

doi 10.26089/NumMet.v26r429

## Two decades of algorithmic Feynman integral reduction

**Alexander V. Smirnov**

Research Computing Center, Moscow State University, Moscow, Russia  
Moscow Center for Fundamental and Applied Mathematics,  
Moscow State University, Moscow, Russia  
ORCID: 0000-0002-7779-6735, e-mail: [asmirnov@srcc.msu.ru](mailto:asmirnov@srcc.msu.ru)

**Vladimir A. Smirnov**

Skobeltsyn Institute of Nuclear Physics of Moscow State University, Moscow, Russia  
Moscow Center for Fundamental and Applied Mathematics,  
Moscow State University, Moscow, Russia  
ORCID: 0000-0001-5573-6267, e-mail: [smirnov@theory.sinp.msu.ru](mailto:smirnov@theory.sinp.msu.ru)

**Abstract:** We present a historiographical review of algorithms and computer codes developed for solving integration-by-parts relations for Feynman integrals. This procedure is one of the key steps in the evaluation of Feynman integrals, since it enables to express integrals belonging to a given family as linear combinations of master integrals. In this review, we restrict ourselves to considering general algorithms which can, in principle, be applied to any family of Feynman integrals.

**Keywords:** Feynman integrals, dimensional regularization, integration by parts, master integrals

**Acknowledgements:** This work was supported by the Moscow Center for Fundamental and Applied Mathematics of Lomonosov Moscow State University under Agreement No. 075–15–2025–345.

**For citation:** A. V. Smirnov, V. A. Smirnov, “Two decades of algorithmic Feynman integral reduction,” Numerical Methods and Programming. 26 (4), 438–448 (2025). doi 10.26089/NumMet.v26r429.

## Два десятилетия алгоритмической редукции фейнмановских интегралов

**А. В. Смирнов**

Московский государственный университет имени М. В. Ломоносова,  
Научно-исследовательский вычислительный центр,  
Москва, Российская Федерация  
Московский государственный университет имени М. В. Ломоносова,  
Московский Центр фундаментальной и прикладной математики,  
Москва, Российская Федерация  
ORCID: 0000-0002-7779-6735, e-mail: [asmirnov@srcc.msu.ru](mailto:asmirnov@srcc.msu.ru)

**В. А. Смирнов**

Московский государственный университет имени М. В. Ломоносова,  
Научно-исследовательский институт ядерной физики имени Д. В. Скобельцына,  
Москва, Российская Федерация  
Московский государственный университет имени М. В. Ломоносова,  
Московский Центр фундаментальной и прикладной математики,  
Москва, Российская Федерация  
ORCID: 0000-0001-5573-6267, e-mail: [smirnov@theory.sinp.msu.ru](mailto:smirnov@theory.sinp.msu.ru)

**Аннотация:** Статья представляет собой историографический обзор алгоритмов и компьютерных программ, разработанных для решения соотношений интегрирования по частям для фей-



фeyнмановских интегралов. Соответствующая процедура решения является одним из ключевых этапов при вычислении этих интегралов, поскольку позволяет выражать интегралы, принадлежащие заданному семейству, в виде линейных комбинаций мастер-интегралов. В данном обзоре рассматриваются универсальные алгоритмы, которые, в принципе, могут быть применены к любому семейству фейнмановских интегралов.

**Ключевые слова:** фейнмановские интегралы, размерная регуляризация, интегрирование по частям, мастер-интегралы.

**Благодарности:** Работа выполнена при поддержке Московского центра фундаментальной и прикладной математики по договору № 075–15–2025–345.

**Для цитирования:** Смирнов А.В., Смирнов В.А. Два десятилетия алгоритмической редукции фейнмановских интегралов // Вычислительные методы и программирование. 2025. **26**, № 4. 438–448. doi 10.26089/NumMet.v26r429.

**1. Introduction.** When performing evaluations in perturbative quantum field theory, one needs to compute dimensionally regularized Feynman integrals from a family of scalar integrals associated with a given diagram. The integrals have the following form

$$I_{a_1, \dots, a_n} = \int \cdots \int \frac{1}{P_1^{a_1} \cdots P_n^{a_n}} \prod_{i=1}^h d^d k_i, \quad (1)$$

where  $d = 4 - 2\varepsilon$  is the parameter of dimensional regularization,  $P_i$  are denominators of propagators which are linear combinations of Lorentz scalar products  $k_i \cdot k_j$  and  $k_i \cdot p_l$ , where  $p_l$  are external momenta and  $k_i$  are  $h$  loop momenta. The  $n$  integer numbers  $a_i$  are called *indices*. After applying Feynman rules to a family of integrals, one obtains integrals with Lorentz indices. Then the usual procedure is to express them in terms of scalar integrals (1), where the indices can take, generally speaking, any integer values. Feynman integrals are rational functions of the regularization parameter  $d$  with poles arising from ultraviolet, infrared and collinear divergences.

The number of resulting integrals can be very large but in the early days of perturbative quantum field theory it was necessary to evaluate every appearing Feynman integral by some method. The situation has significantly changed after the groundbreaking discovery of Chetyrkin and Tkachov [1], who suggested to apply so-called integration by parts relations (IBP) in order to express any given Feynman integral as a linear combination of a finite number of Feynman integrals. These relations are derived from equations

$$\int \cdots \int \left( \frac{\partial}{\partial k_j} \cdot r \right) \left( \frac{1}{P_1^{a_1} \cdots P_n^{a_n}} \right) \prod_{i=1}^h d^d k_i = 0, \quad (2)$$

where  $r$  is a loop momentum  $k_j$  or an external momentum  $p_l$  and  $j = 1, \dots, h$ . After the differentiation, resulting scalar products  $k_j \cdot k_{j'}$  and  $k_j \cdot p_l$  are expressed in terms of the factors in the denominator, and one obtains a system of equations which are called IBP relations. The dependence on  $d$  appears due to the prescription  $\frac{\partial}{\partial k_j} \cdot k_j = d$  of dimensional regularization. When solving IBP relations Feynman integrals are considered as algebraic objects. However, when evaluating a given Feynman integrals analytically, one usually applies Feynman parametric representation given by a multifold parametric integral, where  $d$  enters as a complex parameter in exponents of some polynomials.

