

Hardware acceleration for next-to-leading order event generation within MADGRAPH5_AMC@NLO

Zenny Wettersten^{1,2,*}, Olivier Mattelaer³, Stefan Roiser¹, Andrea Valassi¹, and Marco Zaro⁴

¹European Organization for Nuclear Research (CERN)

²Technical University of Vienna (TU Wien)

³University of Louvain (UCLouvain)

⁴University of Milan (UniMi)

Abstract. As the quality of experimental measurements increases, so does the need for Monte Carlo-generated simulated events — both with respect to the total amount and to their precision. In perturbative methods, this involves the evaluation of higher order corrections to the leading order (LO) scattering amplitudes, including real emissions and loop corrections. Although experimental uncertainties today are larger than those of simulations, at the High Luminosity LHC experimental uncertainties are expected to be smaller than the theoretical uncertainty for events generated below next-to-leading order (NLO) precision. As forecasted hardware resources will not meet CPU requirements for these simulation needs, speeding up NLO event generation is a necessity.

In recent years, collaborators across Europe and the United States have been working on CPU vectorisation of LO event generation within the MADGRAPH5_AMC@NLO framework, as well as porting it to GPUs, to major success. Recently, development has also started on vectorising NLO event generation. Due to the more complicated nature of NLO amplitudes this development faces several difficulties not accounted for in the LO development, but it shows promise. Here, we present these issues as well as the current status of our event-parallel NLO implementation.

1 Introduction

Over the course of the upcoming High-Luminosity LHC (HL-LHC) era, the integrated luminosity is expected to rise by an order of magnitude when compared to the end of Run-3 [1]. The measurements made at the LHC need to be statistically compared with Monte-Carlo (MC) generated simulated events, and the predicted precision of HL-LHC experiments necessitate not only an order of magnitude increase in the number of generated events, but also of their precision [2]. In perturbative quantum field theory (pQFT), this entails moving from predictions made at leading order (LO) to ones at next-to-leading order (NLO).

In response to this issue, several MC event generators have made efforts to port event generation to data-parallel architectures, such as GPUs and vectorised CPUs [2–9]. For the general-purpose event generator MADGRAPH5_AMC@NLO (MGAMC), this has been done at the level of LO *tree-level* scattering amplitudes, where significant speed-up has been found

*e-mail: zenny.wettersten@cern.ch

using *single instructions on multiple data* (SIMD) and *single instructions across multiple threads* (SIMT). Continuing this effort, we have started work on parallelising NLO amplitudes, where many challenges — both old and new — are found. Nevertheless, we have developed a multithreaded prototype for event-parallel fixed order calculations in MGAMC, which is currently in the process of being validated. Efforts to port these developments to event generation and to inject tree-level amplitudes with vector instructions are planned.

Due to the developmental status of the parallel fixed order implementation, we here focus on the fundamental difficulties of NLO calculations. In sections 2 and 3, we detail the theoretical background of pQFT and MC event generation respectively. Section 4 extensively details MGAMC NLO calculations, both in the upstream codebase and in our multithreaded prototype, and elaborates on the technical issues differentiating NLO from LO. Finally, we summarise the concerns and current status in section 5.

2 Perturbation theory

There are many stages between a QFT described in terms of a Lagrangian and the expected experimental observations, but the first one — the one describing the underlying fundamental physics — lies in going from a Lagrangian to a cross section. In section 2.1 we briefly discuss the closed form solution for cross sections using perturbation theory, and in section 2.2 we detail the additional difficulties entailed in beyond-leading order calculations. For a thorough derivation, we refer the reader to a QFT textbooks, e.g. [10].

2.1 Differential cross sections

Particle physics concerns the evaluation of cross sections. A problem arises as there is generally no closed form solution for the cross section. The standard way to handle this is pQFT, where we Taylor expand the interactions about the non-interacting theory. Ignoring the identity operator, the first non-vanishing term is referred to as the *leading order* contribution, the second the *next-to-leading order*, etc.¹

In order to describe amplitudes in a clear and explicit way, we use *Feynman diagrams*, representations of the ways particles in an interaction can be connected to each other; particles are shown as lines, and interactions as nodes. As an example, consider the process $e^+e^- \rightarrow 2\gamma$, electron-positron annihilation producing two photons. In the standard model, the LO Feynman diagrams for this process are²

$$\frac{d\sigma_{LO}}{d\Omega} \simeq \left| \begin{array}{c} e^- \rightarrow \text{---} \text{---} \text{---} \gamma \\ \text{---} \text{---} \text{---} \gamma \\ e^+ \leftarrow \text{---} \text{---} \text{---} \end{array} \right|^2, \quad (1)$$

where σ is the total cross section and $d\Omega$ is the Lorentz invariant phase space differential³. As Feynman diagrams are complex scalars, the differential cross section is given by the absolute square of their sum.

