

# Efficient precision simulation of processes with many-jet final states at the LHC

Enrico Bothmann,<sup>1</sup> Taylor Childers,<sup>2</sup> Christian Gütschow,<sup>3</sup> Stefan Höche,<sup>4</sup>  
Paul Hovland,<sup>2</sup> Joshua Isaacson,<sup>4</sup> Max Knobbe,<sup>1</sup> and Robert Latham<sup>2</sup>

<sup>1</sup>*Institut für Theoretische Physik, Georg-August-Universität Göttingen, 37073 Göttingen, Germany*

<sup>2</sup>*Argonne National Laboratory, Lemont, IL, 60439, USA*

<sup>3</sup>*Centre for Advanced Research Computing, University College London, Gower Street, London, WC1E 6BT, UK*

<sup>4</sup>*Fermi National Accelerator Laboratory, Batavia, IL 60510, USA*

We present a scalable technique for the simulation of collider events with multi-jet final states, based on an improved parton-level event file format. The method is implemented for both leading- and next-to-leading order QCD calculations. We perform a comprehensive analysis of the I/O performance and validate our new framework using Higgs-boson plus multi-jet production with up to seven jets. We make the resulting code base available for public use.

## I. INTRODUCTION

The simulation of events with high-multiplicity final-states in experiments at the Large Hadron Collider (LHC) is a challenging computational problem [1–4]. Using the best available algorithms, the calculation of the integrand for multi-jet processes scales at best exponentially with increasing particle multiplicity. The integration over the many-body phase space calls for suitable importance-sampling techniques, which often also scale exponentially [5–15]. Hence, while being a solved problem in principle, the calculation of cross sections and the production of unweighted events at high jet multiplicity is still a hard problem to date. Computing techniques have remained conceptually identical since their inception four decades ago and typically make use of dynamic programming [15–24].

While the calculation of hard cross sections with full quantum interference effects is a considerable challenge even at tree level, the Markovian methods used in parton showers are often sufficient to describe the dynamics of collider events, and in fact they are necessary to properly account for the all-orders resummation of virtual corrections [1, 25, 26]. The combination of the evolution implemented in parton showers with the exact calculations implemented by hard matrix elements provides the best available physics modeling of LHC events, accounting for both inter-jet correlations and intra-jet evolution by means of NLO matching [27–31] and multi-jet merging [32–45]. However, parton showers and the subsequent hadronization [46–51] and multiple interaction [52–54] models are associated with various free parameters. Varying these parameters is key to assessing the uncertainty of LHC simulations. The large difference in required computation time, and the need to perform the simulation for multiple parton shower, underlying event and hadronization parameters makes it natural to separate the generation of LHC events into the calculation of the hard interaction and the simulation of the remaining physics aspects.

A number of approaches have been proposed to address this problem. The earliest and most widely used include the user process functionality of Pythia [55], and the LesHouches event file format [56–58]. More recently, with the need for high-statistics event simulation leading to the use of high-performance computing facilities [59], the need for improved I/O performance and truly parallel I/O has become apparent. In this manuscript we report on the extension of a new event file format [60] for particle-level events. Firstly, we enable the storage of the standard and hard events needed for next-to-leading matching in the MC@NLO method. Secondly, we propose a new layout of the event file in order to increase the performance in large-scale parallel processing. Thirdly, we implement the new technology in various parton-level and particle-level event generators. At parton level we use Sherpa [61, 62] with the two internal matrix-element generators Amegic [63] and Comix [15], as well as the new GPU-accelerated leading-order parton-level event generator Pepper [14, 64]. At particle level, we use the event generators Pythia [55, 65] and Sherpa [61, 62].

We also provide a first physics application of our new algorithms. We simulate Higgs boson plus multi-jet events with up to seven jets at tree level, and up to two jets at next-to-leading order QCD precision. With the High-Luminosity LHC (HL-LHC) expected to collect  $3 \text{ ab}^{-1}$  of data, these predictions can be used to test QCD associated Higgs production over a large dynamic range. Moreover, Higgs-boson plus multi-jet events play an important role as irreducible backgrounds to more detailed tests of the Higgs sector of the Standard Model, and especially in weak vector boson fusion. Using the MEPS@NLO merging method [41, 42], in particular the reweighting of higher-multiplicity tree-level predictions with the help of Born-local K-factors from the Higgs plus two-jet setup, our new code base enables the most precise predictions of Higgs plus  $\geq 4$  jet events to date. We make the corresponding input event samples available for public use<sup>1</sup> and provide a publicly available version of the parton- and particle-level event generators that can be used to process these event files.

<sup>1</sup> The event samples can be obtained from <https://doi.org/10.5281/zenodo.7751000> and <https://doi.org/10.5281/zenodo.7747376>.

This manuscript is structured as follows: Section II discusses the challenges faced in previous simulations and provides a documentation and performance assessment of our new event file format and parallel event generation techniques, including the changes needed for the simulation of events at next-to-leading order QCD precision. Section III discusses the new GPU-accelerated parton-level simulation pipeline. Section IV presents the first phenomenological application and discusses the impact of NLO matching and multi-jet merging in Higgs plus multi-jet production. We conclude with an outlook in Sec. V.

## II. EXTENDED PARALLELIZATION FRAMEWORK

The first scalable particle-level event generator was based on legacy versions of Alpgen [22] and Pythia 6 [55] and was introduced in Ref. [59]. To make state-of-the-art parton-level and particle-level simulation tools available for use on HPC systems, Ref. [60] proposed a new event generation framework, based on Sherpa 2 [62] and Pythia 8 [65]. This framework is based on a parallelized main routine for Pythia 8, and a new parton-level event file format, using the HDF5 library for parallel I/O.<sup>2</sup> It solved the main problem of making the event production scalable to thousands of MPI ranks, but still suffers from an I/O bottleneck. Various modern high-performance computing systems are not well suited for the fast writing and reading of large amounts of data, as is common in parton-level event simulations. The solution to this problem is discussed in this section.

Reference [60] also did not provide a means to store information for parton-level events simulated at NLO QCD precision. At present, a pure leading-order based event simulation falls short of the precision requirements at the LHC experiments. Therefore, an extension of the previous event file format to NLO QCD precision is an additional problem we will address. At leading order QCD, one can write the expectation value of an arbitrary infrared safe observable,  $O$ , at particle level as

$$\langle O \rangle = \int d\Phi_B B(\Phi_B) \mathcal{F}_{\text{MC}}(O, \Phi_B), \quad (1)$$

where  $B$  is the differential Born cross section, including flux and symmetry factors, as well as the parton luminosity, and  $d\Phi_B$  is the differential phase-space element, including the integration over the light-cone momentum fractions of the initial-state partons. The functional  $\mathcal{F}_{\text{MC}}(O, \Phi_B)$  implements the parton shower and is explained in more detail in [27, 30]. In the MC@NLO [27] or POWHEG [28] NLO QCD matching technique, Eq. (1) becomes

$$\begin{aligned} \langle O \rangle = & \int d\Phi_B \left[ B(\Phi_B) + V(\Phi_B) + I(\Phi_B) + \int dz_1 dz_2 \text{KP}(\Phi_B, z_1, z_2) \right. \\ & \left. + \sum_{ijk} \int d\Phi_{+1,ijk} \left( D_{ijk}^{(A)}(\Phi_B, \Phi_{+1,ijk}) - D_{ijk}^{(S)}(\Phi_B, \Phi_{+1,ijk}) \right) \right] \mathcal{F}_{\text{MC}}(O, \Phi_B) \\ & + \int d\Phi_R \left( R(\Phi_R) - \sum_{ijk} S(\Phi_B, ijk(\Phi_R)) \right) \mathcal{F}_{\text{MC}}(O, \Phi_R), \end{aligned} \quad (2)$$

where  $V$  and  $R$  are the virtual and real-emission corrections,  $S$  and  $I$  are the differential and integrated NLO infrared subtraction counterterms, and  $\text{KP}$  are the factorization scale dependent finite corrections arising from the combination of the integrated infrared subtraction counterterms with collinear mass-factorization counterterms. The infrared subtraction counterterms are most commonly defined in the Frixione-Kunszt-Signer [66] or the Catani-Seymour [67, 68] subtraction scheme and depend on the momenta and flavors of two partons  $i$  and  $j$  that are to be combined, as well as a spectator or recoil momentum,  $k$ . The differential phase-space element for the real-emission process is given by  $d\Phi_R$ . It can be factorized into a differential Born phase-space element and a single-emission phase-space element as  $d\Phi_R = d\Phi_{B,ijk} d\Phi_{+1,ijk}$  [30].

The NLO QCD expression, Eq. (2), shows a number of important differences compared to the LO expression, Eq. (1):

- There are two different classes of events, one with parton-shower starting condition  $\Phi_B$  (the so-called standard events, or  $\mathbb{S}$ -events), and one with parton-shower starting condition  $\Phi_R$  (the so-called hard events, or  $\mathbb{H}$ -events).
- The  $\mathbb{H}$  events require a simple, leading-order like phase-space generator for  $d\Phi_R$ , which implies that the event file format is the same as at leading order.

<sup>2</sup> The source code can be found at <https://gitlab.com/hpcgen/> and <https://gitlab.com/sherpa-team/sherpa/-/tree/rel-2-3-0>.