The very first impressive application of the IBP relations was immediately presented in [1], where an algorithm to evaluate massless three-loop propagator integrals was described. Later this algorithm was implemented as the program MINCER [2, 3] (implemented in FORM [4]) to evaluate this class of Feynman integrals. It was used in numerous applications. We will mention below only one similar result because we are oriented at general algorithms and corresponding computer codes which can, in principle, be applied to any family of Feynman integrals.

The IBP relations following from (2) can be written down as operator equations with the help of *annihilation/creation* operators

$$\mathbf{i}^{\pm} I_{a_1, \dots, a_n} = I_{a_1, \dots, a_{i-1}, a_i \pm 1, a_{i+1}, \dots, a_n}.$$

Starting from these operator equations one obtains an infinite set of equations for functions of integer variables  $a_i$ . From the mathematical point of view, the IBP reduction is solving a huge sparse system of linear equations with polynomial coefficients.

To be more exact, when one fixes a subset of unknown Feynman integrals (that is, the possible values of indices), and writes down all possible IBP relations connecting them one obtains a sparse system of linear equations with polynomial coefficients. This system can be solved reducing all unknowns to integrals that are smallest under some chosen ordering. As it was observed long time ago, if one keeps increasing the number of possible indices, and thus the number of integrals involved, the number to which the others are reduced eventually stabilizes. This number remains constant no matter how much the indices are increased. Such irreducible integrals are called *master* integrals. However, this is not a strict mathematical definition and its more rigorous formulation is beyond the scope of this paper. It was proved in [5] that the number of master integrals is always finite. Nevertheless, this fact does not help in calculations since the proof is purely theoretical.

In practice, a sufficiently large subset of Feynman integrals from a given family is usually considered, for example, all integrals for which the inequality  $\sum_{i=1}^n |a_i| \leq M$  holds. One can refer to these integrals as *seeded* Integrals. Then all IBP relations are written out, that involve only the integrals in this subset, and solved. One introduces an ordering when solving such a system: integrals with small values of  $|a_i|$  are naturally considered to be on the right-hand side of the corresponding solutions. Following Laporta's presentation of the first detailed version of such an algorithm [6], these methods became widely known as "Laporta algorithms".

However, the strategy formulated above is far from being optimal. It turns out that in advanced strategies, the notion of *sector* is used. A sector of integrals of a given family (2) is determined by a set of positive indices  $a_i$  in (1), while the rest of indices are non-positive. It is enough to construct a reduction procedure in each sector. Then a Feynman integral from a given sector can be reduced to master integrals in the same sector up to a linear combination of integrals of *lower* sectors, i.e. for which some of the positive indices become non-positive. There is a variety of ways to construct the reduction in a given sector presented in available algorithms.

The goal of this paper is to present a historiographical review of algorithms and computer programs to solve IBP relations with explanations of their essential features. We follow the chronological order also because this is a story about competition in the field of solving IBP relations. It turns out that this competition was fruitful since existing algorithms were taken into account when constructing new algorithms and/or new versions of existing algorithms. We do not refer at all to papers with applications of these algorithms and computer codes in practical calculations because the number of corresponding applications is exceedingly large. However, they can easily be found by studying citations to a given algorithm. By examining the applications and focusing on the most important ones, it is possible to decide which algorithm is more suitable for a given problem.

**2. AIR.** The first computer implementation of the Laporta algorithm was created by Anastasiou and Lazopoulos [7] and called Automatic Integral Reduction (AIR). To define enlarging sets of seeded integrals, the following functions of integrals of a given family  $I(\{a_i\})$  are introduced:

$$N_{\text{prop}} = \sum_i \Theta(a_i), \quad N_+ = \sum_i \Theta(a_i) (a_i - 1), \quad N_- = - \sum_i \Theta(-a_i) a_i,$$

where  $\Theta(x) = 1$  for  $x > 0$  and 0 otherwise. At each elimination step, the most "complicated" integral (highest  $N_{\text{prop}}$ ,  $N_+$ ,  $N_-$  in lexicographic order) is isolated and expressed in terms of simpler ones. The user specifies ranges of  $(N_{\text{prop}}, N_+, N_-)$  and template IBP/LI relations for a given family. Here and below by LI we mean so-called Lorentz-invariance identities [8], which are used together with IBP relations. This allows AIR to generate explicit "seed equations" for concrete sets of indices  $\{a_i\}$ . The resulting equations are sorted and used in the elimination procedure.

The main reduction cycle of AIR consists of the following steps:

- selecting a seed integral;
- generating IBP/LI equations;
- ordering terms by complexity and isolating the hardest integral;
- solving in terms of simpler ones;
- performing back-substitution into previously derived equations.

This iterative process gradually reduces all integrals to master integrals.



To control the growth of intermediate algebraic expressions, AIR employs also two masking strategies. Equations are stored in a file database, where each one is associated with its target (isolated) integral. Auxiliary index files keep track of where each integral appears. At the end, masked objects are re-expanded through nested substitutions and any given integral  $I$  is expressed as  $I = \sum_j c_j(\{s_l\}, d) M_j$ , where  $M_j$  are master integrals and coefficients  $c_j$  are rational functions of kinematic invariants  $s_l$ , particle masses and space-time dimension.

**3. FIRE 1–3.** Before the public release of FIRE (Feynman Integrals REduction), our work was focused on building prototype systems referred as FIRE1 and FIRE2 (written in *Wolfram Mathematica*), which were never made public.

