

doi 10.26089/NumMet.v24r425

Rational Function Simplification for Integration-by-Parts Reduction and Beyond

Kirill S. Mokrov

Lomonosov Moscow State University, Faculty of Computational Mathematics and Cybernetics, Moscow,
Russia

ORCID: 0009-0001-5036-8503, e-mail: kmokrov@mail.ru

Alexander V. Smirnov

Lomonosov Moscow State University, Research Computing Center, Moscow, Russia

ORCID: 0000-0002-7779-6735, e-mail: asmirnov80@gmail.com

Mao Zeng

Higgs Centre for Theoretical Physics, Edinburgh, United Kingdom

ORCID: 0000-0002-4741-4038, e-mail: mao.zeng@ed.ac.uk

Abstract: We present FUEL (Fractional Universal Evaluation Library), a C++ library for performing rational function arithmetic with a flexible choice of third-party computer algebra systems as simplifiers. FUEL is an outgrowth of a C++ interface to Fermat which was originally part of the FIRE code for integration-by-parts (IBP) reduction for Feynman integrals, now promoted to be a standalone library with access to simplifiers other than Fermat. We compare the performance of various simplifiers for standalone benchmark problems as well as IBP reduction runs with FIRE. A speedup of more than 10 times is achieved for an example IBP problem related to calculation of the off-shell three-particle form factors in $\mathcal{N} = 4$ supersymmetric Yang-Mills theory.

Keywords: Feynman integrals, integration by parts, computer algebra.

Acknowledgements: We thank the authors of FLINT, Nemo and FORM for help with our questions about the software in mailing lists and/or private communications. We especially thank the author of Symbolica, Ben Ruijl, for tirelessly answering our questions and customizing the software to integrate with our library. The work of Alexander Smirnov was supported by the Russian Science Foundation under the agreement No. 21-71-30003. M.Z.'s work is supported in part by the U.K. Royal Society through Grant URF\R1\20109.

For citation: 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.



Упрощение рациональных функций для редукции с использованием соотношений интегрирования по частям и не только

К. С. Мокров

Московский государственный университет имени М. В. Ломоносова,
факультет вычислительной математики и кибернетики, Москва, Российская Федерация
ORCID: 0009-0001-5036-8503, e-mail: kmokrov@mail.ru

А. В. Смирнов

Московский государственный университет имени М. В. Ломоносова,
Научно-исследовательский вычислительный центр, Москва, Российская Федерация
ORCID: 0000-0002-7779-6735, e-mail: asmirnov80@gmail.com

Мао Цзэн

Центр теоретической физики Хиггса,
университет Эдинбурга, Эдинбург, Великобритания
ORCID: 0000-0002-4741-4038, e-mail: mao.zeng@ed.ac.uk

Аннотация: Мы представляем C++ библиотеку FUEL (Fractional Universal Evaluation Library), предназначенную для выполнения арифметики рациональных функций с гибким выбором сторонних систем компьютерной алгебры в качестве упрощителей. FUEL является развитием интерфейса C++ для Fermat, который изначально был частью кода FIRE для редукции соотношений интегрирования по частям (IBP) для интегралов Фейнмана, а теперь стал отдельной библиотекой с доступом к упрощителям, отличным от Fermat. Мы сравниваем производительность различных упрощителей для отдельных задач тестирования, а также для редукции IBP с помощью FIRE. Достигнуто ускорение более чем в 10 раз для примера редукции, связанной с вычислением трехпетлевых формфакторов вне оболочки светового конуса в $\mathcal{N} = 4$ суперсимметричной теории Янга-Миллса.

Ключевые слова: Интегралы Фейнмана, интегрирование по частям, компьютерная алгебра.

Благодарности: Мы благодарим авторов FLINT, Nemo и FORM за помощь в ответах на наши вопросы о программном обеспечении в списках рассылки и/или личных сообщениях. Мы особенно благодарим автора Symbolica Бена Руйла за неустанные ответы на наши вопросы и настройку программного обеспечения для интеграции с нашей библиотекой. Работа А. В. Смирнова поддержана Российским научным фондом по договору № 21-71-30003. Работа Мао Дзена частично поддерживается Королевским обществом Великобритании в виде гранта URF\R1\20109.

Для цитирования: Мокров К.С., Смирнов А.В., Цзэн М., Упрощение рациональных функций для редукции с использованием соотношений интегрирования по частям и не только // Вычислительные методы и программирование. 2023. 24, № 4. 352–367. doi 10.26089/NumMet.v24r425.

1. Introduction. Many problems of high energy physics and quantum field theory are difficult to solve without using a computer. An example of such a problem is the calculation of Feynman integrals in complicated scattering amplitudes and correlation functions. For cutting-edge problems involving a huge number of Feynman integrals, the standard calculation workflow consists of two stages: integration-by-parts (IBP) reduction [1, 2] to so-called master integrals and finding the values of these master integrals. The problem of IBP reduction with the *Laporta algorithm* [2] can be viewed as a problem of solving a huge system of sparse linear equations with polynomial coefficients. The coefficients generally become rational functions, i.e. fractions of polynomials, when solving the linear system via (variants of) Gaussian elimination.

Because of the complex nature of the coefficients, they need to be stored in a special form, and most importantly, the coefficients need to be periodically simplified when solving the linear system. The simplifications include e.g. collecting similar terms in polynomials, writing sums of fractions as a single fraction with a common denominator, and simplifying the numerator and denominator by computing polynomial greatest common denominator (GCD). Without the simplifications, arithmetic operations on the coefficients will take more and more time, and their storage will require more and more memory, eventually making performance unacceptable.

In this paper we consider programs (either standalone programs or libraries), called *simplifiers*, which are used to perform all necessary simplifying transformations of rational function coefficients. The list of simplifiers considered in this paper is: CoCoA [3, 4], Fermat [5], FORM [6], GiNaC [7], Macaulay2 [8], Maple [9], Maxima [10], Nemo [11], PARI/GP [12], Symbolica [13], and Wolfram Mathematica [14]. These programs are compared for three different sets of input data: a large set of rational functions in one variable, a large set of rational functions in three variables, and a set of a few dozen huge rational functions whose lengths range from tens of thousands to several hundred thousand characters when printed. The main performance indicators for comparison are the time spent on simplification and the amount of memory needed.

At the end of the 20th century, the task of IBP reduction was done manually. Later, computer programs appeared that automated and speeded up this process. Some of the publicly available general-purpose programs are: FIRE [15–17], AIR [18], Reduze [19], LiteRed [20], and Kira [21–23]. The public version of FIRE was first published in 2014 and has been used by the scientific community to perform cutting-edge calculations, e.g. in Refs. [24–26]. Initially, Fermat was the only simplifier used by the C++ version of FIRE with the use of the gateToFermat library by M. Tentukov. In this work, several more simplifiers are connected to FIRE for the first time, through the standalone C++ library FUEL which provides access to the simplifiers.

