

FIRE6: Feynman Integral REduction with Modular Arithmetic

A.V. Smirnov^{a,b,*}, F.S. Chukharev^c

^a*Research Computing Center, Moscow State University,
119992 Moscow, Russia*

^b*Institut für Theoretische Teilchenphysik, Karlsruhe Institute of Technology (KIT),
76128 Karlsruhe, Germany*

^c*Faculty of Computational Mathematics and Cybernetics, Moscow State University,
119992 Moscow, Russia*

Abstract

FIRE is a program performing reduction of Feynman integrals to master integrals. The C++ version of FIRE was presented in 2014. There have been multiple changes and upgrades since then including the possibility to use multiple computers for one reduction task and to perform reduction with modular arithmetic. The goal of this paper is to present the current version of FIRE.

Keywords: Feynman diagrams, Multiloop Feynman integrals, Dimensional regularization, Computer algebra, Modular arithmetic, Supercomputers

*Corresponding author

Email addresses: `asmirnov80@gmail.com` (A.V. Smirnov),
`fedor.s.chukharev@gmail.com` (F.S. Chukharev)

PROGRAM SUMMARY

Manuscript Title: FIRE6 Feynman Integral REduction with Modular Arithmetic

Authors: A.V. Smirnov, F.S. Chukharev

Program title: FIRE, version 6 (FIRE6)

Licensing provisions: GPLv2

Programming language: Wolfram Mathematica 6.0 or higher, C++

Computer(s) for which the program has been designed: starting from a desktop PC up to a supercomputer

Operating system(s) for which the program has been designed: Linux 64bit, Mac OS X 10.6 or higher 64bit

RAM required to execute with typical data: depends on the complexity of the problem

Has the code been vectorized or parallelized?: yes

Number of processors used: depending on the mode and the complexity of the task, the program can both run single-threaded on a laptop, or use a supercomputer (tested with up to 2024 cores)

Supplementary material: The article, install instructions, <https://bitbucket.org/feynmanIntegrals/fire>

Keywords: Feynman diagrams, Multiloop Feynman integrals, Dimensional regularization, Computer algebra

CPC Library Classification: 4.4 Feynman diagrams, 4.8 Linear Equations and Matrices, 5 Computer Algebra, 20 Programming and Publication Practice

External routines/libraries used: Wolfram Mathematica [1], Snappy [2], ZStandard [3], KyotoCabinet [4], Fermat [5], LiteRed [6]

Nature of problem: Reducing Feynman integrals to master integrals can be treated as a task to solve a huge system of sparse linear equations with polynomial coefficients.

Solution method: Since the matrix of equations is very specific, none of standard methods of solving linear equations can be applied efficiently. The program approaches solving those equations with a special version of Gauss elimination. In complex cases the direct reduction approach might fail, so the approach with modular arithmetic is used, where the reduction is performed multiple times with different values of variables over large prime number fields, afterwards the coefficients are reconstructed. The data preparation and result analysis is performed in Wolfram Mathematica [1], but the main reduction procedure is written in C++; FIRE compresses data with the use of the Snappy [2] or ZStandard [3] library, stores it on disk with the use of the KyotoCabinet [4] database engine, and performs algebraic simplifications with the Fermat [5] program. The external package LiteRed [6] can

be used to produce additional rules for reduction.

Restrictions: The complexity of the problem is mostly restricted by CPU time required to perform the reduction of integrals and the available RAM. The program has the following limits: maximal number of indices = 22, maximal number of positive indices = 15, maximal number of non-trivial sectors = $128 \times 256 - 3 = 32765$ (global symmetries decrease the number of sectors, indices that cannot be positive do not double the number of sectors). **FIRE6** follows the C++11 standard, so requires gcc 4.8.1 or higher to be compiled, but works with the current gcc 7.3 as well.

Running time: depends on the complexity of the problem

References:

- [1] <http://www.wolfram.com/mathematica/>, commercial algebraic software;
- [2] <https://github.com/google/snappy>, open source;
- [3] <https://github.com/facebook/zstd>, open source;
- [4] <http://fallabs.com/kyotocabinet/>, open source;
- [5] <https://home.bway.net/lewis/>, free-ware with some restrictions for organizations;
- [6] <http://www.inp.nsk.su/~lee/programs/LiteRed/>, open source.

1. Introduction

Feynman integrals are fundamental objects arising in modern elementary particle physics. The problems being considered currently result in millions of Feynman integrals that have to be evaluated. A classical approach is to apply *integration by parts (IBP) relations* [1] (see Chapter 6 of [2] for a review) and reduce all integrals to a smaller set, the so-called *master integrals*¹. The first approach to integration by parts relations was based on solving those equations by hand. This procedure was used in many papers starting from [1] before the appearance of computer codes to solve IBP relations.

Currently there is a number of programs that can perform Feynman integral reduction, some of those are public, some are private. Most of the existing programs are based on the Laporta algorithm [4] (Gauss elimination after choosing an ordering), in particular, the public codes described in [5, 6, 7, 8, 9, 10]. However there were also attempts based on constructing explicit rules when solving IBP relations, both by one of the authors of this paper [11, 12, 13], as well as by R.Lee in his program `LiteRed` [14, 15, 16].

A recent approach to the reduction is based on modular arithmetic. While the direct reduction might be too complex because of the growth of coefficients, one can fix different values of variables and also perform reduction over fields of remainders over large prime numbers instead of rationals. Those numbers fit into machine-size integers, so the reduction can run significantly faster. However, to recover coefficients afterwards one has to run a rather large number of reduction jobs. This approach has already been presented in reduction programs `Finred` [17] and `Kira` [9].

One of the authors of this paper created the program `FIRE` (Feynman integral reduction) about ten years ago. The initial version of `FIRE` [18, 19] was written in `Wolfram Mathematica`, later a C++ version was also published [20].

The goal of this paper is to present the new version of `FIRE` that was developed during recent years. We consider `FIRE6` to be a major release significantly differing from the previous public version. It was already successfully applied in [21, 22, 23, 24, 25, 26] — the result obtained in those papers were most probably unachievable with the previous versions of `FIRE`. In particular, the first example of a calculation which was made real by the use of `FIRE` with modular arithmetic can be found in [27].

The new version brings a performance upgrade, being about twice as fast

¹It has been shown in [3] that the number of master integrals is always finite.

than the previous public version on small tests. However the main power comes at problems being “at the edge of science” since the memory economy and stability features results in a “can be done — cannot be done” difference.

Along with a performance upgrade, the new version contains a number of new features, including but not limited to the following:

- usage of internal sector symmetries from `LiteRed`;
- possibility to recover from system crashes without a restart from the very beginning;
- possibility to use multiples nodes on one reduction task;
- modular arithmetic approach;
- MPI-approach to run multiple modular arithmetic tasks on a super-computer;

In section 2 we recall basic definitions related to Feynman integrals, in section 3 we describe how to install `FIRE` and in section 4 how to use it. Section 5 explains the internals of `FIRE` which might be important for choosing options efficiently, section 6 is devoted to the usage of modular arithmetic, and in the appendix we summarize all options of config files.

This paper can also serve as a user manual for `FIRE6`. Additional information and examples can be found in the `FIRE` distribution. The new version of `FIRE` (starting with 6.1) is also covered with the doxygen documentation which can be useful for contributors, bug finding or those who wish to use some internal functions. Details how to build and read this documentation are also provided below.

2. Basic definitions

Let us remind the basic definitions that will be used in this paper. Feynman integrals are functions of integer variables which are also called indices,

$$F(a_1, \dots, a_n) = \int \dots \int \frac{d^d k_1 \dots d^d k_h}{E_1^{a_1} \dots E_n^{a_n}}. \quad (1)$$

Here the factors E_i in the denominator are linear functions with respect to scalar products of loop momenta k_i and external momenta p_i ; dimensional regularization with $d = 4 - 2\epsilon$ is applied.

One considers the so-called integration by parts relations [1]

$$\int \dots \int d^d k_1 d^d k_2 \dots \frac{\partial}{\partial k_i} \left(p_j \frac{1}{E_1^{a_1} \dots E_n^{a_n}} \right) = 0 \quad (2)$$

They can be rewritten as:

$$\sum \alpha_i F(a_1 + b_{i,1}, \dots, a_n + b_{i,n}) = 0. \quad (3)$$

where $b_{i,j} \in \{-1, 0, 1\}$ and α_i are linear functions of a_j .

