

Two-loop tensor integral coefficients in OpenLoops

Stefano Pozzorini,^a Natalie Schär^b and Max F. Zoller^b

^a*Physik-Institut, Universität Zürich,
Winterthurerstrasse 190, CH-8057 Zürich, Switzerland*

^b*Paul Scherrer Institut,
Forschungsstrasse 111, CH-5232 Villigen PSI, Switzerland*

E-mail: pozzorin@physik.uzh.ch, natalie.schaer@psi.ch,
max.zoller@psi.ch

ABSTRACT: We present a new and fully general algorithm for the automated construction of the integrands of two-loop scattering amplitudes. This is achieved through a generalisation of the open-loops method to two loops. The core of the algorithm consists of a numerical recursion, where the various building blocks of two-loop diagrams are connected to each other through process-independent operations that depend only on the Feynman rules of the model at hand. This recursion is implemented in terms of tensor coefficients that encode the polynomial dependence of loop numerators on the two independent loop momenta. The resulting coefficients are ready to be combined with corresponding tensor integrals to form scattering probability densities at two loops. To optimise CPU efficiency we have compared several algorithmic options identifying one that outperforms naive solutions by two orders of magnitude. This new algorithm is implemented in the OPENLOOPS framework in a fully automated way for two-loop QED and QCD corrections to any Standard Model process. The technical performance is discussed in detail for several $2 \rightarrow 2$ and $2 \rightarrow 3$ processes with up to order 10^5 two-loop diagrams. We find that the CPU cost scales linearly with the number of two-loop diagrams and is comparable to the cost of corresponding real-virtual ingredients in a NNLO calculation. This new algorithm constitutes a key building block for the construction of an automated generator of scattering amplitudes at two loops.

KEYWORDS: Automation, Higher-Order Perturbative Calculations, Higher Order Electroweak Calculations

ARXIV EPRINT: [2201.11615](https://arxiv.org/abs/2201.11615)

Contents

1	Introduction	1
2	Scattering amplitudes and partonic cross sections up to NNLO	4
2.1	Dimensional regularisation and rational terms	6
2.2	Scattering amplitudes and colour factors	7
2.3	Helicity bookkeeping	8
3	Tree and one-loop amplitudes in OPENLOOPS	9
3.1	Tree-level amplitudes	9
3.2	One-loop amplitudes	10
3.3	Born-loop interference	13
4	Reducible two-loop integrands	14
4.1	Reducible two-loop integrands of class Red2	15
4.2	Reducible two-loop integrands of class Red1	19
5	Irreducible two-loop integrands	20
5.1	Generic structure of a recursive two-loop algorithm	22
5.2	Integrand of a single irreducible two-loop diagram	24
5.3	Two-loop tensor coefficients for scattering amplitudes	31
6	Technical performance of the two-loop algorithm	32
6.1	Validation with pseudotree test	32
6.2	Numerical stability	34
6.3	CPU efficiency	35
6.4	Memory usage	38
7	Conclusion	39
A	CPU efficiency measurements	40

1 Introduction

Precise theory predictions based on the Standard Model (SM) play a crucial role for the success of the LHC physics program. In particular, in the absence of striking signals of new physics, the availability of high theory precision for the widest possible range of scattering processes is a key prerequisite in order to maximise the sensitivity of LHC measurements to small effects of physics beyond the SM.

Predictions at next-to-leading order (NLO) in perturbation theory can be obtained through multi-purpose Monte Carlo tools [1–4], and the hard scattering amplitudes at the core of these calculations are computed with automated numerical tools [5–17] at tree and one-loop level.

At next-to-next-to-leading order (NNLO) the required two-loop amplitudes still need to be derived on a process-by-process basis. To date, full NNLO predictions for $2 \rightarrow 2$ processes are widely available (see [18–27] for recent progress¹), while for $2 \rightarrow 3$ processes only a few two-loop amplitudes [29–44] and pioneering NNLO results [45–48] exist. The complexity and the status of NLO calculations for loop-induced processes is similar, with an increasing number of $2 \rightarrow 2$ predictions [49–53] and first results for $2 \rightarrow 3$ processes [54]. Calculations at N³LO are highly challenging and currently restricted to $2 \rightarrow 1$ processes [55–62].

While specialised NNLO Monte Carlo tools [63–65] have been developed, a fully general and automated NNLO tool, including in particular the two-loop scattering amplitudes, is not yet within reach. In light of the high-precision requirements of the LHC and future colliders, such a NNLO tool for arbitrary $2 \rightarrow 2$ and $2 \rightarrow 3$ SM processes is highly desirable. With this objective in mind, in this paper we present a new algorithm that provides an important building block for the automated construction of two-loop amplitudes.

This new algorithm represents the extension of the well-established OPENLOOPS technique [10] from one to two loops. In the OPENLOOPS approach, one-loop scattering amplitudes are constructed as sums of Feynman diagrams. Individual one-loop diagrams Γ have the form

$$\bar{\mathcal{M}}_{1,\Gamma} = C_{1,\Gamma} \int d\bar{q}_1 \frac{\bar{\mathcal{N}}(\bar{q}_1)}{\mathcal{D}(\bar{q}_1)}, \tag{1.1}$$

where $C_{1,\Gamma}$ is a colour factor, and the bar indicates quantities that are defined in D dimensions in order to regularise divergences in loop integrals [66]. The denominator $\mathcal{D}(\bar{q}_1)$ is a product of propagator denominators depending on the loop momentum \bar{q}_1 . In OPENLOOPS the numerator $\bar{\mathcal{N}}(\bar{q}_1)$ is decomposed into loop-momentum tensors,

$$\bar{\mathcal{N}}(\bar{q}_1) = \sum_r \bar{\mathcal{N}}_{\mu_1 \dots \mu_r} \bar{q}_1^{\mu_1} \dots \bar{q}_1^{\mu_r}, \tag{1.2}$$

and the tensor coefficients $\bar{\mathcal{N}}_{\mu_1 \dots \mu_r}$ are constructed through a numerical recursion [10, 67, 68] based on process-independent operations, which depend only on the Feynman rules of the model at hand. This construction is carried out in four space-time dimensions, and the contribution of the missing numerator parts of order $(D - 4)$ are reconstructed via insertion of process-independent rational counterterms [69–72] into tree diagrams. Scattering amplitudes are obtained by combining the tensor coefficients with the associated tensor integrals

$$I^{\mu_1 \dots \mu_r} = \int d\bar{q}_1 \frac{q_1^{\mu_1} \dots q_1^{\mu_r}}{\mathcal{D}(\bar{q}_1)}, \tag{1.3}$$

which can be reduced to scalar integrals using external tools [5, 14]. Alternatively, the construction of tensor coefficients and the reduction of tensor integrals can be combined in a single numerical recursion using the on-the-fly reduction method [68].

¹This list of references includes only a few representative recent results. For a recent review see [28].

Besides all required scattering amplitudes for NLO calculations, the public OPENLOOPS program provides also some of the building blocks of NNLO calculations, namely squared one-loop amplitudes as well as Born one-loop interferences for processes with one unresolved emission, which enter the real-virtual parts of NNLO cross sections.

In this paper we present a new algorithm for the construction of the interference of two-loop and Born amplitudes, which enters the virtual-virtual parts of NNLO calculations. Also in this case, full scattering processes are handled as sums over individual Feynman diagrams. The amplitude of a generic two-loop diagram Γ has the form

$$\bar{\mathcal{M}}_{2,\Gamma} = C_{2,\Gamma} \int d\bar{q}_1 \int d\bar{q}_2 \frac{\bar{\mathcal{N}}(\bar{q}_1, \bar{q}_2)}{\mathcal{D}(\bar{q}_1, \bar{q}_2)}, \quad (1.4)$$

where $C_{2,\Gamma}$ is a colour factor, and the denominator $\mathcal{D}(\bar{q}_1, \bar{q}_2)$ embodies all propagator denominators depending on the two independent loop momenta \bar{q}_1, \bar{q}_2 . Similarly as in (1.2), for the numerator we apply a tensor decomposition in the two loop momenta,

$$\bar{\mathcal{N}}(\bar{q}_1, \bar{q}_2) = \sum_{r,s} \bar{\mathcal{N}}_{\mu_1 \dots \mu_r, \nu_1 \dots \nu_s} \bar{q}_1^{\mu_1} \dots \bar{q}_1^{\mu_r} \bar{q}_2^{\nu_1} \dots \bar{q}_2^{\nu_s}. \quad (1.5)$$

The main challenges in the development of a fully automated two-loop tool are the efficient construction of the tensor coefficients $\bar{\mathcal{N}}_{\mu_1 \dots \mu_r, \nu_1 \dots \nu_s}$, the reduction of the associated two-loop tensor integrals to a set of master integrals, and the evaluation of these master integrals.

At higher loop levels, a fully automated tensor-integral library in the style of the existing one-loop tools [5, 8, 14] is not yet available. However, concerning the integral reduction, the tools and techniques based on integration-by-parts relations [73] have been greatly advanced in the recent years [74–78]. Furthermore, differential equations methods [79, 80] for the computation of Feynman integrals have greatly expanded the range of available two-loop master integrals (see e.g. [81–87]).

The new algorithm presented in this paper deals with the construction of the two-loop tensor coefficients $\bar{\mathcal{N}}_{\mu_1 \dots \mu_r, \nu_1 \dots \nu_s}$ in (1.5). This is achieved through a process-independent numerical recursion, which includes also the effect of the interference with the Born amplitude as well as the summation over all colour and helicity degrees of freedom. This algorithm has been implemented in a fully automated way within the OPENLOOPS framework for two-loop QCD and QED corrections to any SM process. Due to the generality of the algorithm it can easily be extended to other models. Similarly as in the one-loop case, all tensor coefficients are constructed in four space-time dimensions, and the contributions associated with the $\mathcal{O}(D - 4)$ parts of loop numerators can be reconstructed by means of two-loop rational counterterms [88–90]. More precisely, renormalised two-loop amplitudes in D dimensions can be obtained from two-loop amplitudes with four-dimensional loop numerators together with insertions of one-loop and two-loop counterterms of rational and UV kind into one-loop and tree amplitudes. This approach is fully established for rational terms of UV origin, while we expect that rational terms of infrared (IR) origin cancel in IR-subtracted amplitudes. This cancellation mechanism is well understood at one loop [91] and is presently under investigation at the two-loop level. In this case, we expect that all rational terms of IR origin can be cancelled by means of a simple and process-independent modification of the well-established procedures for the subtraction of two-loop IR divergences.

The paper is organised as follows. In section 2 we outline general aspects of the construction of scattering amplitudes and cross sections at NLO and NNLO. In this context we discuss rational counterterms, as well as the bookkeeping of colour and helicity degrees of freedom. In section 3 we review the OPENLOOPS method for the construction of tree and one-loop amplitudes. This sets the stage for the discussion of the new two-loop algorithm. Section 4 deals with reducible two-loop diagrams, which can be decomposed into one-loop subdiagrams. For this kind of two-loop diagrams we present an algorithm that exploits and extends various features of the OPENLOOPS method at one loop. In section 5 we present the main novelty of this paper, a fully general numerical algorithm for the construction of the tensor coefficients of irreducible two-loop diagrams. We first analyse the structure of a generic two-loop algorithm, and then present the most efficient algorithm that we have identified through a systematic CPU cost analysis. The implementation of this new algorithm in the OPENLOOPS framework and its technical performance are discussed in section 6, where we show that the CPU cost of two-loop tensor coefficients scales linearly in the number of Feynman diagrams and is comparable to the cost of the real-virtual building blocks in a NNLO calculation. Our conclusions are presented in section 7.

2 Scattering amplitudes and partonic cross sections up to NNLO

Partonic cross sections are derived from scattering amplitudes, which in perturbation theory are computed up to a fixed loop order,

$$\bar{\mathcal{M}} = \bar{\mathcal{M}}_0 + \bar{\mathcal{M}}_1 + \bar{\mathcal{M}}_2 + \dots \tag{2.1}$$

The bar marks quantities in $D = 4 - 2\epsilon$ dimensions throughout this paper. The L -loop matrix element $\bar{\mathcal{M}}_L$ is computed as the sum over all L -loop Feynman diagrams Γ of the process at hand,

$$\bar{\mathcal{M}}_L(h) = \sum_{\Gamma} \bar{\mathcal{M}}_{L,\Gamma}(h), \tag{2.2}$$

where the argument h corresponds to the helicity configuration of all external particles as described in section 2.3. Loop integrals can exhibit UV divergences, which are cancelled through the renormalisation procedure. We denote a renormalised amplitude as

$$\mathbf{R}\bar{\mathcal{M}} = \bar{\mathcal{M}}_0 + \mathbf{R}\bar{\mathcal{M}}_1 + \mathbf{R}\bar{\mathcal{M}}_2 + \dots, \tag{2.3}$$

where the renormalisation of the fields, couplings and masses in an L -loop amplitude is implemented through counterterm insertions into lower-loop amplitudes. In a renormalisable model, the finite set of UV counterterms is computed once and for all in the chosen renormalisation scheme. Differential cross sections are obtained from the colour- and helicity-summed scattering probability density

$$\mathcal{W} = \frac{1}{N_{\text{hcs}}} \sum_h \sum_{\text{col}} \left| \mathbf{R}\bar{\mathcal{M}}(h) \right|^2, \tag{2.4}$$

where the initial-state average and symmetry factors for identical final-state particles are encoded in

$$N_{\text{hcs}} = \left(\prod_{i \in \mathcal{S}_{\text{in}}} N_{\text{hel},i} N_{\text{col},i} \right) \left(\prod_{p \in \mathcal{P}_{\text{out}}} n_p! \right). \quad (2.5)$$

Here \mathcal{S}_{in} is the set of initial-state particles, and \mathcal{P}_{out} the set of final-state particle and antiparticle types,² while $N_{\text{hel},i}$ and $N_{\text{col},i}$ are the number of helicity and colour states of particle i , and n_p the number of identical final-state particles of type p .

The probability density (2.4) is split into contributions of different orders in perturbation theory. For processes with $\bar{\mathcal{M}}_0 \neq 0$ the leading-order (LO) and higher-order virtual contributions are

$$\mathcal{W} = \underbrace{\mathcal{W}_{00}}_{\text{LO}} + \underbrace{\mathcal{W}_{01}}_{\text{NLO virtual}} + \underbrace{\mathcal{W}_{02} + \mathcal{W}_{11}}_{\text{NNLO virtual}} + \dots, \quad (2.6)$$

with the L -loop squared and the Born L -loop interference terms

$$\mathcal{W}_{LL} = \frac{1}{N_{\text{hcs}}} \sum_h \sum_{\text{col}} |\mathbf{R} \bar{\mathcal{M}}_L(h)|^2 \quad (L = 0, 1), \quad (2.7)$$

$$\mathcal{W}_{0L} = \frac{1}{N_{\text{hcs}}} \sum_h \sum_{\text{col}} 2 \text{Re} \left[\bar{\mathcal{M}}_0^*(h) \mathbf{R} \bar{\mathcal{M}}_L(h) \right] \quad (L = 1, 2). \quad (2.8)$$

For loop-induced processes \mathcal{W}_{11} is the leading order contribution.

Finite partonic cross sections require, in addition to the UV renormalisation, the inclusion of real-emission contributions in order to cancel final-state collinear and soft divergences, as well as the factorisation of initial-state collinear singularities, which are absorbed into the parton distribution functions. For the amplitudes of scattering processes with X additional unresolved partons, in analogy with (2.1) and (2.6), we use the notation

$$\bar{\mathcal{M}}^{(X)} = \bar{\mathcal{M}}_0^{(X)} + \bar{\mathcal{M}}_1^{(X)} + \dots, \quad (2.9)$$

and

$$\mathcal{W}^{(X)} = \mathcal{W}_{00}^{(X)} + \mathcal{W}_{01}^{(X)} + \dots. \quad (2.10)$$

In addition to (2.6), calculations up to NNLO require the NLO real and NNLO real-virtual contributions

$$\mathcal{W}^{(1)} \Big|_{\text{NNLO}} = \underbrace{\mathcal{W}_{00}^{(1)}}_{\text{NLO real}} + \underbrace{\mathcal{W}_{01}^{(1)}}_{\text{NNLO real-virtual}}, \quad (2.11)$$

as well as the NNLO double-real contribution

$$\mathcal{W}^{(2)} \Big|_{\text{NNLO}} = \underbrace{\mathcal{W}_{00}^{(2)}}_{\text{NNLO real-real}}. \quad (2.12)$$

Partonic cross sections up to NNLO are computed as

$$\hat{\sigma} = \hat{\sigma}_{\text{LO}} + \Delta \hat{\sigma}_{\text{NLO}} + \Delta \hat{\sigma}_{\text{NNLO}} + \dots, \quad (2.13)$$

²Note that particles and antiparticles count as different types here.

with

$$\hat{\sigma}_{\text{LO}} = \int d\Phi_N \mathcal{W}_{00}, \tag{2.14}$$

$$\Delta\hat{\sigma}_{\text{NLO}} = \int d\Phi_N \mathcal{W}_{01} + \int d\Phi_{N+1} \mathcal{W}_{00}^{(1)}, \tag{2.15}$$

$$\Delta\hat{\sigma}_{\text{NNLO}} = \int d\Phi_N (\mathcal{W}_{02} + \mathcal{W}_{11}) + \int d\Phi_{N+1} \mathcal{W}_{01}^{(1)} + \int d\Phi_{N+2} \mathcal{W}_{00}^{(2)}, \tag{2.16}$$

where $d\Phi_{N+X}$ contains the integration measure of the Lorentz-invariant phase space with N resolved and X unresolved particles as well as the flux factor. In general, only the full sum in (2.15) and (2.16) is IR-finite. A number of powerful methods for the subtraction of IR divergences at the level of individual terms in (2.15) and (2.16) is available [92–97].

The program OPENLOOPS 2 [17] supports the calculation of tree-level and one-loop amplitudes, i.e. the contributions (2.7) for $L = 0, 1$ and (2.8) for $L = 1$, and hence also the various NLO and NNLO contributions in (2.11) and (2.12), for any process. This is achieved through automated algorithms that feature high CPU efficiency and numerical stability. In particular, thanks to targeted analytic expansions and a hybrid-precision approach [17, 68], OPENLOOPS 2 guarantees a numerically stable evaluation of the real-virtual NNLO contributions $\mathcal{W}_{01}^{(1)}$ in the full phase space, including the regions where the unresolved radiation becomes highly soft or collinear. This high degree of stability in the IR regions was demonstrated through successful applications of OPENLOOPS 2 to state-of-the-art NNLO calculations based on local subtraction methods, such as in the recent NNLO calculations of $pp \rightarrow 3\text{jets}$ [48, 98].

In this paper, we present a new algorithm for the efficient numerical computation of the Born two-loop interference \mathcal{W}_{02} , defined in (2.8), at the level of tensor-integral coefficients.

2.1 Dimensional regularisation and rational terms

In the OPENLOOPS framework, UV and IR divergences are regularised in the 't Hooft-Veltman scheme [66], where external wave functions and momenta are four-dimensional, while loop momenta, metric tensors and Dirac γ -matrices inside the loops live in $D = 4 - 2\epsilon$ dimensions. The integration measure as well as the denominators of loop integrands are kept in D dimensions throughout, while loop-integrand numerators are split into a four-dimensional part and a remainder of $\mathcal{O}(\epsilon)$. Contributions stemming from the former part are referred to as $D_n = 4$ dimensional, where D_n denotes the dimensionality of the loop numerator, and are calculated with numerical algorithms where the loop momenta, metric tensors and γ -matrices in the numerator are handled in four dimensions. The remaining $(D - 4)$ -dimensional numerator parts give rise to additional rational contributions, which originate from their interplay with the divergences of loop integrals and can be reconstructed through process-independent counterterms.