Note here that Feynman diagrams are topological in nature. Consequently, even disregarding the additional complexity in evaluations of singular Feynman diagrams at higher multiplicities, the complexity scaling of these methods are $\sim \mathcal{O}(n!)$ in the number of external partons n .

¹The LO term is not necessarily tree-level. For simplicity, we assume that the LO contribution is tree-level and forego any discussion on loop-induced processes — they would, for our purposes, be considered NLO.

²These are two different diagrams under photon interchange, as photons are indistinguishable.

³See dLIPS formulation in any standard QFT textbook.

2.2 Next-to-leading order corrections

Going on to the NLO contribution, we remind the reader that the perturbative expansion is in the interaction Hamiltonian. NLO thus necessitates an additional vertex. This can be thought of as appending an extra particle, and these contributions are referred to as *real emissions*. Returning to $e^+e^- \rightarrow 2\gamma$ as in eq. (1), some of the real emission diagrams are given by



$$(2)$$

where as in eq. (1) we must sum over all photon orderings. Furthermore, we must consider *all* ways we can add an additional vertex, such as e.g. from a photon splitting into new fermions. Each resulting real emission faces factorial scaling; excluding other electroweak bosons, there are 54 real emission diagrams for this process, in contrast to two at LO.

Furthermore, we reiterate that $d\sigma$ is given by an absolute square, and at NLO we must also consider interference between the LO Born diagrams and *loop amplitudes*. These arise when we add *two* vertices to the LO diagrams with a completely internal particle, e.g.



$$(3)$$

where the leptons first scatter against each other before annihilating. The term “loop” refers to the closed particle loop visible in the diagram; for *tree-level* processes (eqs. (1) and (2)) the momenta of internal particles are uniquely defined by the external partons. For loop processes, the total momentum at each vertex depends on the unobservable internal propagator and must thus be integrated over internal momenta. For $e^+e^- \rightarrow 2\gamma$, again excluding other bosons, there are 19 loop diagrams.

The final complication of NLO amplitude computations arises during phase space integration — $d\sigma_{NLO}$ has singularities across the phase space⁴. In fact, both the real and the loop contributions are singular in their respective phase spaces. However, their *sum* should be convergent. This is proven in general in e.g. [10]. For general evaluations, we employ *subtraction schemes* where we subtract zero in such a way that both the real and the loop contributions are locally convergent. This can be expressed as a *counter term* $d\sigma_{CT}$ of the form

$$\begin{aligned} \sigma_{NLO} - \sigma_{LO} &= \int d\Omega_n \frac{d\sigma_V}{d\Omega_n} + \int d\Omega_{n+1} \frac{d\sigma_R}{d\Omega_{n+1}} \\ &= \int d\Omega_n \left\{ \frac{d\sigma_V}{d\Omega_n} + \int d\Omega_1 \frac{d\sigma_{CT}}{d\Omega_{n+1}} \right\} + \int d\Omega_{n+1} \left\{ \frac{d\sigma_R}{d\Omega_{n+1}} - \frac{d\sigma_{CT}}{d\Omega_{n+1}} \right\}, \end{aligned} \quad (4)$$

with virtual loop (real emission) contributions $d\sigma_V$ ($d\sigma_R$), and the phase space measures $d\Omega$ have been explicitly defined in terms of the n - and $(n+1)$ -parton phase spaces. Two standard subtraction schemes are FKS subtraction [11, 12] and CS dipole subtraction [13], the former of which is used in MGAMC.

⁴Technically, $d\sigma_{LO}$ already has singularities, but they are by definition unobserved and thus excluded by the cuts enforced on the integral.

3 Parton-level event generation

Section 2 provides an overview into the evaluation of observables, but misses one important aspect: QFTs are stochastic. What is observed at particle colliders will not be the expectation value of the theory but a statistical sample of events. To test a theory, we must compare experimental samples with predicted ones. The method used to generate these samples is called *Monte-Carlo event generation* due to its basis in Monte-Carlo (MC) integration.