The initial formulation of FIRE was strongly influenced by our attempt to apply Gröbner basis techniques to the reduction of Feynman integrals [9, 10]. The IBP relations can be represented as polynomial relations in shift operators acting on indices of Feynman integrals. By adapting and modifying the Buchberger algorithm [11], the Gröbner-type bases in this operator algebra were constructed. This perspective provided a systematic framework for solving the reduction problem: an arbitrary integral from a given family can be reduced to a linear combination of a smaller set of master integrals by repeated application of these basis relations. In FIRE1 the user had to provide manually *boundary conditions*, which characterize index regions where integrals take zero values, thus allowing the algorithm to terminate effectively.

The version FIRE2 was an update of FIRE1 that incorporated several new features such as symmetry relations and internal optimizations. Still at this point the authors came to a conclusion that it becomes really impossible to construct Gröbner bases in some sectors. Therefore, to move forward, the Laporta algorithm had to be implemented in FIRE. The first public version of FIRE named FIRE3 still written in *Wolfram Mathematica* and having a hybrid approach, combining Gröbner-basis with Laporta elimination, was published in 2008 [12]. The version FIRE3 possessed the following features:

- Laporta mode: a classical IBP-based elimination in sectors, used when no s-basis is available for a given sector;
- s-basis mode: in sectors where an s-basis has been constructed, any non-master integral can be reduced to lower ones in (essentially) constant time by symbolic substitution, avoiding a full system solve;
- region-bases: a feature allowing to treat shifted indices.

To bring in a proper formalism and to define master-integrals in a mathematical manner, the notion of a global integral ordering was refined in FIRE3. A global ordering is defined first by diagram number, then by sector, and then within the sector by a linear ordering on the shift vector representation. Sectors are defined by sign patterns of indices (positive vs non-positive). Regions generalize sectors allowing some indices to be “passive” (i.e. not fixed sign). The presence of region definitions requires that the ordering be adjusted so that integrals lying in the lower regions appear “lower” in the partial order, and that reduction steps do not violate region consistency. With these definitions FIRE3 formalizes a concept: a *proper expression* for an integral  $I$  is an identity expressing  $I$  as a linear combination of integrals strictly lower in the ordering. The core algorithm ensures that eventually all integrals of interest possess a proper expression and at this stage the backward substitution reduces them to masters integrals. The master integrals can now be defined like integrals for which no a proper expression can exist.

To manage the growth of the number of expressions, FIRE3 uses *masking* techniques (analogous to those in AIR). When an IBP relation involves integrals from lower sectors, the parts corresponding to lower-sector integrals are masked (not substituted immediately). This is called *tail-masking*. The substitution of such expressions is delayed. The tail-masking concept turned out to be important for future reduction development, especially within the FIRE framework working unlike other reduction programs in two steps: first, top-bottom reduction with tails being masked and then, a bottom-top backward substitution.

In analogy to approaches proposed by Lee [16], FIRE3 could also eliminate certain IBP relations deemed redundant and only generate a subset sufficient to express most integrals in a sector. This reduces the number of IBPs one must consider in system solving.

While still being written in *Wolfram Mathematica*, FIRE3 introduced such tools as QLink, an interface to the QDBM disk-based database for managing large algebraic storage within Mathematica, and FLink, a bridge to the Fermat program, which can perform fast polynomial arithmetic and substitution outside Mathematica.

**4. Reduze.** The next implementation of the Laporta algorithm was presented by Studerus [13] in 2009, where some notation was introduced to be used in future papers too.

Let  $T_t$  be the set of integrals of a given family associated with integrals with  $t$  indices. A sub-sector  $T_{t-1}$  of a sector  $T_t$  is a sector where one propagator is removed. There are in general  $t$  different sub-sectors for a sector  $T_t$ . The sub-sector tree of a sector  $T_t$  is the set of all sub-sectors of  $T_t$  and, recursively, all sub-sectors of these sub-sectors. In other articles *sub-sectors* can also be named as *lower sectors*.

The following additional notation is also introduced in [13]. To every  $t$ -propagator sector  $T_t$  with propagators  $P_{j_1}, \dots, P_{j_t}$  belongs a infinite set of  $d$ -dimensionally regularized  $l$ -loop integrals, which all share the same propagators in the denominator of the integrand. They have the generic form

$$\int d^d k_1 \dots \int d^d k_l \frac{P_{j_{t+1}}^{s_1} \dots P_{j_n}^{s_{n-t}}}{P_{j_1}^{r_1} \dots P_{j_t}^{r_t}} \quad (3)$$

with integer exponents  $r_i \geq 1$  and  $s_i \geq 0$ . Then, one defines  $r = \sum_{i=1}^t r_i \geq t$  as the sum of the indices propagators in the denominator and  $s = \sum_{i=1}^{n-t} s_i \geq 0$  as the sum of the indices of propagators in the numerator. These two numbers  $r$  and  $s$  identify the complexity of integrals of a given sector. But since  $r$  is obviously not smaller than the number of positive indices, nowadays a shifted formula can be used with  $r = \sum_{i=1}^t (r_i - 1) \geq t$ , so that  $r$  and  $t$  together can be called a complexity of the shift from the corner of the sector.

To reduce the integrals of a sector up to  $r = r_{\max}$  and  $s = s_{\max}$ , it is necessary to solve the homogeneous system of linear equations, which follow from the corresponding IBP and/or LI identities. Then, one obtains a solution in terms of some integrals of a given sector and integrals of lower sectors. Afterwards, the procedure continues in these lower sectors and finally a solution in terms of true master integrals is obtained.

The reduction scheme of **Reduze** is as follows. For the reduction of a sector and its sub-sectors, **Reduze** first constructs the full sub-sector tree. It begins by reducing the lowest sub-sectors (with the fewest propagators), substitutes their results into higher sectors and continues iteratively until the target sector with the largest number of propagators is reduced. For a single-sector reduction, **Reduze** generates the corresponding equations and inserts available sub-sector results from the default directory.