It is already standard to use the notion of sectors during reduction. There are 2^n sectors, each of those defined by specifying a subset of indices that have to be positive (and the remaining have to be non-positive). Each sector has a unique so-called *corner integral*, that is the one with indices equal to 0 or 1.

We will define a sector to be *lower* than another sector if all indices of the corner integral in the first one are smaller than the corresponding indices of the corner integral in the second one. Basically integrals in lower sectors are simpler, so starting from the time of reduction “by hand” one tries to reduce Feynman integrals to ones in lower sectors.

We will call a sector *trivial* if all integrals corresponding to sets of indices in this sector are equal to zero. It is known that the lowest sector, where all indices are non-positive, is trivial, but in fact a large number of sectors are trivial. The conditions determining whether a sector is trivial are called *boundary conditions*.

3. Installation

FIRE6 is distributed via bitbucket. One can also download a binary package compiled on Ubuntu 16.04 or openSUSE 15.0 from the download section of the repository (<https://bitbucket.org/feynmanIntegrals/fire/>), but the recommended way is to build FIRE from sources.

To do it one has to

1) clone it with `git` (this should be done in a folder that has no space symbols in its full path)

```
git clone https://bitbucket.org/feynmanIntegrals/fire.git
```

As a result a `fire/FIRE6` folder will appear. The `fire` folder name can be later changed if needed, however it is not recommended to change the name of the internal folder `FIRE6` to be able to receive updates with `git pull`. So now one can move to the internal folder with `cd fire/FIRE6`.

To receive updated one can pull them from the repository

```
git pull
```

For details please refer to git manuals since this paper cannot cover all those.

2) configure FIRE. There is a `configure` file in the folder that can be run with

```
./configure
```

This file is not a part of the autoconf system, but a simple script that sets some preprocessor variables and makefiles corresponding to the settings provided to `./configure` There are the following options:

- `--enable_zlib`: enables the zlib compressor shipped with `kyotocabinet`, requires the `zlib/deflate` library (`zlib1g-dev`) to be installed system-wide;
- `--enable_snappy`: tries to build the `snappy` library shipped with `FIRE` to serve as a compressor;
- `--enable_zstd`: tries to build the `zstandard` library shipped with `FIRE` to serve as a compressor;
- `--enable_debug`: adds `-g` to compile options for debugging symbols and also attaches hooks for printing out stack trace in case of crashes;
- `--enable_tcmalloc`: compiles google perf tools in order to make `FIRE` use `tcmalloc` instead of the `libc malloc`;

- `--enable_lthreads`: gives FIRE possibility to use multiple threads inside a sector".

The first two options are related to compressor settings of the database engine. By default FIRE6 uses the lz4 compressor which comes shipped with FIRE and is easiest to build. The `--enable_debug` option is self-explanatory, the `--enable_tcmalloc` option might provide significant memory economy and the `lthreads` option provides a better parallelization resource (details will be explained below).

The options of `configure` are saved in the `previous_options` file.

3) build the libraries that are shipped with FIRE

```
make dep
```

This command builds `kyotocabinet`, `lz4`, `snappy` (in the case where `--enable_snappy` is set), `zstandard` (in the case where `--enable_zstd` is set) and `google perf tools` (if configured with the `--enable_tcmalloc` option). As usual, one can use the `-j N` option where `N` is the number of cores to run compilation faster. As a result, the libraries are “installed” in the `usr` subfolder of `fire/FIRE6`.

4) build FIRE

```
make
```

the simlinks to binaries are located in the `bin` folder. FIRE builds a number of binaries, with or without `p` at the end of the name. The `p` binaries are the modular arithmetic version of FIRE (comes from “prime”). The objects and binaries themselves are created in the `poly` and `prime` folders. Both of them contain the `FIRE6`, `FLAME6` and `FTool6` binaries. The `FIRE6` is the main binary, `FLAME6` is used for sector jobs and `FTool6` is an auxiliary binary that can be used for database analysis.

5) run tests


```
make test
```

The new version of FIRE has a set of tests that are also available at bitbucket via bitbucket pipelines.

6) To build the `mpi` version of FIRE (for use on supercomputers) one has to run

```
make mpi
```

This version can also be tested on a personal computer in case the `openmpi` libraries are installed system-wide. The `mpicxx` binary is called for this compilation, this can be changed by editing `mpi/Makefile`.

7) The development of FIRE now follows the gitflow strategy. The master branch only receives hotfixes and is moved mostly when a new release is being prepared. Both the master and the dev branches are changed only with pull requests, and a bunch of tests is performed via bitbucket pipelines before such a pull request can be merged. The dev branch however can receive changes much faster, so if one is willing to receive updates between releases, the dev branch has to be chosed.

```
git checkout dev
```

After that everything should be recompiled.

8) The `FLink` and `KLink` binaries are no longer a part of the FIRE distribution, so there are no more instructions on how to build them.

9) The doxygen documentation can be built with

```
make doc  
make pdf
```

The first command builds doxygen html and latex documentation. It requires doxygen to be installed. The main generated index file of the html documentation is `documentation/html/index.html`. The second command can also build a pdf manual from the latex version and requires pdflatex to be installed. If successful, it created the `documentation/refman.pdf` file. If the `xdg-open` utility is installed, the documentation in its html format can be also opened with

```
make showdoc
```

4. Basic usage of FIRE

FIRE has a part written in Wolfram *Mathematica* and a much bigger part written in C++. While FIRE is able to perform reduction by the *Mathematica* part, it should be used mainly for tests and for simple tasks. Partially because of that the support of calls to `fermat` and the database engine from *Mathematica* were removed in the current version. If a problem is complex enough that *Mathematica* cannot handle it with ease, one should use the C++ part.

The proper approach to use FIRE for complicated problems is the following. First one should create a start file in *Mathematica*. Then this file is used by the C++ part to run the reduction, afterwards tables are created which can be read by *Mathematica* in order to produce results. Moreover it is strictly recommended to use `LiteRed` in order to produce symmetries before running the reduction.

4.1. Preparing a start file

The *Mathematica* part of FIRE is loaded simply with

```
Get["FIRE6.m"];
```

provided that the current directory is set to `fire/FIRE6`. The recommendation to set the `FIREPath` variable used in previous versions is no longer valid.

Like in previous versions one has to prepare a start file to use `FIRE`. Such a file contains information on the reduction of one family of Feynman integrals — dimension, sectors, integration-by-parts relations. To have backward compatibility we keep the syntax at this stage. One has to set the following variables.

- **Internal** — the list of internal momenta, for example, $\{k\}$;
- **External** — the list of external momenta, for example, $\{p_1, p_2, p_4\}$; one should list only linearly independent external momenta after having used momentum conservation;
- **Propagators** — the list of propagators, for example, $\{-k^2, -(k + p_1)^2, -(k + p_1 + p_2)^2, -(k + p_1 + p_2 + p_4)^2\}$; the propagators should be quadratic in terms of momenta;
- **Replacements** (optional) — the list of replacement rules for kinematic invariants, for example, $\{p_1^2 \rightarrow 0, p_2^2 \rightarrow 0, p_4^2 \rightarrow 0, p_1 p_2 \rightarrow s/2, p_2 p_4 \rightarrow t/2\}$;

There are also two more variables that were usually set in previous versions but can be skipped now.

- **RESTRICTIONS** (optional) — list of boundary conditions. For example if this list has an element $\{-1, -1, -1, 0\}$, this means that the integrals are equal to zero if the first three indices are non-positive. Since currently `FIRE` can detect boundary conditions automatically in most cases, this option can be skipped. Moreover, the usage of `LiteRed` can detect even more boundary conditions at a later stage. Still, if this variable is set, those restrictions are also used;
- **SYMMETRIES** (optional) — list of symmetries (permutations of indices not changing the integrals, each element consists of list positions); in older versions one had to provide the whole symmetry group, but currently it is enough to provide the generators; for example, if this list has an element $\{3, 2, 1, 4\}$, this means that $F[a, b, c, d] = F[c, b, a, d]$. In principle, the use of symmetries can speed up the reduction, however

in case `LiteRed` is used, it can also detect equivalent sectors and create corresponding mapping rules. This will be somewhat slower compared to the use of global symmetries, but we estimate the difference as relatively small.