FIRE can run both on desktop computers and on specialized nodes with 32 or more computing cores and more than 1.5 TB of RAM, on 64-bit versions of the Linux operating system. Program running time and the required amount of RAM depend on the complexity of the task, and the running time can be up to several months for real-world tasks, of which up to 95% can be spent exclusively on the simplification of rational function coefficients when solving linear systems. In this regard, it is important to find the best programs for simplifying the coefficients, which would allow us to optimize this part of FIRE's performance.

There are many other performance considerations relevant for IBP reduction computations with the Laporta algorithm while keeping analytic dependence on kinematic and spacetime dimension variables. Such considerations include e.g. the ordering of integrals and ordering of equations [21, 27], selection of IBP identities and Lorentz-invariance identities [15, 28], the use of reduction rules with abstract propagator powers (see e.g. [20, 29]), the choice of master integral bases that avoid spurious singularities [30, 31], block triangular form [32, 33], syzygy equations [27, 34–37] and the related numerical unitarity method [38–40]. In this work, however, we focus exclusively on the simplification of rational functions in the process of solving linear systems.

The FUEL library from this work is available from the following git repository: <https://bitbucket.org/feynmanIntegrals/fuel/src/main/>.

2. Problem statement. The purpose of this work is to select and test existing third-party programs for simplifying rational functions, and develop a C++ library FUEL for accessing the simplification functionality. The third-party programs under consideration should be able to simplify complicated expressions and be compatible with the Linux operating system. The programs must be tested for the correctness of simplification.

In order to achieve the goal, it is necessary to solve the following tasks:

- Find programs, or *simplifiers*, that meet the requirements, and write the FUEL library for accessing the simplifiers from C++.
- Test and compare the simplifiers in terms of rational function simplification performance, and select the best ones.
- Connect FIRE with these simplifiers via FUEL to perform IBP reduction computations, and check if they work correctly.

Caveat: our results **should not be** interpreted as a performance comparison between computer algebra systems for polynomial-oriented tasks. Rather, we test the computer algebra systems for the overall performance for tasks similar to IBP reduction when interfacing with FIRE. In particular, efficient bi-directional transfer, i.e. **parsing, printing, and transferring** of expressions in a text format is often an important performance bottleneck, and certain programs can be uncompetitive even with their inherent excellent simplification speed.



3. Simplifiers, input data, and connecting to FIRE.

3.1. Overview. In general, when solving a linear system with polynomial coefficients, e.g. by Gaussian elimination, an intermediate coefficient to be simplified is a sum of fractions, whose numerator and denominator can also contain fractions which, in turn, contain polynomials in their numerators and denominators, as written schematically in Eq. (1):

$$\sum_k \frac{\frac{\text{Poly}_{k,1} \cdot \text{Poly}_{k,3}}{\text{Poly}_{k,2} \cdot \text{Poly}_{k,4}}}{\frac{\text{Poly}_{k,5}}{\text{Poly}_{k,6}}}, \quad \text{Poly}_{k,j} = \sum_i q_i \cdot x_1^{n_{i1}} \cdot x_2^{n_{i2}} \cdot \dots \cdot x_m^{n_{im}}, \quad q_i \in \mathbb{Q}, x_h \in X, n_{ih} \in \mathbb{Z}_0^+. \quad (1)$$

A simplification is considered successful if the result is a polynomial or a fraction where the numerator and denominator are polynomials without a nontrivial GCD. The polynomials have to be either in the expanded form or in some nested form such as the Horner form. Throughout the most of this study, we did not consider other forms such as the factorized or partial-fractioned ones for polynomials and rational functions¹. There are three additional characteristics of the problem arising from IBP reduction by FIRE. First, rational numbers can be as large as desired, that is, the simplifier must support arbitrary-precision integer and rational number arithmetic. Second, the set of possible variables is known in advance (which are kinematic variables and the spacetime dimension variable). Third, the maximum level of nested fractions does not exceed three. The second and third points are important for simplifiers that require these parameters to be passed in advance before actual calculations.

3.2. Connecting with the simplifier.

§3.2.1. Method 1: pipe communications. In order to connect the simplifier, the *fork-exec* technique, popular on Unix systems, is used by FUEL to first create a copy of the current process, then run a new executable file in the context of the newly created process. If the simplifier is an executable file or comes with source code from which an executable file can be built, there is no need to do anything apart from downloading or compiling the simplifier before calling *exec*. If the simplifier is a library, we write wrapper code to use this library and then compile the code into an executable. This technique is as universal as possible, that is, it is applicable to almost any program written in any programming language.

Communication with the spawned process is done through two specially created *pipes*, the first of which is used by the parent program (e.g. FIRE) to write messages to the simplifier, and the second of which is used to send messages in the opposite direction, both in the text format. These techniques make it possible to create a single universal interface for the simplifiers and separate their code from the main program which creates the rational function coefficients (e.g. from creating linear systems of IBP equations). The final procedure for connecting the new simplifier consists of the following steps:

1. Determine which command-line arguments should be passed when *exec* is called.
2. Communicate the rational function to be simplified to the simplifier in an appropriate syntax understood by the simplifier. This triggers the simplifier to parse the expression into an internal representation, simplify the expression, and prints out the result which is then sent back as the response.
3. Write the code for parsing the simplifier response.
4. Determine which commands to pass to the simplifier process in order for it to terminate correctly.

The second and third items are the most time-consuming to implement, because they are unique to each of the pluggable simplifiers, and it was necessary for us to study their documentation, examples, and sometimes advice on the Internet.

§3.2.2. Method 2: library interface. Some of the simplifiers are alternatively exposed as C++ libraries to be used directly by FUEL without pipe communications with a separate executable. In this mode of operation, every input expression is supplied as a string to the relevant library function, and a simplified expression represented as a string is returned. Parallel evaluation is possible when the relevant library is thread-safe. Apart from these changes, most of the considerations of Section 3.2.1 still apply. As will be shown later in this paper, a noticeable performance gains are observed by avoiding pipe communications².

¹An exception is Wolfram Mathematica, whose **Together** command may choose to keep part of the expression in factorized form.

²One could also consider exchanging expressions as C++ objects of opaque types defined by the library, to further avoid the

3.3. Connected simplifiers. This is the complete list of simplifiers (and the languages they are written in) which can be accessed by the current version of FUEL: CoCoA (C++), Fermat (C), FORM (C), GiNaC (C++), Macaulay2 (C/C++, Macaulay2), Maple (C, Java, Maple), Maxima (Common Lisp), Nemo (C, Julia), PARI (C), Wolfram Mathematica (C/C++, Java, Wolfram Language). This list contains both open-source and proprietary software solutions. The issue of software licensing is an important concern, because it may affect, for example, the license under which a derived program can be distributed, the right to modify the code, and the right to distribute the code to third parties. Though we initially searched exclusively among open-source programs and libraries, we could not omit widely used proprietary computer algebra systems such as Maple and Wolfram Mathematica.

All simplifiers can be accessed by pipe communications, in which case parallel evaluation is always possible by running multiple processes. Meanwhile, CoCoA and GiNaC, and Symbolica can be accessed as C++ libraries. Among the three, only Symbolica’s library is thread-safe and supports parallel evaluation.

