

HL-LHC Computing Review Stage-2, Common Software Projects: Event Generators

The HSF Physics Event Generator WG

Efe Yazgan (editor)¹ Josh McFayden (editor)² Andrea Valassi (editor)³
 Simone Amoroso⁴ Enrico Bothmann⁵ Andy Buckley⁶ John Campbell⁷
 Gurpreet Singh Chahal^{8,9} Taylor Childers¹⁰ Gloria Corti³ Rikkert Frederix¹¹
 Stefano Frixione¹² Francesco Giuli¹³ Alexander Grohsjean⁴ Stefan Hoeche⁷
 Phil Ilten^{14,15} Frank Krauss⁹ Michal Kreps¹⁶ David Lange¹⁷ Leif Lonnblad¹¹
 Zach Marshall¹⁸ Olivier Mattelaer¹⁹ Stephen Mrenna⁷ Paolo Nason²⁰
 Simon Plaetzer²¹ Christian Preuss²² Emanuele Re²³ Stefan Roiser³
 Marek Schoenherr⁹ Steffen Schumann⁵ Markus Seidel²⁴
 Elizabeth Sexton-Kennedy⁷ Frank Siegert²⁵ Andrzej Siodmok²⁶
 Graeme A. Stewart³ Aravind Thachayath Sugunan²⁷ Zbigniew Was²⁶

¹*National Taiwan University, Taipei, Taiwan*

²*University of Sussex, Brighton, UK*

³*CERN, Geneva, Switzerland*

⁴*DESY, Hamburg, Germany*

⁵*Institute for Theoretical Physics, Georg-August-University Göttingen, Germany*

⁶*University of Glasgow, UK*

⁷*Fermi National Accelerator Laboratory, Batavia, USA*

⁸*Imperial College, London, UK*

⁹*Institute for Particle Physics Phenomenology (IPPP), Durham University, UK*

¹⁰*Argonne National Laboratory, Lemont, USA*

¹¹*Lund University, Sweden*

¹²*INFN Sezione di Genova, Italy*

¹³*Università di Roma Tor Vergata and INFN Sezione di Roma 2, Italy*

¹⁴*University of Birmingham, UK*

¹⁵*University of Cincinnati, USA*

¹⁶*University of Warwick, UK*

¹⁷*Princeton University, USA*

¹⁸*Lawrence Berkeley National Laboratory, USA*

¹⁹*Université Catholique de Louvain, Belgium*

²⁰*Università degli Studi di Milano Bicocca, Italy*

²¹ *University of Vienna, Austria*

²² *Monash University, Melbourne, Australia*

²³ *Laboratoire d'Annecy-le-Vieux de Physique Théorique, France*

²⁴ *University of Maryland, USA*

²⁵ *Institut für Kern- und Teilchenphysik, TU Dresden, Germany*

²⁶ *IFJ PAN and Jagellonian University, Krakow, Poland*

²⁷ *Tata Institute of Fundamental Research-B, Mumbai, India*

E-mail: efe.yazgan@cern.ch, mcfayden@cern.ch, andrea.valassi@cern.ch

Version 1.1 (30 September 2021)

Contents

1	Overview and Executive Summary	1
1.1	Review questions	3
1.2	Review findings	4
2	Funding for generator work	5
3	Review of individual generator packages	6
3.1	MadGraph5_aMC@NLO	7
3.2	Powheg	9
3.3	Sherpa	11
3.4	Herwig7	13
3.5	Pythia8	14
3.6	EvtGen	17
3.7	MC Tools: LHAPDF, HepMC, Rivet, and others	18
4	Opportunities for common projects	22
5	Conclusions and Outlook	24

1 Overview and Executive Summary

This paper has been prepared by the HEP Software Foundation (HSF) Physics Event Generator Working Group (WG), as an input to the second phase of the LHCC review [1–3] of High-Luminosity LHC (HL-LHC) computing, which is due to take place in November 2021. It complements previous documents prepared by the WG in the context of the first phase of the LHCC review in 2020, including in particular the WG paper on the specific challenges in Monte Carlo event generator software for HL-LHC, which has since been updated and published [4], and which we are also submitting to the November 2021 review as an integral part of our contribution.

Physics event generators are essential software components of the data processing and analysis chain of the LHC experiments, and large consumers of their computing resources. Because of the wide range of complex physics processes which must be simulated, the landscape of generator software is extremely varied, including many packages developed by different teams. A structured and comprehensive review of the computational problems posed by generators in the HL-LHC era, including a description of technical challenges such as multi-threading or memory footprint, and

a quantitative estimate of different types of computational inefficiencies that can be addressed by dedicated R&D on both CPUs and GPUs, has already been presented in the published WG paper [4] and we will not repeat it here. In this new document, we mainly go through the generator packages that are most relevant to the HL-LHC experimental program, based on input collected in the WG in Q2 2021 through a series of dedicated meetings [5–13] with the teams who are ensuring the development and maintenance of these packages. For each of these generators, we address the specific questions that we have received as part of the charge for the review [5], while also discussing progress and plans on the issues that we had previously identified as high-priority for the WG.

We would like to start by emphasizing that the event generator area is not a single project (unlike ROOT [14] or GEANT4 [15]), but consists of a variety of independent research and software projects that are managed independently by different teams, which are largely made up of theoretical physicists distributed across several institutes. There are no centralized decisions on policies and strategies. In particular, the different generator and MC tool projects do not fall under the control of the same management structure as the LHC experimental programme, also because in most cases these teams are not directly funded for that goal. That said, there is clearly a (at least partial) convergence of interests between the goals of the generator projects and those of the LHC experiments, and there are several mechanisms for communication between generator teams and experiments. While each generator team has its own agenda, which is largely dictated by several factors relevant to theorists' careers such as the publication of new theoretical research, these agendas are also heavily influenced by the needs of experimental analyses. Many new developments in MC generator software follow specific requirements from the needs of the experimental groups. The generator experts in the experiments (for example, members of the generator or modelling working groups of the experiments, or analyzers who have specific needs) may communicate their issues or needs either directly to MC authors, for instance privately or with presentations or discussions in conferences and workshops, or through people involved in structures like MCnet or (in recent years) the HSF physics generators working group. MC authors address this wide variety of issues or requests as much as possible, but this depends on the person-power and work load in the MC team. Some of the problems do not have immediate solutions and may actually require years of technical and/or theoretical developments. Understanding the current collaboration model between generator teams and the experiments, and finding ways to improve the management and funding (and, to some extent, communication) mechanisms which are needed for the mutual benefit and success of both parties, are one of the main challenges in this area to be discussed in the LHCC review.

The structure of this paper is the following. Section 1 includes an overview of the paper and of the questions we addressed during the review process, a general de-

scription of how the generator projects are internally managed and how they interact with the experiments, and an executive summary of our main findings. Section 2 gives more details about the important issue of funding and careers, which is common to all generator packages. Section 3 goes through various generator packages individually, based on the input collected from the corresponding teams. Section 4 gives a non-exhaustive list of ideas for possible common projects in the generator area, including more than one generator team and/or external partners such as the experiments, HSF, CERN IT or CERN EP. We finally give some conclusions and an outlook for further follow-up in Section 5.

1.1 Review questions

Among others, we address the following questions for each generator:

- Are there plans and funds in place to continue support through HL-LHC?
- What major physics updates are foreseen for HL-LHC? Examples: transition from NLO to NNLO (next-to-leading order to next-to-next-to-leading order) calculations for ME+PS (matrix element plus parton showers), improved shower models, higher jet multiplicities, etc.
- What major software updates are foreseen for HL-LHC? What will be the main bottlenecks? In particular, what are the current CPU performance bottlenecks and how are they being worked on? Are there mitigation strategies to reduce negative weight fractions? Are improvements foreseen in phase-space sampling and unweighting efficiency?
- What work is in progress to adapt the software to GPUs and heterogeneous architectures for HL-LHC?
- Are there issues which may complicate the use of the generator by the experiments in multi-threaded workflows and what is the progress in this area?
- Is there any work in progress to include Machine Learning (ML) tools as part of the generator workflow?
- Are there issues or areas of work where help from HSF or from the experiments may be needed?
- Are there generators/tools not listed below that are expected or wished to become heavily used by the experiments?

1.2 Review findings

The following is a high-level summary of our findings:

- Funding and career opportunities are problematic. While the interpretation and sensitivity of HL-LHC results heavily depend on the development and maintenance of generator software packages, their authors (most often theorists and phenomenologists) are generally not directly funded to work for the HL-LHC experimental program, and often lack the career incentives to provide the software and computing support and improvements needed by the experiments. While the number of core contributors with permanent positions and their personal interest in the experimental program seems currently enough to provide at least basic support for most generator packages, there is limited funding for post-docs and PhD students. Thus, the loss of know-how remains an important risk, and in some cases this has already started to happen.
- Concerning the problem of events with negative weights, work is in progress in both MadGraph5_aMC@NLO (MG5aMC henceforth) and Sherpa on improved NLO matching prescriptions. The recent work on the MC@NLO- Δ technique seems promising, but there is not yet a clear timeline for its use in a new production version of MG5aMC, also because its implementation requires a strict coordination among the MG5aMC and Pythia8 authors. The Sherpa team is also working towards achieving similar goals. It is important to note, however, that improvements in precision calculations, such as new parton shower algorithms (e.g Dire) and NNLO in QCD matrix element calculations, are expected to come with new sources and an increased level of negative weights. Therefore, despite the recent progress, negative weights will likely remain an issue that needs close attention.
- Several teams expressed an interest in porting some of their software to GPUs, either directly or through Machine Learning frameworks. One development that looks especially promising is the ongoing work on reengineering the matrix element calculation in MG5aMC, to exploit not only GPUs but also CPU vectorization. This work is making rapid progress and might hopefully result in a new production version of MG5aMC in the next year, also thanks to a collaboration between theorists and software engineers which was largely promoted by the HSF generator WG.
- phase-space sampling is an important ingredient for efficient unweighted event generation. There are some new ideas and studies to improve unweighting efficiency, such as Machine Learning, for example, neural networks in Sherpa. However, these improvements have not yet been carefully validated or used for large scale productions by the experiments.