Since **Reduze** cannot reduce a single integral directly, it uses all integrals within a user-defined range of  $r$  and  $s$  to form and reduce the full equation system. Because these systems can be very large, **Reduze** automatically divides them into smaller subsets. Equations are sorted by the complexity of their most complicated integral, split into user-defined chunks, and stored in temporary files. The reduction proceeds sequentially: the simplest subset is reduced first, its results are inserted into all others, and the process repeats for each subsequent subset. In the end, all results are combined into a single output file.

The approach of **Reduze** here differs from the approach of **FIRE** since it uses one bottom-top pass when solving equations.

**5. Reduze 2.** The second version of **Reduze** [14], presented by von Manteuffel and Studerus in 2012, was a considerable revision and extension of **Reduze 1**, rearchitected to support the distributed algorithmic framework. The main feature was built around the distributed (MPI-based) approach to reduction, supporting distributed reduction of a single family of integrals (or, a single sector) across multiple processor cores using MPI hence using multiple computers for the same task. It also added a support to a parallel reduction of different families and a load balancer to properly distribute jobs between nodes.

This is a summary of other new features:

- graph and matroid-based algorithms for sector/integral relations. A key improvement is the use of fast graph- and matroid-based methods to identify equivalences (relations) among sectors (families) and integrals, including those across different families, and to match diagrams to canonical sectors. These algorithms help to reduce redundancy and simplify the reduction workload;
- improved handling of sector symmetries and relations. **Reduze 2** enhances the detection and utilization of sector relations (between sectors of the same or different families) and sector symmetries (automorphisms), beyond the symmetry handling in the first version;
- matching of diagrams to canonical sectors. **Reduze 2** implements routines to match a given Feynman diagram (with its labeling) to the canonical representative sectors (via mapping or shifts);
- differential equations for master integrals support. **Reduze 2** includes features to derive differential equations for master integrals (in the kinematic variables) as a built-in option;
- interference-term reduction. The program can reduce interference terms (e.g. integrals appearing in interference of amplitudes) directly.





**6. LiteRed.** In 2012 Lee introduced **LiteRed**, a Mathematica package aimed at automating the heuristic derivation of symbolic IBP reduction rules (rather than repeatedly solving linear systems) [15, 16]. **LiteRed** departs, to some extent, from the traditional Laporta-style approach of solving large IBP systems anew for each integral. Instead, its goal is to precompute reduction rules (in a sector-by-sector manner) that can later be applied directly to any integral in that family. Once the symbolic rules are found, reduction becomes a matter of rule application, which is extremely fast and lightweight. As the author states, symbolic rules are compact and can be stored for reuse, avoiding repeated costly system solving. Thus, in a way **LiteRed** follows the ideas of **FIRE1-2**, but the author developed his own way of precomputing reduction rules not related to Gröbner bases.

However, the trade-off is that the search for symbolic rules is heuristic and not guaranteed to succeed in all cases. The algorithm may fail or stall in particularly challenging sectors, and success often depends on ordering choices or the particular set of irreducible numerators chosen.

**LiteRed**'s workflow broadly splits into two phases: rule search (preprocessing) and rule application (actual reduction). Precomputation consists of the following steps:

- the user defines a basis for integrals (i.e. denominators and optional numerators) via a “NewBasis” command. **LiteRed** checks linear independence and completeness of the basis representation in terms of scalar products;
- IBP and Lorentz-invariance identities are generated via **GenerateIBP**, producing symbolic relations in shift-operator form;
- sectors are analyzed via **AnalyzeSectors**: zero (scaleless) sectors are identified, and “simple” sectors are recognized;
- symmetries between sectors are detected through **FindSymmetries**, which produces mappings between equivalent or isomorphic sectors, generating substitution rules to reduce duplication;
- the central routine **SolvejSector** attempts to construct a full set of symbolic reduction rules in a given sector (for each unique sector after symmetry factoring). The outcome is a set of reduction rules “jRules”, which express non-master integrals in terms of simpler ones.

The sector solving approach of **LiteRed** turned out to be more effective than the one in **FIRE1-2**. Let us note that some important information revealed with **LiteRed** is used in the next versions of **FIRE** (see below).

**7. FIRE4.** Published in 2013 it was the first version [17] of **FIRE** that featured reduction in C++ providing a serious performance upgrade with a direct access to databases and **Fermat** algebraic evaluations. Apart of that it has the following features:

- integration with “tsort” using an algorithm (based on Feynman parameters and described by Pak [18]) to search for equivalences between integrals and detect additional linear relations among integrals (beyond IBP) to reduce the master set;
- use of symmetries and Lorentz-invariance identities to derive further relations among master integrals, which helps to eliminate spurious masters;
- integration with **LiteRed** — the new version of **FIRE** was able to provide input for **LiteRed** and to use its output for the C++ reduction both handling symmetry relations and reduction rules;
- the command “MakeMaster” to allow the user to assign priorities to particular integrals to be chosen as masters in ambiguous sectors.

In the case of **LiteRed** integration it turned out to be always beneficial to use **LiteRed** symmetries between sectors. **FIRE4** could also use **LiteRed** reductions rules resorting to internal reduction methods in the missing sectors in case **LiteRed** does not provide full coverage. However, the further usage has shown that it is beneficial either to build **LiteRed** reduction rules everywhere or not to use them at all, since the rules might greatly increase the complexity of integrals to be reduced in lower sectors and when the rules are missing there, the reduction becomes more complex than the original one without **LiteRed** rules. Let us mention that **LiteRed** was used to construct **FORCER** [19], i.e. a four-loop variant of the **MINCER** code [2, 3]. Here, **LiteRed** provided a starting set of symbolic reduction rules, that were manually optimized for most complicated four-loop diagrams.