At one loop, such rational terms originate only from UV poles [91], and the corresponding counterterms are known for the full SM and for a variety of other models [69–72]. Using rational counterterms, renormalised one-loop amplitudes in D dimensions can be obtained from quantities in $D_n = 4$ dimensions through the formula

$$\mathbf{R}\bar{\mathcal{M}}_1(h) = \mathcal{M}_1(h) + \mathcal{M}_{0,1}^{(\text{CT})}(h). \tag{2.17}$$

Here and in the following, loop amplitudes $\bar{\mathcal{M}}_L$ carrying a bar are fully D -dimensional as introduced in (2.1), while amplitudes \mathcal{M}_L without a bar are in $D_n = 4$ dimensions, i.e. they are computed with four-dimensional integrand numerators and D -dimensional denominators. The first term on the r.h.s. of (2.17) is the unrenormalised amplitude in $D_n = 4$, and $\mathcal{M}_{0,1}^{(\text{CT})}$ stands for the tree-level amplitude with all relevant insertions of UV and rational one-loop counterterms.

At two loops, as recently shown in [88, 89], renormalised amplitudes in D dimensions can be obtained from amplitudes in $D_n = 4$ dimensions through a general formula of the form

$$\mathbf{R}\bar{\mathcal{M}}_2(h) = \mathcal{M}_2(h) + \mathcal{M}_{1,1}^{(\text{CT})}(h) + \mathcal{M}_{0,2}^{(\text{CT})}(h) + \mathcal{M}_{0,1,1}^{(\text{CT})}(h). \quad (2.18)$$

Here the first term on the r.h.s. is the unrenormalised two-loop amplitude in $D_n = 4$, while each of the three additional contributions embodies standard counterterms for the subtraction of UV divergences in combination with rational counterterms for the reconstruction of the contributions of the $(D - 4)$ -dimensional parts of loop numerators. The term $\mathcal{M}_{1,1}^{(\text{CT})}(h)$ denotes the one-loop amplitude with all relevant one-loop counterterm insertions, while $\mathcal{M}_{0,2}^{(\text{CT})}(h)$ and $\mathcal{M}_{0,1,1}^{(\text{CT})}(h)$ correspond, respectively, to the tree-level amplitudes with single two-loop and double one-loop counterterm insertions. Similarly as for the related UV counterterms, also two-loop rational counterterms of UV origin are process-independent [88]. The explicit expressions for all two-loop rational counterterms of UV origin in QED and for QCD corrections to the full SM have been derived in [88–90].

Two-loop rational terms originating from the interplay of $(D - 4)$ -dimensional numerator parts and IR divergences are currently under investigation. As anticipated in the introduction, we expect that — at the level of IR-subtracted two-loop amplitudes — all rational terms of IR origin can be cancelled by means of a simple and process-independent modification of the well-established procedures for the subtraction of two-loop IR divergences.

2.2 Scattering amplitudes and colour factors

In the original OPENLOOPS algorithm, as well as in the new algorithm, L -loop matrix elements \mathcal{M}_L are computed as sums of all Feynman diagrams Γ of the scattering process,

$$\mathcal{M}_L(h) = \sum_{\Gamma} \mathcal{M}_{L,\Gamma}(h). \quad (2.19)$$

The contribution from each diagram is factorised into a colour factor $C_{L,\Gamma}$ and a colour-stripped helicity amplitude $\mathcal{A}_{L,\Gamma}(h)$,

$$\mathcal{M}_{L,\Gamma}(h) = C_{L,\Gamma} \mathcal{A}_{L,\Gamma}(h). \quad (2.20)$$

Since each quartic gluon vertex gives rise to three independent colour structures, Feynman diagrams Γ that involve n_q quartic gluon vertices are decomposed into

$$\mathcal{M}_{L,\Gamma} = \sum_{j=1}^{3^{n_q}} C_{L,\Gamma_j} \mathcal{A}_{L,\Gamma_j}, \quad (2.21)$$

and, in the following, each colour-factorised contribution Γ_j is handled as a different Feynman diagram.

The colour structures $C_{L,\Gamma}$ are algebraically reduced to a standard colour basis $\{\mathcal{C}_i\}$ (see [17] for details),

$$C_{L,\Gamma} = \sum_i a_{L,i}(\Gamma) \mathcal{C}_i, \quad (2.22)$$

and Born amplitudes are cast in the form

$$\mathcal{M}_0(h) = \sum_i \mathcal{A}_0^{(i)}(h) \mathcal{C}_i. \quad (2.23)$$

Colour-summed interferences are built by means of the colour-interference matrix

$$\mathcal{K}_{ij} = \sum_{\text{col}} \mathcal{C}_i^\dagger \mathcal{C}_j. \quad (2.24)$$

For example, the LO probability density is computed as

$$\mathcal{W}_{00} = \frac{1}{N_{\text{hcs}}} \sum_h \sum_{i,j} \left(\mathcal{A}_0^{(i)}(h) \right)^* \mathcal{K}_{ij} \mathcal{A}_0^{(j)}(h). \quad (2.25)$$

The contribution of an L -loop diagram Γ with $L \geq 1$ to (2.8) is computed as

$$\begin{aligned} \mathcal{W}_{0L,\Gamma} &= \frac{1}{N_{\text{hcs}}} \sum_h \sum_{\text{col}} 2 \operatorname{Re} \left[\mathcal{M}_0^*(h) \mathcal{M}_{L,\Gamma}(h) \right] \\ &= \frac{1}{N_{\text{hcs}}} \operatorname{Re} \left(\sum_h \mathcal{U}_{0,\Gamma}(h) \mathcal{A}_{L,\Gamma}(h) \right), \end{aligned} \quad (2.26)$$

where the colour-stripped loop amplitude $\mathcal{A}_{L,\Gamma}(h)$ is factorised from the colour-Born interference term

$$\mathcal{U}_{0,\Gamma}(h) = 2 \sum_{\text{col}} \mathcal{M}_0^*(h) C_{L,\Gamma} = 2 \sum_j \left[\sum_i \left(\mathcal{A}_0^{(i)}(h) \right)^* \mathcal{K}_{ij} \right] a_{L,j}(\Gamma). \quad (2.27)$$

This colour treatment is implemented in the public OPENLOOPS code [17] as well as in the new two-loop algorithm presented in this paper.

2.3 Helicity bookkeeping

In this section we define the helicity labels used in OPENLOOPS at all loop orders, following the notation of [68]. For the bookkeeping of external momenta and helicities in a process with N scattering particles we introduce the set of particle indices

$$\mathcal{E} = \{1, 2, \dots, N\}. \quad (2.28)$$

To characterise the helicity configurations of individual particles $p \in \mathcal{E}$ we use the labels

$$\lambda_p = \begin{cases} 1, 3 & \text{for fermions with helicity } s = -1/2, 1/2 \\ 1, 2, 3 & \text{for gauge bosons with } s = -1, 0, 1 \\ 0 & \text{for scalars with } s = 0 \end{cases} \quad \forall p \in \mathcal{E}. \quad (2.29)$$

The configuration $\lambda_p = 0$ is also used to characterise unpolarised particles, i.e. fermions or gauge bosons whose helicity is still unassigned at a certain stage of the calculation or has already been summed over. We use a helicity numbering scheme based on the labels

$$\bar{h}_p = \lambda_p 4^{p-1}, \tag{2.30}$$

which correspond to a quaternary number with $\lambda_p \in \{0, 1, 2, 3\}$ as p^{th} -last digit and all other digits equal to zero. This scheme allows us to derive the helicity labels of any set of particles as the sum of the helicity labels of its disjoint subsets, and hence as the sum of the helicity labels of all its particles. A set of particles $\mathcal{E}_a = \{p_{a_1}, \dots, p_{a_n}\}$ has the helicity label,

$$h_a = \sum_{p \in \mathcal{E}_a} \bar{h}_p. \tag{2.31}$$

In particular, the global helicity label of the scattering process is

$$h = \sum_{p \in \mathcal{E}} \bar{h}_p. \tag{2.32}$$

For two disjoint sets of particles $\mathcal{E}_b, \mathcal{E}_c$ the helicity label of the combined set $\mathcal{E}_a = \mathcal{E}_b \cup \mathcal{E}_c$ is

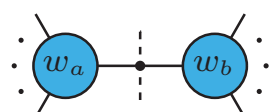
$$h_a = h_b + h_c. \tag{2.33}$$

3 Tree and one-loop amplitudes in OPENLOOPS

In this section we review the algorithm for the construction of tree and one-loop amplitudes that is implemented in the public OPENLOOPS program. This sets the stage for the discussion of the new two-loop algorithm.

3.1 Tree-level amplitudes

At tree level, the colour-stripped amplitude of a Feynman diagram Γ is decomposed into two subtrees connected by a certain off-shell propagator,³

$$\mathcal{A}_{0,\Gamma}(h) = \text{diagram} = w_a^{\sigma_a}(k_a, h_a) \delta_{\sigma_a \sigma_b} \tilde{w}_b^{\sigma_b}(k_b, h_b). \tag{3.1}$$


Here $k_a = -k_b$ and σ_a, σ_b are the off-shell momenta and Lorentz/spinor indices⁴ of the subtrees, while h_a, h_b denote the helicity labels of the subsets of external on-shell particles connected to the respective subtrees, and

$$h = h_a + h_b \tag{3.2}$$

corresponds to the global helicity of all scattering particles. The subtree w_a is bounded by external wave functions and the off-shell propagator, at which the diagram was cut, while

³The Feynman diagrams in this paper are drawn with AXODRAW [99].

⁴Spinor and Lorentz indices have values from 1 to 4 while the indices associated with a scalar propagator only have the value 1.

For later convenience we apply a label (i) to various q_i -dependent building blocks of a loop integrand, where q_i is either a loop momentum or a linear combination of independent loop momenta. At one-loop level, $i = 1$, since there is only a single loop momentum. For the integration measure in loop-momentum space we define the shorthand

$$\int d\bar{q}_i = \mu^{2\epsilon} \int \frac{d^D \bar{q}_i}{(2\pi)^D}, \tag{3.8}$$

where μ is the scale of dimensional regularisation.

In OPENLOOPS the loop numerator is constructed through a numerical recursion that exploits its factorisation into loop segments. A loop segment consists of a loop propagator, one adjacent triple or quartic vertex, and the one or two external subtrees connected to this loop vertex. In the case of a triple vertex the loop segment has the form

$$\left[S_a^{(i)}(q_i, h_a^{(i)}) \right]_{\beta_{a-1}^{(i)}}^{\beta_a^{(i)}} = \begin{array}{c} \text{circle } w_a^{(i)} \\ \downarrow k_{ia} \\ \text{line } \beta_{a-1}^{(i)} \rightarrow \beta_a^{(i)} \end{array} = \left\{ [Y_{ia}^\sigma]_{\beta_{a-1}^{(i)}}^{\beta_a^{(i)}} + [Z_{ia,\nu}^\sigma]_{\beta_{a-1}^{(i)}}^{\beta_a^{(i)}} q_i^\nu \right\} w_{a\sigma}^{(i)}(k_{ia}, h_a^{(i)}), \tag{3.9}$$

where an external subtree $w_a^{(i)}$ is connected to a loop vertex and the adjacent loop propagator. The index $a \in [1, N_i]$ corresponds to the position of the segment along the loop. In renormalisable theories, a segment can be written as a rank-one polynomial in the loop momentum with coefficients Y and Z . The indices $\beta_a^{(i)}$ are, depending on the particle type in the loop propagator, Lorentz indices (gauge bosons) or spinor indices (fermions) with $\beta_a^{(i)} = 1, \dots, 4$. For scalar particles (ghosts, scalars) the index has a fixed value $\beta_a^{(i)} = 1$. Segments associated with a quartic vertex involve two subtrees, $w_{a_1}^{(i)}$ and $w_{a_2}^{(i)}$, and are of rank zero in q_1 ,

$$\left[S_a^{(i)}(q_i, h_a^{(i)}) \right]_{\beta_{a-1}^{(i)}}^{\beta_a^{(i)}} = \begin{array}{c} \text{circle } w_{a_1}^{(i)} \quad w_{a_2}^{(i)} \\ \swarrow k_{ia_1} \quad \searrow k_{ia_2} \\ \text{line } \beta_{a-1}^{(i)} \rightarrow \beta_a^{(i)} \end{array} = [Y_{ia}^{\sigma_1\sigma_2}]_{\beta_{a-1}^{(i)}}^{\beta_a^{(i)}} w_{a_1\sigma_1}^{(i)}(k_{ia_1}, h_{a_1}^{(i)}) w_{a_2\sigma_2}^{(i)}(k_{ia_2}, h_{a_2}^{(i)}) \tag{3.10}$$

with $h_a^{(i)} = h_{a_1}^{(i)} + h_{a_2}^{(i)}$ and $k_{ia} = k_{ia_1} + k_{ia_2}$.

At one loop, the numerator is computed by cut-opening the loop at one propagator, which results in a tree-like object consisting of a product of N_1 loop segments,

$$\left[\mathcal{N}^{(1)}(q_1, h) \right]_{\beta_0^{(1)}}^{\beta_{N_1}^{(1)}} = \begin{array}{c} \text{circle with } w_{N_1}^{(1)} \text{ and } w_1^{(1)} \\ \text{cut at } \beta_0^{(1)} \text{ and } \beta_{N_1}^{(1)} \end{array} \tag{3.11}$$

$$= \left[S_1^{(1)}(q_1, h_1^{(1)}) \right]_{\beta_0^{(1)}}^{\beta_1^{(1)}} \left[S_2^{(1)}(q_1, h_2^{(1)}) \right]_{\beta_1^{(1)}}^{\beta_2^{(1)}} \cdots \left[S_{N_1}^{(1)}(q_1, h_{N_1}^{(1)}) \right]_{\beta_{N_1-1}^{(1)}}^{\beta_{N_1}^{(1)}},$$

where $\beta_0^{(1)}, \beta_{N_1}^{(1)}$ are the Lorentz/spinor indices of the cut propagator, and $h_a^{(1)}$ describes the helicity configuration of the a -th subtree,

$$h_a^{(1)} = \sum_{p \in \mathcal{E}_a^{(i)}} \bar{h}_p, \tag{3.12}$$

where $\mathcal{E}_a^{(i)} \subseteq \mathcal{E}$ is the corresponding set of external particles.

The loop numerator is constructed through recursive matrix multiplications,

$$\mathcal{N}_n^{(1)}(q_1, \hat{h}_n^{(1)}) = \mathcal{N}_{n-1}^{(1)}(q_1, \hat{h}_{n-1}^{(1)}) S_n^{(1)}(q_1, h_n^{(1)}), \tag{3.13}$$

which are applied for $n = 1, \dots, N_1$, starting from the initial condition $\mathcal{N}_0 = \mathbb{1}$. The label

$$\hat{h}_n^{(1)} = \sum_{a=1}^n h_a^{(1)} \tag{3.14}$$

describes the helicity configuration of the external legs entering the first n segments, and $\hat{h}_n^{(1)} = \hat{h}_{n-1}^{(1)} + h_n^{(1)}$. The operations (3.13) are referred to as dressing steps, and the partially dressed numerator

$$\mathcal{N}_n^{(1)}(q_1, \hat{h}_n^{(1)}) = \prod_{a=1}^n S_a^{(1)}(q_1, h_a^{(1)}) = \begin{array}{c} \text{blue blob } w_1^{(1)} \quad \text{blue blob } w_2^{(1)} \quad \text{blue blob } w_n^{(1)} \quad \text{grey blob } w_{n+1}^{(1)} \quad \text{grey blob } w_{N_1}^{(1)} \\ | \quad | \quad | \quad | \quad | \\ \beta_0^{(1)} \text{---} D_1^{(1)} \text{---} D_2^{(1)} \text{---} \dots \text{---} D_n^{(1)} \text{---} \beta_n^{(1)} \text{---} D_{n+1}^{(1)} \text{---} \dots \text{---} D_{N_1}^{(1)} \text{---} \beta_{N_1}^{(1)} \end{array} \tag{3.15}$$

is called an open loop. In this schematic representation, blue and grey blobs correspond to the dressed loop segments and those that remain to be dressed, respectively. Each open loop is a polynomial in q_1 ,

$$\mathcal{N}_n^{(1)}(q_1, \hat{h}_n^{(1)}) = \sum_{r=0}^R \mathcal{N}_{n; \mu_1 \dots \mu_r}^{(1)}(\hat{h}_n^{(1)}) q_1^{\mu_1} \dots q_1^{\mu_r}, \tag{3.16}$$

with rank $R \leq n$, and the dressing recursion (3.13) is implemented at the level of the tensor coefficients $\mathcal{N}_{n; \mu_1 \dots \mu_r}^{(1)}$. The explicit form of a dressing step for a segment (3.9) with a three-point vertex is

$$\begin{aligned} \left[\mathcal{N}_{n; \mu_1 \dots \mu_r}^{(1)}(\hat{h}_n^{(1)}) \right]_{\beta_0^{(1)}}^{\beta_n^{(1)}} &= \left\{ \left[\mathcal{N}_{n-1; \mu_1 \dots \mu_r}^{(1)}(\hat{h}_{n-1}^{(1)}) \right]_{\beta_0^{(1)}}^{\beta_{n-1}^{(1)}} \left[Y_{1n}^\sigma \right]_{\beta_{n-1}^{(1)}}^{\beta_n^{(1)}} \right. \\ &\quad \left. + \left[\mathcal{N}_{n-1; \mu_2 \dots \mu_r}^{(1)}(\hat{h}_{n-1}^{(1)}) \right]_{\beta_0^{(1)}}^{\beta_{n-1}^{(1)}} \left[Z_{1n, \mu_1}^\sigma \right]_{\beta_{n-1}^{(1)}}^{\beta_n^{(1)}} \right\} w_{n\sigma}^{(1)}(k_n, h_n^{(1)}). \end{aligned} \tag{3.17}$$

For an efficient implementation the $\mu_1 \dots \mu_r$ indices are symmetrised. In the final step of the dressing algorithm, the trace is taken over the indices $\beta_0^{(1)}, \beta_{N_1}^{(1)}$,

$$\mathcal{N}_{\mu_1 \dots \mu_r}^{(1)}(h) = \text{Tr} \left[\mathcal{N}_{N_1; \mu_1 \dots \mu_r}^{(1)}(h) \right], \tag{3.18}$$

and for the amplitude of the colour-stripped Feynman diagram (3.5) we obtain

$$\mathcal{A}_{1, \Gamma}(h) = \sum_{r=0}^N \mathcal{N}_{\mu_1 \dots \mu_r}^{(1)}(h) \int d\bar{q}_1 \frac{q_1^{\mu_1} \dots q_1^{\mu_r}}{D_0 \dots D_{N-1}}. \tag{3.19}$$

The tensor integrals on the r.h.s. can be reduced with external libraries such as COLLIER [14]. Alternatively, they can be reduced with the on-the-fly method of [68], where dressing steps are interleaved with reduction steps in such a way that the tensor rank remains low at all stages of the calculation. The on-the-fly reduction is the default method in OPENLOOPS 2 [17].

