

Integration-by-parts reductions of Feynman integrals using Singular and GPI-Space

Dominik Bendle,^{a,b} Janko Böhm,^a Wolfram Decker,^a Alessandro Georgoudis,^c
Franz-Josef Pfreundt,^b Mirko Rahn,^b Pascal Wasser^d and Yang Zhang^{e,f}

^a*Department of Mathematics, Technische Universität Kaiserslautern,
67663 Kaiserslautern, Germany*

^b*Fraunhofer Institute for Industrial Mathematics ITWM,
Fraunhofer-Platz 1, 67663 Kaiserslautern, Germany*

^c*Department of Physics and Astronomy, Uppsala University,
SE-75108 Uppsala, Sweden*

^d*PRISMA+ Cluster of Excellence, Johannes Gutenberg University,
D-55128 Mainz, Germany*

^e*Interdisciplinary Center for Theoretical Study,
University of Science and Technology of China, Hefei, Anhui 230026, China*

^f*Max-Planck-Institut für Physik, Werner-Heisenberg-Institut,
D-80805 München, Germany*

E-mail: bendle@rhrk.uni-kl.de, boehm@mathematik.uni-kl.de,
decker@mathematik.uni-kl.de, Alessandro.Georgoudis@physics.uu.se,
franz-josef.pfreundt@itwm.fhg.de, mirko.rahn@itwm.fhg.de,
wasserp@uni-mainz.de, yzhang@mpp.mpg.de

ABSTRACT: We introduce an algebro-geometrically motivated integration-by-parts (IBP) reduction method for multi-loop and multi-scale Feynman integrals, using a framework for massively parallel computations in computer algebra. This framework combines the computer algebra system SINGULAR with the workflow management system GPI-SPACE, which are being developed at the TU Kaiserslautern and the Fraunhofer Institute for Industrial Mathematics (ITWM), respectively. In our approach, the IBP relations are first trimmed by modern tools from computational algebraic geometry and then solved by sparse linear algebra and our new interpolation method. Modelled in terms of Petri nets, these steps are efficiently automatized and automatically parallelized by GPI-SPACE. We demonstrate the potential of our method at the nontrivial example of reducing two-loop five-point non-planar double-pentagon integrals. We also use GPI-SPACE to convert the basis of IBP reductions, and discuss the possible simplification of master-integral coefficients in a uniformly transcendental basis.

KEYWORDS: Scattering Amplitudes, Differential and Algebraic Geometry

ARXIV EPRINT: [1908.04301](https://arxiv.org/abs/1908.04301)

Contents

1	Introduction	1
2	The module intersection method reloaded	4
2.1	Module intersection	4
2.2	Linear reduction	7
3	Massively parallel computations using Singular and GPI space	8
4	Parallel matrix reduction as a Petri net	10
4.1	General structure of the algorithm	10
4.2	Managing the interpolation	11
4.3	Description of the Petri net	11
4.4	Parallel timings	15
5	IBP conversion between different integral bases	18
5.1	Dlog bases and the dlog algorithm	18
5.2	IBP reduction with a dlog basis	20
6	The two-loop five-point nonplanar double pentagon example	20
6.1	Module intersection with cuts	22
6.2	IBP reduction	24
6.3	Master-integral coefficient conversion to a dlog basis	24
7	Summary	25
A	Rational function interpolation	26

1 Introduction

With the success of Large Hadron Collider (LHC) Run II and the upcoming LHC run III, high precision background computation, especially next-to-next-to-leading-order (NNLO) scattering computation, is crucial for the interpretation of experimental results. In recent years, great progress has been made in multi-loop scattering amplitude calculations, for instance, in the case of $2 \rightarrow 3$ processes [1–15]. The progress is due to modern developments of scattering amplitudes, like the integrand construction method [16, 17], canonical integrals [18, 19], numeric unitarity [20, 21], bootstrap methods [22–29], reconstruction using finite fields [30–33] and new ideas in the integration-by-parts (IBP) reduction. The latter is the main topic to be discussed in this paper.

Frequently, when computing scattering amplitudes, IBP reduction is a crucial and bottleneck step. It is a fundamental tool for both the reduction of integrals to master integrals (MIs), and for computing the master integrals themselves using the differential equation method. IBP relations (IBPs) are derived from integrating a total derivative in a Feynman loop integral [34],

$$0 = \int \frac{d^D \ell_1}{i\pi^{D/2}} \cdots \frac{d^D \ell_L}{i\pi^{D/2}} \sum_{j=1}^L \frac{\partial}{\partial \ell_j^\mu} \frac{v_j^\mu}{D_1^{n_1} \cdots D_m^{n_m}}, \quad (1.1)$$

where ℓ 's are dimensionally regularized loop momenta and D is the spacetime dimension. The D_i 's are the denominators of propagators in Feynman integrals. The vectors v_i^μ are linearly spanned by the loop momenta as well as the external momenta. By expanding the integrand inside (1.1), we are able to obtain IBP relations between different Feynman integrals.

The standard approach to obtain IBP reductions, by which we are able to express an integral as a linear combination of a finite number of MIs, is to generate sufficiently many IBP relations, and then use the Laporta algorithm [35] to solve the associated linear system. The algorithm works by imposing an ordering on the different integral families and solving recursively. There exist multiple public and private implementations of this approach [32, 36–41], which usually generates a large linear system to be solved.

In the case of a system of IBPs which does not have double propagators [42–44], however, we obtain a much smaller linear system. The IBPs without double propagators are physically related to dual conformal symmetries [45]. A significant simplification can be made by using unitarity methods, where by considering a spanning set of generating cuts it is possible to reduce the size of the IBP system. This requires prior knowledge of a basis of MIs. Such a basis can be obtained by running the Laporta algorithm with constant kinematics, or by using specialized programs such as MINT [46] or AZURITE [47]. (Note that the dimension of a basis of integrals can also be obtained by studying the parametric annihilators [48].) There is also the important technique [49] of simultaneously nullifying all master integrals except one, which often makes large-scale linear reductions feasible.

Besides the advances in purely analytical methods in recent years, there has been a lot of work towards numerical implementations of the generation of IBPs. The idea is to utilize either integer values or finite-field values for the kinematical invariants [30, 31, 38], depending on the difficulty of the problem, and then to run the same reduction several times for reconstruction. This method has been very successful in tackling difficult problems. Furthermore, it is possible to numerically generate and reduce the IBP relations, and, while skipping the master-integral coefficient reconstruction in IBP identities, directly carry out an amplitude reconstruction. (For examples, see [9, 10, 13, 50]). In this paper, we in particular present our own implementation of a semi-numeric rational interpolation method, see appendix A for more details.

Furthermore, new approaches were developed recently to obtain the reduction directly, without generating IBP relations from total derivatives. In [51], the direct solution method was presented to derive recurrence relations for Feynman integrals and solve them analytically.

ically with arbitrary numerator degree. One very promising progress is based on the intersection theory of differential forms in the Baikov representation [52–54]. This approach calculates the master integral coefficients from intersection numbers. There is also a very intuitive approach to reduce Feynman integrals by considering the η expansion of the Feynman prescription [55–57]. Using this approach, the scaling of the reduction computation depends only linearly on the number of master integrals. Furthermore, it is possible to determine two-loop planar diagram master-integral coefficients directly from the Baikov representation [58].

In this paper, we present our new powerful IBP reduction method based on:

1. Computational algebraic geometry. Computational algebraic geometry method was used in [42] for generating IBP without double propagators via syzygy computations. The geometric integrand decomposition method [20] has been successfully developed for constructing multi-loop integrand and IBP reduction [4, 6, 21].

In this paper, we apply the module intersection method from [59–61], modified by using a suitably chosen degree bound for the Gröbner basis computation, to efficiently generate a small IBP system, without double propagators (or IBPs with a given bound on the propagator exponents). Like [42], this method also restrict the degree of propagator. However module intersection method uses the Baikov presentation instead of the Feynman momentum representation. We find it is easy to generate trimmed IBPs in Baikov representation and apply cuts, with our module intersection representation.

2. A modern framework of massively parallel computations in computer algebra which combines the computer algebra system SINGULAR [62] with the workflow management system GPI-SPACE [63]. We have completely automatized our approach and make our algorithms run automatically in parallel on high performance computing clusters. In this way, IBP results can be obtained in an efficient, reliable and scalable way. Our implementation can automatically determine the minimal number of points needed for interpolating the master-integral coefficients, it can identify possible “bad” points, add more points, if necessary, and interpolate the final result.

We demonstrate the power of our method by reducing the two-loop five-point nonplanar double pentagon diagram *analytically*, up to numerator degree 4. This is a nontrivial test since the diagram has a complicated topology and there are five symbolic Mandelstam variables as well as the spacetime variable D .

Furthermore, we start to look at the possible simplification of master-integral coefficients by converting the master integral basis. In this paper, we test the conversion to a “dlog” basis [64], a special case of the canonical basis [19]. We find that for the double pentagon diagram above, the size of the master integral coefficients reduces significantly from the byte size $\sim 2.0\text{G}$ in the Laporta basis to $\sim 0.48\text{G}$ in the dlog basis on disk, that is, by 76%. The master integral basis conversion computation is also automated by the SINGULAR-GPI-SPACE framework.

Our paper is structured as follows. In section 2, we present the general background on how to generate simple and trimmed IBP systems using computational algebraic geometry and finite-field methods, as well as the improvement on the algorithm in [60]. In section 3, we give a short overview on how we use SINGULAR in conjunction with GPI-SPACE. In section 4, we describe how to model our algorithm in the SINGULAR-GPI-SPACE framework, and discuss timings and scaling of the algorithm, focusing on the double pentagon diagram. This, in particular, demonstrates the potential of the SINGULAR-GPI-SPACE framework for applications in high-energy physics. In section 5, we review the algorithmic computation of a dlog basis which has uniform transcendental weight, and we comment on how to convert coefficients from the Laporta basis to the dlog basis. In section 6, we study the working example of our implementation, the double pentagon graph, in detail. We discuss the analytic IBP reduction and the conversion of master integral coefficients to the dlog basis. Finally we present a summary and conclusion of this paper.

The result of our IBP reductions can be downloaded from the following links: whereas

https://www.dropbox.com/s/1ubdhcyhe8e4pwy/IBPmatrix_Laporta_basis.tar.gz

provides the master integral coefficients in the Laporta basis with the scale $s_{12} = 1$,

https://www.dropbox.com/s/e6t4evftkfo95pr/IBPmatrix_dlog_basis.tar.gz

contains the master-integral coefficients in the dlog basis with the scale $s_{12} = 1$.

For the convenience of the reader, we also present the master-integral coefficients in the dlog basis with the full scale dependence:

https://www.dropbox.com/s/dnkr6h5t3vik2r0/IBPmatrix_dlog_basis_scaled.tar.gz

We encourage researchers in the high energy community to send us IBP reduction problems (mailto: alessandro.georgoudis@physics.uu.se) for cutting-edge precision calculations and the further sharpening of our new reduction method.

2 The module intersection method reloaded

In this section, we present a refined version of the approach of using module intersections to trim IBP systems. We refer to [42] for the original idea of generating IBP without double propagator, and [6] for the recent development of the geometric integrand decomposition and numeric unitarity method.

For the detailed account of the module intersection IBP reduction method, we refer to [60].

2.1 Module intersection

The Feynman integrals with dimensional regularization under consideration are labeled as

$$I[n_1, \dots, n_m] = \int \frac{d^D \ell_1}{i\pi^{D/2}} \cdots \frac{d^D \ell_L}{i\pi^{D/2}} \frac{1}{D_1^{n_1} \cdots D_m^{n_m}}, \quad n_i \in \mathbb{Z}, \quad i = 1, \dots, m \quad (2.1)$$

where L is the loop order and the l_i 's are the dimensional regularized loop momenta. D_i 's are the denominators of propagators in the momentum representation. We have E linearly independent external momenta that we label as p_1, \dots, p_E . We assume that the Feynman integrals have been reduced on the integrand level, and thus set $m = LE + L(L + 1)/2$ which equals the number of reducible + irreducible scalar products in the configuration.