**8. FIRE5.** Then **FIRE5** [20] was released in 2014. This was a major milestone in the **FIRE** lineage: the core reduction engine was rewritten in C++, while Mathematica remained as the user-facing frontend and orchestration layer. This structural shift allowed to overcome many performance and scalability limitations of a

pure Mathematica implementation. The whole scheme introduced at this stage was the following (and remained true for more recent versions):

- **Mathematica** frontend handles problem setup (diagram definition, generation of IBP relations, symmetry declarations, sector definitions, boundary conditions), programming interface, and final result collection;
- **C++** backend performs the heavy reduction steps: equation solving, substitution into integrals, table building;
- interfacing between them is handled via files/tables produced by **Mathematica** that the **C++** engine consumes, and result tables emitted by **C++** that can be loaded back into **Mathematica**;
- the design is backward-compatible: reductions produced by **FIRE5** can be imported into earlier **FIRE** versions, and the same user workflow (e.g. **Mathematica** commands like `Burn[]`, `LoadTables`) remains available.

By moving the bottleneck parts to **C++**, the version **FIRE5** achieved orders-of-magnitude improvements in performance for nontrivial reductions.

The **C++** reduction implementation brought the following features:

- multi-threading/parallelism: **FIRE5** supports parallel reduction tasks over multiple cores;
- **Fermat** integration: for symbolic polynomial arithmetic (simplification, GCDs, rational function operations) the **C++** engine delegates to the external **Fermat** program;
- KyotoCabinet database backend: used for efficient on-disk storage of tables and sparse relations, allowing large intermediate data beyond in-memory capacity; the database settings can adjust how aggressively it caches in memory or spills to disk depending on available RAM, to optimize performance;
- compression with Snappy: intermediate tables and data structures are compressed to reduce disk space and I/O load.

**9. Kira.** The first version of **Kira** [21] by Maierhöfer, Usovitsch and Uwer published in 2017 was designed to handle problems with multiple mass and kinematic scales, mitigating symbolic blowup in intermediate expressions. **Kira** is written in **C++**, targeting 64-bit Linux systems, with support for multi-core parallelism. It uses external algebraic support from **Fermat** (via a “gateToFermat” interface) and **GiNaC**, and employs libraries like `yaml-cpp`, `zlib`, and `SQLite3` for configuration and storage.

The reduction workflow is as follows:

- generate a system of linear relations (IBP + Lorentz invariance + symmetry relations) among integrals in chosen seed sets;
- perform a preprocessing elimination step using modular arithmetic to detect and remove linearly dependent equations before symbolic elimination. This step also helps to identify master integrals early;
- using the remaining independent equations, perform forward elimination (triangularization) according to a chosen ordering of integrals and equations;
- execute a back-substitution phase to express dependent integrals in terms of master integrals; the algorithm is optimized to delay full expansions and reduce intermediate expression swell;
- post-process result tables, output mappings to master integrals, and provide user interfaces for queries.

A given set of Feynman integrals are ordered lexicographically by a tuple involving (topology ID, sector ID, the sum of positive indices  $r$ , the sum of negative indices  $s$ , and detailed index vectors). Equations are sorted by their leading integral (largest in the ordering), then by equation length. Symmetry relations among integrals and zero (trivial) sectors are detected and exploited to reduce system size.

The modular-arithmetic preprocessing is the principal novelty: it avoids needless symbolic operations on redundant equations and thus alleviates memory and time overheads especially in multi-scale settings. Benchmark comparisons to **Reduze 2** [14] and **FIRE5** [20] showed that **Kira** is highly competitive, sometimes outperforming them in challenging multi-scale reductions. Nevertheless, for extremely large systems the memory and CPU demands remain significant; the success of **Kira** depends on effective seeding, symmetry detection, and the ability to eliminate redundant relations early.

**10. FIRE6.** The release of **FIRE6** [22] in 2019 marked a major revision of the **FIRE** framework, integrating modular arithmetic into its **C++** core and introducing large-scale parallel capabilities via MPI. While **FIRE5** already provided a fast **C++** engine, **FIRE6** significantly extended both its computational reach and stability for extremely large reductions.



The most important innovation was *the modular arithmetic reduction mode*, inspired in part by Kira’s finite-field preprocessing. Instead of performing symbolic elimination over rationals, FIRE6 presents a mode to run multiple reductions for many integer substitutions of variables modulo large primes and later reconstructs the rational coefficients. This method, already explored in Finred (see a short description in Conclusion) and Kira, was implemented here using MPI parallelization across compute nodes to run reductions in parallel over distinct primes.

Key features of FIRE6 include:

- full support for distributed reduction across multiple computers and nodes (MPI mode), allowing scalability up to thousands of cores;
- implementation of modular arithmetic reconstruction to mitigate coefficient blowup and accelerate large-scale reductions;
- automatic recovery from system crashes — reduction tasks can resume from checkpoints instead of restarting;
- integration of internal sector symmetries and rule generation from LiteRed, improving elimination efficiency and master-integral identification;
- upgraded I/O and compression subsystems (using Snappy, ZStandard, KyotoCabinet), improving disk efficiency and database handling;
- enhanced C++11 compliance, improved documentation, and doxygen-based internal API descriptions for contributors.

Conceptually, FIRE6 incorporates Kira’s finite-field strategy but remains distinct in its execution model: Kira uses modular arithmetic primarily for dependency detection and reconstruction, while FIRE6 applies modular arithmetic directly as the main reduction engine. Its implementation is therefore more deeply parallelized and designed for heterogeneous environments — from desktops to supercomputers.

However, FIRE6 had a significant disadvantage — the reconstruction could be performed only within Wolfram Mathematica. (This was later improved in FIRE7). Nevertheless, there was a number of successful projects, where the reduction was carried out using supercomputers with subsequent reconstruction of coefficients.

**11. Kira2.** The main novelty that was introduced in the second version of Kira [23] presented by Klappert, Lange, Maierhöfer and Usovitsch in 2020 was Finite field reduction using integrated FireFly package (first published as a standalone code [24, 25]) with a built in MPI support for distributing reduction jobs. Unlike FIRE6, Kira came to a complete solution making the reconstruction possible without proprietary software such as Wolfram Mathematica. The FireFly usage also enables Kira to separate known denominators (prefactors) from coefficients prior to interpolation, thereby simplifying the interpolation task.