Name	Data type	Contents
VERSION	$3 \times \text{int}$	Version ID
INIT	$10 \times \text{double}$	beamA, beamB, energyA, energyB, PDFgroupA, PDFgroupB, PDFsetA, PDFsetB, weightingStrategy, numProcesses
PROCINFO	$6 \times \text{double}$	procId, npLO, npNLO, xSection, error, unitWeight
EVENTS	$9 \times \text{double}$	pid, nparticles, start, trials, scale, fscale, rscale, aqed, aqcd
PARTICLES	$9 \times \text{double}$	id, status, mother1, mother2, color1, color2, px, py, pz, e, m, lifetime, spin
CTEVENTS	$9 \times \text{double}$	ijt, kt, i, j, k, z1, z2, bbpsw, tlpsw
CTPARTICLES	$4 \times \text{double}$	px, py, pz, e

TABLE I. Data sets in the LHEH5 event format.

- The  $\mathbb{S}$  events require two additional integrations, one for the KP counterterms, and one for the one-emission phase space  $d\Phi_{+1,ijk}$ . They also require a sampling of the indices  $i$ ,  $j$  and  $k$ .

In order to make MC@NLO  $\mathbb{S}$ -events reproducible and enable a reweighting of events to arbitrary PDF sets and/or scales at NLO QCD, the event format therefore requires the storage of the indices  $i$ ,  $j$  and  $k$ , as well as the phase-space point  $(\Phi_B, \Phi_{+1,ijk})$  and the momentum fractions  $z_1$  and  $z_2$ . We will define the corresponding structures in Sec. II A.

### A. Event file format

In this subsection we describe the new event file layout, which includes the optimizations that will be described in Sec. II B as well as extensions for event simulation at NLO QCD. Both are inspired by Les Houches Event Files (LHEF) standard [56], which is widely used in the high-energy physics community. It is based on XML, which makes it flexible enough to add any desired feature, but poses a challenge for I/O operations at scale. In contrast the HDF5 format uses a computing model similar to databases, making it rigid, but highly efficient in parallel workflows.

The LHEF format comprises global properties as well as event-wise properties [56]. Global properties include process information (i.e. the type of collisions), total cross-sections as well as reweighting information. The event-wise properties are the process ID, the event weight, the scale of the hard process as well as the values of  $\alpha_{\text{QCD}}$ ,  $\alpha_{\text{QED}}$  and the list of particles generated. The latter contain information about momentum four-vectors, particle ID, charge, spin and lifetime, as well as production history. We collect this information in consolidated datasets, reflecting the global, event-wise and particle properties. In addition, we introduce two new datasets, which include the event-wise and particle-wise information needed to simulate NLO QCD events in the MC@NLO matching scheme. We will call this structure the LHEH5 event file format.<sup>3</sup>

We follow the naming scheme of Ref. [60] and define datasets called INIT and PROCINFO that are used to store basic information about the entirety of events contained in the file. We also add a new dataset, VERSION that identifies the version of the event file format. Event-wise properties for leading-order events are stored in the dataset EVENTS and for MC@NLO  $\mathbb{S}$ -events in the dataset CTEVENTS. Particle-wise properties for leading-order events are stored in the dataset PARTICLES and for MC@NLO  $\mathbb{S}$  events in the dataset CTPARTICLES.

Each dataset is a two-dimensional array and has an HDF5 attribute PROPERTIES that identifies the individual columns of the dataset in order of appearance. For example the PROCINFO dataset has the properties PROCID, NPLO, NPNLO, XSECTION, ERROR and UNITWEIGHT. In future updates of the event file format, these attributes can be used to communicate the content of the individual entries to the user of the file, similar, although not quite as flexible as in the case of XML-based Les Houches event files. The content of all datasets is summarized in Tab. I. In addition to Ref. [60], we introduce the following entries:

- Process properties NPLO and NPNLO. If the process is computed at leading order QCD, we set NPLO to the final-state particle multiplicity. If the process is computed at next-to-leading order QCD, we instead set NPNLO to the final-state particle multiplicity at Born level.

<sup>3</sup> Together with this publication, we provide a set of simple tools to parse event files written in the new format, to merge two event files, and to filter event files for overweight and zero weight events. The source code can be found at <https://gitlab.com/hpcgen/tools>.

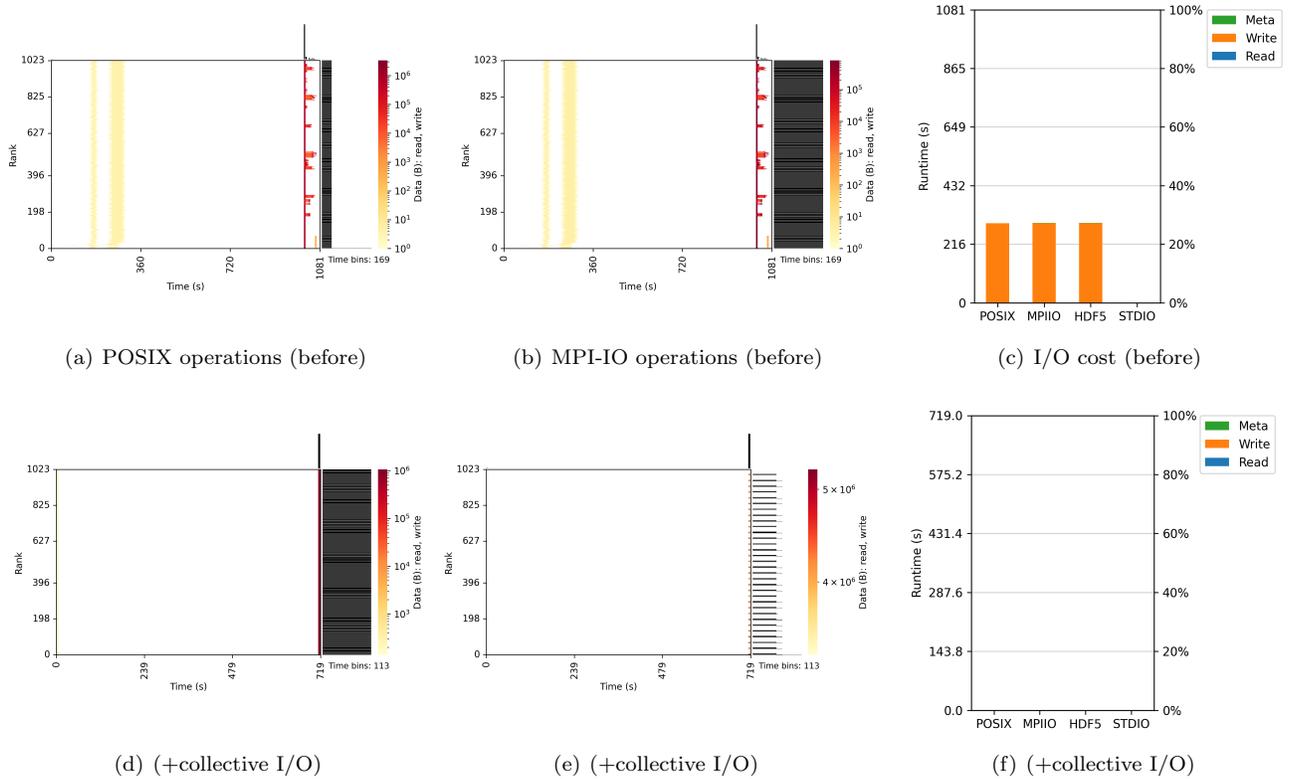


FIG. 1. Darshan graphs depicting the I/O behavior of Sherpa during parton-level event generation for H+3 jets at leading order QCD on 1024 MPI ranks. Left: POSIX I/O behavior. Middle: MPI-IO behavior. Right: Overall I/O cost. Top row: Before optimizations. I/O operations are fragmented and uncoordinated among processes. Bottom row: Including collective I/O, improved file layout to reduce metadata operations and limiting stat calls to the master rank. The number of events was held constant for this test, and the total amount of data written was 1.01 GiB.

- At NLO QCD, we include the minimal information needed to reconstruct the complete event weight in the MC@NLO matching method. For hard remainder events (H-events), the leading-order type information is sufficient. For standard events (S-events) we add the following:
  - Counterterm properties in the CTEVENTS dataset: IJT and KT refer to the Born-level QCD dipole used to generate a real-emission phase-space point in Eq. (2), I, J and K correspond to the respective particle IDs at real-emission level. TLPSW is the phase-space weight  $d\Phi_B$  and BBPSW is the corresponding single-emission phase-space weight  $d\Phi_{+1,ijk}$ . The variables z1 and z2 are the MC points of the integration variables in the KP contribution. See Eq. (2) for details.
  - Counterterms properties in the CTPARTICLES dataset: PX, PY, PZ and E store the momenta of all particles in the phase-space point  $(\Phi_B, \Phi_{+1,ijk})$ . See Eq. (2) for details.

## B. I/O operations at scale

Our event simulation frameworks read and write event data through a multi-layered I/O software stack based on HDF5 [69], an array-oriented library and data model, which in turn uses MPI-IO [70]. Within Sherpa and Pythia, HDF5 is accessed through the HighFive header library [71]. Each layer typically provides tuning parameters. Optimal performance can in principle be achieved with the help of sophisticated I/O tuning systems [72]. However, selecting the right parameters with the help of subject expertise is often more efficient and more reliable. In our case, some straightforward changes to the I/O layer resulted in large performance gains, reducing the I/O for particle-level events to a very small fraction of the overall runtime.