For us, it is convenient to use the Baikov representation [46, 65] for IBP reductions,

$$I[n_1, \dots, n_m] = C_E^L U^{\frac{E-D+1}{2}} \int_{\Omega} dz_1 \cdots dz_m P^{\frac{D-L-E-1}{2}} \frac{1}{z_1^{n_1} \cdots z_m^{n_m}}. \quad (2.2)$$

Here, we redefined the scalar products $z_i \equiv D_i$, $i = 1, \dots, m$. P is the Baikov polynomial, which can be written as a Gram determinant,

$$P = \det G \begin{pmatrix} l_1, \dots, l_L, p_1, \dots, p_E \\ l_1, \dots, l_L, p_1, \dots, p_E \end{pmatrix}. \quad (2.3)$$

Since P only consists of reducible and irreducible scalar products, P can be reformulated as a polynomial explicitly in z_i 's.

Moreover, U and C_E^L are respectively the Gram determinant and a constant factor below:

$$U = \det G \begin{pmatrix} p_1, \dots, p_E \\ p_1, \dots, p_E \end{pmatrix}, \quad C_E^L = J \frac{\pi^{\frac{L-m}{2}}}{\Gamma(\frac{D-E-L+1}{2}) \cdots \Gamma(\frac{D-E}{2})}, \quad (2.4)$$

where J is a constant Jacobian. The factors U and C_E^L are irrelevant for deriving the IBP relations.

As in [20, 44, 66], the IBP relations in the Baikov representation are of type

$$0 = \int dz_1 \cdots dz_m \sum_{i=1}^m \frac{\partial}{\partial z_i} \left(a_i(z) P^{\frac{D-L-E-1}{2}} \frac{1}{z_1^{n_1} \cdots z_m^{n_m}} \right), \quad (2.5)$$

where each $a_i(z)$ is a polynomial in the variables z_1, \dots, z_m . Note that P vanishes on the boundary of the Baikov integration domain, so this form of IBP identities does not have surface terms.

Suppose we wish to reduce an integral family with $n_j \leq 0$, $j = \kappa + 1, \dots, m$, for some κ . That is, we face integrals with the inverse propagator product $1/(D_1 \dots D_{\kappa})$ and the sub-topology integrals. We use the idea of restricting to IBP systems without double propagators [42], choosing suitable $a_i(z)$ to prevent the appearance of double propagators in (2.5). In the Baikov representation, we also need to avoid total derivatives with dimension shifts [20, 44]. These constraints translate into syzygy equations of the following type:

$$\left(\sum_{i=1}^m a_i(z) \frac{\partial P}{\partial z_i} \right) + b(z)P = 0, \quad (2.6)$$

$$a_i(z) = b_i(z)z_i, \quad i = 1, \dots, \kappa, \quad (2.7)$$

where $b(z)$ and the $b_i(z)$ are also polynomials in z_i 's. Relation (2.6) avoids dimension shifts of the integrals, while (2.7) ensures that there is no double propagator for D_i if the initial

index $n_i = 1$ in (2.5). The goal is to find such polynomials $a_i(z)$, $b(z)$, and $b_i(z)$. Since we require polynomial solutions, this is not a linear algebra problem, but a computational algebraic geometry problem.

We use the module intersection method from [59, 60] to solve (2.6) and (2.7) simultaneously. Note that the analytic generators of all solutions of (2.6) can be directly written down via either the canonical IBP vector method [20] or the Gram matrix Laplace expansion method [61].¹ We remark that the Gram matrix Laplace expansion method does not provide more relations beyond the IBP formalism with loop momenta and other vectors. This Laplace expansion procedure is to convert the standard IBP to the Baikov representation. Ref. [6] achieved the conversion without resorting to a Laplacian representation.

The relations in (2.7) can be trivially expressed as a module membership condition. Hence without any algorithmic computation, we know the individual solutions for (2.6) and (2.7), respectively.

These form polynomial submodules M_1 respectively M_2 of R^m over the polynomial ring $R = \mathbb{Q}(c_1, \dots, c_k)[z_1, \dots, z_m]$ (where the variables c_1, \dots, c_k collect the Mandelstam variables and the mass parameters). The task is then to compute

$$M_1 \cap M_2. \tag{2.8}$$

This module intersection can be obtained by computing a module Gröbner basis in a particular ordering [60]. One decisive strategy is the *localization* technique, which allows us to compute $M_1 \cap M_2$ over the polynomial ring $\tilde{R} = \mathbb{Q}[c_1, \dots, c_k, z_1, \dots, z_m]$. In this manner, we treat kinematic variables in the same way as the Baikov variables. This greatly speeds up the intersection computation for multi-scale problems, but results in a redundant generating system. The latter can be trimmed further by importing the result back to R^m and removing redundant generators by checking the leading monomials. This is powered by SINGULAR's command `simplify`. Once $M_1 \cap M_2$ is obtained, we know all simultaneous solutions for (2.6) and (2.7), and can use (2.5) to get IBPs without double propagators.

We emphasize that, although (2.6) and (2.7) were originally designed for IBPs without double propagators, the solutions of (2.6) and (2.7) can be used to simplify IBP systems *with* double or multiple propagators. Using these solutions $a_i(z)$, the resulting IBP system does not introduce integrals with higher powers of propagators, and hence also greatly decreases the size of the IBP system.

Frequently, instead of computing IBPs directly, we compute IBPs on spanning cuts and assemble the full IBPs afterwards. This amounts to setting some of the z_i to zero in (2.6) and (2.7). For details on IBPs on cuts using the Baikov representation, we refer to [60].

Compared to the approach in [60], in this paper we use the following improvements for the module intersection method. These ideas have been used in [6].

- When we compute the intersection $M_1 \cap M_2$, instead of finding a full generating system, we heuristically impose a polynomial degree bound in the computation. Then we reduce the resulting IBPs over finite fields to test if we already have all the IBP

¹We learned the Laplace expansion method from Roman Lee, and proved its completeness via the Gulliksen-Negard/Jozefiak exact sequence in [61].

relations needed. If the IBP relations are insufficient, we increase the degree bound and repeat the computation. This approach speeds up the intersection computation dramatically in many cases. In practice, we use the option `degbound` in the computer algebra software SINGULAR to set the degree bound.

- In the approach of [60], the module intersection was only computed for the top sector, which, for the hexagon box diagram, turned out to be sufficient for reducing integrals to a master integral basis. However, in this paper, we compute the module intersection for the top sector and also all subsectors. This approach may, in general, generate more IBP relations. Via linear algebra trimming as discussed in the next subsection, this approach eventually gives a block triangular linear system and makes the linear reduction easier.

2.2 Linear reduction

For the simplified IBP system arising from the module-intersection method, we use our own linear reduction algorithm to reduce the IBP system. The steps are:

1. Trim the linear system in two stages: (a) Set all the kinematic variables to integer values, and use linear algebra over a finite field to find the independent IBP relations. (b) Again over a finite field, carry out the reduction. From the intermediate steps, determine a sufficient subset of IBP relations for reducing the target integrals. These operations are powered by the finite field computation tool SPASM [67].
2. Remove the overlap between two different cuts and simplify the linear system: if two cuts have a common master integral, use the idea from [49] to set the master integral to zero in the IBP system of one of the two cuts. This will later on dramatically simplify the IBP reduction for the cut.
3. For the linear system simplified by the first two steps, we use our own SINGULAR row reduce echelon form (RREF) code over function fields to reduce the target integrals to master integrals. Our code applies both row and column swaps for finding the optimal pivots. Note that column swaps change the set of master integrals. After the RREF computation, we convert the new master integrals to the original master integrals. We have observed that this approach is in general much faster than fixing the column ordering and directly reducing the target integrals to the original master integrals.

For difficult IBP reduction computations, we use a “semi-numeric” approach: this approach sets several but usually not all of the kinematic variables for the reduction computation to numeric values (that is, to constant integers). Without loss of generality, for the kinematic variables (c_1, \dots, c_k) , we set

$$c_i \mapsto a_i, \quad 1 \leq i \leq k_1, \quad (2.9)$$

for some $k_1 < k$ and some $a_i \in \mathbb{Z}$.

The actual degree of the coefficients in these variables can be decided by a univariate analytic computation (that is, we set all but one of the c_i to constant values). For example, we may pick the dimension D and all parameters c_i except c_1 as random integers, and then carry out the reduction. This computation is much easier than the actual IBP reduction with fully analytic parameters. From the reduction, we determine the degree of c_1 in the final IBP reduction coefficients. Proceeding similarly for each i , we find the degree of each c_i . This determines the minimal number of semi-numeric points for the subsequent interpolation step. (See [31] for an alternative way of finding the degree of each parameter in a rational function.)

After accumulating enough points, we collect the semi-numeric reduction results and interpolate to get the final IBP reduction coefficients. To do this, we first run step 3 above for a semi-numeric set of parameters, find the optimal pivots and record the row/column swap history as a *trace* of our computation. For other numeric values, we always use the same trace to ensure the relatively uniform running time of the computation.

In practice, we use our rational function interpolation algorithm described in appendix A. We do a reduction computation, with a carefully chosen semi-numeric reference point,

$$c_j \mapsto b_j, \quad b_j \in \mathbb{Z}, \quad k_1 < j \leq k, \quad (2.10)$$

and c_1, \dots, c_{k_1} symbolic. Using the reference point result, we convert the rational function interpolation problem to individual polynomial interpolation problems for the numerators and denominators. With this approach, the number of “semi-numeric” computations is

$$(d_1 + 1) \times (d_2 + 1) \times \dots \times (d_{k_1} + 1), \quad (2.11)$$

where the d_i , for $1 \leq i \leq k_1$, are the maximal degrees of the c_i in the numerator and denominator polynomials in the RREF matrix. This algorithm is also implemented in SINGULAR.

For the semi-numerical reduction and interpolation, we need to parallelize our computations in an efficient way. Furthermore, with semi-numeric points, we may have some bad points in the reduction or interpolation. In order to make use of massively parallel computations in an efficient way, and to automatize the workflow for the replacement of bad points, we use the modern workflow management system GPI-SPACE, in conjunction with the computer algebra system SINGULAR. We will discuss the ideas behind this approach in the subsequent section.

3 Massively parallel computations using Singular and GPI space

Large scale calculations such as row reductions of IBP identities in the case of Feynman diagrams which are relevant to current research in high-energy physics, are only feasible by using parallel computing on high-performance clusters. The computer algebra methods applied in this context require to model algorithms which rely on sub-computations with time and memory requirements that are difficult to predict. This is due, for example, to the behaviour of Buchberger’s algorithm for finding Gröbner bases: although this algorithm




Figure 1. An enabled transition and its firing.




Figure 2. A transition which is not enabled.

performs well in many practical examples of interest, its worst case complexity is doubly exponential in the number of variables [68]. Nevertheless it turned out recently [69, 70] that massively parallel methods, which have been a standard tool in numerical simulation for many years, can also be applied successfully in symbolic computation. Proposing the general use of massively parallel methods in computer algebra, we describe our ongoing effort in this direction which is based on connecting the computer algebra system SINGULAR for polynomial calculations with the workflow management GPI-SPACE (for an introduction see [71]). The latter consists of a scheduler distributing the actual computations to workers in the cluster, a virtual memory layer to facilitate communication between the workers, and a workflow management system which relies on modeling algorithms in terms of Petri nets.

In its basic form, a Petri net is a directed bipartite graph with two kinds of nodes: while a *place* can hold a number of indistinguishable (structure-less) *tokens*, a *transition* may *fire* if each input place contains at least one token (we then say that the transition is *enabled*). When fired, a transition consumes one token from each input place and puts one token on each output place. See figure 1 for an enabled transition and its firing, and figure 2 for a transition which is not enabled. In the figures, places are shown as circles, transitions as rectangles, and tokens as black dots.