- Many analyses at the LHC depend on reweighted MC events to do systematic uncertainty calculations. PDF variations, parton shower initial- and final-state radiation variations, ME renormalization and factorization scale variations can be handled through computing alternative weights for the same physical event, saving CPU cycles. However, this cannot be done in some cases, such as for parton shower starting scale and underlying event tune variations (including color reconnections and intrinsic transverse momentum). This results in extra load in MC production systems of the experiments because of the additional CPU cycles to produce each sample with different parameters. Better flexibility in (re-)deriving scale and PDF weights can be obtained by writing more information to the event record (at least in Sherpa), but the CPU and storage implications of this have not yet been assessed.
- Since the previous stage of the LHCC review, and having been highlighted strongly in the published WG paper, the generator groups are performing more in-depth profiling of the CPU costs and possible computational inefficiencies of their software. By clearly identifying some performance bottlenecks, this has already allowed major improvements in several generators, such as helicity recycling in MG5aMC, an optimised program flow for weight variation calculations in Sherpa, and a new interface to MCFM’s analytic one-loop amplitudes; more details for each of these are available in the corresponding generator subsections in Section 3. In parallel, a more detailed study of the CPU costs of MC generation campaigns in the experiments is progressing, also in terms of CPU times normalised to the HEP-SPEC06 (HS06) benchmark. Establishing better mechanisms to collect CPU costs and generation efficiencies (e.g. due to sampling, merging or filtering) from the production systems of experiments would allow for an easier comparison between different experiments and different generators.
- Several teams noted that a collaboration with experts in the HSF would be very useful and welcome to help them in the optimization or reengineering of their software. One group mentioned that a software review, like that organised by the HSF for the GeantV project in the past [16], may eventually be an interesting option. We also believe that redesigning and harmonising software interfaces, to improve package modularity and ease the integration of plugins from different projects, would be a useful development, a view that was shared by some of the generator groups.

2 Funding for generator work

Funding and career perspectives are a major common issue reported by many generator experts. Those working in the development, optimisation or execution of

generator software provide essential contributions to the success of the (HL-)LHC physics program and it is critical that they get the right recognition and motivation. However, theorists get recognition on published papers, and may not be motivated to work on software optimisations that do not have enough theoretical physics content to advance their careers. Generator support tasks in the experiments may also not be enough to secure jobs or funding for experimentalists pursuing a career in research. It is still not clear for people who are doing work on generators or related tools whether it is really possible to have a dedicated career path from now throughout HL-LHC. This is slowly being acknowledged from funding bodies, but the community has low expectations especially because there is no sign that hiring policies will be modified to provide reasonable funding for generator work. The delay caused by the lack of funding for generator work has already started to cause a loss of know-how in some generator packages.

In general, MC tools (see Section 3.7) are even less supported than generators. The work involved for these tools has a lot of technical parts but little physics content. One needs to find a good model to support this type of work. Such tasks may be ideal for Research Software Engineers (RSEs), an approach which seems to work well in the UK for similar situations: the idea is that these are relatively small, self-contained, tasks which do not require a large number of dedicated FTEs over a long time, but could be dealt with for example by allocating 10% of the time of one person from a pool of RSEs with the required technical competencies. Additional support from CERN EP-SFT for MC tools, which was possible in the past, would be very welcome, for instance through the GENSER project that is already supporting common builds and distributions of generator packages for the LHC experiments.

It should also be noted that the EU ITN MCnet [17] network has been used extensively for more than a decade by many generator teams, especially for more optimisation-focused work, but this project was not renewed and expires in few months. This funding was particularly important for MC tools, and needs to be replaced with another source to ensure their continued support.

On a positive note, some recent funding initiatives have explicitly targeted work in the event generator area. This is the case for instance for the DOE funding of the High-Energy Physics Center for Computational Excellence (HEP-CCE) [18] and of the STFC funding of the SWIFT-HEP [19] project, both of which could be interesting models also for other countries to fund software evolution in HEP.

3 Review of individual generator packages

In this section, we individually go through the generators that are most relevant to the HL-LHC experimental program. Each subsection is based on the input, and in most cases on the text contribution, from the team ensuring the development and maintenance of the corresponding package.

In particular, we were explicitly asked [5] by the reviewers to cover MG5aMC, Powheg, Sherpa, Pythia8, EvtGen and HepMC. We then added Herwig7, LHAPDF and other MC tools. There are also other generator-related software packages that have not been discussed in this document, even if they already are widely used, like FastJet [20], or may eventually become heavily used by the experiments, like MATRIX [21] and Geneva [22, 23] or dedicated resummation frameworks built in event generator spirit like CVolver [24]. We have also not covered any dedicated heavy-ion generators, such as EPOS [25] or HIJING [26]. All of these are likely candidates for topical discussions in the upcoming meetings that we plan to organize in our WG.

3.1 MadGraph5_aMC@NLO

MadGraph5_aMC@NLO [27] (MG5aMC) is a multipurpose hard-event generator that carries out automated calculations of SM and BSM processes at the leading- and next-to-leading order (henceforth LO and NLO respectively) in QCD. It may deal with any theory whose Lagrangian can be specified by the user through an UFO [28] model, including EFTs. Such computations may be either at fixed order, or matched and/or merged to parton shower MCs. Fixed-order LO and NLO results are also available (see in particular Ref. [29]) in the combined perturbative expansions in the QCD (α_S) and EWK (α) couplings, namely for any matrix element that factorises the quantity $\alpha_S^n \alpha^m$, with $n + m$ a constant. As such, the program features a significant flexibility, with extensive user-definable run and configuration settings. In keeping with its scope of hard-event generator, MG5aMC only produces parton-level events, and delegates parton showering and hadronisation to external tools such as Pythia8 and Herwig7.

The MG5aMC collaboration is a composite team of theorists with different expertise. Most members hold permanent positions and one is in a software-focused role. Software development is shared semi-equally across several members of the team. But the collaboration are of the opinion, shared by other groups, that time to work on performance improvements, even by significant factors, is not well recognised by funding agencies.

The version control of MG5aMC is done through Bazaar. Its webpage [30] in launchpad has extensive information and support.

Support through HL-LHC. There are no dedicated funds for HL-LHC. The effort comes through different funding schemes with typically a theoretical emphasis. MCnet funding was used in the past for more optimisation-focused work but this project was not renewed and expires in few months.

Major physics updates for HL-LHC. Physics improvements are foreseen especially for higher-order EWK corrections, in particular for the extension of fixed-order capabilities in a QCD+EWK mixed expansion to include their parton-shower matched/merged counterparts.

Major software updates for HL-LHC and main bottlenecks. The main consumer of CPU cycles is, as expected, the ME calculation both at LO and NLO. Many optimisations are already in place for the loop calculations, so the NLO real emission is the main bottleneck there. Recent improvements, such as helicity recycling, have been implemented, achieving speed-ups [31] ranging from factors 2-3 for standard processes such as W +jets and $t\bar{t}$, up to four orders of magnitude for some VBF-like processes, but these still need to be ported to NLO. When the issues above will be addressed or mitigated, MG5aMC authors believe that one of the most time-consuming operations will be the computation of the colour algebra in an exact manner.

Negative weights are an issue at NLO as in many generators. Processes such as V +jets can have a negative weight fraction as high as 15-30% for configurations with multiple additional partons. However, the negative weight fraction can depend more on the process than on the number of additional partons, for example, the negative weight fraction in $t\bar{t}$ production is larger than V +jets for the same final state multiplicity. Significant work has recently been undertaken to improve this in the MC@NLO- Δ technique [32], which builds upon the MC@NLO matching formalism [33]. The improvements from MC@NLO- Δ vary significantly depending on the configuration being considered, but factors of 2-3 reduction in number of events required for a given statistical power can be obtained even for standard processes like W +jet and $t\bar{t}$. The final implementation of this requires a strict coordination among the MG5aMC and Pythia8 authors.

There has been much effort on improving the LO phase-space integrator [31]. This mainly addresses VBF/VBS processes, where the gain in efficiency is very large. Neutral network approaches to phase-space integration, in collaboration with the VegasFlow [34] team, are also starting.

Significant effort is being invested in the reengineering of the matrix element calculation, targeting GPUs (using CUDA but also abstraction layers like Kokkos, Alpaka or SYCL) as well as CPU vectorization in C++, in collaboration with members of the HSF Generators Working Group from CERN IT and other institutes. With respect to the current Fortran implementation of the ME calculation, speed-ups by factors of 3 to 4 on CPUs through C++ vectorization were observed [35] for a simple electroweak process, while the CUDA implementation on a full V100 GPU was observed to be 1000 times faster than Fortran on a single core of a CPU. An additional speedup by a factor ~ 2 was observed in both cases when moving the ME calculation from double to single floating point precision. While this progress is very promising, many essential developments like the integration of the new ME calculation in the existing Fortran framework for phase-space sampling and unweighted event generation, the backport to code-generating code to allow the ME computation of any LHC process, and a more careful analysis of QCD-specific computations like color algebra, are all still work in progress. In addition, this work is restricted to calculation of LO matrix elements, and extending this to NLO will also require

significant effort, in part due to added complexities, such as loop diagrams, some of which may require quadruple precision. Just like most of the other speed-ups described in this document also for other generators, it should be stressed that a realistic extrapolation to achievable speed-ups and improvements for the LHC experiment productions will only be possible after many additional careful studies of software integration and physics validation.

Issues or areas of work where help from HSF or experiments may be needed. There is already a lot of effective help in the GPU project [35]. More coordination with experiments, either via HSF or directly, would be very welcome, for example in defining a better release strategy and on improving the communication between the generator authors and the experiments. MG5aMC authors have specifically pointed out [9] that the feedback cycle between generator teams and experiments on software updates, although much improved in recent years, would benefit from being made more efficient.