A brief introduction to each simplifier is given in Appendix 1.

4. Benchmark tests with small to moderately large expressions.

4.1. Testing method. The simplifiers under consideration are used as sequential programs when accessed via pipe communications, i.e. the simplifier process does not treat more than one rational function expressions simultaneously. If parallel evaluation is desired, the user should spawn the required number of simplifier processes and organize parallel sending and receiving of rational function expressions³. For simplifiers accessed as C++ libraries, parallel evaluation is possible only when the library is thread-safe (which is the case for Symbolica, but not CoCoA or GiNaC).

In order to test how fast the simplifiers work, a special benchmark program was written. It reads rational function expressions from a data file into memory, then, with the help of a pre-selected simplifier, processes them individually: sending each expression to the simplifier, waiting for a response, receiving and parsing the response, and finally printing out the total time taken to simplify all expressions from the file.

The initialization of the simplifier (creating a process, loading its context and, for some simplifiers, passing some configuration parameters) is done once at the very beginning and takes at most several seconds. Then the process can in principle run for many hours, so the initialization cost of the simplifier is not included in the benchmark results.

Information about memory usage by the simplifier process was collected by the utility program in Ref. [41], which monitors the memory usage every half a second.

The main machine used for testing has the AMD Ryzen 7 3750H CPU with base frequency 2.3 GHz, boost frequency 4.0 GHz, and 24 GB of RAM.

4.2. Description of test data. Testing was performed on three sets of input data. The rational function expressions from the first set have one variable, and the rational function expressions from the second set have three variables. The number of variables is important because it can significantly affect the running time of some simplifiers. While the first two sets consist of small expressions, the third one is made up of moderately large expressions; the expression size is also an important parameter that directly affects running time. The Table 1 shows the main quantitative characteristics of the utilized sets: the minimum, maximum, and average length of expressions, and the number of expressions in the file.

An even larger expression will be considered eventually, but in a very different setting in Section 5.

Table 1. Parameters of sets with rational function expressions, on which testing was conducted.

No.	No. of variables	Min. expr. length	Max. expr. length	Avg. expr. length	Number of coefs.
1	1	1	5341	~ 29	692584
2	3	1	2133	~ 33	971330
3	2	232971	465943	~ 310628	12

overhead of conversion from/to strings. However, this would not be compatible with the current version of FIRE which relies on database storage that accepts strings as data.

³For example, FIRE has one main execution thread and several additional threads (named FLAME) that communicate with simplifiers.



4.3. Results for running time. The results obtained for the running times of the simplifiers, *as accessed from FUEL*, on different test sets on the two machines are presented in Table 2. All values are specified in seconds and are rounded to one decimal place. We stress again that the running time includes not only the simplification itself, but also parsing, printing, and possibly pipe communications, so the results **should not be** considered as indicators of the inherent quality of the tested simplifiers, especially when they are used in workflows different from ours.

Table 2. Simplifier running times (in seconds) for each of the machines for three sets of input data. A simplifier is accessed via pipe communications unless the library mode is used as indicated by “(lib)”.

Set. No. \ Simplifier	1	2	3
CoCoA	70.6	127.5	84.5
CoCoA (lib)	56.2	100.3	83.3
Fermat	13.9	15.6	1.8
FORM	62.0	107.2	1967.8
GiNaC	21.7	43.3	4.6
GiNaC (lib)	11.9	27.2	4.5
Macaulay2	81.9	244.7	-
Maple	192.1	276.0	85.4
Maxima	106.6	182.8	10.9
Nemo	19.6	35.5	3.4
PARI/GP	10.3	18.3	736.2
Symbolica	10.3	16.5	1.8
Symbolica (lib)	1.9	4.2	1.8
Wolfram Mathematica	349.4	882.0	44.2

We have checked that on different machines, though the absolute running times are different, the relative performance of different simplifiers remains largely unchanged.

Within each set of expressions, the simplifiers are divided into three groups, from the ones with the best performance to the ones with the worst performance: the first group includes those whose running times differ from the minimum on a given set by no more than five times, the second group includes those whose times differ by no more than 10 times, and the third group includes the rest. For clarity, the cells of the Table 2 are colored according to this division: the first group in green, the second in yellow, and the third in red.

Let us now consider each of the test sets in more detail:

1. The first set is characterized by the fact that there is not more than one variable in each of the expressions. Based on performance in simplifying expressions from this test set, the first group includes Fermat, GiNaC, Nemo, PARI/GP, and Symbolica, the second group involves CoCoA and FORM, and the third one includes all others: Macaulay2, Maple, Maxima, and Wolfram Mathematica. Notice that for CoCoA, GiNaC, the library modes run faster but fall into the same groups as their pipe versions.
2. The second set differs from the previous one in that the expressions can now contain up to three variables. As can be seen from the results, the distribution by groups has not changed compared to the distribution for the first set.
3. The third set of expressions is very different from the previous two sets, in that the average expression length has increased by a factor of about 10000, although there is a smaller total number of expressions to keep the total running time manageable. We comment on the performance of a few simplifiers: CoCoA has moved from the second group to the third, showing a slight deterioration. FORM and PARI/GP have the poorest performance (as usual, with the caveat that only in the workflow considered here). While in the previous sets FORM was about 20 to 30 times worse than the best simplifier, and PARI/GP was about 5 times worse than the best one, now FORM is worse than the best-performing simplifier by about

1100 times, and PARI / GP by about 410 times. It turns out that for these simplifiers, the expression length is a defining characteristic and with its growth their speed rapidly degrades. We have not been able to test Macaulay2 on this set due to technical problems with pipe communications. For the third set, the division into groups is as follows: the first includes Fermat, GiNaC, Nemo, and Symbolica, the second includes Maxima, and the third includes CoCoA, FORM, Macaulay2, Maple, PARI/GP, and Wolfram Mathematica.

To draw a conclusion based on these time measurements for simplifying small to moderately large expressions, we apply the following heuristic: if a simplifier is in the first group (i.e. the fastest group) for all three test sets, then we will call it “best”. If it is in the third group, that is, in the group with the slowest performance, then we will call it “bad”. If neither is the case, we call it “good”. It is important to stipulate that these labels should be understood only in conjunction with the phrase “for this class of tasks”. According to the results of testing, the best simplifiers are Fermat, GiNaC, Nemo, and Symbolica, the group of good ones is empty, and the bad ones are CoCoA, FORM, Macaulay2, Maple, Maxima, PARI/GP, and Wolfram Mathematica. When choosing a program or library for simplifying polynomial expressions, e.g. for IBP reduction computations with FIRE, you should first choose from the “best” programs; if for some reason none of them is suitable, only then one should consider others. Furthermore, as was mentioned previously, while the library mode (when available) of a simplifier always runs faster than its pipe mode in our single-threaded benchmark, only Symbolica supports parallel evaluation when used in the library mode.

It is also noted that two proprietary programs, Maple and Wolfram Mathematica, are grouped with the slowest simplifiers for all the three test sets. This suggests that they are unlikely to be good candidates for use as simplifiers in FIRE runs. However, it would be misleading to conclude that these two programs are inherently poor in simplifying rational functions in general, as they perform well in a different benchmark problem in Section 5 which mainly measures inherent simplification time with less overhead in other tasks like parsing text.