3.3 Born-loop interference

As pointed out in [68], for the efficient construction of the helicity- and colour-summed Born-loop interference defined in (2.26) it is convenient to absorb the colour-Born interference factor (2.27) into the loop numerator. In this approach, instead of the original helicity-dependent numerator we construct the helicity summed quantity

$$\mathcal{U}(q_1) = \sum_h \mathcal{U}_0(h) \mathcal{N}^{(1)}(q_1, h) = \sum_h \mathcal{U}_0(h) \prod_{a=1}^{N_1} S_a^{(1)}(h_a^{(1)}) \quad (3.20)$$

with the colour-Born interference

$$\mathcal{U}_0(h) = 2 \sum_{\text{col}} \mathcal{M}_0^*(h) C_{1,\Gamma}, \quad (3.21)$$

where, again, the label Γ is kept implicit. Exploiting the factorisation of $\mathcal{U}(q_1)$ into $\mathcal{U}_0(h)$ and loop segments, one can write

$$\mathcal{U}(q_1) = \sum_{h_{N_1}^{(1)}} \left[\dots \sum_{h_2^{(1)}} \left[\sum_{h_1^{(1)}} \mathcal{U}_0(h) S_1^{(1)}(h_1^{(1)}) \right] S_2^{(1)}(h_2^{(1)}) \dots \right] S_{N_1}^{(1)}(h_{N_1}^{(1)}), \quad (3.22)$$

where helicity sums are partially factorised at the level of individual segments. More precisely, each time the numerator is dressed with a new segment $S_n^{(1)}(q_1, h_n^{(1)})$, the helicity d.o.f. of the associated subtree $w_n^{(1)}(h_n^{(1)})$ can be summed, and the subsequent dressing steps depend only on the helicities of the yet undressed segments. This way of constructing the loop numerator (3.20) corresponds to a recursive dressing algorithm

$$\mathcal{U}_n(q_1, \check{h}_n^{(1)}) = \sum_{h_n^{(1)}} \mathcal{U}_{n-1}(q_1, \check{h}_{n-1}^{(1)}) S_n^{(1)}(q_1, h_n^{(1)}), \quad (3.23)$$

where the initial condition (3.21) is used, and the helicity label

$$\check{h}_n^{(1)} = h - \hat{h}_n^{(1)} = \sum_{a=n+1}^{N_1} h_a^{(1)} \quad (3.24)$$

corresponds to the helicity configuration of all undressed segments. To keep track of the polynomial dependence of the partially dressed numerators on the loop momentum q_1 , a tensorial representation similar to (3.16) is used,

$$\mathcal{U}_n(q_1, \check{h}_n^{(1)}) = \sum_{r=0}^R \mathcal{U}_{n; \mu_1 \dots \mu_r}(\check{h}_n^{(1)}) q_1^{\mu_1} \dots q_1^{\mu_r}, \quad (3.25)$$

and the dressing recursion is implemented at the level of tensor coefficients in the same form as in (3.17), but with an additional sum over $h_n^{(1)}$ as in (3.23). This on-the-fly helicity summation approach guarantees the smallest possible number of helicity configurations in the last dressing steps, where the rank in q_1 , and hence the number of tensor coefficients and required matrix multiplications, is highest.

At the end of the dressing recursion, i.e. after step $n = N_1$, in analogy with (3.18) we take the trace

$$\mathcal{U}_{\mu_1 \dots \mu_r} = \text{Tr} \left[\mathcal{U}_{N_1; \mu_1 \dots \mu_r}(\check{h}_{N_1}^{(1)}) \right], \tag{3.26}$$

where $\check{h}_{N_1}^{(1)} = 0$ since no undressed segments are left. Finally, combining the tensor coefficients with the associated tensor integrals, one arrives at

$$\mathcal{W}_{01, \Gamma} = \frac{1}{N_{\text{hcs}}} \text{Re} \left[\sum_{r=0}^N \mathcal{U}_{\mu_1 \dots \mu_r} \int d\bar{q}_1 \frac{q_1^{\mu_1} \dots q_1^{\mu_r}}{D_0 \dots D_{N-1}} \right], \tag{3.27}$$

which corresponds to the one-loop virtual scattering probability density as defined in (2.26).

The reduction of the tensor integrals can be performed with external libraries, such as COLLIER [14], or with the on-the-fly reduction method of [68].

4 Reducible two-loop integrands

In this and the following section we present a new automated algorithm that extends the construction of the tensor coefficients (3.26) to two loops. Specifically, we will focus on the construction of tensor coefficients for unrenormalised two-loop amplitudes in $D_n = 4$ dimensions, i.e. with four-dimensional numerators, while all other ingredients of renormalised two-loop amplitudes (2.18) can be obtained from one-loop and tree amplitudes with UV and rational counterterm insertions. All amplitudes are split into individual Feynman diagrams, and their colour structures are factorised as described in (2.19)–(2.21). Thus we will concentrate on the construction of tensor coefficients for colour-stripped unrenormalised two-loop diagrams.

At two loops we categorise Feynman diagrams into reducible and irreducible ones as illustrated in figure 1. This categorisation is based on the structure of loop chains. Such chains consist of loop propagators that are linked to each other and to external subtrees through connecting vertices. More precisely, a loop chain \mathcal{C}_i contains the propagators that depend only on a certain loop momentum q_i when all external momenta are set to zero.

Two-loop diagrams that involve only two loop chains $\mathcal{C}_1, \mathcal{C}_2$, with two independent loop momenta q_1, q_2 , are categorised as reducible. This category is further split into two-loop diagrams of type Red2 and Red1 depending on how the two chains are connected to each other. As indicated in figure 1, in reducible two-loop diagrams of type Red2 the chains \mathcal{C}_1 and \mathcal{C}_2 are connected through a tree structure P , which is referred to as bridge, while the type Red1 corresponds to the case where the two chains are connected through a single quartic vertex \mathcal{V}_4 . Since reducible two-loop diagrams factorise into two one-loop subdiagrams, for their calculation one can exploit various functions of the one-loop OPENLOOPS framework. An algorithm for the construction of reducible two-loop integrands of type Red2 and Red1 is presented in sections 4.1–4.2.




Figure 1. Categorisation of two-loop topologies into reducible (Red) and irreducible (Irred) ones. Reducible two-loop topologies are further split into two subcategories, where two one-loop subdiagrams are either connected to a bridge through two vertices (Red2) or are attached to each other through a common quartic vertex (Red1). Similarly as in section 3 the blue blobs denote tree structures that are connected to internal and external lines (the latter are not shown). In the Red2 topology the blue blob labelled P can be connected only to the two visible internal lines or also to additional external lines.

Two-loop diagrams that involve three loop chains $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$ are categorised as irreducible. In this case the three chains are connected to each other by two vertices $\mathcal{V}_0, \mathcal{V}_1$, and of the three loop momenta q_1, q_2, q_3 only two are linearly independent. The loop momenta can always be chosen in such a way that $q_1 + q_2 + q_3 = 0$. An algorithm for the construction of irreducible two-loop integrands is presented in section 5.

In general, the algorithms presented in sections 4.1–4.2 and 5 compute tensor coefficients that corresponds to the helicity and colour-summed interference (2.26) of the full Born amplitude with the integrand of individual two-loop diagrams.

4.1 Reducible two-loop integrands of class Red2

A reducible two-loop diagram of class Red2 factorises into two chains \mathcal{C}_i and a bridge P as depicted in figure 1. The loop chains \mathcal{C}_i each consist of N_i loop propagators ($i = 1, 2$), and the external subtrees $w_1^{(i)}, \dots, w_{N_i-1}^{(i)}$ connected to the triple or quartic loop vertices as in (3.9) and (3.10). The bridge P is a tree structure connected to each of the two chains with a corresponding vertex, which we call the bridge vertex of that chain. In general, the bridge is also connected to a subset of the external particles of the scattering process. The colour-stripped amplitude of a reducible diagram Γ has the generic form

$$\begin{aligned}
 \mathcal{A}_{2,\Gamma}(h) &= \int d\bar{q}_1 \frac{\text{Tr}[\mathcal{N}^{(1)}(q_1, h^{(1)})]^{\alpha_1}}{\mathcal{D}^{(1)}(\bar{q}_1)} P_{\alpha_1 \alpha_2}(h^{(B)}) \int d\bar{q}_2 \frac{\text{Tr}[\mathcal{N}^{(2)}(q_2, h^{(2)})]^{\alpha_2}}{\mathcal{D}^{(2)}(\bar{q}_2)}, \quad (4.1)
 \end{aligned}$$




Figure 2. Cut open reducible two-loop diagram of type Red2.

with the denominator chains

$$\mathcal{D}^{(i)}(\bar{q}_i) = D_0^{(i)}(\bar{q}_i) \cdots D_{N_i-1}^{(i)}(\bar{q}_i), \quad D_a^{(i)}(\bar{q}_i) = (\bar{q}_i + p_{ia})^2 - m_{ia}^2, \quad (4.2)$$

where p_{ia} and m_{ia} are the external momenta and internal masses along the chain \mathcal{C}_i . The Lorentz/spinor indices α_1, α_2 connect the bridge $P_{\alpha_1\alpha_2}$ to the chains $\mathcal{C}_1, \mathcal{C}_2$. The external subtrees $w_a^{(i)}$ depicted in (4.1) play the two-fold role of single subtrees or pairs of subtrees connected to the loop chains via triple and quartic vertices, respectively. The global helicity configuration is given by

$$h = h^{(1)} + h^{(B)} + h^{(2)}, \quad (4.3)$$

with the chain helicities

$$h^{(i)} = \sum_{a=1}^{N_i} h_a^{(i)} \quad (i = 1, 2), \quad (4.4)$$

and the bridge helicity $h^{(B)}$, which corresponds to the helicity configuration of all external particles connected to P .

Reducible two-loop diagrams can be efficiently constructed by the following algorithm, which uses and extends functions of the one-loop OPENLOOPS algorithm.

Ordering and cutting rule. Chain \mathcal{C}_1 is chosen to be the chain with less loop propagators.⁷ As illustrated in figure 2, the construction of the loop numerator is organised by cut-opening the two one-loop subdiagrams between the bridge vertices and the $D_0^{(i)}$ propagators. As usual the corresponding external momenta are set to $p_{i0} = 0$. The position of the cut (left or right of the bridge vertex) is determined in a unique way based on the properties of the subtrees. The dressing directions of the loop chains correspond to the direction of the loop momenta q_i .

⁷In case there is more than one chain with the minimum number of propagators, then the one with the minimum number of helicity d.o.f. is chosen to be \mathcal{C}_1 . If in turn there is more than one chain with the minimum number of propagators and helicity d.o.f., then \mathcal{C}_1 is selected based on the particle content of the external subtrees that are connected to the various chains.

Step 1 — construction of chain \mathcal{C}_1 . The numerator of the chain \mathcal{C}_1 factorises into loop segments,

$$\text{Tr} \left[\mathcal{N}^{(1)}(q_1, h^{(1)}) \right]^{\alpha_1} = S_0^{(1)}(q_1) S_1^{(1)}(q_1, h_1^{(1)}) \cdots S_{N_1-1}^{(1)}(q_1, h_{N_1-1}^{(1)}) \left[S_{N_1}^{(1)}(q_1, h_{N_1}^{(1)}) \right]^{\alpha_1}, \quad (4.5)$$

where the last segment carries the Lorentz/spinor index α_1 that connects it to the bridge, while the indices connecting the segments to each other are kept implicit.

The recursive construction of (4.5) starts from the segment $S_0^{(1)}$, which is simply the propagator numerator corresponding to $D_0^{(1)}$. The segments $S_a^{(1)}$ for $a = 1, \dots, N_1 - 1$ consist of a loop vertex together with the attached subtree(s) $w_a^{(1)}$, and the attached propagator numerator corresponding to $D_a^{(1)}$. The last segment $S_{N_1}^{(1)}$ consists of the bridge vertex on chain \mathcal{C}_1 , and in the case of a four-vertex a subtree $w_{N_1}^{(1)}$. This decomposition ensures that the open index α_1 is only introduced in the last step of the construction of \mathcal{C}_1 .

The first chain is constructed starting from the initial condition $\mathcal{N}_{-1}^{(1)} = \mathbb{1}$ with dressing steps

$$\mathcal{N}_n^{(1)}(q_1, \hat{h}_n^{(1)}) = \mathcal{N}_{n-1}^{(1)}(q_1, \hat{h}_{n-1}^{(1)}) \cdot S_n^{(1)}(q_1, h_n^{(1)}) \quad \text{for } n = 0, \dots, N_1, \quad (4.6)$$

where we use the partial chain helicities defined in (3.14) with $\hat{h}_n^{(1)} = \hat{h}_{n-1}^{(1)} + h_n^{(1)}$. Similarly as in (3.13)–(3.18) the recursion is implemented in terms of tensor coefficients and, taking the trace after the last step, yields the coefficients

$$\left[\mathcal{N}_{\mu_1 \dots \mu_r}^{(1)}(h^{(1)}) \right]^{\alpha_1} = \text{Tr} \left[\mathcal{N}_{N_1; \mu_1 \dots \mu_r}^{(1)}(\hat{h}_{N_1}^{(1)}) \right]^{\alpha_1}, \quad (4.7)$$

where $h^{(1)} = \hat{h}_{N_1}^{(1)}$. Contracting all coefficients with corresponding one-loop tensor integrals results into the closed first loop

$$\begin{aligned} \left[P_{-1}(h^{(1)}) \right]^{\alpha_1} &= \int d\bar{q}_1 \frac{\text{Tr} \left[\mathcal{N}^{(1)}(q_1, h^{(1)}) \right]^{\alpha_1}}{\mathcal{D}^{(1)}(\bar{q}_1)} \\ &= \sum_{r=0}^{N_1} \left[\mathcal{N}_{\mu_1 \dots \mu_r}^{(1)}(h^{(1)}) \right]^{\alpha_1} \int d\bar{q}_1 \frac{q_1^{\mu_1} \cdots q_1^{\mu_r}}{\mathcal{D}^{(1)}(\bar{q}_1)}, \end{aligned} \quad (4.8)$$

which serves as the starting point for the bridge construction. Choosing \mathcal{C}_1 to be the shorter chain, i.e. $N_1 \leq N_2$, ensures that the number of helicity states $h^{(1)}$, for which the tensor coefficients need to be computed and contracted with a tensor integral, is minimal. Step 1 is performed for all diagrams, recycling partially or entirely constructed chains wherever possible.

Step 2 — bridge construction. The bridge involves various segments $S_a^{(B)}$ with $a = 0, \dots, N_B$ as depicted in figure 2. As a starting point for its construction, the segment $S_0^{(B)}$, which consists solely of the first bridge propagator, is attached to the one-loop subdiagram (4.8),

$$P_0(\hat{h}_0^{(B)}) = P_{-1}(\hat{h}_0^{(B)}) \cdot S_0^{(B)}. \quad (4.9)$$

Here the global helicity $h^{(1)}$ of the subdiagram (4.8) has been renamed as $\hat{h}_0^{(B)}$. The remaining bridge segments $S_n^{(B)}(h_n^{(B)})$ with $n > 0$ consist of a vertex and a bridge propagator together with one or two external subtrees $w_n^{(B)}$. Note that, in contrast to loop segments,

these tree segments contain also the associated propagator denominators, which depend solely on external momenta and internal masses. The label $h_n^{(B)}$ corresponds to the helicity configuration of the external particles in $w_n^{(B)}$. These subtrees are recursively attached to (4.9) through dressing steps

$$P_n(\hat{h}_n^{(B)}) = P_{n-1}(\hat{h}_{n-1}^{(B)}) \cdot S_n^{(B)}(h_n^{(B)}) \tag{4.10}$$

for $n = 1, \dots, N_B$, which are implemented in the OPENLOOPS tree-level algorithm. The global helicity configuration for this partially dressed bridge contracted with the first one-loop subdiagram corresponds to

$$\hat{h}_n^{(B)} = \hat{h}_{n-1}^{(B)} + h_n^{(B)} = h^{(1)} + \sum_{a=1}^n h_a^{(B)}. \tag{4.11}$$

Recycling opportunities are systematically exploited whenever partially dressed bridges connected to the first loop occur in multiple diagrams.

The final result of the recursion (4.10) corresponds to the first two building blocks on the r.h.s. of (4.1), i.e.

$$\left[P_{N_B}(\hat{h}^{(B)}) \right]_{\alpha_2} = \int d\bar{q}_1 \frac{\text{Tr} \left[\mathcal{N}^{(1)}(q_1, h^{(1)}) \right]^{\alpha_1}}{\mathcal{D}^{(1)}(\bar{q}_1)} P_{\alpha_1 \alpha_2}(h^{(B)}), \tag{4.12}$$

where we define

$$\hat{h}^{(B)} = \hat{h}_{N_B}^{(B)} = h^{(1)} + h^{(B)}. \tag{4.13}$$

Step 3 — construction of chain \mathcal{C}_2 . The second chain is constructed starting from the loop segment associated with the propagator denominator $D_1^{(2)}$. In this loop segment we include, as an effective external subtree, the full bridge-loop combination (4.12). If the bridge and the chain \mathcal{C}_2 are connected by a triple vertex, according to (3.9) the first loop segment has the form

$$S_1^{(2)}(q_2, h_1^{(2)}) = \left(Y_{21}^{\alpha_2} + Z_{21,\nu}^{\alpha_2} q_2^\nu \right) \left[P_{N_B}(\hat{h}^{(B)}) \right]_{\alpha_2}, \tag{4.14}$$

with $h_1^{(2)} = \hat{h}^{(B)}$. Here Y and Z embody the connecting triple vertex together with the propagator numerator associated with $D_1^{(2)}$. For a quartic vertex, similarly as in (3.10) we have

$$S_1^{(2)}(q_2, h_1^{(2)}) = Y_{21}^{\alpha_2 \alpha'_2} \left[P_{N_B}(\hat{h}^{(B)}) \right]_{\alpha_2} \left[w_{1'}^{(2)}(h_{1'}^{(2)}) \right]_{\alpha'_2}, \tag{4.15}$$

where $w_{1'}^{(2)}$ denotes an additional subtree that is connected to the quartic vertex, and $h_{1'}^{(2)} = \hat{h}^{(B)} + h_{1'}^{(2)}$. As a result of (4.14)–(4.15) the index α_2 is saturated at the beginning of the construction of the chain \mathcal{C}_2 , which renders the subsequent operations more efficient. The subsequent loop segments $S_a^{(2)}(q_2, h_a^{(2)})$ with $a > 1$ along \mathcal{C}_2 are defined in the same way as in the one-loop case and have the form (3.9) and (3.10). In this way, the chain of segments $S_a^{(2)}$ corresponds to

$$\prod_{a=1}^{N_2} S_a^{(2)}(q_2, h_a^{(2)}) = \int d\bar{q}_1 \frac{\text{Tr} \left[\mathcal{N}^{(1)}(q_1, h^{(1)}) \right]^{\alpha_1}}{\mathcal{D}^{(1)}(\bar{q}_1)} P_{\alpha_1 \alpha_2}(h^{(B)}) \left[\mathcal{N}^{(2)}(q_2, h^{(2)}) \right]^{\alpha_2}. \tag{4.16}$$

The dressing steps for this chain are carried out at the level of the Born-loop interference using the one-loop algorithm of section 3.3. To this end, using as initial condition the colour-Born interference

$$\mathcal{U}_0^{(2)}(h) = 2 \left(\sum_{\text{col}} \mathcal{M}_0^*(h) C_{2,\Gamma} \right), \quad (4.17)$$

the various segments are recursively combined via

$$\mathcal{U}_n^{(2)}(q_2, \check{h}_n^{(2)}) = \sum_{h_n^{(2)}} \mathcal{U}_{n-1}^{(2)}(q_2, \check{h}_{n-1}^{(2)}) \cdot S_n^{(2)}(q_2, h_n^{(2)}), \quad (4.18)$$

for $n = 1, \dots, N_2$. In this way helicity states are efficiently summed on-the-fly, and

$$\check{h}_n^{(2)} = h - \sum_{a=1}^n h_a^{(2)} = \sum_{a=n+1}^{N_2} h_a^{(2)} \quad (4.19)$$

corresponds to the helicity configuration of the segments that are still undressed. Note that $\check{h}_0^{(2)} = h$. The dressing recursion is again implemented in terms of tensor coefficients, and after step $n = N_2$ one arrives at the two-loop scattering probability density along the same lines as in (3.26)–(3.27), i.e. by taking the trace

$$\mathcal{U}_{\mu_1 \dots \mu_r}^{(2)} = \text{Tr} \left[\mathcal{U}_{N_2; \mu_1 \dots \mu_r}^{(2)}(\check{h}_{N_2}^{(2)}) \right], \quad (4.20)$$

where $\check{h}_{N_2}^{(2)} = 0$, and combining the tensor coefficients with the associated tensor integrals. This leads to

$$\mathcal{W}_{02,\Gamma} = \frac{1}{N_{\text{hcs}}} \text{Re} \left[\sum_{r=0}^{N_2} \mathcal{U}_{\mu_1 \dots \mu_r}^{(2)} \int d\bar{q}_2 \frac{q_2^{\mu_1} \dots q_2^{\mu_r}}{\mathcal{D}^{(2)}(\bar{q}_2)} \right], \quad (4.21)$$

which corresponds to the two-loop scattering probability density as defined in (2.26).