The execution of a Petri net is non-deterministic: at each step, a single random enabled transition is chosen to fire. We have observed that the randomized reformulation of deterministic algorithms in computer algebra in terms of Petri nets can lead to a more consistent and predictable behavior throughout the course of the computation.

In our approach, we model the coarse-grained structure of an algorithm in terms of a Petri net. The transitions call procedures from the C-library version of SINGULAR to do the actual computations. The result of this setup is a flexible framework for massively parallel computations in computational algebraic geometry (similar setups are possible using C-libraries of computer algebra systems aiming at possibly different application areas). Our framework has, for example, already been used to implement a non-singularity test for algebraic varieties [69, 72], the computation of combinatorial objects in geometric invariant theory [73], and the computation of tropical varieties associated to algebraic varieties [74].

For the efficient use in practical programming, the basic concept of a Petri net has to be extended. Here, GPI-SPACE provides multiple additional features:




Figure 3. Task and data parallelism in a Petri net.

- Modeling complex algorithms just by the use of structure-less tokens is not very efficient. In GPI-SPACE, tokens can have a data type and hold actual data. In fact, it is often more efficient if the tokens just hold a reference to a storage place for the data (in memory or in the file system). Using the shared memory subsystem of GPI-SPACE or the powerful file systems of modern high-performance clusters, computations can then scale far beyond the limitations of a single machine.
- The firing of a transition may be subject to conditions which have to be fulfilled by the input tokens.
- Transitions in practice involve computations which take time. The properties of Petri nets allow us to execute different enabled transitions at the same time (*task parallelism*) and to execute multiple instances of the same transition in parallel, provided the input places hold multiple tokens (*data parallelism*). In figure 3, the transitions f_1 and f_2 can fire in parallel, and, if the input place of f_i holds multiple tokens, then f_i can fire in multiple instances.

We have observed that some algorithms in computer algebra scale in a superlinear way when implemented in parallel as a Petri net. The reason is that then, at run time, the algorithms can automatically determine from a given set of paths a path which leads to the solution in the fastest possible way (see [69, section 6.2]).

In the next section, we illustrate the use of the SINGULAR-GPI-SPACE framework for applications in high-energy physics by modeling our IBP reduction algorithm.

4 Parallel matrix reduction as a Petri net

In this section, we first describe how to model the parallel IBP reduction algorithm in terms of a Petri net. Focusing on the cut $\{1, 3, 4, 5\}$ of the two-loop five-point nonplanar double pentagon diagram, we then discuss timings and scaling of the algorithm to indicate the practical use and significant potential of the SINGULAR-GPI-SPACE framework for algorithmic problems in high-energy physics.

4.1 General structure of the algorithm

Our approach includes a massively parallel execution of row-reductions over function fields, where a number of parameters has been replaced by integers, followed by a parallel interpolation step to reconstruct the dependency on these parameters.

So the task is to find the reduced row-echelon form M_{red} of a large linear system of equations, given as a matrix M over the rational function field $\mathbb{Q}(c_1, \dots, c_k)$. Since

applying Gaussian elimination directly is not feasible, we instead proceed by substituting, say, the first r parameters by the coordinates of a point $a \in \mathbb{Z}^r$, and then by computing the reduction

$$(M|_{c_1 \mapsto a_1, \dots, c_r \mapsto a_r})_{\text{red}}.$$

We refer to section 2.2 above for details on how we handle this reduction step. To determine the number of interpolation points required to reconstruct the dependency on c_1, \dots, c_r , we find bounds for the degrees of numerators and denominators for each parameter by doing a univariate row reduction (that is, all but one of the parameters are set to be numeric). After the reduction, we check that the resulting matrix is equal to the desired result

$$M_{\text{red}}|_{c_1 \mapsto a_1, \dots, c_r \mapsto a_r}$$

by normalizing it relative to a previously computed reference matrix with c_{r+1}, \dots, c_k constant, and performing degree checks using the exact degrees obtained from the univariate calculations. These steps are described in more detail in appendix A. The final result M_{red} is then found by iteratively combining the reduced matrices via univariate interpolation (see again appendix A).

Let d_1, \dots, d_r be degree bounds for the entries of M_{red} in the parameters c_1, \dots, c_r , respectively. To obtain M_{red} by interpolation, we need $d_1 + 1$ matrices over $\mathbb{Q}(c_2, \dots, c_k)$ of the form

$$M_{\text{red}}|_{c_1 \mapsto a_1^{(0)}, \dots}, M_{\text{red}}|_{c_1 \mapsto a_1^{(d_1)}}, \tag{4.1}$$

for $d_1 + 1$ values $a_1^{(0)}, \dots, a_1^{(d_1)} \in \mathbb{Z}$. Similarly, to obtain any one of the above matrices, we need $d_2 + 1$ matrices over $\mathbb{Q}(c_3, \dots, c_k)$. Continuing inductively, this process ends with matrices defined over $\mathbb{Q}(c_{r+1}, \dots, c_k)$, which are then computed by reduction with c_1, \dots, c_r numeric. This tree-like dependency structure is depicted in figure 4.

4.2 Managing the interpolation

We model the current status of the interpolation process in a tree-like data structure corresponding to that from figure 4, with references to the reduction results at the leaves, and references to the interpolation results at the other nodes. Within GPI-SPACE, reductions and interpolations are executed according to this data structure. The tree is generated as soon as the degree bounds d_1, \dots, d_r are known, and it is extended if the algorithm requires additional data points due to the occurrence of bad interpolation points.

4.3 Description of the Petri net

Figure 5 depicts the Petri net that implements the complete reduction algorithm. Going beyond the standard syntax introduced in section 3, dashed arrows stand for read-only access, that is, the data in the respective places is not consumed. The dotted arrows illustrate read and write access to the interpolation tree described in section 4.2. A transition can be annotated by conditions which indicate that the transition can only fire by con-

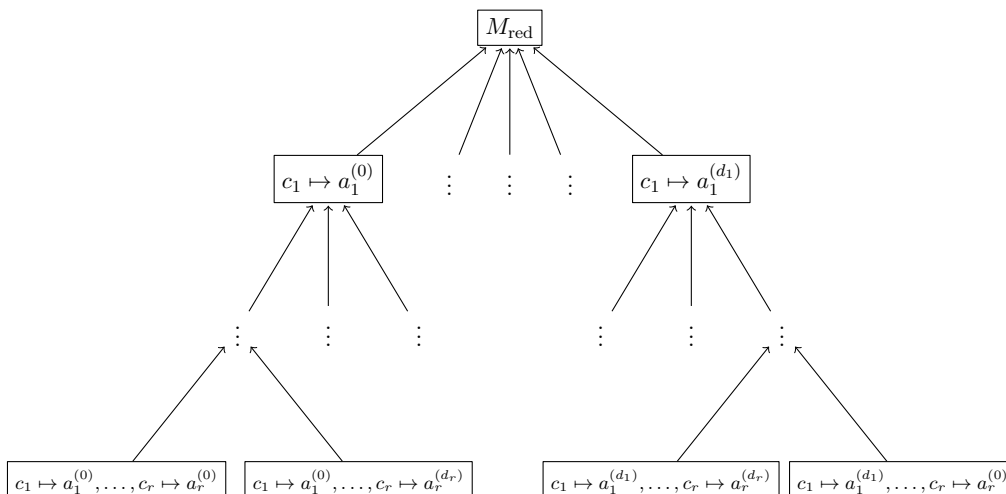


Figure 4. The structure of the interpolation tree.

suming tokens for which the conditions evaluate to true.² In the following, we describe the individual structures of the net:

Input token: The net is initialized with one token: a token on the place I , which holds references to the following input data:

- The input linear relations, which are given as a matrix M over the rational function field $\mathbb{Q}(c_1, \dots, c_k)$.
- The vector of indices of the parameters which will be interpolated (in the following we assume that these indices are $1, \dots, r$).
- The vector of indices of the target variables.
- Optionally: a precomputed trace for the reduction step (consistent with the targets). In the Petri net, the trace is referred to as $I.\text{trace}$ (we use the usual dot-notation for sub-data structures). Note that the trace fixes the variables corresponding to the master integrals.

Transition trace: If the token on I does not contain a trace, then `trace` is enabled, computes a trace for the linear reduction (recall that `trace` is the optimal pivot strategy and the row/column swap history, and see section 2.2 for the discussion there) and returns a copy of I with the trace included.

Transition copy: If the token on I already contains a trace, then `copy` is enabled and simply passes the token on I through.

Transition init: This transition takes the input token, which was produced by either `trace` or `copy`, and pushes it onto I_t . This way, the input data on I_t is guaranteed to contain trace data. It additionally enables the transitions `degrees` and `reference`.

²When formulating conditions in the Petri net, we use the name of a place and a token on the place interchangeably.

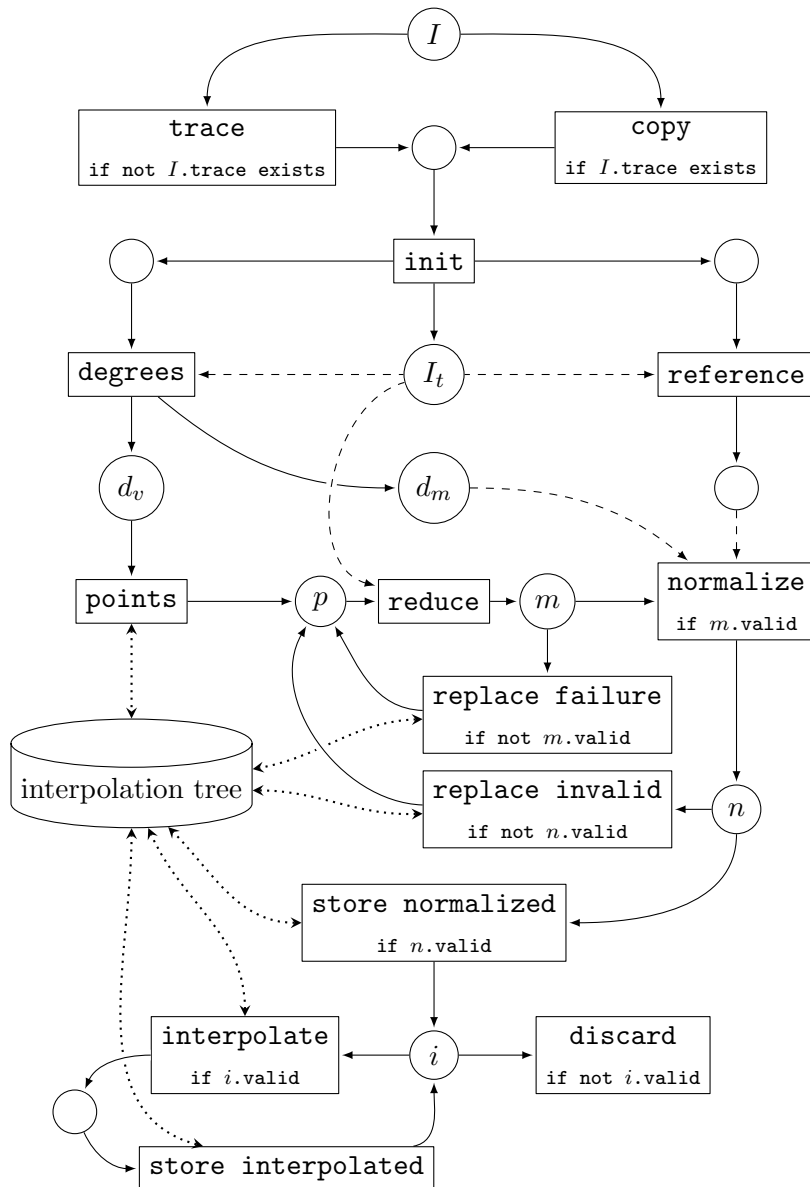


Figure 5. The Petri net for row reduction via interpolation. A description of the syntax is given in the first paragraph of section 4.3.