Figure 1 shows profiling results obtained with the help of Darshan [73, 74] for a parton-level simulation of Higgs

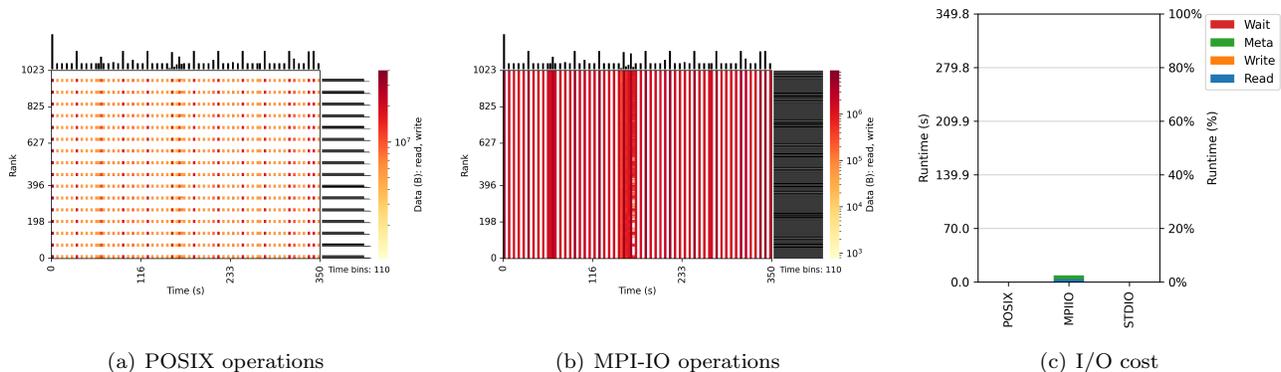


FIG. 2. Darshan graphs depicting the I/O behavior of Sherpa during particle-level event generation for H+4 jet leading order multi-jet merging on 1024 MPI ranks. Left: POSIX I/O behavior. Middle: MPI-IO behavior. Right: Overall I/O cost. All figures after including collective I/O, improved file layout to reduce metadata operations and limiting stat calls to the master rank. The number of events was held constant for this test, and the total amount of data read was 128.85 GiB.

plus four jets at leading order, run on 1024 ranks of the Cori system at NERSC.<sup>4</sup> Here we make use of a new feature of Darshan, depicting which ranks perform I/O at which times. The color in the heatmap represents the transferred data volume. Time runs along the x axis, MPI ranks along the y axis. A histogram along the x axis reports the total data volume over time. A histogram along the y axis reports the total data volume per rank. The top panels depict the initial performance of our simulation framework for parton-level event production [60]. The left and middle panels show the data transfer at the POSIX and MPI-IO level. The right panels show the fraction of runtime consumed by the I/O operations. As described in [60], the near lockstep event production implies that I/O also occurs in locksteps, making it a good candidate for collective operations. The top panels in Fig. 1 show that the output operations start at the same time on the various ranks, but they end at very different times, being responsible for a large variation in runtime overall. The bottom panels in the same figure show the I/O with collective operations turned on, for which we make use of an updated HighFive library, exposing HDF5’s collective data and collective metadata features. At the POSIX level, there are now very few ranks participating in actual write operations, and the start and end times are nearly identical on all ranks. An additional timing improvement is obtained from consolidating the individual HDF5 data sets proposed in [60] into a single data set. We also limit stat calls performed by the program to the master rank and broadcast the results of the call via MPI. The overall I/O time is reduced to a negligible amount, below 1s per rank. This concludes our optimization of the parton-level simulation.

Figure 2 shows profiling results obtained with the help of Darshan [73, 74] for a particle-level simulation of Higgs plus four jets with leading-order multi-jet merging, run on 1024 ranks of the Perlmutter system at NERSC.<sup>5</sup> For this test, we used the CPU-only nodes and did not access the scratch file system in order to give a more reliable estimate of the expected I/O time on typical computing clusters and HPC machines. The total amount of data read during the test was 128.85 GiB, and the time spent in I/O operations was less than 5% of the runtime. The POSIX-level data rate was 103.44 GiB/s and the MPI-IO level data rate was 14.43 GiB/s. Figure 2 shows that the I/O operations in our improved code are spread evenly over the runtime of the simulation, leading to more file access operations, but smaller data transfers per operation. While the file system would support larger transfer rates, storing the data for processing in the simulation program would require larger RAM allocations, leading to slower overall execution times. This effect becomes particularly important at larger scales, of the order of 1000 ranks and beyond, where the aggregate time needed for heap allocation would constitute a substantial part of the total runtime and break the strong scaling. This concludes our optimization of the particle-level simulation.

Figure 3 shows the strong scaling tests for the parton-level and the particle-level component of the simulation. The test is performed on the Perlmutter system at NERSC. We begin at a scale similar to the upper end of the tested range in [60]. For parton-level calculations, we observe good scaling properties up to 8192 MPI ranks. We note that the work for this test was selected such that the minimum runtime would correspond to about 30s, below which the initialization time of the executable consumes a significant fraction of the overall runtime. In practice, one would rather

<sup>4</sup> Cori was a Cray XC40 system, comprised of 2388 Intel Xeon “Haswell” processor nodes, 9688 Intel Xeon Phi “Knight’s Landing” (KNL) nodes and a Cray Aries network with Dragonfly topology with  $> 45\text{TB/s}$  global peak bisection bandwidth. <https://nersc.gov/systems/cori>

<sup>5</sup> Perlmutter is a Cray Shasta system, using AMD “Milan” EPYC CPUs, a novel HPE Slingshot high-speed network, and a 35-petabyte FLASH scratch file system. In total, it is comprised of 3,072 CPU-only and 1,792 GPU-accelerated nodes. <https://nersc.gov/systems/perlmutter>

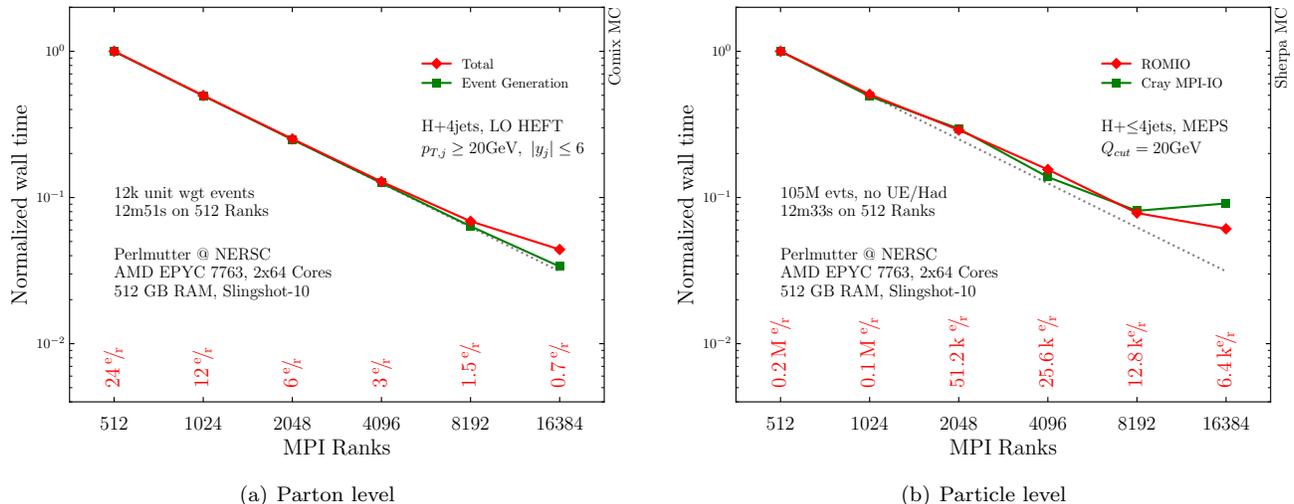


FIG. 3. Strong scaling test of the simulations. Left: Parton level. Right: Particle level. At particle level, we do not include the simulation of multiple interactions and hadronization, and we do not process events further. All times are normalized to the individual results obtained on 512 MPI ranks.

choose runtimes that are significantly longer, in order to minimize the impact of initialization. At the particle level, we observe scaling up to about 1024 MPI ranks, above which the behavior depends on the particular implementation of MPI-IO. Figure 3 compares the Cray implementation to ROMIO. Above 16384 ranks, the Cray implementation suffers from a problem that prevents collective open calls through HDF5. The ROMIO implementation does allow collective open calls, but does not reach the full performance of the Cray MPI-IO library in data transfer. However, we note that at this scale only about 6400 events are processed per rank, leading to an overall runtime of about one minute for a total of 105 million events. There is no practical need to perform a calculation of this scale in less than an hour, therefore our example should be seen as a test of the absolute limits of the code. We believe that further optimization is not needed at this stage. In addition, we note that the particle-level simulation was limited to the perturbative event phases, i.e. we did not include hadronization, underlying event simulation and hadron decays. Due to the reduced event processing time in this scenario, any scaling violations observed in our test are more severe than in practical applications.

We would like to conclude this section with a seemingly obvious but practically very important remark on the limits of scalability. The aim of an efficient parallel code is to maximize the effective computation time per worker node, i.e. the time spent in useful computations between I/O operations, with the I/O contributing an insignificant fraction of the overall runtime. One of the main reasons for scaling violations to occur is that the time between I/O operations becomes too short because of the limited size of input files. This can lead to significant problems at very large scales, where the input files must then be tens or hundreds of Gigabytes in size. Therefore, it is not practical for us to attempt scaling tests for particle-level simulations that go beyond  $\mathcal{O}(10^4)$  MPI ranks. We note that this intermediate scale parallelism is actually advantageous, because it allows to access backfill queues at large computing centers.