4.2 Reducible two-loop integrands of class Red1

The above algorithm can be easily extended to reducible diagrams of type Red1. In this case the one-loop chains $\mathcal{C}_1, \mathcal{C}_2$ are connected by a single four-gluon vertex \mathcal{V}_4 (see figure 1). The various colour-stripped amplitudes that result from the splitting (2.21) of \mathcal{V}_4 and any other quartic vertices are handled as independent diagrams. For each of them the first one-loop subdiagram is constructed as in step 1 of section 4.1. The only exception is that the open index α_1 , which connects the chain \mathcal{C}_1 to the bridge, is replaced by two indices $\beta_0^{(2)}, \beta_1^{(2)}$. The latter correspond to the Lorentz indices that result from cut-opening the two propagators of the chain \mathcal{C}_2 that are connected to \mathcal{V}_4 . In practice, the quartic vertex \mathcal{V}_4 is included in the last segment of the chain \mathcal{C}_1 , and the first loop (4.8) assumes the form

$$\left[P_{-1}(h^{(1)}) \right]^{\alpha_1} \rightarrow \left[P_{-1}(h^{(1)}) \right]_{\beta_0^{(2)}}^{\beta_1^{(2)}}. \quad (4.22)$$

Since topologies of type Red1 feature a trivial bridge (no bridge segments, $N_B = h^{(B)} = 0$), step 2 can be by-passed. In step 3 the starting point, i.e. the first loop segment of chain \mathcal{C}_2 , is simply given by

$$\left[S_1^{(2)}(q_2, h_1^{(2)}) \right]_{\beta_0^{(2)}}^{\beta_1^{(2)}} = -i \left[P_{-1}(h^{(1)}) \right]_{\beta_0^{(2)}}^{\beta_1^{(2)}}, \quad (4.23)$$

where $h_1^{(2)} = h^{(1)}$, and the factor $-i$ originates from the numerator $-ig^{\beta\beta'}$ of the gluon propagator associated with $D_1^{(2)}$. The rest of step 3 can be implemented as in section 4.1.

5 Irreducible two-loop integrands

The colour-stripped amplitude $\mathcal{A}_{2,\Gamma}(h)$ of an irreducible two-loop diagram Γ has the form

$$\mathcal{A}_{2,\Gamma}(h) = \int d\bar{q}_1 \int d\bar{q}_2 \frac{\mathcal{N}(q_1, q_2, h)}{\mathcal{D}^{(1)}(\bar{q}_1) \mathcal{D}^{(2)}(\bar{q}_2) \mathcal{D}^{(3)}(\bar{q}_3)} \Big|_{\bar{q}_3 = -(\bar{q}_1 + \bar{q}_2)}, \quad (5.1)$$

with the denominator chains

$$\mathcal{D}^{(i)}(\bar{q}_i) = D_0^{(i)}(\bar{q}_i) \cdots D_{N_i-1}^{(i)}(\bar{q}_i) \quad \text{with} \quad D_a^{(i)}(\bar{q}_i) = (\bar{q}_i + p_{ia})^2 - m_{ia}^2, \quad (5.2)$$

where p_{ia} and m_{ia} are the external momenta and internal masses along the chain \mathcal{C}_i . The three chains are connected by two vertices $\mathcal{V}_0, \mathcal{V}_1$, and each connecting vertex \mathcal{V}_a can either be a three-point vertex, as depicted in (5.1), or a quartic vertex. In the latter case \mathcal{V}_a is attached to an external subtree $w_a^{(V)}$, which is not shown in (5.1). Along each chain \mathcal{C}_i flows a single loop momentum q_i in the direction from \mathcal{V}_0 to \mathcal{V}_1 , with the boundary condition $q_3 = -(q_1 + q_2)$. If \mathcal{V}_0 is a triple vertex the external momenta associated with the propagators $D_i^{(0)}$ connected to \mathcal{V}_0 are set to $p_{i0} = 0$ for $i = 1, 2, 3$. Instead, if \mathcal{V}_0 is a quartic vertex we choose $p_{10} = p_{20} = 0$ and $p_{30} = k_{V0}$, where k_{V0} is the external momentum entering through the external subtree $w_0^{(V)}$ connected to \mathcal{V}_0 .

The numerator of (5.1) factorizes into three numerator chains and the two vertices $\mathcal{V}_0, \mathcal{V}_1$, each connecting all three chains,

$$\begin{aligned} \mathcal{N}(q_1, q_2, h) = & \prod_{i=1}^3 \left[\mathcal{N}^{(i)}(q_i, h^{(i)}) \right]_{\beta_0^{(i)}}^{\beta_{N_i}^{(i)}} \left[\mathcal{V}_0(q_1, q_2, q_3, h_0^{(V)}) \right]_{\beta_0^{(1)} \beta_0^{(2)} \beta_0^{(3)}}^{\beta_0^{(1)} \beta_0^{(2)} \beta_0^{(3)}} \\ & \times \left[\mathcal{V}_1(q_1, q_2, q_3, h_1^{(V)}) \right]_{\beta_{N_1}^{(1)} \beta_{N_2}^{(2)} \beta_{N_3}^{(3)}}^{\beta_{N_1}^{(1)} \beta_{N_2}^{(2)} \beta_{N_3}^{(3)}} \Big|_{q_3 = -(q_1 + q_2)}, \end{aligned} \quad (5.3)$$

where the tensors $\mathcal{V}_a(q_1, q_2, q_3, h_a^{(V)})$ represent the connecting vertices \mathcal{V}_a . In case of quartic vertices they also embody the related external subtree $w_a^{(V)}(h_a^{(V)})$, thereby inheriting the

dependence on the helicity states $h_a^{(V)}$. For triple vertices \mathcal{V}_a there is no helicity dependence, and $h_a^{(V)} = 0$.

The global helicity configuration defined by all external particles of the two-loop diagram is hence decomposed into

$$h = \sum_{i=1}^3 h^{(i)} + h_1^{(V)} + h_0^{(V)}, \quad (5.4)$$

where $h^{(i)}$ denotes the helicity configuration of the chain \mathcal{C}_i and is simply given by the sum of the corresponding segment helicities, as defined in (3.12),

$$h^{(i)} = \sum_{a=1}^{N_i-1} h_a^{(i)}. \quad (5.5)$$

Each chain numerator in (5.3) factorises into loop segments

$$\mathcal{N}^{(i)}(q_i, h^{(i)}) = S_0^{(i)}(q_i) S_1^{(i)}(q_i, h_1^{(i)}) \cdots S_{N_i-1}^{(i)}(q_i, h_{N_i-1}^{(i)}). \quad (5.6)$$

The segment $S_0^{(i)}(q_i, h_0^{(i)}) \equiv S_0^{(i)}(q_i)$ consists only of the corresponding propagator numerator, while all other $S_a^{(i)}(q_i, h_a^{(i)})$ with $a \geq 1$ are standard loop segments of the form (3.9) or (3.10), i.e. they consist of a triple or quartic loop vertex connected to one or two external subtrees and to the numerator of a loop propagator adjacent to the loop vertex in the dressing direction. In loop segments with a quartic vertex, the external subtrees $w_a^{(i)}$ depicted in (5.1) should be understood as pairs of subtrees.

The multiplications in (5.6) should be understood as matrix multiplications,

$$\left[\mathcal{N}^{(i)}(q_i, h^{(i)}) \right]_{\beta_0^{(i)}}^{\beta_{N_i}^{(i)}} = \left[S_0^{(i)}(q_i) \right]_{\beta_0^{(i)}}^{\beta_1^{(i)}} \left[S_1^{(i)}(q_i, h_1^{(i)}) \right]_{\beta_1^{(i)}}^{\beta_2^{(i)}} \cdots \left[S_{N_i-1}^{(i)}(q_i, h_{N_i-1}^{(i)}) \right]_{\beta_{N_i-1}^{(i)}}^{\beta_{N_i}^{(i)}}. \quad (5.7)$$

For simplicity we will suppress the Lorentz/spinor indices $\beta_a^{(i)}$ wherever possible.

Ultimately, we are interested in the construction of the scattering probability density (2.8), and in the following we focus on the contribution of a single unrenormalised two-loop diagram as defined in (2.26). To account for the effect of the colour-Born interference, similarly as in the one-loop relations (3.20) and (3.21) we introduce

$$\mathcal{U}_0(h) = 2 \sum_{\text{col}} \mathcal{M}_0^*(h) C_{2,\Gamma}, \quad (5.8)$$

and, as discussed in the following subsections, we combine it with the two-loop numerator in order to obtain a single helicity-summed object

$$\mathcal{U}(q_1, q_2) = \sum_h \mathcal{U}_0(h) \mathcal{N}(q_1, q_2, h). \quad (5.9)$$

We note that possible symmetry factors can be included in the two-loop numerator \mathcal{N} or, alternatively, in the associated colour factor $C_{2,\Gamma}$.

5.1 Generic structure of a recursive two-loop algorithm

Thanks to its factorised structure, the colour and helicity summed two-loop numerator (5.9) can be constructed through a numerical recursion of the form

$$\hat{\mathcal{U}}_n(q_1, q_2) = \hat{\mathcal{U}}_{n-1}(q_1, q_2) \cdot \mathcal{K}_n(q_1, q_2) \tag{5.10}$$

which starts with $\hat{\mathcal{U}}_0 = \mathbb{1}$ and terminates, after a certain number of steps N_r , with $\hat{\mathcal{U}}_{N_r}(q_1, q_2) = \mathcal{U}(q_1, q_2)$. We refer to the operations (5.10) as dressing steps. The building blocks \mathcal{K}_n that are attached in each step are either the colour-Born interference \mathcal{U}_0 , a connecting vertex, individual loop segments, or full chain numerators, i.e.

$$\mathcal{K}_n \in \left\{ \mathcal{U}_0, \mathcal{V}_a, S_k^{(i)}, \mathcal{N}^{(i)} \right\}. \tag{5.11}$$

The chain numerators $\mathcal{N}^{(i)}$ can in turn be constructed through a similar recursion from their loop segments. To this end we define the partially dressed chains

$$\mathcal{N}_n^{(i)}(q_i, \hat{h}_n^{(i)}) = S_0^{(i)}(q_i) \cdots S_n^{(i)}(q_i, h_n^{(i)}) \quad \text{and} \quad \mathcal{N}_{-1}^{(i)} = \mathbb{1}, \tag{5.12}$$

with helicities

$$\hat{h}_n^{(i)} = \sum_{a=1}^n h_a^{(i)}, \tag{5.13}$$

where $\hat{h}_0^{(i)} = 0$ and $\hat{h}_{N_i-1}^{(i)} = h^{(i)}$.

In order to control the polynomial dependence on the loop momenta in an efficient way, similarly as in the one-loop case, the dressing steps (5.10) and all relevant building blocks are implemented at the level of tensor coefficients.

For instance, for the partially dressed chains (5.12) we use the tensorial representation

$$\mathcal{N}_n^{(i)}(\hat{h}_n^{(i)}) = \sum_{r_i=0}^n \mathcal{N}_{n;\mu_1 \dots \mu_{r_i}}^{(i)}(\hat{h}_n^{(i)}) q_i^{\mu_1} \cdots q_i^{\mu_{r_i}}, \tag{5.14}$$

and the full chains $\mathcal{N}^{(i)}(h^{(i)}) = \mathcal{N}_{N_i}^{(i)}(h^{(i)})$ are handled in a similar way. The individual loop segments in (3.13) have the same form as in (3.9) and (3.10). Three-point vertices \mathcal{V}_0 and \mathcal{V}_1 can be written as

$$\left[\mathcal{V}_a(q_1, q_2, q_3, h_a^{(V)}) \right]_{\beta_{a_1}^{(1)} \beta_{a_2}^{(2)} \beta_{a_3}^{(3)}} = \left[\hat{Y}_a \right]_{\beta_{a_1}^{(1)} \beta_{a_2}^{(2)} \beta_{a_3}^{(3)}} + \sum_{i=1}^3 \left[\hat{Z}_{ia,\nu} \right]_{\beta_{a_1}^{(1)} \beta_{a_2}^{(2)} \beta_{a_3}^{(3)}} q_i^\nu, \tag{5.15}$$

where $a_i = 0$ for $a = 0$, $a_i = N_i - 1$ for $a = 1$, and $h_a^{(V)} = 0$. As observed above, triple vertices are helicity-independent, while the coefficients \hat{Y}_a and $\hat{Z}_{ia,\nu}$ encode the linear dependence on the loop momenta. Quartic vertices are independent of the loop momenta and have the form

$$\left[\mathcal{V}_a(q_1, q_2, q_3, h_a^{(V)}) \right]_{\beta_{a_1}^{(1)} \beta_{a_2}^{(2)} \beta_{a_3}^{(3)}} = \left[\hat{Y}_a^\sigma \right]_{\beta_{a_1}^{(1)} \beta_{a_2}^{(2)} \beta_{a_3}^{(3)}} w_{a\sigma}^{(V)}(h_a^{(V)}). \tag{5.16}$$

The dependence of the two-loop numerator (5.9) on the two independent loop momenta q_1, q_2 is encoded in the tensorial representation

$$\mathcal{U}(q_1, q_2) = \sum_{r=0}^{R_1} \sum_{s=0}^{R_2} \mathcal{U}_{\mu_1 \dots \mu_r, \nu_1 \dots \nu_s} q_1^{\mu_1} \dots q_1^{\mu_r} q_2^{\nu_1} \dots q_2^{\nu_s}, \quad (5.17)$$

and its contribution to the two-loop scattering probability density is

$$\mathcal{W}_{02,\Gamma} = \frac{1}{N_{\text{hcs}}} \sum_{r=0}^{R_1} \sum_{s=0}^{R_2} \text{Re} \left[\mathcal{U}_{\mu_1 \dots \mu_r, \nu_1 \dots \nu_s} I^{\mu_1 \dots \mu_r, \nu_1 \dots \nu_s} \right], \quad (5.18)$$

with the two-loop tensor integrals

$$I^{\mu_1 \dots \mu_r, \nu_1 \dots \nu_s} = \int d\bar{q}_1 \int d\bar{q}_2 \frac{q_1^{\mu_1} \dots q_1^{\mu_r} q_2^{\nu_1} \dots q_2^{\nu_s}}{\mathcal{D}^{(1)}(\bar{q}_1) \mathcal{D}^{(2)}(\bar{q}_2) \mathcal{D}^{(3)}(-\bar{q}_1 - \bar{q}_2)}. \quad (5.19)$$

The CPU efficiency and the memory footprint of the dressing recursion (5.10) depend in a critical way on the tensorial and helicity structure of the various building blocks, which depend in turn on the order in which they are attached to each other. In general, the most relevant aspects for the efficiency of the algorithm are the number of independent tensor coefficients, and the structure of the Feynman rules involved in the chain segments $S_k^{(i)}$ and in the connecting vertices \mathcal{V}_a .

At step n of the construction, the structure of the first term on the r.h.s. of (5.10) is

$$\hat{\mathcal{U}}_{n-1}(q_1, q_2, h_{n-1}) = \sum_{r=0}^{R_1} \sum_{s=0}^{R_2} \left[\hat{\mathcal{U}}_{n-1}(h_{n-1}) \right]_{\mu_1 \dots \mu_r, \nu_1 \dots \nu_s}^{\beta_1 \dots \beta_{N_l}} q_1^{\mu_1} \dots q_1^{\mu_r} q_2^{\nu_1} \dots q_2^{\nu_s}. \quad (5.20)$$

Here the number of independent tensor structures $q_1^{\mu_1} \dots q_2^{\nu_s}$ grows exponentially with the rank in q_1 and q_2 as illustrated in table 1, and the tensorial rank increases during the construction. In renormalisable theories, attaching a loop segment $S_k^{(i)}$ or a connecting vertex \mathcal{V}_a augments the rank in a loop momentum q_i by either 0 or 1. In addition, the number of tensor coefficients in (5.20) increases by a factor 4 for each one of the open Lorentz/spinor indices $\beta_1, \dots, \beta_{N_l}$. Such indices are associated with the various loop chains, and their number is typically $N_l = 0, 2$ or 3 in the phases of the construction that have a significant CPU cost. Finally, the number of tensor coefficients is also proportional to the number of independent helicity configurations h_{n-1} in (5.20). Analogous considerations hold for the term \mathcal{K}_n on the r.h.s. of (5.10).

In general, we observe that the construction of a single chain has the same complexity as for a one-loop chain, due to the dependence on a single loop momentum and two open indices, while the steps involving the two-loop vertices \mathcal{V}_a are much more expensive, since they involve three indices and they typically depend on two loop momenta.

Based on these qualitative considerations we have selected from all possible algorithms of the form (5.10) the most promising candidates. The algorithm presented in the following section has been identified as the most efficient candidate based on a CPU cost simulation for a wide range of QED and QCD Feynman diagrams. In this simulation the CPU cost

$R_1 \backslash R_2$	0	1	2	3
0	1	5	15	35
1	5	25	75	175
2	15	75	225	525
3	35	175	525	1225
4	70	350	1050	2450
5	126	630	1890	4410

Table 1. Number of independent tensor structures $q_1^{\mu_1} \dots q_1^{\mu_r} q_2^{\nu_1} \dots q_2^{\nu_s}$ symmetrised in the q_1 and q_2 indices up to maximal ranks $R_{1,2}$ in $q_{1,2}$, i.e. for $r \leq R_1$ and $s \leq R_2$.

was approximated by the number of numerical multiplications, which represent the most expensive operations.

It is worth mentioning that the algorithm presented in section 5.2 is roughly two orders of magnitude faster than a naive implementation, in which all three chains are first constructed independently of each other and are then contracted with the vertices \mathcal{V}_0 and \mathcal{V}_1 . The inefficiency of this naive approach is due to the fact that the expensive \mathcal{V}_0 and \mathcal{V}_1 contractions are carried out when the number of active helicities, Lorentz/spinor indices and tensor rank in the loop momenta are all maximal.

5.2 Integrand of a single irreducible two-loop diagram

In this section we present a dressing algorithm of type (5.10) for the construction of the integrand of a single irreducible two-loop diagram of the form (5.1) and its interference with the Born amplitude. The organisation of the full set of irreducible two-loop diagrams is discussed in section 5.3.

As a starting point, given an irreducible two-loop diagram, its chains and connecting vertices are ordered in a well-defined way that will correspond to the sequence of dressing steps.

Ordering rules. The ordering is implemented by assigning the labels $\mathcal{V}_0, \mathcal{V}_1$ and $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$ to specific vertices and chains as follows.

- I) The three chains are ordered by their number of segments in such a way that $N_1 \geq N_2 \geq N_3$. In case of equality, the number of helicity configurations along the chain is used as a criterion. In case of equality there, the minimal external particle index on the chain is used, as defined in (2.28).
- II) The role of \mathcal{V}_0 and \mathcal{V}_1 is assigned to the connecting vertices according to the following ordered set of criteria:
 - (i) A single three-gluon vertex or a ghost-gluon-vertex with a rank increment in q_1 or q_3 becomes \mathcal{V}_0 .⁸

⁸Since ghosts are not allowed as external particles, ghost-gluon vertices always appear in pairs as $\mathcal{V}_{0,1}$. Hence, this rule covers all cases involving ghost-gluon-vertices.

- (ii) A single two-loop four-gluon vertex becomes \mathcal{V}_1 .
- (iii) Use the external-particle indices, as defined in (2.28), and the position of these particles along the chains: particle 1 is required to be closer to \mathcal{V}_0 , in case of equality particle 2, etc.