If one wishes to set the integration by parts relations manually, there is a special variable `startinglist` that should be set (see [20] for details).

Now one has to run the `PrepareIBP[]` and then `Prepare[]` commands. The second one leads also to the autodetection of boundary conditions. This can be turned off by running

```
Prepare[AutoDetectRestrictions -> False]
```

instead. The command also has a parallel option run with

```
Prepare[Parallel -> True]
```

Now a start file can be saved with `SaveStart["filename"]`. It is recommended to quit the kernel afterwards.




Figure 1: Massless on-shell doublebox

For example, for a double box diagram with massless lines, incoming momenta p_1 , p_2 and p_3 and $p_1^2 = 0$, $p_2^2 = 0$, $p_3^2 = 0$, $(p_1+p_2+p_3)^2 = 0$, $s = (p_1 + p_2)^2$, $t = (p_1 + p_3)^2$ one has:

```

Get["FIRE6.m"];
Internal = {k1, k2};
External = {p1, p2, p3};
Propagators = {-k12, -(k1 + p1 + p2)2, -k22, -(k2 + p1 + p2)2, -(k1 + p1)2, -(k1 - k2)2, -(k2 - p3)2, -(k2 + p1)2, -(k1 - p3)2};
Replacements = {p12 -> 0, p22 -> 0, p32 -> 0, p1 p2 -> s/2, p1 p3 -> t/2, p2 p3 -> -1/2 (s + t)};
PrepareIBP[];
Prepare[];
SaveStart["doublebox"];
Quit[];

```

As a result one has a file "doublebox.start" containing all the required information to proceed.

Please note that start files created by old versions of FIRE (public versions up to 5.2 and some private versions) cannot be used by the current C++ part. They have to be either recreated or converted by the `ConvertStart[oldfile, newfile]` command.

4.2. Reduction in Wolfram Mathematica

The start file can be loaded in Mathematica with the `LoadStart` command. For example,

```

Get["FIRE6.m"];
LoadStart["doublebox", 1];
Burn[]

```

The second argument is the number assigned to the current family of Feynman integrals. It is recommended to have different numbers for different families, this will allow one to distinguish integrals belonging to different families and to find equivalents between them later. This should be a positive number fitting an unsigned short integer (less than 2^{16}).

In complicated situations `Burn[]` might work slowly. Hence one can run `SaveData["filename"]` after `Burn[]`, then quit the kernel and later load everything with `LoadData["filename"]` (without `LoadStart` and `Burn`).

Now one can perform the reduction with

```
F[1, {1, 1, 1, 1, 1, 1, 1, -1, -1}]
```

Here `F` is a call to perform the reduction, “1” is the same number as in `LoadStart` and `{1, 1, 1, 1, 1, 1, 1, -1, -1}` is the set of indices. As a result one has an expression of the form

```
3/2 s G[1,{1,1,1,1,1,1,1,-1,0}] + 1/2 s t
G[1,{1,1,1,1,1,1,1,0,0}] + ...
```

where the dots refer to further terms that have been omitted for brevity. In the resulting expression `G` follows the same notation as `F` but is treated as an irreducible (master) integral and does not lead to new reduction.

As it has been mentioned earlier, running complicated reductions in `Mathematica` is not recommended. But in case this reduction is performed it is strictly recommended to avoid the `F` command which reduced integrals one by one and to use the more similar to the `C++` mode command `EvaluateAndSave`. It has two arguments: the list of integrals that has to be reduced and the file where to save tables. For example, to reduce two integrals by this command one should run

```
EvaluateAndSave[{{1, {1, 1, 1, 1, 1, 1, 1, -1, -1}},
{1, {1, 1, 1, 1, 1, 1, 1, 0, -2}}}, "doublebox.tables"]
```

This command creates the tables file which can be loaded, and then the `F` command can be used to produce results.

4.3. Loading the tables

The tables created with the `EvaluateAndSave` or with the C++ version can be loaded with the `LoadTables` command. In case multiple tables are to be loaded one should pass all file names in a list to a single call to `LoadTables`.

```
LoadTables[{file1,file2,...,fileN}]
```

Now a call to `F` will take the result from tables. Sometimes this can be slow because `FIRE` tries to factorize coefficients in order to present a nice-looking result. This behavior can be switched off by `FactorCoefficients = False`.

In case one prefers to use the tables without loading start file and the `F` syntax, they can be converted into `Mathematica` rules. The command for this is

```
Tables2Rules[filename, Func: Identity, JoinTerms: True]
```

It is not required to load start files to use this command and get a list of `Mathematica` rules as a result. The first argument is the file name with tables. The second (optional) argument is the function that is applied to all coefficients, where default value is `Identity` meaning no function and a reasonable choice can be `Factor` or `Together`. The third (optional) argument is the indication whether the right-hand side is transformed into a sum of integrals with coefficients. Setting it to `False` provides a list of pairs (an integral and a coefficient) which can be useful in case of further automatic processing of results.

4.4. C++ reduction

As it has been mentioned a few times above, the reduction should be run with the C++ part of the program. As a result it created tables that can be loaded into Mathematica afterwards.

To run the C++ version, one has to create a configuration file (with the `config` extension). All options of config files will be discussed below, but let us first illustrate the syntax with this example:

```
#variables d, s, t
#start
#folder examples/
#problem 1 doublebox.start
#integrals doublebox.m
#output doublebox.tables
```

The spaces are insignificant, each line should start from `#` (nowadays called the hash symbol). `#variables` should list all variables that can appear. If variables are set incorrectly, the reduction will freeze due to fermat.

The `#folder` command is optional and provides a path to a folder where files listed by the following commands reside (unless specified with an absolute path). If the folder instruction is missing, the paths are considered absolute or from the “current” directory. `#problem` followed by the number corresponding to the current family of Feynman integrals and the path to a start file. `#integrals` points to a file containing a list of integrals that have to be reduced and `#output` points to a file where the resulting tables are to be saved. The input is a Mathematica list of pairs, for example, $\{\{1, \{1, 1, 1, 1, 1, 1, 1, -1, -1\}\}, \{1, \{1, 1, 1, 1, 1, 1, 1, 1, 0, -2\}\}\}$.

Now the reduction can be launched with

```
bin/FIRE6 -c examples/doublebox
```

It should be stated once more that this reduction does not need Mathematica. As a result one gets a file containing reduction results — expressions of integrals via master integrals.

4.5. Using LiteRed rules in Mathematica

LiteRed is a program by R. Lee [14, 15, 16] that aims at solving the IBP relations before the substitution of indices. This is an alternative approach

to the IBP reduction. However even without reduction rules the `LiteRed` program can provide important information that should be used in `FIRE`. This information consists of

- Boundary conditions (the knowledge all integrals in some sectors are equal to zero);
- External symmetries (mappings between different sectors);
- Internal symmetries (mappings inside sectors);
- Reduction rules in some sectors (we normally avoid this, but sometimes it can help).

Currently `FIRE` comes shipped with the version 1.8 of `LiteRed`. It is recommended not to use `LiteRed` downloaded elsewhere since there can be a difference in file formats.

If one wishes to upgrade to the most recent version of `LiteRed` it can be downloaded from <http://www.inp.nsk.su/~lee/programs/LiteRed/>, however we cannot guarantee their compatibility.

Please note that `LiteRed` is not a part of `FIRE` but a program having another author. Therefore, when using `LiteRed` options one should give proper credit to `LiteRed` program by citing [14, 15, 16].

To use `LiteRed` rules one has to construct them first. There are multiple examples shipped with the `LiteRed` package, let us consider one of them, the two-loop massless onshell vertex.




Figure 2: Two-loop massless onshell vertex.

