ORIGINAL ARTICLE

Challenges in Monte Carlo Event Generator Software for High-Luminosity LHC

The HSF Physics Event Generator WG · Andrea Valassi[1] · Efe Yazgan[2] · Josh McFayden[1,3,4] · Simone Amoroso[5] · Joshua Bendavid[1] · Andy Buckley[6] · Matteo Cacciari[7,8] · Taylor Childers[9] · Vitaliano Ciulli[10] · Rikkert Frederix[11] · Stefano Frixione[12] · Francesco Giuli[13] · Alexander Grohsjean[5] · Christian Gütschow[14] · Stefan Höche[15] · Walter Hopkins[9] · Philip Ilten[16,17] · Dmitri Konstantinov[18] · Frank Krauss[19] · Qiang Li[20] · Leif Lönnblad[11] · Fabio Maltoni[21,22] · Michelangelo Mangano[1] · Zach Marshall[3] · Olivier Mattelaer[22] · Javier Fernandez Menendez[23] · Stephen Mrenna[15] · Servesh Muralidharan[1,9] · Tobias Neumann[14,24] · Simon Plätzer[25] · Stefan Prestel[11] · Stefan Roiser[1] · Marek Schönherr[19] · Holger Schulz[17] · Markus Schulz[1] · Elizabeth Sexton-Kennedy[15] · Frank Siegert[26] · Andrzej Siódmok[27] · Graeme A. Stewart[1]

Abstract

We review the main software and computing challenges for the Monte Carlo physics event generators used by the LHC experiments, in view of the High-Luminosity LHC (HL-LHC) physics programme. This paper has been prepared by the HEP Software Foundation (HSF) Physics Event Generator Working Group as an input to the LHCC review of HL-LHC computing, which has started in May 2020.

Keywords Monte Carlo · Physics event generator · LHC experiments · WLCG · High-luminosity LHC

Andrea Valassi, Efe Yazgan and Josh McFayden: Editors.

✉ Andrea Valassi
andrea.valassi@cern.ch

Efe Yazgan
efe.yazgan@cern.ch

Josh McFayden
mcfayden@cern.ch

[1] CERN, Geneva, Switzerland

[2] National Taiwan University, Taipei, Taiwan

[3] Lawrence Berkeley National Laboratory, Berkeley, USA

[4] University of Sussex, Brighton, UK

[5] DESY, Hamburg, Germany

[6] University of Glasgow, Glasgow, UK

[7] LPTHE, Sorbonne Université and CNRS, Paris, France

[8] Université de Paris, Paris, France

[9] Argonne National Laboratory, Lemont, USA

[10] INFN and Università di Firenze, Florence, Italy

[11] Lund University, Lund, Sweden

[12] INFN and Università di Genova, Genoa, Italy

[13] INFN and Università Tor Vergata, Roma, Italy

[14] University College London, London, UK

[15] Fermi National Accelerator Laboratory, Batavia, USA

[16] University of Birmingham, Birmingham, UK

[17] University of Cincinnati, Cincinnati, USA

[18] NRC "Kurchatov Institute"-IHEP, Protvino, Russia

[19] University of Durham, Durham, UK

[20] Peking University, Beijing, China

[21] Università di Bologna, Bologna, Italy

[22] Université Catholique de Louvain, Louvain-la-Neuve, Belgium

[23] Universidad de Oviedo, Oviedo, Spain

[24] Illinois Institute of Technology, Chicago, USA

[25] University of Vienna, Vienna, Austria

[26] Technische Universität, Dresden, Germany

[27] IFJ PAN and Jagiellonian University, Krakow, Poland

Introduction

Physics event generators are one of the computational pillars of any High Energy Physics (HEP) experiment, and in particular of the Large Hadron Collider (LHC) experiments. In this paper, we review the main software and computing challenges for the physics event generators used by the ATLAS [1] and CMS [2] experiments, in view of the high-luminosity running phase of the LHC experimental programme (HL-LHC), which should be operational from the end of 2027 [3]. This document has been prepared by the Physics Event Generators Working Group (WG) [4] of the HEP Software Foundation (HSF), as an input to the review of the HL-LHC computing strategy by the LHC Experiments Committee (LHCC) [5], which has started in May 2020 [6], as previously planned [7–10]. As is the case for the LHCC review, this paper focuses on ATLAS and CMS, but it also contains important considerations for ALICE [11] and LHCb [12]. The HSF has also prepared a more general document [13] for the LHCC, which covers the status and challenges in the broader area of common tools and community software.

This paper gives an overview of the many challenges in the generator area, and of the work that can be done to address them. Its outline is the following. Section "The HSF Physics Event Generator WG" gives an overview of the role and challenges of physics event generators in LHC computing, summarising the steps which led to the creation of the HSF generator WG, and its current activities. Section "Collaborative Challenges" describes the collaborative challenges in the development, use and maintenance of generator software for LHC physics. Section "Technical Challenges" gives more details about the computational anatomy of physics event generators, and the technical challenges in their development and performance optimization. Section "Physics Challenges (Increasing Precision)" summarizes some of the main open questions about the required physics accuracy of event generators at HL-LHC, and their impact on computational costs. Finally, in section "Conclusions" we compile a list of high-priority items on which we propose that the R&D on the computational aspects of generators (and in particular the activities of the HSF generator WG) should focus, in view of the more in-depth LHCC review of HL-LHC software that is currently scheduled for Q3 2021 [10].

It should be stressed that this paper focuses on the software and computing aspects of event generators, rather than on the underlying physics. To be able to describe the overall computational goals and structure of these software applications and put them in context, many of the relevant physics concepts are in any case mentioned and briefly explained. This is done using a language that tries to be somewhat accessible also to software engineers and computing experts with no background in particle physics, even if the resulting text is not meant to be an exhaustive overview of these complex issues from a theoretical point of view.

One should also note that, to some extent, some of the issues described in this paper, such as the collaborative challenges and human resource concerns related to the development and support of generator software, have already been raised in previous community efforts. These include, in particular, the HSF Community White Paper (CWP) [14] and the document [15] that was submitted as an input to the Open Symposium [16] on the Update of European Strategy for Particle Physics.

The HSF Physics Event Generator WG

Physics event generators are an essential component of the data processing and analysis chain of the LHC experiments, and a large consumer of resources in the Worldwide LHC Computing Grid (WLCG) [17]. All of the scientific results of the LHC experiments, such as precision measurements of physics parameters and searches for new physics, depend significantly on the comparison of experimental measurements to theoretical predictions, in most cases computed using generator software.

Using Monte Carlo (MC) techniques, generators allow both the calculation of differential and total cross sections and the generation of weighted or unweighted events for experimental studies (this is explained in more detail in section "Computational Anatomy of a MC Event Generator", where these concepts are briefly defined). Within the experiments, generators are used primarily to produce large samples of (mostly unweighted) events: this is the first step in the production chain for simulating LHC collisions, which is followed by detector simulation and event reconstruction. In each of the two general purpose LHC experiments, ATLAS and CMS, the overall number of events that are generated by the central production teams and passed through full detector simulation and event reconstruction, across all relevant physics processes, is of the order of magnitude of $O(10^{10})$ events for every year of LHC data taking. Typically, the sizes of these samples of simulated events are approximately a factor of 3 larger than the overall number of data events collected during the corresponding time range. These large-scale event generation campaigns have a computational cost, mainly in terms of the "compute" (i.e. CPU) resources used, the majority of which are provided by the WLCG infrastructure. The limited size of the simulated samples that can be produced under resource constraints is a source of major uncertainty in many analyses (for example, in Higgs boson measurements of both ATLAS [18] and CMS [19]). This

is an issue which is limiting the potential physics output of the LHC programme, and may get significantly worse at HL-LHC, where the projected computing needs of the experiments exceed the resources that are expected to be available [14], despite the fact that the most aggressive HL-LHC physics projections [20–23] assume no uncertainty due to the limited size of simulated samples.

When the HEP Software Foundation prepared its CWP [14] in 2017, the fraction of the ATLAS CPU resources in WLCG used for event generation was estimated [24] at around 20%. Beyond the existing projections, which assume the same level of theoretical precision as in the current event generation campaigns, concern was also raised that event generation would become computationally more expensive at the HL-LHC, where more complex calculations (e.g. beyond next-to-leading-order or with higher jet multiplicities) will be needed [25]. It was thus clear that speedups in generator software are needed to address the overall computing resource problem expected at the HL-LHC. This is of course also the case for the other big consumers of CPU (detector simulation and reconstruction), but until now these areas have had more focus, and significant speedups are already expected on the HL-LHC timescales, which has not been the case for generators. Other issues in the generator area, both technical and non-technical (e.g. funding, training and careers) also became obvious while preparing the CWP.

For these reasons, the HSF organised a three-day Workshop [26, 27] at the end of 2018 to focus on the software and computing aspects of event generators. Their usage in the experiments was reviewed, revealing a large discrepancy in the CPU budgets quoted by ATLAS and CMS, 14% and 1%, respectively, for 2017 [28]. This was attributed, at least partly, to the different packages and parameter settings used by the two experiments, but it was clear that further studies were needed.

A Working Group of the HSF on Physics Event Generators [4] was therefore set up at the beginning of 2019. The main focus of the WG so far has been to get a better understanding of the current usage of generators in the experiments, and to identify and prioritise the areas where computing costs can be reduced. In particular, the ATLAS and CMS compute budgets have been analysed in detail: currently, it is estimated that the fractions of WLCG compute allocations used for generation today are around 12% for ATLAS and 5% for CMS. In terms of absolute CPU time spent for event generation, the ratio between ATLAS and CMS is actually larger, as the overall ATLAS budget for compute resources is larger than that of CMS. To understand what causes this difference, detailed benchmarking of the computational costs of Sherpa [29] and MadGraph5_aMC@NLO [30] (in the following abbreviated as MG5_aMC) have also started [31, 32], as these are the two generators used for some of the most expensive event generation productions in ATLAS and CMS, respectively. The WG has also been active in other areas, such as in discussing the possible sharing of common parton-level samples by ATLAS and CMS [33], and in reviewing and supporting the efforts for porting generators to modern architectures, notably GPUs. This last activity is particularly important, as it has become increasingly clear that being able to run compute-intensive WLCG software workloads on GPUs [34] would allow the exploitation of modern GPU-based supercomputers at High Performance Computing (HPC) centers, and generators look like a natural candidate for this, as discussed later on in section "Modernisation of Generator Software".

Looking forward, the WG plans to continue its activities in the areas described above, but also to expand it in a few other directions. One of the goals of this paper is that of dissecting and analysing the many different challenges, both technical and non-technical, in the generator domain, to identify the specific areas where work is most urgently needed, or where the largest improvements are expected to be possible to reduce the gap between required and available computing resources at the time of HL-LHC. It should also be pointed out that the role of the WG in this context is mainly that of providing a forum for information exchange, and possibly supporting and coordinating common activities involving the collaboration of several teams or the comparison of their results, but most of the concrete work is generally expected to be done by the individual experiments or theoretical physicist teams.

Collaborative Challenges

In this section, we give an overview of the collaboration challenges in the development, use and maintenance of generator software for LHC. By and large, these are mainly non-technical challenges that concern human resources, i.e. actual people, and their organisation, training and motivation, rather than computing resources, software modules or theoretical physics models.

A Very Diverse Software Landscape

The landscape of generator software is extremely varied, even more than in detector simulation, event reconstruction or analysis workloads. For a review, see for instance Refs. [15, 35–37]. Different generators (Sherpa, MG5_aMC, the POWHEG BOX [38], Pythia [39], Herwig [40–42], Alpgen [43], etc.) are used in the community, mainly for two reasons: firstly, one needs multiple independent calculations with potentially different approximations to cross-check one another; and secondly, the different generators vary in their features (for example, some might simulate only a subset of the physics processes of interest). A given process may be

simulated with a different physics precision, e.g. leading-order (LO), next-to-leading-order (NLO), or next-to-next-to-leading-order (NNLO) in a power series expansion in the strong-force "coupling constant". Generating a sample also involves choices of hadronization and parton shower (PS) models (Pythia [39], Herwig [40–42], Ariadne [44], etc.), underlying event tunes [45–48], prescriptions for matching/merging[1] (MC@NLO [50], POWHEG [51], KrkNLO [52], CKKW [53], CKKW-L [54], MLM [55, 56], MEPS@NLO [57], MINLO [58], FxFx [49], UNLOPS [59], Herwig7 Matchbox [60–62], etc.), "afterburner" tools for simulating particle decays and quantum electrodynamics (QED) radiative corrections (EvtGen [63], Tauola [64], Photos [65], etc.), and other input parameters such as parton distribution functions (PDFs) [66], primarily via the LHAPDF library [67].