3.2 Powheg

Powheg [36–38] is a framework that hosts the code to simulate many physics processes. The main focus of Powheg is matching of accurate fixed-order predictions with PS for SM processes, mostly for LHC physics. Several BSM applications also exist, but BSM is not the main focus so far. Powheg is publicly available at <http://powhegbox.mib.infn.it> and is open to all theorists that want to contribute. There are about 100 processes from contributions of about 100 physicists. There are some “core” developers of Powheg but it is not a strictly well-defined collaboration. Powheg has fully supported interfaces for OpenLoops [39], GoSam [40, 41], Madgraph4, and MG5aMC. The version control of Powheg is done through SVN. Its webpage has extensive reports for bugfixes and revisions. It has two main releases: Powheg BOX V2 which is the main release containing almost all processes, and Powheg BOX RES which is the most recent release created to deal with processes with resonances in particular.

Support through HL-LHC. Although there are no funds specifically dedicated to the maintenance of the framework, support and development will continue through HL-LHC by physicists with permanent positions who are already working on Powheg, as long as there is interest in the LHC community.

Major physics updates for HL-LHC. Major directions for physics updates are NNLO QCD+PS (MiNNLO_{PS} [42, 43]), NLO EWK+PS, and interplay with modern PS. All of these updates, in particular the interplay with modern PS codes, may require major rethinking, restructuring, or re-coding of core parts of the software. Although it is difficult to predict, there are some possible processes that could be

addressed to increase precision, such as loop-induced processes, for which some results already exist, and extremely high-multiplicity processes.

Major software updates for HL-LHC and main bottlenecks. ATLAS, CMS and LHCb have raised issues relating to difficulties due to the nature of independent implementations of the various Powheg processes. This requires a lot of non-user-friendly by-hand work to get compilation to run smoothly in the experiments' software environments. An improved central infrastructure would greatly improve this situation. Considering the person power, the core software is expected to remain Fortran based with modern f90 structures. There is growing interest in using GPUs for some aspects which will be detailed later.

Issues or areas of work where help from HSF or experiments may be needed. The Powheg authors already find it very useful to have contacts in the experimental community, in particular to test new developments especially before public releases, and also to test large scale production. HSF could coordinate refactoring and reworking that could be useful as a common work. MG5aMC authors and others in CERN IT have been working together to make MG5aMC work with GPUs. For Powheg, a specific example could be identified, and this could start a collaboration between HSF, CERN IT, Powheg core authors and Powheg contributors who provide the ME calculations.

ME+PS generator specific considerations. As indicated above, the major directions for physics updates are to increase the precision of the predictions with NNLO QCD+PS (MiNNLO_{PS}), NLO EWK+PS, and interplay with modern PS (e.g. PanScales [44]). LHC experiments have already started using some of the NNLO simulations from Powheg for the available processes. Until now, Powheg dealt with CPU bottlenecks through reweighting (for example with multiple PDFs, one or two loop amplitudes, scale variations), or through parallelization. In reweighting, calls to CPU-intensive routines are minimized to avoid re-computing. In parallelization, only using multicore has been sufficient without any complicated arrangements. It was noted by the experiments that some aspects of the reweighting functionality could be improved, such as the extension to physics parameter reweighting, m_W , m_t and couplings, and also standard scale and PDF weights in MiNNLO_{PS}.

Some initial attempts have been made using MPI and there are plans to explore using GPUs or Machine Learning inspired techniques, for example for efficient phase-space sampling and for generation of underlying Born level events. The latter is the typically the more computationally delicate aspect of the Powheg core algorithm. It requires large statistics, and multidimensional integrals. Machine learning is being considered as a way to provide efficiency improvements in this area, but no work has been done in this direction so far.

Negative weight events do not seem to be a major issue with respect to other generators, except at NNLO. There are already many processes available at NNLO QCD and they have negative weight fractions ranging from about 10 to 30% depending on the process, which are much higher than the negative weights at NLO QCD in Powheg. However, at least for Drell-Yan with $\text{MiNNLO}_{\text{PS}}$, the fraction of negative weights can be substantially reduced (to a few percent) through folding. Larger negative weight fractions at NNLO are mostly due to the difference in scales in $\text{MiNNLO}_{\text{PS}}$ and complications with photon emissions. It may also be possible to implement a positive resampler [45, 46].

In Powheg, having a harmonised framework to benefit from automated tools, even just for NLO, would be very useful especially for some processes, e.g. for $W + c$ which has large negative weights in MG5aMC but doesn't exist in Powheg. The implementation of new processes in Powheg could be made simpler through improvements of the existing interface to MG5aMC. In particular, it would be useful if the interface exposed single amplitude contributions from the matrix element calculation, which are available in MG5aMC (where they are used by MadEvent), but not yet to external generators like Powheg.

3.3 Sherpa

Sherpa [47] is a general-purpose generator that simulates hard interactions, radiative corrections, multiparticle interactions, hadronization, and decays. For the hard interaction it can simulate events up to NNLO in QCD (for selected processes) and NLO QCD with one-loop EWK.

Support through HL-LHC. Several of the Sherpa team are in permanent positions, and furthermore post-docs and PhD students have regularly contributed to Sherpa development and support over the last decades. In the UK, Sherpa posts at IPPP are seen as part of IPPP's core mission and therefore are expected to be treated as core in future funding reviews.

Major physics updates for HL-LHC. The major physics updates foreseen for HL-LHC are an NLO parton shower, NLO multi-jet merging in decays, resonance-aware matching and NLO subtraction, NNLO matching for important processes, a full SM parton shower and NLO EW matching, SHRiMPS (inclusive QCD scattering simulation), improved bottom- and charm-quark decays, a fully validated Standard Model as an Effective Field Theory (SMEFT) and improved heavy-flavour parton evolution and NLO matching schemes.

Major software updates for HL-LHC and main bottlenecks. The major software updates foreseen for HL-LHC are restructuring of event generation framework for improved efficiency, separation into ME and PS generator, and GPU support in ME generation. For restructuring of the framework and porting of code to GPUs,

help from HSF and experiments would be extremely valuable. Performance analysis of the code has been done with the Intel Vtune profiler, MAQAO, MUST, and with Valgrind. A large part of the time is required to compute all weight variations. When weight variations are ignored, a typical multijet process spends 96% of its time in the ME step, including phase-space and matching. Without variations, it takes 1 s per event, versus more than 10 s per event with variations. And when variations are computed, more than 80% of the CPU time goes into LHAPDF. A significant bottleneck that has recently been identified and fixed is that many potentially large numbers of systematic variations (mainly for PDF uncertainties) were being calculated before the unweighting procedure, so for many events they would not in the end be used. The restructuring gives a factor of 5 speedup for a representative V +jets setup similar to the one used by ATLAS; the eventual speedup achievable by ATLAS in production, however, can only be assessed by the experiment through further tests in its own software framework and configuration. At the cost of theoretical precision, some computationally expensive aspects of NLO event generation can be simplified. For example, in some cases, spin correlations and color connections in the NLO matching procedure could be ignored [48]. This is optional and already available in Sherpa 2.8.8. Some inefficiencies in the clustering algorithm are also being addressed. After these optimizations, there are still three bottlenecks: one loop ME, PS clustering, and maths library. To improve one loop ME calculations, a generic interface to the MCFM [49,50] library has been constructed. This sped up the computation of one-loop MEs by at least a factor of 6-10 with respect to automated one-loop providers, but is available for selected processes only. Another improvement for computing performance is expected to come from better phase-space sampling, to improve unweighting efficiency. To this end, the Sherpa collaboration explores variations of multi-channel integration techniques as well as Machine Learning [51, 52]. The usage of Neural Network surrogates for the full event weight has shown significant potential to accelerate unweighted event generation [48, 53]. An ongoing effort to calculate many-gluon tree amplitudes on GPUs [54] is being extended to the full SM at present. A further efficiency improvement will be the reduction of the fraction of negative weights in NLO QCD calculations. There has already been significant progress in this direction [48], and there are several other ideas left to investigate. NLO parton showers will induce additional negative weights, whose rate can not be predicted reliably at present.

ME and PS specific considerations. In the experimental analyses, some uncertainties such as PDF or ME scale can be calculated using weights stored in the samples. This saves experiments generating many extra samples for systematic variations. Sherpa provides standard scale and PDF variation weights along with parton shower renormalization and factorization scale variations and uncertainties on NLO EWK corrections. For HL-LHC, due to the long timescales involved, it

may be desirable to have the ability to change PDF sets some time after a sample has been generated to avoid re-producing samples for systematic uncertainty calculations and modeling studies. Sherpa has the ability to write-out all coefficients needed to do a posteriori arbitrary scale and PDF variations into the HepMC event record [55]. This offers more flexibility than the on-the-fly reweighting methods. However, for a MEPS@NLO [56, 57] event of a non-trivial multiplicity, or a high-multiplicity MEPS@LO event, this can easily be >50 to ~ 100 coefficients. The preferred usage mode is a strategic choice to be considered by the experiments for future large productions.

Unlike the scale and PDF variations mentioned above, there are some variations which, in all generators, require alternative samples as they cannot be derived through event weights (for example, parton shower starting scale, tune variations, colour reconnections or hadronisation models). However, an example of another variation typically needing alternative samples, the merging scale, is being developed as a weight variation in Sherpa 3.0.0.

NNLO simulations are computationally very expensive, and the method of choice is likely to sacrifice precision in a controlled manner. For example, at NLO, color and spin correlations are typically dropped in PS matching. In some cases there is very little difference between the predictions from an NNLO calculation and from an NLO calculation with additional corrections, so one could push to detector simulation only with NLO and apply some corrections for NNLO a posteriori. We need to understand the cases when we think we really need NNLO at the fully differential level.