These ordering rules fix the direction in which the various chains are constructed, namely from vertex \mathcal{V}_0 to \mathcal{V}_1 as depicted in (5.1). Moreover they determine the order and the way in which the three chains and the two connecting vertices are constructed and attached to each other. For this reason, the ordering rules affect the efficiency of the dressing algorithm in a significant way. The following algorithm has been optimised by testing a wide range of ordering rules, using for example different hierarchies of the criteria in I) and/or interchanging the roles of the three chains, e.g. making \mathcal{C}_3 the chain with the most helicity configurations. Furthermore the vertex rules have been optimised for each combination of QED and QCD vertices as $\mathcal{V}_{0,1}$.

The reasoning behind the above ordering rules will be explained together with the steps of the algorithm for which they are relevant. This algorithm consists of two main parts.

Part 1: construction of chain \mathcal{C}_3 . The shortest chain, i.e. \mathcal{C}_3 , is fully dressed through the recursion steps

$$\mathcal{N}_n^{(3)}(q_3, \hat{h}_n^{(3)}) = \mathcal{N}_{n-1}^{(3)}(q_3, \hat{h}_{n-1}^{(3)}) \cdot S_n^{(3)}(q_3, h_n^{(3)}) \quad \text{for } n = 0, \dots, N_3 - 1, \quad (5.21)$$

with the initial condition $\mathcal{N}_{-1}^{(3)} = \mathbb{1}$. The partially dressed chain numerator $\mathcal{N}_n^{(3)}$ depends on the helicity configuration of the already dressed segments,

$$\hat{h}_n^{(3)} = \hat{h}_{n-1}^{(3)} + h_n^{(3)}, \quad (5.22)$$

and the number of such helicity states grows in each recursion step (5.21) by a factor equal to the number of helicity configurations $h_n^{(3)}$ of the currently dressed segment $S_n^{(3)}$. Similarly as in (3.13)–(3.17), the above recursion is implemented in terms of tensor coefficients, and the result

$$\mathcal{N}^{(3)}(q_3, h^{(3)}) = \mathcal{N}_{N_3-1}^{(3)}(q_3, \hat{h}_{N_3-1}^{(3)}), \quad (5.23)$$

where $h^{(3)} = \hat{h}_{N_3-1}^{(3)}$, is used as a building block in part 2 of the algorithm.

Part 2: construction of the full diagram. This part of the algorithm deals with the construction of the colour- and helicity-summed interference of the two-loop integrand with the full Born amplitude. This is achieved through a dressing recursion of the form (5.10), where all dressing operations are implemented at the level of tensor coefficients as described in section 5.1.

The starting point of the recursion is the colour-Born interference term $\hat{U}_0 = \mathcal{U}_0$, as defined in (5.8), and the kernels of the subsequent dressing steps are

$$\mathcal{K}_n \in \left\{ S_0^{(1)}, \dots, S_{N_1-1}^{(1)}, \mathcal{V}_1, \mathcal{N}^{(3)}, \mathcal{V}_0, S_0^{(2)}, \dots, S_{N_2-1}^{(2)} \right\}, \quad (5.24)$$

i.e. the full integrand is built by starting from the chain \mathcal{C}_1 , which is subsequently connected to the vertex \mathcal{V}_1 and the chain \mathcal{C}_3 . Finally, the vertex \mathcal{V}_0 and the chain \mathcal{C}_2 are attached.

Step 1 — construction of chain \mathcal{C}_1 . The goal of this step is the construction of the numerator of the longest chain \mathcal{C}_1 interfered with the colour-Born factor \mathcal{U}_0 , defined in (5.8) combined with the colour-Born interference factor (5.8) and summed over the relevant helicities, i.e.

$$\mathcal{U}^{(1)}(q_1, \check{h}^{(1)}) = \sum_{h^{(1)}} \mathcal{U}_0(h) \mathcal{N}^{(1)}(q_1, h^{(1)}). \quad (5.25)$$

Here $h^{(1)}$ and h correspond, respectively, to the helicity configurations of the full chain \mathcal{C}_1 and of the whole process. Upon summation over $h^{(1)}$, the quantity on the l.h.s. depends only on

$$\check{h}^{(1)} = h - h^{(1)}, \quad (5.26)$$

which corresponds to the helicity state of the still undressed part of the two-loop diagram. The quantity (5.25) is constructed by connecting the various segments of \mathcal{C}_1 through a sequence of dressing steps

$$\mathcal{U}_n^{(1)}(q_1, \check{h}_n^{(1)}) = \sum_{h_n^{(1)}} \mathcal{U}_{n-1}^{(1)}(q_1, \check{h}_{n-1}^{(1)}) \cdot S_n^{(1)}(q_1, h_n^{(1)}) \quad (5.27)$$

for $n = 0, \dots, N_1 - 1$. In the n -th dressing step, the helicity d.o.f. $h_n^{(1)}$ of the n -th segment are summed on-the-fly with the technique of section 3.3. As a result, $\mathcal{U}_n^{(1)}$ depends only on the helicity state of the undressed part of the two-loop diagram,

$$\check{h}_n^{(1)} = h - \hat{h}_n^{(1)}, \quad (5.28)$$

where $\hat{h}_n^{(1)}$ is defined in (5.13) and corresponds to the dressed part of \mathcal{C}_1 . Each time that a new segment is attached, the number of remaining helicity configurations decreases as $\check{h}_n^{(1)} = \check{h}_{n-1}^{(1)} - h_n^{(1)}$. The initial condition for the recursion (5.27) is

$$\mathcal{U}_{-1}^{(1)}(q_1, \check{h}_{-1}^{(1)}) = \mathcal{U}_0(h), \quad (5.29)$$

with $\check{h}_{-1}^{(1)} = h$, and the last step results into

$$\mathcal{U}_{N_1-1}^{(1)}(q_1, \check{h}_{N_1-1}^{(1)}) = \mathcal{U}^{(1)}(q_1, \check{h}^{(1)}), \quad (5.30)$$

where $\check{h}_{N_1-1}^{(1)} = \check{h}^{(1)}$.

Since \mathcal{C}_1 is the longest chain, a large number of helicities is already summed during this part of the algorithm, and a large number of segments is constructed with dependence on a single loop momentum.

Step 2 — connecting the \mathcal{V}_1 vertex and chain \mathcal{C}_3 . In this twofold step the chain \mathcal{C}_1 interfered with \mathcal{U}_0 is connected to \mathcal{V}_1 and to the previously constructed numerator of \mathcal{C}_3 , and \mathcal{C}_1 interfered with \mathcal{U}_0 ,

$$\begin{aligned} \left[\mathcal{U}_1^{(13)}(q_1, q_3, h^{(2)} + h_0^{(V)}) \right]_{\beta_0^{(1)} \beta_0^{(3)} \beta_{N_2}^{(2)}} &= \sum_{h^{(3)}} \sum_{h_1^{(V)}} \left[\mathcal{U}^{(1)}(q_1, \check{h}^{(1)}) \right]_{\beta_0^{(1)}}^{\beta_{N_1}^{(1)}} \left[\mathcal{N}^{(3)}(q_3, h^{(3)}) \right]_{\beta_0^{(3)}}^{\beta_{N_3}^{(3)}} \\ &\times \left[\mathcal{V}_1(q_1, q_2, q_3, h_1^{(V)}) \right]_{\beta_{N_1}^{(1)} \beta_{N_2}^{(2)} \beta_{N_3}^{(3)}} \Big|_{q_2 = -(q_1 + q_3)}. \end{aligned} \quad (5.31)$$

The helicity d.o.f. of \mathcal{C}_3 and \mathcal{V}_1 are summed over on the fly, and the result $\mathcal{U}_1^{(13)}$ depends only on the helicity states $h^{(2)} + h_0^{(V)}$ of the undressed parts \mathcal{C}_2 and \mathcal{V}_0 . Note that $\mathcal{U}_1^{(13)}$ can involve a high number of coefficients due to three open indices $\beta_0^{(1)}\beta_0^{(3)}\beta_{N_2}^{(2)}$ and two independent loop momenta. Thus, for efficiency reasons, when implementing (5.31) it is useful to anticipate the dependence of \mathcal{V}_0 on the indices $\beta_0^{(1)}$ and $\beta_0^{(3)}$, and to ignore from the beginning all $\beta_0^{(1)}\beta_0^{(3)}$ combinations for which \mathcal{V}_0 vanishes.

The step (5.31) is implemented using a tensorial representation of the form (5.20) with q_1 and q_3 as independent loop momenta.

Step 3 — connecting the \mathcal{V}_0 vertex. In this step \mathcal{V}_0 is connected to the previously constructed object (5.31), which consists of the numerators of \mathcal{C}_1 , \mathcal{V}_1 and \mathcal{C}_3 interfered with \mathcal{U}_0 ,

$$\left[\mathcal{U}_{10}^{(13)}(q_1, q_2, h^{(2)})\right]_{\beta_{N_2}^{(2)}}^{\beta_0^{(2)}} = \sum_{h_0^{(V)}} \left\{ \left[\mathcal{U}_1^{(13)}(q_1, q_3, h^{(2)} + h_0^{(V)})\right]_{\beta_0^{(1)}\beta_0^{(3)}\beta_{N_2}^{(2)}} \times \left[\mathcal{V}_0(q_1, q_2, q_3, h_0^{(V)})\right]_{q_3=-(q_1+q_2)}^{\beta_0^{(1)}\beta_0^{(2)}\beta_0^{(3)}} \right\}. \quad (5.32)$$

Here the number of open indices is reduced to two. For quartic vertices \mathcal{V}_0 the associated helicity d.o.f. are summed on the fly, and the output depends only on the helicity $h^{(2)}$ of the still undressed chain \mathcal{C}_2 .

Also this step is implemented using a tensorial representation, but now we switch to the independent loop momenta q_1, q_2 by replacing $q_3 = -(q_1 + q_2)$. For a generic polynomial

$$A(q_1, q_3) = \sum_{r=0}^{r_1} \sum_{s=0}^{r_3} A_{\mu_1 \dots \mu_r, \nu_1 \dots \nu_s} q_1^{\mu_1} \dots q_1^{\mu_r} q_3^{\nu_1} \dots q_3^{\nu_s}, \quad (5.33)$$

with maximum ranks $r_{1,3}$ in $q_{1,3}$, the replacement $q_3 = -(q_1 + q_2)$ results into a polynomial

$$\tilde{A}(q_1, q_2) = A(q_1, -q_1 - q_2) = \sum_{r=0}^{R_1} \sum_{s=0}^{R_2(r)} \tilde{A}_{\mu_1 \dots \mu_r, \nu_1 \dots \nu_s} q_1^{\mu_1} \dots q_1^{\mu_r} q_2^{\nu_1} \dots q_2^{\nu_s}, \quad (5.34)$$

with maximum ranks $R_1 = r_1 + r_3$ and $R_2(r) = 0, \dots, r_3$ in q_1 and q_2 . In our implementation of (5.32), by default this replacement is applied only upon contraction of $\mathcal{U}_1^{(13)}$ with \mathcal{V}_0 and summation over $h_0^{(V)}$. In some cases it can be more efficient to apply $q_3 = -(q_1 + q_2)$ to $\mathcal{U}_1^{(13)}$ already at the end of step 2. In QCD this is the case if the vertex \mathcal{V}_1 in step 2 is a triple gluon vertex or a ghost vertex that does not increase the rank in q_3 .

For an efficient implementation of the entire algorithm it is important to note that — as explained in the following — the CPU cost of step 3 is much lower as compared to step 2. For example, let us consider step 3 for the case of a triple-gluon vertex,

$$\mathcal{V}_0^{\beta_0^{(1)}\beta_0^{(2)}\beta_0^{(3)}} = g \beta_0^{(1)}\beta_0^{(2)}(q_1 - q_2) \beta_0^{(3)} + \text{permutations}, \quad (5.35)$$

where the prefactor with the coupling constant has been suppressed and can be restored as an overall factor at the end of the construction. When this triple vertex is attached to the

tensor $\mathcal{U}_1^{(13)}$ via (5.32), the various metric tensors in (5.35) give either rise to a summation over the components of $\mathcal{U}_1^{(13)}$ with $\beta_0^{(1)} = \beta_0^{(3)}$, or to the identification of the open index $\beta_0^{(2)}$ with one of the summed indices $\beta_0^{(1)}, \beta_0^{(3)}$. As for the loop momenta q_i in (5.35), their effect is simply to raise the tensor rank in q_i . More explicitly, let us assume that $\mathcal{U}_1^{(13)}$ is parametrised in terms of q_1, q_3 as in (5.33). In practice, exploiting the symmetry of tensor integrals w.r.t. permutations of the indices associated with the individual loop momenta, such tensors can be implemented in the symmetrised form

$$A(q_1, q_3) = \sum_{n_0, \dots, m_3} A_{n_0, \dots, n_3; m_0, \dots, m_3} \prod_{\mu=0}^3 (q_1^\mu)^{n_\mu} \prod_{\nu=0}^3 (q_3^\nu)^{m_\nu}, \quad (5.36)$$

where the tensor-coefficient indices n_μ and m_ν correspond to the powers in q_1^μ and q_3^ν , respectively. In this symmetrised representation, multiplying $A(q_1, q_3)$ by a component of q_1 or q_3 amounts to a trivial reshuffling of the tensor coefficients. For instance,

$$B(q_1, q_3) = q_1^\beta A(q_1, q_3) \quad (5.37)$$

corresponds simply to

$$B_{n_0, \dots, n_\beta, \dots, m_3} = A_{n_0, \dots, n_{\beta-1}, \dots, m_3}. \quad (5.38)$$

In summary, in case of a triple-gluon vertex \mathcal{V}_0 step 3 does not involve any multiplication, but only sums and reshuffling operations, which are by far less time consuming.

Let us now consider step 2, where the vertex \mathcal{V}_1 is connected. This operation can involve a very large number of multiplications, which can be estimated as follows. The combination of the chains \mathcal{C}_1 and \mathcal{C}_3 , i.e. the tensors $\mathcal{U}^{(1)}$ and $\mathcal{N}^{(3)}$ in (5.31), requires $4^4 = 256$ multiplications due to the four open Lorentz/spinor indices $\beta_0^{(1)}\beta_{N_1}^{(1)}$ and $\beta_0^{(3)}\beta_{N_3}^{(3)}$. Taking into account the number of helicity configurations $N_{\text{hel}}^{(i)}$ and the number of tensor coefficients $N_{\text{tens}}^{(i)}$ of both chains, $\mathcal{C}_i = \mathcal{C}_1, \mathcal{C}_3$, yields $256 \times N_{\text{hel}}^{(1)} N_{\text{tens}}^{(1)} N_{\text{hel}}^{(3)} N_{\text{tens}}^{(3)}$ multiplications. The remaining operations depend on the form of the Feynman rules for the connecting vertex \mathcal{V}_1 in (5.31). Let us consider again a triple-gluon vertex,

$$\mathcal{V}_1^{\beta_{N_1}^{(1)}\beta_{N_2}^{(2)}\beta_{N_3}^{(3)}} = g^{\beta_{N_1}^{(1)}\beta_{N_2}^{(2)}}(k_1 - k_2)^{\beta_{N_3}^{(3)}} + \text{permutations}, \quad (5.39)$$

with $k_i = q_i + p_{i, N_i-1}$, where p_{i, N_i-1} is the external momentum of the last propagator along \mathcal{C}_i . Similarly as for step 3, the metric tensors in (5.39) do not give rise to any additional multiplication. The same holds for q_i terms, while the various p_{i, N_i-1} terms can give rise to $16 \times 256 \times N_{\text{hel}}^{(1)} N_{\text{tens}}^{(1)} N_{\text{hel}}^{(3)} N_{\text{tens}}^{(3)}$ additional multiplications. In total, thanks to additional optimisations, our implementation of step 2 for a triple-gluon vertex involves a number of multiplications slightly above $700 \times N_{\text{hel}}^{(1)} N_{\text{tens}}^{(1)} N_{\text{hel}}^{(3)} N_{\text{tens}}^{(3)}$. Given the potentially large number of helicity and tensorial components, this number can range from order 10^5 to 10^7 for nontrivial two-loop diagrams.

In general, step 2 is the most expensive operation of the entire algorithm, and its exact cost depends on the nature of the connecting vertex \mathcal{V}_1 . For this reason, the CPU efficiency of the entire algorithm depends in a critical way on the ordering rule II, which assigns the roles of \mathcal{V}_1 and \mathcal{V}_0 to the actual connecting vertices of two-loop diagrams. Regarding

point (i) of the ordering rule II, complex vertices are preferably chosen to be \mathcal{V}_0 . In QCD the triple-gluon vertex is by far the most CPU expensive one, and as a consequence it is only assigned the role of \mathcal{V}_1 if both connecting vertices $\mathcal{V}_{0,1}$ are of this type. The rules for ghost-gluon vertices are guided by which rank of a loop momentum q_i is increased, and hence at which point in the algorithm the replacement $q_3 = -(q_1 + q_2)$ is performed. Regarding point (ii), a quartic vertex is preferred to be inserted as \mathcal{V}_1 , since in this way the on-the-fly summation of the helicities $h_1^{(V)}$ of its external subtree can be performed at an earlier stage.

Finally we note that when both \mathcal{V}_0 and \mathcal{V}_1 are fermion-gauge-boson vertices, whose tensorial structure consists of γ -matrices containing only a few non-vanishing components, the most efficient implementation available in our framework is to combine \mathcal{V}_1 and \mathcal{V}_0 into a single tensor $\Gamma(\mathcal{V}_1, \mathcal{V}_0)$ and to merge steps 2 and 3 in a single operation,

$$\begin{aligned} \left[\mathcal{U}_{10}^{(13)}(q_1, q_2, h^{(2)}) \right]_{\beta_{N_2}^{(2)}}^{\beta_0^{(2)}} &= \sum_{h^{(3)}} \left\{ \left[\mathcal{U}^{(1)}(q_1, \check{h}^{(1)}) \right]_{\beta_0^{(1)}}^{\beta_{N_1}^{(1)}} \left[\mathcal{N}^{(3)}(q_3, h^{(3)}) \right]_{\beta_0^{(3)}}^{\beta_{N_3}^{(3)}} \right. \\ &\quad \left. \times \left[\Gamma(\mathcal{V}_1, \mathcal{V}_0) \right]_{\beta_{N_1}^{(1)} \beta_{N_2}^{(2)} \beta_{N_3}^{(3)}}^{\beta_0^{(1)} \beta_0^{(2)} \beta_0^{(3)}} \right\}_{q_3 = -(q_1 + q_2)}, \end{aligned} \quad (5.40)$$

where the index sums run only over the non-zero components. This is particularly efficient in pure QED calculations, where all vertices are proportional to γ^μ .

Step 4 — connecting \mathcal{C}_2 . The construction of the two-loop numerator is completed by attaching the chain \mathcal{C}_2 . The various segments of \mathcal{C}_2 are connected through a sequence of dressing steps

$$\mathcal{U}_n^{(123)}(q_1, q_2, \tilde{h}_n^{(2)}) = \sum_{h_n^{(2)}} \mathcal{U}_{n-1}^{(123)}(q_1, q_2, \tilde{h}_{n-1}^{(2)}) \cdot S_n^{(2)}(q_2, h_n^{(2)}) \quad (5.41)$$

for $n = 0, \dots, N_2 - 1$. At step n the helicity d.o.f. $h_n^{(2)}$ of the n -th segment are summed on the fly, and the result depends only on the helicities of the still undressed segments,

$$\tilde{h}_n^{(2)} = \sum_{a=n+1}^{N_2-1} h_a^{(2)} = \tilde{h}_{n-1}^{(2)} - h_n^{(2)}. \quad (5.42)$$

The initial condition for the recursion is given by the outcome of step 3, i.e.

$$\mathcal{U}_{-1}^{(123)}(q_1, q_2, \tilde{h}_{-1}^{(2)}) = \mathcal{U}_{10}^{(13)}(q_1, q_2, h^{(2)}), \quad (5.43)$$

where $\tilde{h}_{-1}^{(2)} = h^{(2)}$ corresponds to the helicity configuration of the full chain \mathcal{C}_2 .