Various combinations of software libraries are thus possible, often written by different authors and some dating back many years, reflecting theoretical research within different teams. For a given process, the LHC experiments often use different software packages and settings from one another, and a single experiment can generate events using more than one choice. Many different packages and configurations may therefore need to be studied and improved to get cumulative CPU cost reductions. The large number of packages also complicates their long-term maintenance and integration in the experiments software and workflows, sometimes leading to Grid job failures and computing inefficiencies. Other packages are also absolutely critical for the whole generator community and must be maintained, even if their CPU cost is relatively low (Rivet [68], Professor [69], HepMC [70, 71], FastJet [72], etc.).

A Very Diverse Human Environment

A broad spectrum of skills and profiles are needed for the development and support of event generators: theorists (who create fundamental physics models, and design, develop and optimize most generator code); experimentalists working on research (who determine which types of event samples are required, and of which size); experimentalists working on computing (who implement, monitor and account execution of workflows on computing resources); software engineers and system performance experts (who may help to analyse and improve the efficiency of software applications and deployment models). This is a richness and opportunity, as some technical problems are best addressed by people with specific skills, but it also poses some challenges as these

technical problems are best addressed by bringing all these people together. Facilitating this cross-collaboration is one of the main goals of the WG.

Training challenges. Theorists and experimentalists often lack formal training in software development and optimization. Software engineers, but also many experimentalists, are not experts in the theoretical physics models implemented in MC codes.

Communication challenges. It is difficult to find a shared terminology and set of concepts to understand one another: notions and practices that are taken for granted in one domain may be obscure for others. An example: there are many articles about the physics in generators, but software engineers would need papers describing the main software modules and overall data and control flow. Similarly, there are only very few articles where the experiments describe the software and computing workflows of their large scale MC productions (Ref. [73] is one such example for LHCb).

Career challenges. Those working in the development, optimization and execution of generator software provide essential contributions to the success of the HL-LHC physics programme and it is critical that they get the right recognition and motivation. However, theorists, in general, get recognition from the papers they publish and from the citations on these, and they may not be motivated to work on software optimizations that do not have enough theoretical physics content to advance their careers. Generator support tasks in the experiments may also not be valued enough to secure jobs or funding to experimentalists pursuing a career in research.

Mismatch in usage patterns and in optimization focus. The way generators are built and used by their authors is often different from the way in which they are deployed and integrated by the experiments in their software frameworks and computing infrastructure. The goals and metrics of software optimization work may also differ, as discussed more in detail in section "Technical Challenges". Theorists, who typically work with weighted events and fast detector parameterizations if any, are mainly interested in calculating cross sections and focus on minimising the phase space integration time for a given statistical precision. The LHC experiments typically run large scale productions for generating fully exclusive events, which are mostly unweighted as they must be processed through expensive detector simulation and event reconstruction steps: therefore, they need to maximize the throughput of events generated per unit time on a given computing system.

Programming languages. Attracting collaborators with a computer science background to work on generators, especially students, may also be complicated by the fact that critical components of some generator packages are written in Fortran, which is rarely used in industry and less popular among developers than other programming languages. Some

---

[1] In this paper, we use the definitions of matching and merging given in Ref. [49], which are briefly hinted at in section "Inefficiencies in Unweighted Event Generation".

of the generators also do not use industry standard version control systems, making it harder to contribute code.

Technical Challenges

In this section, we give more details about the technical challenges in the software development and performance optimization of MC physics event generator codes. To this end, it is useful to first give a brief, high-level, reminder of their computational goals and internal data flows, and of the typical production workflows used by the experiments.

Computational Anatomy of a MC Event Generator

Particle physics is based on quantum mechanics, whose description of Nature is intrinsically probabilistic. The predictions of HEP theoretical models that are numerically computed in event generators (through a combination of quantum field theory methodologies and phenomenological approximations), and which can be compared to experimental measurements, ultimately consist of probabilities and probability density functions.

In particular, the probability that a collision "event" with a given "final state", i.e. including n particles of given types, is observed in the collision of the LHC proton beams, is expressed in HEP in terms of the concept of a "cross section". In general terms, a cross section σ represents the number of events $N_{exp} = \sigma \mathcal{L}$ that are expected per unit "integrated luminosity" $\mathcal{L}$ of the colliding beams (a parameter that depends on their intensities and geometries, and on the overall duration of data-taking time). More in detail, a differential cross section, $\frac{d\sigma}{dO}$, with respect to an observable O (such as a rapidity or a transverse momentum), refers to the observation of the desired final state at different points dO of the observable "phase space"; conversely, its integral $\sigma = \int_{\Omega_O} \frac{d\sigma}{dO} dO$ is referred to as the total cross section, if over the entire phase space, or as a fiducial cross section, if over a well delimited region Ω_O of the phase space (the so-called acceptance).

In this context, the computational core of a physics event generator is the code that numerically calculates, from first principles, the fully differential cross section $\frac{d\sigma}{d\Phi_n}(\mathbf{x})$ for the "hard scattering" process that leads to the desired n-particle final state; this is computed as a function of the complete kinematical configuration $\mathbf{x}$ of the elementary particles, or "partons", involved in this "hard interaction" for an individual collision event. In the majority of cases, the calculation of $\frac{d\sigma}{d\Phi_n}$ is implemented by identifying all Feynman diagrams contributing to this process, and calculating the "invariant amplitude" or "matrix element" (ME) for all of these diagrams combined (although there are also generators

where matrix elements are computed using algorithms not based on Feynman diagrams [43, 74]).

For LHC processes, the kinematical configuration $\mathbf{x} = \{x_1, x_2, \Phi_n\}$ of a collision event essentially consists of a vector Φ_n, including four real numbers (related to their energy, mass and directions) for each of the n outgoing (final state) partons, and of two real numbers x_1 and x_2 representing the momentum fractions of the two incoming (initial state) partons. As described later on in Eq. 1, $\frac{d\sigma}{d\Phi_n}(\mathbf{x})$ is, together with two parton distribution functions $p(x_1)$ and $p(x_2)$, the central ingredient in the computation of a function $f(\mathbf{x})$, which essentially describes the probability distribution in the space of all possible kinematical configurations $\mathbf{x}$, and from whose integral in this space other relevant cross sections may be computed, $\sigma = \int f(\mathbf{x}) d\mathbf{x}$.

Integration and unweighted event generation. Given the function $f(\mathbf{x})$, physics event generators are commonly used in HEP to solve two types of computational problems, which are related to each other and generally addressed within a same execution of the software, as discussed more in detail later on. The first goal ("phase space integration") is to compute a cross section as the integral of $f(\mathbf{x})$ over the relevant phase space region. The second goal ("unweighted event generation") is to draw random samples of events whose kinematical configurations $\mathbf{x}$ are distributed according to the theoretical prediction $f(\mathbf{x})$.

Both of these goals are achieved using Monte Carlo (MC) methods (see Refs. [75, 76] for early reviews of this technique in HEP). The distinctive feature of MC methods is their reliance on the generation of random numbers (or, more precisely, of "pseudo-random" [77] numbers).[2] In particular, the starting point of both MC phase space integration and MC unweighted event generation is the calculation of $f(\mathbf{x})$ for a large sample of events $\mathbf{x}_i \in \{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$, drawn at random from a known probability density function $g(\mathbf{x})$. More specifically:

---

[2] As discussed in Ref. [78], phase space integration may also be performed using classical numerical methods (which do not belong to the MC category), or "quasi-MC" methods based on "quasi-random" numbers [76, 79]. The classical methods, such as Newton-Cotes formulas and Gaussian quadrature rules (one example is the Gauss-Kronrod algorithm which is used for numerical integration in the TOP++ [80] program), work well for one-dimensional problems, but tend to be inefficient for multi-dimensional integrals, where MC methods using pseudo-random numbers converge much faster, and quasi-MC methods using quasi-random numbers even faster. Unlike phase space integration, however, unweighted event generation can only be addressed by MC methods using pseudo-random numbers, as neither of the other approaches is applicable: classical integration methods because they do not involve random numbers, and quasi-MC techniques because quasi-random numbers in a sample are highly correlated to one another.

1. MC phase space integration consists in drawing a random sample of events $\mathbf{x}_i$ from the sampling function $g(\mathbf{x})$, and in numerically calculating an estimator of the integral $\sigma = \int f(\mathbf{x}) d\mathbf{x}$, as the average of the "weight" $w_i = w(\mathbf{x}_i) = f(\mathbf{x}_i)/g(\mathbf{x}_i)$ for all the events $\mathbf{x}_i$ in the sample. It should be noted that this is not a deterministic approach, in the sense that the result of the calculation may change if a different random sample is used: it is easy to show, however, that the estimator is unbiased, and that its variance decreases as $1/N$ if the number of events N in the sample is increased. From a software point of view, the output of MC phase space integration[3] is essentially only one number, the estimate of the integral $\sigma = \int f(\mathbf{x}) d\mathbf{x}$ over the acceptance Ω_O; alternatively, several numbers may also be calculated, representing the values of $\frac{d\sigma}{dO}$ computed as the MC integrals of $f(\mathbf{x})$ over different regions of phase space within Ω_O.

2. MC unweighted event generation consists in drawing a random sample of events $\mathbf{x}_i$ from the sampling function $g(\mathbf{x})$, and in randomly rejecting some of them depending on the ratio of $w(\mathbf{x}_i)$ to the maximum weight $w_{\max}$ over the phase space. For each event, an accept-or-reject (or "hit-or-miss") decision is taken by drawing a random number R uniformly distributed between 0 and 1: the event is accepted if $R < w(\mathbf{x}_i)/w_{\max}$, and rejected otherwise. The resulting events, whose distribution is now described by $f(\mathbf{x})$ rather than by $g(\mathbf{x})$, are referred to as "unweighted" in the sense that they all have the same weight, which by convention is equal to 1. A special case of unweighting, producing events whose weights can be either $+1$ or -1, exists for calculations leading to events with negative weights: this is described later on. From a software point of view, the output of MC unweighted event generation is a sample of events, i.e. essentially a sample of vectors $\mathbf{x}_i$.

The choice of the sampling algorithm (e.g. VEGAS [81, 82]), or equivalently of the function $g(\mathbf{x})$, is very important. The closer $g(\mathbf{x})$ is to $f(\mathbf{x})$, that is to say the more constant the weight $f(\mathbf{x})/g(\mathbf{x})$ is over the entire phase space, the more precise is the integration (i.e. the lower the variance on the result) for a given sample size, and the more efficient is the unweighting procedure (i.e. the lower the fraction of events rejected).

---

[3] To avoid misunderstandings, it should be noted that, in an inconsistent way, the term "phase space integration" is also commonly used to indicate the computational step in the software before unweighted event generation, which is needed not only to compute a first coarse estimate of cross sections, but also to iteratively optimize the sampling algorithm (i.e. the choice of the sampling function $g(\mathbf{x})$), and to compute the maximum value $w_{\max}$ of $w(\mathbf{x})$ over the relevant region of phase space. This is further discussed below.

It should be noted that the experiments also do physics analysis with samples of weighted events, which they produce for instance through "biasing" techniques, as discussed in section "Inefficiencies in Unweighted Event Generation". Wherever possible, however, unweighted events (and in particular events with a positive weight $+1$) are preferred, as smaller event samples are required than when using events with non-uniform weights, resulting in overall savings of compute and storage resources.

Internal software workflow. Schematically, the internal software workflow of a typical generator is the following: first, when necessary (i.e. when the process is too complex to be manually hardcoded in advance), the source code to compute the differential cross section $\frac{d\sigma}{d\Phi_n}$ of the hard process, which is needed to derive $f(\mathbf{x})$, is produced through automatic code generation, after identifying the relevant Feynman diagrams; a "phase space integration" step follows, where event samples are iteratively drawn not only to provide a first coarse estimate of the relevant cross sections, but also to optimize the sampling function $g(\mathbf{x})$ and to estimate the maximum weight $w_{\max}$; parton-level unweighted events are then generated using the final, frozen, $g(\mathbf{x})$ and $w_{\max}$; parton showers, hadronization and hadron decays to stable particles are finally applied on top of those "parton-level" events. During the unweighted event generation step, "merging" prescriptions (described in more detail in section "Inefficiencies in Unweighted Event Generation") may also need to be applied, after parton showers and before hadronization; experiment-level filters and other techniques such as forced decays or forced fragmentation may also be applied [83, 84], for instance to produce event samples containing specific decays of B hadrons.