3.4 Herwig7

Herwig7 [58] is a general purpose particle physics event generator. It computes any observable at NLO precision in perturbation theory matched/merged to a PS. It has two main shower modules: angular ordered and dipole-type, both including PS uncertainty estimated. It includes cluster hadronization and Eikonal MPI models and colour reconnection. It also allows automated BSM simulations using UFO model files. The release of Herwig7 was a significant update with respect to the previous Herwig++ interface. On top of the previous angular-ordered shower there was the addition of a dipole shower as well as a new Matchbox tool [59] allowing an interface to a large range of ME and loop providers and libraries, such as MadGraph and OpenLoops. In addition, there has been work on improving the logarithmic accuracy and including reweighting functionality for systematic uncertainties in the parton shower, and significant development on the minimum bias and underlying event simulation.

Support through HL-LHC. The Herwig7 team is comprised of a core of permanent staff and a semi-regular cycle of post-docs and PhD students. The plans and funds depend on institution and strategic goals. Funding is a serious issue. Any

funding for generator development must be strongly linked to a very strong physics case. In some cases this may not be what the experiments need. To fund the bread and butter of what the experiments need, one would need different specific funding.

Major physics updates for HL-LHC. The major physics improvements foreseen are to accommodate higher-order EWK corrections and matched NNLO in QCD predictions. A full implementation of NNLO+PS matching requires developments of physics understanding in, and therefore changes to, the parton shower algorithms.

An interface to string hadronisation model is under development. This would complement the existing cluster hadronisation model and allow for a much more robust estimate of the systematic uncertainties associated with these algorithms, which are projected to be dominant in some important physics goals at HL-LHC.

Major software updates for HL-LHC and main bottlenecks. To accommodate many physics improvements, such as those in the understanding of the shower, a refactoring/rebuilding of some parts of the ThePEG framework, upon which Herwig7 is built, is needed.

Negative weight fractions in Herwig7 for NLO processes have been prohibitively high for significant adoption in the experiments. Developments are in progress to improve these negative weight fractions, such as re-ordering contributions and re-sampling algorithms. This work also relies on addressing the structural problems in ThePEG previously mentioned. KrkNLO [60], an alternative matching scheme that can be used in Herwig7, has no negative weights, but it is not automated and is limited to only two processes for the moment. In addition, it cannot be used for merging and its systematics are not yet understood to the same level as in the default Herwig7 matching prescription.

At this point, there is no work in progress to adapt the software of Herwig7 to GPUs and heterogeneous architectures. Things like showering are clearly a bottleneck for single instruction, multiple data (SIMD) processing. However, the ME calculation is vectorizable and Matchbox actually already heavily relies on interfacing to MG5aMC for MEs, which will allow porting to GPUs.

There is no work in progress to include ML tools as part of the generator workflow at this point.

3.5 Pythia8

Pythia8 [61, 62] is an event generator for high-energy particle collisions. It contains physics models for hard interactions, radiative corrections, multi-parton interactions, hadronization, decays, beyond the standard physics, γ -beams and collisions, heavy-ion (HI) collisions, and other phenomenon. A hallmark of the Pythia8 collaboration is providing good documentation, prompt responses to user questions and problems, and a rapid response to coding errors. It has no dependencies on external packages

like ThePEG or ROOT, is easily extendible with a UserHooks class, and has interfaces to many external packages and inputs, such as the (S)LHA interfaces. The Pythia8 Collaboration (<https://pythia.org>) avoids duplicating efforts of other groups unless it is necessitated by physics or computing.

Support through HL-LHC. Pythia8 is an international and multi-institutional collaboration with no formal constitution. Like the other projects discussed in this report, the Pythia team has a few senior researchers, some junior researchers, and a rotation of graduate and post-doctoral students. As such, the continued development of Pythia8 throughout the HL-LHC era depends on funding agencies and laboratory priorities, though these are not expected to change. The biggest challenge to the collaboration is finding permanent positions for junior people so that they can continue their valuable contributions.

Major physics updates for HL-LHC. Major physics updates foreseen for HL-LHC are NLO parton showers, resonance-aware matching and NLO subtraction, fully differential NNLO matching for important processes, full SM parton shower and NLO EW matching, automated ME-corrections interfacing with ME generators, forced hadronization and rare hadron production, and QED showers that also work for decays.

Major software updates for HL-LHC and main bottlenecks. In general, Pythia8 is not a computing bottleneck. As a result, there has not been extensive or systematic profiling nor dedicated attempts to optimize the code. Performance improvements could be expected in the time and memory usage in some merging configurations, parton shower variations, multi-threading, and heavy-ion collisions.

Few of the proposed physics updates are trivial, and they will require some large code additions and a thorough validation process. VINCIA [63] and DIRE [64–67] do not have a large user-base at this point, but are now integrated directly to the core code, which was already a major update. Ref. [68] presents the implementation of CKKW-L [69] merging with VINCIA sector showers. This approach mitigates the bottlenecks in the standard multi-jet merging implementations. Ref. [68] shows that the time spent constructing sector shower histories has an approximately linear scaling behavior with increasing number of hard jets and the parton-level event generation and the total memory usage do not increase. With the merging scheme in Ref. [68], it is already possible to have more than 9 hard jets from the shower. Using this approach, the remaining main bottleneck in simulating merged configurations would be on the fixed-order calculations.

To estimate uncertainties related to the simulation of the parton shower, many CMS and some ATLAS simulations use the event-by-event parton shower variations available in Pythia8 through event weights. In the CMSSW [70] environment, parton shower variations increase the run time by a factor of ~ 3 -4 on average. But here,

since one run corresponds to many independent runs, there is a net gain in CPU time, memory and storage. However, it may still be necessary to find methods to improve performance.

Pythia8 is "thread friendly" in the sense that multiple instances of the Pythia8 class can be used concurrently on different threads. Of course, this is still dependent on the thread-safety of any external packages that are interfaced to Pythia8 at run time. For example, EvtGen is known to be *not* thread safe. CMS has made some simple tests, demonstrating that the use of multiple threads increases the speed of the simulation. However, serializing and de-serializing data products takes considerable time, lessening somewhat the impact. There is still significant room for improved MP/GPU/concurrency.

It is not entirely clear how to leverage GPUs and GPU-based HPCs. One option, which would still require some developments, is to have Pythia8 delegate matrix element calculation to the GPUs, for instance either by external calls of the MG5aMC ME API or by reading HDF5 events.

ALICE, CMS and LHCb use the Pythia8 heavy-ion option "Angantyr" [71]. Simulating HI collisions takes much longer compared to pp collisions. Simulating PbPb collisions (at $\sqrt{s_{NN}} = 5.02$ TeV) is two to five orders of magnitude slower than simulating pp collisions depending on the configuration. For example, hadronic re-scattering in HI collisions is a source of significant slowdown. More efficient reading of LHE files could also be investigated.

ME+PS generator specific considerations. A major use case of Pythia8 is in conjunction with external ME calculations that are matched or merged with the parton shower. To understand better the current CPU performance, Pythia8 needs more dedicated profiling. The acceptance efficiency is something that could be improved. The veto method is currently delayed until late in the algorithm, which allows Pythia8 to calculate the no-branching probability and to conserve unitarity at a loss of CPU efficiency, but the extent of this inefficiency is currently not known quantitatively. Little is known about the large scale performance of VINCIA and DIRE. What *is* known is that the VINCIA shower is an order of magnitude slower than the standard shower. Promising areas of improvement are known, but no one is currently working on this. NNLO matching in Pythia8 is possible with DIRE and being worked on with VINCIA. At this point, no issues are foreseen for this, but it should be investigated more thoroughly. Currently, Pythia8 does not offer any actual functionality for NNLO+PS, neither in DIRE nor in VINCIA.

For all updates, the main bottleneck is finding time and people. Help from HSF and experiments may be needed in restructuring of code, porting of code and algorithms to GPUs (if there is a use case), profiling for things people actually use for the HL-LHC era, and HDF5 usage and exploitation. Pythia8 would benefit from HSF support on issues that increase the common good in the generator world. Basic

profiling for relevant use cases would be beneficial for improving Pythia8 computing performance. VINCIA and DIRE are recent major updates that need to be tuned to data and to be test-driven at a larger scale by experiments to provide feedback both on physics and computing performance. One foreseeable issue with using DIRE is the significant negative weight fractions, especially when DIRE is used with NLO ME that already has negative weights. VINCIA does not have negative weights and tries to avoid them as far as possible in future developments. The Pythia8 authors expressed an interest in embarking on a GPU port through collaborations with HSF members and the experiments, similar to the ongoing work with MG5aMC.

3.6 EvtGen

EvtGen [72] is a package specialised for heavy-flavour hadron decays. It consists of about 130 decay models which implement the specific dynamics of various decays. EvtGen maintains a detailed decay table with a large number of processes. Known decays do not add up to 100% branching fraction. What is missing is filled up by generating quark configurations and passing those to Pythia8 for fragmentation. The τ decays are done through Tauola [73], and Photos [74] is used for radiative corrections. The code has been largely stable over past 10 years with most of the changes due to added models. Some modernisation and cleanup was done rather recently.

Support through HL-LHC. The core EvtGen team has 4 members who have a fraction of their time funded for EvtGen through an UK STFC grant. Part of the funding is attached to LHCb as part of the responsibility (for general maintenance) within the experiment. One of the team members is funded by Warwick university as part of a collaboration with Monash. Beyond the core team, others contribute by writing new decay models driven by the needs of measurements. There is some funding for a computer engineer to work on code redesign for multithreading.