Note that for challenging IBP problems, significantly larger expressions can be encountered. Therefore, in the next section, we will supplement the picture we have by additional tests involving a huge expression, while setting the tests in a different manner to shed light on computational overhead unrelated to the simplification itself.

5. Additional tests with huge expression with low parsing overhead. The CPU time consumed by a simplifier process consists of three parts:

1. Parsing the mathematical expression in a text format passed from an external program such as FIRE.
2. Simplifying the expression.
3. Printing out the expression.

The “inherent performance” of the simplifier is measured by the time spend on part (2) above, but depending on usage pattern, this may not be the performance bottleneck. For example, part (1), the parsing process, can often be a bottleneck, considering that when performing IBP reduction with FIRE, hundreds of thousands of expressions may be sent to the simplifier, some of which could be rather small and not inherently difficult to simplify. We will not carry out a full investigation of this issue. However, to shed light on the impact of parsing performance, we present results from another set of test data, where the task is to simplify a single expression in 6 variables:

$$\frac{(a + b + c + d + f + g)^{14} + 3}{(2a + b + c + d + f + g)^{14} + 4} - \frac{(3a + b + c + d + f + g)^{14} + 5}{(4a + b + c + d + f + g)^{14} + 6}. \quad (2)$$

The expression is given in the test data file as the following line:

$$((a+b+c+d+f+g)^{14+3})/((2*a+b+c+d+f+g)^{14+4}) - ((3*a+b+c+d+f+g)^{14+5})/((4*a+b+c+d+f+g)^{14+6}).$$

The time taken to parse this short expression is negligible, but simplification of this expression, involving e.g. expanding polynomials and finding polynomial GCDs, is computationally demanding due to high powers of sub-expressions⁴. Not all simplifiers finish the test in reasonable time; for those that do finish, the running times are summarized in Table 3 below.

This test is drastically different from the tests in Section 4 since it artificially involves negligible parsing overhead, while the simplification itself is very demanding. The test still includes the time taken for the

⁴Note that such expressions do not arise from FIRE: even though high-degree expressions can be generated when solving IBP linear systems, the intermediate expressions are always simplified so that polynomials are in an expanded form or a nested form, and therefore there will be no explicit appearances of a single expression raised to a high power.



Table 3. Time taken, in seconds, by various simplifier to run the test of this section. The numbers are rounded to the nearest integer, or one decimal place if it is less than 10. Only 6 simplifiers are shown in the table. The remaining ones, when accessed from FUEL, are not able to finish the test within 1200 seconds. *For Symbolica and CoCoA, the pipe mode has a similar performance as the library mode for this test since the parsing overhead is low.

Simplifier	Time (seconds)
Symbolica* (lib)	5.2
Nemo	6.9
Maple	7.9
Fermat	98.3
Maxima	112.8
Wolfram Mathematica	168.8
CoCoA* (lib)	365

simplifiers to print out the results to be read by the simplifier, but printing usually has a smaller CPU footprint than parsing when large expressions are involved⁵. The results are also very different from those discussed in Section 4. For example, Maple is now among the most performant programs in this test, either because it suffered from significant parsing overhead in previous tests or because Maple may have a relative advantage in simplifying very large rational expressions.

6. Memory usage. An important statistic is the maximum memory usage, which determines whether or not the *out-of-memory killer* of the operating system will terminate the simplifier. This is especially a concern in the process of simplifying complicated expressions. We compare the maximum memory usage for the simplest test set 1 in Section 4 and the most complicated expression in Section 5. The results are given in Table 4, only for simplifiers that can finish the latter test in reasonable time⁶. All values are rounded to integers.

Table 4. The maximum amount of used RAM by simplifiers when accessed via FUEL, in megabytes, for three sets of input expressions. The “small” test set refers to set 1 in Section 4, and the “huge” test set refers to the expression in Section 5.

Stat. Simplifier	Max	
	Small	Huge
CoCoA	3	242
Fermat	21	22
Maple	20	34
Maxima	931	953
Nemo	409	410
Symbolica	2	20
Wolfram Mathematica	121	225

Except for CoCoA and Wolfram Mathematica, the simplifiers in the table exhibit memory usage that depends rather mildly on the complexity of the expressions. One should keep in mind that even standard laptops now have at least 8GB of RAM and server systems can have several hundred GB of memory. For IBP reduction with FIRE, the memory consumption from storing a large number of reduction rules (in which the rational functions are already simplified) can be a much more serious concern.

⁵For Nemo, we found that the running time is reduced by only about 15% when the printing of the result is turned off. For Mathematica and Maple, we found that the time used in printing is completely overwhelmed by the rest of the running time.

⁶For the convenience of measurement, we have also restricted our attention to simplifiers that can run as separate processes accessed via pipes.




Figure 1. The double box diagram to be tested for IBP reduction in FIRE6 calling various different simplifiers.

7. Tests with FIRE.

7.1. A simple example. We run the double box IBP reduction example in FIRE6. This IBP problem is very simple by current standards and should be considered as a preliminary test, as the main focus of this work is presenting FUEL and standalone benchmark tests. The double box diagram is shown in Fig. 1. We reduce a rank-2 tensor integral with numerator

$$(k_2 + p_1)^2(k_1 - p_3)^2$$

using only one worker thread. The statistics printed out at the end of FIRE runs are presented in Table 5. FIRE runs involve two stages, forward elimination and backward substitution. While back substitution consumes a small percentage of the total running time in this simple problem, it can become the dominant part in more complicated IBP problems. Therefore we have shown the “substitution time” as a separate column in Table 5. In this test, the performance of the simplifiers relative to each other is very similar to the situation in test

Table 5. Performance of various simplifier when used by FIRE to reduce a rank-2 tensor integral for the massless two-loop double box. The running times are in seconds, while the memory usage is in units of MB.

Simplifier	Total Time	Substitution time	Memory usage (FIRE + simplifier)
Symbolica (lib)	14.6	1.1	15.6
Fermat	19.8	1.8	23.0
Pari/GP	21.6	2.3	14.1
Nemo	32.8	2.4	431.7
GiNaC (lib)	27.6	6.0	14.6
CoCoA (lib)	73.2	4.7	16.0
FORM	78.5	5.8	14.4
Maxima	131.2	6.7	955.8
Macaulay	152.5	12.1	350.9
Maple	177.6	6.5	105.1
Wolfram Mathematica	581.0	22.2	146.1

sets 1 and 2 in Section 4. Based on the data for test 3 in Section 4, it is likely that the situation can change dramatically for highly demanding IBP reduction problems.

7.2. A complicated example. Let us turn now to consideration of a demanding IBP problem for non-planar double box integrals with massless internal lines and off-shell external lines. Three of the four external lines have the same virtuality. This problem of computing off-shell three-particle form factors arises in $\mathcal{N} = 4$ super-Yang-Mills theory as a natural extension of the similar problem of calculating two-particle form factors, which was studied previously in [42, 43]. The corresponding diagram is shown in Fig. [2].