In the context of `FIRE` it is convenient to call `LiteRed` rules construction with

```

SetDirectory["extra/LiteRed/Setup/"];
Get["LiteRed.m"];
SetDirectory["../../../../../"];
Get["FIRE6.m"];
Internal = {1, r};
External = {p, q};
Propagators = {-(1 - r)2, -12, -r2, -(-1 + p)2, -(q - r)2,
-(-1 + p + r)2, -(1 + q - r)2};
Replacements = {p2 -> 0, q2 -> 0; p q -> -1/2};
CreateNewBasis[v2, Directory -> "temp/v2.dir"];
GenerateIBP[v2];
AnalyzeSectors[v2, {0, __}];
FindSymmetries[v2, EMs->True];
DiskSave[v2];
Quit[];

```

Note: `CreateNewBasis` is not a command of `LiteRed`, it is a command in `FIRE` that translates `FIRE` input into `LiteRed` input.

Here the `extra/LiteRed/Setup/` is inside the `FIRE6` folder, `v2` is a name for the diagram used inside `LiteRed`.

The `AnalyzeSectors` command finds trivial sectors. The last parameter in this example assumes the one is considering integrals with non-positive first index. In other cases one should move the position of the zeros and underscore wildcards.

The `FindSymmetries` command finds symmetric sectors meaning those that can be mapped as sums to other sectors, the `EMs->True` option stands for the autodetection of external symmetries. For other syntax variants please refer to `LiteRed` examples.

Before `DiskSave` one can also invoke the `SolvejSector` command in order to construct rules for some sectors. However our experience shows that one is advised to be careful with it because partial rules can degrade performance of the Laporta approach.

As a result one gets a folder with multiple files containing information on the chosen family of Feynman integrals. They can be loaded in `Mathematica` with

```
Get["FIRE6.m"];
LoadStart["examples/v2", 2];
LoadLRules["temp/v2.dir", 2];
Burn[];
```

Here “2” stands for the problem number that should be the same in both lines. The `LoadLRules` command reads the directory with `LiteRed` rules and loads everything it can use out of there. Then the reduction can be performed as before.

4.6. Using `LiteRed` rules in C++

The C++ version is not capable of parsing the `LiteRed` folder directly, they have to be converted first. It can be done with

```
Get["FIRE6.m"];
LoadStart["examples/v2"];
TransformRules["temp/v2.dir", "examples/v2.lbases", 2];
SaveSBases["examples/v2"];
```

Here 2 is the problem number. It is important to mention that the result of the `TransformRules` command is both a `lbases` file and a new file with the `sbases` extension that should be used instead of the original start file. The syntax of those files is currently the same (`sbases` comes from a deprecated idea to use Groebner bases in order to run the reduction), however we keep different extensions in order to distinguish the original start file and the result of `TransformRules`. The new file contains some extra information compared with the original start file such as sector priorities, orderings, extra boundary conditions.

Now the config file can be changed and should contain the following lines:

```
#problem 2 v2.sbases
#lbases v2.lbases
```

It is worth to mention that lbases files created by old versions of FIRE (up to public 5.2 and some private versions) do not contain internal symmetries and should be re-created.

4.7. Finding equivalent between master integrals

Extra master integrals that are equivalent to each other can degrade performance a lot for final reductions. Normally the usage of `Litered` should result in having no equivalent masters. However sometimes they can still exist. In order to find equivalents one should use the `FindRules` and `WriteRules` commands.

In the doublebox example the `MasterIntegrals []` command can be used to obtain the list of masters:

```
{1, {0, 0, 0, 0, 1, 1, 1, 0, 0}}, {1, {0, 0, 1, 1, 1, 1, 1, 0, 0},
0, 0}}, {1, {0, 0, 1, 1, 1, 1, 1, 0, 0}}, {1, {0, 1, 1, 0,
0, 1, 0, 0, 0}}, {1, {0, 1, 1, 0, 1, 1, 1, 0, 0}}, {1, {1,
0, 0, 1, 0, 1, 0, 0, 0}}, {1, {1, 0, 0, 1, 1, 1, 1, 0, 0}},
{1, {1, 1, 0, 0, 0, 1, 1, 0, 0}}, {1, {1, 1, 0, 0, 1, 1, 1,
0, 0}}, {1, {1, 1, 1, 1, 0, 0, 0, 0, 0}}, {1, {1, 1, 1, 1,
1, 1, 1, 0, 0}}, {1, {1, 1, 1, 1, 1, 1, 1, 1, -1, 0}}
```

It is easy to see that some of the integrals in lower sectors are equivalent. This can be found with

```
Internal = {k1, k2};
External = {p1, p2, p3};
```

```

Propagators = {-k12, -(k1 + p1 + p2)2, -k22, -(k2 + p1 +
p2)2, -(k1 + p1)2, -(k1 - k2)2, -(k2 - p3)2, -(k2 + p1)2,
-(k1 - p3)2};
Replacements = {p12 -> 0, p22 -> 0, p32 -> 0, p1 p2 -> s/2,
p1 p3 -> t/2, p2 p3 -> -1/2 (s + t)};
FindRules[MasterIntegrals[]]

```

or saved to a file with

```

WriteRules[MasterIntegrals[], "examples/doublebox"];

```

The resulting file has a syntax like

```

G[1, {0, 0, 1, 1, 1, 1, 1, 0, 0}] -> {{1, G[1, {1, 1, 0, 0,
1, 1, 1, 0, 0}]}};
G[1, {1, 0, 0, 1, 1, 1, 1, 0, 0}] -> {{1, G[1, {0, 1, 1, 0,
1, 1, 1, 0, 0}]}};
G[1, {1, 1, 0, 0, 0, 1, 1, 0, 0}] -> {{1, G[1, {0, 0, 1, 1,
1, 1, 0, 0, 0}]}};
G[1, {1, 0, 0, 1, 0, 1, 0, 0, 0}] -> {{1, G[1, {0, 1, 1, 0,
0, 1, 0, 0, 0}]}};

```

In case one wishes to provide rules mapping an integral to multiple integrals the right-hand sides should be presented as lists of pairs containing a coefficient and an integral. All rules in the file should be separated by two new lines. This format can be also created by FIRE automatically by the `SaveRulesToFile[rules,filename]`, where `rules` should be in a Mathematica rules format.

Moreover, the `FindRules` command can find equivalents between integrals of different families of Feynman integrals. To do this one has to set

the `Problems` variable to a list of different diagram numbers, then set the `Internal[i]`, `External[i]`, `Propagators[i]` and `Replacements[i]` for each `i` from this set and call the `FindRules` command on a set of integrals under consideration. In this case `FIRE` tries to map integrals to the ones having higher problem numbers when possible.

The resulting file contains rules and can be loaded in `Mathematica` with `LoadRules[filename]` or in `C++` with the `#rules` line.

5. Internals of FIRE

In order to be able to set options of the `C++ FIRE` properly, one should understand to some extent how it works internally. `FIRE6` launches several processes. The first among them is the `FIRE6` binary itself. It reads the config file and the files listed there, enumerates the sectors and starts running individual sector jobs. Each sector has its own database, and unlike in older versions, there is no shared database for different sectors, so this leads to a better parallelization. Since `FIRE` needs to launch other binaries in its folder, it needs to know their location. So it should be run with a relative or full path, but copying the binary to other location or creating simlinks might lead to crashes.

5.1. Reduction processes

The `FIRE6` binary first reads the list of requested integrals and puts them into corresponding sector databases. Then it launches individual sectors jobs (`FLAME6` processes) starting from highest sectors and going down. Multiple sector jobs can be launched at the same time in case they share the same level (number of positive indices in a sector) and sublevel (in case of `LiteRed` external symmetries the sectors which are mapped to another sector are considered to be in a higher sublevel). The maximal number of the simultaneous sector jobs is controlled by the `#threads` option.

When `FLAME6` works in a sector, it masks all integrals belonging to lower sectors to prevent growth of substitutions. They will be substituted later.

When a level (or sublevel) is over, the `FIRE6` process reads the databases which were modified by the sector jobs and obtains the lists of integrals required in lower sectors. It also opens the lower databases and puts the required lists there, then proceeds with lower levels.

After `FIRE` finishes with the lowest possible level, it turns back to substitutions and starts increasing levels launching substitution processes (also

FLAME6). In case the `#sthreads` option is missing, the `#threads` option controls the maximum number of simultaneous jobs. Between different levels it copies required results from lower databases to higher databases, which does not take much computing time.

One should keep in mind that the memory usage of FIRE is almost linearly proportional to the number of sectors reduced (or substituted) at the same time, so the options should be set accordingly.

