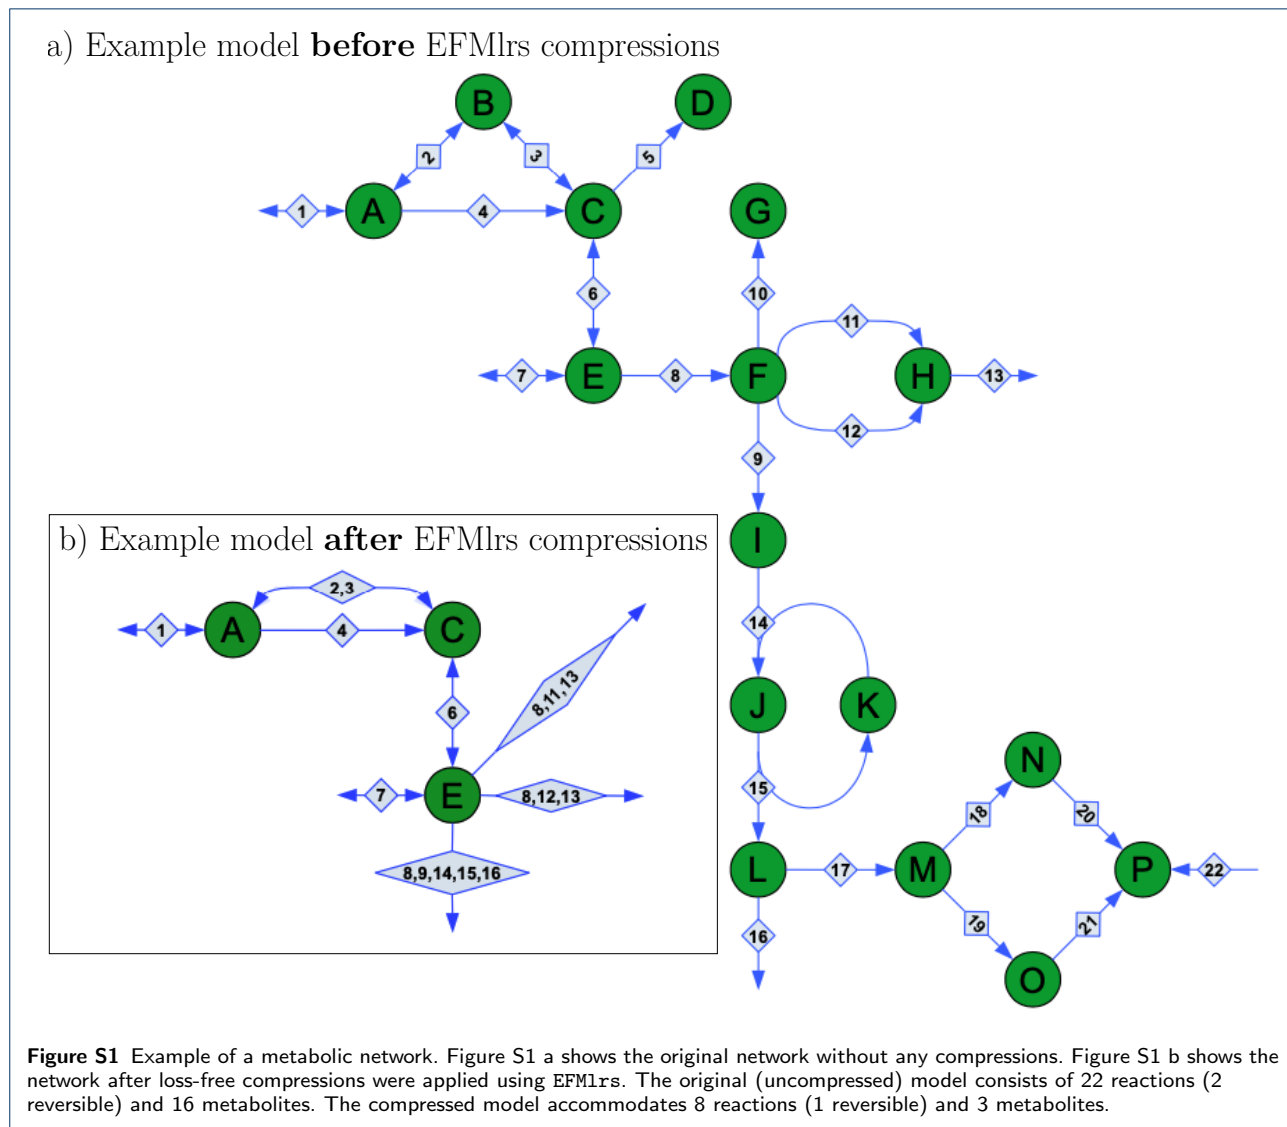


Supplementary material

EFM1rs

Compressions

The compression algorithms used by EFM1rs are already known in the metabolic modeling community, have been discussed e.g. by Gagneur et al. [22] and have been implemented in e.g. *efmtool* [12]. Table S1 compares the compression results of EFM1rs and *efmtool*. Since the compressions of both tools are based on the same algorithms, their results are very similar as well. However, for 3 out of 4 models (Table 1) used in this paper, EFM1rs could achieve a *stronger* compression and *efmtool* was not able to further compress any models compressed by EFM1rs. Only for the *EColiCore2* [29] model the compression results of both tools were equivalent. A compression comparison the other way round - first compressions by *efmtool* and then by EFM1rs could not be done, since *efmtool* does not output the compressed files needed for this.



We attribute the slightly different compression results to the different order and amount of iteration steps found in the respective implementations of the compression algorithms. However, at this moment the exact reason is unknown and subject of future investigations. It should be noted that the *stronger* compressions of EFM1rs together with the implementation in Python and the usage of *SymPy* lead to longer compression times compared to *efmtool* [12] which is implemented in Java - a statically typed and compiled and therefore faster programming language compared to Python. However, since in contrast to *efmtool* the compressions of EFM1rs are not directly coupled with the following calculations and the compressions have to be done only once, the overall time loss is not a big factor when calculating large models. For a detailed description of EFM1rs workflow and the compression algorithms see the supplementary material.

EFM1rs workflow

pre-processing

The EFM1rs workflow begins with parsing a *sbml* input file that contains the model of interest. Therefore, EFM1rs utilizes *COBRApy* functions for parsing and preparing metabolic models and thereby features model editing if necessary or desired e.g. excluding specific compartments or specifying if non-default reaction bounds should be taken into account, thus resulting in computing EFVs instead of EFM.

