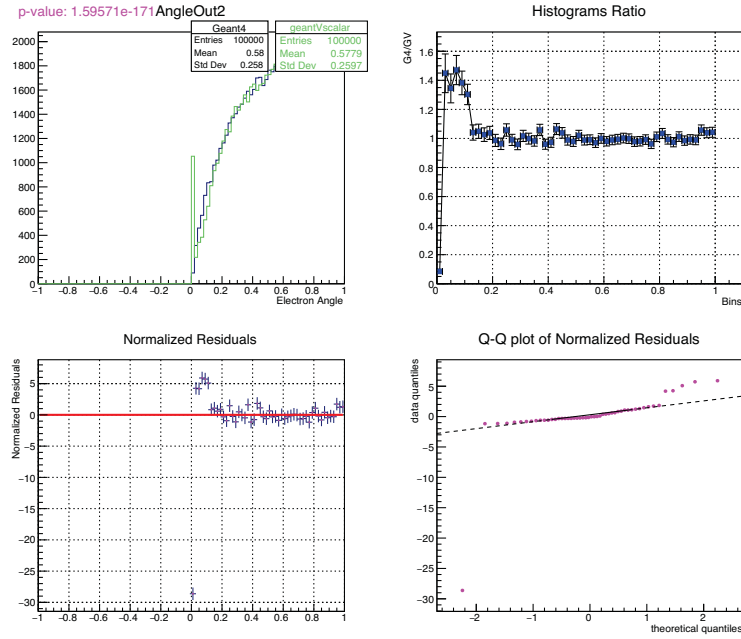


Figure 2. Example of problems/bugs detection obtained running the statistical validation suite. The verification plots are referring to the outgoing e^- scattering angle for the Klein-Nishina model for a photon input energy $E_{\gamma}^{in} = 0.01\text{MeV}$.

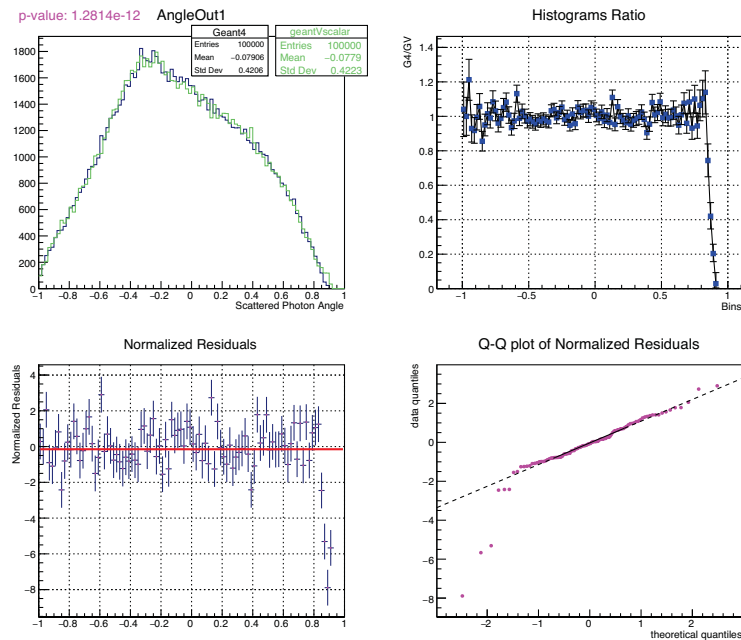


Figure 3. Example of problems/bugs detection obtained running the statistical validation suite. The verification plots are referring to the outgoing γ scattering angle for the Klein-Nishina model for a photon input energy $E_{\gamma}^{in} = 0.01\text{MeV}$.