### III. PARTON-LEVEL EVENT GENERATION ON GPUS

We have added support for the generation of the proposed parton-level HDF5 event files to the Pepper parton-level event generator. This recently presented generator, previously called BlockGen [64, 75], is being developed to deliver performance portability for standard-candle processes, which currently include  $V$ +jets,  $t\bar{t}$ +jets and pure jet production at tree level. This is achieved by focusing on this subset of processes, and choosing the right algorithm for such parallelization [64], for example by making use [75] of a minimal color basis [76–79]. Furthermore, parallelized execution on accelerators like GPUs is supported. Preliminary results have been reported in [64, 75] for various processes. Detailed speed comparisons e.g. to Comix and comparing GPU with CPU evaluation will be reported in a forthcoming publication that will also mark the first public release of Pepper, in combination with the Chili phase-space generator [14]. As in the Sherpa case (see Sec. II B), we added MPI-IO HDF5 support to Pepper via the HighFive header library, and enabled HDF5’s collective data and collective metadata features.

Process	Tot. unc. [%]	Sherpa (Comix)			Pepper+Chili			Speed-up
		Walltime [s]	Mem. (USS) [MB]	Eff. [%]	Walltime [s]	Mem. (USS) [MB]	Eff. [%]	
Z+0j	0.089	68	62	22	10	40	43	6.8
Z+1j	0.19	76	66	5.3	31	33	10	2.5
Z+2j	0.99	92	64	0.28	10	35	1.4	9.2
Z+3j	3.8	95	65	0.037	36	43	0.097	2.6
Z+4j	14	122	115	0.0050	71	133	0.016	1.7

TABLE II. Benchmarks for the production and HDF5 writeout of  $pp \rightarrow Z + \text{jets}$  events, comparing Sherpa’s Comix with Pepper+Chili, on a single core of an Intel(R) Core(TM) i3-8300 CPU at 3.70GHz and 8MB L3 cache. Event samples are generated with a given target for the total cross section uncertainty (“Tot. unc.”). “Speed-up” gives the walltime gain factor of Pepper+Chili vs. Sherpa (Comix). For Pepper+Chili, the lower multiplicities Z+0j and Z+1j are generated using helicity summing, while the higher ones are generated using helicity sampling, in order to achieve the best possible performance in each case.

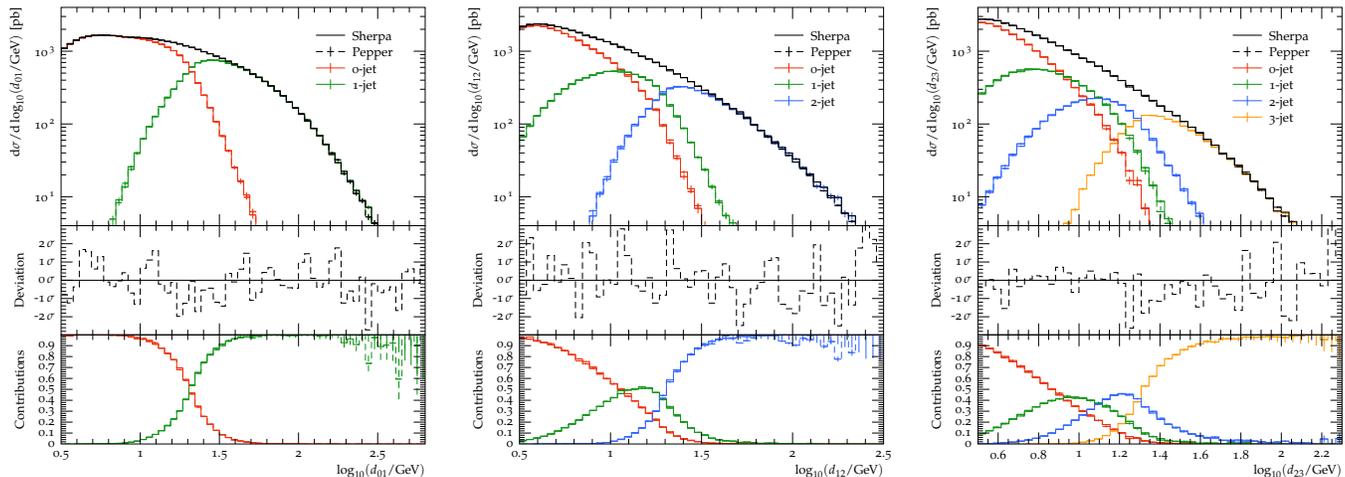


FIG. 4. Differential jet rates for the leading, sub-leading and sub-sub-leading jet clustering in  $Z + \text{jets}$  production at the LHC. Simulations have been performed with up to 1-jet, 2-jet and 3-jet matrix elements at leading order QCD. The colored lines represent the contributions from the parton-level inputs with the specified multiplicity.

Table II shows benchmarks for the production of parton-level HDF5 event files on a single CPU thread for the  $pp \rightarrow Z + n \text{ jets}$  process with  $n = 0, \dots, 4$  for a given uncertainty target of the total cross section (“Tot. unc.”), comparing Sherpa’s Comix generator with Pepper, where the latter uses Chili for the phase-space sampling. Prior to event generation, the different phase-space generators are optimized until a given accuracy target is reached to ensure a fair comparison. The benchmark metrics are the walltime for the generation of the event sample, the memory consumption in terms of the applications’ unique set size (USS) in RAM, and the fraction of the number of non-zero events over the total number of events generated, i.e. the measured combined efficiency of the phase-space sampling and unweighting (“Eff.”). For Pepper+Chili, we switch from using helicity summing for the  $n = 0, 1$  multiplicities to using helicity sampling for the  $n = 2, 3, 4$  ones, in order to achieve the best performance. We find that the walltimes are significantly lower for Pepper+Chili for the given multiplicities, with the speed-up factor ranging between 2 and 10. For the higher multiplicities,  $n = 3, 4$ , the speed-up becomes smaller, but is still significant with factors of 2.6 and 1.7, respectively. The efficiencies for Pepper+Chili are also better by factors between 2 and 5. We remind the reader that Pepper is based on explicit color sums to achieve a perfect lock-step parallel processing, which is essential for achieving excellent performance on a GPU. The computational complexity for these explicit sums scales factorially and will eventually cause Pepper to become slower (and more memory-consuming) than Comix, which uses an algorithm with overall exponential scaling. For gluon scattering, this transition occurs at seven final-state gluons [64]. This is also the reason for the larger memory usage growth going from  $n = 3$  to  $n = 4$  for Pepper+Chili compared to Sherpa.

Figure 4 shows a cross-check of differential distributions of  $k_T$  jet rates between Pepper and Comix, after leading-order multi-jet merging [38] with Sherpa 2.2 [62]. The first ratio panel compares the predictions obtained with Pepper+Sherpa to the results from Comix+Sherpa, normalized to the statistical uncertainty of the latter. The second ratio panel shows the relative contributions from the event samples with  $Z + 0$  jets,  $Z + 1$  jet,  $Z + 2$  jets and  $Z + 3$  jets

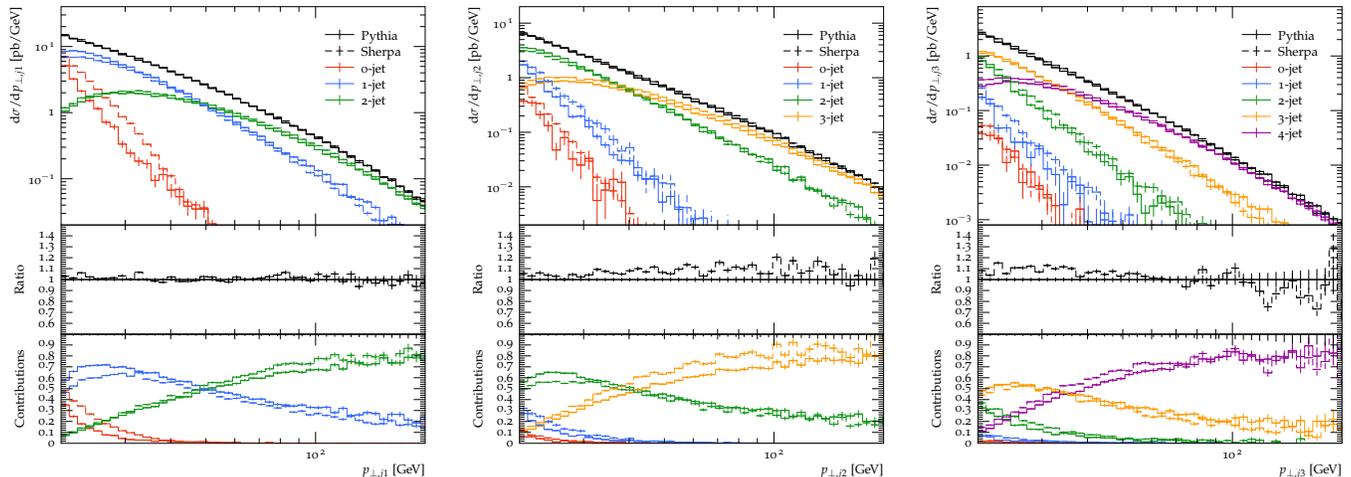


FIG. 5. Transverse momentum spectra of the leading, sub-leading and sub-sub-leading jet in  $Z$ +jets production at the LHC. Simulations have been performed with up to 2-jet, 3-jet and 4-jet matrix elements at leading order QCD. The colored lines represent the contributions from the parton-level inputs with the specified multiplicity.

to the overall prediction. It can be seen that the results are in agreement up to statistical fluctuations, which are typically at or below the  $1\sigma$  level, as expected.