Transition reference: This generates a random substitution point $q = (q_{r+1}, \dots, q_k)$ with values for all parameters which will *not* be interpolated, substitutes the q_i for the c_i , and runs the row reduction step (see section 2.2), that is, computes

$$M_{\text{ref}} := (M|_{c_{r+1} \mapsto q_{r+1}, \dots, c_k \mapsto q_k})_{\text{red}}.$$

The transition then stores the actual result in the file-system and produces an output token which contains both a reference to the result and the point q . The stored data will be used later in the normalization step of the interpolation (see above).

Transition degree: This generates a substitution point $p^{(j)} \in \mathbb{Z}^{\{1, \dots, j-1, j+1, \dots, k\}}$ for each $1 \leq j \leq k$ yielding a matrix

$$M^{(j)} := M|_{c_1 \mapsto p_1^{(j)}, \dots, c_{j-1} \mapsto p_{j-1}^{(j)}, c_{j+1} \mapsto p_{j+1}^{(j)}, \dots, c_k \mapsto p_k^{(j)}}$$

over the field $\mathbb{Q}(c_j)$. After applying the row reduction, $M_{\text{red}}^{(j)}$ can be used to determine degree bounds for the numerator and denominator of each entry of the final result M_{red} as a polynomial in c_j .

For $j \leq r$, we need a global degree bound to determine the number of interpolation points. We thus take the maximum of all numerator and denominator degrees of entries of $M_{\text{red}}^{(j)}$, and store these as a vector in $\mathbb{N}_0^{\{1, \dots, r\}}$, which is put on the place d_v .

If $j > r$, two integer matrices will be produced, which store the degrees of the numerators and denominators of each entry of M_{red} , respectively. This information will be used later to filter out bad interpolation points, that is, points at which polynomial cancellation occurs (see appendix A). The result is stored in the file system and a token with a reference to the result is put on the place d_m .

Note that **degree** is in fact modeled by a sub-Petri net which behaves in a hierarchical manner as a transition. In practice, we actually compute multiple matrices $M^{(j)}$ per parameter to reduce the probability of a bad point producing wrong degree bounds.

Transition points: This transition takes the degree data in d_v and initializes the interpolation tree described in section 4.2 and depicted in figure 4. This, in turn, produces the corresponding set of interpolation points, which are put as separate tokens on the place p .

Transition reduce: This transition consumes a point $p' \in \mathbb{Z}^{\{1, \dots, r\}}$ from the place p and computes

$$(M_{c_1 \mapsto p'_1, \dots, c_m \mapsto p'_m})|_{\text{red}}.$$

The resulting matrix together with its interpolation point are put on the place m . Since **reduce** performs parameter substitutions in rational function expressions, the computation may fail due to division by zero. If this happens, $m.\text{valid}$ is set to **false**, otherwise it is set to **true**.

Transition replace failure: An input token for which $m.\text{valid}$ is **false** is consumed by the transition **replace failure**, which marks the respective interpolation point as failed in the interpolation tree. If necessary, the interpolation tree is extended by additional interpolation points, which are also put on the place p .

Transition normalize: An input token for which $m.\text{valid}$ is **true** is consumed by the transition **normalize**. This transition reads M_{ref} and multiplies the input matrix referenced by m with a suitable constant factor. It also compares the entries with the degree matrices in d_m to identify bad interpolation points. The result is put on the place n . If the corresponding point was bad, $n.\text{valid}$ is set to **false**, otherwise to **true**.

Transition replace invalid: For an input token for which `n.valid` is `false`, the transition generates new interpolation points in a fashion similar to that in `replace failure`.

Transition store normalized: For an input token for which `n.valid` is `true`, the transition marks the corresponding interpolation point as successful in the external storage. If enough interpolation points for a given parameter have been marked as successful, the storage produces a token on place i , which triggers the respective interpolation. If the point (p'_1, \dots, p'_r) triggers the interpolation (which will then use further points of the form $(p'_1, \dots, p'_{r-1}, p''_r)$), the result of the interpolation will be associated to the point (p'_1, \dots, p'_{r-1}) in the interpolation tree. If there are not yet enough interpolation points, the transition produces a token which only contains `i.valid` with value `false`.

Transition discard: This transition discards tokens with `i.valid` equal to `false`.

Transition interpolate: Tokens with `i.valid` equal to `true` are consumed by this transition, which then retrieves the references to the input data for the interpolation from the interpolation tree, loads the respective data from the file system, and executes the interpolation. If (in the above notation) the token holds (p'_1, \dots, p'_{r-1}) , then for $(d_v)_r + 1$ many points the corresponding row reduced matrices are retrieved from the storage. Note that due to the tree structure of the interpolation tree, all these points must have the first $r - 1$ coordinates equal to (p'_1, \dots, p'_{r-1}) . The interpolation is then performed entry-wise as explained in appendix A.

Transition store interpolated: This transition marks the current point (p'_1, \dots, p'_{r-1}) in the interpolation tree as processed. If $r > 1$, just like in `store normalized`, the transition produces an interpolation token for the next parameter. If $r = 1$, we have arrived at the final result, and a token with `i.valid` equal to `false` is produced, which will then be discarded.

The Petri net contains additional infrastructure (not described here) which terminates the execution once no tokens exist any more on the places i and p .

4.4 Parallel timings

To illustrate the efficiency of our approach, we consider the cut $\{1, 3, 4, 5\}$ of the double pentagon diagram (see section 6 for a discussion of all possible cuts). Choosing this particular cut, which is less complex than others, our computations finish even when only a small number of cores is involved. This is necessary to analyze the scaling of our algorithm. In table 1, we give timings for different numbers of cores. All timings are in seconds, taken on the high performance compute cluster at the Fraunhofer Institute for Industrial Mathematics (ITWM). Each compute node provides two Xeon E5-2670 processors, which amounts to 16 cores³ running at a base clock speed of 2.6 GHz. Each node has 64 GB of memory. For all runs with more than 15 cores, on each node we ran 15 compute jobs and one job for interfacing with the storage system. Since the storage jobs use negligible computation time, we omit them from the CPU core count when determining speedup and efficiency.

³Hypertexting is disabled.

nodes	cores	runtime	relative	
			speedup	efficiency
1	1	122857.6	1.201	1.201
1	15	9837.8	15.000	1.000
2	30	4954.8	29.7822	0.992
4	60	2625.4	56.2058	0.936
8	120	1341.3	110.014	0.916
14	210	952.3	154.958	0.737
15	225	705.6	209.132	0.929
16	240	694.3	212.514	0.885
29	435	611.8	241.199	0.554
30	450	385.4	382.856	0.850
32	480	379.9	388.336	0.809
40	600	367.7	401.307	0.668
48	720	363.2	406.195	0.564

Table 1. Timings and efficiency for the cut $\{1, 3, 4, 5\}$. We use the same algorithm for all core counts. The single core run serves as a reference.

Apart from the running time $T(n)$ of the algorithm on a total of n cores, we also give the speedup $S(n) = \frac{T(1)}{T(n)}$ and the efficiency $E(n) = \frac{T(1)}{nT(n)}$, which measure how “well” the algorithm parallelizes with increasing core counts. Note that the single-core timing is somewhat special: as experiments have shown, the performance per core decreases with the number of cores used on a given node. This effect has been investigated in [69] (see in particular [69, figure 5]). Thus, for the analysis of the expected run-time below, we rather consider the *relative* speedup and efficiency with respect to the 15-core timing. This in particular makes the assumption that the 15-core speedup is 15.

The saw-tooth shape of the efficiency graph in figure 6 (and the corresponding behavior in the timing and speedup graphs) is due to the fact that the number of reductions to execute is usually not divisible by the number of cores utilized. Since in our test problem approximately 450 reductions are required to enable the final interpolation, the running time of the full algorithm is roughly

$$\left\lceil \frac{450}{\text{number of CPUs}} \right\rceil. \tag{4.2}$$

This effect can be avoided by a more fine-grained structuring of the problem (for instance by interpolating more parameters). Note, however, that increasing the number of processes in this way will lead to more overhead via inter-process communication and disk accesses. Thus, dividing the algorithm into very small parts may in fact slow down the overall computation.

Figure 6 also depicts the ideal expected runtime, speedup and efficiency. These ideal graphs stem from the simple assumption, called *Amdahl’s law*, that an algorithm can be divided up into a part that is ideally parallelizable and a part which is not parallelizable at all. Denoting the parallelizable fraction by f , the expected runtime $T_{\text{ideal}}(n)$ on n cores




Figure 6. Running time, relative speedup and efficiency graphs for cut $\{1, 3, 4, 5\}$.

is not $\frac{T(1)}{n}$, but rather

$$T_{\text{ideal}}(n) = (1 - f) \cdot T(1) + \frac{f \cdot T(1)}{n}, \quad (4.3)$$

which yields the ideal speedup and efficiency

$$S_{\text{ideal}}(n) = \frac{n}{(1 - f) \cdot n + f}, \quad E_{\text{ideal}}(n) = \frac{1}{(1 - f) \cdot n + f}. \quad (4.4)$$

Using the experimental values for 15 and 30 cores, we arrive at a value $f \approx 0.999748$, that is, only 0.025 % of the algorithm is not parallelizable.

As we can see, the ideal curves give a fairly tight bound on the actual timings, at least in the cases where the core count is properly aligned to the number of reductions. This indicates that our approach for parallelization not only provides an automatic and fast solution to a tedious and complicated task, but stays highly efficient even when used with a large amount of computing power.

5 IBP conversion between different integral bases

It is well known that after the IBP reduction, the master-integral coefficient size may vary significantly if we choose different master integral bases.

We prefer the IBP reduction to a uniformly transcendental (UT) basis as introduced in [18], for several reasons: a) The differential equations satisfied by a UT integral basis have a particularly simple form [18] which allows for the integrals to be solved analytically in terms of polylogarithms. There is also evidence that for numerical computations, a UT basis is more convenient to evaluate.⁴ So the IBP reduction to a UT basis greatly simplifies the amplitude computations after the IBP reduction. b) We observe that, in the case of the double pentagon, the master-integral coefficients in a UT basis are significantly simpler than those in a traditional Laporta basis. This makes the IBP relations easier to use.

In practice, we consider special forms of UT bases, the so-called dlog bases, which will be introduced in the next subsection.

5.1 Dlog bases and the dlog algorithm

We call a Feynman integral a dlog integral if its four dimensional integrand can be written as a sum of dlog forms [75] with constant coefficients. For the construction of a dlog basis we additionally demand the integrals to be UT functions. Depending on the integral family under consideration this additional requirement is automatically fulfilled for Feynman integrals up to a certain numerator power which makes the dlog criterium a very strong tool to construct UT integrals. Beyond this numerator power an extended approach as in [76] using the Baikov parametrization can be applied to construct UT integrals. The approach outlined here can be applied to Feynman integrals evaluating to multiple polylogarithms. For an integrand

$$dI = dx_1 \wedge dx_2 \wedge \dots \wedge dx_n R(x_1, \dots, x_n), \quad (5.1)$$

⁴Private communication with Yang Gang.

where $R(x_1, \dots, x_2)$ is a rational function in x_1, \dots, x_n we say dI has a dlog form if it can be written as

$$dI = \sum_i c_i d\log f_{i,1} \wedge \dots \wedge d\log f_{i,n}, \tag{5.2}$$

with $f_{i,j}$ being algebraic functions of x_1, \dots, x_n . This is only possible if the integrand has at most simple poles, including points at infinity. For example, both forms $\frac{dx}{x^2}$ and dx admit no dlog form because of the double poles at zero respectively infinity.