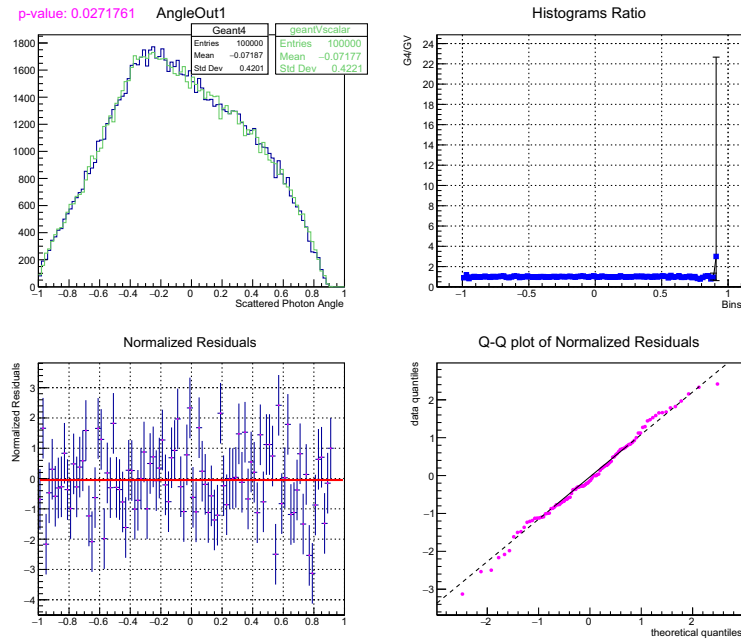


Figure 4. Verification plots of the outgoing γ energy for the Klein-Nishina model for a photon input energy $E_{\gamma}^{in} = 0.1 \text{ MeV}$.

bins contribute substantially to the final p-value. The presence of the q-q plot of normalized residuals completes the analysis since it helps understanding the correlation between the two distributions. In particular, if the hypothesis of statistical compatibility is accepted, according to the theory the normalized residuals of the χ -squared test should be distributed like a normal distribution $N(0,1)$. The q-q plot therefore is a normal probability plot that compares the results of the quantiles of the first distribution (normalized residual) against the quantiles of the normal distribution, identifying departures from normality. When points lie on the diagonal of the quadrant it means that the two distributions are linearly related and that H_0 can be accepted. The points mutual distribution provides additional statistical information on the form, outliers, skewness, kurtosis and other features of the distributions.

3.1. Test case: Verification of the Klein-Nishina Compton Model

The goal of the verification phase is to check the results of new physics models against pre-existent/already tested models. Data are typically stored in histograms which represent the input files for the verification tests. When comparing two histograms, the null hypothesis H_0 usually consists of the assumption that the two histograms represent samples coming from the same population. To decide whether to accept or reject the null hypothesis, a p-value should be calculated and a significance level threshold must be derived. The hypothesis of identity is rejected if the p-value is lower than this threshold. Typically used values are for example 0.1, 0.05 or 0.01. In our analysis we chose a threshold value of 0.05. Simulated quantities such as the final energy and the scattering angle of the primary particle as well as the kinematic of secondary particles have been compared and verified with respect to results obtained running Geant4. Figures 2 and 3 show some kinematic distributions of outgoing particles related to Compton scattering as a validation example. In particular, plots are referring respectively to the outgoing γ and e^- scattering angles for the Klein-Nishina model for a photon input energy