Major physics and software updates for HL-LHC and main bottlenecks. No updates for physics is foreseen in the near future. Currently, the team is working on some code consolidation to unify coding style and for C++ modernization. There is a plan to decrease code duplication within the decay models and to have an updated documentation to be in Doxygen, as well as a paper or a manual. There seem to be difficulties with the multithreaded environments that experiments are moving to, due to issues with thread safety, but the plan is to allow event level multithreading. In this context, event means the particle to be decayed through the full decay chain. The main bottlenecks are Tauola and Photos, neither of which allows multithreading. There are also structural limitations in EvtGen for running multithreaded mode, such as a global instance of random number generators. Structural changes can be done, but this requires a major re-write of the code. The issues in Tauola and Photos could be fixed with more time and help. It is also possible that, for the

dependencies, Tauola might be replaced by Pythia8 and the interface to Photos may be improved or instead an alternative such as SOPHTY [75] may be used. Photos is used in practically in every decay in the usual use cases. It is seen from profiling that significant CPU is spent in Photos itself and in EvtGen the conversion to and from HepMC is significant. Pythia8 provides τ decays using the helicity formalism. EvtGen already depends on Pythia8 for other reasons, but that functionality could potentially be used for τ decays. The main interface is being finalized. Unlike Tauola, Pythia8 provides amplitude for particular spin states and sums over spin states are done by EvtGen. One challenge is passing the spin states from EvtGen to the tau decayer. This is being finalized as well. The decay table is one of the things that is difficult to update because of unobserved (i.e., explicitly listed) decays and consistency of particle properties between generators. It is important to agree with other generators about particle properties. In LHCb, the new version of the simulation under development uses a single EvtGen instance which is locked as needed for multithreading. LHCb, by default, runs generator and simulation in the same job. This is not the same in ATLAS or CMS in all cases.

3.7 MC Tools: LHAPDF, HepMC, Rivet, and others

Tools to support and build on MCs are an important part of the ecosystem of collider phenomenology. They are socially distinct from typical event-generator development, the teams working on them being as much or more rooted in the experimental than theory community. Sustainability is clearly an issue since most of this work is not funded, particularly with the end of more than a decade of funding through the EU ITN MCnet network. MCnet has provided funding in particular for MC-development PhD students and for theory and experiment PhD students to contribute to generator-project initiatives, and for travel and meetings to enable collaboration between developer and user groups. It is undoubtedly the case that development of shower+hadronisation MC generator features, and in particular the tools that connect them to experiment and phenomenology, will be curtailed by the end of MCnet-supported activities. This funding needs to be replaced with another source to ensure continued support and development of these essential tools.

LHAPDF. LHAPDF [76, 77] is a library for PDF-value access, and associated metadata including α_s calculations consistent with the PDF fit. LHAPDF was redeveloped from v5 to v6 in 2012/13 as a purely interpolation-based tool rather than recalculating DGLAP evolutions as with some older PDF sets in v5 and earlier. It defines the standard data format for all PDF-set data, is the standard PDF interface for LHC-experiment MC production, and is MPI-friendly.

However, its CPU cost is unnecessarily high: profiling of Sherpa+LHAPDF in experiment-like NLO V +jets setups has shown that the main CPU consumers are calls to LHAPDF and transcendental functions, e.g. log. Regarding the latter, there

could be a 10% gain in LHAPDF performance simply from runtime use of the Intel math library; this is already in use by ATLAS, with other benefits throughout the software chain. In 2020, LHAPDF 6.3 introduced thread-safe caching as a response to these studies, in particular to benefit from thousands of identical (x, Q^2) phase-space-point queries across the set of parton flavours, or large PDF-fit “error sets”. Efficient use of this caching, however, depends on the generator codes using a commensurate call strategy: repeated calls to the same phase-space point are often not batched together, but separated by many other PDF evaluations, forcing cache misses.

A final MCnet short-term studentship project has focused further on optimisations such as data restructuring and precomputation of repeatedly used functional derivatives, and will prototype a GPU-based interpolation feature for use with GPU-based matrix-element sampling. A Google Summer of Code (GSoC) project on implementation of higher-order interpolators further revealed improvements to the interpolation heuristic, and clarified that for N3LO calculations the precision of the interpolation may in fact be less important than its global smoothness: this motivates future development of alternative interpolator strategies which prioritise higher-derivative continuity over either precision or CPU efficiency. A UK SWIFT-HEP part-time research assistant on LHAPDF and corresponding MC efficiency will extend this work starting from October 2021. CERN EP-SFT provides part-time staff support for adding new PDF data to the EOS-based PDF-set repository, previously hosted at Durham IPPP.

In summary, LHAPDF and HepMC are crucial to MC production in experiments. There is a huge PDF-interpolation cost, which can be largely addressed by concerted efforts to restructure generators to make more efficient PDF calls; trivially, by lazily evaluating error sets in ME sampling, and less trivially by reorganising call orders to make better use of caching, either in the generator or in LHAPDF. Further improvements to the interpolation library, for focuses on efficiency, precision, and smoothness are also required, as are extensions to support generalised PDFs (in particular resolved-photon PDFs for EIC and FCC-eh, and GPDFs and DPDFs for small- x hadron-collider physics). All these developments, technical and physics-oriented, require explicit community support.

HepMC, LHE and other MC event formats. HepMC [78] is the standard library for event-graph representation. HepMC v1 had undergone problematic forking between the CLHEP project and the LHC experiments, which was resolved with v2 before the start of LHC data-taking. However, this version retained other structural problems, including an unmaintainable API, poorly organised code, ROOT I/O unfriendliness, logical inconsistencies in e.g. const-correctness, and limitations in handling of heavy-ion and other requirements from specialised areas of collider physics. HepMC3 [79] is a major ground-up rewrite, addressing these issues and making the library more extensible. CERN EP-SFT helped with the v3 rewrite,

project management and some funding, until the rewrite was finished.

There is still much work needed to optimise the efficiency of the graph algorithms used to write and explore HepMC event-graphs, and technical C++ optimisation of the data structures. Event copying, graph-walking, and filtering/rewriting are CPU-expensive, which can likely be reduced by expert technical developer time. New features such as the “attributes” have not been optimised for I/O or CPU efficiency, and are similarly amenable to improvements. Usability is another area ripe for development: more expressive search and filtering functionality in the API, and provision of a set of command-line tools for event-file splitting, merging, and filtering would enable more efficient workflows for experiment and phenomenology users alike. This character of “utility development” is unlikely to occur without paid, dedicated effort, as there is little academic reward for such contributions.

Smaller graphs which are physics subsets (size smaller by a factor of 10) could offer significant benefits to users both in storage and CPU time. A discussion on defining such subsets and ensuring their implementation in parton-shower event generations should take place, via the MCnet community, HSF, and/or the Les Houches workshops.

In terms of I/O formats, HepMC3 supports (zipped) textual formats from v2 and v3, and ROOT. Read and write support for further formats can be easily added, if other formats are necessary (e.g. from the heavy-ion community as it moves to increasingly use LHC standard tools in the EIC and elsewhere. Format improvements are possible, but must be undertaken judiciously to avoid breaking read-compatibility with previous versions: this can be handled explicitly in the custom formats, but schema evolution is more complicated with binary formats such as ROOT and HDF5.

Handling parton-level events is also important. The LHE format is textual and not fully standardized. It may be required to move to a more efficient format such as HDF5. HDF5 support was added by the Sherpa and Pythia8 teams, but it is also not standardized. HepMC3 can embed XML LHE, however is it not clear how ME and showering and hadronization event generators (SHG) event-record should be linked for HDF5. None of these required developments are institutionally supported. This is fully a technical question that has no career incentive but it requires a good degree of software design and engineering experience. In Pythia 8.304, a new versioning scheme for HDF5 event samples exist, LHAHDF5 [80]. This is based on the structure defined in Ref. [81].

We need better graph algorithms, copying, manipulation tools and small event-graph options from MCs.

Rivet. Rivet [82] is an analysis-preservation tool, used in MC validation, development, tuning, and BSM reinterpretation. Rivet v3 was released in June 2019, with the current version being 3.1.4. This latest release series notably provides automatic MC multi-weight handling for systematic uncertainty calculations, improved

run-merging (that allows things like ratios to be calculated correctly without any assumptions), machinery for heavy-ion observables, and the ability to pass optional parameters to analyses.

Possible development areas are statistical-object handling, HepData integration, detector folding and smearing, systematic uncertainties and bootstrap, as well as technical developments like HDF5 formats to enable speed-up, and embeddability in adaptive-sampling physics tools, e.g. for MC-tuning and BSM limit-setting.

The code efficiency was profiled in GSoC 2020, and found to be limited by HepMC in I/O, event copying for unit standardisation, and event-graph traversal. Multi-weighted events provide opportunities to cache histogram operations, which currently occupy the next largest fraction of CPU time. This and thread safety require detailed technical work, as there are few simple opportunities for efficiency improvements.

Many suitable LHC experimental results are preserved using Rivet, with semi-active preservation programs in experiments. The official requirements in the experiments are not currently strong, but they exist and are getting better. The important role of Rivet in MC development, validation, and BSM fitting requires better standards-compliance in the routines submitted to Rivet. The experiment-published analysis coverage (mostly of ATLAS and CMS for now) is tracked in a Rivet webpage [83], using publication metrics from Inspire. ALICE has initiated and strongly contributed to the RIVET for HI implementation. Their analysis contribution is ramping up and also non-LHC HI experiments have started to contribute. If analysis preservation is not foreseen from the beginning of experimental measurements, it is more likely to encounter issues with Rivet integration or downstream usage.

There is also a critical issue that the analyses submitted with reference data are sometimes incompatible with that submitted to HepData. The untangling of this incompatible portion of the Rivet data is draining the Rivet resources. Rivet no longer accepts analyses that are not HepData-compatible, therefore, this won't be a problem in the long term.

Rivet development was greatly impacted by the Covid19 pandemic, with fewer patch releases than usual and no development toward the planned major updates. Much of the development effort needed is in areas involving complex statistical tools, e.g., to handle event weights.

YODA, Professor, MCUtils/HEPUtils, HepData. These are other, “small” but important, tools that are even less supported than those described above. This is a work it is hard to get rewarded on, in particular it is difficult to write papers on these tools.