Kira 2.0 significantly extended the flexibility, scalability, and efficiency of Feynman integral reductions. Apart of that Kira2 provides a number of new features:

- user-defined systems of equations — an improved interface for supplying one’s own linear systems (e.g. not necessarily standard IBP systems). Support for multiple input files, compressed files, on-the-fly solving, and custom integral weights. Also it allows more flexible definitions of propagators (e.g. combinations with Feynman parameters, bilinear forms);
- iterative reduction — automating splitting reductions over master integrals or sectors (setting other masters or sectors to zero), reducing memory footprint by iterating over subsets. Works both with algebraic and finite-field modes. Also Kira2 allows marking entire sectors as trivial (all integrals zero) so that they are excluded from reduction;
- master equations — allows treating linear combinations of integrals as master integrals by adding equations equating them to such combinations. Useful for choosing convenient bases, e.g. for differential equations;
- preferred-master sector inclusion. If a master integral lies in a sector not explicitly requested, Kira ensures that its sector is still generated to support that master integral. This helps find “magic” relations via higher sectors.
- export and re-import of reduction rules. With `kira2file` one can export reduction rules in a format compatible with user-defined input systems; results may also be exported to Mathematica or FORM via `kira2math`, `kira2form`.



**12. FIRE6.5.** In 2023 the intermediate release FIRE6.5 [26] introduced a new library, FUEL [27], which makes it possible to simplify algebraic coefficients not with a single Computer Algebra System (CAS) but allows integration with multiple simplifier backends. FUEL provides a unified interface to multiple simplifier backends (e.g. FLINT, Symbolica [28] and Fermat). The open-source FLINT [29] backend was newly added to this release and often outperforms Fermat in problems involving many kinematic variables. The Symbolica simplifier (by Ruijl) is also integrated as a potential successor to FORM. Thus, FIRE6.5 becomes more flexible: the user may choose which simplifier to use, and the core engine is no longer bound to one CAS. To maintain backward compatibility, the default behavior remains Fermat.

**13. Kira3.** Kira3 [30], published in 2025 by Lange, Usovitsch and Wu, introduced optimized seeding and equation selection algorithms, significantly improving performance for multi-loop and multi-scale problems. The following features can be marked as most important:

- optimized seeding strategy — the selection of seeds is improved. Kira3 inherits bounds ( $r_{\max}, s_{\max}$ ) (see (3)) more flexibly across subsectors, and allows *truncation* of seeds so that only a subset of seeds are generated that suffice for full reduction;
- better equation selection — after generating linearly independent IBP equations, Kira3 employs a refined selection algorithm: it performs a numerical finite-field check post forward elimination to detect hidden zeros and drop irrelevant equations. This reduces the number of selected equations and, thus, speeds up the solving step;
- support for extra relations — users can now include additional linear relations (e.g. from nontrivial symmetries or kinematic constraints) in the standard reduction workflow by specifying `extra_relations` in the job file;
- numerical sampling mode — the user may provide finite-field sample points (via `numerical_points`) to drive reductions numerically, possibly bypassing full analytic expressions or guiding the interpolation step. Kira ensures that the seed generation is consistent with the sampled phase space to avoid loss of information;
- symbolic IBP reductions — Kira3 can treat one or more propagator powers symbolically (i.e. as formal symbols) and derive recurrence (lowering/raising) relations by Laporta’s approach. This allows to generate analytic recursion relations in favorable cases;
- master basis checking — with the option `check_masters`, Kira3 verifies whether a user-specified preferred master integral basis is sufficient. If not, it reports missing or superfluous masters and aborts.

**14. FIRE7.** The release of FIRE7 [31] marks a significant evolution of the FIRE program, focused on optimizing the performance of IBP reduction under modular arithmetic, introducing new pre-solve strategies, enhanced ordering control, and a suite of auxiliary tools that streamline reduction workflows. Within FIRE7 one can completely automatize the reduction and reconstruction process with modular arithmetic on supercomputers while still keeping the original FIRE idea that the user is supposed to provide the list of integrals requiring reduction and not care about levels of integrals being seeded.

FIRE7 extends the command-line toolset with utilities for direct manipulation of reduction tables with everything being implemented in C++:

- `reconstruct` — runs rational or modular reconstruction with different methods;
- `tables2rules` — converts reduction tables into Mathematica-readable rule format without invoking Mathematica;
- `add` — merges tables from parallel or split-master runs and supports subtraction mode for consistency checks;
- `combine` — reduces linear combinations of integrals using precomputed tables, useful for form-factor and amplitude applications;
- `substitute` — replaces analytic variables by numerical values (modular probes), ensuring full consistency with direct modular runs;
- `diff` — efficiently compares tables to verify identical results across different configurations or numerical probes.

Together, these utilities make it possible to conduct complete modular reductions, merging, validation, and reconstruction workflows entirely from the command line, independent of Mathematica. This separation of tasks facilitates large-scale, automated reductions on high-performance computing systems.

Aside of that **FIRE7** comes with a bunch of new improvements:

- two-stage Gaussian elimination pre-solve before seeding sampling points;
- configurable integral orderings (like in **Kira**);
- multiprime modular mode for parallel reconstruction allowing to run reduction in multiple prime points at the same time;
- split-master reduction and table combination workflow;
- extended support for **LiteRed2** integration.

**15. Conclusion.** Notably, various private algorithms and codes have been and continue to be used in the field. Let us mention some of them because references to corresponding private versions can be found in multiple papers. These are, for example, **IdSolver** by Czakon and **Crusher** by Marquard and Seidel. We cannot describe them in detail in this review, as they were never made available to a broad audience; consequently, their key features remain unknown. Let us also point out that the use of modular arithmetic for solving IBP relations was originally advocated by von Manteuffel and Schabinger [32]. The corresponding private code by von Manteuffel is called **Finred**. It is also worth mentioning methods based on syzygy equations, which can be considered a modern alternative to Gröbner basis techniques and the **LiteRed** approach. This is because they aim to generate better relations before substituting indices. We did not write a section on these methods since such approaches have not yet resulted in public standalone reduction programs. Nevertheless, a program based on this approach, **Neat IBP** [33], was recently integrated into the **Kira** framework [34] to serve as a generator of improved equations.