3.1 Monte-Carlo integration

MC integration is a numerical integration method based on the notion that an integral is necessarily equal to the mean value of the integrand times the volume of the integrated measure,

$$I = \int dx f(x) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N f(x_n), \quad (5)$$

assuming that x_n are randomly sampled across the integrated measure. The reason to use MC methods in HEP is twofold: Firstly, MC integration accuracy scales independently of the dimensionality of the integrand [14]; and secondly, importance sampling the integral measure is exactly the same thing as determining the shape of a probability distribution. Thus, we can use MC techniques to first determine the total cross section of a process and then sample the phase space with the same likelihood we expect to see in an experiment. This process of phase space sampling is known as *unweighted event generation*.

3.2 Data-parallel LO event generation

Many event generators are employed for prediction and analysis at the LHC experiments; for an overview of some of the primary ones, see e.g. [15]. Our focus will be on MGAMC [16, 17], a hard-event generator⁵ supporting generic QFTs expressed in the UFO format [18, 19] with LO and NLO routines generated by ALOHA [20]. MGAMC event generation generates amplitude routines for relevant parton configurations and determines the different phase space distributions before sampling them for unweighted events.

At LO, each parton configuration is identical across phase space. Each event evaluates the exact same function (eq. (1)) for many phase space points. Consequently, “*event-level parallelism looks like an appropriate approach to try and exploit efficiently [GPU and vectorized CPU architectures]*” [2]. This was the motivation behind development of CUDACPP, a plugin for MGAMC which generates event-level data-parallel amplitude routines and interfaces them with MGAMC to run event generation on GPUs and vectorised CPUs [4–7]. In our effort to utilise hardware acceleration also for hard NLO event generation, we expect opportunities to recycle much of the work done for LO, as elaborated on in sections 4 and 5.

4 Event-level parallelism at NLO

To increase the precision of theoretical predictions in pQFT, we must evaluate higher order terms. Given that the CUDACPP plugin presented in section 3.2 enables hardware accelerated event generation at LO, the obvious next step is to port also NLO event generation to accelerators; we justify our decision to maintain an event-level parallel structure at NLO in section 4.1. We then continue with the novel complications that arise when moving onto NLO in section 4.2, and finally describe our efforts in sections 4.3 and 4.4.

⁵By a “hard-event generator” we mean a generator of events according to the parton-level hard scattering process. The term event generator is somewhat ambiguous, but generally refers to the production of showered events.

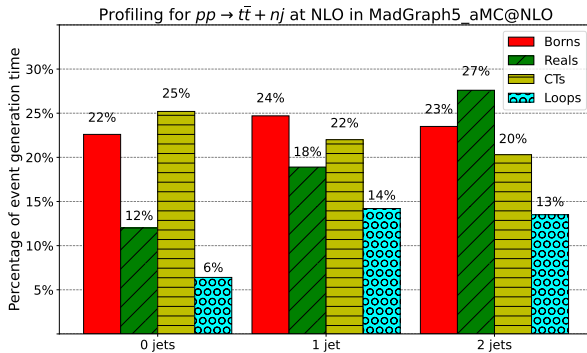


Figure 1. Runtime profile, generating 10 000 QCD NLO events in MGAMC. Borns refers to LO contributions; Reals (Loops) to real emissions (loop diagrams); and counter-terms (CTs) collectively to FKS subtraction terms and MC counter-events. Only amplitude evaluation runtimes shown.

4.1 Compute bottlenecks in MadGraph5_aMC@NLO

In fig. 1, a runtime profile for pure-QCD NLO event generation in MGAMC is presented, considering the process $pp \rightarrow t\bar{t} + nj$, $n = 0, 1, 2$. We only show the runtime fraction spent in amplitudes — summing, we find that these make up 65 – 83% of total runtime. While less dominant than at LO, amplitudes are still the primary bottleneck and still make up an increasing fraction of runtime as process complexity increases. From these numbers, the expected Amdahl's acceleration ranges from 2.4 – 3.3 (2.9 – 5.9) assuming loop integrations cannot (can) be parallelised. Event-level parallelism still seems the appropriate approach.

4.2 Complications in higher order evaluations

While section 4.1 illustrates that parallelising scattering amplitudes still makes for the most obvious way to accelerate NLO event generation, certain fundamental differences between LO and NLO need to be addressed. There are a plethora of details that could be discussed, but we will limit ourselves to a couple key ones that need to be considered for hardware acceleration. In particular, these are phase space asymmetry and loop integrals.

As mentioned in section 3.2, amplitudes are identical across phase space for LO event generation. At NLO we need to consider the singular regions of the amplitudes as well as the local subtraction terms to ensure their convergence. Using FKS subtraction this is done by splitting the phase space into sectors with different kinematic parametrisations. As such, amplitudes are no longer identical; for SIMD and SIMT, parallelism should be across individual FKS sectors rather than the full phase space.