Table S1 Compression comparison of EFM1rs versus efmtool. Values in table are reactions (reversible reactions) / metabolites.

	uncompressed model	EFM1rs compressions	efmtool compressions
EColiCentral [28]	71 (15) / 53	44 (11) / 21	44 (11) / 26
EColiCore2 [29]	82 (22) / 54	58 (18) / 30	58 (18) / 30
iPS189 [30]	277 (21) / 271	63 (13) / 35	67 (13) / 42
JCVI-syn3A [31]	316 (8) / 286	100 (7) / 48	100 (7) / 50

loss-free compressions of N

Parsing the model is followed by the most important step during pre-processing – the loss-free compression of the stoichiometric matrix N . Iteratively, four different loss-free compressions are applied on N . All compression steps are executed one after another in a large outer loop, with each compression step itself being executed within a smaller loop. The inner loops of the single compression steps are executed until no more redundancies are found during this step. Only then the next compression step starts. After all compression steps have been run through, the large outer loop starts again until no redundancies can be found during any compression step. The complete sequence of pre-processing is illustrated in the top part of Figure 3. A pseudo-code snippet for each compression step is shown in the Algorithms 1, 2, 3 and 4.

The targets of the first compression step (Algorithm 1) are so-called dead-end metabolites. The metabolites D , G and P in Figure S1 are examples for such metabolites. They can be identified by analyzing the rows of the stoichiometric matrix N . Rows that contain only values with the same sign, indicate that the corresponding metabolites are only produced (positive values) respectively consumed (negative values). Therefore they can not be in a steady-state and consequently, they, as well as their contributing, irreversible reactions, can be removed from the network [12]. Thus the metabolites D , G and P together with their corresponding reactions can be removed from the metabolic network.