There is also a mode of FIRE that does not perform backward substitutions and creates tables listing irreducible integrals after the forward stage. This is turned on by using `#masters` instead of `#output`.

The numbering of the sectors starts from 2 since 1 is reserved for the points that are of highest priority during the reduction, meaning the right-hand of rules or integrals listed by the `#rules` setting. Whenever an integral from the right-hand side is read, FIRE creates a relation that maps it to its virtual analogue in sector 1. Among other things it means that integrals is right-hands side of rules should not appear in left-hand sides for consistency.

The `#preferred` setting as well as `#rules` setting influence the integral priority inside a sector. All integrals listed as preferred as well as all integrals appearing in both sides of the rules have a priority in their sectors — FIRE tries to express other integrals through them if possible. Note: unlike in FIRE5 the corner integral in each sector is not automatically considered to be preferred. This might be useful in some cases but will require adding it to preferred list if this behavior was assumed.

The sectors symmetric to lower sectors by the global symmetries do not have their own numbers, instead the requested integrals are mapped to corresponding lowest points in orbits before reduction starts.

5.2. Database usage and multiple computers

FIRE6 has two ways to work with databases, the disk mode (default) and the memory mode (turned on by the `#memory` setting). In the default mode FIRE uses disk databases which are open by sector jobs or the master job for data transfer. In the other case on-memory databases are used, and their snapshots are saved to disk when work in the sector is done. In both cases the databases are stored in the folder specified by the `#database` setting. By default it is equal to `temp/db`, so one should consider whether this path is suitable for large files. Moreover multiple jobs using the same data path might conflict with each other.

In order to be able to recover from system crashes or timeouts, one should keep the databases. However when FIRE starts working in a sector, the sector database might end in a non-consistent stage after a crash. Moreover, often one might use temporary locations for the databases. Therefore we introduce a concept of a storage folder, another folder where FIRE stores copies of sector databases and copies the required databases to the database folder when needed, then copies them back when work is done. If a storage folder is set (by the `#storage` setting), one can restart the job after a crash, and it can recover relatively fast.

The storage folder also gives the possibility to get FIRE running on multiple nodes for the same reduction. To do this one should start a FLAME6 binary on the slave machine with the same options. The process will parse the configuration and then start waiting for an IP file to appear in the storage folder. Then the FIRE6 job starts, it opens a socket, writes its IP address into the IP file and one of the threads start accepting connections. The FLAME6 process from another machine connects to the open socket, and now FIRE6 is capable of distributing jobs both locally and remote. For this to work the storage folder should be on a network drive with access from both nodes, however the access to this folder is infrequent. Also the nodes should be able to connect to each other via the TCP/IP protocol. The default port is 8080, but can be set by the `#port` setting. If the port is set to zero, this possibility is turned off.

5.3. Parallelization

There is a number of ways to get FIRE to use parallelization. Some of them have been mentioned above and are related to multiple sector jobs, both on one or multiple machines. But those are not the only options.

FIRE6 uses `fermat` for calculations and sends expressions to it to get them simplified. A number of `fermat` processes are launched, defined by the `#fthreads` setting (by default it is equal to `#threads`). There are two different ways how FIRE can handle `fermat` jobs.

The default mode is a shared pool of `fermat` workers for all sectors. The child FLAME6 processes communicate with FIRE6 via pipes and send expressions to be simplified. FIRE6 has a pool of `fermat` workers and distributes received expressions to this pool. An alternative way is to have a separate pool of `fermat` workers for each sector job. This mode is turned on by adding the letter `s` before the number in the `fthreads` setting. Please note that in

this case the `fthreads` is divided by the `threads` setting in order to determine the number of `fermat` workers launched by a sector job. It depends on the problem which mode works more efficiently.

Another resource for parallelization is the usage of level threads. To use them one first needs to compile `FIRE` and dependent libraries with support of level threads (`enable_lthreads`). Then the `lthreads` setting should be used. Level threads are threads that can run in parallel in a given sector working almost independently for integration by parts relations of similar complexity. By complexity of an integral inside a sector we mean a pair of numbers, the sum of positive indices minus the number of positive indices (= number of “dots”) and minus the number of non-positive indices (= total power of irreducible numerators). Integrals with equal sum of numbers in this pair can be treated almost independently, and here level threads can be applied. Again, the possibility to gain speedup from this feature greatly depends on the nature of the problem and on the structure of the computers in use. The level threads work only when the separate `fermat` mode is set.

Please note that the `lthreads` mode (even if configured and not turned on) is slightly slower due to database locks. This change is negligible in the polynomial version but can be noticed in the prime version, so there might be a reason not to configure `lthreads` support.

6. Modular arithmetic

The approach with the use of modular arithmetic² for IBP reduction was first suggested in [17]. Basically there are two ideas here. The first one follows from the fact that we know that there are too many integration by parts relations but before solving them we do not know any minimally required subset. So what if one first runs the reduction with substituted variables, gets the information of what relations were really needed and then runs the real reduction? Of course, one has to choose values for d and other variables such that they do not fall into possible zeros of denominator factors.

The other idea follows from the fact that intermediate coefficients when solving sets of linear equations can be huge compared to resulting coefficients. So one would wish to avoid this intermediate stage since it might fail to fit RAM limits. Hence the approach is to perform reduction with different values

² Modular arithmetic is in use not only when solving IBP relations but also in other problems in high-energy particle physics and quantum field theory – see, e.g., [28]

of variables and then to reconstruct the coefficients. However simply giving integer or rational values to all variables is normally not enough — instead of polynomial growth one gets number growth. To avoid this one also sets a large prime number p and used the so-called modular arithmetic meaning all integers are replaced with their remainders from division by p . The set of those remainders is a field meaning every non-zero remainder has a unique inverse. Hence moving to such a field leads to a limit for all numbers. A prime has to be chosen big enough so it makes sense to take biggest prime numbers that fit into machine-sized integers (2^{64}) for faster calculations.

To run FIRE with modular arithmetic one has to do the following:

- Use the FIRE6p instead of the FIRE6 binary;
- Set values for all variables with the `#variables` setting (the syntax is explained in the Appendix);
- Set a prime number with the `#prime` setting (it is not the prime number itself that is provided but its index in the set of hard-coded primes close to 2^{64}).

A single run with the modular arithmetic might be used together with the `#hint` option pointing at a folder where the hint files are to be saved. Those files list the integration by parts relations that were really used during the reduction. A consequent run with the `#hint` option will get use of those files.

Now if one is going to use the modular arithmetic and reconstruction of the results, a reasonable number of FIRE runs is to be made. To avoid using multiple config files, one can use the following syntax to run FIRE6p:

```
bin/FIRE6p d-p examples/doublebox
```

in the case when d is the only variable and

```
bin/FIRE6p d-x-p examples/doublebox
```

in the case of one extra variable. Here p is the number of the large prime number to be used, d is the value of d and x is the value of the other variable. With such a syntax one can avoid editing config files and changing the `#variables` and `#prime` settings.

More variable substitutions can also be used although there is no stable reconstruction procedure at this point. The usage is the following:

- Variable values should be separated by `-`;
- No negative values allowed;
- Last number indicates the prime number;
- Variables are replaced in the order they appear in the `#variables` section, first N variables are replaced where N is the number of values provided in the command line minus one;
- Replacements like `d->57` in the configuration file get overwritten if they are among the first N variables;
- Replacements like `d->/57` in the configuration file get joined and form a fraction;
- Extra variables after the first N should have replacement values in the `#variables` section.

One might need a huge number of jobs like that, and consequent runs can take too much time. Therefore this approach is most suitable when a supercomputer is used. To use massive parallelization one should prepare a config file and then use the `FIRE6_MPI` binary. It has the following syntax:

```
bin/FIRE6_MPI -c config p d_min d_max
```

in the case when d is the only variable and

```
bin/FIRE6_MPI -c config p d_min d_max x_min x_max
```

in the case of one extra variable. Here p is the number of different primes that will be used (starting from the biggest possible prime number corresponding to the setting `#prime 1` and going on), `d_min` is the minimal value of d , `d_max` is the maximal value of d , similarly for `x_min` and `x_max`. The name of the other variable in problem files is unimportant.