The second issue are the loop integrals. Firstly, they are computationally expensive⁶. This is not an issue in and of itself, but for the fact that MGAMC offloads loop evaluations to external libraries. Consequently, loop integration cannot be accelerated directly through MGAMC. Secondly, due to instabilities loop integration at times requires quadruple floating point precision (FP128). No major manufacturers have hardware-level support for FP128 instructions, neither for vectorised CPUs nor for GPUs.

One saving grace to read from fig. 1 is that loop evaluations, at least for these pure QCD processes, are not a major bottleneck. A majority of the runtime is spent in LO Born and NLO real emission amplitudes and in CTs. It is important to note that *all of these are tree-level*; they use the routines that have been vectorised for the CUDACPP plugin. We expect to be able to recycle work on tree-level amplitudes from CUDACPP for NLO event generation.

⁶Figure 1 might give the impression that loops are not a problem, but the reason for this is that MGAMC generally uses an approximation of the full loop integral rather than explicit evaluation. For details, see section 2.4.3 of [16].

Algorithm 1 Current event loop

```

1:  $FKS \leftarrow \text{SETFKSSECTOR}()$ 
2:  $p \leftarrow \text{GENMOMENTA}(FKS)$ 
3: if  $\text{PASSCUTS\_NBODY}(p)$  then
4:    $M(\text{Born}) \leftarrow \text{BORNMATRIX}(p)$ 
5:    $M(\text{Loop}) \leftarrow \text{LOOPMATRIX}(p)$ 
6: end if
7: if  $\text{PASSCUTS\_REAL}(p)$  then
8:    $M(\text{Real}) \leftarrow \text{REALMATRIX}(p)$ 
9: end if
10: if  $\text{PASSCUTS\_SOFT}(p)$  then
11:    $M(\text{Soft}) \leftarrow \text{SOFTMATRIX}(p)$ 
12: end if
13: if  $\text{PASSCUTS\_COLL}(p)$  then
14:    $M(\text{Coll}) \leftarrow \text{COLLMATRIX}(p)$ 
15: end if

```

Algorithm 2 Branch-free event loop

```

1:  $FKS \leftarrow \text{SETFKSSECTOR}()$ 
2:  $p \leftarrow \text{GENMOMENTA}(FKS)$ 
3: while  $\text{!PASSANYCUTS}(p)$  do
4:    $p \leftarrow \text{GENMOMENTA}(FKS)$ 
5: end while
6:  $M(\text{Born}) \leftarrow \text{BORNMATRIX}(p)$ 
7:  $M(\text{Real}) \leftarrow \text{REALMATRIX}(p)$ 
8:  $M(\text{Soft}) \leftarrow \text{SOFTMATRIX}(p)$ 
9:  $M(\text{Coll}) \leftarrow \text{COLLMATRIX}(p)$ 
10:  $M(\text{Loop}) \leftarrow \text{LOOPMATRIX}(p)$ 
11: for  $a \in [B, R, S, C]$  do
12:   if  $\text{!PASSCUTS}(a, p)$  then
13:      $M(a) \leftarrow 0$ 
14:   end if
15: end for

```

4.3 A data parallel algorithm

Before describing our event-parallel NLO algorithm, we note that the codebases for LO and NLO in MGAMC are completely distinct, and the work on developing a multi-event interface at LO could not be reused. However, as mentioned in section 4.2, the vectorised tree-level amplitudes in CUDACPP *can* be recycled for the brunt of amplitude calls at NLO. Now that an NLO multi-event interface exists, there is no need to develop these from scratch.

In algorithms 1 and 2, we give an overview of the program flow for a phase space point at NLO in the MGAMC codebase (left) and our “naive” branch-free algorithm (right). Algorithm 1 is highly branching, making it unsuited for SIMD/SIMT architectures. We forego this branching in algorithm 2 at the cost of superfluous calculations, allowing for parallelism over lines 6 through 10 (omitting the MC over loop evaluations for now). While this implies wasted compute, on heterogeneous systems we expect the loss to be outweighed by the speed-up from remaining on the device for longer. A smarter (albeit memory-intensive and latency-bound) algorithm might store momenta passing different cuts separately and evaluate asynchronously one type of amplitude in parallel once sufficiently many points pass. We defer this discussion until we have an implementation with SIMD/SIMT amplitudes running.