YODA (<https://yoda.hepforge.org>) is the statistics library underpinning Rivet. It is lightweight and focused on the issues encountered in MC event-analysis and analysis reinterpretation, and it is easy to integrate in multi-threaded applications because it has no global state. It is being extended for generic-object binned contain-

ers, coherent many-dimensional histograms, correlations, and to have a new binary format. However, there is again no funding support or dedicated person-power.

Professor [84] is an MC-tuning and general parametrization and optimization tool used to make many of the main LHC MC-generator tunes. It has been rewritten as Apprentice [85] with added support for rational approximants, better handling of correlations, and other improvements. Again, it is unsupported, although it has started to be used in MC teams of LHC experiments at least of ATLAS and CMS.

MCUtils/HEPUtils are general tools to help with MC analysis, including HepMC graph filtering and PDG ID decoding. HepMC3 development may be able to integrate the search and graph-reduction algorithm into the HepMC tool suite.

HepData (<https://www.hepdata.net>) is another important service that is not strictly software run by the experiments, but that holds an important place to store and preserve data used as input to tools such as Rivet. We note that continued support of HepData is also essential for the event generators ecosystem. Operation-wise, it is supported but there is little development. LHC data-reinterpretation studies make heavy use of HepData and have demands which would be best solved by feature development in the database system, particularly around provision of correlation information, semantic metadata, and preservation of theory predictions.

4 Opportunities for common projects

Proposal for modular MC event generator framework. There was appetite from a number, but not all, of the generator groups to create a modular framework with well-defined software interfaces where different teams can plug in different ME generators or different parton showers, even different shower algorithms and different hadronisation models, etc.

There are several places where such an interface would not be a giant leap from what is already in place. For example: Sherpa is able to load Pythia8 libraries for hadronisation and hadron decays and choose between BlackHat [86, 87], MadLoop [88], MCFM, OpenLoops, and Recola [89–92] for the loop ME provider; Herwig7 is working on a string hadronisation model based upon an old version of Pythia; and Powheg already has an interface to MG5aMC. Pythia8 has a new generic ME interface framework, currently able to load C++ ME libraries generated with MG5aMC, but currently being extended by a runtime interface to the Comix ME generator in Sherpa and the new C++ interface [50] to MCFM. This was also the initial idea behind ThePEG, which Herwig7 continues to use, but as noted significant redevelopment would be needed to make this more widely useful.

Such a framework could be very beneficial from both a physics and software point of view. For physics studies it would make the estimation of uncertainties in a factorised way much simpler. From a software point of view it would allow greater possibilities to improve and/or port code when a) the workflow is clearer to

outsiders and b) there is a common structure. In order for this to happen, HSF could help identify and agree on some interfaces, but such a project would likely only be successful with support from the vast majority of generator groups.

More generally, taking the time to redesign some internal and external software interfaces of the generator packages would also be a good opportunity to improve the readiness of the code for modern software paradigms and hardware architectures. During the ongoing reengineering of the ME calculation in MG5aMC [35], for instance, it has become clear that the redefinition of some interfaces to deal with multiple events in parallel (rather than one event at a time), as well as a clearer encapsulation of the inputs and outputs of each software block (avoiding possible pitfalls from Fortran common blocks), are required in order to exploit event-level data parallelism on GPUs and vector CPUs. This would also simplify the implementation of task parallelism via multi-threaded approaches.

Help in the effort of porting MC generators to GPUs. Experiments have GPU resources where they can test new MC workflows and validate them with running large scale validations, and can identify technical difficulties. HSF can help MC authors and experiments identifying, refactoring, profiling and reworking that could be useful as a common work to make ME calculations with GPUs. HSF, MG5aMC authors, and CERN IT are working together on GPUs for MG5aMC. This could be generalized to include other MC generators and HSF can also help in the coordination with experiments as well.

Basic profiling and improving computing performance. Generators may benefit from HSF support on issues that increase the common good in the generator world. Basic profiling for relevant use cases would be beneficial for improving computing performance. Thanks to a large effort from the generator teams in both experiments and the WG, a lot of insight into the settings used to support each experiment's physics program was gained. A more detailed study of the different strategies is ongoing, in particular by analysing individually the CPU costs using standard benchmark suites such as HEP-SPEC06 (HS06). The work on establishing better mechanisms to calculate CPU costs using standard benchmarks from the production systems of experiments is ongoing. This will allow for more fair and easier comparison between different experiments and help identify bottlenecks in production. For profiling, another important development would be to have the MC generators report on sampling, merging, filtering and other relevant efficiencies.

For performance profiling and comparisons, it would be useful to have a standalone reproducible common setup for ATLAS and CMS. This is being worked on for the $t\bar{t}$ final state by the LHC Top WG, in collaboration with the HSF Physics Generators WG.

5 Conclusions and Outlook

In this note, we have discussed the progress and plans for some of the main event generators and tools used by the LHC experiments, taking the perspective of considering them as common software projects that are essential to the success of the HL-LHC programme. This complements the recent published paper [4], also prepared in the context of the LHCC review of HL-LHC computing, where our WG had provided a comprehensive discussion of the challenges in event generator software, and which we are also submitting to the November 2021 review as an integral part of our contribution.

The ongoing and planned developments and the main issues for the various generators have been presented. It is found that funding and career opportunities are still problematic. There is progress in the problematic area of events with negative weights, however new higher order predictions are expected to yield an increased level of negative weights: therefore, negative weights will continue to remain an issue and require even more developments. There is rapid progress in the porting of matrix element calculations in MG5aMC to GPUs thanks to a collaboration between MC authors and software engineers, which has been promoted by the HSF generator WG. There are some developments in improving phase-space sampling for efficient unweighted event generation and more are expected to come. The generator and tools authors are now paying much more attention to CPU efficiency and major gains have already been made to improve their software. Possible improvements are also discussed in event reweighting, which would reduce the load on the MC production systems of the experiments. Improvements in benchmarking and profiling in the experiments have also been discussed. Some opportunities for common projects have been highlighted. Other known sources of inefficiencies in the experiment workflows, such as those due to complicated filtering needs, still need to be worked on to identify possible ways to mitigate them with new developments in generators and tools. The increase in computational costs for MC generators in the HL-LHC era due to the need for more precise NNLO calculations, which we started to discuss in our WG [4, 93], also needs to be carefully assessed.

Finally, we note that defining a specific roadmap for R&D in the generator area for the next few years is a very complex and long iterative process, for which we believe that we are now only at the first step. As we already discussed extensively in our published WG paper, this is particularly difficult as this work mainly needs to be done by several independent teams of theorists, who are in most cases not directly funded to contribute to the HL-LHC experimental programme. An important next step is to communicate with the generator/tools teams to obtain more specific prioritisation of the different development streams, the person effort required to achieve them and the performance gains expected. Following this, further iterations with the experiments will be needed to try to ensure that these aims are aligned with their priorities.

References

- [1] LHC Experiments Committee (LHCC), *HL-LHC Computing Review*.
<https://committees.web.cern.ch/review-hl-lhc-computing>
- [2] *Minutes of the 145th Meeting of the LHCC*, CERN-LHCC-2021-005.
<https://cds.cern.ch/record/2754112>
- [3] *Charge for Stage 2 of the HL-LHC Computing Review*, CERN-LHCC-2021-014.
<https://cds.cern.ch/record/2777575>
- [4] The HSF Physics Event Generator Working Group, *Challenges in Monte Carlo Event Generator Software for High-Luminosity LHC*, Comput. Softw. Big Sci. 5, 12 (2021).
[doi:10.1007/s41781-021-00055-1](https://doi.org/10.1007/s41781-021-00055-1)
- [5] E. Sexton-Kennedy, G. Stewart, *LHCC review of HL-LHC Computing*, HSF generator WG meeting #10 (February 2021). <https://indico.cern.ch/event/1019213>
- [6] S. Hoeche, *Sherpa report for the LHCC review*, HSF generator WG meeting #11 (April 2021). <https://indico.cern.ch/event/1030889>
- [7] M. Kreps, *EvtGen report for the LHCC review*, HSF generator WG meeting #12 (May 2021). <https://indico.cern.ch/event/1032036>
- [8] S. Mrenna, *Pythia8 report for the LHCC review*, HSF generator WG meeting #13 (May 2021). <https://indico.cern.ch/event/1036317>
- [9] O. Mattelaer, *Madgraph5_aMC@NLO report for the LHCC review*, HSF generator WG meeting #14 (May 2021). <https://indico.cern.ch/event/1032037>
- [10] S. Platzer, *Herwig7 report for the LHCC review*, HSF generator WG meeting #15 (June 2021). <https://indico.cern.ch/event/1042656>
- [11] E. Re, P. Nason, *Powheg report for the LHCC review*, HSF generator WG meeting #16 (June 2021). <https://indico.cern.ch/event/1042657>
- [12] A. Buckley, *LHAPDF/Rivet report for the LHCC review*, HSF generator WG meeting #16 (June 2021). <https://indico.cern.ch/event/1042657>
- [13] HSF Meeting Minutes.
<https://hepsoftwarefoundation.org/organization/minutes.html>
- [14] R. Brun and F. Rademakers, *ROOT — An object oriented data analysis framework*, Nucl. Instr. Meth. A **389** (1997) 81. [doi:10.1016/S0168-9002\(97\)00048-X](https://doi.org/10.1016/S0168-9002(97)00048-X)
- [15] S. Agostinelli et al., *Geant4 — a simulation toolkit*, Nucl. Instr. Meth. A **506** (2003) 250. [doi:10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8)