The coefficients c_i in equation (5.2) are called leading singularities [77]. For Feynman integrals, that are not of the elliptic type, they are in general algebraic functions of the external variables. By choosing an appropriate parametrization of the external variables, the leading singularities are typically rational functions. This is, in particular, true for the two-loop five-point integrals that are discussed in the next section. The leading singularities can also be understood as integrals over the original integrand where the integration contour is localized around the poles of the integrand. Leading singularities and the integrals integrated on the real contour have analytic properties in common. So, integrals with leading singularities that are just constant numbers are particularly useful, most importantly because they fulfill differential equations in the canonical form [18]. This implies that they have the property of uniform transcendental weight, which means that if the series is expanded in ϵ , the parameter of dimensional regularization, the coefficients have homogeneous transcendental weight and the weight increases by one for each order in ϵ .

Next, we recall from [64] how to transform a given integrand into dlog form, in case this is possible. Given an integrand in n integration variables, we choose, if possible, one variable x that is linear in all denominator factors and do a partial fraction decomposition while treating all other variables as constants. In this way, we obtain a sum of integrands of the form

$$\frac{dx}{x-a} \wedge \Omega = d\log(x-a) \wedge \Omega, \tag{5.3}$$

where Ω is an $(n-1)$ -form, independent of x , and a is a polynomial that may depend on the other integration variables. Then we iterate this procedure taking Ω as our new integrand until no integration variables are left. If in any intermediate step a pole of degree two or higher is encountered, then the integrand does not admit a dlog form. There are cases where no variable exists that is linear in all denominator factors. One way to proceed in such a case is to make a variable transformation such that at least one of the new variables is linear in all denominator factors.

With the help of an algorithm based on the approach described in this section a complete basis of dlog master integrals with constant leading singularities for all two-loop five point integral families were constructed [76, 78]. The denominator structure for each integral family is given by the propagators. To construct the dlog integrals we make a general numerator ansatz. We write the numerator as a linear combination of terms that are products of inverse propagators and irreducible scalar products. Each term is multiplied by a free parameter, and by applying the algorithm to this general integrand, we can determine values of the free parameters such that the integrand has a dlog form

and constant leading singularities. In this way, we obtain a set of dlog integrals that form a basis of dlog master integrals.

5.2 IBP reduction with a dlog basis

Given a dlog basis, we discuss the IBP reduction in two settings:

1. When both the master-integral coefficients in the Laporta basis and the dlog basis are needed, we first compute the reduction in the Laporta basis I with our module intersection and GPI-SPACE reduction algorithm,

$$F = AI, \tag{5.4}$$

where F is the list of target integrals as a column vector. Then we reduce the dlog basis \tilde{I} to the Laporta basis I ,

$$\tilde{I} = TI. \tag{5.5}$$

Note that since the dlog basis construction has a restriction on the numerator degree, this reduction is usually easy. Terms exceeding the allowed numerator degree have double poles at infinity. This can be seen by inverting the loop momenta $k_i^\mu \rightarrow k_i^\mu/k_i^2$. Using our SINGULAR RREF code, with a good pivot strategy, we can analytically find the inverse T^{-1} . The matrix product AT^{-1} contains the coefficients of an IBP reduction to the dlog basis.

We remark that the product AT^{-1} can be difficult to calculate even if T^{-1} has a relative small size. Instead of computing the product directly, we again use the semi-numerical approach, setting several of the kinematic values to be integers, computing the product several times, and then using our interpolation program to get the fully analytical matrix product AT^{-1} . This is again implemented using our SINGULAR-GPI-SPACE framework.

2. When only the master-integral coefficients in a dlog basis are needed, we apply our semi-numerical reduction method to a set of numeric master-integral coefficients in the Laporta basis. Instead of interpolating these coefficients, we use the semi-numeric points to interpolate the product AT^{-1} , not calculating the analytic form of A .

In the next section, we illustrate our approach by considering a non-trivial example, the two-loop five-point nonplanar double pentagon diagram. This includes the IBP generation via the module intersection method, the massively parallel reduction of the IBP system and the basis conversion.

6 The two-loop five-point nonplanar double pentagon example

In this section, we illustrate our IBP reduction method by applying it to a nontrivial example, the two-loop five-point nonplanar double pentagon. Note that a symbolic UT basis for this example was found in [5, 11]. Furthermore, UT bases in terms of polylogarithm




Figure 7. We depict the two-loop five-point nonplanar double pentagon diagram, writing z_i for the Baikov variables, which are equal to the inverse propagators. In particular, $z_1 = l_1^2$ and $z_4 = l_2^2$. We also draw the 11 spanning cuts of this integral family. These correspond to the non-collapsible master integrals, before using symmetries.

functions for the double pentagon and other two-loop five-point nonplanar massless integral families were analytically calculated in [76].

For the diagram in figure 7, we chose the following labeling for the propagators:

$$\begin{aligned}
 D_1 &= l_1^2 & D_2 &= (l_1 - p_1)^2 & D_3 &= (l_1 - p_{12})^2 & D_4 &= l_2^2 \\
 D_5 &= (l_2 - p_{123})^2 & D_6 &= (l_2 - p_{1234})^2 & D_7 &= (l_1 - l_2)^2 & D_8 &= (l_1 - l_2 + p_3)^2 \\
 D_9 &= (l_1 - p_{1234})^2 & D_{10} &= (l_2 - p_1)^2 & D_{11} &= (l_2 - p_{12})^2, & &
 \end{aligned} \tag{6.1}$$

where the l_i represent the loop momenta, the p_i represent external momenta, and $p_{i\dots j} = \sum_i^j p_i$. The first 8 propagators represent the topology and the last three ones the irreducible scalar products.

This is a complicated integral family for IBP reduction, due to the number of independent scalars, which are s12, s23, s34, s45, s45, s15 and the spacetime dimension D, and

$$\begin{aligned}
 & I[1, 1, 1, 1, 1, 1, 1, 1, 0, 0, -4], \quad I[1, 1, 1, 1, 1, 1, 1, 1, 0, -1, -3], \quad I[1, 1, 1, 1, 1, 1, 1, 1, 0, -2, -2], \\
 & I[1, 1, 1, 1, 1, 1, 1, 1, 0, -3, -1], \quad I[1, 1, 1, 1, 1, 1, 1, 1, 0, -4, 0], \quad I[1, 1, 1, 1, 1, 1, 1, 1, -1, 0, -3], \\
 & I[1, 1, 1, 1, 1, 1, 1, 1, -1, -1, -2], \quad I[1, 1, 1, 1, 1, 1, 1, 1, -1, -2, -1], \quad I[1, 1, 1, 1, 1, 1, 1, 1, -1, -3, 0], \\
 & I[1, 1, 1, 1, 1, 1, 1, 1, -2, 0, -2], \quad I[1, 1, 1, 1, 1, 1, 1, 1, -2, -1, -1], \quad I[1, 1, 1, 1, 1, 1, 1, 1, -2, -2, 0], \\
 & I[1, 1, 1, 1, 1, 1, 1, 1, -3, 0, -1], \quad I[1, 1, 1, 1, 1, 1, 1, 1, -3, -1, 0], \quad I[1, 1, 1, 1, 1, 1, 1, 1, -4, 0, 0], \\
 & I[1, 1, 1, 1, 1, 1, 1, 1, 0, 0, -3], \quad I[1, 1, 1, 1, 1, 1, 1, 1, 0, -1, -2], \quad I[1, 1, 1, 1, 1, 1, 1, 1, 0, -2, -1], \\
 & I[1, 1, 1, 1, 1, 1, 1, 1, 0, -3, 0], \quad I[1, 1, 1, 1, 1, 1, 1, 1, -1, 0, -2], \quad I[1, 1, 1, 1, 1, 1, 1, 1, -1, -1, -1], \\
 & I[1, 1, 1, 1, 1, 1, 1, 1, -1, -2, 0], \quad I[1, 1, 1, 1, 1, 1, 1, 1, -2, 0, -1], \quad I[1, 1, 1, 1, 1, 1, 1, 1, -2, -1, 0], \\
 & I[1, 1, 1, 1, 1, 1, 1, 1, -3, 0, 0], \quad I[1, 1, 1, 1, 1, 1, 1, 1, 0, 0, -2]
 \end{aligned}$$

Figure 8. Integrals up to numerator degree 4 without double propagators for the nonplanar double pentagon diagram.

due to the nonplanar topology with two pentagons inside. We demonstrate our method by reducing the 26 integrals listed in figure 8 to a master integral basis in the fashion of Laporta. Furthermore, we convert the master-integral coefficients to the coefficients of a dlog basis given in [76]. In this base change, we observe a significant coefficient size reduction.

6.1 Module intersection with cuts

First, we use AZURITE [47] to find an integral basis. Without considering symmetries, there are 113 irreducible integrals, and with symmetries, there are 108 master integrals. Note that due to the number of master integrals, this IBP reduction is significantly more complicated than the reduction of the hexagon-box diagram in [60], which has only 73 master integrals.

We can then construct the set of spanning cuts of this integral family. Each spanning cut corresponds to a “non-collapsible” master integral [47]. There are 11 spanning cuts (without considering symmetries),

$$\begin{aligned}
 & \{1, 5, 7\}, \{1, 5, 8\}, \{1, 6, 8\}, \{2, 4, 8\}, \{2, 5, 7\}, \{2, 6, 7\}, \\
 & \{2, 6, 8\}, \{3, 4, 7\}, \{3, 4, 8\}, \{3, 6, 7\}, \{1, 3, 4, 5\},
 \end{aligned} \tag{6.2}$$

where the numbers indicate the propagators on cut. For example, $\{3, 4, 7\}$ means that $D_3 \mapsto 0, D_4 \mapsto 0, D_7 \mapsto 0$.

For each cut, we can apply our module intersection method to generate IBPs without double propagators. In [60], the IBPs are generated only from the top sector. Here, for each cut, we generate the IBPs from both the top sector and lower sectors. For example, for the cut $\{1, 5, 7\}$, we consider the 32 sectors supported on cut $\{1, 5, 7\}$, and compute all the module intersections. This approach will generate more IBPs but make the IBP system more block-triangular and easier for linear reduction.

With the `degbound` option in SINGULAR, it is easy to generate all the module intersections. For this integral family, choosing the degree bound 5, and using one core for each cut, it takes less than 5 minutes in total to solve all the module intersection problems *analytically*. Later on, by finite-field methods, we find that with this choice of degree bound, we obtain sufficiently many IBPs for our problem.

Cut	# relations	# integrals	size	d_2	d_3
{1,5,7}	1134	1182	0.77 MB	21	22
{1,5,8}	1141	1192	0.85 MB	18	18
{1,6,8}	1203	1205	1.1 MB	19	30
{2,4,8}	1245	1247	1.1 MB	35	24
{2,5,7}	1164	1211	0.84 MB	26	18
{2,6,7}	1147	1206	0.62 MB	16	17
{2,6,8}	1126	1177	0.83 MB	16	18
{3,4,7}	1172	1221	0.78 MB	19	18
{3,4,8}	1180	1226	1.0 MB	19	22
{3,6,7}	1115	1165	0.82 MB	21	28
{1,3,4,5}	721	762	0.43MB	14	14

Table 2. The IBP relations generated on each cut by the module intersection method. We used finite-field methods to pick linearly independent and necessary IBP relations to reduce all target integrals. The size is the output file size on disk before reduction. The numbers d_2 and d_3 are the maximal degrees in the reduced IBP relations for c_2 and c_3 , respectively.

After generating the IBPs, we use the two-step trimming process described in section 2.2 to select necessary IBPs for our targets. This computation is via finite-field methods and powered by the package SPASM.