4.4 Implementation status

Recently, development of event-level data-parallel NLO computations in MGAMC was officially started. The original work plan can briefly be summarised in four stages:

1. Write a multi-event interface for fixed order NLO computations⁷ as per algorithm 2.
2. Make an initial data-parallel implementation by multithreading over events.
3. Inject CUDACPP-generated tree-level amplitudes to utilise SIMD and SIMT hardware.
4. Port functionality over to event generation.

⁷The interface for fixed order computations (cross section integration) and unweighted event generation at NLO are disparate but largely identical in MGAMC. The former simply omits some complications of full event generation.

As of now, we have a working prototype at step 2, i.e. an event-parallel fixed order implementation using multithreading. This prototype is currently being validated, but appears to be data-safe for tested processes and yields cross sections in agreement with MGAMC to (1) numerical precision when run with the same seed and no vectorisation, and (2) statistical precision when run with vectorisation, due to a shift in random number generation calls. At present, this validation primarily concerns cross section validation across minimal 2-to-2 processes such as $pp \rightarrow l^+v_l$, $pp \rightarrow l^+l^-$, and $pp \rightarrow t\bar{t}$, excluding loop contributions.

While validation efforts are taking place and code refactoring and optimisations are considered, we are planning for the injection of CUDACPP-generated tree-level amplitudes into NLO executables. Although there are slight differences in how amplitudes are called, the interface is sufficiently similar to LO that only minor modifications need to be made on both sides to enable this. We expect to have a prototype also of this within the near future.

5 Outlook

In lieu of the need for higher-precision theoretical HEP predictions and the success in porting hard LO event generation to data-parallel architectures, parallelising also NLO computations to GPUs and vectorised CPUs appears to be an obvious next step. We have illustrated that implementing hardware acceleration at NLO using the same methodology as used for LO event generation (i.e. focusing initially on porting scattering amplitudes to SIMD and SIMT architectures) is an appropriate measure, although the expected speed-up is less than what can be observed for large-multiplicity processes at LO. Furthermore, we are currently working on implementing this for tree-level amplitudes. Additional complications arise when considering e.g. loop amplitudes, which generally may require FP128 instructions, but moving tree-level amplitudes to SIMD/SIMT hardware alone should provide a significant speed-up.

Furthermore, we have an extensive work plan for development of event-level parallelism for NLO event generation in the MGAMC software suite, which we are already tackling. So far, we have implemented a multi-event interface for fixed order NLO computations and tested the data safety of the algorithm by multithreading over phase space points. The next steps of development involve recycling data-parallel tree-level scattering amplitude routines from the CUDACPP plugin within the NLO codebase and porting these developments to event generation. We expect to have presentable results in the near future.

Acknowledgements

We thank Stefano Frixione for useful discussion regarding FKS subtraction and the relevance of FKS sector parametrisation to SIMD/SIMT event-parallelism at NLO. We also thank Rikkert Frederix for extensive assistance in disentangling the MGAMC NLO codebase.

References

- [1] B. Schmidt, The High-Luminosity upgrade of the LHC: Physics and Technology Challenges for the Accelerator and the Experiments, *Journal of Physics: Conference Series* **706**, 022002 (2016). [10.1088/1742-6596/706/2/022002](https://doi.org/10.1088/1742-6596/706/2/022002)
- [2] A. Valassi, E. Yazgan, J. McFayden 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), 2004. 13687. [10.1007/s41781-021-00055-1](https://doi.org/10.1007/s41781-021-00055-1)
- [3] E. Yazgan, J. McFayden, A. Valassi, S. Amoroso, E. Bothmann, A. Buckley, J. Campbell, G.S. Chahal, T. Childers, G. Corti et al. (HSFPhysicsEventGeneratorWG), *Tech. rep.* (2021), 2109. 14938, <https://cds.cern.ch/record/2783019>