#### IV. PHENOMENOLOGICAL APPLICATIONS

In this section we present first phenomenological applications of our new framework. We show that it can be used to perform simulations with two of the main particle-level simulation tools for LHC physics, Pythia 8 and Sherpa 2. We also present the first computation of Higgs with up to seven final-state jets, and we make the corresponding parton-level event samples available for public use. The source codes for our study can be found at <https://gitlab.com/hpcgen> and at <https://gitlab.com/sherpa-team/sherpa/-/tree/rel-2-3-0>.

##### A. Systematic comparison of particle-level simulations

The systematic assessment of uncertainties in particle-level simulations has been vital for the success of the LHC physics program. It is particularly important in cases where the uncertainty is not of parametric type, such as when switching between two formally equivalent, but practically different, NLO QCD matching schemes [30]. When used correctly, the residual variations between event generator predictions give the best possible non-parametric estimate of perturbative (and non-perturbative) uncertainties in the simulation. Our new event generation framework allows to obtain such uncertainty estimates based on the same parton-level input configurations and at minimal computational cost. This contributes to creating a sustainable computing model for high-energy physics research.

We carry out an example study of this type for the standard candle process of  $Z$ -boson plus multi-jet production. We consider proton-proton collisions at the high-luminosity LHC at  $\sqrt{s} = 14$  TeV. The complete setup has been described in [60]. In particular, we use the CT14 NNLO PDF set [80] and define the strong coupling accordingly. Our modified parton-level event generator is based on Comix [15] as included in Sherpa version 2.2.4 [61, 62]. Our modified particle-level event generators are based on Pythia 8 [65] and Sherpa 2.2 [62], including the improvements reported in [81, 82]. Jets are defined using the  $k_T$  clustering algorithm with  $R = 0.4$ ,  $p_{T,j} > 20$  GeV and  $|\eta_j| < 6$ . Following the good agreement between parton-level and particle-level results established in [83, 84], and the good agreement between fixed-order and MINLO [85] results established in [86, 87], the renormalization and factorization scales are set to  $\hat{H}'_T/2$ , where  $\hat{H}'_T = \sum_{j \in \text{jets}} p_{t,j} + \sqrt{m_U^2 + p_{T,U}^2}$ .

Figure 5 shows the transverse momentum spectra of the leading, sub-leading and sub-sub-leading jet in the simulation. The colored lines correspond to the contributions from the individual parton-level input samples after the full simulation. The upper ratio panel shows the ratio between the Sherpa and the Pythia predictions. This ratio is of the order of 10%, which can be ascribed to differences in the parton-shower algorithm used in the two different generators.

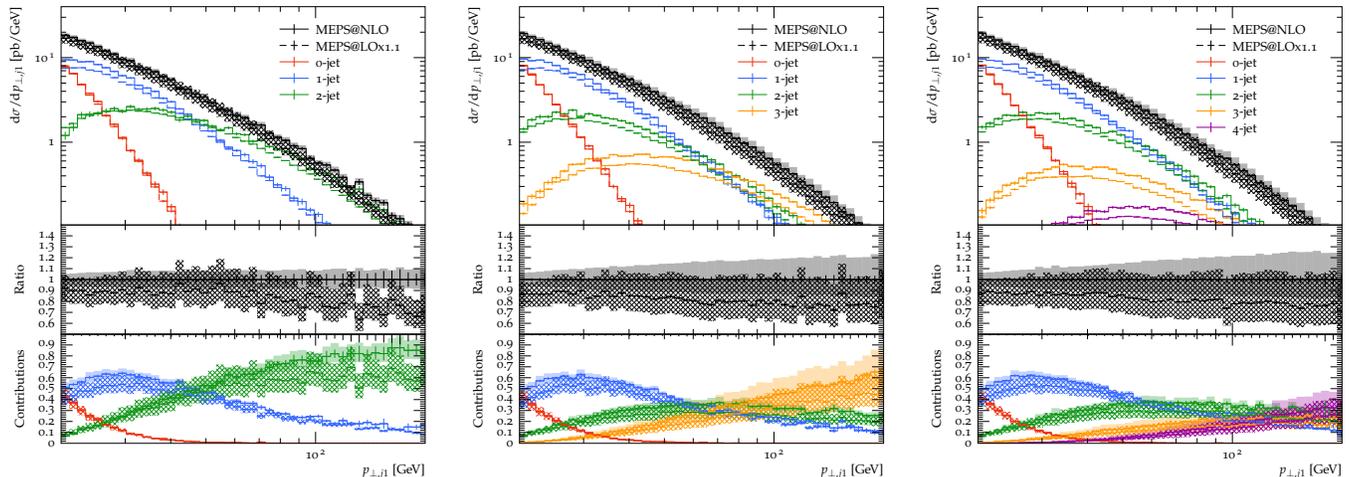


FIG. 6. Transverse momentum spectrum of the leading jet in  $Z$ +jets production at the LHC, simulated using multi-jet merging for up to two jets at NLO with up to zero, one and two additional jets at leading-order precision (from left to right), compared to a purely LO multi-jet merged prediction with the same overall multiplicities. The colored lines represent the contributions from the parton-level inputs with the specified multiplicity, and the hatched and solid bands indicate the uncertainties from renormalization and factorization scale variations at leading- and at next-to-leading order.

This uncertainty should be added as a variation to the parametric scale uncertainties, which we investigate next.

Figure 6 shows the transverse momentum spectrum of the leading jet in a multi-jet merged setup with up to two jets computed at next-to-leading order precision, and with up to zero, one and two additional jets computed at leading order precision.<sup>6</sup> For reference, we also show the prediction from a leading-order multi-jet merged event sample with identical jet multiplicity (dashed lines). The leading-order predictions have been scaled such as to reproduce the total cross section of the next-to-leading order predictions. The colored lines correspond to the contributions from the individual parton-level input samples after the full simulation. The hatched bands indicate the scale uncertainties from a seven-point scale variation at leading order, and the solid bands represent the corresponding uncertainties at next-to-leading order precision. Note that the scale uncertainties increase with increasing jet multiplicity in the merging. This is an artifact of the method to estimate the scale uncertainty in the complete calculation, and is due to the fact that scales are varied in the computation of the hard matrix elements alone. It also indicates the importance of higher-multiplicity final states for the experimental observable. To obtain a comprehensive picture of the uncertainty, the renormalization scale dependent terms of the parton-shower resummation at higher logarithmic order should be taken into account. This is the topic of active research elsewhere [88, 89], and we will therefore not discuss the effect in this publication. We emphasize, however, that the simulation of additional radiation at tree level is necessary for a proper physics modeling of high-multiplicity final states, and it is therefore not sufficient to limit the fixed-order perturbative calculations to low multiplicity. This is where the increased efficiency of our event generation framework becomes relevant for practical applications at the LHC.

## B. Higgs boson plus multi-jet production as an example

With an anticipated  $3 \text{ ab}^{-1}$  at the high-luminosity LHC, Higgs-boson plus multi-jet events will be copiously produced, and even the six jet final state will be measurable at good precision. While not a discovery channel in its own right, the Higgs-boson plus multi-jet signature can be used to test the dynamics of the Standard Model, and it also provides the background to a number of Higgs-boson related measurements and searches, such as Di-Higgs production. In anticipation of these analyses it behooves us to provide precision simulations. In this subsection, we therefore present the first study of Higgs-boson production through gluon fusion at the LHC, with up to seven additional jets computed at LO QCD and up to two jets computed at NLO QCD in the Higgs effective theory [90, 91]. We use the MEPS@NLO algorithm [41, 42] to merge these calculations into an inclusive event sample. The parton-level inputs are generated using Amegic [63], Comix [15] and MCFM [92–98]<sup>7</sup>.

<sup>6</sup> The event files for the NLO parton-level input can be found at <https://doi.org/10.5281/zenodo.8226865>.

<sup>7</sup> The event files can be found at <https://doi.org/10.5281/zenodo.7751000> and <https://doi.org/10.5281/zenodo.7747376>.