Note that different cuts can support the same master integral. We then speak of a cut overlap. For example, the integral $I[1, 1, 1, 1, 1, 0, 1, 1, 0, 0, 0]$ is supported by both cuts $\{1, 5, 7\}$ and $\{2, 4, 8\}$. The IBP reductions on these two cuts should give the same coefficients. To avoid redundant computations, we can apply the master integral removal method [49], setting $I[1, 1, 1, 1, 1, 0, 1, 1, 0, 0, 0] \mapsto 0$ either in cut $\{1, 5, 7\}$ or in cut $\{2, 4, 8\}$. Clearly there are many different removal strategies for master integral overlapping in different cuts, and different strategies result in different performances. In our computational problem, we find that the cuts $\{1, 6, 8\}$ and $\{2, 4, 8\}$ are relatively “difficult” cuts for IBP reduction. Hence, we set as many master integrals as possible in $\{1, 6, 8\}$ and $\{2, 4, 8\}$ to zero, and later on recover the remaining master integral coefficients from other cuts via cut overlap.

We compute the module intersections analytically. For the purpose of linear reduction, we further set

$$s_{12} \mapsto 1, \quad c_2 \equiv s_{23}/s_{12}, \quad c_3 \equiv s_{34}/s_{12}, \quad c_4 \equiv s_{45}/s_{12}, \quad c_5 \equiv s_{15}/s_{12} \quad (6.3)$$

to dehomogenize the IBP relations and speed up the computation. The s_{12} dependence can be recovered in the final step.

The resulting IBPs are summarized in table 2. Note that for the cut $\{1, 6, 8\}$, there are 1203 independent relations and 1205 integrals after applying the idea of [49] to set most master integrals supported on the cut $\{1, 6, 8\}$ to zero. As a result we only have to compute just two master integral coefficients.

6.2 IBP reduction

We apply our reduction method via SINGULAR and GPI-SPACE to reduce the linear systems in table 2. We use a semi-numeric approach, choosing c_4 , c_5 and the space-time dimension D to be symbolic, and compute the linear reduction with integer-valued c_2 and c_3 .

By a linear reduction with c_2 (respectively c_3) symbolic and all the other parameters numeric, we easily determine the maximal degree of c_2 (respectively c_3) in the reduced IBP relations. The degrees are listed in table 2 as d_2 and d_3 , respectively. From this information, we get the minimal number $(d_2 + 1) \times (d_3 + 1)$ of semi-numeric computations for interpolating the analytic reduction result. For example, for the cut $\{1, 5, 7\}$, we need to run semi-numeric computations at least 506 times. Of course, the cuts exhibit different running times when performing the reductions: for instance, cut $\{1, 3, 4, 5\}$, which we already considered as an example in section 4.4, is the easiest in terms of running time, taking only about 11 minutes when using 384 CPU cores. In contrast, the cut $\{3, 4, 8\}$ is much more complex: its reduction took 12 hours and 21 minutes, using 384 cores.

After getting the analytic reduction of all the cuts, we merge them to get the full IBP reduction to a 113-integral basis. Furthermore, we apply the symmetries, obtained from [47],

$$I[0, 0, 1, 1, 0, 0, 0, 1, 0, 0, 0] = I[1, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0] \quad (6.4)$$

$$I[0, 0, 1, 1, 0, 0, 1, 0, 0, 0, 0] = I[1, 0, 0, 0, 1, 0, 0, 1, 0, 0, 0] \quad (6.5)$$

$$I[1, 0, 1, 0, 1, 0, 1, 0, 0, 0, 0] = I[1, 0, 1, 1, 0, 0, 0, 1, 0, 0, 0] \quad (6.6)$$

$$I[0, 0, 1, 1, 1, 0, 1, 0, 0, 0, 0] = I[1, 0, 0, 1, 1, 0, 0, 1, 0, 0, 0] \quad (6.7)$$

$$I[1, 0, 1, 1, 1, 0, 0, 1, 0, 0, 0] = I[1, 0, 1, 1, 1, 0, 1, 0, 0, 0, 0] \quad (6.8)$$

to reduce the 26 target integrals to a 108-integral Laporta basis I . We note that the resulting file is large, with a size of ~ 2.0 GB on disk.

By setting all Mandelstam variables to integers, we have verified that our result is consistent with FIRE6 [38].

6.3 Master-integral coefficient conversion to a dlog basis

In this subsection, we discuss converting the master-integral coefficients for the Laporta basis to the master-integral coefficients of the dlog basis found in [76].

For this conversion, we again use the semi-numeric approach, taking integer-valued c_2 , c_3 , and symbolic c_4 , c_5 and D , converting the coefficients and then interpolating. It is easy to determine that the coefficients in the dlog basis have the following maximal degrees for c_2 and c_3 , respectively,

$$d'_2 = 20, \quad d'_3 = 20. \quad (6.9)$$

By comparing with table 2, where d_2 can be as high as 35, we find that the maximal degree drops. For the basis conversion, we carry out a semi-numeric matrix multiplication with subsequent interpolation using SINGULAR and GPI-SPACE.

After the computation, we see that the IBP reduction coefficients of figure 8 in this dlog basis have size 480 MB on disk, which shows a significant 76% reduction of the

master-integral coefficient size compared to what we have for the Laporta basis. On the other hand, if only the master-integral coefficients in the dlog basis are needed, we can skip the interpolation for the Laporta basis master-integral coefficients, and directly convert the intermediate numerical results to dlog basis master-integral coefficients. Because of the maximal degree drop, this shortcut reduces the required number of semi-numeric computations.

For convenience, we also provide the master-integral coefficients in the dlog basis, with the s_{12} scalar recovered. All these analytic results can be obtained via the links presented in the introduction of this paper. Note that all files provided under the links contain 26×108 matrices. For each matrix, the entry in the i th row and j th column is the corresponding master-integral coefficient for the i th target integral in figure 8, expanded on the j th master integral. The Laporta basis and the dlog basis are included in the auxiliary files of this paper.

7 Summary

In this paper, we present our powerful new IBP reduction method, which is based on computational algebraic geometry powered by the computer algebra system SINGULAR in conjunction with the taskflow management system GPI-SPACE. Our method is suitable for large scale IBP reduction problems with complicated Feynman diagrams and multiple variables. We demonstrate the power of the new method by the analytic two-loop five-point nonplanar double pentagon IBP computation. The computational result has been cross-checked numerically using state-of-the-art IBP programs.

Our method is flexible and can be adapted in various different scenarios:

1. Modern methods for amplitude computation often follow the approach of numerically or semi-numerically calculating the IBP relations in order to interpolate the amplitude coefficient under consideration directly, instead of interpolating the analytic IBP relations. Our method can efficiently compute the reduced numeric or semi-numeric IBP relations and, hence, perfectly fits into this purpose.
2. Our module intersection method can also be used for integrals with double propagators or multiple-power propagators since this IBP generating method avoids the increase of propagator exponents and significantly reduces the size of the IBP system.
3. Although our method is currently based on semi-numerical parallelizations with integer-valued numerics, it clearly can be extended to finite-field linear reduction, if necessary.
4. More generally, our linear reduction parallelization method can be used for computational problems other than IBP reduction. For example, in recent years, it was found that the Bethe Ansatz equation of integrable spin chains can be analytically computed by algebraic geometry methods [79, 80]. Often, this involves large-scale linear algebra computations with symbolic parameters, and our parallelization via the SINGULAR-GPI-SPACE framework can greatly speed up the computation. We

also expect that our reduction method can be used more generally for Gröbner basis computations with parameters.

In the future, we will develop our code into an automated software package, powered by SINGULAR and GPI-SPACE, for solving large-scale IBP or amplitude problems. The possible simplification of master integral coefficients in a UT/dlog basis will be further investigated. We expect that this method will provide a boost for the current NNLO precision computations.

Acknowledgments

YZ thanks Johann Usovitsch for many enlightening discussions on IBP reduction algorithms and their implementation, and Johannes Henn for cutting-edge techniques for finding uniformly transcendental and dlog integrals. We acknowledge Kasper Larsen for his work in the starting stage of the analysis of the two-loop nonplanar five-point double pentagon diagram. We thank Simon Badger, Rutger Boels, Christian Bogner, Hjalte Frellesvig, Gudrun Heinrich, Harald Ita, Stephan Jahn, David Kosower, Roman Lee, Hui Luo, Yanqing Ma, Andreas von Manteuffel, Pierpaolo Mastrolia, Alexander Mitov, Erik Panzer, Tiziano Peraro, Robert Schabinger, Hans Schönemann, Alexander Smirnov, Vladimir Smirnov, Peter Uwer, Gang Yang and Mao Zeng for their help on this research project.

The research leading to these results has received funding from the Swiss National Science Foundation (Ambizione grant PZ00P2 161341), from the National Science Foundation under Grant No. NSF PHY17-48958, from the NSF of China with Grant No. 11235010, from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 725110). The work of AG is supported by the Knut and Alice Wallenberg Foundation under grant #2015-0083. The work of DB, JB, and WD was supported by Project II.5 of SFB-TRR 195 *Symbolic Tools in Mathematics and their Application* of the German Research Foundation (DFG). The authors would also like to express a special thanks to the Mainz Institute for Theoretical Physics (MITP) of the Cluster of Excellence PRISMA+ (Project ID 39083149) for its hospitality and support during the “Mathematics of Linear Relations between Feynman Integrals” workshop.

A Rational function interpolation

In this appendix, we introduce our simple approach to rational function interpolation. Although this algorithm is rather straight-forward compared to other more involved techniques available (see, for example, [81–83]), we have found that it is more suitable for our setup. The idea is to heuristically convert a rational function interpolation problem to a polynomial function interpolation problem.

We focus on a general computational process with symbolic variables x_1, \dots, x_k which would give the final result as a rational function

$$\frac{F(x_1, \dots, x_k)}{G(x_1, \dots, x_k)}. \quad (\text{A.1})$$

Here, F and G are integer valued polynomials with $\gcd(F, G) = 1$. Suppose we are in a situation where it is difficult to perform this process with all parameters x_1, \dots, x_k symbolic, but where it is feasible to obtain results when some of the parameters are set to be general integer values, and only the remaining parameters are symbolic. Then our approach is to repeat such semi-numeric computations sufficiently often, and interpolate to get (A.1).

To be specific, in such a situation, and for a fixed $k_1 < k$, we refer to the computation with x_{k_1+1}, \dots, x_k symbolic, and the other parameters substituted by random integers,

$$x_1 \mapsto a_1^{(i)}, \quad \dots, \quad x_{k_1+1} \mapsto a_{k_1+1}^{(i)}, \quad (\text{A.2})$$

as the i th semi-numeric computation. We write the result of the i th semi-numeric computation as

$$\frac{f_i(x_{k_1+1}, \dots, x_k)}{g_i(x_{k_1+1}, \dots, x_k)}. \quad (\text{A.3})$$

Note that, although computer algebra software can cancel the fraction to get polynomials f_i and g_i with $\gcd(f_i, g_i) = 1$, the relation between G and g_i is not clear a priori. For example, it may happen that

$$g_i(x_{k_1+1}, \dots, x_k) \neq G(a_1^{(i)}, \dots, a_k^{(i)}, x_{k_1+1}, \dots, x_k). \quad (\text{A.4})$$

Similarly, we may have that

$$f_i(x_{k_1+1}, \dots, x_k) \neq F(a_1^{(i)}, \dots, a_k^{(i)}, x_{k_1+1}, \dots, x_k). \quad (\text{A.5})$$

The reason for this is that after taking integer values $a_1^{(i)}, \dots, a_k^{(i)}$ for the first k variables, there can be additional cancellations between F and G . This phenomenon makes the direct polynomial interpolation of the g_i and f_i inapplicable.

We solve this cancellation problem in a heuristic way.