- [16] HEP Software Community Meeting on GeantV R&D (October 2016).
<https://indico.cern.ch/event/570876>
- [17] MCnet — a European Union funded Marie Curie Initial Training Network.
<https://www.montecarlonet.org/about>
- [18] High Energy Physics Center for Computational Excellence.
<https://www.anl.gov/hep-cce>
- [19] A. Buckley and M. Schoenherr, *SWIFT-HEP Generators area & projects*, SWIFT-HEP and EXCALIBUR-HEP Workshop (January 2021).
<https://indico.cern.ch/event/976081/#3-event-generators>
- [20] M. Cacciari, G. P. Salam, G. Soyez, *FastJet user manual*, Eur. Phys. J. C **72** (2012) 1896. doi:10.1140/epjc/s10052-012-1896-2
- [21] M. Grazzini, S. Kallweit and M. Wiesemann, *Fully differential NNLO computations with MATRIX*, Eur. Phys. J. C **78** (2018) 537. doi:10.1140/epjc/s10052-018-5771-7 [arXiv:1711.06631]
- [22] S. Alioli, C. W. Bauer, C. J. Berggren, A. Hornig, F. J. Tackmann, C. K. Vermilion, J. R. Walsh and S. Zuberi, *Combining Higher-Order Resummation with Multiple NLO Calculations and Parton Showers in GENEVA*, JHEP09(2013)120. doi:10.1007/JHEP09(2013)120 [arXiv:1211.7049].
- [23] S. Alioli, C. W. Bauer, C. Berggren, F. J. Tackmann and J. R. Walsh, *Drell-Yan production at NNLL'+NNLO matched to parton showers*, Phys. Rev. D **92** (2015) 094020. doi:10.1103/PhysRevD.92.094020 [arXiv:1508.01475].
- [24] M. De Angelis, J. R. Forshaw and S. Plätzer, *Resummation and Simulation of Soft Gluon Effects beyond Leading Color*, Phys. Rev. Lett. **126** (2021) 112001. doi:10.1103/PhysRevLett.126.112001 [arXiv:2007.09648].
- [25] S. Porteboeuf, T. Pierog and K. Werner, *Producing Hard Processes Regarding the Complete Event: The EPOS Event Generator*, Proc. Moriond 2010. [arXiv:1006.2967]
- [26] G. Bíró, G. Papp, G. G. Barnaföldi, D. Nagy, M. Gyulassy, P. Lévai, X.-N. Wang and B.-W. Zhang, *HIJING++, a Heavy-Ion Jet Interaction Generator for the High-luminosity Era of the LHC and Beyond*, Proc. Hot Quarks 2018. [arXiv:1811.02131]
- [27] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli and M. Zaro, *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, JHEP07(2014)079. doi:10.1007/JHEP07(2014)079 [arXiv:1405.0301]

- [28] C. Degrande, C. Duhr, B. Fuks, D. Grellscheid, O. Mattelaer and T. Reiter, *UFO — The Universal FeynRules Output*, Comput. Phys. Commun. **183** (2012) 1201-1214. doi:10.1016/j.cpc.2012.01.022 [arXiv:1108.2040]
- [29] R. Frederix, S. Frixione, V. Hirschi, D. Pagani, H. S. Shao and M. Zaro, *The automation of next-to-leading order electroweak calculations*, JHEP07(2018)185. doi:10.1007/JHEP07(2018)185 [arXiv:1804.10017]
- [30] MadGraph5_aMC@NLO web site. <https://launchpad.net/mg5amcnlo>
- [31] O. Mattelaer and K. Ostrolenk, *Speeding up MadGraph5_aMC@NLO*, Eur. Phys. J. C **81** (2021) 435. doi:10.1140/epjc/s10052-021-09204-7 [arXiv:2102.00773]
- [32] R. Frederix, S. Frixione, S. Prestel and P. Torrielli, *On the reduction of negative weights in MC@NLO-type matching procedures*, JHEP07(2020)238. doi:10.1007/JHEP07(2020)238 [arXiv:2002.12716]
- [33] S. Frixione and B. R. Webber, *Matching NLO QCD computations and parton shower simulations*, JHEP06(2002)029. doi:10.1088/1126-6708/2002/06/029 [arXiv:0204244].
- [34] S. Carrazza, J. M. Cruz-Martinez, *VegasFlow: accelerating Monte Carlo simulation across multiple hardware platforms*, Comp. Phys. Comm. **254** (2020) 107376. doi:10.1016/j.cpc.2020.107376
- [35] A. Valassi, S. Roiser, O. Mattelaer, S. Hageboeck, *Design and engineering of a simplified workflow execution for the MG5aMC event generator on GPUs and vector CPUs*, Proc. vCHEP2021, EPJ Web of Conf. **251** (2021) 03045. doi:10.1051/epjconf/202125103045 [arXiv:2106.12631]
- [36] P. Nason, *A New method for combining NLO QCD with shower Monte Carlo algorithms*, JHEP11(2004)040. doi:10.1088/1126-6708/2004/11/040 [arXiv:hep-ph/0409146]
- [37] S. Frixione, P. Nason and C. Oleari, *Matching NLO QCD computations with Parton Shower simulations: the POWHEG method*, JHEP11(2007)070. doi:10.1088/1126-6708/2007/11/070 [arXiv:0709.2092]
- [38] S. Alioli, P. Nason, C. Oleari and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, JHEP06(2010)043. doi:10.1007/JHEP06(2010)043 [arXiv:1002.2581]
- [39] F. Buccioni, J. N. Lang, J. M. Lindert, P. Maierhöfer, S. Pozzorini, H. Zhang and M. F. Zoller, *OpenLoops 2*, Eur. Phys. J. C **79** (2019) 866. doi:10.1140/epjc/s10052-019-7306-2 [arXiv:1907.13071]
- [40] G. Cullen, H. van Deurzen, N. Greiner, G. Heinrich, G. Luisoni, P. Mastrolia, E. Mirabella, G. Ossola, T. Peraro and J. Schlenk, *et al. GOSAM-2.0: a tool for*

- automated one-loop calculations within the Standard Model and beyond*, Eur. Phys. J. C **74** (2014) 3001. doi:10.1140/epjc/s10052-014-3001-5 [arXiv:1404.7096]
- [41] G. Cullen, N. Greiner, G. Heinrich, G. Luisoni, P. Mastrolia, G. Ossola, T. Reiter and F. Tramontano, *Automated One-Loop Calculations with GoSam*, Eur. Phys. J. C **72** (2012) 1889. doi:10.1140/epjc/s10052-012-1889-1 [arXiv:1111.2034]
- [42] P. F. Monni, P. Nason, E. Re, M. Wiesemann and G. Zanderighi, *MiNNLO_{PS}: a new method to match NNLO QCD to parton showers*, JHEP05(2020)143. doi:10.1007/JHEP05(2020)143 [arXiv:1908.06987]
- [43] P. F. Monni, E. Re and M. Wiesemann, *MiNNLO_{PS}: optimizing $2 \rightarrow 1$ hadronic processes*, Eur. Phys. J. C **80** (2020) 1075. doi:10.1140/epjc/s10052-020-08658-5 [arXiv:2006.04133]
- [44] M. Dasgupta, F. A. Dreyer, K. Hamilton, P. F. Monni, G. P. Salam and G. Soyez, *Parton showers beyond leading logarithmic accuracy*, Phys. Rev. Lett. **125** (2020) 052002. doi:10.1103/PhysRevLett.125.052002 [arXiv:2002.11114]
- [45] J. R. Andersen, C. Gütschow, A. Maier and S. Prestel, *A Positive Resampler for Monte Carlo events with negative weights*, Eur. Phys. J. C **80** (2020) 1007. doi:10.1140/epjc/s10052-020-08548-w [arXiv:2005.09375]
- [46] J. R. Andersen and A. Maier, *Unbiased Elimination of Negative Weights in Monte Carlo Samples*, MCNET-21-14 (2021). [arXiv:2109.07851]
- [47] E. Bothmann *et al.* [Sherpa], *Event Generation with Sherpa 2.2*, SciPost Phys. **7** (2019) 034. doi:10.21468/SciPostPhys.7.3.034 [arXiv:1905.09127]
- [48] K. Danziger, *Efficiency Improvements in Monte Carlo Algorithms for High-Multiplicity Processes*, Master Thesis (2020). <https://cds.cern.ch/record/2715727>
- [49] J. Campbell and T. Neumann, *Precision Phenomenology with MCFM*, JHEP12(2019)034. doi:10.1007/JHEP12(2019)034 [arXiv:1909.09117]
- [50] J. M. Campbell, S. Höche and C. T. Preuss, *Accelerating LHC phenomenology with analytic one-loop amplitudes: 1 C++ interface to MCFM*. [arXiv:2107.04472]
- [51] E. Bothmann, T. Janßen, M. Knobbe, T. Schmale and S. Schumann, *Exploring phase-space with Neural Importance Sampling*, SciPost Phys. **8** (2020) 069. doi:10.21468/SciPostPhys.8.4.069 [arXiv:2001.05478]
- [52] C. Gao, J. Isaacson and C. Krause, *i-flow: High-dimensional Integration and Sampling with Normalizing Flows*, Mach. Learn. Sci. Tech. **1** (2020) 045023. doi:10.1088/2632-2153/abab62 [arXiv:2001.05486]