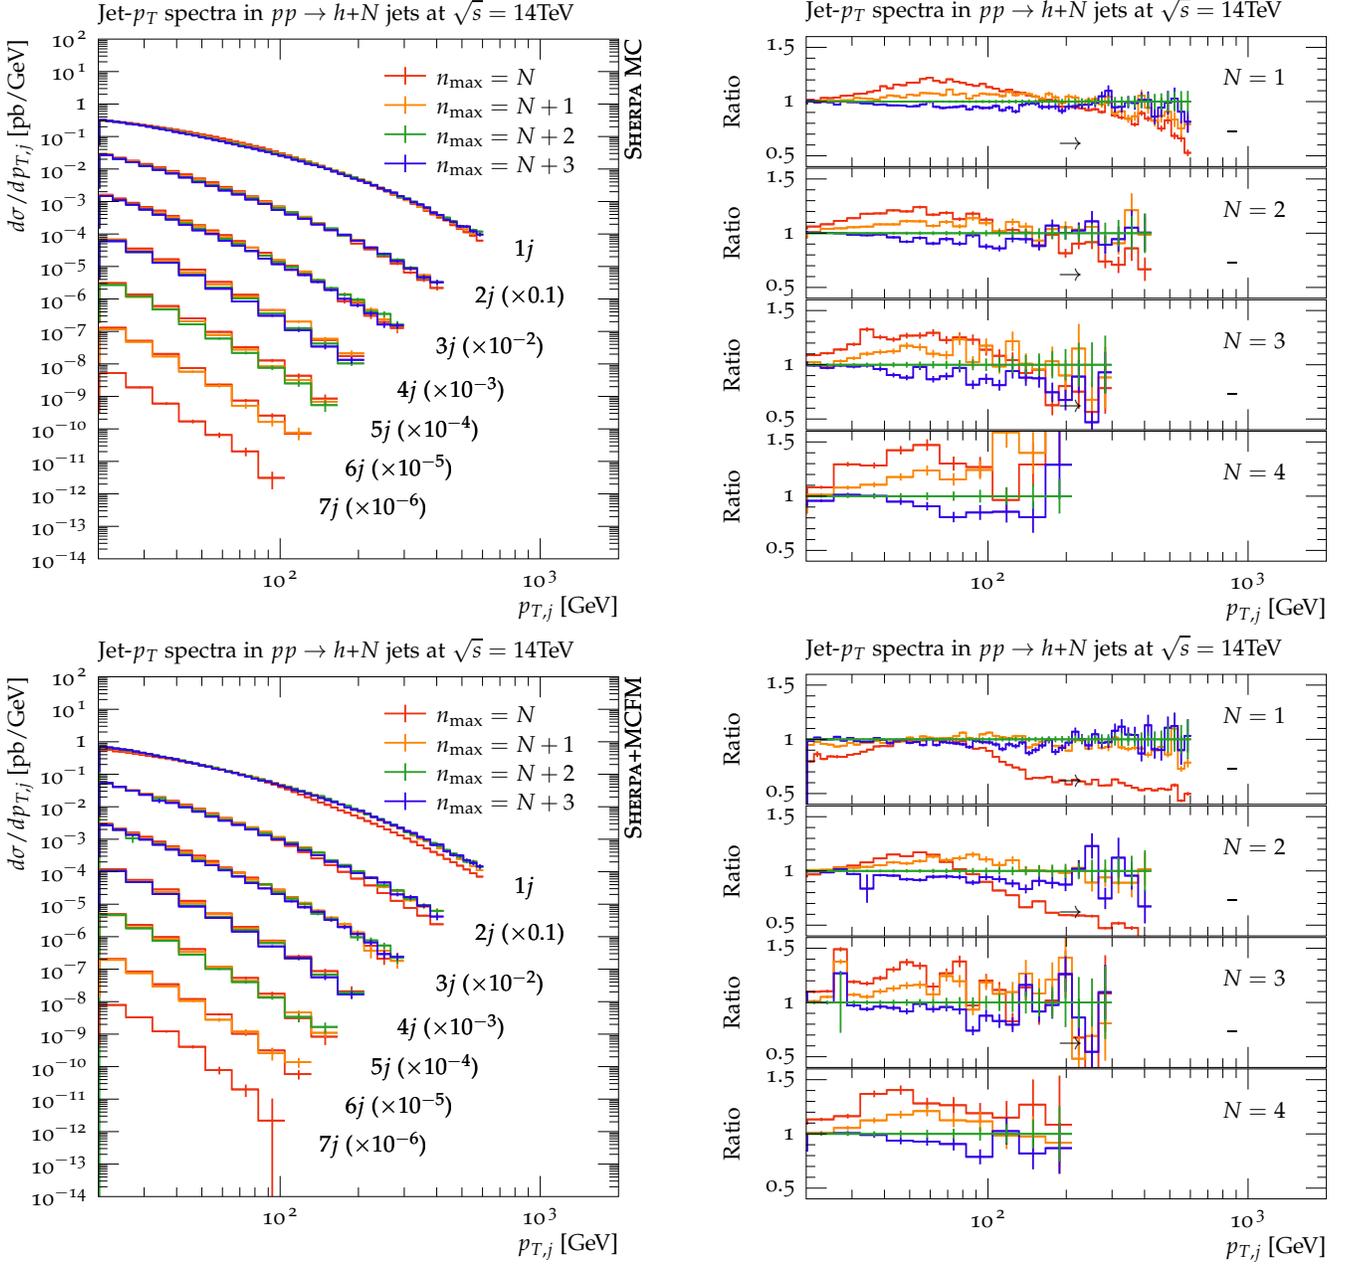


FIG. 7. Jet transverse momentum distributions in Higgs-boson+jets events, computed using multi-jet merging with maximum jet multiplicity equal to  $N$  (red),  $N + 1$  (green),  $N + 2$  (blue) and  $N + 3$  (purple), with  $N$  the number of measured jets. The top panels show leading-order, the bottom panels show next-to-leading order merged results with  $n_{\max, \text{NLO}} = 2$ .

We consider proton-proton collisions at the high-luminosity LHC at  $\sqrt{s} = 14 \text{ TeV}$ . The basic setup has been described in [60]. We use the CT14 NNLO PDF set [80] and define the strong coupling accordingly. Our modified parton-level event generator is based on Comix [15] as included in Sherpa version 2.2.4 [61, 62]. Our modified particle-level event generator is based on Sherpa version 2.2 [61, 62]. Jets are defined using the  $k_T$  clustering algorithm with  $R = 0.4$ ,  $p_{T,j} > 20 \text{ GeV}$  and  $|\eta_j| < 6$ . Following the good agreement between parton-level and particle-level results established in [83], the renormalization and factorization scales are set to  $\hat{H}_{T,m}/2$ , where  $\hat{H}_{T,m} = \sum_{j \in \text{jets}} p_{t,j} + \sqrt{m_H^2 + p_{T,H}^2}$ .

Figure 7 shows the jet transverse momentum spectra at leading order (top panels) and at next-to-leading order (bottom panels). We compare multi-jet merged simulations where the maximum jet multiplicity,  $n_{\max}$ , is set to the number of measured jets,  $N$  (red), to  $N + 1$  (green),  $N + 2$  (blue) and  $N + 3$  (purple). For NLO merged simulations,

the maximum number of jets computed at NLO precision is  $n_{\text{max,NLO}} = 2$ , and we apply local  $K$ -factors based on this calculation to higher jet multiplicities. The panels on the right show the ratio between different predictions, normalized to the result for  $n_{\text{max}} = N + 2$ . It can be seen that the NLO merged predictions are more stable with respect to variations of  $n_{\text{max}}$  at  $N = 1$ , as expected from the higher precision of the calculation at low jet multiplicity. At  $N = 2$ , this effect is diluted by higher-multiplicity tree-level contributions, as explained in Sec. IV A, Fig. 6. NLO accurate predictions for 3 jets at parton level would help to alleviate this problem [99, 100]. However, we could not generate the corresponding unweighted event samples within the limited computing budget for this publication, and we therefore leave a detailed investigation to future work.

### C. Implementation in state of the art experimental simulations

We have validated our new event processing framework in the ATLAS benchmark setups described in [81]<sup>8</sup>. For practical applications where multiple particle-level simulations are generated with the same parton-level input, the LHEH5 event file technology will result in a significantly reduced overall production cost. There may however be remaining obstacles to implementing the method in large-scale event production for the LHC experiments, in particular the access of sub-samples and the synchronization of sub-samples across various sites of the WLCG. The solution to this problem must be found in collaboration with experts from the LHC experiments, who are proficient in WLCG workflows. We therefore postpone the discussion to a future publication.

## V. CONCLUSIONS

We have presented a new framework for the precise and efficient simulation of events in collider experiments, with particular emphasis on the high-luminosity Large Hadron Collider. The new technique is especially suited for the physics modeling of high-multiplicity final states as it allows to match parton-level calculations at next-to-leading order QCD precision to parton showers and merge multiple exclusive calculations into inclusive predictions. Parametric uncertainty estimates can be computed on the fly, using the techniques from [101]. There are no restrictions on the variations that can be performed, and the variations do not need to be included at the time of parton-level event production. We have demonstrated scalability of our approach on a state of the art high-performance computer at a leadership class computing facility. With the computing demands of the LHC experiments becoming an ever more pressing problem due to increased precision in the measurements, our new framework presents an important step towards a more flexible as well as economically and ecologically sustainable approach to event generation in the high-luminosity era. We have validated the new technology against previous simulation programs and enabled event production with a modern, portable parton-level event generator.

## ACKNOWLEDGMENTS

We are grateful to Rui Wang for collaboration in the initial stages of the project. We would like to thank Christian Preuss for his help in porting the HDF5 technology to the most recent version of the Pythia event generator. This research was supported by the Fermi National Accelerator Laboratory (Fermilab), a U.S. Department of Energy, Office of Science, HEP User Facility. Fermilab is managed by Fermi Research Alliance, LLC (FRA), acting under Contract No. DE-AC02-07CH11359. The work of T.C., S.H., P.H., J.I. and R.L. was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through Advanced Computing (SciDAC) program, grant “HPC framework for event generation at colliders”. E.B. and M.K. acknowledge support from BMBF (contract 05H21MGCAB). Their research is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 456104544; 510810461. C.G. is supported by the STFC SWIFT-HEP project (grant ST/V002627/1). This research used the Fermilab Wilson Institutional Cluster for code development, testing and validation. We are grateful to James Simone for his support.

---

<sup>8</sup> The event files can be found at <https://doi.org/10.5281/zenodo.8298371> and <https://doi.org/10.5281/zenodo.8298334>.