The dressing steps (5.41) are similar to the ones for the construction of \mathcal{C}_1 , since the segments $S_n^{(2)}$ depend on a single loop momentum q_2 . However, the dressing of chain \mathcal{C}_2 features a higher complexity due to the dependence of $\mathcal{U}_n^{(123)}$ on two loop momenta, which leads to a larger number of tensor components.




Figure 3. Example of the interdependence of the tensorial ranks. The coloured area shows the ranks (r, s) in (q_1, q_2) that result from the $q_3 = -(q_1 + q_2)$ replacement applied to a polynomial with ranks $r_1, r_2, r_3 \leq 2$ in q_1, q_2, q_3 . In this case the maximum rank in q_2 obeys $R_2(r) \leq \min(4, 6 - r)$. Note that the excluded region with $r, s \leq 4$ and $s + r > 6$ corresponds to the highest-rank configurations, which involve a large number of tensor coefficients.

During the last step the loop is closed by contracting the index $\beta_{N_2}^{(2)}$ between vertex \mathcal{V}_1 and chain \mathcal{C}_2 ,

$$\begin{aligned} \mathcal{U}(q_1, q_2) &= \text{Tr} \left[\mathcal{U}_{N_2-1}^{(123)}(q_1, q_2, \tilde{h}_{N_2-1}^{(2)}) \right] \\ &= \sum_{h_{N_2-1}^{(2)}} \left[\mathcal{U}_{N_2-2}^{(123)}(q_1, q_2, \tilde{h}_{N_2-2}^{(2)}) \right]_{\beta_{N_2}^{(2)}}^{\beta_{N_2-1}^{(2)}} \cdot \left[S_{N_2-1}^{(2)}(q_2, h_{N_2-1}^{(2)}) \right]_{\beta_{N_2-1}^{(2)}}^{\beta_{N_2}^{(2)}}, \end{aligned} \quad (5.44)$$

where $\tilde{h}_{N_2-1}^{(2)} = 0$ and $\tilde{h}_{N_2-2}^{(2)} = h_{N_2-1}^{(2)}$.

Efficient implementation of steps 2–4. The steps 2 to 4 involve polynomials of the form (5.33) and (5.34), with a large number of tensor coefficients. For their efficient implementation, besides using symmetrised representations like (5.36), we systematically exploit the correlation of tensorial ranks as explained in the following. Let us consider polynomials in the two independent loop momenta q_1 and q_2 , which consist of various monomials of rank (r, s) in (q_1, q_2) . As a result of the particular form of the q_1, q_2 dependence of the involved Feynman rules, the maximum ranks $R_1 = \max(r)$ and $R_2 = \max(s)$ are not independent. For instance, attaching a triple-gluon vertex can increase the rank (r, s) of a monomial to $(r + 1, s)$ or $(r, s + 1)$ but not to $(r + 1, s + 1)$. In general, as a result of such correlations, for the subset of tensor coefficients with fixed rank r in q_1 , the maximum rank in q_2 depends on r , i.e. $R_2 = R_2(r)$, and vice versa. As discussed in (5.34), this kind of correlation arises also from the $q_3 = -(q_1 + q_2)$ replacement. An example of the corresponding $R_2(r)$ dependence is shown in figure 3.

In our implementation, for each single step of the algorithm the values of R_1 and $R_2(r)$ are determined based on the actual q_1, q_2 dependence of the involved Feynman rules, and all operations are restricted to the tensor components with $r \leq R_1$ and $s \leq R_2(r)$. In particular, we systematically avoid memory allocation for irrelevant tensor components. As a result the overall memory usage is typically reduced by more than a factor two as compared to an implementation with independent R_1 and R_2 .

5.3 Two-loop tensor coefficients for scattering amplitudes

For a set of Feynman diagrams, in particular for all irreducible two-loop diagrams of a full scattering amplitude, the algorithm is structured in such a way that the overall number of dressing steps is minimised. This is achieved by identifying all equivalent dressing steps that are required in multiple diagrams and performing such steps only once. To this end, all diagrams are ordered as discussed in section 5.2, and the steps described above for a single diagram are performed for the full set of diagrams in the following way.

Part 1. The chains \mathcal{C}_3 of all diagrams are grouped according to their number of segments N_3 , and the various groups are dressed starting from the lowest N_3 , i.e. the chains with only a propagator segment $S_0^{(3)}$. All fully or partially dressed chains that have already been constructed are stored and can be recycled whenever they appear as part of another chain. Since the various chains are ordered such that $N_3 \leq N_2 \leq N_1$, all chains \mathcal{C}_3 involve only a few segments,⁹ which require a rather small number of dressing equations.

Part 2. For the main algorithm, where the full two-loop amplitudes are constructed, we employ an extension of the on-the-fly diagram merging procedure, first introduced in [68] for one-loop amplitudes. This approach exploits the factorisation of the numerators of Feynman diagrams into segments. Consider for instance two diagrams with the same loop topology, i.e. the same scalar propagator denominators, which have the same undressed segments after a certain number n of recursion steps, e.g.

$$\begin{aligned} \mathcal{U}_A &= \mathcal{U}_{A,n}^{(1)} \cdot S_{n+1}^{(1)} \cdots S_{N_1-1}^{(1)} \cdot \mathcal{N}^{(3)} \cdot \mathcal{V}_1 \cdot \mathcal{V}_0 \cdot S_0^{(2)} \cdots S_{N_2-1}^{(2)}, \\ \mathcal{U}_B &= \mathcal{U}_{B,n}^{(1)} \cdot S_{n+1}^{(1)} \cdots S_{N_1-1}^{(1)} \cdot \mathcal{N}^{(3)} \cdot \mathcal{V}_1 \cdot \mathcal{V}_0 \cdot S_0^{(2)} \cdots S_{N_2-1}^{(2)}. \end{aligned} \quad (5.45)$$

Here $\mathcal{U}_{A,n}^{(1)}$ and $\mathcal{U}_{B,n}^{(1)}$ can differ due, for instance, to a different external subtree in the n -th segment of the chain \mathcal{C}_1 . Since ultimately the sum of all two-loop diagrams interfered with the full Born amplitude is needed, the partial numerators $\mathcal{U}_{A,n}^{(1)}$ and $\mathcal{U}_{B,n}^{(1)}$ can be summed, and all subsequent steps, here from segment $S_{n+1}^{(1)}$ to segment $S_{N_2-1}^{(2)}$, are performed only once, i.e.

$$\mathcal{U}_A + \mathcal{U}_B = \left(\mathcal{U}_{A,n}^{(1)} + \mathcal{U}_{B,n}^{(1)} \right) \cdot S_{n+1}^{(1)} \cdots S_{N_1-1}^{(1)} \cdot \mathcal{N}^{(3)} \cdot \mathcal{V}_1 \cdot \mathcal{V}_0 \cdot S_0^{(2)} \cdots S_{N_2-1}^{(2)}. \quad (5.46)$$

To this end, the full set of two-loop diagrams is constructed *segment by segment*, i.e. each of the steps 1 to 4 in section 5.2 is performed on all diagrams before continuing with the next step. For example, the first dressing step $\mathcal{U}_{-1}^{(1)} S_0^{(1)}$ is performed for all diagrams before all diagrams are dressed with segment $S_1^{(1)}$. After each dressing step along chain \mathcal{C}_1 and \mathcal{C}_2 the algorithm recognises and performs all possible merging operations before proceeding with the next segment. After the final dressing and merging steps all diagrams with the same topology, and hence corresponding to a single type of tensor integral, are summed into a single object.

The last criterion in ordering rule I, see section 5.2, namely the ordering of chains by external particle indices in case of $N_i = N_j$ and the same number of helicities, is designed in

⁹For a $2 \rightarrow 3$ process \mathcal{C}_3 does not have more than two segments.

such a way that on-the-fly merging opportunities during the construction of sets of diagrams are maximised. In general, merging steps can be performed for any number of diagrams, independently of their ranks. The overall CPU efficiency benefit of the on-the-fly merging is 10 to 20% for most processes.

Alternatively, to the *segment-by-segment* approach, the full set of recursion steps 1 to 4 can be performed *diagram by diagram*, without merging. This makes it possible to free the memory of previously computed partially dressed diagrams, thereby minimising memory usage. Since memory is usually not the bottleneck of the calculation, we chose the segment-by-segment approach as default.

6 Technical performance of the two-loop algorithm

The two-loop algorithms presented in sections 4–5 have been implemented in the OPENLOOPS framework in a fully automated way. The present implementation supports QED and QCD corrections to any SM process, and the full generality of our method allows for extensions to any other model.

This section is devoted to the validation and the technical performance of our implementation. In particular, we discuss the numerical stability, CPU efficiency and memory footprint for a variety of $2 \rightarrow 2$ and $2 \rightarrow 3$ processes at two loops in QED and QCD. Since the construction of reducible two-loop diagrams is largely based on well-established one-loop techniques, in the following we restrict ourselves to the contributions stemming from irreducible two-loop diagrams, which are constructed with the algorithm of section 5.

6.1 Validation with pseudotree test

For the validation of the two-loop algorithm we have implemented a so-called *pseudotree test*, which was employed to validate every new numerical routine, e.g. all possible two-loop vertices $\mathcal{V}_{0,1}$, in single diagrams, as well as the full algorithm for a wide range of processes.

The idea of the pseudotree test is to cut-open L -loop Feynman diagrams by “cutting” L loop lines, such as to obtain tree diagrams that can be evaluated with well established tree-level algorithms. In the case of irreducible two-loop diagrams, as illustrated in figure 4, we cut the loop lines associated with $D_0^{(1)}$ and $D_0^{(2)}$ on the side of the vertex \mathcal{V}_0 . The resulting open Lorentz/spinor indices are contracted with four pseudo wave functions e_1, \dots, e_4 , which are filled with random numbers. The loop momenta q_1 and q_2 are also fixed at random values, and the incoming external momenta associated with e_1, e_2, e_3, e_4 are $-q_1, q_1, -q_2, q_2$.

In the pseudotree test the cut-open diagrams of figure 4 are constructed in two different ways, as detailed in the following.

First, the colour- and helicity-summed interference of the cut-open two-loop diagram with the Born amplitude is computed using the well-tested tree-level algorithm (t) in OPENLOOPS. This calculation is carried out at fixed external momenta q_1, q_2 , and is implemented in double precision (DP) and quadruple precision (QP), where the latter is intended as a benchmark. The Born two-loop interferences computed in this way are




Figure 4. Insertion of pseudo wave functions into an irreducible two-loop diagram.

denoted $\widehat{\mathcal{W}}_{02}^{(t,DP)}$ and $\widehat{\mathcal{W}}_{02}^{(t,QP)}$, and

$$\widehat{\mathcal{W}}_{02} = \sum_{\Gamma} \widehat{\mathcal{W}}_{02,\Gamma}, \tag{6.1}$$

where the sum runs over the full set of irreducible two-loop diagrams Γ of the process at hand, and $\widehat{\mathcal{W}}_{02}$ denotes the loop integrand of the $\mathcal{W}_{02,\Gamma}$ defined in (2.26).

Second, we compute the pseudo trees with the two-loop algorithm (2L), constructing the original two-loop diagram with pseudo segments $e_1^{\gamma_1} e_2^{\gamma_2}$ and $e_3^{\gamma_3} e_4^{\gamma_4}$ inserted between \mathcal{V}_0 and the cut propagators $D_0^{(1)}$ and $D_0^{(2)}$. The resulting loop numerator $U(q_1, q_2)$, defined in (5.9), is obtained in the tensorial representation (5.17), which encodes the full q_1, q_2 dependence in the form of tensor coefficients. Including all denominators, we obtain

$$\widehat{\mathcal{W}}_{02,\Gamma}^{(2L)} = \frac{U(q_1, q_2)}{\mathcal{D}^{(1)}(\bar{q}_1) \mathcal{D}^{(2)}(\bar{q}_2) \mathcal{D}^{(3)}(\bar{q}_3)} \Big|_{q_3 = -(q_1 + q_2)}, \tag{6.2}$$

which is evaluated at fixed q_1, q_2 and compared against the result obtained with the tree algorithm. The Born two-loop interference computed in this way is denoted $\widehat{\mathcal{W}}_{02,\Gamma}^{(2L,DP)}$ and $\widehat{\mathcal{W}}_{02,\Gamma}^{(2L,QP)}$, for a single diagram in DP and QP respectively, and for a full process we compute $\widehat{\mathcal{W}}_{02}^{(2L,DP)}$ and $\widehat{\mathcal{W}}_{02}^{(2L,QP)}$ according to (6.1).

To validate our algorithm with the pseudotree test we have determined the relative numerical uncertainty

$$\mathcal{A}_p^{(t)} := \log_{10} \left(\frac{|\widehat{\mathcal{W}}_{02}^{(t,p)} - \widehat{\mathcal{W}}_{02}^{(2L,p)}|}{\text{Min}(|\widehat{\mathcal{W}}_{02}^{(t,p)}|, |\widehat{\mathcal{W}}_{02}^{(2L,p)}|)} \right), \tag{6.3}$$

where $p=DP, QP$. This test was applied to a wide range of single Feynman diagrams and full SM processes, considering two-loop QED or QCD corrections.

In particular, we have investigated the $2 \rightarrow 2$ and $2 \rightarrow 3$ processes listed in table 2, which involve all possible incarnations of the QED and QCD Feynman rules in the various steps of the algorithm of section 5. The double-precision comparison $\mathcal{A}_{DP}^{(t)}$ showed excellent agreement for all tested diagrams and processes with typical accuracy at the level of 10^{-15} .

corrections	process type	massless fermions	massive fermions	process
QED	2 → 2	e	–	$e^+e^- \rightarrow e^+e^-$
	2 → 3	e	–	$e^+e^- \rightarrow e^+e^-\gamma$
QCD	2 → 2	u	–	$gg \rightarrow u\bar{u}$
		u, d	–	$d\bar{d} \rightarrow u\bar{u}$
		u	–	$gg \rightarrow gg$
		u	t	$u\bar{u} \rightarrow t\bar{t}g$
		u	t	$gg \rightarrow t\bar{t}$
		u	t	$gg \rightarrow t\bar{t}g$
	2 → 3	u, d	–	$d\bar{d} \rightarrow u\bar{u}g$
		u	–	$gg \rightarrow ggg$
		u, d	–	$u\bar{d} \rightarrow W^+gg$
		u, d	–	$u\bar{u} \rightarrow W^+W^-g$
		u	t	$u\bar{u} \rightarrow t\bar{t}H$
		u	t	$gg \rightarrow t\bar{t}H$

Table 2. Processes considered for the validation of the two-loop algorithm and for the assessment of its CPU efficiency and memory footprint (see sections 6.3–6.4). The massive and massless states that enter closed fermion loops are listed in the third and fourth columns. Light quarks are implemented as one or two first-generation fields. In processes with a single internal quark field u , light-fermion loops are multiplied with the number of light quarks N_l in order to restore all active massless quark flavours. If the doublet (u, d) is used, light-fermion loops are multiplied with the number of light generations $N_l/2$.

6.2 Numerical stability

In order to demonstrate the numerical stability of the two-loop algorithm we have assessed the relative deviation between numerical evaluations of (6.2) in DP and QP,

$$\mathcal{A}_{\text{DP}} = \log_{10} \left(\frac{|\widehat{\mathcal{W}}_{02}^{(2\text{L},\text{DP})} - \widehat{\mathcal{W}}_{02}^{(2\text{L},\text{QP})}|}{\text{Min}(|\widehat{\mathcal{W}}_{02}^{(2\text{L},\text{DP})}|, |\widehat{\mathcal{W}}_{02}^{(2\text{L},\text{QP})}|)} \right). \quad (6.4)$$

To this end, we have first established the QP result $\widehat{\mathcal{W}}_{02}^{(2\text{L},\text{QP})}$ as a benchmark by confirming that the pseudotree test (6.3) in QP yields a typical accuracy at the level of 10^{-30} and always much better than 10^{-17} , which is the limit accuracy in DP.

The numerical accuracy (6.4) was tested for selected processes using samples of 10^5 uniformly distributed phase-space points. Results for a $2 \rightarrow 2$ and a $2 \rightarrow 3$ process are reported in figure 5. In both cases the bulk of the points has a relative uncertainty of 10^{-16} to 10^{-14} , and the upper bounds for the relative uncertainties are of order 10^{-12} and 10^{-11} for the $2 \rightarrow 2$ and the $2 \rightarrow 3$ process, respectively.

In conclusion, the two-loop algorithm for the construction of tensor coefficients is fully implemented and validated for QED and QCD corrections. Its numerical stability is




Figure 5. Numerical stability of the two-loop algorithm for irreducible diagrams in pseudotree mode. The plots show the fraction of phase-space points with a double-precision instability $\mathcal{A}_{DP} > \mathcal{A}_{min}$ as a function of \mathcal{A}_{min} . The distributions correspond to $gg \rightarrow t\bar{t}$ (left) and $d\bar{d} \rightarrow u\bar{u}g$ (right) samples with 10^5 events.

similarly good as for the tree-level and one-loop OPENLOOPS algorithm. This is an important prerequisite for a full two-loop calculation, where the dominant numerical instabilities are expected to arise from the reduction of tensor integrals, and should not be further enhanced by the calculation of tensor coefficients.

6.3 CPU efficiency

In this section we assess the CPU cost of the two-loop algorithm of section 5 and — in order to judge its impact in the context of realistic applications — we compare it to the cost of the related real-virtual contributions to a full NNLO calculation.

The CPU cost of the algorithm of section 5 is influenced by several aspects, such as the number of Feynman diagrams, the presence of massive propagators, the number of external helicity configurations (2 for fermions, 3 for W and 1 for H bosons) and the complexity of the vertices, in particular the connecting vertices $\mathcal{V}_{0,1}$, as well as the maximum tensor ranks. The overall impact of these different aspects is illustrated in figure 6, where we present the total CPU cost for a wide range of processes with two-loop QED and QCD corrections. Again we consider the $2 \rightarrow 2$ and $2 \rightarrow 3$ processes listed in table 2, which involve all relevant building blocks of our algorithm in different combinations. Colour and helicity sums are included throughout, and all two-loop timings in figure 6 correspond to the *segment-by-segment* approach.

In the upper frame of figure 6 we present the virtual-virtual timings per phase-space point, t_{VV} , for the construction of the two-loop tensor coefficients. Plotting these timings versus the number of irreducible two-loop diagrams N_{diags} we observe clusters of processes of similar number of diagrams, which also have similar timings. The simplest $2 \rightarrow 2$ processes, such as $e^+e^- \rightarrow e^+e^-$, take around 10 ms/point. Next there is a group of more complex $2 \rightarrow 2$ and simple $2 \rightarrow 3$ processes, which require about 100 ms/point, the fastest one being $u\bar{u} \rightarrow t\bar{t}H$ at 65 ms/point. The next group consists of more complex $2 \rightarrow 3$ processes, the fastest one being




Figure 6. Runtimes per phase-space point for the calculation of the two-loop tensor coefficients on a single core Intel(R) Core(TM) i7-6600U @ 2.6 GHz with 16GB RAM. The timings correspond to the QED and QCD processes in table 2 and are plotted versus the number N_{diags} of irreducible two-loop diagrams. The upper frame displays the absolute two-loop timings t_{VV} in ms. The middle frame shows the ratio $t_{\text{VV}}/t_{\text{RV}}^{\text{full}}$, where $t_{\text{RV}}^{\text{full}}$ is the time for computing the full one-loop probability density $\mathcal{W}_{01}^{(1)}$ (including tensor integrals) for a related process with one extra gluon or photon which enters the real-virtual part of the same NNLO calculation. The lower frame shows the ratio $t_{\text{VV}}/t_{\text{RV}}$, where t_{RV} is the cost of computing the one-loop tensor coefficients of $\mathcal{W}_{01}^{(1)}$ without tensor integrals. Colour and helicity sums are included throughout, and the latter are always carried out on-the-fly. All timings have been taken in DP using the GNU Fortran 7.5.0 compiler and averaging over 1000 points.