- [53] K. Danziger, T. Janßen, S. Schumann and F. Siegert, *Accelerating Monte Carlo event generation – rejection sampling using neural network event-weight estimates*, MCNET-21-13 (2021). [[arXiv:2109.11964](#)]
- [54] E. Bothmann, W. Giele, S. Hoeche, J. Isaacson and M. Knobbe, *Many-gluon tree amplitudes on modern GPUs: A case study for novel event generators*. [[arXiv:2106.06507](#)]
- [55] *A posteriori scale and PDF variations using the HepMC GenEvent Output* <https://sherpa.hepforge.org/doc/SHERPA-MC-2.2.10.html#A-posteriori-scale-and-PDF-variations-using>
- [56] S. Hoeche, F. Krauss, M. Schonherr and F. Siegert, *QCD matrix elements + parton showers: The NLO case*, JHEP04(2013)027. doi:10.1007/JHEP04(2013)027 [[arXiv:1207.5030](#)]
- [57] T. Gehrmann, S. Hoche, F. Krauss, M. Schonherr and F. Siegert, *NLO QCD matrix elements + parton showers in $e^+e^- \rightarrow$ hadrons*, JHEP01(2013)144. doi:10.1007/JHEP01(2013)144 [[arXiv:1207.5030](#)]
- [58] J. Bellm, S. Gieseke, D. Grellscheid, S. Plätzer, M. Rauch, C. Reuschle, P. Richardson, P. Schichtel, M. H. Seymour and A. Siódmok, *et al. Herwig 7.0/Herwig++ 3.0 release note*, Eur. Phys. J. C **76** (2016) 196. doi:10.1140/epjc/s10052-016-4018-8 [[arXiv:1512.01178](#)]
- [59] S. Plätzer and S. Gieseke, *Dipole Showers and Automated NLO Matching in Herwig++*, Eur. Phys. J. C **72** (2012) 2187. doi:10.1140/epjc/s10052-012-2187-7 [[arXiv:1109.6256](#)]
- [60] S. Jadach, G. Nail, W. Placzek, S. Sapeta, A. Siódmok, M. Skrzypek, *KrkNLO in Herwig 7*, AcPPB 48 (2017) 1121. doi:10.5506/APhysPolB.48.1121
- [61] T. Sjostrand, S. Mrenna and P. Z. Skands, *A Brief Introduction to PYTHIA 8.1*, Comput. Phys. Commun. **178** (2008) 852-867. doi:10.1016/j.cpc.2008.01.036 [[arXiv:0710.3820](#)]
- [62] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen and P. Z. Skands, *An introduction to PYTHIA 8.2*, Comput. Phys. Commun. **191** (2015) 159-177. doi:10.1016/j.cpc.2015.01.024 [[arXiv:1410.3012](#)]
- [63] N. Fischer, S. Prestel, M. Ritzmann and P. Skands, *Vincia for Hadron Colliders*, Eur. Phys. J. C **76** (2016) 589. doi:10.1140/epjc/s10052-016-4429-6 [[arXiv:1605.06142](#)]
- [64] S. Höche and S. Prestel, *Triple collinear emissions in parton showers*, Phys. Rev. D **96** (2017) 074017. doi:10.1103/PhysRevD.96.074017 [[arXiv:1705.00742](#)]
- [65] S. Höche, F. Krauss and S. Prestel, *Implementing NLO DGLAP evolution in Parton Showers*, JHEP10(2017)093. doi:10.1007/JHEP10(2017)093 [[arXiv:1705.00982](#)]

- [66] F. Dulat, S. Höche and S. Prestel, *Leading-Color Fully Differential Two-Loop Soft Corrections to QCD Dipole Showers*, Phys. Rev. D **98** (2018) 074013. doi:10.1103/PhysRevD.98.074013 [arXiv:1805.03757]
- [67] N. Fischer and S. Prestel, *Combining states without scale hierarchies with ordered parton showers*, Eur. Phys. J. C **77** (2017) 601. doi:10.1140/epjc/s10052-017-5160-7 [arXiv:1706.06218]
- [68] H. Brooks and C. T. Preuss, *Efficient multi-jet merging with the Vincia sector shower*, Comput. Phys. Commun. **264** (2021) 107985. doi:10.1016/j.cpc.2021.107985 [arXiv:2008.09468]
- [69] L. Lönnblad, *Correcting the colour-dipole cascade model with fixed order matrix elements*, JHEP05(2002)046. doi:10.1088/1126-6708/2002/05/046
- [70] CMSSW - CMS Offline Software. <https://github.com/cms-sw/cmssw>
- [71] C. Bierlich, G. Gustafson, L. Lönnblad and H. Shah, *The Angantyr model for Heavy-Ion Collisions in PYTHIA8*, JHEP10(2018)134. doi:10.1007/JHEP10(2018)134 [arXiv:1806.10820]
- [72] D. J. Lange, *The EvtGen particle decay simulation package*, Nucl. Instrum. Meth. A **462** (2001) 152-155. doi:10.1016/S0168-9002(01)00089-4
- [73] N. Davidson, G. Nanava, T. Przedzinski, E. Richter-Was and Z. Was, *Universal Interface of TAUOLA Technical and Physics Documentation*, Comput. Phys. Commun. **183** (2012) 821-843. doi:10.1016/j.cpc.2011.12.009 [arXiv:1002.0543]
- [74] E. Barberio and Z. Was, *PHOTOS: A Universal Monte Carlo for QED radiative corrections. Version 2.0*, Comput. Phys. Commun. **79** (1994) 291-308. doi:10.1016/0010-4655(94)90074-4
- [75] K. Hamilton and P. Richardson, *Simulation of QED radiation in particle decays using the YFS formalism*, JHEP07(2006)010. doi:10.1088/1126-6708/2006/07/010 [arXiv:hep-ph/0603034]
- [76] A. Buckley, J. Ferrando, S. Lloyd, K. Nordström, B. Page, M. Rüfenacht, M. Schönherr and G. Watt, *LHAPDF6: parton density access in the LHC precision era*, Eur. Phys. J. C **75** (2015) 132. doi:10.1140/epjc/s10052-015-3318-8 [arXiv:1412.7420]
- [77] J. R. Andersen, S. Badger, L. Barze, J. Bellm, F. U. Bernlochner, A. Buckley, J. Butterworth, N. Chanon, M. Chiesa and A. Cooper-Sarkar, *et al. Les Houches 2013: Physics at TeV Colliders: Standard Model Working Group Report*. [arXiv:1405.1067]
- [78] M. Dobbs and J. B. Hansen, *The HepMC C++ Monte Carlo event record for High Energy Physics*, Comput. Phys. Commun. **134** (2001) 41-46. doi:10.1016/S0010-4655(00)00189-2

- [79] A. Buckley, P. Ilten, D. Konstantinov, L. Lönnblad, J. Monk, W. Pokorski, T. Przedzinski and A. Verbytskyi, *The HepMC3 event record library for Monte Carlo event generators*, Comput. Phys. Commun. **260** (2021) 107310. doi:10.1016/j.cpc.2020.107310 [arXiv:1912.08005].
- [80] Pythia8.3 Manual, *HDF5 Event File Format*. <https://pythia.org/manuals/pythia8306/LHAHDF5.html>
- [81] S. Höche, S. Prestel and H. Schulz, *Simulation of Vector Boson Plus Many Jet Final States at the High Luminosity LHC*, Phys. Rev. D **100** (2019) 014024. doi:10.1103/PhysRevD.100.014024 [arXiv:1905.05120].
- [82] C. Bierlich, A. Buckley, J. Butterworth, C. H. Christensen, L. Corpe, D. Grellscheid, J. F. Grosse-Oetringhaus, C. Gutsche, P. Karczmarczyk and J. Klein, *et al. Robust Independent Validation of Experiment and Theory: Rivet version 3*, SciPost Phys. **8** (2020) 026. doi:10.21468/SciPostPhys.8.2.026 [arXiv:1912.05451]
- [83] *Rivet analysis coverage*. <https://rivet.hepforge.org/rivet-coverage>.
- [84] A. Buckley, H. Hoeth, H. Lacker, H. Schulz and J. E. von Seggern, *Systematic event generator tuning for the LHC*, Eur. Phys. J. C **65** (2010) 331. doi:10.1140/epjc/s10052-009-1196-7 [arXiv:0907.2973]
- [85] M. Krishnamoorthy, H. Schulz, X. Ju, W. Wang, S. Leyffer, Z. Marshall, S. Mrenna, J. Müller and J. B. Kowalkowski, *Apprentice for Event Generator Tuning*. [arXiv:2103.05748]
- [86] C. F. Berger, Z. Bern, L. J. Dixon, F. Febres Cordero, D. Forde, H. Ita, D. A. Kosower and D. Maître, *One-Loop Calculations with BlackHat*, Nucl. Phys. B Proc. Suppl. **183** (2008) 313-319. doi:10.1016/j.nuclphysbps.2008.09.123 [arXiv:0807.3705]
- [87] Z. Bern, L. J. Dixon, F. Febres Cordero, S. Höche, H. Ita, D. A. Kosower, D. Maître and K. J. Ozeren, *The BlackHat Library for One-Loop Amplitudes*, Proc. ACAT 2013, J. Phys. Conf. Ser. **523** (2014) 012051. doi:10.1088/1742-6596/523/1/012051 [arXiv:1310.2808]
- [88] V. Hirschi, R. Frederix, S. Frixione et al., *Automation of one-loop QCD computations*, JHEP05(2011)044. doi:10.1007/JHEP05(2011)044 [arXiv:1103.0621]
- [89] S. Actis, A. Denner, L. Hofer, A. Scharf and S. Uccirati, *Recursive generation of one-loop amplitudes in the Standard Model*, JHEP04(2013)037. doi:10.1007/JHEP04(2013)037 [arXiv:1211.6316].
- [90] S. Actis, A. Denner, L. Hofer, J. N. Lang, A. Scharf and S. Uccirati, *RECOLA: REcursive Computation of One-Loop Amplitudes*, Comput. Phys. Commun. **214** (2017) 140-173. doi:10.1016/j.cpc.2017.01.004 [arXiv:1605.01090].

- [91] A. Denner, J. N. Lang and S. Uccirati, *NLO electroweak corrections in extended Higgs Sectors with RECOLA2*, JHEP07(2017)087. doi:10.1007/JHEP07(2017)087 [arXiv:1705.06053].
- [92] A. Denner, J. N. Lang and S. Uccirati, *Recola2: REcursive Computation of One-Loop Amplitudes 2*, Comput. Phys. Commun. **224** (2018) 346-361. doi:10.1016/j.cpc.2017.11.013 [arXiv:1711.07388].
- [93] M. Grazzini, *Status of NNLO ttbar calculations*, HSF generator WG meeting #9 (July 2020). <https://indico.cern.ch/event/936481>