Currently the MPI version is incapable of running jobs with more than two variables to be substituted with different values.

As a result $p * (d_{\max} - d_{\min} + 1) * (x_{\max} - x_{\min} + 1)$ jobs are to be run one after another. However the `FIRE6_MPI` should not be called directly but through the MPI system of the supercomputer. For instructions one should refer to the instructions of the particular cluster. As a result one runs a large number of jobs at the same time.

It is safe to use one config file and have multiple jobs running at the same machine through MPI: `FIRE` creates subfolders in the database folder with names containing process id's and node names for databases. However in case of crashes (which is normal for supercomputers) those folders might be remaining, so sometimes one should clean then not to run out of quota.

As a result each `FIRE` job creates a tables file, but with a changed name: `name-d-p.tables` or `name-d-p-x.tables`. Moreover, each job "reserves" tables files when starts working by creating empty tables with such names, so that one can launch multiple MPI jobs at the same time. Again, in case of supercomputer node crashes one should delete empty tables files. This can be done by

```
find FOLDERNAME -size 0 -print0 |xargs -0 rm --
```

The jobs should be run until all tables are created. And then one can move to the coefficient reconstruction.

One can also run multiple tasks without MPI but for example with the use of GPU parallel, for example,

```
parallel "bin/FIRE6p -variables {1}-{2} -c config" :::  
$(seq 100 110) ::: $(seq 1 3)
```

However a direct call to such a command will result in a crash because the jobs conflict with other (databases and semaphores) and also can provide too much output. So the command should be changed to

```
parallel "bin/FIRE6p -variables {1}-{2} -c config -parallel  
-silent" ::: $(seq 100 110) ::: $(seq 1 3)
```

to enable the concurrent mode with no conflict between individual jobs (process id's are added to the database paths and POSIX named semaphores are also modified).

The recommended options for running multiple modular jobs are:

- `#memory` for faster runs;
- `#compressor none` in case there is enough RAM;
- `#bucket N` where `N` is the smallest number that does not result in “reopening database with bucket...” messages;
- `#wrap` if there is a quota on the number of files;
- `#small` to save some RAM.

Still before running jobs one does not know how many tables will be required. We will try to give our “algorithm” on how those jobs should be run, but let us first explain the reconstruction process.

6.1. Reconstruction

The reconstruction of coefficients consists of two stages. First there is the reconstruction of rational coefficients from modular numbers (while variables have still fixed values). It searches for the rational number with minimal absolute values of the numerator and denominator that is mapped to the obtained remainders over large prime numbers. When increasing the number of primes no longer changes the result, some more checks are made and the rational number is reconstructed. FIRE can perform this operation on all coefficients in tables, this is done by the `Mathematica` part of FIRE6 with:

```
RationalReconstructTables[filename, pnum]
```

command. Here `filename` is the name of the target filename (without the `p`-part) and `pnum` is the number of different prime numbers in use, starting from the first (largest). The result is the requested tables file if successful. `FIRE` produces messages which can be turned off with adding the `True` option:

```
RationalReconstructTables[filename, pnum, True]
```

6.2. One-dimensional reconstruction

When the rational reconstruction is done, one can move to the reconstruction of coefficients. Here we recall classical approaches to the reconstruction of formulas. Let us start with the case of one variable. One of the approaches is based on the so-called Newton reconstruction which uses the the following formula:

```
a[0] + (x - x[0]) (a[1] + (x - x[1]) ( ... a[n] )... ))
```

There are well-defined algorithms that can perform this reconstruction for any polynomial (not a rational function). If a reconstruction of a rational function is needed, then one should use the Thiele formula:

```
a[0] + (x - x[0]) / (a[1] + (x - x[1]) / ( ... a[n] ) ... ))
```

One should keep in mind that this method needs more values. For example, reconstructing a polynomial with the Newton formula needs the number

of points close to its power, but reconstructing the same polynomial with the Thiele formula needs about twice that many points.

One can perform the reconstruction with the

```
ThieleReconstructTables[filename, dlist]
```

Here `filename` is the reconstruction target filename and `dlist` is the list of different values of d that will be used. As a result one obtains the file with tables in the same format, but containing rational polynomial coefficients. FIRE produces messages which can be turned off with adding the `True` option.

Hence the general instructions to perform a one-dimensional reconstruction are the following

1. Choose a relatively big number of d and run modular arithmetic reductions until the rational numbers can be reconstructed (for example, in [27] we required 13 values);
2. Increase the number of primes used a bit to be safe and run massive reductions at the supercomputer with different values of prime numbers and d ;
3. Run the rational reconstruction for each value of d in use, if it fails, increase the number of primes and return to step 2;
4. Run the Thiele reconstruction, if it fails, increase the number of values of d and return to step 2.

6.3. Two-dimensional reconstruction

Now when considering the case of two variables there is nothing known to work better than a mixture of the two approaches. One builds a Newton or Thiele formula in one variable and the coefficients are built by the Newton or Thiele formula in another variable. Now the problem is that a Thiele-Thiele formula might make the reconstruction too complex, both computationally and also requiring a large number of tables, but other variants do not cover the general case of a rational function of two variables. However, we can take into account that the denominators in IBPs correspond to singularities of Feynman integrals which correspond to thresholds and poles whose positions are described in terms of kinematical invariants and masses, so that the

dependence on d should not mix with the dependence on these variables. In practice, it can happen however that these variables do mix with d . Then the natural procedure is to find a basis of master integrals such that the denominators are split into a product of a function of d and a function of x . In our experience this is always possible. We will call such a basis *factorising*. It can be checked that a basis is factorising on a simpler task that can be reduced with the classical approach.

So now supposing we have a factorising basis, we can continue in the following way. First, after the numerical reconstruction of rational functions is complete, one fixes a value of x and recovers the tables with the one-dimensional Thiele function. Then the same is done with fixing a value of d . Now knowing the structure of denominators we understand that we can take the least common multiple of tables with d and tables with x , multiply those factors, and afterwards all coefficients multiplied by this factor become polynomials both in x and d . Then one can run a Newton-Newton reconstruction.

Now let us proceed to instructions on how one uses this approach provided a factorising basis has been found.

1. Fix a value of d and a value of x (we will use this name for the variable) and build a number of tables until rational reconstruction can be performed with

```
RationalReconstructTables["TASK-" <> ToString[d0] <>
  "-" <> ToString[x0] <> ".tables", p0, True]
```

$p0$ stands for the number of different big primes in tables, `True` stands for silent, giving only a summary for each table reconstruction. If "Rational reconstruction unstable" is encountered, $p0$ has to be increased. After that one knows the number of prime numbers needed for the reconstruction. In the following steps this should be increased by 1 or 2 to be on the safe side;

2. Fix a value of x and create a number of tables with different values of d and p until Thiele reconstruction over d succeeds with


```

For[x = x0, x <= x1, ++x,
RationalReconstructTables["TASK-" <> ToString[d0] <>
 "-" <> ToString[x] <> ".tables", p0, True]
];
ThieleReconstructTables["TASK-"<>ToString[d0]<>
 ".tables", Range[x0, x1], True]

```

Here d_0 is the chosen value for d , x_0 is the minimal x and x_1 is the maximal x for which there are tables, `True` again stands for the silent mode. If the reconstruction is unstable, x_1 should be increased. When the Thiele reconstruction is stable, the required denominator factor can be recovered from new tables by

```

DenominatorFactor["TASK-"<>ToString[d0]<>.tables"]
/.d->x

```

The replacement here is needed because the Thiele reconstruction originally is intended for the case of one variable and assumes it is d . The answer returned is the worst denominator and will work as the factor later. To be safe one can recheck the factor with another d value reconstruction.

Note that the Thiele reconstruction in x needs more values than the later appearing Newton reconstruction after multiplying by a coefficient. So when there is the factor, the reconstruction should be rerun with

```

NewtonReconstructTables[TASK <> "-" <> ToString[d0] <>
 ".tables", Range[x0, x1], factorX, True];

```

This gives the information on how many values of x for the final table creation, and it is normally much smaller than the original number needed for the Thiele reconstruction.