The internal workflow of a generator application is actually more complex than described above, because many different hard interactions may contribute to the simulated process. To start with, for hadron colliders like the LHC, the hard interactions take place not between two protons, but between two of the partons in their internal substructure (quarks of different flavors, and gluons): this implies that separate integrals for all possible types of initial state partons, using different sets of diagrams and of functions $f(\mathbf{x})$, must be considered. Using the factorisation theorem [85], which allows separating perturbative (i.e. ME) and non-perturbative (parton distribution function) calculations in quantum chromodynamics (QCD), the total cross section may be written [35] as

$$\sigma = \sum_{a,b} \int dx_1 p_a(x_1) \int dx_2 p_b(x_2) \int d\Phi_n \frac{d\sigma_{ab}}{d\Phi_n}(x_1, x_2, \Phi_n), \quad (1)$$

i.e. as the convolution, by the appropriate parton distribution functions $p_a(x_1)$ and $p_b(x_2)$, of the differential cross section $\frac{d\sigma_{ab}}{d\Phi_n}$ for the production of n final state particles with proper-

ties Φ_n, in the hard interaction of two partons of types a and b with momentum fractions x_1 and x_2, respectively.

In addition, NLO calculations imply the need to compute two separate classes of integrals, which involve two different classes of Feynman diagrams and of functions $f(\mathbf{x})$, because matrix elements need to be separately computed for standard "$\mathbb{S}$-events" and hard "$\mathbb{H}$-events" [50], i.e. for final states with n body kinematics Φ_n (at tree level and one loop) and $n + 1$ body kinematics Φ_{n+1} (at tree level), respectively; "matching" prescriptions are then needed to ensure that parton showers are used appropriately in both types of events (see also for instance Refs. [86–88] for detailed presentations that include a graphical representation of these issues). In NNLO calculations, the situation is similar to that of NLO calculations, and even more complex.

Experiment production workflows. Phase space integration (i.e. the optimization of the phase space sampling algorithm) is a resource intensive step, but in many cases it is only executed once in a given experiment production; this is known as the creation of "gridpacks" in MG5_aMC and POWHEG, or "sherpacks" or "integration grids" in Sherpa[4]. For instance, creating a typical MG5_aMC gridpack for V+jets (i.e. a W or Z vector boson produced in association with quarks or gluons) at NLO may take up to several weeks on one multi-core node, or up to several days in a typical cluster usage scenario; see also Ref. [92] for further details about how gridpacks are used in CMS. The generation of unweighted event samples, conversely, is where the LHC experiments spend essentially all of their yearly generator CPU budgets: when pre-computed integration grids are available, this typically involves many Grid jobs submitted in parallel with different random number seeds and thus unrelated to one another, all of them reading the same integration grids as an input and storing events on their own output files. In principle, every Grid job could also go through the whole event generation chain, including both phase space integration and unweighted event generation, but this is an inefficient workflow which the experiments only use in specific cases, e.g. for productions involving simple physics processes or few events, where phase space integration is relatively fast and inexpensive and where the overhead from

repeating it in each Grid job is negligible with respect to the overall CPU cost of the production, or for generators lacking the option to create integration grids. It should be noted, in any case, that also the workflows involving one gridpack creation and several unweighted event generation jobs can be somewhat inefficient, if the initialisation phase of each Grid job is not negligible with respect to its overall duration; this may happen, for example, if the pre-computed integration grids are very large and take a long time to load [93].

Computational costs. The computational cost of a MC application roughly scales with the number of points $\mathbf{x}$ where the function $f(\mathbf{x})$ is computed. This is true both for gridpack creation, where the cost scales with the number of events sampled during phase space integration (itself a function of the accuracy required for this step), and for unweighted event generation, where the cost scales with the overall number of events drawn prior to rejection by the unweighting algorithm. As a consequence, the most obvious approach to reduce the overall computational cost of event generation is simply to try and decrease the number of points $\mathbf{x}$ for which $f(\mathbf{x})$ is computed. This is described in detail in section "Inefficiencies in Unweighted Event Generation", where the possible reduction of many large inefficiencies in unweighted event generation is discussed, as well as possible strategies for reusing events for more than one goal.

In addition, the intrinsic cost per event of computing $f(\mathbf{x})$ approximately scales itself with the number of Feynman diagrams contributing to that process. In particular, with respect to LO calculations for a given process, NLO and especially NNLO calculations for the same process involve much higher numbers of diagrams, some of which ("loop diagrams") are also intrinsically more complex to compute. Matrix element calculations are in fact performed as a power series expansion in terms of the strong-force coupling constant α_s (which is smaller than 1); the difference between LO, NLO and NNLO calculations is primarily that of considering the following level in this power series expansion, which leads to a roughly factorial increase in computational complexity. It should be pointed out, nevertheless, that NLO calculations for simple processes with low final state multiplicities may be computationally cheaper than LO calculations for complex processes with high final state multiplicities. In summary, it would thus seem that the intrinsic cost per event $\mathbf{x}$ of computing $f(\mathbf{x})$ is to some extent incompressible, because of the relatively fixed amount of arithmetic calculations that this involves. One of the only obvious strategies for reducing this cost consists in improving the efficiency with which these arithmetic operations are performed on modern computing systems, for instance through the use of parallel programming techniques such as vectorization or GPU programming, as discussed later in section "Modernisation of Generator Software". In addition, radically new approaches are also being worked on,

---

[4] While "gridpacks" and "integration grids" serve essentially the same purpose in different generators and contain similar information (the parametrization of an optimized and frozen tuning of the sampling algorithm, as well as an estimate of the maximum event weight over the phase space), it should be noted that the word "grid" in these two terms alludes to two very different concepts. The term gridpacks [89, 90], or "Grid packages" refers in MG5_aMC to packages suitable to be sent over for event generation on Grid nodes (e.g. on the nodes provided by WLCG computing sites). The term "integration grid", conversely, refers to the partitioning of multi-dimensional phase space into hypercubes in the VEGAS adaptive sampling algorithm, which is used by default in Sherpa [82, 91].

involving for example the approximation of matrix element calculations using Machine Learning (ML) regression methods [94, 95].

Inefficiencies in Unweighted Event Generation

The complex workflow described above presents several challenges and opportunities for improvement. To start with, there are many sources of inefficiency in unweighted event generation, as discussed in the following.

Phase space sampling inefficiency. The algorithm used for phase space sampling is the most critical ingredient for efficient unweighted event generation. Some basic techniques, such as stratified sampling, which essentially consists in binning the phase space, and importance sampling, which is often implemented as a change of variables to parametrize the phase space, date back to more than 40 years ago [76]. Many algorithms, most notably VEGAS [81, 82] or MISER [96], are adaptive, i.e. recursive, in that their parameters are tuned iteratively as the shape of $f(\mathbf{x})$ is learnt by randomly drawing more and more phase space points. Adaptive multi-channel algorithms [97, 98] are often used to address the complex peaking structures of LHC processes, by defining the sampling function $g(\mathbf{x})$ as a weighted sum of functions, each of which essentially describes a different peak. Many generic sampling algorithms exist, including very simple ones like RAMBO [99], others derived from VEGAS such as BASES/SPRING [100, 101] or MINT [102], and cellular algorithms like FOAM [103]. Other sampling algorithms have been developed specifically for a given generator: examples include MadEvent [104] and VAMP [105], which are based on modified versions of VEGAS and are used in the MG5_aMC and WHIZARD [106] generators, respectively, as well as COMIX [107], which is used in Sherpa.

In general, the larger the dimensionality of the phase space, the lower the unweighting efficiency that can be achieved: in W+jets at LO, for instance, the Sherpa efficiency [108] is 30% for W + 0 jets and 0.1% for W + 4 jets. This is an area where research is very active, and should be actively encouraged, as significant cost reductions in WLCG compute budgets could be achieved. Improvements in this area can only start from physics-motivated approaches based on the knowledge of phase space peaks, but they can be complemented by brute-force ML algorithmic methods [108–113], therefore people with different profiles can contribute to this area. The use of one of these ML tools, Generative Adversarial Networks (GAN), is being investigated [114] not only as a way to provide a more efficient phase space sampling, but also as a possible replacement for unweighted event generation altogether, for example when complemented with maximum mean discrepancy methods [115].

In this context, it is useful to point out that maximizing the efficiency of unweighted event generation and minimizing the variance on total cross section predictions by MC integration represent two different, even if closely related, strategies for the optimization of the phase space sampling algorithm. The two strategies imply the use of different loss metrics during the learning phase of an algorithm, and result in different weight distributions. This is discussed in detail in Ref. [103], and to some extent also in Ref. [105]. A completely different optimization strategy [116] for the sampling algorithm has also been recently proposed, where the goal is that of giving priority to populating the regions of phase space which are most sensitive to the presence of a signal or to the value of a parameter.

There are several reasons for the very large set of sampling codes. Many of them represent evolutions of VEGAS, others are completely different algorithms (like FOAM), and many of the modern ones are based on ML techniques. Some of these codes exist for historical reasons, because of the different choices adopted over time by each time. Possible work on a "common" integrator is sometimes mentioned in the community. Another possible way forward would consist in trying to harmonise the software interfaces of these packages, so that each generator could plug in different sampling algorithms and implementations. Discussions in this direction have already started in the context of the ongoing developments on GPU ports and on ML algorithms.

Slicing and biasing. A further issue [117], somewhat related to sampling inefficiencies, is that jet production cross sections fall very sharply as the transverse momenta (p_T) of the leading jets increase, and generating events with uniform weight generally fails to give a reasonable yield in the high-p_T regions of phase space. One approach to solving this problem ("slicing") is to produce several independent samples of events, using different generation cuts in each one, in order to populate all the regions of interest. An additional approach ("biasing" or "enhancement"), available for instance in POWHEG [117], MG5_aMC [118, 119], Sherpa [120], Pythia8 [121] and Herwig7.1 [42], consists in generating samples of events with non uniform weights, the shape of whose distribution can however be controlled by user-defined suppression factors. Both approaches are used in practice by the LHC experiments, as each has its pros and cons, and both reduce the resources required to populate the low-statistics tails of distributions. With additional work, these methods could help reduce the overall event generation resource requirements at HL-LHC.

Merging inefficiency. Merging prescriptions (e.g. MLM, CKKW-L at LO, and FxFx, MEPS@NLO at NLO) imply the rejection of some events to avoid double counting, between events produced with n+1 jets in the matrix element, and events produced with n jets in the matrix element and one jet from the parton shower [56]. This is only needed

if the required final state includes a variable number of jets n_{jets} between 0 and n, i.e. for so-called "merged" or "multi-leg" setups. The resulting inefficiencies can be relatively low depending on the process, but they are unavoidable in the algorithmic strategy used by the underlying physics modeling. The merging efficiency of the MLM prescription, for instance, is discussed in in Ref. [122], which shows how this can be improved using a method like shower-k_T MLM.

Filtering inefficiency. An additional large source of inefficiency is due to the way the experiments simulate some processes, where they generate large inclusive event samples, which are then filtered on final-state criteria to decide which events are passed on to detector simulation and reconstruction (e.g. CMS simulations of specific Λ_B decays have a 0.01% efficiency, and ATLAS B-hadron filtering in a V+jets sample has ~10% efficiency [123]). This inefficiency could be reduced by developing filtering tools within the generators themselves, designed for compatibility with the requirements of the experiments. A particularly wasteful example is where events are separated into orthogonal subsamples by filtering, in which case the same large inclusive sample is generated many times, once for each filtering stream: allowing a single inclusive event generation to be filtered into several orthogonal output streams would improve efficiency. Filtering is an area where the LHCb collaboration has a lot of experience [83] and has already obtained significant speedups through various techniques. In this context, one should also note that the speed of color reconnection algorithms [124, 125] is a limiting factor for simulating rare hadron decays in LHCb.

Sample sharing. In addition to removing inefficiencies, other ways could be explored to make maximal use of the CPU spent for generation by reusing samples for more than one purpose. Sharing parton-level, or even particle-level, samples between ATLAS and CMS is being discussed for some physics analyses. However, the implications of the statistical correlations that this would introduce need further investigation in the context of combinations of results across experiments.

Sample reweighting. Another way to re-use samples is through event reweighting. Recently, there have been major improvements in available tools in this area [30, 126–130], which have made it possible to obtain systematic uncertainty variations as well as reweighting to alternative model parameters. The latter may be useful for example in new physics searches, but also in the optimization of experimental measurements of model parameters [131]. This machinery is particularly important because in the past obtaining these variations would have required multiple samples to go through detector simulation and reconstruction, whereas the reweighting only requires this overhead for a single sample that can then be reused in multiple ways. This significantly reduces the CPU and storage requirements for the

same end result. However, this issue can still be explored further as in some areas there are limitations to the validity of these reweighting schemes [128–130, 132, 133]. In addition, some systematic uncertainty variations, such as merging scale variations, are not yet available as weights but there is work ongoing. There are also systematic variations such as changes of the hadronization model which are not well suited to the type of event reweighting discussed here, but for which alternative approaches using ML techniques to train an ad-hoc reweighting between samples are under investigation [134–138].