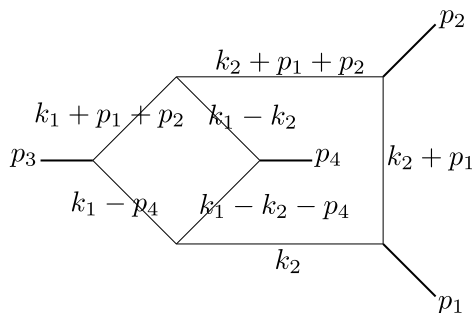


Figure 2. A nonplanar double box diagram topology with massless internal lines and off-shell external lines.

The kinematics is

$$\begin{aligned}
 p_1^2 &= p_2^2 = p_3^2 = -m^2, \\
 (p_1 + p_2)^2 &= -u, \\
 (p_2 + p_3)^2 &= -v, \\
 (p_1 + p_3)^2 &= -w.
 \end{aligned}$$

We perform the IBP reduction of top-level integrals with the following three different numerators:

1. the unit numerator 1;
2. $(k_1 + p_1)^2$;
3. $(k_2 - p_4)^2$;
4. $(k_1 + p_1)^2(k_2 - p_4)^2$.

There are 97 master integrals after IBP reduction. As a form of parallelization, FIRE allows the user to set all master integrals excepting one equal to zero and compute the coefficient of the one chosen master integral. For the purpose of benchmarking, we only calculate the coefficient of the bottom-level “sunrise” master integral which contains the propagators $(k_1 - p_4)^2$, $(k_1 - k_2)^2$ and $(k_2 + p_1)^2$. Five variables are involved in the IBP reduction, including the spacetime dimension d and the kinematic variables m, u, v, w . The large number of variables makes the simplification of rational functions computationally demanding during the IBP reduction. The time needed to complete the task for Fermat, Nemo, and Symbolica are shown in Table 6⁷.

Table 6. Performance of various simplifier when used by FIRE to reduce sample integrals for the nonplanar double box with massless internal lines and off-shell external lines, involving four kinematic variables and the spacetime variable in rational function simplification.

Simplifier	Total Time (× 1000 seconds)
Symbolica (lib)	8.9
Nemo	11.9
Fermat	109.7

We can see that Symbolica and Nemo offers significant speedups compared to Fermat, as in the case of the artificial benchmark in Section 5 outside of FIRE.

8. Conclusions. In light of ongoing efforts to improve the performance of integration-by-parts reduction for complicated Feynman integral calculations, we have presented a new C++ library FUEL for simplifying rational function expressions. As a standalone library, FUEL can also potentially find applications in other areas. Under a universal interface, FUEL allows a flexible choice of *simplifiers*, i.e. existing computer algebra programs or libraries, as the underlying computation engine. FUEL grew out of FIRE’s original interface to Fermat, the latter being a computer algebra system written by Robert Lewis. FUEL has two modes of operation.

⁷A different computer is used for this benchmark since the previous tests were carried out at an earlier time. The CPU used here is the Intel Xeon Gold 6240 Processor (24.75M Cache, 2.60 GHz). We allowed up to 16 simplifier processes to run in parallel, but due to limitations of parallelization capabilities when only a single master integral is targeted, only a single thread is active for almost the entire duration of the run, regardless of the utilized simplifier.

The first of them is the pipe mode, originally used in FIRE, based on inter-process communication over unix pipes, sending and receiving text expressions to a third-party simplifier program in a text format. The other mode is the library mode, i.e. directly linking with third-party libraries without spawning child processes. With both mode available, the setup allows for maximum flexibility in connecting with any simplifier written in any programming language. In the pipe mode, parallel computation is achieved by running multiple processes of the same simplifier (or even different simplifiers catering to different types of expressions, experimentally supported by FUEL). In the library mode, parallel computation is possible when the third-party library is thread-safe (which is the case only for Symbolica so far). Good performance under the pipe mode requires the simplifier to be fast both in the key task of simplifying mathematical expressions and in overhead tasks such as parsing text inputs, which makes certain simplifiers (such as recent versions of Maple) uncompetitive for our purpose even when they have a reputation for being fast manipulations of polynomials and rational functions.

In the current version of FUEL, we have implemented connections with 11 different simplifiers, including CoCoA (pipe or library mode), Fermat, FORM, GiNaC (pipe or library mode), Macaulay2, Maple, Maxima, Nemo, PARI/GP, Symbolica (pipe or library mode), and Wolfram Mathematica. Nemo has been augmented by a dedicated Julia package we wrote (distributed with FUEL) to enable fast parsing of mathematical expressions. Artificial benchmark tests with small to moderately large expressions are presented in Section 4. Symbolica is the fastest simplifier (or in a tie with the fastest simplifiers) for all the test sets, while Fermat, Nemo and GiNaC are also consistently among the fastest simplifiers. Pari/GP is very fast for the first two sets of test data involving shorter expressions (which likely mimic less demanding FIRE runs), but performs very poorly in the third data set involving moderately large expressions.

An additional special-purpose test is presented in Section 5. Compared with the main tests discussed above, this test minimizes the overhead of text parsing but is extremely computationally intensive in the simplification itself. Here Fermat has dropped to the third place in the ranking of the fast programs, led by Symbolica, Nemo, and Maple. Both Pari/GP and GiNaC (among others) have failed to complete the test before the 20-minute timeout. We plan to explore using more than one simplifier in a single C++ program, e.g. FIRE, given their different performance characteristics for different problems.

A private experimental version of FIRE has been linked with FUEL to perform IBP reduction of Feynman integrals. We have first demonstrated a very simple IBP reduction example for the two-loop double box, which can be completed by FIRE with any of the 10 connected simplifiers. The time required by the run for each simplifier has been tabulated, and the results are broadly consistent with those from the simpler test sets in the artificial benchmarks of Section 4. Next, we have tested the IBP reduction performance for a non-planar double box family of integrals which involve four kinematic variables besides the spacetime variable. In this case, Symbolica and Nemo significantly outperform Fermat, while GiNaC and Pari/GP cannot finish the test in reasonable time, mirroring the demanding standalone benchmark of Section 5. We leave further improvements to further work, and a first application to a concrete physics problem using our software will soon be carried out by other authors.

License statement: for the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) license to any Author Accepted Manuscript version arising from this submission.



Appendix 1

Overviews of connected simplifiers

8.1. CoCoA. CoCoA [3] is a computer algebra system for computing in polynomial rings. The development of the system began in 1987 in Pascal, hence its Pascal-like syntax, and it was later rewritten in C, and yet a little later a C++ library, CoCoALib [4], appeared, and the latter library was used in our work. CoCoA allows to perform calculations in rings of polynomials of many variables with rational or integer coefficients, as well as over ideals of these rings. The user can redefine the polynomial ring used as well as various homomorphisms for converting elements from one ring to another. According to CoCoA’s authors, the Gröbner basis is used as the key mechanism for efficient computations in commutative algebra.