- [4] 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, EPJ Web Conf. **251**, 03045 (2021), 2106.12631. [10.1051/epjconf/202125103045](https://doi.org/10.1051/epjconf/202125103045)
- [5] A. Valassi, T. Childers, L. Field, S. Hageboeck, W. Hopkins, O. Mattelaer, N. Nichols, S. Roiser, D. Smith, Developments in Performance and Portability for MadGraph5_aMC@NLO, PoS **ICHEP2022**, 212 (2022), 2210.11122. [10.22323/1.414.0212](https://doi.org/10.22323/1.414.0212)
- [6] A. Valassi, T. Childers, L. Field, S. Hageböck, W. Hopkins, O. Mattelaer, N. Nichols, S. Roiser, D. Smith, J. Teig et al., Speeding up Madgraph5 aMC@NLO through CPU vectorization and GPU offloading, in *21th International Workshop on Advanced Computing and Analysis Techniques in Physics Research* (2023), 2303.18244
- [7] S. Hageboeck, T. Childers, W. Hopkins, O. Mattelaer, N. Nichols, S. Roiser, J. Teig, A. Valassi, C. Vuosalo, Z. Wettersten, Madgraph5_aMC@NLO on GPUs and vector CPUs Experience with the first alpha release, EPJ Web Conf. **295**, 11013 (2024), 2312.02898. [10.1051/epjconf/202429511013](https://doi.org/10.1051/epjconf/202429511013)
- [8] A. Valassi, Madgraph on GPUs and vector CPUs: towards production (2024), <https://doi.org/10.5281/zenodo.14940556>
- [9] E. Bothmann, T. Childers, W. Giele, S. Höche, J. Isaacson, M. Knobbe, A portable parton-level event generator for the high-luminosity LHC, SciPost Phys. **17**, 081 (2024), 2311.06198. [10.21468/SciPostPhys.17.3.081](https://doi.org/10.21468/SciPostPhys.17.3.081)
- [10] M.E. Peskin, D.V. Schroeder, An Introduction to Quantum Field Theory, Frontiers in Physics (Addison-Wesley Pub. Co, 1995), ISBN 978-0-201-50397-5
- [11] S. Frixione, Z. Kunszt, A. Signer, Three jet cross-sections to next-to-leading order, Nucl. Phys. B **467**, 399 (1996), hep-ph/9512328. [10.1016/0550-3213\(96\)00110-1](https://doi.org/10.1016/0550-3213(96)00110-1)
- [12] S. Frixione, A General approach to jet cross-sections in QCD, Nucl. Phys. B **507**, 295 (1997), hep-ph/9706545. [10.1016/S0550-3213\(97\)00574-9](https://doi.org/10.1016/S0550-3213(97)00574-9)
- [13] S. Catani, M.H. Seymour, A General algorithm for calculating jet cross-sections in NLO QCD, Nucl. Phys. B **485**, 291 (1997), [Erratum: Nucl.Phys.B 510, 503–504 (1998)], hep-ph/9605323. [10.1016/S0550-3213\(96\)00589-5](https://doi.org/10.1016/S0550-3213(96)00589-5)
- [14] W. Press, Numerical Recipes 3rd Edition: The Art of Scientific Computing (Cambridge University Press, 2007), ISBN 9780521880688
- [15] A. Buckley et al., General-purpose event generators for LHC physics, Phys. Rept. **504**, 145 (2011), 1101.2599. [10.1016/j.physrep.2011.03.005](https://doi.org/10.1016/j.physrep.2011.03.005)
- [16] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.S. Shao, T. Stelzer, P. Torrielli, M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, JHEP **07**, 079 (2014), 1405.0301. [10.1007/JHEP07\(2014\)079](https://doi.org/10.1007/JHEP07(2014)079)
- [17] R. Frederix, S. Frixione, V. Hirschi, D. Pagani, H.S. Shao, M. Zaro, The automation of next-to-leading order electroweak calculations, JHEP **07**, 185 (2018), [Erratum: JHEP 11, 085 (2021)], 1804.10017. [10.1007/JHEP11\(2021\)085](https://doi.org/10.1007/JHEP11(2021)085)
- [18] C. Degrande, C. Duhr, B. Fuks, D. Grellscheid, O. Mattelaer, T. Reiter, UFO - The Universal FeynRules Output, Comput. Phys. Commun. **183**, 1201 (2012), 1108.2040. [10.1016/j.cpc.2012.01.022](https://doi.org/10.1016/j.cpc.2012.01.022)
- [19] L. Darmé et al., UFO 2.0: the ‘Universal Feynman Output’ format, Eur. Phys. J. C **83**, 631 (2023), 2304.09883. [10.1140/epjc/s10052-023-11780-9](https://doi.org/10.1140/epjc/s10052-023-11780-9)
- [20] P. de Aquino, W. Link, F. Maltoni, O. Mattelaer, T. Stelzer, ALOHA: Automatic Libraries Of Helicity Amplitudes for Feynman Diagram Computations, Comput. Phys. Commun. **183**, 2254 (2012), 1108.2041. [10.1016/j.cpc.2012.05.004](https://doi.org/10.1016/j.cpc.2012.05.004)