Negative weights. In NLO calculations, matching prescriptions (e.g. MC@NLO, POWHEG, etc.) are required to avoid double counting between phase space configurations that may come both from ℍ-events and from 𝕊-events with parton showers. The solution of this issue becomes technically even more complex at the NNLO. A widely used NLO matching prescription, MC@NLO [50], is implemented by using a "modified subtraction method" that may lead to the appearance of events with negative weights. A MC unweighting procedure is still applied, but the resulting events are "unweighted" in the sense that their weight can only be $+1$ or -1. This is a source of (possibly large) inefficiency, as larger event samples must be generated and passed through the experiment simulation, reconstruction and analysis software, increasing the compute and storage requirements. For a fraction r of events with weight -1, the number of events to generate increases by a factor $1/(1-2r)^2$, because the statistical error on MC predictions is a factor $1/(1-2r)$ higher; for a more detailed explanation of these formulas, see for instance Ref. [139]. For example, negative weight fractions equal to $r=25\%$ and $r=40\%$, which may be regarded as worst-case scenarios occurring in $t\bar{t}$ and $H b\bar{b}$ production [139], respectively, imply the need to generate 4 times and 25 times as many events to achieve the same statistical precision on MC predictions.

Negative weights can instead be almost completely avoided, by design, in another popular NLO matching prescription, POWHEG [51], which however is only available for a limited number of processes. POWHEG describes the relevant physics in a different way with respect to MC@NLO, so that predictions which have formally the same level of accuracy may visibly differ in the two codes, and are associated with different systematics (see Ref. [139] for an in-depth discussion). Negative weights can also be avoided in the KrkNLO [52] matching prescription, which is based on a very different approach from those used by MC@NLO and POWHEG; this method however is only available for a limited number of processes, and so far has been rarely used in practice by the LHC experiments.

Progress in this area at the fundamental physics level can only be achieved by theorists, and research is active in this area. For instance, a modified MC@NLO matching

procedure with reduced negative weights, known as MC@NLO-Δ, has recently been proposed [139]. Similarly, techniques to significantly reduce the negative weight fraction are also available in Sherpa [95]. Negative weights also exist for NNLO calculations, for instance in the UN2LOPS prescription [140].

It should be stressed that negative weights due to matching are absent in LO calculations. One possibility for avoiding negative weights, while possibly still achieving a precision beyond LO, could then be to generate LO multi-leg setups and reweight them to higher order predictions; a careful evaluation of the theoretical accuracy of this procedure would however be needed in this case. In addition, negative weights can also happen at LO because of not-positive-definite parton distribution function sets and interference terms, which is particularly relevant for effective field theory calculations.

One should also note that developments to incorporate contributions in parton shower algorithms beyond the currently adopted approximations, see e.g. Refs. [141–143], very often necessitate weighted evolution algorithms. In the future, this may represent another mechanism leading to the appearance of events with negative weights, in addition to and distinct from NLO matching prescriptions. Overcoming the prohibitively broad weight distributions is subject to an ongoing development and might necessitate structural changes in the event generation workflow. An example is the resampling approach proposed in Ref. [144], which also contains a useful historical review of sampling/importance resampling [145] techniques in the broader context of Monte Carlo simulations.

The term resampling has also been used to indicate the unrelated technique of positive resampling [146], which has recently been proposed as a new approach for reducing the impact of negative weights introduced by NLO matching prescriptions. The idea behind this method, which addresses negative weights as a statistics problem without looking at the underlying theoretical physics, is to remove negative weights using a quasi-local weight rebalancing scheme. While positive resampling uses histograms to determine bin-by-bin reweighting factors, a neural resampling [147] approach has later been proposed as an extension of this method, using neural networks to determine reweighting factors in the unbinned high-dimensional phase space.

Accounting of Compute Budgets for Generators

While progress has been made in the HSF generator WG to better understand which areas of generator software have the highest computational cost, more detailed accounting of the experiment workloads and profiling of the main generator software packages would help to further refine R&D priorities.

Accounting of CPU budgets in ATLAS/CMS. Thanks to a large effort from the generator teams in both experiments, a lot of insight into the settings used to support each experiment's physics programme was gained within the WG. It is now clear that the fraction of CPU that ATLAS spends for event generation is somewhat higher than that in CMS, although the difference is lower than previously thought: the latest preliminary estimates of these numbers are 12% and 5%, respectively. A more detailed study of the different strategies is ongoing, in particular by analysing individually the CPU costs of the main processes simulated (notably, V+jets, $t\bar{t}$, diboson and multijet). This effort aims at providing these figures also as absolute numbers in normalized HEP-SPEC06 (HS06) seconds [148, 149], to allow a more meaningful comparison of the compute budgets for event generation in ATLAS and CMS.

A practical issue is that these figures had to be harvested from logs and production system databases a posteriori. Deriving precise numbers for CMS has been particularly difficult, requiring significant person hours to extract the required information, as until recently the event generation (GEN) and detector simulation (SIM) steps were mixed in a single software application, and no separate accounting figures for GEN and SIM could be recovered from past job logs, therefore Grid costs had to be extrapolated from ad-hoc local tests. CMS is now also using workflows including GEN-only applications, like that used in ATLAS, which makes GEN CPU accounting easier. In addition, job monitoring information in CMS is presently kept for only 18 months, which complicates the analysis of past productions, and does not always include reliable HEP-SPEC06 metrics. For the future, it would be important to establish better mechanisms to collect all this information, to allow for an easy comparison between different experiments. It would also help if the various types of efficiencies described above (sampling, merging and filtering) could be more easily retrieved for all simulated processes.

Profiling of generators using production setups. Another area where the WG has been active, but more work is needed, is the definition and profiling of standard generator setups, reproducing those used in production. This has been used to compare the speeds of Sherpa and MG5_aMC in the configurations used by ATLAS and CMS, respectively. For instance, Sherpa was found to be 3 to 8 times slower than MG5_aMC in the generation of NLO $W + (0 - 2)$ jets, but the exact ratio depends on some of the model parameters used in Sherpa, e.g. the dynamical scale choice of Sherpa, which results in taking about 50% of the total CPU time for generation: when modifying Sherpa to use an equivalent scale to MG5_aMC, the CPU consumption for this process was reduced by over a factor of two. The choice of a scale, however, has important consequences not only on computational costs, but also on physics accuracy: an in-depth

discussion of this important issue, which has been described in many research papers by different teams of theorists (see, for instance, Refs. [53, 54, 139, 150, 151]), is beyond the scope of this paper, but the WG will continue to investigate the computing and physics implications of such choices.

Detailed profiling of different generator setups has also already helped to assess the CPU cost of external PDF libraries, and to optimise their use [152, 153]. The profiling of the memory footprint of the software would also be very useful, and may motivate in some cases a move to multithreading or multiprocessing approaches.

Modernisation of Generator Software

More generally, as is the case for many software packages in other areas of HEP, some R&D on generators would certainly be needed to modernise the software and make it more efficient, or even port it to more modern computing architectures (see also the discussion of these issues in the Snowmass 2013 report [154] and in the HSF CWP [14]).

Data parallelism, GPUs and vectorization. The data flow of an MC generator, where the same function $f(\mathbf{x})$, corresponding to the matrix element for the simulated HEP process, has to be computed over and over again at many phase space points $\mathbf{x}_i$ (i.e. for many different events), should, in principle, lend itself naturally to the data parallel approaches found in GPU compute kernels, and possibly to some extent in CPU vectorized code. In other words, event-level parallelism looks like an appropriate approach to try and exploit efficiently these architectures [155]. In this respect, generators should be somewhat easier to reengineer efficiently for GPUs than detector simulation software (notably Geant4 [156]), where the abundance of conditional branching of a stochastic nature may lead to "thread divergence" and poor software performance (see, for examples, Refs. [157–162]).

Porting and optimizing generators on GPUs is especially important to be able to exploit modern GPU-based HPCs (such as SUMMIT [163], where 95% of the compute capacity comes from GPUs [164]). Some work in this direction was done in the past on MG5_aMC, including both a port to GPUs (HEGET [165–167]) of the library that was used in MG5_aMC, before ALOHA [168] was introduced, for the automatic generation of matrix element code (HELAS [169, 170]), and a port to GPUs of VEGAS and BASES (gVEGAS and gBASES [171, 172]). This effort, which unfortunately never reached production quality, is now being revamped by the WG [155, 173], in collaboration with the MG5_aMC team, and represents one of the main R&D priorities of the WG. This work is presently focusing on Nvidia CUDA, but abstraction libraries like Alpaka [174, 175] or oneAPI [176] will also be investigated.

GPUs may also be relevant to the ML-based phase space sampling algorithms discussed in section "Inefficiencies in Unweighted Event Generation"; some recent work in this area has targeted GPUs explicitly [177, 178]. Similar techniques involving ML techniques on GPUs have recently been used for the computation of parton distribution functions [179], which are another essential building block of the event generator software chain for LHC processes. Finally, work is also ongoing [180] on the efficient exploitation of GPUs in the pseudo-random number generation libraries that are used in all MC generators (see Ref. [77] for a recent review of these components).

Task parallelism, multithreading, multiprocessing. The use of concurrency mechanisms based on multiprocessing (MP) or multithreading (MT) within event generators has increased in recent years, but it is not yet a common practice. Most often, the use of single-threaded (ST), single-process (SP), executables is not a problem, as the memory footprint of event generation is small and usually fits within the 2 GB per core available on WLCG nodes, which makes it possible to exploit all of the available cores by running many independent applications in parallel. However, there are cases (e.g. diboson production, or Z and $Z\gamma$+jets production with electroweak corrections, all with up to 4 additional jets) where more than 2 GB of memory, and even as much as 4 GB, may be needed; this leads to inefficiencies as some CPU cores remain unused, which could be avoided using MT or MP approaches to reduce memory footprints. Very often, the experiments do not use generators as standalone applications, but instead embed them in their own event processing frameworks, which may themselves implement MP or MT approaches. For MT frameworks, the fact that some generator packages (such as EvtGen [181]) are not thread safe may also lead to inefficiencies, as this often implies that access to these components must be locked and their methods can only be used by one thread at a time. Many different use cases and approaches exist in the various experiments, as described more in detail below.

In ATLAS, the most commonly used workflow for event generation currently consists in executing several independent ST/SP applications based on the experiment's event processing framework, Athena [182], each running as an independent Grid job on a different CPU core. Less frequently, ATLAS also uses multi-core jobs, using either ad-hoc features in Athena or, in the specific case of MG5_aMC, the internal concurrency mechanism provided by this generator; neither of these options, however, leads to an overall reduction in memory footprint, as they both ultimately consist in spawning several independent ST applications on the available cores. The ATLAS event processing framework, Athena [182], does have a MP extension based on forking and copy-on-write, AthenaMP [183, 184], which is routinely used to reduce per-core memory footprints of the ATLAS simulation and reconstruction workflows, but currently this is not used for any GEN workflows. ATLAS is also making progress in

the development of a fully multi-threaded event processing framework, AthenaMT [185], but this effort is also mainly focusing on simulation [186] and reconstruction workflows rather than on event generation. In particular, as in the similar MT developments in LHCb and CMS, described below, one of the main aims of this work is the integration into the experiment's workflows of the recent multi-threaded version of the Geant4 detector simulation toolkit, Geant4-MT [187].

In LHCb, event generation currently proceeds only via ST/SP Grid jobs. A notable difference with respect to the ATLAS GEN-only jobs is that LHCb uses a GEN-SIM workflow where the same application, Gauss [188], based on the Gaudi [189, 190] event processing framework, performs both the event generation and detector simulation steps. To improve the efficiency of these workflows, LHCb is gradually moving away from ST/SP Gauss. An MP framework using forking and copy-on-write as in AthenaMP, GaussMP, was recently used for some MC productions, but this was only a temporary ad-hoc solution for the low-memory many-core Intel Knights Landing (KNL) CPUs deployed on the Marconi HPC system at CINECA [191]. In the future, LHCb plans to replace its current SP/ST application by a fully multi-threaded version of Gauss based on an experiment-independent GEN-SIM framework, Gaussino [192], which is built directly on Gaudi and is interfaced to Geant4-MT. Gaussino, whose development is making rapid progress, has the potential for a much better optimization of memory usage, especially in the SIM step. To achieve thread safety, Gaussino uses a single-threaded locking instance of Evt-Gen to handle decays; special care is also taken in the way Gaussino is interfaced to Pythia8, as described in Ref. [192].