In order for CoCoA to handle a factor of the form Eq. (1) passed to it, we need to specify in the C++ constructor a field that it belongs to. To do this, we set the appropriate *ring of integers*, *ring of fractions* and *ring of polynomials*, combining and substituting one into the other to get the desired field. Then it is possible to supply a string representation of the rational function and get a simplified representation from it.

The fully expanded form is utilized for polynomial representations in our usage of CoCoA. In this paper, we use the version CoCoA 0.99715.

8.2. Fermat. Fermat [5] is a computer algebra system, developed by Robert Lewis, aiming to be fast and memory efficient, covering “arithmetic of arbitrarily long integers and fractions, multivariate polynomials, symbolic calculations, matrices over polynomial rings, graphics, and other numerical calculations”. Fermat has influenced research in fast rational function arithmetic in computer algebra [44]. Until recently, Fermat was the only simplifier used by the C++ version of FIRE. Fermat is also the main simplifier in two other IBP reduction programs, Kira and Reduze.

The output of Fermat expresses polynomials in the Horner form. In this paper, we use the version Fermat 5.17.

8.3. Form. FORM [6] is a computer algebra system, which is sometimes called a system for formula conversions. It has been in development since 1984, and its original goal was to simplify calculations in quantum field theory. FORM is written in C and accepts input in a special programming language, which is then interpreted and executed. In other words, FORM is not interactive but operates as a batch-processing program. FORM’s language has many features: it has an advanced preprocessor with more than sixty commands, several types of variables: symbols, vectors, indices, functions, sets, more than a hundred commands controlling the execution, output and properties of variables, and more than eighty functions. Besides the regular version of the program, there are also two parallelized versions: ParFORM, which runs on a system with independent nodes, each using its own processor, memory and disk, and TFORM, which uses POSIX threads to better expose multiprocessor capabilities on shared-memory machines. FORM is distributed under the GNU GPL license.

For some special cases it may be necessary to override the standard settings that control how FORM works, such as the maximum size of the substitution tree, the maximum size of a term that does not require additional allocations, the size of the I/O buffers, and other settings. In order to simplify a large expression, we need to redefine several settings, otherwise the program would stop due to insufficient memory.

The fully expanded form is utilized for polynomial representations in our usage of FORM. In this paper, we use the version FORM 4.2.1.

8.4. GiNaC. GiNaC [7] is a C++ library for computer algebra, initially designed for Feynman diagram calculations. On the contrary, to many other computer algebra systems which come with their own proprietary interactive languages, GiNaC emphasizes programmatic use, extensibility and interoperability with other programs within a statically-typed compiled language (C++). Besides features commonly found in most systems, such as big integers and polynomial simplification, GiNaC offers functionalities useful for Feynman diagram calculations, such as handling of expressions involving Lorentz, Dirac, and color indices. In high energy physics research, GiNaC is perhaps most well known for its support for numerical evaluations of special functions known as multiple polylogarithms. A fork of GiNaC, PyNaC [45], was used as a core component of SageMath [46], a flagship open-source computer algebra system.

The fully expanded form is utilized for polynomial representations in our usage of GiNaC. In this paper, we use the version GiNaC 1.8.2.

8.5. Macaulay2. Macaulay2 [8] is system for computation in algebraic geometry and commutative algebra, covering functionalities such as Groebner bases, free resolutions of modules, Betti numbers, primary

decomposition of ideals, etc. Macaulay2 has been used in research in applying computational algebraic geometry to IBP reduction [47].

The fully expanded form is utilized for polynomial representation in our usage of Macaulay2. In this paper, we use the Macaulay2 version 1.19.1+ds-6.

8.6. Maple. Maple is a general-purpose computer algebra system. It is widely used in high energy physics. For example, the IBP reduction program AIR [18] is written in Maple. In comparison with its competitor Wolfram Mathematica, Maple offers a more conventional ALGOL/C-like programming language. High-performance algorithms relevant for polynomials and rational functions are continually and actively developed in Maple (see e.g. [44, 48, 49]). To simplify rational functions, Maple utilizes the *normal* function with the option *expanded* to prevent storing polynomials in factorized form. Therefore the fully expanded form is used for polynomial representation in our usage of Maple.

In this paper, we apply the version Maple 2022.

8.7. Maxima. Maxima [10] is a computer algebra system, a descendant of Macsyma, which allows many different operations on symbolic and numeric expressions. Maxima is self-described as a “fairly complete computer algebra system” and can be used to e.g. differentiate, integrate, solve Laplace transforms, and construct graphs. Maxima, like its ancestor, is written in Lisp. It is distributed under the GNU GPL license. SageMath [46] uses Maxima internally for certain nontrivial computations such as symbolic integration and taking limits.

Maxima has rich functionality for simplifications; it offers several functions for user to choose from: *rat*, *ratsimp*, *fullratsimp*, *radscan* and many flags that affect how the functions work. Some functions are mainly intended to simplify rational expressions, while others can simplify expressions containing logarithmic, exponential and power functions. Some perform simplification once over the expression, while others do it until the resulting expression stops to change.

In addition, several flags have been included to make it easier to parse the rational functions simplified by Maxima: *display2d:false* disables 2D output, *stardisp:true* removes unnecessary multiplication signs, and *nolabels:true* allows to remove unnecessary I/O labels for entered and resulting expressions.

A partial expanded form is utilized for polynomial representation in our use of Maxima. In this paper, we apply the version Maxima 5.45.1.

8.8. Nemo (with a custom parser and printer). Nemo [11] is a computer algebra system for the Julia programming language [50], and it aims to “provide highly performant commutative algebra, number theory, group theory and discrete geometry routines.” It provides a fast Julia interface to C/C++ libraries such as FLINT [51]. FLINT provides efficient operations for polynomials over a variety of number fields such as rational numbers and prime fields. Benefiting from the EU-funded OpenDreamKit project for open-source computer algebra, FLINT gained fast code for multivariate polynomials. Meanwhile, the Julia code in Nemo provides, among other functionalities, operations for rational functions that build upon polynomial operations of FLINT. In the current version of FUEL, we always use Nemo’s sparse multivariate polynomials and associated rational functions. In the univariate case, specialized routines from Nemo and FLINT can be faster but have not been used in our work due to a lack of implementation effort on our side.

A previous internal version of FUEL calls Nemo from a Julia REPL session (i.e. an interactive user session), and the performance was poor due to the overhead of parsing. The parser of the Julia REPL is designed to process arbitrary syntax in the Julia language and is relatively slow for our special purpose of parsing rational function expressions. Fortunately, taking advantage of Julia’s JIT compilation, we are able to write a fast parser for mathematical expressions based on a variation of the well-known shunting-yard algorithm. The parser is included in a Julia package, *RationalCalculator*, bundled with FUEL. Additionally, the aforementioned package supports printing out calculation results in a format that is not human-readable but instead optimized for transfer of expressions between the simplifier and FUEL. Human readable output can be re-enabled by calling a routine in FUEL, e.g. when FIRE writes out IBP reduction tables, so that the final IBP reduction table from the FIRE run is unaffected.