- 
- [1] J. M. Campbell *et al.*, Event Generators for High-Energy Physics Experiments, in *2022 Snowmass Summer Study* (2022) arXiv:2203.11110 [hep-ph].
- [2] S. Amoroso *et al.* (HSF Physics Event Generator WG), Challenges in Monte Carlo Event Generator Software for High-Luminosity LHC, *Comput. Softw. Big Sci.* **5**, 12 (2021), arXiv:2004.13687 [hep-ph].
- [3] S. Amoroso *et al.* (HSF Physics Event Generator WG), Challenges in Monte Carlo event generator software for High-Luminosity LHC, arXiv:2004.13687 [hep-ph] (2020).
- [4] E. Yazgan *et al.* (HSF Physics Event Generator WG), HL-LHC Computing Review Stage-2, Common Software Projects: Event Generators, arXiv:2109.14938 [hep-ph] (2021).
- [5] F. James, Monte-Carlo phase space (1968), CERN-68-15.
- [6] E. Byckling and K. Kajantie, Reductions of the phase-space integral in terms of simpler processes, *Phys. Rev.* **187**, 2008 (1969).
- [7] E. Byckling and K. Kajantie, N-particle phase space in terms of invariant momentum transfers, *Nucl. Phys.* **B9**, 568 (1969).
- [8] R. Kleiss, W. J. Stirling, and S. D. Ellis, A new Monte Carlo treatment of multiparticle phase space at high energies, *Comput. Phys. Commun.* **40**, 359 (1986).
- [9] C. G. Papadopoulos, PHEGAS: A phase-space generator for automatic cross-section computation, *Comput. Phys. Commun.* **137**, 247 (2001), hep-ph/0007335.
- [10] F. Maltoni and T. Stelzer, MadEvent: automatic event generation with MadGraph, *JHEP* **02**, 027, hep-ph/0208156.
- [11] A. van Hameren and C. G. Papadopoulos, A hierarchical phase space generator for QCD antenna structures, *Eur. Phys. J. C* **25**, 563 (2002), hep-ph/0204055.
- [12] A. van Hameren, Kaleu: A General-Purpose Parton-Level Phase Space Generator, arXiv:1003.4953 [hep-ph] (2010).
- [13] S. Plätzer, RAMBO on diet, arXiv:1308.2922 [hep-ph] (2013).
- [14] E. Bothmann, T. Childers, W. Giele, F. Herren, S. Höche, J. Isaacsson, M. Knobbe, and R. Wang, Efficient phase-space generation for hadron collider event simulation, arXiv:2302.10449 [hep-ph] (2023).
- [15] T. Gleisberg and S. Höche, Comix, a new matrix element generator, *JHEP* **12**, 039, arXiv:0808.3674 [hep-ph].
- [16] F. A. Berends, R. Kleiss, P. De Causmaecker, R. Gastmans, and T. T. Wu, Single Bremsstrahlung Processes in Gauge Theories, *Phys. Lett. B* **103**, 124 (1981).
- [17] F. A. Berends and W. Giele, The six-gluon process as an example of Weyl-van der Waerden spinor calculus, *Nucl. Phys.* **B294**, 700 (1987).
- [18] M. L. Mangano, S. J. Parke, and Z. Xu, Duality and multi-gluon scattering, *Nucl. Phys.* **B298**, 653 (1988).
- [19] F. A. Berends and W. T. Giele, Recursive calculations for processes with  $n$  gluons, *Nucl. Phys.* **B306**, 759 (1988).
- [20] F. A. Berends, W. T. Giele, and H. Kuijf, Exact expressions for processes involving a vector boson and up to five partons, *Nucl. Phys.* **B321**, 39 (1989).
- [21] A. Kanaki and C. G. Papadopoulos, HELAC: A package to compute electroweak helicity amplitudes, *Comput. Phys. Commun.* **132**, 306 (2000), hep-ph/0002082.
- [22] M. L. Mangano, M. Moretti, F. Piccinini, R. Pittau, and A. D. Polosa, ALPGEN, a generator for hard multiparton processes in hadronic collisions, *JHEP* **07**, 001, arXiv:hep-ph/0206293.
- [23] A. Cafarella, C. G. Papadopoulos, and M. Worek, Helac-Phegas: A generator for all parton level processes, *Comput. Phys. Commun.* **180**, 1941 (2009), arXiv:0710.2427 [hep-ph].
- [24] A. Lifson and O. Mattelaer, Improving colour computations in MadGraph5\_aMC@NLO and exploring a  $1/N_c$  expansion, *Eur. Phys. J. C* **82**, 1144 (2022), arXiv:2210.07267 [hep-ph].
- [25] B. Webber, Monte Carlo Simulation of Hard Hadronic Processes, *Ann. Rev. Nucl. Part. Sci.* **36**, 253 (1986).
- [26] A. Buckley *et al.*, General-purpose event generators for LHC physics, *Phys. Rept.* **504**, 145 (2011), arXiv:1101.2599 [hep-ph].
- [27] S. Frixione and B. R. Webber, Matching NLO QCD computations and parton shower simulations, *JHEP* **06**, 029, hep-ph/0204244.
- [28] P. Nason, A new method for combining NLO QCD with shower Monte Carlo algorithms, *JHEP* **11**, 040, hep-ph/0409146.
- [29] S. Frixione, P. Nason, and C. Oleari, Matching NLO QCD computations with parton shower simulations: the POWHEG method, *JHEP* **11**, 070, arXiv:0709.2092 [hep-ph].
- [30] S. Höche, F. Krauss, M. Schönherr, and F. Siegert, A critical appraisal of NLO+PS matching methods, *JHEP* **09**, 049, arXiv:1111.1220 [hep-ph].
- [31] S. Plätzer and S. Gieseke, Dipole Showers and Automated NLO Matching in Herwig++, *Eur. Phys. J. C* **72**, 2187 (2012), arXiv:1109.6256 [hep-ph].
- [32] S. Catani, F. Krauss, R. Kuhn, and B. R. Webber, QCD matrix elements + parton showers, *JHEP* **11**, 063, hep-ph/0109231.
- [33] M. L. Mangano, M. Moretti, and R. Pittau, Multijet matrix elements and shower evolution in hadronic collisions:  $W\bar{b}b+n$ -jets as a case study, *Nucl. Phys.* **B632**, 343 (2002), hep-ph/0108069.
- [34] F. Krauss, Matrix elements and parton showers in hadronic interactions, *JHEP* **08**, 015, hep-ph/0205283.
- [35] L. Lönnblad, Correcting the colour-dipole cascade model with fixed order matrix elements, *JHEP* **05**, 046, hep-ph/0112284.
- [36] N. Lavesson and L. Lönnblad, Merging parton showers and matrix elements – back to basics, *JHEP* **04**, 085, arXiv:0712.2966 [hep-ph].

- [37] J. Alwall *et al.*, Comparative study of various algorithms for the merging of parton showers and matrix elements in hadronic collisions, *Eur. Phys. J.* **C53**, 473 (2008), arXiv:0706.2569 [hep-ph].
- [38] S. Höche, F. Krauss, S. Schumann, and F. Siegert, QCD matrix elements and truncated showers, *JHEP* **05**, 053, arXiv:0903.1219 [hep-ph].
- [39] K. Hamilton, P. Richardson, and J. Tully, A modified CKKW matrix element merging approach to angular-ordered parton showers, *JHEP* **11**, 038, arXiv:0905.3072 [hep-ph].
- [40] L. Lönnblad and S. Prestel, Matching Tree-Level Matrix Elements with Interleaved Showers, *JHEP* **03**, 019, arXiv:1109.4829 [hep-ph].
- [41] S. Höche, F. Krauss, M. Schönherr, and F. Siegert, QCD matrix elements + parton showers: The NLO case, *JHEP* **04**, 027, arXiv:1207.5030 [hep-ph].
- [42] T. Gehrmann, S. Höche, F. Krauss, M. Schönherr, and F. Siegert, NLO QCD matrix elements + parton showers in  $e^+e^- \rightarrow$  hadrons, *JHEP* **01**, 144, arXiv:1207.5031 [hep-ph].
- [43] L. Lönnblad and S. Prestel, Merging Multi-leg NLO Matrix Elements with Parton Showers, *JHEP* **03**, 166, arXiv:1211.7278 [hep-ph].
- [44] R. Frederix and S. Frixione, Merging meets matching in MC@NLO, *JHEP* **12**, 061, arXiv:1209.6215 [hep-ph].
- [45] J. Bellm, S. Gieseke, and S. Plätzer, Merging NLO Multi-jet Calculations with Improved Unitarization, *Eur. Phys. J.* **C78**, 244 (2018), arXiv:1705.06700 [hep-ph].
- [46] H.-U. Bengtsson and T. Sjöstrand, The Lund Monte Carlo for Hadronic Processes: Pythia Version 4.8, *Comput. Phys. Commun.* **46**, 43 (1987).
- [47] B. Andersson, *The Lund model*, Vol. 7 (Cambridge University Press, 2005).
- [48] T. D. Gottschalk, A realistic model for  $e^+e^-$  annihilation including parton bremsstrahlung effects, *Nucl. Phys.* **B214**, 201 (1983).
- [49] T. D. Gottschalk, An improved description of hadronization in the QCD cluster model for  $e^+e^-$  annihilation, *Nucl. Phys.* **B239**, 349 (1984).
- [50] B. R. Webber, A QCD model for jet fragmentation including soft gluon interference, *Nucl. Phys.* **B238**, 492 (1984).
- [51] J.-C. Winter, F. Krauss, and G. Soff, A modified cluster-hadronisation model, *Eur. Phys. J.* **C36**, 381 (2004), hep-ph/0311085.
- [52] T. Sjöstrand and M. van Zijl, A multiple-interaction model for the event structure in hadron collisions, *Phys. Rev.* **D36**, 2019 (1987).
- [53] T. Sjöstrand and P. Z. Skands, Transverse-momentum-ordered showers and interleaved multiple interactions, *Eur. Phys. J.* **C39**, 129 (2005), hep-ph/0408302.
- [54] M. Bähr, J. M. Butterworth, and M. H. Seymour, The Underlying Event and the Total Cross Section from Tevatron to the LHC, *JHEP* **01**, 065, arXiv:0806.2949 [hep-ph].
- [55] T. Sjöstrand, S. Mrenna, and P. Skands, PYTHIA 6.4 physics and manual, *JHEP* **05**, 026, hep-ph/0603175.
- [56] J. Alwall *et al.*, A standard format for Les Houches Event Files, *Comput. Phys. Commun.* **176**, 300 (2007), hep-ph/0609017.
- [57] *Les Houches 2017: Physics at TeV Colliders Standard Model Working Group Report* (2018) arXiv:1803.07977 [hep-ph].
- [58] E. Bothmann *et al.*, A standard convention for particle-level Monte Carlo event-variation weights, arXiv:2203.08230 [hep-ph] (2022).
- [59] J. T. Childers, T. D. Uram, T. J. LeCompte, M. E. Papka, and D. P. Benjamin, Adapting the serial Alpgen parton-interaction generator to simulate LHC collisions on millions of parallel threads, *Comput. Phys. Commun.* **210**, 54 (2017), arXiv:1511.07312 [hep-ph].
- [60] 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**, 014024 (2019), arXiv:1905.05120 [hep-ph].
- [61] T. Gleisberg, S. Höche, F. Krauss, M. Schönherr, S. Schumann, F. Siegert, and J. Winter, Event generation with Sherpa 1.1, *JHEP* **02**, 007, arXiv:0811.4622 [hep-ph].
- [62] E. Bothmann *et al.* (Sherpa), Event Generation with Sherpa 2.2, *SciPost Phys.* **7**, 034 (2019), arXiv:1905.09127 [hep-ph].
- [63] F. Krauss, R. Kuhn, and G. Soff, AMEGIC++ 1.0: A Matrix Element Generator In C++, *JHEP* **02**, 044, hep-ph/0109036.
- [64] E. Bothmann, W. Giele, S. Höche, J. Isaacson, and M. Knobbe, Many-gluon tree amplitudes on modern GPUs: A case study for novel event generators, *SciPost Phys. Codebases* **3** (2022), arXiv:2106.06507 [hep-ph].
- [65] 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, arXiv:1410.3012 [hep-ph] (2014).
- [66] S. Frixione, Z. Kunszt, and A. Signer, Three-jet cross-sections to next-to-leading order, *Nucl. Phys.* **B467**, 399 (1996), arXiv:hep-ph/9512328.
- [67] S. Catani and M. H. Seymour, A general algorithm for calculating jet cross sections in NLO QCD, *Nucl. Phys.* **B485**, 291 (1997), arXiv:hep-ph/9605323.
- [68] S. Catani, S. Dittmaier, M. H. Seymour, and Z. Trocsanyi, The dipole formalism for next-to-leading order QCD calculations with massive partons, *Nucl. Phys.* **B627**, 189 (2002), hep-ph/0201036.
- [69] The HDF Group, Hierarchical Data Format, version 5 (1997-2023), <https://www.hdfgroup.org/HDF5/>.
- [70] The MPI Forum, MPI-2: Extensions to the Message-Passing Interface (1997).
- [71] HighFive - HDF5 header-only C++ Library, <https://bluebrain.github.io/HighFive/>.
- [72] B. Behzad, H. V. T. Luu, J. Huchette, S. Byna, Prabhat, R. Aydt, Q. Koziol, and M. Snir, Taming parallel i/o complexity with auto-tuning, in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, SC '13 (Association for Computing Machinery, New York, NY, USA, 2013).