Another interesting modern approach involves improving how the sampling points are seeded for the reduction. A package that performs this task was published last year [35].

To summarize, we described known public general algorithms and computer codes for solving IBP relations for Feynman integrals. Nevertheless, even after two decades of algorithmic development for Feynman integral reduction, new ideas and methods must still be implemented to achieve optimal performance. We therefore hope that the work in this field continues, leading to the development and practical application of new powerful algorithms and improved versions of existing ones.

## References

1. K. G. Chetyrkin, F. V. Tkachov, “Integration by parts: The algorithm to calculate  $\beta$ -functions in 4 loops,” Nucl. Phys. B **192** (1), 159–204 (1981). doi [10.1016/0550-3213\(81\)90199-1](https://doi.org/10.1016/0550-3213(81)90199-1).
2. S. G. Gorishny, S. A. Larin, L. R. Surguladze, F. V. Tkachov, “Mincer: Program for Multiloop Calculations in Quantum Field Theory for the Schoonschip System,” Comput. Phys. Commun. **55** (3), 381–408 (1989). doi [10.1016/0010-4655\(89\)90134-3](https://doi.org/10.1016/0010-4655(89)90134-3).
3. S. A. Larin, F. V. Tkachov, J. A. M. Vermaseren, “*The FORM version of MINCER*,” Preprint NIKHEF-H/91–18 (National Institute for Subatomic Physics (NIKHEF), Amsterdam, Netherlands, 1991). <https://inis.iaea.org/records/y68hj-8d895>.
4. J. A. M. Vermaseren, “New features of FORM,” arXiv:math-ph/0010025 (2000). doi [10.48550/arXiv.math-ph/0010025](https://doi.org/10.48550/arXiv.math-ph/0010025).
5. A. V. Smirnov and A. V. Petukhov, “The Number of Master Integrals is Finite,” Lett. Math. Phys. **97** (1), 37–44 (2011). doi [10.1007/s11005-010-0450-0](https://doi.org/10.1007/s11005-010-0450-0).
6. S. Laporta, “High-precision calculation of multiloop Feynman integrals by difference equations,” Int. J. Mod. Phys. A **15** (32), 5087–5159 (2000). doi [10.1142/S0217751X00002159](https://doi.org/10.1142/S0217751X00002159).
7. C. Anastasiou and A. Lazopoulos, “Automatic integral reduction for higher order perturbative calculations,” J. High Energy Phys. No. 7, Article Number 046 (2004). doi [10.1088/1126-6708/2004/07/046](https://doi.org/10.1088/1126-6708/2004/07/046).
8. T. Gehrmann, E. Remiddi, “Differential equations for two-loop four-point functions,” Nucl. Phys. B **580** (1), 485–518 (2000). doi [10.1016/S0550-3213\(00\)00223-6](https://doi.org/10.1016/S0550-3213(00)00223-6).
9. A. V. Smirnov, V. A. Smirnov, “Applying Gröbner bases to solve reduction problems for Feynman integrals,” J. High Energy Phys. No. 1, Article Number 001 (2006). doi [10.1088/1126-6708/2006/01/001](https://doi.org/10.1088/1126-6708/2006/01/001).
10. A. V. Smirnov, “An Algorithm to construct Gröbner bases for solving integration by parts relations,” J. High Energy Phys. No. 4, Article Number 026 (2006). doi [10.1088/1126-6708/2006/04/026](https://doi.org/10.1088/1126-6708/2006/04/026).
11. B. Buchberger, “Bruno Buchberger’s PhD thesis 1965: An algorithm for finding the basis elements of the residue class ring of a zero dimensional polynomial ideal,” J. Symbolic Comput. **41**, 475–511 (2006). doi [10.1016/j.jsc.2005.09.007](https://doi.org/10.1016/j.jsc.2005.09.007).
12. A. V. Smirnov, “Algorithm FIRE — Feynman Integral REduction,” J. High Energy Phys. No. 10, Article Number 107 (2008). doi [10.1088/1126-6708/2008/10/107](https://doi.org/10.1088/1126-6708/2008/10/107).