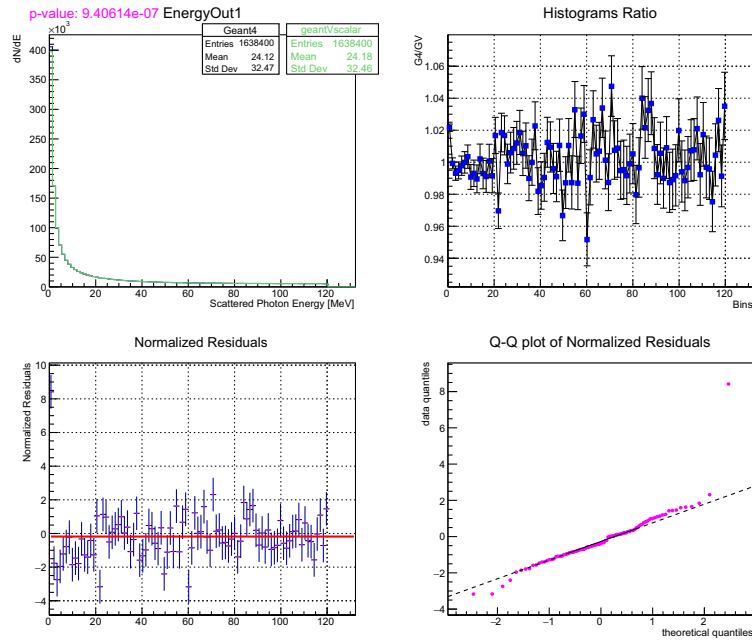


Figure 5. Verification plots of the outgoing γ energy for the Klein-Nishina model for a photon input energy $E_{\gamma}^{in} = 120.226 \text{ MeV}$.

$$E_{\gamma}^{in} = 0.01 \text{ MeV}.$$

The graphic analysis helps in the investigation of a particular p-value. For example, in Figure 2 is shown a case where we have a very low p-value, $p(t) = 1.59571e^{-171}$. Looking at the plot is clearly visible a problem within the first bin, which significantly contributes with a normalized residual of about -28.5 and this is visible already in the comparison between histograms. But the overall value also depends on the following 5 bins each of which has a residual that is around 5. This is evident in the normalized residuals plot but mostly clear looking at the q-q plot and at the outliers that deviate substantially from the normality. Figure 3 instead, shows a situation in which, even if there is a very low p-value, the problem is not immediately identifiable from the comparison between the two histograms. But if we look at the other diagrams we can easily isolate the residuals that are contributing with their weight and that are located in the final part of the distribution. Through this analysis it was possible to identify some implementation problems and correct them. As a result in Figure 4 we see the diagrams related to the same variable shown in Figure 3, this time with totally different behaviour. However, it is not always possible to easily solve problems highlighted by statistical analysis and in some cases they requires further studies. For instance, the alias sampling method introduces some discretization operations that have relevant consequences when the pdf has a strongly non-linear trend within the bin. It is the situation shown in Figure 5 where the pdf is very peaked to the left and the error introduced by the discretization cannot be trivially eliminated. In the case of Klein-Nishina model this pattern occurs for γ input energies above 100 MeV . We are implementing solutions to overcome such problems using different approaches depending on the input particle energy range. For energies above the threshold we are combining for example an adaptive binning approach with useful variable transformations. All the adopted solutions depend on the shape of the pdf from which we want to sample, and they will be presented in a separate paper.

4. Conclusions

In this paper we introduced a statistical analysis suite designed for the GeantV project to verify and validate physics models implemented for parallel architectures within the project context. The suite has a generic character and consists of different automated regression analysis tests that can be executed on two pairs of samples stored in histograms. The tool helps in interpreting results of the statistical tests through some graphical output consisting of p-values tables, analysis of residuals and normal probability plots. As a case of study the verification plots of the Klein-Nishina vectorized model part of the VecPhys library against the corresponding Geant4 model have been presented. The suite has showed to be a useful tool to check the consistency of the compared data thanks to different outputs coming from the tests run. The development of the suite is still in progress and ongoing work is related to the extension of the pool of available tests and the identification of appropriate multivariate statistics that will be particularly useful to perform regression analysis between different releases of the code. Once the verification phase will be successfully concluded, the future work foresees the use of the suite for the validation of VecPhys EM models against the corresponding experimental data.