$d\bar{d} \rightarrow u\bar{u}g$ at 261 ms/point, as well as the most expensive $2 \rightarrow 2$ process, $gg \rightarrow gg$. The most time consuming processes we studied are $gg \rightarrow t\bar{t}H$, which takes 2.5 s/point, and $gg \rightarrow ggg$, which takes 9.2 s/point. The timings plotted in figure 6 are presented in appendix A.

In general, we observe that the two-loop runtimes grow only linearly with the number of diagrams, i.e. the ratio $t_{\text{VV}}/N_{\text{diags}}$ is approximately a process-independent constant. On the employed CPU we find

$$t_{\text{VV}} \approx N_{\text{diags}} \times 150 \mu\text{s}. \quad (6.5)$$

We also observe that adding an extra external gluon (photon) to a process in QCD (QED) increases the number of diagrams as well as the computation time t_{VV} by approximately one order of magnitude, e.g. in $e^+e^- \rightarrow e^+e^- (+\gamma)$ and $gg \rightarrow t\bar{t} (+g)$. These process-independent scaling features are comparable to the behaviour of one-loop calculations in OPENLOOPS.

In the lower frames of figure 6 we show the ratio of virtual-virtual and real-virtual timings for partonic subprocesses that contribute to the same NNLO calculation. More precisely, the virtual-virtual time t_{VV} , which corresponds to the two-loop tensor coefficients for a given $2 \rightarrow N$ process, is compared to the real-virtual time t_{RV} corresponding to the one-loop correction to the associated $2 \rightarrow N + 1$ process with one extra gluon or photon. Helicity d.o.f are always summed one-the-fly, both at one and two loops.

The real-virtual timings t_{RV}^{full} in the middle frame correspond to the cost of the full real-virtual density $\mathcal{W}_{01}^{(1)}$ in (2.11), i.e. in t_{RV}^{full} we include also the cost of the required one-loop tensor integrals.¹⁰ In this way we gain an idea of the share of the two-loop tensor coefficient construction in a full NNLO calculation. As shown in figure 6, in spite of the fact that t_{VV} and t_{RV}^{full} vary by more than three orders of magnitude, their ratio amounts to

$$\frac{t_{VV}}{t_{RV}^{\text{full}}} \approx 4 \pm 1, \tag{6.6}$$

and is approximately constant. In other words, independently of the type of process and its particle multiplicity, the cost of two-loop tensor coefficients per phase-space point is only a few times higher as compared to the related real-virtual amplitudes.

The lower frame of figure 6 shows the same ratio for real-virtual timings t_{RV} that do not include the contribution of tensor integrals. This provides a more direct comparison of the cost of one-loop and two-loop NNLO building blocks at the level of tensor coefficients. For this ratio we find

$$\frac{t_{VV}}{t_{RV}} \approx 9 \pm 3, \tag{6.7}$$

and again we observe an approximately process-independent behaviour. This is related to the fact that the number of two-loop diagrams for a given process usually lies between the number of one-loop diagrams for the corresponding process with one and two extra partons (gluon or photons). We also note that the total tensor rank in the loop momenta is the same at two-loop level and at one-loop level with two extra partons, and comparing the cost of the corresponding tensor coefficients we find that t_{VV} is a factor 3 to 8 less expensive w.r.t. the one-loop coefficients for processes with two extra partons (see appendix A). Considering the much higher complexity of two-loop diagrams as compared to one-loop diagrams, these are very promising findings.

Finally, let us compare the efficiency of the *diagram-by-diagram* and *segment-by-segment* approaches discussed in section 5.3. The former uses less memory and tends to be slower, while the latter requires more memory and tends to be faster thanks to the on-the-fly merging of diagrams. For simple processes the two approaches feature similar timings, their ratio being 1 ± 0.15 , and in most cases the *segment-by-segment* approach being faster. For more complex processes, however, the *segment-by-segment* approach outperforms the *diagram-by-diagram* approach by a factor 1.15 to 1.35, with the highest ratio of 1.35 for $gg \rightarrow ggg$. This behaviour is expected, since the efficiency gain due the on-the-fly merging is larger for more complex processes with a higher number of Feynman diagrams. As

¹⁰In this case the full integral reduction and evaluation is performed after the construction of the tensor coefficients with COLLIER [14].

hard process	virtual-virtual memory [MB]		real-virtual [MB]	
	segment-by-segment	diagram-by-diagram	coefficients	full
$e^+e^- \rightarrow e^+e^-$	18	8	6	23
$e^+e^- \rightarrow e^+e^-\gamma$	154	25	22	54
$gg \rightarrow u\bar{u}$	75	31	10	26
$gg \rightarrow t\bar{t}$	94	35	15	34
$gg \rightarrow t\bar{t}g$	2000	441	152	213
$u\bar{d} \rightarrow W^+gg$	563	143	54	90
$u\bar{u} \rightarrow W^+W^-g$	264	67	36	67
$u\bar{u} \rightarrow t\bar{t}H$	82	28	14	40
$gg \rightarrow t\bar{t}H$	604	145	50	90
$u\bar{u} \rightarrow t\bar{t}g$	323	83	41	74
$gg \rightarrow gg$	271	94	41	55
$d\bar{d} \rightarrow u\bar{u}$	18	10	9	20
$d\bar{d} \rightarrow u\bar{u}g$	288	85	39	68
$gg \rightarrow ggg$	6299	1597	623	683

Table 3. Memory required by the full set of tensor coefficients for the processes of table 2. The two central columns correspond to the construction of two-loop coefficients in the *segment-by-segment* and *diagram-by-diagram* approaches (see section 5.3). The last two columns report the memory footprint of the one-loop coefficients for the real-virtual contributions to the same NNLO calculation, as well as the memory footprint of the full real-virtual calculation, including the reduction and evaluation of the tensor integrals with COLLIER. For these the public OPENLOOPS 2 program with on-the-fly helicity summation was used.

the default option for our two-loop algorithm the *segment-by-segment* approach is the clear choice.

In conclusion, for the same number of phase-space points the CPU cost of two-loop tensor coefficients is of the same order as the cost of the related real-virtual parts of a NNLO calculation. However, the number of required points for the integration of the two-loop contributions is typically much lower. In the light of these observations, the presented two-loop timings are quite promising. In practice we expect that the cost of two-loop tensor coefficients should play a subleading role in the CPU budget of complete $2 \rightarrow 3$ NNLO calculations.

6.4 Memory usage

The construction of two-loop tensor coefficients with the algorithm of section 5 can require an important amount of memory. This is quantified in table 3 for the full set of processes listed in table 2. In addition to the memory footprint of the *segment-by-segment* and *diagram-by-diagram* approaches at two loops, for comparison we also report the corresponding memory

usage for real-virtual contributions, i.e. for the one-loop corrections to the related processes with one extra parton. At two loops only the memory for the storage of tensor coefficients is considered, while at one loop we present both the memory consumption for the construction of the tensor coefficients and for the full calculation.

We observe that the memory requirements for the two-loop tensor coefficients in the *diagram-by-diagram* and *segment-by-segment* approaches are, respectively, a factor 1–3 and 2–13 higher w.r.t. the real-virtual case. Including the one-loop tensor integrals, the *diagram-by-diagram* and *segment-by-segment* require, respectively, a factor 0.3–2.3 and 0.8–9.4 of the memory for the full real-virtual calculation. The *segment-by-segment* approach uses 2 to 6 times more memory than the *diagram-by-diagram* approach for the construction of the tensor coefficients. Thus the latter may be preferable when memory consumption plays a critical role. However, this is never the case for the considered processes, and both two-loop approaches are well within the scope of current computing environments.

7 Conclusion

We have presented a new and fully general algorithm for the efficient construction of two-loop integrands. To this end we have designed and implemented a generalisation of the open-loops method to two loops. In this approach, the polynomial dependence of two-loop integrand numerators on the loop momenta q_1, q_2 is encoded in a set of tensor coefficients that are constructed through a numerical recursion. The interference with the Born amplitude as well as helicity and colour sums are systematically included, and the resulting tensor coefficients are ready to be combined with the relevant two-loop tensor integrals to form scattering probability densities at two loops.

For reducible two-loop diagrams, which factorise into one-loop subdiagrams, we have presented an algorithm that exploits and extends various aspects of the pre-existing OPEN-LOOPS program at one loop, while for irreducible two-loop diagrams we have developed a largely new and more sophisticated algorithm, which represents the main novelty of this paper.

Irreducible two-loop diagrams involve three loop chains that depend, respectively, on the loop momenta $q_1, q_2, q_3 = -q_1 - q_2$, and are attached to each other through two connecting vertices. Each loop chain corresponds to a sequence of so-called loop segments, each consisting of a loop propagator attached to one or two external subtrees through a vertex. The numerator of an irreducible two-loop diagram is constructed through a numerical recursion where the various loop segments and connecting vertices are attached to each other one after the other. The single steps of the recursion correspond to process-independent operations that depend only on the Feynman rules of the model at hand, and the entire recursion is implemented in terms of tensor coefficients. Due to the large number of tensorial structures and helicity degrees of freedom, the most expensive individual steps can require of the order of 10^6 multiplications per two-loop diagram. Thanks to a systematic comparison of several algorithmic options and a detailed cost analysis we have identified an efficient algorithm, which turned out to outperform naive approaches by two orders of magnitude. This high efficiency is achieved through an optimal ordering of the construction

steps, the factorisation of colour structures, the on-the-fly summation of helicity degrees of freedom, and various other tricks.

The algorithm has been implemented in a fully automated way in the OPENLOOPS framework, and this first implementation supports two-loop QED and QCD corrections to any scattering process within the Standard Model. Its correctness and numerical stability have been verified by means of a so-called pseudo-tree test in combination with quadruple-precision benchmarks. In order to assess the efficiency of the new algorithm we have studied a set of $2 \rightarrow 2$ and $2 \rightarrow 3$ processes with numbers of two-loop diagrams ranging from order 10^2 to 10^5 per process. We have observed that the cost of the two-loop tensor coefficients per phase-space point for an entire process is approximately proportional to the number of (irreducible) two-loop diagrams. On a single Intel i7 core this cost is around $150 \mu\text{s}$ per two-loop diagram, and for typical $2 \rightarrow 3$ partonic processes at two loops in QCD it is of the order of one second. Finally we have shown that the cost of the two-loop tensor coefficients that enter the virtual-virtual part of a NNLO calculation is comparable to the one of the one-loop ingredients that enter the corresponding real-virtual part.

In conclusion, the presented algorithm provides a key building block for the construction of an automated generator of scattering amplitudes at two loops, and the observed technical performance guarantees its applicability to arbitrary $2 \rightarrow 2$ and $2 \rightarrow 3$ processes.

Acknowledgments

We would like to thank F. Buccioni for useful discussions. This research was supported by the Swiss National Science Foundation (SNSF) under the SNSF Ambizione grant PZ00P2-179877. The work of S.P. was supported through contract BSCGI0-157722.

A CPU efficiency measurements

In table 4 we report the explicit timings that have been discussed in section 6.3 and summarised in figure 6. The first column corresponds to the full list of hard scattering processes defined in table 2. The second and third columns indicate, respectively, the number of irreducible two-loop diagrams, N_{diags} , and the CPU time t_{VV} per phase space point for the construction of two-loop tensor coefficients, as defined in section 6.3. This time t_{VV} corresponds the default variant of the two-loop algorithm, i.e. the *segment-by-segment* approach (see section 5.3). The fourth column shows the ratio of t_{VV} to the corresponding time $t_{\text{VV}}^{\text{dbd}}$ required by the *diagram-by-diagram* approach. The next two columns give the ratios of t_{VV} to the timings for the one-loop tensor coefficients (t_{RV}) or full one-loop amplitude including tensor integrals ($t_{\text{RV}}^{\text{full}}$) for the corresponding process with one additional parton (gluons in QCD and photons in QED). The last two columns give the ratios of t_{VV} and the timings for the one-loop tensor coefficients (t_{RRV}) or full one-loop calculation ($t_{\text{RRV}}^{\text{full}}$) for the corresponding process with two extra partons. All time measurements have been carried out as detailed in the caption of figure 6.

process	N_{diags}	t_{VV} [ms]	$\frac{t_{\text{VV}}}{t_{\text{dbd}}^{\text{VV}}}$	$\frac{t_{\text{VV}}}{t_{\text{RV}}}$	$\frac{t_{\text{VV}}}{t_{\text{full}}^{\text{RV}}}$	$\frac{t_{\text{VV}}}{t_{\text{RRV}}}$	$\frac{t_{\text{VV}}}{t_{\text{full}}^{\text{RRV}}}$
$e^+e^- \rightarrow e^+e^-$	94	7.18	1.15	10.1	3.47	0.633	0.280
$e^+e^- \rightarrow e^+e^-\gamma$	848	114	1.10	10.0	4.44	0.578	0.267
$gg \rightarrow u\bar{u}$	835	66.3	0.994	8.67	5.35	0.309	0.200
$d\bar{d} \rightarrow u\bar{u}$	197	12.7	1.04	12.9	4.08	0.471	0.231
$d\bar{d} \rightarrow u\bar{u}g$	2690	261	0.940	9.71	4.76	0.339	0.203
$gg \rightarrow gg$	3633	315	0.845	5.86	5.01	0.185	0.129
$gg \rightarrow ggg$	63105	9250	0.738	5.26	3.68	—	—
$u\bar{u} \rightarrow t\bar{t}g$	2690	309	0.964	7.46	3.97	0.259	0.165
$gg \rightarrow t\bar{t}$	955	91.9	1.04	7.10	4.14	0.251	0.158
$gg \rightarrow t\bar{t}g$	15462	2541	0.957	6.93	4.37	—	—
$u\bar{d} \rightarrow W^+gg$	3450	455	0.956	8.96	3.25	0.334	0.123
$u\bar{u} \rightarrow W^+W^-g$	765	194	0.979	8.68	3.40	0.362	0.115
$u\bar{u} \rightarrow t\bar{t}H$	724	64.7	1.03	11.8	3.49	0.408	0.165
$gg \rightarrow t\bar{t}H$	4304	636	0.923	10.9	3.38	0.361	0.151

Table 4. Virtual-virtual and real-virtual timings for the $2 \rightarrow 2$ and $2 \rightarrow 3$ processes discussed in section 6.3.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] T. Gleisberg et al., *Event generation with SHERPA 1.1*, *JHEP* **02** (2009) 007 [[arXiv:0811.4622](https://arxiv.org/abs/0811.4622)] [[INSPIRE](#)].
- [2] S. Alioli, P. Nason, C. Oleari and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, *JHEP* **06** (2010) 043 [[arXiv:1002.2581](https://arxiv.org/abs/1002.2581)] [[INSPIRE](#)].
- [3] G. Bevilacqua et al., *HELAC-NLO*, *Comput. Phys. Commun.* **184** (2013) 986 [[arXiv:1110.1499](https://arxiv.org/abs/1110.1499)] [[INSPIRE](#)].
- [4] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079 [[arXiv:1405.0301](https://arxiv.org/abs/1405.0301)] [[INSPIRE](#)].
- [5] G. Ossola, C.G. Papadopoulos and R. Pittau, *CutTools: A Program implementing the OPP reduction method to compute one-loop amplitudes*, *JHEP* **03** (2008) 042 [[arXiv:0711.3596](https://arxiv.org/abs/0711.3596)] [[INSPIRE](#)].
- [6] C.F. Berger et al., *An Automated Implementation of On-Shell Methods for One-Loop Amplitudes*, *Phys. Rev. D* **78** (2008) 036003 [[arXiv:0803.4180](https://arxiv.org/abs/0803.4180)] [[INSPIRE](#)].

- [7] A. van Hameren, C.G. Papadopoulos and R. Pittau, *Automated one-loop calculations: A Proof of concept*, *JHEP* **09** (2009) 106 [[arXiv:0903.4665](#)] [[INSPIRE](#)].
- [8] A. van Hameren, *OneLOop: For the evaluation of one-loop scalar functions*, *Comput. Phys. Commun.* **182** (2011) 2427 [[arXiv:1007.4716](#)] [[INSPIRE](#)].
- [9] V. Hirschi, R. Frederix, S. Frixione, M.V. Garzelli, F. Maltoni and R. Pittau, *Automation of one-loop QCD corrections*, *JHEP* **05** (2011) 044 [[arXiv:1103.0621](#)] [[INSPIRE](#)].
- [10] F. Cascioli, P. Maierhofer and S. Pozzorini, *Scattering Amplitudes with Open Loops*, *Phys. Rev. Lett.* **108** (2012) 111601 [[arXiv:1111.5206](#)] [[INSPIRE](#)].
- [11] S. Badger, B. Biedermann, P. Uwer and V. Yundin, *Numerical evaluation of virtual corrections to multi-jet production in massless QCD*, *Comput. Phys. Commun.* **184** (2013) 1981 [[arXiv:1209.0100](#)] [[INSPIRE](#)].
- [12] G. Cullen 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 [[arXiv:1404.7096](#)] [[INSPIRE](#)].
- [13] T. Peraro, *Ninja: Automated Integrand Reduction via Laurent Expansion for One-Loop Amplitudes*, *Comput. Phys. Commun.* **185** (2014) 2771 [[arXiv:1403.1229](#)] [[INSPIRE](#)].
- [14] A. Denner, S. Dittmaier and L. Hofer, *Collier: a fortran-based Complex One-Loop LIBrary in Extended Regularizations*, *Comput. Phys. Commun.* **212** (2017) 220 [[arXiv:1604.06792](#)] [[INSPIRE](#)].
- [15] 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 [[arXiv:1605.01090](#)] [[INSPIRE](#)].
- [16] S. Carrazza, R.K. Ellis and G. Zanderighi, *QCDLoop: a comprehensive framework for one-loop scalar integrals*, *Comput. Phys. Commun.* **209** (2016) 134 [[arXiv:1605.03181](#)] [[INSPIRE](#)].
- [17] F. Buccioni et al., *OpenLoops 2*, *Eur. Phys. J. C* **79** (2019) 866 [[arXiv:1907.13071](#)] [[INSPIRE](#)].
- [18] R. Boughezal, F. Caola, K. Melnikov, F. Petriello and M. Schulze, *Higgs boson production in association with a jet at next-to-next-to-leading order*, *Phys. Rev. Lett.* **115** (2015) 082003 [[arXiv:1504.07922](#)] [[INSPIRE](#)].
- [19] G. Heinrich, S. Jahn, S.P. Jones, M. Kerner and J. Pires, *NNLO predictions for Z-boson pair production at the LHC*, *JHEP* **03** (2018) 142 [[arXiv:1710.06294](#)] [[INSPIRE](#)].
- [20] A. Gehrmann-De Ridder, T. Gehrmann, E.W.N. Glover, A. Huss and J. Pires, *Triple Differential Dijet Cross Section at the LHC*, *Phys. Rev. Lett.* **123** (2019) 102001 [[arXiv:1905.09047](#)] [[INSPIRE](#)].
- [21] D. de Florian, I. Fabre and J. Mazzitelli, *Triple Higgs production at hadron colliders at NNLO in QCD*, *JHEP* **03** (2020) 155 [[arXiv:1912.02760](#)] [[INSPIRE](#)].
- [22] M. Grazzini, S. Kallweit, J.M. Lindert, S. Pozzorini and M. Wiesemann, *NNLO QCD + NLO EW with Matrix+OpenLoops: precise predictions for vector-boson pair production*, *JHEP* **02** (2020) 087 [[arXiv:1912.00068](#)] [[INSPIRE](#)].
- [23] M. Czakon, A. Mitov, M. Pellen and R. Poncelet, *NNLO QCD predictions for W+c-jet production at the LHC*, *JHEP* **06** (2021) 100 [[arXiv:2011.01011](#)] [[INSPIRE](#)].
- [24] P. Banerjee, T. Engel, N. Schalch, A. Signer and Y. Ulrich, *Bhabha scattering at NNLO with next-to-soft stabilisation*, *Phys. Lett. B* **820** (2021) 136547 [[arXiv:2106.07469](#)] [[INSPIRE](#)].