In CMS, event generation is embedded within the multi-threaded C++ event processing framework, CMSSW [193–195]. All CMS workflows for event generation involve multi-core Grid jobs (either GEN-SIM or, more recently, GEN-only), where a single instance of the CMSSW application simultaneously uses several CPU cores. One of the primary motivations for CMS to use a MT framework is to reduce the amount of memory used per CPU core [194], a goal that is particularly important, and has been achieved in production (also thanks to the use of the new GEANT4-MT), for GEN-SIM workflows [196]. Some examples of the integration of event generators in the CMS worflow, and more particularly in the MT CMSSW framework, are described in Refs. [197, 198]. The simplest use case is that where a general-purpose generator (like Pythia8, Herwig7, or Sherpa) is used both for the generation of parton-level events and for their hadronization and decay. In this case, the CMSSW main thread starts many separate worker processes on the available CPU cores, ensuring that each worker receives the appropriate random numbers to process the event assigned to it (a mechanism which has some similarities to that used in GaussMP by LHCb). A second important use case is that

where event generation is split into two steps, the generation of parton-level events in the Les Houches Event File (LHEF) format [199, 200] using a matrix-element generator (like MG5_aMC or POWHEG), and the hadronization and decay to stable particles of those parton-level events using other tools (like Pythia8, Herwig7, EvtGen or Tauola). To execute the LHEF event generation step concurrently on the available cores, CMSSW provides a generic mechanism where several independent processes are spawned on the available cores; in the case of MG5_aMC, its internal concurrency mode may also be used, but this also consists in spawning several ST applications, as already mentioned. The concurrent execution of the hadronization and decay step, conversely, is always handled by CMSSW using its internal multi-threading: the main challenge in this context is that decayers like EvtGen and Tauola are not thread-safe and may only be used to process one event at a time [198].

Multiprocessing approaches involving several nodes may also be useful to speed up the integration and optimization step for complex high-dimensional final states. In particular, Sherpa workflows based on the Message Passing Interface (MPI) [201], which have been available for quite a long time, have been found very useful by ATLAS and CMS to speed up the preparation of integration grids on local batch clusters. A lot of work has also been done in recent years to implement and benchmark MPI-based workflows on HPC systems. For instance, the Sherpa LO-based generation of merged many-jet samples has been successfully tested [202] on the Cori [203] system at NERSC, both on traditional Intel Haswell CPUs and on many-core Intel KNL CPUs. This work has used a technique similar to that previously developed [204] for testing and benchmarking the scaling of the parallel execution of Alpgen on Mira [205] at ALCF, a supercomputer based on IBM PowerPC CPUs. New event formats, migrating LHEF to HDF5, have also been instrumental in enabling the execution of the Sherpa tests at Cori. MPI integration has also been completed for MG5_aMC [206]. In this context, one should note that, although HPCs offer very high-speed inter-node connectivity, HPC resources can be exploited efficiently even without using this connectivity: in particular, WLCG workflows, including generators, generally use these systems as clusters of unrelated nodes, because the computational workflow can be split up into independent tasks on those nodes.

Hybrid parallelization approaches are also possible, where multithreading or multiprocessing techniques are used internally on a single multi-core node, while the MPI protocol is used to manage the communication between distinct computing nodes. This approach is implemented for example in the WHIZARD [105] and MCFM [207] codes, both of which combine OpenMP [208] multithreading on individual multi-core nodes with MPI message passing between them.

Generic code optimizations. A speedup of generators may also be achievable by more generic optimizations, not involving concurrency. For instance, one could study if different compilers and build strategies [206] may lead to any improvements. Another possible strategy is to search for redundant calculations, i.e. to investigate if some numerical results can be reused more than once, for instance via data caching. Recent studies [152] on the way LHAPDF6 is used in Pythia8 have indeed resulted in significant speedups through better data caching, and similar studies are in progress for Sherpa [153]. Ongoing studies [209] on MG5_aMC have similarly shown that important speedups may obtained through "helicity recycling", i.e. by avoiding the recomputation of some building blocks of Feynman diagrams which are needed in more than one matrix element calculation.

Physics Challenges (Increasing Precision)

In addition to software issues, important physics questions should also be addressed about more accurate theoretical predictions, above all NNLO QCD calculations, but also electroweak (EW) corrections, and their potential impact on the computational cost of event generators at HL-LHC. For a recent review of these issues, see for example Ref. [25]. Some specific NNLO calculations are already available and used today by the LHC experiments in their data analysis. For example, the measurements of fiducial $t\bar{t}$ cross sections, extrapolated to the full phase space, are compared to the predictions of TOP++ [80], accurate to NNLO: this program, however, does not use MC methods and cannot be used to generate unweighted events. Research on NNLO matching has also made significant progress, for example on the NNLOPS [210], GENEVA [211], UN2LOPS [140] and MINNLOPS [212, 213] prescriptions. In addition, samples of unweighted events are routinely generated for Higgs boson final states using the POWHEG/MINLO NNLOPS approach [210, 214]. With a view to HL-LHC times, however, some open questions remain to be answered, as discussed below.

NNLO: status of theoretical physics research. The first question is for which processes QCD NNLO precision will be available at the time of the HL-LHC. For example, first results for triphoton results at NNLO have recently been published [215]: when would NNLO be expected for other $2\rightarrow3$ processes or even higher multiplicity final states? Also, for final states such as $t\bar{t}$, differential NNLO predictions exist [216, 217] and the first matched computation for NNLO+PS was very recently achieved [213], but the software for generating unweighted NNLO+PS events using the latter is not yet publicly available for use in the experiments, when can this be expected? In particular, it would be important to clarify which are the theoretical

and more practical challenges in these calculations, and the corresponding computational strategies and predicted impact on CPU time needs (e.g. more complex definition of matching procedures, higher fraction of negative weights, and more complex 2-loop MEs?).

The accuracy of shower generators is also important in this context. Current shower generators rely on first order splitting kernels, together with an appropriate scheme to handle soft emissions. Recent work aims at improving parton showers by increasing their accuracy either by developing novel shower schemes within the standard parton or dipole branching, such as DIRE [218] and Vincia [219] or by going beyond the typical probabilistic approach [220] and by incorporating higher order splitting functions [52, 141, 221, 222]. In addition, very recently, significant theoretical advance opening the way to showers with next-to-leading logarithmic (NLL) accuracy has been achieved [223].

To match NNLO accuracy in QCD, EW corrections must also be included. Recently, much progress has been achieved on the automation of the computation of EW corrections [224–227], to the point that fixed-order NLO QCD and EW corrections are readily available for any process of interest at the LHC. A general interface of these calculations to shower generators that correctly account for QED radiation for these computations, however, is not yet available.

An additional concern, in general but especially in higher-order phenomenology, is the control of numerical and methodological errors at the sub-percent level. This is relevant for processes where high-precision measurements and predictions are available, but also to efficiently and precisely test the input parameter dependence (PDFs, α_s, etc.). These issues, and the way in which they are addressed in the MCFM parton-level code, are discussed in detail in Ref. [207]. A key component of this code is a fully parallelized phase space integration, using both OpenMP and MPI on multi-core machines and cluster setups, where technical cutoffs can be controlled at the required level of precision.

NNLO: experimental requirements at HL-LHC. The second question is for which final states unweighted event generation with NNLO precision would actually be required ($t\bar{t}$ production is a clear candidate), and how many events would be needed. One should also ask if reweighting LO event samples to NNLO would not be an acceptable cheaper alternative to address the experimental needs, and what would be the theoretical accuracy reached by this procedure.

Size of unweighted event samples required at HL-LHC. Another question to be asked, unrelated to NNLO, is in which regions of phase space the number of unweighted events must be strictly proportional to the luminosity. For example, in the bulk (low p_T) regions of W boson production it is probably impossible to keep up with the data, due to the huge cross section. Alternative techniques could be

investigated, to avoid the generation of huge samples of unweighted events.

Conclusions

This paper has been prepared by the HSF Physics Event Generator Working Group as an input to the LHCC review of HL-LHC computing, which has started in May 2020. We have reviewed the main software and computing challenges for the Monte Carlo physics event generators used by the LHC experiments, in view of the HL-LHC physics programme.

Out of the many issues that we have described, we have identified the following five as the main priorities on which the R&D on the computational aspects of generators, and in particular the activities of our WG, should focus:

1. Gain a more detailed understanding of the current CPU costs by accounting and profiling.
2. Survey generator codes to understand the best way to move to GPUs and vectorized code, and prototype the port of the software to GPUs using data-parallel paradigms.
3. Support efforts to optimize phase space sampling and integration algorithms, including the use of Machine Learning techniques such as neural networks.
4. Promote research on how to reduce the cost associated with negative weight events, using new theoretical or experimental approaches.
5. Promote collaboration, training, funding and career opportunities in the generator area.

Additional material about these and the other issues described in this paper, including detailed plots and diagrams, may be found in the recent presentation by the HSF Generator WG to the LHCC [155]. We plan to report about the progress in these areas during the more in-depth LHCC review of HL-LHC software, which is currently scheduled in Q3 2021, and reassess the WG priorities for future activities at that point in time.

References

1. ATLAS Collaboration (2008) The ATLAS experiment at the CERN Large Hadron Collider. J Instrum 3:S08003. https://doi.org/10.1088/1748-0221/3/08/S08003
2. CMS Collaboration (2008) The CMS experiment at the CERN LHC. J Instrum 3:S08004. https://doi.org/10.1088/1748-0221/3/08/S08004
3. High-Luminosity Large Hadron Collider (HL-LHC). https://home.cern/science/accelerators/high-luminosity-lhc
4. HSF Physics Event Generator Working Group. https://hepsoftwarefoundation.org/workinggroups/generators.html
5. LHC Experiments Committee. https://committees.web.cern.ch/lhcc
6. A. Boehnlein et al. (2020) HL-LHC Software and Computing Review Panel, 1st Report, CERN-LHCC-2020-012. https://cds.cern.ch/record/2725487
7. Minutes of the 136th Meeting of the LHCC, CERN-LHCC-2018-033. https://cds.cern.ch/record/2649242
8. Minutes of the 139th Meeting of the LHCC, CERN-LHCC-2019-010. https://cds.cern.ch/record/2689443
9. Minutes of the 140th Meeting of the LHCC, CERN-LHCC-2019-016. https://cds.cern.ch/record/2702745
10. Minutes of the 141st Meeting of the LHCC, CERN-LHCC-2020-003. https://cds.cern.ch/record/2711432
11. ALICE Collaboration (2008) The ALICE experiment at the CERN LHC. J Instrum 3:S08002. https://doi.org/10.1088/1748-0221/3/08/S08002
12. LHCb Collaboration (2008) The LHCb detector at the LHC. J Instrum 3:S08005. https://doi.org/10.1088/1748-0221/3/08/S08005
13. HEP Software Foundation (2020) Common Tools and Community Software, input for the LHCC review of HL-LHC computing, HSF-DOC-2020-1. https://doi.org/10.5281/zenodo.3779249
14. HEP Software Foundation (2019) A roadmap for HEP software and computing R&D for the 2020s. Comput Softw Big Sci 3:7. https://doi.org/10.1007/s41781-018-0018-8
15. Buckley A, et al. (2019) Monte Carlo event generators for high energy particle physics event simulation, MCnet-19-02. arxiv:1902.01674

16. CERN Council Open Symposium on the Update of European Strategy for Particle Physics (2019) Granada. https://cafpe.ugr.es/eppsu2019

17. Worldwide LHC Computing Grid. https://wlcg.web.cern.ch

18. ATLAS Collaboration (2020) Combined measurements of Higgs boson production and decay using up to 80 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment. Phys Rev D 101:12002. https://doi.org/10.1103/PhysRevD.101.012002

19. CMS Collaboration (2018) Observation of Higgs boson decay to bottom quarks. Phys Rev Lett 121:121801. https://doi.org/10.1103/PhysRevLett.121.121801

20. Azzi P, et al. (2018) Report from Working Group 1: Standard Model Physics at the HL-LHC and HE-LHC, CERN-LPCC-2018-03, HL/HE-LHC Workshop, CERN. arxiv:1902.04070

21. Cepeda M, et al. (2018) Report from Working Group 2: Higgs physics at the HL-LHC and HE-LHC, CERN-LPCC-2018-04, HL/HE-LHC Workshop, CERN. arxiv:1902.00134

22. Cid Vidal X, et al. (2018) Report from Working Group 3: Beyond the Standard Model Physics at the HL-LHC and HE-LHC, CERN-LPCC-2018-05, HL/HE-LHC Workshop, CERN. arxiv:1812.07831