5. References

- [1] Apostolakis J, Brun R, Carminati F, Gheata A and Wenzel S 2014 *Journal of Physics: Conference Series* **523** 012004 URL <http://stacks.iop.org/1742-6596/523/i=1/a=012004>
- [2] Amadio G, Apostolakis J, Bandieramonte M, Bhattacharyya A, Bianchini C, Brun R, Canal P, Carminati F, Duhem L, Elvira D, de Fine Licht J, Gheata A, Iope R L, Lima G, Mohanty A, Nikitina T, Novak M, Pokorski W, Seghal R, Shadura O, Vallecorsa S and Wenzel S 2015 *Journal of Physics: Conference Series* **664** 072006 URL <http://stacks.iop.org/1742-6596/664/i=7/a=072006>
- [3] Apostolakis J, Bandieramonte M, Bitzes G, Cosmo G, de Fine Licht J, Duhem L, Gheata A, Lima G, Nikitina T and Wenzel S 2015 *GPU Computing in High-Energy Physics, Pisa (Italy), 10 Sep 2014 - 12 Sep 2014* 95–102 URL <http://bib-pubdb1.desy.de/record/291334>
- [4] Portable SIMD and the VecCore library <https://indico.cern.ch/event/570876/contributions/2347250/attachments/1359720/2057229/Portable-SIMD-and-the-VecCore-Library-2016-10-24.pdf>
- [5] Apostolakis J, Brun R, Carminati F, Gheata A and Wenzel S 2014 *Journal of Physics: Conference Series* **513** 052038 URL <http://stacks.iop.org/1742-6596/513/i=5/a=052038>
- [6] Apostolakis J, Bandieramonte M, Bitzes G, Brun R, Canal P, Carminati F, Cosmo G, Licht J C D F, Duhem L, Elvira V D, Gheata A, Jun S Y, Lima G, Nikitina T, Novak M, Sehgal R, Shadura O and Wenzel S 2015 *J. Phys.: Conf. Ser.* **608** 012023. 6 p URL <http://cds.cern.ch/record/2159061>
- [7] Amadio G, Apostolakis J, Bandieramonte M, Bianchini C, Bitzes G, Brun R, Canal P, Carminati F, de Fine Licht J, Duhem L, Elvira D, Gheata A, Jun S Y, Lima G, Novak M, Presbyterian M, Shadura O, Seghal R and Wenzel S 2015 *Journal of Physics: Conference Series* **664** 092013 URL <http://stacks.iop.org/1742-6596/664/i=9/a=092013>
- [8] Amadio G, Ananya A, Apostolakis J, Aurora A, Bandieramonte M, Bhattacharyya A, Bianchini C, Brun R, Canal P, Carminati F, Duhem L, Elvira D, Gheata A, Gheata M, Goulas I, Iope R, Jun S Y, Lima G, Mohanty A, Nikitina T, Novak M, Pokorski W, Ribon A, Seghal R, Shadura O, Vallecorsa S, Wenzel S and Zhang Y 2016 *Journal of Physics: Conference Series* **762** 012014 URL <http://stacks.iop.org/1742-6596/762/i=1/a=012014>
- [9] Walker A J 1977 *ACM Trans. Math. Softw.* **3** 253–256 ISSN 0098-3500 URL <http://doi.acm.org/10.1145/355744.355749>
- [10] High-luminosity LHC <http://hilumilhc.web.cern.ch>
- [11] Future Circular Collider <https://fcc.web.cern.ch>
- [12] Pearson K 1922 *Biometrika* **14** 186–191 ISSN 00063444 URL <http://www.jstor.org/stable/2331860>
- [13] Smirnov N 1948 *Ann. Math. Statist.* **19** 279–281 URL <http://dx.doi.org/10.1214/aoms/1177730256>
- [14] Anderson T W and Darling D A 1952 *Ann. Math. Statist.* **23** 193–212 URL <http://dx.doi.org/10.1214/aoms/1177729437>
- [15] Chambers J 1983 *Graphical methods for data analysis* Chapman & Hall statistics series (Wadsworth International Group) ISBN 9780534980528 URL <https://books.google.ch/books?id=I-tQAAAAAAJ>