- [73] P. Carns, R. Latham, R. Ross, K. Iskra, S. Lang, and K. Riley, 24/7 Characterization of Petascale I/O Workloads, in *Proceedings of 2009 Workshop on Interfaces and Architectures for Scientific Data Storage* (2009).
- [74] P. Carns, K. Harms, W. Alcock, C. Bacon, S. Lang, R. Latham, and R. Ross, Understanding and Improving Computational Science Storage Access through Continuous Characterization, *Trans. Storage* **7** (2011).
- [75] *QCD tree amplitudes on modern GPUs: A case study for novel event generators*, Vol. ICHEP2022 (2022).
- [76] T. Melia, Dyck words and multi-quark primitive amplitudes, *Phys. Rev. D* **88**, 014020 (2013), arXiv:1304.7809 [hep-ph].
- [77] T. Melia, Dyck words and multi-quark amplitudes, *PoS RADCOR2013*, 031 (2013).
- [78] H. Johansson and A. Ochirov, Color-Kinematics Duality for QCD Amplitudes, *JHEP* **01**, 170, arXiv:1507.00332 [hep-ph].
- [79] T. Melia, Proof of a New Colour Decomposition for QCD Amplitudes, *JHEP* **12**, 107, arXiv:1509.03297 [hep-ph].
- [80] S. Dulat, T.-J. Hou, J. Gao, M. Guzzi, J. Huston, P. Nadolsky, J. Pumplin, C. Schmidt, D. Stump, and C. P. Yuan, New parton distribution functions from a global analysis of quantum chromodynamics, *Phys. Rev.* **D93**, 033006 (2016), arXiv:1506.07443 [hep-ph].
- [81] E. Bothmann, A. Buckley, I. A. Christidi, C. Gütschow, S. Höche, M. Knobbe, T. Martin, and M. Schönherr, Accelerating LHC event generation with simplified pilot runs and fast PDFs, *Eur. Phys. J. C* **82**, 1128 (2022), arXiv:2209.00843 [hep-ph].
- [82] K. Danziger, S. Höche, and F. Siegert, Reducing negative weights in Monte Carlo event generation with Sherpa, (2021), arXiv:2110.15211 [hep-ph].
- [83] J. Bellm *et al.*, Jet Cross Sections at the LHC and the Quest for Higher Precision, *Eur. Phys. J. C* **80**, 93 (2020), arXiv:1903.12563 [hep-ph].
- [84] A. Buckley *et al.*, A comparative study of Higgs boson production from vector-boson fusion, *JHEP* **11**, 108, arXiv:2105.11399 [hep-ph].
- [85] K. Hamilton, P. Nason, and G. Zanderighi, MINLO: Multi-scale improved NLO, *JHEP* **1210**, 155, arXiv:1206.3572 [hep-ph].
- [86] S. Höche, P. Maierhöfer, N. Moretti, S. Pozzorini, and F. Siegert, Next-to-leading order QCD predictions for top-quark pair production with up to three jets, *Eur. Phys. J. C* **77**, 145 (2017), arXiv:1607.06934 [hep-ph].
- [87] F. R. Anger, F. Febres Cordero, S. Höche, and D. Maitre, Weak vector boson production with many jets at the LHC  $\sqrt{s} = 13$  TeV, *Phys. Rev. D* **97**, 096010 (2018), arXiv:1712.08621 [hep-ph].
- [88] F. Dulat, S. Höche, and S. Prestel, Leading-Color Fully Differential Two-Loop Soft Corrections to QCD Dipole Showers, *Phys. Rev.* **D98**, 074013 (2018), arXiv:1805.03757 [hep-ph].
- [89] S. Ferrario Ravasio, K. Hamilton, A. Karlberg, G. P. Salam, L. Scyboz, and G. Soyez, A parton shower with higher-logarithmic accuracy for soft emissions, (2023), arXiv:2307.11142 [hep-ph].
- [90] S. Dawson, Radiative corrections to Higgs boson production, *Nucl. Phys.* **B359**, 283 (1991).
- [91] A. Djouadi, M. Spira, and P. Zerwas, Production of Higgs bosons in proton colliders: QCD corrections, *Phys. Lett.* **B264**, 440 (1991).
- [92] J. M. Campbell and R. K. Ellis, Update on vector boson pair production at hadron colliders, *Phys. Rev.* **D60**, 113006 (1999), arXiv:hep-ph/9905386.
- [93] J. M. Campbell, R. K. Ellis, and G. Zanderighi, Next-to-Leading order Higgs + 2 jet production via gluon fusion, *JHEP* **10**, 028, arXiv:hep-ph/0608194.
- [94] S. Badger, J. M. Campbell, R. K. Ellis, and C. Williams, Analytic results for the one-loop NMHV Hqgg amplitude, *JHEP* **12**, 035, arXiv:0910.4481 [hep-ph].
- [95] J. M. Campbell, R. K. Ellis, and C. Williams, Vector boson pair production at the LHC, *JHEP* **07**, 018, arXiv:1105.0020 [hep-ph].
- [96] J. M. Campbell, R. K. Ellis, and W. T. Giele, A Multi-Threaded Version of MCFM, arXiv:1503.06182 [physics.comp-ph] (2015).
- [97] J. Campbell and T. Neumann, Precision Phenomenology with MCFM, *JHEP* **12**, 034, arXiv:1909.09117 [hep-ph].
- [98] J. M. Campbell, S. Höche, and C. T. Preuss, Accelerating LHC phenomenology with analytic one-loop amplitudes: A C++ interface to MCFM, *Eur. Phys. J. C* **81**, 1117 (2021), arXiv:2107.04472 [hep-ph].
- [99] N. Greiner, S. Höche, G. Luisoni, M. Schönherr, J.-C. Winter, and V. Yundin, Phenomenological analysis of Higgs boson production through gluon fusion in association with jets, *JHEP* **01**, 169, arXiv:1506.01016 [hep-ph].
- [100] N. Greiner, S. Höche, G. Luisoni, M. Schönherr, and J.-C. Winter, Full mass dependence in Higgs boson production in association with jets at the LHC and FCC, *JHEP* **01**, 091, arXiv:1608.01195 [hep-ph].
- [101] E. Bothmann, M. Schönherr, and S. Schumann, Reweighting QCD matrix-element and parton-shower calculations, *Eur. Phys. J.* **C76**, 590 (2016), arXiv:1606.08753 [hep-ph].