23. Cerri A, et al. (2018) Report from Working Group 4: Opportunities in Flavour Physics at the HL-LHC and HE-LHC, CERN-LPCC-2018-06, HL/HE-LHC Workshop, CERN. arxiv:1812.07638

24. Boughezal R, et al. (2017) Generator and Theory Working Group Chapter for CWP, unpublished draft. https://github.com/HSF/documents/tree/master/CWP/papers/HSF-CWP-2017-11_generators

25. Maltoni F, Schönherr M, Nason P. Monte Carlo generators, in Ref. [20]

26. HSF Physics Event Generator Computing Workshop (2018). CERN. https://indico.cern.ch/event/751693

27. Buckley A (2020) Computational challenges for MC event generation. In: Proc. ACAT2019, Saas Fee. J Phys Conf Ser 1525:012023. https://doi.org/10.1088/1742-6596/1525/1/012023

28. Sexton-Kennedy L, Stewart G. CWP challenges and workshop aims, in Ref. [26]. https://indico.cern.ch/event/751693/contributions/3182927

29. Bothmann E, et al. (2019) Event generation with Sherpa 2.2. SciPost Phys 7:34. https://doi.org/10.21468/SciPostPhys.7.3.034

30. Alwall J, et al. (2014) The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, JHEP07(2014)079. https://doi.org/10.1007/JHEP07(2014)079

31. McFayden J. ATLAS needs and concerns, in Ref. [26]. https://indico.cern.ch/event/751693/contributions/3182932

32. Yazgan E. CMS needs and concerns, in Ref. [26]. https://indico.cern.ch/event/751693/contributions/3182936

33. ATLAS and CMS Collaborations (2019) Comparison of ATLAS and CMS nominal $t\bar{t}$ Monte Carlo simulation for Run 2, CMS-DP-2019-011, CERN. https://cds.cern.ch/record/2678959

34. Valassi A (2019) Overview of the GPU efforts for WLCG production workloads, Pre-GDB on benchmarking, CERN. https://indico.cern.ch/event/739897/contributions/3559134

35. Buckley A et al. (2011) General-purpose event generators for LHC physics. Phys Rep 504:145. https://doi.org/10.1016/j.physrep.2011.03.005

36. Sjöstrand T (2012) Introduction to Monte Carlo techniques in High Energy Physics, CERN Summer Student Lectures. https://indico.cern.ch/event/190076

37. Sjöstrand T (2016) Status and developments of event generators, Fourth LHC Physics Conference (LHCP2016), Lund. arxiv:1608.06425

38. Alioli S, et al. (2010) A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX, JHEP06(2010)043. https://doi.org/10.1007/JHEP06(2010)043

39. Sjöstrand T, et al. (2015) An introduction to PYTHIA 8.2. Comput Phys Comm 191:159. https://doi.org/10.1016/j.cpc.2015.01.024

40. Bähr M et al. (2008) Herwig++ physics and manual. Eur Phys J C 58:639. https://doi.org/10.1140/epjc/s10052-008-0798-9

41. Bellm J et al. (2016) HERWIG 7.0/HERWIG++ 3.0 release-note, Eur Phys J C 76:196. https://doi.org/10.1140/epjc/s10052-016-4018-8

42. Bellm J et al. (2017) Herwig 7.1 Release Note CERN-PH-TH-2017-109. arxiv:1705.06919

43. Mangano ML, et al. (2003) Alpgen, a generator for hard multiparton processes in hadronic collisions, JHEP07(2003)001. https://doi.org/10.1088/1126-6708/2003/07/001

44. Lönnblad L (1992) Ariadne version 4: a program for simulation of QCD cascades implementing the colour dipole model. Comp Phys Comm 71:15. https://doi.org/10.1016/0010-4655(92)90068-A

45. ATLAS Collaboration (2012) Summary of ATLAS Pythia 8 tunes, ATL-PHYS-PUB-2012-003. https://cds.cern.ch/record/1474107

46. CMS Collaboration (2020) Extraction and validation of a new set of CMS Pythia8 tunes from underlying-event measurements. Eur Phys J C 80:4. https://doi.org/10.1140/epjc/s10052-019-7499-4

47. Ju X, et al. A novel workflow of generator tunings in HPC for LHC new physics searches, in Ref. [26]. https://indico.cern.ch/event/751693/contributions/3183027

48. Bellm J, Gellersen L (2020) High dimensional parameter tuning for event generators. Eur Phys J C 80:54. https://doi.org/10.1140/epjc/s10052-019-7579-5

49. Frederix R, Frixione S (2012) Merging meets matching in MC@NLO. JHEP12(2012)061. https://doi.org/10.1007/JHEP12(2012)061

50. Frixione S, Webber BR (2002) Matching NLO. QCD computations and parton shower simulations. JHEP06(2002)029. https://doi.org/10.1088/1126-6708/2002/06/029

51. Frixione S, Nason P, Oleari C (2007) Matching NLO QCD computations with parton shower simulations: the POWHEG method. JHEP11(2007)070. https://doi.org/10.1088/1126-6708/2007/11/070

52. Jadach S, et al. (2015) Matching NLO QCD with parton shower in Monte Carlo scheme—the KrkNLO method. JHEP10(2015)052. https://doi.org/10.1007/JHEP10(2015)052

53. Catani S, Krauss F, Kuhn R, Webber BR (2001) QCD matrix elements + parton showers. JHEP11(2001)063. https://doi.org/10.1088/1126-6708/2001/11/063

54. Lönnblad L (2002) Correcting the colour-dipole cascade model with fixed order matrix elements. JHEP05(2002)046. https://doi.org/10.1088/1126-6708/2002/05/046

55. Mangano ML, Moretti M, Pittau R (2002) Multijet matrix elements and shower evolution in hadronic collisions: $Wb\bar{b}$ + n-jets as a case study. Nucl Phys B 632:343. https://doi.org/10.1016/S0550-3213(02)00249-3

56. Alwall J et al. (2008) Comparative study of various algorithms for the merging of parton showers and matrix elements in hadronic collisions. Eur Phys J C 53:473. https://doi.org/10.1140/epjc/s10052-007-0490-5

57. Höche S, Krauss F, Schönherr M (2014) Uncertainties in MEPS@NLO calculations of h+jets. Phys Rev D 90:014012. https://doi.org/10.1103/PhysRevD.90.014012

58. Hamilton K, Nason P, Zanderighi G (2012) MINLO: multi-scale improved NLO. JHEP10(2012)155. https://doi.org/10.1007/JHEP10(2012)155

59. Lönnblad L, Prestel S (2013) Merging multi-leg NLO matrix elements with parton showers. JHEP03(2013)166. https://doi.org/10.1007/JHEP03(2013)166

60. Plätzer S, Gieseke S (2012) Dipole showers and automated NLO matching in Herwig++. Eur Phys J C 72:2187. https://doi.org/10.1140/epjc/s10052-012-2187-7

61. Plätzer S (2013) Controlling inclusive cross sections in parton shower + matrix element merging, JHEP08(2013)114. https://doi.org/10.1007/JHEP08(2013)114

62. Bellm J, Gieseke S, Plätzer S (2018) Merging NLO multi-jet calculations with improved unitarization. Eur Phys J C 78:244. https://doi.org/10.1140/epjc/s10052-018-5723-2

63. Lange DJ (2001) The EvtGen particle decay simulation package. Nucl Instrum Meth A 462:152. https://doi.org/10.1016/S0168-9002(01)00089-4

64. Jadach S, Was Z, Decker R, Kühn JH (1993) The τ decay library TAUOLA, version 2.4. Comp Phys Comm 76:361. https://doi.org/10.1016/0010-4655(93)90061-G

65. Golonka P, Was Z (2006) PHOTOS Monte Carlo: a precision tool for QED corrections in Z and W decays. Eur Phys J C 45:97. https://doi.org/10.1140/epjc/s2005-02396-4

66. Botje M (2011) The PDF4LHC Working Group Interim Recommendations. arxiv:1101.0538v1

67. Buckley A et al. (2015) LHAPDF6: parton density access in the LHC precision era. Eur Phys J C 75:132. https://doi.org/10.1140/epjc/s10052-015-3318-8

68. Buckley A et al. (2013) Rivet user manual. Comp Phys Comm 184:2803. https://doi.org/10.1016/j.cpc.2013.05.021

69. Buckley A et al. (2010) Systematic event generator tuning for the LHC. Eur Phys J C 65:331. https://doi.org/10.1140/epjc/s10052-009-1196-7

70. Dobbs M, Hansen JB (2001) The HepMC C++ Monte Carlo event record for High Energy Physics. Comp Phys Comm 134:41. https://doi.org/10.1016/S0010-4655(00)00189-2

71. Buckley A et al. (2021) The HepMC3 event record library for Monte Carlo event generators. Comp Phys Comm 260:107310. https://doi.org/10.1016/j.cpc.2020.107310

72. Cacciari M, Salam GP, Soyez G (2012) FastJet user manual. Eur Phys J C 72:1896. https://doi.org/10.1140/epjc/s10052-012-1896-2

73. Roiser S, et al. (2015) The LHCb Distributed Computing Model and Operations during LHC Runs 1, 2 and 3, Proc. ISGC2015, Taipei. https://doi.org/10.22323/1.239.0005

74. Caravaglios F, Moretti M (1995) An algorithm to compute Born scattering amplitudes without Feynman graphs. Phys Lett B 358:332. https://doi.org/10.1016/0370-2693(95)00971-M

75. James F (1968) Monte Carlo phase space, CERN Yellow Report CERN-68-15. https://doi.org/10.5170/CERN-1968-015

76. James F (1980) Monte Carlo theory and practice. Rep Progr Phys 43:1145. https://doi.org/10.1088/0034-4885/43/9/002

77. James F, Moneta L (2020) Review of high-quality random number generators. Comput Softw Big Sci 4:2. https://doi.org/10.1007/s41781-019-0034-3

78. Weinzierl S (2000) Introduction to Monte Carlo methods, NIKHEF-00-012. arxiv:hep-ph/0006269

79. James F, Hoogland J, Kleiss R (1997) Multidimensional sampling for simulation and integration: measures, discrepancies, and quasi-random numbers. Comp Phys Comm 99:180. https://doi.org/10.1016/S0010-4655(96)00108-7

80. Czakon M, Mitov A (2014) Top++: A program for the calculation of the top-pair cross-section at hadron colliders. Comp Phys Comm 185:2930. https://doi.org/10.1016/j.cpc.2014.06.021

81. Lepage GP (1978) A new algorithm for adaptive multidimensional integration. J Comp Phys 27:192. https://doi.org/10.1016/0021-9991(78)90004-9

82. Lepage GP (1980) VEGAS: an adaptive multi-dimensional integration program, Cornell report CLNS-447. https://cds.cern.ch/record/123074

83. Ilten P. LHCb needs and concerns, in Ref. [26]. https://indico.cern.ch/event/751693/contributions/3182938

84. Davis A (2020) Fast Simulation in LHCb, Workshop in Efficient Computing for High Energy Physics (ECHEP), Edinburgh. https://indico.ph.ed.ac.uk/event/66/contributions/844

85. Collins J, Soper D, Sterman G (1989) Factorization of Hard Processes in QCD. Adv Ser Direct High Energy Phys 5:1 arxiv:hep-ph/0409313v1

86. Weinzierl S (2012) NLO Calculations, Monte Carlo School, Hamburg. https://indico.desy.de/indico/event/5064/session/8

87. Zaro M (2015) MadGraph5_aMC@NLO tutorial, QCD and event simulation for the LHC lectures, Pavia. https://cp3.irmp.ucl.ac.be/projects/madgraph/wiki/Pavia2015

88. Luisoni G (2017) An introduction to POWHEG, Dartmouth-UW Experimental/Theory discussion. https://indico.cern.ch/event/602457/contributions/2435408

89. Alwall J, et al. (2009) New Developments in MadGraph/MadEvent, Proc. SUSY08, Seoul. AIP Conf Proc 1078:84. https://doi.org/10.1063/1.3052056

90. MadGraph, Technical details for setting up and running the Grid Package. https://cp3.irmp.ucl.ac.be/projects/madgraph/wiki/GridDevelopment

91. Sherpa Integration, Sherpa 2.0.0 Manual (2013) https://sherpa.hepforge.org/doc/SHERPA-MC-2.0.0.html#Integration

92. Yazgan E (2018) Event Generators in CMS, CMS Heavy Flavor Tagging Workshop, Brussels. https://indico.cern.ch/event/695320/contributions/2850950