1. First, we compute a “reference” result with symbolic x_1, \dots, x_k and random integer values for the other parameters, $x_{k_1+1} \mapsto b_{k_1+1}, \dots, x_k \mapsto b_k$:

$$\frac{p(x_1, \dots, x_{k_1})}{q(x_1, \dots, x_{k_1})}, \quad (\text{A.6})$$

where p and q are integer valued polynomials with $\gcd(p, q) = 1$. Generally, except for a statistically very small set of points $b = (b_{k_1+1}, \dots, b_k)$, we can assume that the two polynomials $F(x_1, \dots, x_{k_1}, b_{k_1+1}, \dots, b_k)$ and $G(x_1, \dots, x_{k_1}, b_{k_1+1}, \dots, b_k)$ are coprime. Then, since

$$\frac{p(x_1, \dots, x_{k_1})}{q(x_1, \dots, x_{k_1})} = \frac{F(x_1, \dots, x_{k_1}, b_{k_1+1}, \dots, b_k)}{G(x_1, \dots, x_{k_1}, b_{k_1+1}, \dots, b_k)} \quad (\text{A.7})$$

by the unique factorization domain (UFD) property,

$$G(x_1, \dots, x_{k_1}, b_{k_1+1}, \dots, b_k) = \mu \cdot q(x_1, \dots, x_{k_1}), \quad (\text{A.8})$$

$$F(x_1, \dots, x_{k_1}, b_{k_1+1}, \dots, b_k) = \mu \cdot p(x_1, \dots, x_{k_1}), \quad (\text{A.9})$$

for some integer μ . Note that our algorithm relies only on the existence of such a μ .

2. Second, we do a “majority vote” selection from among the semi-numeric results: for each of the polynomials, we record the leading exponent in x_{k+1}, \dots, x_k , and determine the most frequently occurring exponent \mathbf{r} . Then we drop all semi-numeric result whose exponent is not equal to \mathbf{r} . This step ensures that the gcd of $F(a_1^{(i)}, \dots, a_k^{(i)}, c_{k_1+1}, \dots, c_k)$ and $G(a_1^{(i)}, \dots, a_k^{(i)}, x_{k_1+1}, \dots, x_k)$ is just an integer instead of a non-constant polynomial in x_{k_1+1}, \dots, x_k . Again by the UFD property,

$$G(a_1^{(i)}, \dots, a_k^{(i)}, x_{k_1+1}, \dots, x_k) = \lambda_i \cdot g_i(x_{k_1+1}, \dots, x_k), \quad (\text{A.10})$$

$$F(a_1^{(i)}, \dots, a_k^{(i)}, x_{k_1+1}, \dots, x_k) = \lambda_i \cdot f_i(x_{k_1+1}, \dots, x_k), \quad (\text{A.11})$$

where each λ_i is an integer. Setting $x_1 \mapsto a_1^{(i)}, \dots, x_{k_1+1} \mapsto a_{k_1+1}^{(i)}$ in (A.8) and $x_{k_1+1} \mapsto b_{k_1+1}, \dots, x_k \mapsto b_k$ in (A.10), we determine that

$$\lambda_i = \mu \frac{q(a_1^{(i)}, \dots, a_{k_1}^{(i)})}{g_i(b_{k_1+1}, \dots, b_k)}. \quad (\text{A.12})$$

3. Define

$$\tilde{g}_i(x_{k_1+1}, \dots, x_k) \equiv \frac{\lambda_i}{\mu} g_i(x_{k_1+1}, \dots, x_k) = \frac{q(a_1^{(i)}, \dots, a_{k_1}^{(i)})}{g_i(b_{k_1+1}, \dots, b_k)} g_i(x_{k_1+1}, \dots, x_k), \quad (\text{A.13})$$

$$\tilde{f}_i(x_{k_1+1}, \dots, x_k) \equiv \frac{\lambda_i}{\mu} f_i(x_{k_1+1}, \dots, x_k) = \frac{q(a_1^{(i)}, \dots, a_{k_1}^{(i)})}{g_i(b_{k_1+1}, \dots, b_k)} f_i(x_{k_1+1}, \dots, x_k). \quad (\text{A.14})$$

Then interpolate both polynomials $\tilde{f}_i(x_{k_1+1}, \dots, x_k)$ and $\tilde{g}_i(x_{k_1+1}, \dots, x_k)$ for the selected semi-numeric points by standard polynomial interpolation algorithms, say by using Newton polynomials. From (A.12), we see that the resulting two polynomials are

$$\frac{1}{\mu} F(x_1, \dots, x_k, x_{k_1+1}, \dots, x_k), \quad \frac{1}{\mu} G(x_1, \dots, x_k, x_{k_1+1}, \dots, x_k), \quad (\text{A.15})$$

so that the ratio of the polynomials gives the desired fraction F/G , while the integer factor μ is canceled out.

So with the compensation factors λ_i/μ , we can use simple polynomial interpolation algorithms to interpolate rational functions at the extra cost of computing only one additional reference point. If the degree of F and G in x_j is d_j ($1 \leq j \leq k_1$), for a general choice of the interpolation points, we need to compute the semi-numeric result

$$(d_1 + 1) \times \dots \times (d_{k_1} + 1) \quad (\text{A.16})$$

times, plus one computation for the reference point.

The resulting algorithm is implemented using SINGULAR in conjunction with GPI-SPACE.

In practice, this algorithm is extendable in many ways. For example, instead of splitting the variables x_1, \dots, x_k into two groups (x_1, \dots, x_{k_1}) and (x_{k_1+1}, \dots, x_k) , we can split the variables in more groups and use our algorithm recursively. The algorithm can also be combined with finite field reconstruction.

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] S. Badger, H. Frellesvig and Y. Zhang, *A two-loop five-gluon helicity amplitude in QCD*, *JHEP* **12** (2013) 045 [[arXiv:1310.1051](https://arxiv.org/abs/1310.1051)] [[INSPIRE](#)].
- [2] T. Gehrmann, J.M. Henn and N.A. Lo Presti, *Analytic form of the two-loop planar five-gluon all-plus-helicity amplitude in QCD*, *Phys. Rev. Lett.* **116** (2016) 062001 [Erratum *ibid.* **116** (2016) 189903] [[arXiv:1511.05409](https://arxiv.org/abs/1511.05409)] [[INSPIRE](#)].
- [3] S. Badger, C. Brønnum-Hansen, H.B. Hartanto and T. Peraro, *First look at two-loop five-gluon scattering in QCD*, *Phys. Rev. Lett.* **120** (2018) 092001 [[arXiv:1712.02229](https://arxiv.org/abs/1712.02229)] [[INSPIRE](#)].
- [4] S. Abreu et al., *Planar two-loop five-gluon amplitudes from numerical unitarity*, *Phys. Rev. D* **97** (2018) 116014 [[arXiv:1712.03946](https://arxiv.org/abs/1712.03946)] [[INSPIRE](#)].
- [5] S. Abreu et al., *The two-loop five-point amplitude in $\mathcal{N} = 4$ super-Yang-Mills theory*, *Phys. Rev. Lett.* **122** (2019) 121603 [[arXiv:1812.08941](https://arxiv.org/abs/1812.08941)] [[INSPIRE](#)].
- [6] S. Abreu et al., *Planar two-loop five-parton amplitudes from numerical unitarity*, *JHEP* **11** (2018) 116 [[arXiv:1809.09067](https://arxiv.org/abs/1809.09067)] [[INSPIRE](#)].
- [7] R.H. Boels, Q. Jin and H. Lüo, *Efficient integrand reduction for particles with spin*, [arXiv:1802.06761](https://arxiv.org/abs/1802.06761) [[INSPIRE](#)].
- [8] T. Gehrmann, J.M. Henn and N.A. Lo Presti, *Pentagon functions for massless planar scattering amplitudes*, *JHEP* **10** (2018) 103 [[arXiv:1807.09812](https://arxiv.org/abs/1807.09812)] [[INSPIRE](#)].
- [9] 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](https://arxiv.org/abs/1811.11699)] [[INSPIRE](#)].
- [10] S. Abreu et al., *Analytic form of planar two-loop five-gluon scattering amplitudes in QCD*, *Phys. Rev. Lett.* **122** (2019) 082002 [[arXiv:1812.04586](https://arxiv.org/abs/1812.04586)] [[INSPIRE](#)].
- [11] D. Chicherin et al., *Analytic result for a two-loop five-particle amplitude*, *Phys. Rev. Lett.* **122** (2019) 121602 [[arXiv:1812.11057](https://arxiv.org/abs/1812.11057)] [[INSPIRE](#)].
- [12] D. Chicherin et al., *The two-loop five-particle amplitude in $\mathcal{N} = 8$ supergravity*, *JHEP* **03** (2019) 115 [[arXiv:1901.05932](https://arxiv.org/abs/1901.05932)] [[INSPIRE](#)].
- [13] S. Abreu et al., *The two-loop five-point amplitude in $\mathcal{N} = 8$ supergravity*, *JHEP* **03** (2019) 123 [[arXiv:1901.08563](https://arxiv.org/abs/1901.08563)] [[INSPIRE](#)].
- [14] S. Abreu et al., *Analytic Form of the Planar Two-Loop Five-Parton Scattering Amplitudes in QCD*, *JHEP* **05** (2019) 084 [[arXiv:1904.00945](https://arxiv.org/abs/1904.00945)] [[INSPIRE](#)].
- [15] H.B. Hartanto et al., *A numerical evaluation of planar two-loop helicity amplitudes for a W -boson plus four partons*, *JHEP* **09** (2019) 119 [[arXiv:1906.11862](https://arxiv.org/abs/1906.11862)] [[INSPIRE](#)].
- [16] Y. Zhang, *Integrand-level reduction of loop amplitudes by computational algebraic geometry methods*, *JHEP* **09** (2012) 042 [[arXiv:1205.5707](https://arxiv.org/abs/1205.5707)] [[INSPIRE](#)].