The fully expanded form is utilized for polynomial representation in our usage of Nemo. In this paper, we apply the versions Nemo 0.33.1 and Julia 1.8.5.

8.9. PARI/GP. PARI/GP is a computer algebra system focused on number theory, developed at the University of Bordeaux. It can compute “factorizations, algebraic number theory, elliptic curves, modular forms, L functions”, etc. [12]. PARI is a C library, while GP is the front-end that allows interactive use. We utilize GP as the expression simplifier.



A Horner-like form is used for polynomial representation by PARI/GP. In this paper, we apply the version GP 2.13.3.

8.10. Symbolica. Symbolica is a new computer algebra system being developed by Ben Ruijl [13], which aims to be a modernized incarnation of FORM [6], e.g. by using iterators in modern programming languages to express the stream processing of large expressions that cannot be packed into RAM. Symbolica is written in Rust, with interfaces for C/C++ and Python. Although Symbolica is at the early stage of development, it already contains highly optimized code for polynomial GCD and rational function arithmetic.

The fully expanded form is utilized for polynomial representation in our use of Symbolica. In this paper, we apply the version a8f72f (the commit hash) from the GitHub repository <https://github.com/benruijl/symbolica.git>.

8.11. Wolfram Mathematica. Wolfram Mathematica is the most widely used general-purpose computer algebra system as of today, at least in theoretical high energy physics research. It was initially developed by Stephen Wolfram, with influences from Maxima and Wolfram’s earlier system, SMP. Mathematica offers a high-level language emphasizing functional programming and term rewriting (called *Replacement* in Mathematica). As of today, Mathematica encompasses a huge range of functionalities in symbolic and numerical computing. Many software packages and research data in high energy physics are published in Mathematica’s formats. FIRE, even when utilized in the C++ mode for the main computation, uses Mathematica for pre-processing user-supplied integral family information and various post-processing tasks such as loading reduction tables and finding symmetry rules relating master integrals. We apply the *Together* function in Mathematica to simplify rational functions, and the output may consider either factorized or expanded polynomials.

In this paper, we use the version Mathematica 13.0.

References

1. K. G. Chetyrkin and F. V. Tkachov, “Integration by Parts: The Algorithm to Calculate β -Functions in 4 Loops,” Nucl. Phys. B **192** (1), 159–204 (1981). doi 10.1016/0550-3213(81)90199-1.
2. 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.
3. J. Abbott, A. M. Bigatti, and L. Robbiano, CoCoA System: Computations in Commutative Algebra. <http://cocoa.dima.unige.it>. Cited October 4, 2023.
4. J. Abbott and A. M. Bigatti, CoCoALib: A C++ Library for Doing Computations in Commutative Algebra. <https://cocoa.dima.unige.it/cocoa/cocoalib/>. Cited October 4, 2023.
5. R. H. Lewis, Fermat: A Computer Algebra System for Polynomial and Matrix Computation. <http://home.bway.net/lewis/>. Cited September 30, 2023.
6. J. A. M. Vermaseren, The Main Directory for FORM. <https://www.nikhef.nl/~form/maindir/maindir.html>. Cited September 30, 2023.
7. GiNaC is Not a CAS. <https://ginac.de/>. Cited September 30, 2023.
8. D. R. Grayson and M. E. Stillman. Macaulay2Doc: Macaulay2 Documentation. <https://macaulay2.com/doc/Macaulay2/share/doc/Macaulay2/Macaulay2Doc/html/>. Cited October 4, 2023.
9. Maplesoft: The Essential Tool for Mathematics. <https://www.maplesoft.com/products/Maple/>. Cited October 4, 2023.
10. Maxima: A Computer Algebra System. Version 5.43.2. <https://maxima.sourceforge.io/>. Cited September 30, 2023.
11. C. Fieker, W. Hart, T. Hofmann, and F. Johansson, “Nemo/Hecke: Computer Algebra and Number Theory Packages for the Julia Programming Language,” in *Proc. 2017 ACM Int. Symposium on Symbolic and Algebraic Computation, Kaiserslautern, Germany, July 23–28, 2017* (ACM Press, New York, 2017), pp. 157–164. doi 10.1145/3087604.3087611.
12. PARI/GP Version 2.11.2, Univ. Bordeaux, 2022. <http://pari.math.u-bordeaux.fr/>. Cited October 1, 2023.
13. Symbolica. <https://symbolica.io/>. Cited October 1, 2023.
14. Wolfram Mathematica, Version 13.1. <https://www.wolfram.com/mathematica>. Cited October 1, 2023.
15. A. V. Smirnov and 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.
16. 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.