13. C. Studerus, “Reduze — Feynman Integral Reduction in C++,” *Comput. Phys. Commun.* **181** (7), 1293–1300 (2010). doi [10.1016/j.cpc.2010.03.012](https://doi.org/10.1016/j.cpc.2010.03.012).
14. A. von Manteuffel, C. Studerus, “Reduze 2 - Distributed Feynman Integral Reduction,” arXiv:1201.4330 [hep-ph] (2012). doi [10.48550/arXiv.1201.4330](https://doi.org/10.48550/arXiv.1201.4330).
15. R. N. Lee, “LiteRed 1.4: a powerful tool for reduction of multiloop integrals,” *J. Phys. Conf. Ser.* **523**, Article Number 012059 (2014). doi [10.1088/1742-6596/523/1/012059](https://doi.org/10.1088/1742-6596/523/1/012059).
16. R. N. Lee, “Presenting LiteRed: a tool for the Loop InTEgrals REDuction,” arXiv:1212.2685 [hep-ph] (2012). doi [10.48550/arXiv.1212.2685](https://doi.org/10.48550/arXiv.1212.2685).
17. A. V. Smirnov, V. A. Smirnov, “FIRE4, LiteRed and accompanying tools to solve integration by parts relations,” *Comput. Phys. Commun.* **184** (12), 2820–2827 (2013). doi [10.1016/j.cpc.2013.06.016](https://doi.org/10.1016/j.cpc.2013.06.016).
18. A. Pak, “The toolbox of modern multi-loop calculations: novel analytic and semi-analytic techniques,” *J. Phys. Conf. Ser.* **368**, Article Number 012049 (2012). doi [10.1088/1742-6596/368/1/012049](https://doi.org/10.1088/1742-6596/368/1/012049).
19. B. Ruijl, T. Ueda and J. A. M. Vermaseren, “Forcer, a Form program for the parametric reduction of four-loop massless propagator diagrams,” *Comput. Phys. Commun.* **253**, Article Number 107198 (2020). doi [10.1016/j.cpc.2020.107198](https://doi.org/10.1016/j.cpc.2020.107198).
20. A. V. Smirnov, “FIRE5: A C++ implementation of Feynman Integral REDuction,” *Comput. Phys. Commun.* **189**, 182–191 (2015). doi [10.1016/j.cpc.2014.11.024](https://doi.org/10.1016/j.cpc.2014.11.024).
21. P. Maierhöfer, J. Usovitsch and P. Uwer, “Kira — A Feynman integral reduction program,” *Comput. Phys. Commun.* **230**, 99–112 (2018). doi [10.1016/j.cpc.2018.04.012](https://doi.org/10.1016/j.cpc.2018.04.012).
22. A. V. Smirnov, F. S. Chukharev, “FIRE6: Feynman Integral REDuction with modular arithmetic,” *Comput. Phys. Commun.* **247**, Article Number 106877 (2020). doi [10.1016/j.cpc.2019.106877](https://doi.org/10.1016/j.cpc.2019.106877).
23. J. Klappert, F. Lange, P. Maierhöfer, and J. Usovitsch, “Integral reduction with Kira 2.0 and finite field methods,” *Comput. Phys. Commun.* **266**, Article Number 108024 (2021). doi [10.1016/j.cpc.2021.108024](https://doi.org/10.1016/j.cpc.2021.108024).
24. J. Klappert and F. Lange, “Reconstructing rational functions with FireFly,” *Comput. Phys. Commun.* **247**, Article Number 106951 (2020). doi [10.1016/j.cpc.2019.106951](https://doi.org/10.1016/j.cpc.2019.106951).
25. J. Klappert, S. Y. Klein, and F. Lange, “Interpolation of dense and sparse rational functions and other improvements in FireFly,” *Comput. Phys. Commun.* **264**, Article Number 107968 (2021). doi [10.1016/j.cpc.2021.107968](https://doi.org/10.1016/j.cpc.2021.107968).
26. A. V. Smirnov and Mao Zeng, “FIRE 6.5: Feynman integral reduction with new simplification library,” *Comput. Phys. Commun.* **302**, Article Number 109261 (2024). doi [10.1016/j.cpc.2024.109261](https://doi.org/10.1016/j.cpc.2024.109261).
27. K. S. Mokrov, A. V. Smirnov, and Mao Zeng, “Rational Function Simplification for Integration-by-Parts Reduction and Beyond,” *Numerical Methods and Programming* **24** (4), 352–367 (2023). doi [10.26089/NumMet.v24r425](https://doi.org/10.26089/NumMet.v24r425).
28. Symbolica | Modern Computer Algebra. <https://symbolica.io>. Cited October 27, 2025.
29. FLINT: Fast Library for Number Theory. <https://flintlib.org>. Cited October 27, 2025.
30. F. Lange, J. Usovitsch, and Z. Wu, “Kira 3: integral reduction with efficient seeding and optimized equation selection,” arXiv:2505.20197 [hep-ph] (2025). doi [10.48550/arXiv.2505.20197](https://doi.org/10.48550/arXiv.2505.20197).
31. A. V. Smirnov, Mao Zeng, “FIRE 7: Automatic Reduction with Modular Approach,” arXiv:2510.07150 [hep-ph] (2025). doi [10.48550/arXiv.2510.07150](https://doi.org/10.48550/arXiv.2510.07150).
32. A. von Manteuffel and R. M. Schabinger, “A novel approach to integration by parts reduction,” *Phys. Lett. B* **744**, 101–104 (2015). doi [10.1016/j.physletb.2015.03.029](https://doi.org/10.1016/j.physletb.2015.03.029).
33. Z. Wu, J. Boehm, R. Ma, H. Xu, and Y. Zhang, “NeatIBP 1.0, a package generating small-size integration-by-parts relations for Feynman integrals,” *Comput. Phys. Commun.* **295**, Article Number 108999 (2024). doi [10.1016/j.cpc.2023.108999](https://doi.org/10.1016/j.cpc.2023.108999).
34. Z. Wu, J. Böhm, R. Ma, J. Usovitsch, Y. Xu, and Y. Zhang, “Performing integration-by-parts reductions using NeatIBP 1.1 + Kira,” arXiv:2502.20778 [hep-ph] (2025). doi [10.48550/arXiv.2502.20778](https://doi.org/10.48550/arXiv.2502.20778).
35. X. Guan, X. Liu, Y.-Q. Ma, and W.-H. Wu, “Blade: A package for block-triangular form improved Feynman integrals decomposition,” *Comput. Phys. Commun.* **310**, Article Number 109538 (2025). doi [10.1016/j.cpc.2025.109538](https://doi.org/10.1016/j.cpc.2025.109538).

Received  
October 12, 2025

Accepted  
October 22, 2025

Published  
October 29, 2025

### Information about the authors

*Alexander V. Smirnov* — Dr. Sci., Laboratory head; 1) Research Computing Center, Moscow State University, Leninskie Gory, 1, building 4, 119991, Moscow, Russia; 2) Moscow Center for Fundamental and Applied Mathematics, Moscow State University, Leninskie Gory, 1, 119991, Moscow, Russia.

*Vladimir A. Smirnov* — Dr. Sci., Leading research fellow; 1) Skobeltsyn Institute of Nuclear Physics of Moscow State University, Leninskie Gory, 1, building 2, 119991, Moscow, Russia; 2) Moscow Center for Fundamental and Applied Mathematics, Moscow State University, Leninskie Gory, 1, 119991, Moscow, Russia.