93. Lange D. Practical computing considerations, in Ref. [26]. https://indico.cern.ch/event/751693/contributions/3182940

94. Bishara F, Montull M (2019) (Machine) Learning amplitudes for faster event generation, DESY 19-232. arxiv:1912.11055

95. Danziger K (2020) Efficiency Improvements in Monte Carlo Algorithms for High-Multiplicity Processes, Master-Arbeit Thesis, TU Dresden, CERN-THESIS-2020-024. https://cds.cern.ch/record/2715727

96. Press WH, Farrar GR (1990) Recursive stratified sampling for multidimensional Monte Carlo Integration. Comput Phys 4:190. https://doi.org/10.1063/1.4822899

97. Kleiss R, Pittau R (1994) Weight optimization in multichannel Monte Carlo. Comput Phys Comm 83:141. https://doi.org/10.1016/0010-4655(94)90043-4

98. Ohl T (1999) Vegas revisited: adaptive Monte Carlo integration beyond factorization. Comput Phys Comm 120:13. https://doi.org/10.1016/S0010-4655(99)00209-X

99. Kleiss RH, Stirling WJ, Ellis SD (1986) A new Monte Carlo treatment of multiparticle phase space at high energies. Comput Phys Comm 40:359. https://doi.org/10.1016/0010-4655(86)90119-0

100. Kawabata S (1986) A new Monte Carlo event generator for high energy physics. Comput Phys Comm 41:127. https://doi.org/10.1016/0010-4655(86)90025-1

101. Kawabata S (1995) A new version of the multi-dimensional integration and event generation package BASES/SPRING. Comput Phys Comm 88:309. https://doi.org/10.1016/0010-4655(95)00028-E

102. Nason P (2007) MINT: a Computer Program for Adaptive Monte Carlo Integration and Generation of Unweighted Distributions. Bicocca-FT-07-13. arxiv:0709.2085

103. Jadach S (2003) Foam: a general-purpose cellular Monte Carlo event generator. Comput Phys Comm 152:55. https://doi.org/10.1016/S0010-4655(02)00755-5

104. Maltoni F, Stelzer T (2003) MadEvent: automatic event generation with MadGraph, JHEP02(2003)027. https://doi.org/10.1088/1126-6708/2003/02/027

105. Brass S, Kilian W, Reuter J (2019) Parallel adaptive Monte Carlo integration with the event Generator WHIZARD. Eur Phys J C 79:344. https://doi.org/10.1140/epjc/s10052-019-6840-2

106. Kilian W, Ohl T, Reuter J (2011) WHIZARD: simulating multi-particle processes at LHC and ILC. Eur Phys J C 71:1742. https://doi.org/10.1140/epjc/s10052-011-1742-y

107. Gleisberg T, Höche S (2008) Comix, a new matrix element generator, JHEP12(2008)039. https://doi.org/10.1088/1126-6708/2008/12/039

108. Gao C et al. (2020) Event generation with normalizing flows. Phys Rev D 101:076002. https://doi.org/10.1103/PhysRevD.101.076002

109. Gao C, et al. (2020) i-flow: High-dimensional Integration and Sampling with Normalizing Flows. Mach Learn 1:045023. https://doi.org/10.1088/2632-2153/abab62

110. Bendavid J (2017) Efficient Monte Carlo integration using boosted decision trees and generative deep neural networks. arxiv:1707.00028

111. Klimek MD, Perelstein M (2020) Neural network-based approach to phase space integration. SciPost Phys 9:53. https://doi.org/10.21468/SciPostPhys.9.4.053

112. Gleyzer S, Seyfert P, Schramm S (eds.) et al. (2019) Machine Learning in High Energy Physics Community White Paper. arxiv:1807.02876

113. Bothmann E, et al. (2020) Exploring phase space with Neural Importance Sampling. SciPost Phys 8:69. https://doi.org/10.21468/SciPostPhys.8.4.069

114. Generative Models session, ML4Jets2020 workshop, NYU. https://indico.cern.ch/event/809820/sessions/329213

115. Butter A, Plehn T, Winterhalder R (2019) How to GAN LHC events. SciPost Phys 7:075. https://doi.org/10.21468/SciPostPhys.7.6.075

116. Matchev KT, Shyamsundar P (2020) OASIS: optimal analysis-specific importance sampling for event generation. arxiv:2006.16972

117. Alioli S, et al. (2011) Jet pair production in POWHEG, JHEP04(2011)081. https://doi.org/10.1007/JHEP04(2011)081

118. MadGraph, Biasing the generation of unweighted partonic events at LO, https://cp3.irmp.ucl.ac.be/projects/madgraph/wiki/LOEventGenerationBias

119. Frederix et al. (2016) Heavy-quark mass effects in Higgs plus jets production. JHEP08(2016)006. https://doi.org/10.1007/JHEP08(2016)006

120. Sherpa Enhance_Function, Sherpa 2.0.0 Manual (2013). https://sherpa.hepforge.org/doc/SHERPA-MC-2.0.0.html#Enhance_005fFunction

121. Pythia8 Sample Main Programs, http://home.thep.lu.se/~torbjorn/pythia81html/SampleMainPrograms.html

122. Alwall J, de Visscher S, Maltoni F (2009) QCD radiation in the production of heavy colored particles at the LHC. JHEP02(2009)017. https://doi.org/10.1088/1126-6708/2009/02/017

123. ATLAS Collaboration (2017) ATLAS simulation of boson plus jets processes in Run 2, ATL-PHYS-PUB-2017-006. http://cds.cern.ch/record/2261937

124. Gieseke S, Röhr C, Siódmok A (2012) Colour reconnections in Herwig++. Eur Phys J C 72:2225. https://doi.org/10.1140/epjc/s10052-012-2225-5

125. Gieseke S, et al. (2018) Colour reconnection from soft gluon evolution. JHEP11(2018)149. https://doi.org/10.1007/JHEP11(2018)149

126. Gainer JS, et al. (2014) Exploring theory space with Monte Carlo reweighting. JHEP10(2014)78. https://doi.org/10.1007/JHEP10(2014)078

127. Mrenna S, Skands P (2016) Automated parton-shower variations in Pythia8. Phys Rev D 94:074005. https://doi.org/10.1103/PhysRevD.94.074005

128. Mattelaer O (2016) On the maximal use of Monte Carlo samples: re-weighting events at NLO accuracy. Eur Phys J C 76:674. https://doi.org/10.1140/epjc/s10052-016-4533-7

129. Bothmann E, Schönherr M, Schumann S (2016) Reweighting QCD matrix-element and parton-shower calculations. Eur Phys J C 76:590. https://doi.org/10.1140/epjc/s10052-016-4430-0

130. Bendavid J, et al. (2017) Les Houches 2017: Physics at TeV Colliders Standard Model Working Group Report. Proc Les Houches. arxiv:1803.07977

131. Valassi A (2020) Optimising HEP parameter fits via Monte Carlo weight derivative regression, Proc. CHEP2019, Adelaide. EPJ Web of Conf 245:6038. https://doi.org/10.1051/epjconf/202024506038

132. Bellm J et al. (2016) Reweighting parton showers. Phys Rev D 94:034028. https://doi.org/10.1103/PhysRevD.94.034028

133. Bellm J et al. (2016) Parton-shower uncertainties with Herwig 7: benchmarks at leading order. Eur Phys J C 76:665. https://doi.org/10.1140/epjc/s10052-016-4506-x

134. ATLAS collaboration (2020) Measurements of WH and ZH production in the $H \to b\bar{b}$ decay channel in pp collisions at 13 TeV with the ATLAS detector, ATLAS-CONF-2020-006. http://cdsweb.cern.ch/record/2714885

135. Cranmer K, Pavez J, Louppe G (2015) Approximating Likelihood Ratios with Calibrated Discriminative Classifiers. arxiv:1506.02169

136. Andreassen A, Nachman B (2020) Neural networks for full phase-space reweighting and parameter tuning. Phys Rev D 101:091901. https://doi.org/10.1103/PhysRevD.101.091901

137. Peyré G, Cuturi M (2019) Computational optimal transport foundations and trends in machine. Learning 11:355. arxiv:1803.00567

138. Rogozhnikov A (2015) Reweighting with Boosted Decision Trees. http://arogozhnikov.github.io/2015/10/09/gradient-boosted-reweighter.html

139. Frederix R, Frixione S, Prestel S. Torrielli P (2020) On the reduction of negative weights in MC@NLO-type matching procedures. JHEP07(2020)238. https://doi.org/10.1007/JHEP07(2020)238

140. Höche S, Li Y, Prestel S (2015) Drell-Yan lepton pair production at NNLO QCD with parton showers. Phys Rev D 91:074015. https://doi.org/10.1103/PhysRevD.91.074015

141. Höche S, Prestel S (2017) Triple collinear emissions in parton showers. Phys Rev D 96:074017. https://doi.org/10.1103/PhysRevD.96.074017

142. Plätzer S, Sjodahl M, Thorén J (2018) Color matrix element corrections for parton showers. JHEP11(2018)009. https://doi.org/10.1007/JHEP11(2018)009

143. Á. Martínez R, et al. (2018) Soft gluon evolution and non-global logarithms. JHEP05(2018)044. https://doi.org/10.1007/JHEP05(2018)044

144. Olsson J, Plätzer S, Sjodahl M (2020) Resampling algorithms for high energy physics simulations. Eur Phys J C 80:934. https://doi.org/10.1140/epjc/s10052-020-08500-y

145. Rubin DB (1987) A noniterative sampling/importance resampling alternative to the data augmentation algorithm for creating a few imputations when the fraction of missing information is modest: the SIR algorithm (comment on an article by Tanner and Wong). J Am Statist Assoc 82:543. https://doi.org/10.2307/2289460

146. Andersen JR, Gutschow C, Maier A, Prestel S (2020) A positive resampler for Monte Carlo Events with negative weights. Eur Phys JC 80:1007. https://doi.org/10.1140/epjc/s10052-020-08548-w

147. Nachman B, Thaler J (2020) A neural resampler for Monte Carlo reweighting with preserved uncertainties. Phys Rev D 102:076004. https://doi.org/10.1103/PhysRevD.102.076004

148. Michelotto M, et al. (2010) A comparison of HEP code with SPEC benchmarks on multi-core worker nodes, Proc. CHEP2009, Prague. J Phys Conf Ser 219:052009. https://doi.org/10.1088/1742-6596/219/5/052009

149. Valassi A, et al. (2020) Using HEP experiment workflows for the benchmarking and accounting of WLCG computing resources. Proc CHEP2019, Adelaide, EPJ Web of Conf 245:07035. https://doi.org/10.1051/epjconf/202024507035

150. Höche S, et al. (2012) A critical appraisal of NLO+PS matching methods. JHEP09(2012)049. https://doi.org/10.1007/JHEP09(2012)049

151. Höche S, et al. (2013) QCD matrix elements + parton showers: the NLO case. JHEP04(2013)027. https://doi.org/10.1007/JHEP04(2013)027

152. Konstantinov D (2020) Optimization of Pythia8, EP-SFT group meeting, CERN. https://indico.cern.ch/event/890670

153. Martin T (2020) Computational bottlenecks, ECHEP/ Excalibur Workshop. https://indico.cern.ch/event/928965/contributions/3933234

154. Bauer C, et al. (2013) Computing for perturbative QCD: a Snowmass White Paper SLAC-PUB-15740. arxiv:1309.3598

155. Valassi A, Yazgan E, McFayden J (2020) Monte Carlo generators challenges and strategy towards HL-LHC. WLCG Meeting with LHCC Referees. https://doi.org/10.5281/zenodo.4028834

156. Agostinelli S et al. (2003) Geant4—a simulation toolkit. Nucl Instr Meth A 506:250. https://doi.org/10.1016/S0168-9002(03)01368-8

157. Seiskari O, Kommeri J, Niemi T (2012) GPU in physics computation: case Geant4 navigation. arxiv:1209.5235

158. Murakami K, et al. (2013) Geant4 based simulation of radiation dosimetry in CUDA. Proc IEEE NSS/MIC, Seoul. https://doi.org/10.1109/NSSMIC.2013.6829452

159. Corti G, et al. (2016) HEP software community meeting on GeantV R&D Panel Report. https://hepsoftwarefoundation.org/assets/GeantVPanelReport20161107.pdf

160. Canal P (2019) Geant Exascale Pilot Project, Geant4 R&D Meeting, CERN. https://indico.cern.ch/event/809393/contributions/3441114

161. Gheata A (2019) Design, implementation and performance results of the GeantV prototype, Outcome of the GeantV prototype HSF meeting, CERN. https://indico.cern.ch/event/818702/contributions/3559124