17. A. V. Smirnov and 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).
18. C. Anastasiou and A. Lazopoulos, “Automatic Integral Reduction for Higher Order Perturbative Calculations,” *J. High Energy Phys.* No. 7, Article Identifier 046 (2004). doi [10.1088/1126-6708/2004/07/046](https://doi.org/10.1088/1126-6708/2004/07/046).
19. A. von Manteuffel and C. Studerus, “Reduze 2 — Distributed Feynman Integral Reduction,” <https://arxiv.org/abs/1201.4330>. Cited October 1, 2023.
20. R. N. Lee, “LiteRed 1.4: A Powerful Tool for Reduction of Multiloop Integrals,” *J. Phys. Conf. Ser.* **523** (1), Article Identifier 012059 (2014). doi [10.1088/1742-6596/523/1/012059](https://doi.org/10.1088/1742-6596/523/1/012059).
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. P. Maierhöfer and J. Usovitsch, “Kira 1.2 Release Notes,” <https://arxiv.org/abs/1812.01491>. Cited October 1, 2023.
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. M. Henn, A. V. Smirnov, V. A. Smirnov, and M. Steinhauser, “A Planar Four-Loop Form Factor and Cusp Anomalous Dimension in QCD,” *J. High Energy Phys.* No. 5, Article Identifier 066 (2016). doi [10.1007/JHEP05\(2016\)066](https://doi.org/10.1007/JHEP05(2016)066).
25. J. Henn, R. N. Lee, A. V. Smirnov, et al., “Four-Loop Photon Quark Form Factor and Cusp Anomalous Dimension in the Large- N_c Limit of QCD,” *J. High Energy Phys.* No. 3, Article Identifier 139 (2017). doi [10.1007/JHEP03\(2017\)139](https://doi.org/10.1007/JHEP03(2017)139).
26. R. N. Lee, A. V. Smirnov, V. A. Smirnov, and M. Steinhauser, “ n_f^2 Contributions to Fermionic Four-Loop Form Factors,” *Phys. Rev. D* **96** (1), Article Identifier 014008 (2017). doi [10.1103/PhysRevD.96.014008](https://doi.org/10.1103/PhysRevD.96.014008).
27. D. Bendle, J. Böhm, W. Decker, et al., “Integration-by-Parts Reductions of Feynman Integrals Using Singular and GPI-Space,” *J. High Energy Phys.* No. 2, Article Identifier 079 (2020). doi [10.1007/JHEP02\(2020\)079](https://doi.org/10.1007/JHEP02(2020)079).
28. R. N. Lee, “Group Structure of the Integration-by-Part Identities and Its Application to the Reduction of Multiloop Integrals,” *J. High Energy Phys.* No. 7, Article Identifier 031 (2008). doi [10.1088/1126-6708/2008/07/031](https://doi.org/10.1088/1126-6708/2008/07/031).
29. 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).
30. A. V. Smirnov and V. A. Smirnov, “How to Choose Master Integrals,” *Nucl. Phys. B* **960**, Article Number 115213 (2020). doi [10.1016/j.nuclphysb.2020.115213](https://doi.org/10.1016/j.nuclphysb.2020.115213).
31. J. Usovitsch, “Factorization of Denominators in Integration-by-Parts Reductions,” <http://arxiv.org/abs/2002.08173>. Cited October 1, 2023.
32. X. Liu and Y.-Q. Ma, “Determining Arbitrary Feynman Integrals by Vacuum Integrals,” *Phys. Rev. D* **99** (7), Article Identifier 071501 (2019). doi [10.1103/PhysRevD.99.071501](https://doi.org/10.1103/PhysRevD.99.071501).
33. X. Guan, X. Liu, and Y.-Q. Ma, “Complete Reduction of Integrals in Two-Loop Five-Light-Parton Scattering Amplitudes,” *Chinese Phys. C* **44** (9), Article Identifier 093106 (2020). doi [10.1088/1674-1137/44/9/093106](https://doi.org/10.1088/1674-1137/44/9/093106).
34. J. Gluza, K. Kajda, and D. A. Kosower, “Towards a Basis for Planar Two-Loop Integrals,” *Phys. Rev. D* **83** (4), Article Identifier 045012 (2011). doi [10.1103/PhysRevD.83.045012](https://doi.org/10.1103/PhysRevD.83.045012).
35. R. M. Schabinger, “A New Algorithm for the Generation of Unitarity-Compatible Integration by Parts Relations,” *J. High Energy Phys.* No. 1, Article Identifier 077 (2012). doi [10.1007/JHEP01\(2012\)077](https://doi.org/10.1007/JHEP01(2012)077).
36. K. J. Larsen and Y. Zhang, “Integration-by-Parts Reductions from Unitarity Cuts and Algebraic Geometry,” *Phys. Rev. D* **93** (4), Article Identifier 041701 (2016). doi [10.1103/PhysRevD.93.041701](https://doi.org/10.1103/PhysRevD.93.041701).
37. J. Böhm, A. Georgoudis, K. J. Larsen, et al., “Complete Integration-by-Parts Reductions of the Non-Planar Hexagon-Box via Module Intersections,” *J. High Energy Phys.* No. 9, Article Identifier 024 (2018). doi [10.1007/JHEP09\(2018\)024](https://doi.org/10.1007/JHEP09(2018)024).
38. H. Ita, “Two-Loop Integrand Decomposition into Master Integrals and Surface Terms,” *Phys. Rev. D* **94** (11), Article Identifier 116015 (2016). doi [10.1103/PhysRevD.94.116015](https://doi.org/10.1103/PhysRevD.94.116015).
39. S. Abreu, F. Febres Cordero, H. Ita, et al., “Two-Loop Four-Gluon Amplitudes from Numerical Unitarity,” *Phys. Rev. Lett.* **119** (14), Article Identifier 142001 (2017). doi [10.1103/PhysRevLett.119.142001](https://doi.org/10.1103/PhysRevLett.119.142001).
40. S. Abreu, F. Febres Cordero, H. Ita, et al., “Planar Two-Loop Five-Gluon Amplitudes from Numerical Unitarity,” *Phys. Rev. D* **97** (11), Article Identifier 116014 (2018). doi [10.1103/PhysRevD.97.116014](https://doi.org/10.1103/PhysRevD.97.116014).
41. A utility to accurately report the core memory usage for a program. https://raw.githubusercontent.com/pixelb/ps_mem/master/ps_mem.py. Cited October 1, 2023.



42. A. V. Belitsky, L. V. Bork, A. F. Pikelner, and V. A. Smirnov, “Exact Off Shell Sudakov Form Factor in $\mathcal{N} = 4$ Supersymmetric Yang–Mills Theory,” *Phys. Rev. Lett.* **130** (9), Article Identifier 091605 (2023). doi [10.1103/PhysRevLett.130.091605](https://doi.org/10.1103/PhysRevLett.130.091605).
43. A. V. Belitsky, L. V. Bork, and V. A. Smirnov, “Off-Shell Form Factor in $\mathcal{N} = 4$ sYM at Three Loops,” <https://arxiv.org/abs/2306.16859>. Cited October 1, 2023.
44. M. Monagan and R. Pearce, “Fermat Benchmarks for Rational Expressionals in Maple,” *ACM Commun. Comput. Algebra* **50** (4), 188–190 (2016). doi [10.1145/3055282.3055299](https://doi.org/10.1145/3055282.3055299).
45. Pynac — symbolic computation with Python objects, Sage Math support library. <https://github.com/pynac/pynac>. Cited October 1, 2023.
46. P. Zimmermann, A. Casamayou, N. Cohen, et al., *Computational Mathematics with SageMath* (SIAM Press, Philadelphia, 2018).
47. Y. Zhang, *Lecture Notes on Multi-Loop Integral Reduction and Applied Algebraic Geometry*. <https://arxiv.org/abs/1612.02249>. Cited October 1, 2023.
48. M. Monagan and R. Pearce, “Poly: A New Polynomial Data Structure for Maple 17,” *ACM Commun. Comput. Algebra* **46** (3/4), 164–167 (2013). doi [10.1145/2429135.2429173](https://doi.org/10.1145/2429135.2429173).
49. M. Monagan, “Speeding up Polynomial GCD, a Crucial Operation in Maple,” *Maple Trans.* **2** (1), Article Number 14457 (2022). doi [10.5206/mt.v2i1.14452](https://doi.org/10.5206/mt.v2i1.14452).
50. J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, “Julia: A Fresh Approach to Numerical Computing,” *SIAM Rev.* **59** (1), 65–98 (2017). doi [10.1137/141000671](https://doi.org/10.1137/141000671).
51. W. B. Hart, “Flint Fast Library for Number Theory,” *Computeralgebra-Rundbrief*, No. 49, 15–17 (2011). <https://fachgruppe-computeralgebra.de/data/CA-Rundbrief/car49.pdf>. Cited October 3, 2023.

Received
September 22, 2023

Accepted for publication
September 25, 2023

Information about the authors

Kirill S. Mokrov — Student, Lomonosov Moscow State University, Faculty of Computational Mathematics and Cybernetics, Leninskie Gory, 1, building 52, 119991, Moscow, Russia.

Alexander V. Smirnov — Dr. Sci., Laboratory head; Lomonosov Moscow State University, Research Computing Center, Leninskie Gory, 1, building 4, 119234, Moscow, Russia.

Mao Zeng — Ph.D., Royal Society University Research Fellow; Higgs Centre for Theoretical Physics, University of Edinburgh, James Clark Maxwell Building, Peter Guthrie Tait Road, EH9 3FD, Edinburgh, United Kingdom.