3. A similar procedure should be repeated with a fixed value of x and running

```

For[d = d0, d <= d1, ++d,
RationalReconstructTables["TASK-" <> ToString[d] <> "-"
<> ToString[x0] <> ".tables", p0, True]
];
ThieleReconstructTables[TASK, Range[d0, d1], True, "-"
<> ToString[x0] <> ".tables"];

```

Note that there is a change in the `ThieleReconstructTables` syntax. There is an extra argument, that is appended to the tables name after the running index.

After the reconstruction goes through, the worst factor in d is obtained with `DenominatorFactor["TASK.tables"]`. The factor and the number of different x values can be checked and obtained with

```

NewtonReconstructTables[TASK <> ".tables", Range[d0,
d1], factorD, True, "-" <> ToString[x0] <> ".tables"];

```

- Now one knows the worst possible denominator which is a product of the factor in d and the factor in x , one also knows the number of different values of d and the number of different values of x required for the reconstruction. This number of tables must now be created. The final reconstruction is run with

```

For[d = d0, d <= d1, ++d,
For[x = x0, x <= x1, ++x,
RationalReconstructTables[
"TASK-"<>ToString[d]<>"-"<>ToString[x]<>".tables",
p0,True] ] ];
NewtonNewtonReconstructTables["TASK.tables", Range[d0,
d1], Range[x0, x1], x, denominator, False];

```

As a result the tables with reconstructed coefficients are created.

Appendix: config files for C++

Most of FIRE options are provided via the config files. But a few can come via the command line:

- `-c filename` — the path to the config file with no extension
- `-variables vars` (only prime mode) — lists the substitution values for variables separated by “-”, the last one is the index of the prime number to be used;
- `-parallel` — makes FIRE use modified paths and named semaphores, so that multiple jobs do not conflict with each other;
- `-silent` — turns off most of the output;
- `-database path` — same with the `#database` option in the config file; usefull if a dynamic path is to be specified; this option overrided the `#database` entry in the config file if both are present;
- `-sector s` (option of FLAME6) — a direct way to make the FLAME6 work in a particular sector; positive values mean forward pass, negative values mean backward pass;

Let us now summarize all options of config files. The order of some of the options is important, so it is recommended to keep the order as described here or in supplementary examples (especially the order of the options that are listed after the `#start` command). A line starting with `##` is considered as a comment and is ignored by FIRE.

- `#fermat` (optional) — the path to the `fermat` binary. By default the binary shipped with FIRE is used, but one might wish to change it;
- `#compressor` (optional) — compressor choice for the database engine. The possible compressors depend on compilation options. The full set of values is `lz4`, `lz4fast`, `lz4hc`, `zlib`, `snappy`, `zstd`, `none`. The first three are from the `lz4` family with `lz4hc` being the slowest and compressing most and `lz4fast` being the fastest among all compressors shipped with FIRE. `lz4fast` is the suggested option for the prime version, probably together with the `#small` option. For a long time we used `snappy` as the best choice for the polynomial version as

being able to compress well enough but reasonably fast. However there is a recent addition to **FIRE**, the **zstd** option standing for the ZStandard compressor that claims to compress better than **snappy** at the same speed. The **none** setting means no compressor and can provide a noticeable speedup for the prime version in case there is enough RAM;

- **#threads** — the number of threads launched for parallel reduction of sectors of same level;
- **#fthreads** (optional) — the number of **fermat** processes launched. By default it is equal to the number of threads, but it might make sense to increase it; an **s** before the number turns on the separate **fermat** mode;
- **#stthreads** (optional) — the number of threads used during the substitution stage. By default it is equal to the number of threads, but it might make sense to decrease it in case of problems with RAM;
- **#lthreads** (optional) — the number of level threads (**FIRE** should be compiled with support of level threads for this option);
- **#port** (optional) — the port for the main job to listen to accept child connections. 0 is the default value meaning the setting if off, value 8080 is recommended to enable;
- **#variables** — comma-separated list of variables used during the reduction; also this setting can provide variable replacements like **d->57**, the right-hand sides of rules should be numbers;
- **#pos_pref** (optional) — a setting that allows to tweak the choice of master integrals. By default **FIRE** chooses such an ordering that having one positive shift has a higher priority than having a negative shift, meaning that obtains master integrals with one 2 instead of one -1. The integrals with positive indices are easier in a sense because the equivalent between them can be found. But in cases where there are many masses sometimes this priority is not enough, and one gets integrals with two or more -1. To set a priority for integrals with two (or more) “dots”, one has to increase the **pos_pref** setting (by default 1). This also leads to internal symmetries being applied to a larger subset of integrals;

- `#database` (optional) — path to the place where FIRE stores the data; by default it points to the `temp/db` directory relative to the current folder, but this can be changed; it is especially important to be sure that this setting does not point to a network drive in case one does not use the `memory` setting;
- `#storage` (optional) — sets a folder for the storage of database files; by default, the storage is used only at the forward pass, to use it also during substitutions precede the path with an exclamation mark (“!”), no quotes); please note that this option is incompatible with different `#threads` and `#stthreads` – FIRE cannot distinguish properly local and remote threads when some are to be stopped;
- `#bucket` (database tuning) — an integer number equal to 20 by default and related to the database engine; small values make FIRE auto-increase the bucket during the reduction, and this can slow things down, large values can make FIRE use too much RAM; for complicated tasks consider a bucket value equal to 27–30; the old values are remaining here for backward compatibility, but currently FIRE uses the shifted (by -4) value and prints it in the log;
- `#wrap` (database tuning) — results in sector databases being stored in a general database, only the active databases remain being files; this greatly reduces the number of files used which can be important in case of file limits when running multiple jobs on a cluster;
- `#prime` (prime version) — is used and should be used only in the “prime” version of FIRE. It sets the prime number to be used in evaluations. If the option is set to zero, the prime number 2017 is used which is useful for tests (to see small numbers in results). Any other value uses one of the maximal prime numbers fitting into a 64-bit unsigned integer (currently there are 128 hard-coded primes in FIRE);
- `#memory` (optional) — if this line exists, FIRE uses the “RAM mode” instead of the “disk mode” — it stores active databases in RAM; this setting makes FIRE use more RAM, but it becomes less vulnerably to freezing because of network drives;
- `#clean` (optional) — cleans the POSIX named semaphores opened by older runs of FIRE that crashed; safe unless one user launches multiple

FIRE jobs at the same machine simultaneously; this might be required when FIRE starts crashing at random places at a particular machine that there is no way to reboot;

- `#keepall` (debugging) — keeps all entries in all databases, no removal;
- `#small` (database tuning) — decreases the key size for database entries, which is especially efficient in case of the number of indices about 10. This might make sense to decrease database size in the “prime” mode;
- `#allIBP` (for non-standard IBPs) — by default FIRE does not use some of the integration by parts relations that are guaranteed consequences of others; this logic is valid in case those are real IBPs in the classical case but might result in missing some relations in special cases. If this is the situation, then this line has to be a part of the config file;
- `#nolock` (database tuning) — if this line exists, FIRE does not lock databases in “disk mode”; this might be required to get FIRE work on some file systems;
- `#start` — just a command following the previous lines;
- `#folder` (optional) — if this path is given, all following paths will be considered relative to this folder unless they are absolute paths (starting with /);
- `#problem` — the instruction to load a start or sbases file; the syntax is `#problem pn filename` or `#problem pn |maxpos|filename` or `#problem pn |minpos,maxpos|filename`; `pn` here is the diagram number; if `maxpos` is provided, then indices bigger than `maxpos` cannot be positive; if `minpos` is provided, then indices smaller than `minpos` cannot be positive;
- `#hint` (optional) — points to a folder with hint files which list the relations that should be used (or created if missing); might be useful if one first makes a “prime” version run, then uses the obtained hint files;
- `#lbases` (optional) — a command to load with LiteRed rules obtained by `TransformRules`;

- `#output` or `#masters` — the path where FIRE will store the resulting tables; if one chooses the `#masters` syntax, then FIRE only aims at finding master integrals, and this can be much faster than the whole reduction; this might be needed to find master integrals, then one can use `WriteRules` to find equivalents between them, so that afterwards the full reduction can be run with the use of those rules;
- `#preferred` (optional) — this file can list integrals that are preferred as master integrals; this might be needed if one does not like the automatic choice; however, one should keep in mind that it is just a hint for FIRE; the syntax of the file is the same with the input list;
- `#rules` (optional) — a command to load a file with rules for some integrals; the syntax is explained in section 4.7;
- `#integrals` — the file with integrals to be reduced.