- [25] P. Banerjee, T. Engel, N. Schalch, A. Signer and Y. Ulrich, *Møller scattering at NNLO*, *Phys. Rev. D* **105** (2022) L031904 [[arXiv:2107.12311](#)] [[INSPIRE](#)].
- [26] J.M. Campbell, G. De Laurentis, R.K. Ellis and S. Seth, *The $pp \rightarrow W(\rightarrow \ell\nu) + \gamma$ process at next-to-next-to-leading order*, *JHEP* **07** (2021) 079 [[arXiv:2105.00954](#)] [[INSPIRE](#)].
- [27] F. Buccioni et al., *Mixed QCD-electroweak corrections to dilepton production at the LHC in the high invariant mass region*, [arXiv:2203.11237](#) [[INSPIRE](#)].
- [28] G. Heinrich, *Collider Physics at the Precision Frontier*, *Phys. Rept.* **922** (2021) 1 [[arXiv:2009.00516](#)] [[INSPIRE](#)].
- [29] S. Badger, C. Brønnum-Hansen, H.B. Hartanto and T. Peraro, *Analytic helicity amplitudes for two-loop five-gluon scattering: the single-minus case*, *JHEP* **01** (2019) 186 [[arXiv:1811.11699](#)] [[INSPIRE](#)].
- [30] S. Abreu, J. Dormans, F. Febres Cordero, H. Ita and B. Page, *Analytic Form of Planar Two-Loop Five-Gluon Scattering Amplitudes in QCD*, *Phys. Rev. Lett.* **122** (2019) 082002 [[arXiv:1812.04586](#)] [[INSPIRE](#)].
- [31] S. Abreu, J. Dormans, F. Febres Cordero, H. Ita, B. Page and V. Sotnikov, *Analytic Form of the Planar Two-Loop Five-Parton Scattering Amplitudes in QCD*, *JHEP* **05** (2019) 084 [[arXiv:1904.00945](#)] [[INSPIRE](#)].
- [32] S. Badger et al., *Analytic form of the full two-loop five-gluon all-plus helicity amplitude*, *Phys. Rev. Lett.* **123** (2019) 071601 [[arXiv:1905.03733](#)] [[INSPIRE](#)].
- [33] H.B. Hartanto, S. Badger, C. Brønnum-Hansen and T. Peraro, *A numerical evaluation of planar two-loop helicity amplitudes for a W-boson plus four partons*, *JHEP* **09** (2019) 119 [[arXiv:1906.11862](#)] [[INSPIRE](#)].
- [34] H.A. Chawdhry, M. Czakon, A. Mitov and R. Poncelet, *Two-loop leading-color helicity amplitudes for three-photon production at the LHC*, *JHEP* **06** (2021) 150 [[arXiv:2012.13553](#)] [[INSPIRE](#)].
- [35] H.A. Chawdhry, M. Czakon, A. Mitov and R. Poncelet, *Two-loop leading-colour QCD helicity amplitudes for two-photon plus jet production at the LHC*, *JHEP* **07** (2021) 164 [[arXiv:2103.04319](#)] [[INSPIRE](#)].
- [36] B. Agarwal, F. Buccioni, A. von Manteuffel and L. Tancredi, *Two-loop leading colour QCD corrections to $q\bar{q} \rightarrow \gamma\gamma g$ and $qg \rightarrow \gamma\gamma q$* , *JHEP* **04** (2021) 201 [[arXiv:2102.01820](#)] [[INSPIRE](#)].
- [37] B. Agarwal, F. Buccioni, A. von Manteuffel and L. Tancredi, *Two-Loop Helicity Amplitudes for Diphoton Plus Jet Production in Full Color*, *Phys. Rev. Lett.* **127** (2021) 262001 [[arXiv:2105.04585](#)] [[INSPIRE](#)].
- [38] S. Abreu, F. Febres Cordero, H. Ita, B. Page and V. Sotnikov, *Leading-color two-loop QCD corrections for three-jet production at hadron colliders*, *JHEP* **07** (2021) 095 [[arXiv:2102.13609](#)] [[INSPIRE](#)].
- [39] S. Abreu, F. Febres Cordero, H. Ita, M. Klinkert, B. Page and V. Sotnikov, *Leading-color two-loop amplitudes for four partons and a W boson in QCD*, *JHEP* **04** (2022) 042 [[arXiv:2110.07541](#)] [[INSPIRE](#)].
- [40] S. Badger, H.B. Hartanto and S. Zoia, *Two-Loop QCD Corrections to Wbb^- Production at Hadron Colliders*, *Phys. Rev. Lett.* **127** (2021) 012001 [[arXiv:2102.02516](#)] [[INSPIRE](#)].

- [41] S. Badger, E. Chaubey, H.B. Hartanto and R. Marzucca, *Two-loop leading colour QCD helicity amplitudes for top quark pair production in the gluon fusion channel*, *JHEP* **06** (2021) 163 [[arXiv:2102.13450](#)] [[INSPIRE](#)].
- [42] S. Badger et al., *Virtual QCD corrections to gluon-initiated diphoton plus jet production at hadron colliders*, *JHEP* **11** (2021) 083 [[arXiv:2106.08664](#)] [[INSPIRE](#)].
- [43] S. Badger, H.B. Hartanto, J. Kryś and S. Zoia, *Two-loop leading-colour QCD helicity amplitudes for Higgs boson production in association with a bottom-quark pair at the LHC*, *JHEP* **11** (2021) 012 [[arXiv:2107.14733](#)] [[INSPIRE](#)].
- [44] S. Badger, H.B. Hartanto, J. Kryś and S. Zoia, *Two-loop leading colour helicity amplitudes for $W^{\pm}\gamma + j$ production at the LHC*, *JHEP* **05** (2022) 035 [[arXiv:2201.04075](#)] [[INSPIRE](#)].
- [45] S. Kallweit, V. Sotnikov and M. Wiesemann, *Triphoton production at hadron colliders in NNLO QCD*, *Phys. Lett. B* **812** (2021) 136013 [[arXiv:2010.04681](#)] [[INSPIRE](#)].
- [46] H.A. Chawdhry, M.L. Czakon, A. Mitov and R. Poncelet, *NNLO QCD corrections to three-photon production at the LHC*, *JHEP* **02** (2020) 057 [[arXiv:1911.00479](#)] [[INSPIRE](#)].
- [47] H.A. Chawdhry, M. Czakon, A. Mitov and R. Poncelet, *NNLO QCD corrections to diphoton production with an additional jet at the LHC*, *JHEP* **09** (2021) 093 [[arXiv:2105.06940](#)] [[INSPIRE](#)].
- [48] M. Czakon, A. Mitov and R. Poncelet, *Next-to-Next-to-Leading Order Study of Three-Jet Production at the LHC*, *Phys. Rev. Lett.* **127** (2021) 152001 [[arXiv:2106.05331](#)] [[INSPIRE](#)].
- [49] F. Caola, J.M. Lindert, K. Melnikov, P.F. Monni, L. Tancredi and C. Wever, *Bottom-quark effects in Higgs production at intermediate transverse momentum*, *JHEP* **09** (2018) 035 [[arXiv:1804.07632](#)] [[INSPIRE](#)].
- [50] J. Davies et al., *Double Higgs boson production at NLO: combining the exact numerical result and high-energy expansion*, *JHEP* **11** (2019) 024 [[arXiv:1907.06408](#)] [[INSPIRE](#)].
- [51] M. Grazzini, S. Kallweit, M. Wiesemann and J.Y. Yook, *W^+W^- production at the LHC: NLO QCD corrections to the loop-induced gluon fusion channel*, *Phys. Lett. B* **804** (2020) 135399 [[arXiv:2002.01877](#)] [[INSPIRE](#)].
- [52] M. Bonetti, E. Panzer, V.A. Smirnov and L. Tancredi, *Two-loop mixed QCD-EW corrections to $gg \rightarrow Hg$* , *JHEP* **11** (2020) 045 [[arXiv:2007.09813](#)] [[INSPIRE](#)].
- [53] M. Becchetti, F. Moriello and A. Schweitzer, *Two-loop amplitude for mixed QCD-EW corrections to $gg \rightarrow Hg$* , *JHEP* **04** (2022) 139 [[arXiv:2112.07578](#)] [[INSPIRE](#)].
- [54] S. Badger, T. Gehrmann, M. Marcoli and R. Moodie, *Next-to-leading order QCD corrections to diphoton-plus-jet production through gluon fusion at the LHC*, *Phys. Lett. B* **824** (2022) 136802 [[arXiv:2109.12003](#)] [[INSPIRE](#)].
- [55] C. Anastasiou, C. Duhr, F. Dulat, F. Herzog and B. Mistlberger, *Higgs Boson Gluon-Fusion Production in QCD at Three Loops*, *Phys. Rev. Lett.* **114** (2015) 212001 [[arXiv:1503.06056](#)] [[INSPIRE](#)].
- [56] C. Anastasiou et al., *High precision determination of the gluon fusion Higgs boson cross-section at the LHC*, *JHEP* **05** (2016) 058 [[arXiv:1602.00695](#)] [[INSPIRE](#)].
- [57] L. Cieri, X. Chen, T. Gehrmann, E.W.N. Glover and A. Huss, *Higgs boson production at the LHC using the q_T subtraction formalism at N^3LO QCD*, *JHEP* **02** (2019) 096 [[arXiv:1807.11501](#)] [[INSPIRE](#)].

- [58] C. Duhr, F. Dulat and B. Mistlberger, *Higgs Boson Production in Bottom-Quark Fusion to Third Order in the Strong Coupling*, *Phys. Rev. Lett.* **125** (2020) 051804 [[arXiv:1904.09990](#)] [[INSPIRE](#)].
- [59] X. Chen, T. Gehrmann, E.W.N. Glover, A. Huss, B. Mistlberger and A. Pelloni, *Fully Differential Higgs Boson Production to Third Order in QCD*, *Phys. Rev. Lett.* **127** (2021) 072002 [[arXiv:2102.07607](#)] [[INSPIRE](#)].
- [60] C. Duhr, F. Dulat and B. Mistlberger, *Charged current Drell-Yan production at N^3 LO*, *JHEP* **11** (2020) 143 [[arXiv:2007.13313](#)] [[INSPIRE](#)].
- [61] C. Duhr and B. Mistlberger, *Lepton-pair production at hadron colliders at N^3 LO in QCD*, *JHEP* **03** (2022) 116 [[arXiv:2111.10379](#)] [[INSPIRE](#)].
- [62] S. Camarda, L. Cieri and G. Ferrera, *Drell-Yan lepton-pair production: q_T resummation at N^3 LL accuracy and fiducial cross sections at N^3 LO*, *Phys. Rev. D* **104** (2021) L111503 [[arXiv:2103.04974](#)] [[INSPIRE](#)].
- [63] M. Grazzini, S. Kallweit and M. Wiesemann, *Fully differential NNLO computations with MATRIX*, *Eur. Phys. J. C* **78** (2018) 537 [[arXiv:1711.06631](#)] [[INSPIRE](#)].
- [64] R. Gauld, N. Glover, A. Huss, I. Majer and A. Gehrmann-De Ridder, *LHC observables with NNLOJET*, *PoS RADCOR2019* (2019) 002 [[INSPIRE](#)].
- [65] P. Banerjee, T. Engel, A. Signer and Y. Ulrich, *QED at NNLO with McMule*, *SciPost Phys.* **9** (2020) 027 [[arXiv:2007.01654](#)] [[INSPIRE](#)].
- [66] G. 't Hooft and M.J.G. Veltman, *Regularization and Renormalization of Gauge Fields*, *Nucl. Phys. B* **44** (1972) 189 [[INSPIRE](#)].
- [67] A. van Hameren, *Multi-gluon one-loop amplitudes using tensor integrals*, *JHEP* **07** (2009) 088 [[arXiv:0905.1005](#)] [[INSPIRE](#)].
- [68] F. Buccioni, S. Pozzorini and M. Zoller, *On-the-fly reduction of open loops*, *Eur. Phys. J. C* **78** (2018) 70 [[arXiv:1710.11452](#)] [[INSPIRE](#)].
- [69] G. Ossola, C.G. Papadopoulos and R. Pittau, *On the Rational Terms of the one-loop amplitudes*, *JHEP* **05** (2008) 004 [[arXiv:0802.1876](#)] [[INSPIRE](#)].
- [70] P. Draggiotis, M.V. Garzelli, C.G. Papadopoulos and R. Pittau, *Feynman Rules for the Rational Part of the QCD 1-loop amplitudes*, *JHEP* **04** (2009) 072 [[arXiv:0903.0356](#)] [[INSPIRE](#)].
- [71] M.V. Garzelli, I. Malamos and R. Pittau, *Feynman rules for the rational part of the Electroweak 1-loop amplitudes*, *JHEP* **01** (2010) 040 [Erratum *ibid.* **10** (2010) 097] [[arXiv:0910.3130](#)] [[INSPIRE](#)].
- [72] R. Pittau, *Primary Feynman rules to calculate the epsilon-dimensional integrand of any 1-loop amplitude*, *JHEP* **02** (2012) 029 [[arXiv:1111.4965](#)] [[INSPIRE](#)].
- [73] K.G. Chetyrkin and F.V. Tkachov, *Integration by Parts: The Algorithm to Calculate β -functions in 4 Loops*, *Nucl. Phys. B* **192** (1981) 159 [[INSPIRE](#)].
- [74] A. von Manteuffel and C. Studerus, *Reduze 2 — Distributed Feynman Integral Reduction*, [arXiv:1201.4330](#) [[INSPIRE](#)].
- [75] A. von Manteuffel and R.M. Schabinger, *A novel approach to integration by parts reduction*, *Phys. Lett. B* **744** (2015) 101 [[arXiv:1406.4513](#)] [[INSPIRE](#)].

- [76] H.A. Chawdhry, M.A. Lim and A. Mitov, *Two-loop five-point massless QCD amplitudes within the integration-by-parts approach*, *Phys. Rev. D* **99** (2019) 076011 [[arXiv:1805.09182](#)] [[INSPIRE](#)].
- [77] A.V. Smirnov and F.S. Chuharev, *FIRE6: Feynman Integral REduction with Modular Arithmetic*, *Comput. Phys. Commun.* **247** (2020) 106877 [[arXiv:1901.07808](#)] [[INSPIRE](#)].
- [78] J. Klappert, F. Lange, P. Maierhöfer and J. Usovitsch, *Integral reduction with Kira 2.0 and finite field methods*, *Comput. Phys. Commun.* **266** (2021) 108024 [[arXiv:2008.06494](#)] [[INSPIRE](#)].
- [79] T. Gehrmann and E. Remiddi, *Differential equations for two loop four point functions*, *Nucl. Phys. B* **580** (2000) 485 [[hep-ph/9912329](#)] [[INSPIRE](#)].
- [80] J.M. Henn, *Multiloop integrals in dimensional regularization made simple*, *Phys. Rev. Lett.* **110** (2013) 251601 [[arXiv:1304.1806](#)] [[INSPIRE](#)].
- [81] C.G. Papadopoulos, D. Tommasini and C. Wever, *The Pentabox Master Integrals with the Simplified Differential Equations approach*, *JHEP* **04** (2016) 078 [[arXiv:1511.09404](#)] [[INSPIRE](#)].
- [82] T. Gehrmann, J.M. Henn and N.A. Lo Presti, *Pentagon functions for massless planar scattering amplitudes*, *JHEP* **10** (2018) 103 [[arXiv:1807.09812](#)] [[INSPIRE](#)].
- [83] D. Chicherin, T. Gehrmann, J.M. Henn, N.A. Lo Presti, V. Mitev and P. Wasser, *Analytic result for the nonplanar hexa-box integrals*, *JHEP* **03** (2019) 042 [[arXiv:1809.06240](#)] [[INSPIRE](#)].
- [84] D. Chicherin and V. Sotnikov, *Pentagon Functions for Scattering of Five Massless Particles*, *JHEP* **20** (2020) 167 [[arXiv:2009.07803](#)] [[INSPIRE](#)].
- [85] S. Abreu, H. Ita, F. Moriello, B. Page, W. Tschernow and M. Zeng, *Two-Loop Integrals for Planar Five-Point One-Mass Processes*, *JHEP* **11** (2020) 117 [[arXiv:2005.04195](#)] [[INSPIRE](#)].
- [86] S. Abreu, H. Ita, B. Page and W. Tschernow, *Two-loop hexa-box integrals for non-planar five-point one-mass processes*, *JHEP* **03** (2022) 182 [[arXiv:2107.14180](#)] [[INSPIRE](#)].
- [87] C. Duhr, V.A. Smirnov and L. Tancredi, *Analytic results for two-loop planar master integrals for Bhabha scattering*, *JHEP* **09** (2021) 120 [[arXiv:2108.03828](#)] [[INSPIRE](#)].
- [88] S. Pozzorini, H. Zhang and M.F. Zoller, *Rational Terms of UV Origin at Two Loops*, *JHEP* **05** (2020) 077 [[arXiv:2001.11388](#)] [[INSPIRE](#)].
- [89] J.-N. Lang, S. Pozzorini, H. Zhang and M.F. Zoller, *Two-Loop Rational Terms in Yang-Mills Theories*, *JHEP* **10** (2020) 016 [[arXiv:2007.03713](#)] [[INSPIRE](#)].
- [90] J.-N. Lang, S. Pozzorini, H. Zhang and M.F. Zoller, *Two-loop rational terms for spontaneously broken theories*, *JHEP* **01** (2022) 105 [[arXiv:2107.10288](#)] [[INSPIRE](#)].
- [91] A. Bredenstein, A. Denner, S. Dittmaier and S. Pozzorini, *NLO QCD corrections to $t\bar{t}b\bar{b}$ production at the LHC: 1. Quark-antiquark annihilation*, *JHEP* **08** (2008) 108 [[arXiv:0807.1248](#)] [[INSPIRE](#)].
- [92] A. Gehrmann-De Ridder, T. Gehrmann and E.W.N. Glover, *Antenna subtraction at NNLO*, *JHEP* **09** (2005) 056 [[hep-ph/0505111](#)] [[INSPIRE](#)].
- [93] S. Catani and M. Grazzini, *An NNLO subtraction formalism in hadron collisions and its application to Higgs boson production at the LHC*, *Phys. Rev. Lett.* **98** (2007) 222002 [[hep-ph/0703012](#)] [[INSPIRE](#)].

- [94] G. Somogyi, Z. Trócsányi and V. Del Duca, *Matching of singly- and doubly-unresolved limits of tree-level QCD squared matrix elements*, *JHEP* **06** (2005) 024 [[hep-ph/0502226](#)] [[INSPIRE](#)].
- [95] M. Czakon, *A novel subtraction scheme for double-real radiation at NNLO*, *Phys. Lett. B* **693** (2010) 259 [[arXiv:1005.0274](#)] [[INSPIRE](#)].
- [96] R. Boughezal, X. Liu and F. Petriello, *N-jettiness soft function at next-to-next-to-leading order*, *Phys. Rev. D* **91** (2015) 094035 [[arXiv:1504.02540](#)] [[INSPIRE](#)].
- [97] M. Cacciari, F.A. Dreyer, A. Karlberg, G.P. Salam and G. Zanderighi, *Fully Differential Vector-Boson-Fusion Higgs Production at Next-to-Next-to-Leading Order*, *Phys. Rev. Lett.* **115** (2015) 082002 [*Erratum ibid.* **120** (2018) 139901] [[arXiv:1506.02660](#)] [[INSPIRE](#)].
- [98] X. Chen, T. Gehrmann, N. Glover, A. Huss and M. Marcoli, *Automation of antenna subtraction in colour space: gluonic processes*, [arXiv:2203.13531](#) [[INSPIRE](#)].
- [99] J.A.M. Vermaseren, *Axodraw*, *Comput. Phys. Commun.* **83** (1994) 45 [[INSPIRE](#)].