- [17] P. Mastrolia, E. Mirabella, G. Ossola and T. Peraro, *Scattering amplitudes from multivariate polynomial division*, *Phys. Lett. B* **718** (2012) 173 [[arXiv:1205.7087](#)] [[INSPIRE](#)].
- [18] J.M. Henn, *Multiloop integrals in dimensional regularization made simple*, *Phys. Rev. Lett.* **110** (2013) 251601 [[arXiv:1304.1806](#)] [[INSPIRE](#)].
- [19] J.M. Henn, *Lectures on differential equations for Feynman integrals*, *J. Phys. A* **48** (2015) 153001 [[arXiv:1412.2296](#)] [[INSPIRE](#)].
- [20] H. Ita, *Two-loop integrand decomposition into master integrals and surface terms*, *Phys. Rev. D* **94** (2016) 116015 [[arXiv:1510.05626](#)] [[INSPIRE](#)].
- [21] S. Abreu et al., *Two-loop four-gluon amplitudes from numerical unitarity*, *Phys. Rev. Lett.* **119** (2017) 142001 [[arXiv:1703.05273](#)] [[INSPIRE](#)].
- [22] L.J. Dixon, J.M. Drummond and J.M. Henn, *Bootstrapping the three-loop hexagon*, *JHEP* **11** (2011) 023 [[arXiv:1108.4461](#)] [[INSPIRE](#)].
- [23] L.J. Dixon, J.M. Drummond, M. von Hippel and J. Pennington, *Hexagon functions and the three-loop remainder function*, *JHEP* **12** (2013) 049 [[arXiv:1308.2276](#)] [[INSPIRE](#)].
- [24] L.J. Dixon and M. von Hippel, *Bootstrapping an NMHV amplitude through three loops*, *JHEP* **10** (2014) 065 [[arXiv:1408.1505](#)] [[INSPIRE](#)].
- [25] S. Caron-Huot, L.J. Dixon, A. McLeod and M. von Hippel, *Bootstrapping a Five-Loop Amplitude Using Steinmann Relations*, *Phys. Rev. Lett.* **117** (2016) 241601 [[arXiv:1609.00669](#)] [[INSPIRE](#)].
- [26] L.J. Dixon, M. von Hippel and A.J. McLeod, *The four-loop six-gluon NMHV ratio function*, *JHEP* **01** (2016) 053 [[arXiv:1509.08127](#)] [[INSPIRE](#)].
- [27] L.J. Dixon et al., *Heptagons from the Steinmann cluster bootstrap*, *JHEP* **02** (2017) 137 [[arXiv:1612.08976](#)] [[INSPIRE](#)].
- [28] D. Chicherin, J. Henn and V. Mitev, *Bootstrapping pentagon functions*, *JHEP* **05** (2018) 164 [[arXiv:1712.09610](#)] [[INSPIRE](#)].
- [29] S. Caron-Huot et al., *Six-Gluon amplitudes in planar $\mathcal{N} = 4$ super-Yang-Mills theory at six and seven loops*, *JHEP* **08** (2019) 016 [[arXiv:1903.10890](#)] [[INSPIRE](#)].
- [30] 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](#)].
- [31] T. Peraro, *Scattering amplitudes over finite fields and multivariate functional reconstruction*, *JHEP* **12** (2016) 030 [[arXiv:1608.01902](#)] [[INSPIRE](#)].
- [32] J. Klappert and F. Lange, *Reconstructing rational functions with FireFly*, *Comput. Phys. Commun.* **247** (2020) 106951 [[arXiv:1904.00009](#)] [[INSPIRE](#)].
- [33] T. Peraro, *FiniteFlow: multivariate functional reconstruction using finite fields and dataflow graphs*, *JHEP* **07** (2019) 031 [[arXiv:1905.08019](#)] [[INSPIRE](#)].
- [34] K. Chetyrkin and F. Tkachov, *Integration by parts: the algorithm to calculate β -functions in 4 loops*, *Nucl. Phys. B* **192** (1981) 159.
- [35] S. Laporta, *High precision calculation of multiloop Feynman integrals by difference equations*, *Int. J. Mod. Phys. A* **15** (2000) 5087 [[hep-ph/0102033](#)] [[INSPIRE](#)].
- [36] A.V. Smirnov, *Algorithm FIRE — Feynman Integral REduction*, *JHEP* **10** (2008) 107 [[arXiv:0807.3243](#)] [[INSPIRE](#)].

- [37] A.V. Smirnov, *FIRE5: a C++ implementation of Feynman Integral REduction*, *Comput. Phys. Commun.* **189** (2015) 182 [[arXiv:1408.2372](#)] [[INSPIRE](#)].
- [38] A.V. Smirnov and F.S. Chuharev, *FIRE6: Feynman Integral REduction with modular arithmetic*, [arXiv:1901.07808](#) [[INSPIRE](#)].
- [39] P. Maierhöfer, J. Usovitsch and P. Uwer, *Kira — A Feynman integral reduction program*, *Comput. Phys. Commun.* **230** (2018) 99 [[arXiv:1705.05610](#)] [[INSPIRE](#)].
- [40] P. Maierhöfer and J. Usovitsch, *Kira 1.2 release notes*, [arXiv:1812.01491](#) [[INSPIRE](#)].
- [41] A. von Manteuffel and C. Studerus, *Reduze 2 — Distributed Feynman integral reduction*, [arXiv:1201.4330](#) [[INSPIRE](#)].
- [42] J. Gluza, K. Kajda and D.A. Kosower, *Towards a basis for planar two-loop integrals*, *Phys. Rev. D* **83** (2011) 045012 [[arXiv:1009.0472](#)] [[INSPIRE](#)].
- [43] R.M. Schabinger, *A new algorithm for the generation of unitarity-compatible integration by parts relations*, *JHEP* **01** (2012) 077 [[arXiv:1111.4220](#)] [[INSPIRE](#)].
- [44] K.J. Larsen and Y. Zhang, *Integration-by-parts reductions from unitarity cuts and algebraic geometry*, *Phys. Rev. D* **93** (2016) 041701 [[arXiv:1511.01071](#)] [[INSPIRE](#)].
- [45] Z. Bern, M. Enciso, H. Ita and M. Zeng, *Dual conformal symmetry, integration-by-parts reduction, differential equations and the nonplanar sector*, *Phys. Rev. D* **96** (2017) 096017 [[arXiv:1709.06055](#)] [[INSPIRE](#)].
- [46] R.N. Lee and A.A. Pomeransky, *Critical points and number of master integrals*, *JHEP* **11** (2013) 165 [[arXiv:1308.6676](#)] [[INSPIRE](#)].
- [47] A. Georgoudis, K.J. Larsen and Y. Zhang, *Azurite: an algebraic geometry based package for finding bases of loop integrals*, *Comput. Phys. Commun.* **221** (2017) 203 [[arXiv:1612.04252](#)] [[INSPIRE](#)].
- [48] T. Bitoun, C. Bogner, R.P. Klausen and E. Panzer, *Feynman integral relations from parametric annihilators*, *Lett. Math. Phys.* **109** (2019) 497 [[arXiv:1712.09215](#)] [[INSPIRE](#)].
- [49] 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](#)].
- [50] 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](#)].
- [51] D.A. Kosower, *Direct solution of integration-by-parts systems*, *Phys. Rev. D* **98** (2018) 025008 [[arXiv:1804.00131](#)] [[INSPIRE](#)].
- [52] P. Mastrolia and S. Mizera, *Feynman integrals and intersection theory*, *JHEP* **02** (2019) 139 [[arXiv:1810.03818](#)] [[INSPIRE](#)].
- [53] H. Frellesvig et al., *Decomposition of Feynman integrals on the maximal cut by intersection numbers*, *JHEP* **05** (2019) 153 [[arXiv:1901.11510](#)] [[INSPIRE](#)].
- [54] H. Frellesvig et al., *Vector space of Feynman integrals and multivariate intersection numbers*, *Phys. Rev. Lett.* **123** (2019) 201602 [[arXiv:1907.02000](#)] [[INSPIRE](#)].
- [55] X. Liu and Y.-Q. Ma, *Determining arbitrary Feynman integrals by vacuum integrals*, *Phys. Rev. D* **99** (2019) 071501 [[arXiv:1801.10523](#)] [[INSPIRE](#)].

- [56] X. Liu, Y.-Q. Ma and C.-Y. Wang, *A systematic and efficient method to compute multi-loop master integrals*, *Phys. Lett. B* **779** (2018) 353 [[arXiv:1711.09572](#)] [[INSPIRE](#)].
- [57] Y. Wang, Z. Li and N. Ul Basat, *Direct reduction of amplitude*, [arXiv:1901.09390](#) [[INSPIRE](#)].
- [58] A. Kardos, *A new reduction strategy for special negative sectors of planar two-loop integrals without Laporta algorithm*, [arXiv:1812.05622](#) [[INSPIRE](#)].
- [59] Y. Zhang, *Lecture notes on multi-loop integral reduction and applied algebraic geometry*, [arXiv:1612.02249](#) [[INSPIRE](#)].
- [60] J. Böhm et al., *Complete integration-by-parts reductions of the non-planar hexagon-box via module intersections*, *JHEP* **09** (2018) 024 [[arXiv:1805.01873](#)] [[INSPIRE](#)].
- [61] J. Böhm et al., *Complete sets of logarithmic vector fields for integration-by-parts identities of Feynman integrals*, *Phys. Rev. D* **98** (2018) 025023 [[arXiv:1712.09737](#)] [[INSPIRE](#)].
- [62] W. Decker, G.-M. Greuel, G. Pfister and H. Schönemann, *Singular 4-1-1 — A computer algebra system for polynomial computations*, <http://www.singular.uni-kl.de> (2018).
- [63] F.J. Pfreundt and M. Rahn, *GPI-Space*, Fraunhofer ITWM Kaiserslautern, <http://www.gpi-space.de/> (2018).
- [64] P. Wasser, *Analytic properties of Feynman integrals for scattering amplitudes*, M.Sc. thesis, Johannes Gutenberg-Universität Mainz, Mainz, Germany (2016).
- [65] P.A. Baikov, *Explicit solutions of the three loop vacuum integral recurrence relations*, *Phys. Lett. B* **385** (1996) 404 [[hep-ph/9603267](#)] [[INSPIRE](#)].
- [66] R.N. Lee, *Modern techniques of multiloop calculations*, in the proceedings of the 49th *Rencontres de Moriond on QCD and High Energy Interactions*, March 22–29, La Thuile, Italy (2014), [arXiv:1405.5616](#) [[INSPIRE](#)].
- [67] The_SpaSM_group, *SpaSM: a Sparse direct Solver Modulo p*, v1.2 ed., <http://github.com/cbouilla/spasm> (2017).
- [68] E.W. Mayr and A.R. Meyer, *The complexity of the word problems for commutative semigroups and polynomial ideals*, *Adv. Math.* **46** (1982) 305.
- [69] J. Böhm et al., *Towards massively parallel computations in algebraic geometry*, [arXiv:1808.09727](#).
- [70] C. Jordan, M. Joswig and L. Kastner, *Parallel enumeration of triangulations*, *Electron. J. Combin.* **25** (2018) 27.
- [71] M. Rahn, *GPI-Space whitepaper*, Fraunhofer ITWM Kaiserslautern, <http://gpi-space.com/wp-content/uploads/2014/06/GPISpaceWhitepaper.pdf> (2014).
- [72] L. Ristau, *Using Petri nets to parallelize algebraic algorithms*, Ph.D. Thesis, TU Kaiserslautern, Kaiserslautern, Germany (2019).
- [73] C. Reinbold, *Computation of the GIT-fan using a massively parallel implementation*, Master’s Thesis (2018).
- [74] D. Bendle, *Massively parallel computation of tropical varieties*, Bachelor’s Thesis (2018).
- [75] Z. Bern et al., *Logarithmic singularities and maximally supersymmetric amplitudes*, *JHEP* **06** (2015) 202 [[arXiv:1412.8584](#)] [[INSPIRE](#)].
- [76] D. Chicherin et al., *All master integrals for three-jet production at next-to-next-to-leading order*, *Phys. Rev. Lett.* **123** (2019) 041603 [[arXiv:1812.11160](#)] [[INSPIRE](#)].

- [77] N. Arkani-Hamed, J.L. Bourjaily, F. Cachazo and J. Trnka, *Local integrals for planar scattering amplitudes*, *JHEP* **06** (2012) 125 [[arXiv:1012.6032](#)] [[INSPIRE](#)].
- [78] D. Chicherin et al., *Analytic result for the nonplanar hexa-box integrals*, *JHEP* **03** (2019) 042 [[arXiv:1809.06240](#)] [[INSPIRE](#)].
- [79] J. Lykke Jacobsen, Y. Jiang and Y. Zhang, *Torus partition function of the six-vertex model from algebraic geometry*, *JHEP* **03** (2019) 152 [[arXiv:1812.00447](#)] [[INSPIRE](#)].
- [80] Y. Jiang and Y. Zhang, *Algebraic geometry and Bethe ansatz. Part I. The quotient ring for BAE*, *JHEP* **03** (2018) 087 [[arXiv:1710.04693](#)] [[INSPIRE](#)].
- [81] R. Zippel, *Probabilistic algorithms for sparse polynomials*, in *International Symposium on Symbolic and Algebraic Manipulation*, E.W. Ng ed., Lecture Notes in Computer Science volume 72, Springer, Germany (1979).
- [82] M. Ben Or and P. Tiwari, *A deterministic algorithm for sparse multivariate polynomial interpolation*, in the proceedings of the 20th annual ACM symposium on Theory of computing (STOC'88), Chicago, U.S.A. (1988).
- [83] E. Kaltofen, W.S. Lee and A. Lobo, *Early Termination in Ben-Or/Tiwari Sparse Interpolation and a Hybrid of Zippel's Algorithm**, in the proceedings of the *International Symposium on Symbolic and Algebraic Computation (ISSAC01)*, London, Canada (2001).