Conclusion

We presented a new version of the FIRE program performing Feynman integral reduction. The new version has been developed in private for a few years and was successfully used in multiple projects. Currently it is made public of the bitbucket. The new version offers a lot of functionality requested by the users and by the needs of reduction problems — the possibility to recover from crashes, usage of multiple machines for one reduction, parallelization inside sectors, modular arithmetic, support of supercomputers and other features. We hope that the new public version will be widely used in modern research in elementary particle physics. We have more plans on how FIRE can be improved and we are ready to accept improvement ideas and error reports.

Acknowledgements

The work is partially supported by RFBR, grant 17-02-00175. The authors would like to thank Vladimir Smirnov, Mathias Steinhauser, Joshua Davies and Go Mishima for useful advices, discussions and testing the program in its development stage. Testing of the MPI version of FIRE was performed at the MSU Lomonosov supercomputer[29].

References

- [1] K. G. Chetyrkin, F. V. Tkachov, Integration by Parts: The Algorithm to Calculate beta Functions in 4 Loops, *Nucl. Phys. B*192 (1981) 159–204. doi:10.1016/0550-3213(81)90199-1.
- [2] V. A. Smirnov, Analytic tools for Feynman integrals, *Springer Tracts Mod. Phys.* 250 (2012) 1–296. doi:10.1007/978-3-642-34886-0.
- [3] A. V. Smirnov, A. V. Petukhov, The number of master integrals is finite, *Lett. Math. Phys.* 97 (2011) 37–44. doi:10.1007/s11005-010-0450-0. arXiv:1004.4199.
- [4] S. Laporta, High-precision calculation of multi-loop Feynman integrals by difference equations, *Int. J. Mod. Phys. A*15 (2000) 5087–5159. doi:10.1016/S0217-751X(00)00215-7. arXiv:hep-ph/0102033.
- [5] C. Anastasiou, A. Lazopoulos, Automatic integral reduction for higher order perturbative calculations, *JHEP* 07 (2004) 046. doi:10.1088/1126-6708/2004/07/046. arXiv:hep-ph/0404258.
- [6] C. Studerus, Reduze - Feynman Integral Reduction in C++, *Comput. Phys. Commun.* 181 (2010) 1293–1300. doi:10.1016/j.cpc.2010.03.012. arXiv:0912.2546.
- [7] A. von Manteuffel, C. Studerus, Reduze 2 - Distributed Feynman Integral Reduction (2012). arXiv:1201.4330.
- [8] A. von Manteuffel, C. Studerus, Top quark pairs at two loops and Reduze 2 (2012). doi:10.22323/1.151.0059. arXiv:1210.1436, [PoSLL2012,059(2012)].
- [9] P. Maierhöfer, J. Usovitsch, P. Uwer, Kira—A Feynman integral reduction program, *Comput. Phys. Commun.* 230 (2018) 99–112. doi:10.1016/j.cpc.2018.04.012. arXiv:1705.05610.
- [10] P. Maierhöfer, J. Usovitsch, Kira 1.2 Release Notes (2018). arXiv:1812.01491.
- [11] A. V. Smirnov, An algorithm to construct Groebner bases for solving integration by parts relations, *JHEP* 04 (2006) 026. doi:10.1088/1126-6708/2006/04/026. arXiv:hep-ph/0602078.

- [12] A. V. Smirnov, V. A. Smirnov, S-bases as a tool to solve reduction problems for Feynman integrals, Nucl. Phys. Proc. Suppl. 160 (2006) 80–84. doi:10.1016/j.nuclphysbps.2006.09.032. arXiv:hep-ph/0606247.
- [13] A. V. Smirnov, V. A. Smirnov, On the reduction of Feynman integrals to master integrals, PoS ACAT2007 (2007) 085. arXiv:0707.3993.
- [14] R. N. Lee, Group structure of the integration-by-part identities and its application to the reduction of multiloop integrals, JHEP 07 (2008) 031. doi:10.1088/1126-6708/2008/07/031. arXiv:0804.3008.
- [15] R. N. Lee, Presenting LiteRed: a tool for the Loop InTEgrals REDuction (2012). arXiv:1212.2685.
- [16] R. N. Lee, LiteRed 1.4: a powerful tool for reduction of multiloop integrals, J.Phys.Conf.Ser. 523 (2014) 012059. doi:10.1088/1742-6596/523/1/012059. arXiv:1310.1145.
- [17] A. von Manteuffel, R. M. Schabinger, Quark and gluon form factors to four-loop order in QCD: the N_f^3 contributions, Phys. Rev. D95 (2017) 034030. doi:10.1103/PhysRevD.95.034030. arXiv:1611.00795.
- [18] A. V. Smirnov, Algorithm FIRE – Feynman Integral REDuction, JHEP 10 (2008) 107. doi:10.1088/1126-6708/2008/10/107. arXiv:0807.3243.
- [19] A. V. Smirnov, V. A. Smirnov, FIRE4, LiteRed and accompanying tools to solve integration by parts relations, Comput.Phys.Commun. 184 (2013) 2820–2827. doi:10.1016/j.cpc.2013.06.016. arXiv:1302.5885.
- [20] A. V. Smirnov, FIRE5: a C++ implementation of Feynman Integral REDuction, Comput. Phys. Commun. 189 (2015) 182–191. doi:10.1016/j.cpc.2014.11.024. arXiv:1408.2372.
- [21] J. M. Henn, A. V. Smirnov, V. A. Smirnov, M. Steinhauser, A planar four-loop form factor and cusp anomalous dimension in QCD, JHEP 05 (2016) 066. doi:10.1007/JHEP05(2016)066. arXiv:1604.03126.

- [22] J. Henn, A. V. Smirnov, V. A. Smirnov, M. Steinhauser, R. N. Lee, Four-loop photon quark form factor and cusp anomalous dimension in the large- N_c limit of QCD, *JHEP* 03 (2017) 139. doi:10.1007/JHEP03(2017)139. arXiv:1612.04389.
- [23] R. N. Lee, A. V. Smirnov, V. A. Smirnov, M. Steinhauser, The n_f^2 contributions to fermionic four-loop form factors, *Phys. Rev. D* 96 (2017) 014008. doi:10.1103/PhysRevD.96.014008. arXiv:1705.06862.
- [24] A. Kurz, T. Liu, P. Marquard, A. V. Smirnov, V. A. Smirnov, M. Steinhauser, Light-by-light-type corrections to the muon anomalous magnetic moment at four-loop order, *Phys. Rev. D* 92 (2015) 073019. doi:10.1103/PhysRevD.92.073019. arXiv:1508.00901.
- [25] P. Marquard, A. V. Smirnov, V. A. Smirnov, M. Steinhauser, D. Wellmann, $\overline{\text{MS}}$ -on-shell quark mass relation up to four loops in QCD and a general $\text{SU}(N)$ gauge group, *Phys. Rev. D* 94 (2016) 074025. doi:10.1103/PhysRevD.94.074025. arXiv:1606.06754.
- [26] P. Marquard, A. V. Smirnov, V. A. Smirnov, M. Steinhauser, Four-loop wave function renormalization in QCD and QED, *Phys. Rev. D* 97 (2018) 054032. doi:10.1103/PhysRevD.97.054032. arXiv:1801.08292.
- [27] R. N. Lee, A. V. Smirnov, V. A. Smirnov, M. Steinhauser, Four-loop quark form factor with quartic fundamental colour factor (2019). arXiv:1901.02898.
- [28] T. Peraro, Scattering amplitudes over finite fields and multivariate functional reconstruction, *JHEP* 12 (2016) 030. doi:10.1007/JHEP12(2016)030. arXiv:1608.01902.
- [29] V. Sadovnichy, A. Tikhonravov, V. Voevodin, V. Opanasenko, "lomonosov": Supercomputing at moscow state university, in: *Contemporary High Performance Computing: From Petascale toward Exascale*, Chapman & Hall/CRC Computational Science, Boca Raton, United States, Boca Raton, United States, 2013, pp. 283–307.