162. Amadio G, et al. (2020) GeantV: results from the prototype of concurrent vector particle transport simulation in HEP. arxiv:2005.00949

163. Oak Ridge Leadership Computing Facility, Summit. https://www.olcf.ornl.gov/summit

164. Feldman M (2018) New GPU-Accelerated Supercomputers Change the Balance of Power on the TOP500, Top500 news. https://www.top500.org/news/new-gpu-accelerated-supercomputers-change-the-balance-of-power-on-the-top500

165. Hagiwara K et al. (2010) Fast calculation of HELAS amplitudes using graphics processing unit (GPU). Eur Phys J C 66:477. https://doi.org/10.1140/epjc/s10052-010-1276-8

166. Hagiwara K et al. (2010) Calculation of HELAS amplitudes for QCD processes using graphics processing unit (GPU). Eur Phys J C 70:513. https://doi.org/10.1140/epjc/s10052-010-1465-5

167. Hagiwara K et al. (2013) Fast computation of MadGraph amplitudes on graphics processing unit (GPU). Eur Phys J C 73:2608. https://doi.org/10.1140/epjc/s10052-013-2608-2

168. de Aquino P, Link W, Maltoni F, Mattelaer O, Stelzer T (2012) ALOHA: Automatic libraries of helicity amplitudes for Feynman diagram computations. Comput Phys Comm 183:2254. https://doi.org/10.1016/j.cpc.2012.05.004

169. Murayama H, Watanabe I, Hagiwara K (1992) HELAS: HELicity Amplitude Subroutines for Feynman Diagram Evaluations, KEK-Report 91-11. https://lib-extopc.kek.jp/preprints/PDF/1991/9124/9124011.pdf

170. Watanabe I, Murayama H, Hagiwara K (1992) Evaluating Cross Sections at TeV Energy Scale by HELAS, KEK preprint 92-39. https://lib-extopc.kek.jp/preprints/PDF/1992/9227/9227039.pdf

171. Kanzaki J (2011) Monte Carlo integration on GPU. Eur Phys J C 71:1559. https://doi.org/10.1140/epjc/s10052-011-1559-8

172. Kanzaki J (2011) Application of graphics processing unit (GPU) to software in elementary particle/high energy physics field. Proc Comput Sci 4:869. https://doi.org/10.1016/j.procs.2011.04.092

173. Roiser S (2020) Progress on porting MadGraph5_aMC@NLO to GPUs, HSF/WLCG Virtual Workshop. https://indico.cern.ch/event/941278/contributions/4101793

174. Zenker E, et al. (2016) Alpaka—An Abstraction Library for Parallel Kernel Acceleration, Proc. IEEE IPDPSW 2016, Chicago. https://doi.org/10.1109/IPDPSW.2016.50

175. Alpaka—Abstraction Library for Parallel Kernel Acceleration. https://github.com/alpaka-group/alpaka

176. Intel oneAPI Toolkits (Beta). https://software.intel.com/en-us/oneapi

177. Wu H-Z, Zhang J-J, Pang L-G, Wang Q (2019) ZMCintegral: a package for multi-dimensional Monte Carlo integration on multi-GPUs. Comput Phys Comm 248:106962. https://doi.org/10.1016/j.cpc.2019.106962

178. Carrazza S, Cruz-Martinez JM (2020) VegasFlow: accelerating Monte Carlo simulation across multiple hardware platforms. Comput Phys Comm 254:107376. https://doi.org/10.1016/j.cpc.2020.107376

179. Carrazza S, Cruz-Martinez JM, Rossi M (2020) PDFFlow: parton distribution functions on GPU. arxiv:2009.06635

180. Jun SY, et al. (2019) Vectorization of random number generation and reproducibility of concurrent particle transport simulation. Proc ACAT2019, Saas Fee. https://inspirehep.net/literature/1754423

181. Kreps M. EvtGen status and plans, in Ref. [26]. https://indico.cern.ch/event/751693/contributions/3182956

182. Calafiura P, et al. (2004) The Athena Control Framework in Production, New Developments and Lessons Learned. Proc CHEP2004, Interlaken. https://doi.org/10.5170/CERN-2005-002.456

183. Calafiura P, et al. (2015) Running ATLAS workloads within massively parallel distributed applications using Athena Multi-Process framework (AthenaMP). J Phys Conf Ser 664:072050. https://doi.org/10.1088/1742-6596/664/7/072050

184. Elmsheuser J, et al. (2019) ATLAS Grid Workflow Performance Optimization, Proc. CHEP2018, Sofia. EPJ Web of Conf 214:3021. https://doi.org/10.1051/epjconf/201921403021

185. Leggett C, et al. (2017) AthenaMT: upgrading the ATLAS software framework for the many-core world with multi-threading, Proc. CHEP2016, San Francisco. J Phys Conf Ser 898:042009. https://doi.org/10.1088/1742-6596/898/4/042009

186. Bandieramonte M, et al. (2020) Multi-threaded simulation for ATLAS: challenges and validation strategy, Proc. CHEP2019, Adelaide. EPJ Web of Conf 245:02001. https://doi.org/10.1051/epjconf/202024502001

187. Allison J et al. (2016) Recent developments in Geant4. Nucl Instr Meth A 835:186. https://doi.org/10.1016/j.nima.2016.06.125

188. Clemencic M, et al. (2011) The LHCb Simulation Application, Gauss: Design, Evolution and Experience, Proc. CHEP2010, Taipei. J Phys Conf Ser 331:032023. https://doi.org/10.1088/1742-6596/331/3/032023

189. Barrand G et al. (2001) GAUDI—A software architecture and framework for building HEP data processing applications. Comput Phys Comm 140:45. https://doi.org/10.1016/S0010-4655(01)00254-5

190. Clemencic M, et al. (2010) Recent developments in the LHCb software framework Gaudi, Proc. CHEP2009, Prague. J Phys Conf Ser 219:042006. https://doi.org/10.1088/1742-6596/219/4/042006

191. Stagni F, Valassi A, Romanovskiy V (2020) Integrating LHCb workflows on HPC resources: status and strategies, Proc. CHEP2019, Adelaide. EPJ Web of Conf 245:09002. https://doi.org/10.1051/epjconf/202024509002

192. Siddi BG, Müller D (2019) Gaussino—a Gaudi-Based Core Simulation Framework. Proc IEEE NSS/MIC 2019, Manchester. https://doi.org/10.1109/NSS/MIC42101.2019.9060074

193. Sexton-Kennedy E, Gartung P, Jones CD, Lange D (2015) Implementation of a Multi-threaded Framework for Large-scale Scientific Applications, Proc. ACAT2014, Prague. J Phys Conf Ser 608:012034. https://doi.org/10.1088/1742-6596/608/1/012034

194. Jones CD, et al. (2015) Using the CMS threaded framework in a production environment. J Phys Conf Ser 664:072026. https://doi.org/10.1088/1742-6596/664/7/072026

195. Jones CD (2017) CMS event processing multi-core efficiency status, Proc. CHEP2016, San Francisco. J Phys Conf Ser 898:042008. https://doi.org/10.1088/1742-6596/898/4/042008

196. Hildreth M, Ivanchenko VN, Lange DJ (2017) Upgrades for the CMS simulation, Proc. CHEP2016, San Francisco. J Phys Conf Ser 898:042040. https://doi.org/10.1088/1742-6596/898/4/042040

197. Bendavid J (2016) CMS experience with current generators, Argonne and Fermilab Workshop on Beyond Leading Order Calculations on HPCs, Fermilab. https://indico.cern.ch/event/557731/contributions/2309458

198. Li C (2021) The CMS Offline WorkBook: multithreading in generators, CMS Public Web. https://twiki.cern.ch/twiki/bin/view/CMSPublic/WorkBookGenMultithread

199. Boos E et al. (2001) Generic user process interface for event generators. Proc. Physics at TeV colliders Workshop, Les Houches arxiv:hep-ph/0109068

200. Alwall J et al. (2007) A standard format for Les Houches Event Files. Comput Phys Comm 176:300. https://doi.org/10.1016/j.cpc.2006.11.010

201. Dongarra JJ et al. (1996) A message passing standard for MPP and workstations. Comm ACM. https://doi.org/10.1145/233977.234000

202. Höche S, Prestel S, Schulz H (2019) Simulation of vector boson plus many final jets at the high luminosity LHC. Phys Rev D 100:0140124. https://doi.org/10.1103/PhysRevD.100.014024

203. NERSC, Cori. https://docs.nersc.gov/systems/cori

204. Childers JT et al. (2017) Adapting the serial Alpgen parton-interaction generator to simulate LHC collisions on millions of parallel threads. Comp Phys Comm 210:54. https://doi.org/10.1016/j.cpc.2016.09.013

205. Argonne Leadership Computing Facility, Mira. https://www.alcf.anl.gov/alcf-resources/mira

206. Mattelaer O. MG5aMC status and plans, in Ref. [26]. https://indico.cern.ch/event/751693/contributions/3182951

207. Campbell J, Neumann T (2019) Precision phenomenology with MCFM, JHEP12(2019)034. https://doi.org/10.1007/JHEP12(2019)034

208. The OpenMP API specification for parallel programming, https://www.openmp.org

209. Mattelaer O, Ostrolenk K (2021) Speeding up MadGraph5_aMC@NLO, MCNET-21-01. arxiv:2102.00773

210. Hamilton K, et al. (2013) Merging H/W/Z + 0 and 1 jet at NLO with no merging scale: a path to parton shower + NNLO matching. JHEP05(2013)082. https://doi.org/10.1007/JHEP05(2013)082

211. Alioli S, et al. (2014) Matching fully differential NNLO calculations and parton showers. JHEP06(2014)089. https://doi.org/10.1007/JHEP06(2014)089

212. Monni PF, et al. (2020) MiNNLOPS: A new method to match NNLO QCD to parton showers. JHEP05(2020)143. https://doi.org/10.1007/JHEP05(2020)143

213. Mazzitelli J, et al. (2020) Next-to-next-to-leading order event generation for top-quark pair production, CERN-TH-2020-219. arxiv:2012.14267

214. Hamilton K, et al. (2013) NNLOPS simulation of Higgs boson production. JHEP10(2013)222. https://doi.org/10.1007/JHEP10(2013)222

215. Chawdhry HA, et al. (2020) NNLO QCD corrections to three-photon production at the LHC. JHEP02(2020)57. https://doi.org/10.1007/JHEP02(2020)057

216. Czakon M, Fiedler P, Heymes D, Mitov A (2016) NNLO QCD predictions for fully-differential top-quark pair production at the Tevatron. JHEP05(2016)034. https://doi.org/10.1007/JHEP05(2016)034

217. Catani S, et al. (2019) Top-quark pair production at the LHC: fully differential QCD predictions at NNLO. JHEP07(2019)100. https://doi.org/10.1007/JHEP07(2019)100

218. Höche S, Prestel S (2015) The midpoint between dipole and parton showers. Eur Phys J C 75:461. https://doi.org/10.1140/epjc/s10052-015-3684-2

219. Fischer N, Prestel S, Ritzmann M, Skands P (2016) VINCIA for hadron colliders. Eur Phys J C 76:589. https://doi.org/10.1140/epjc/s10052-016-4429-6

220. Nagy Z, Soper DE (2018) Jets and threshold summation in deductor. Phys Rev D 98:014035. https://doi.org/10.1103/PhysRevD.98.014035

221. Höche S, Krauss F, Prestel S (2017) Implementing NLO DGLAP evolution in parton showers. JHEP10(2017)093. https://doi.org/10.1007/JHEP10(2017)093

222. Dulat F, Höche S, Prestel S (2018) Leading-color fully differential two-loop soft corrections to QCD dipole showers. Phys Rev D 98:074013. https://doi.org/10.1103/PhysRevD.98.074013

223. Dasgupta M et al. (2020) Parton showers beyond leading logarithmic accuracy. Phys Rev Lett 125:052002. https://doi.org/10.1103/PhysRevLett.125.052002

224. Actis S, et al. (2013) Recursive generation of one-loop amplitudes in the Standard Model. JHEP04(2013)037. https://doi.org/10.1007/JHEP04(2013)037

225. Kallweit S, et al. (2015) NLO electroweak automation and precise predictions for W+multijet production at the LHC. JHEP04(2015)012. https://doi.org/10.1007/JHEP04(2015)012

226. Schönherr M (2018) An automated subtraction of NLO EW infrared divergences. Eur Phys J C 78:119. https://doi.org/10.1140/epjc/s10052-018-5600-z

227. Frederix R, et al. (2018) The automation of next-to-leading order electroweak calculations. JHEP07(2018)185. https://doi.org/10.1007/JHEP07(2018)185

Springer