Algorithm 1 Pseudo-code snippet: *deadend*

```

1: function FINDS AND REMOVES REDUNDANCIES DUE TO DEADEND METABOLITES( $N_{m,n}$ ,  $rev$ )
  ▷ finds and removes deadend metabolites and corresponding irreversible reactions from  $N$ 
  ▷  $N$  is the stoichiometric matrix that denotes  $m \times n$ 
  ▷  $rev$  is a vector containing the reaction reversibilities
2:   for  $m \in N$  do
3:     if all  $n \in m \geq 0$  then
4:       if all  $rev_{i \in n_i > 0} = \text{False}$  then
5:         remove  $m$  from  $N$ 
6:       end if
7:     else if all  $n \in m \leq 0$  then
8:       if all  $rev_{i \in n_i < 0} = \text{False}$  then
9:         remove  $m$  from  $N$ 
10:      end if
11:    end if
12:  end for
13:  for  $n \in N$  do
14:    if  $|n| = 0$  then
15:      remove  $n$  from  $N$ 
16:      remove  $rev_n$  from  $rev$ 
17:    end if
18:  end for
19:  return  $\Delta N$ ,  $\Delta rev$ 
20: end function

```

The next compression step (Algorithm 2) is called *many2one* as during this step, reactions with unique fluxes are merged together. To further illustrate what unique fluxes mean, let's have a look at metabolite F in Figure S1 a. F is only produced by reaction $R8$ and afterward consumed by the reactions 9, 11, 12 (note that $R10$ has already been removed during dead-end compression). So unique fluxes can easily be identified by analyzing the rows of N . Thus the rows corresponding to uniquely produced (respectively consumed) metabolites have one positive entry and otherwise only negative entries, respectively one negative entry and otherwise only positive entries. Since we operate under steady-state conditions, $R8$ has to carry a non-zero flux, whenever the consuming reactions of F are active. Hence, it's possible to merge the sequential consuming (respectively producing) reactions and thereby further decrease the dimensions of the stoichiometric matrix N [22]. So, during a first iteration step of the *many2one* compression, we can lump the reactions $R8 + R11$, $R8 + R12$ and $R8 + R9$ together and remove the metabolite F . Note that this procedure is executed in loops and already merged reactions can be merged further with other reactions if during a next iteration further redundancies are found and the preconditions for the compressions are met. This applies to all compression steps.

Redundancies that only can be detected by analyzing the kernel K of the stoichiometric matrix N are the targets of the next compression step (Algorithm 3), the *nullspace* compression. If two rows in the kernel matrix K only differ from each other by a constant factor, also their fluxes only differ by this factor. If there is a flux through one of these reactions, there has to be a flux through the other, respectively if one reaction has no flux the other reaction is also passive. In Figure S1 a $R2$ and $R3$ are examples for such reactions that are called coupled reactions as they are co-regulated [12, 35]. Thus they are *forced*

Algorithm 2 Pseudo-code snippet: *many2one*

```

1: function FINDS REDUNDANCIES DUE TO UNIQUE FLUXES( $N_{m,n}, rev$ )
  ▷ finds and merges reactions with unique fluxes in  $N$ 
  ▷  $N$  is the stoichiometric matrix that denotes  $m \times n$ 
  ▷  $rev$  is a vector containing the reaction reversibilities
2:   for  $m \in N$  do
3:     skip  $\leftarrow$  False
4:     for  $n \in m$  do
5:       if  $n \neq 0$  and  $rev_n = 1$  then
6:         skip  $\leftarrow$  True
7:       end if
8:     end for
9:     if skip = False then
10:      for  $n \in m$  do
11:        if  $n \geq 0$  then
12:           $a \leftarrow a$  append  $n$ 
13:        else
14:           $b \leftarrow b$  append  $n$ 
15:        end if
16:      end for
17:      pos  $\leftarrow |a|$ 
18:      neg  $\leftarrow |b|$ 
19:      if (neg = 1 or pos = 1) and
        (neg + pos  $\geq$  2) then
20:        if neg = 1 then
21:          swap  $a, b$ 
22:        end if
23:        for  $i \in b$  do
24:           $n_a \leftarrow n_i / |i| + n_a / |a|$ 
25:        end for
26:        remove  $n_a$  from  $N$ 
27:        remove  $rev_a$  from  $rev$ 
28:        for  $m \in N$  do
29:          if  $|m| = 0$  then
30:            remove  $m$  from  $N$ 
31:          end if
32:        end for
33:      end if
34:    end if
35:  end for
36:  return  $\Delta N, \Delta rev$ 
37: end function

```

to work together, they can also be lumped into one reaction by applying their coupling factor. In the example network (S1 a), we can merge $R2$ and $R3$ into one single reaction and thereby remove one reaction and the metabolite B from the network.

Algorithm 3 Pseudo-code snippet: *nullspace*

```

1: function FINDS AND REMOVES REDUNDANCIES DUE TO COUPLED REACTIONS( $N_{m,n}, rev$ )
  ▷ finds coupled reactions in  $K$  and merges them in  $N$ 
  ▷  $N$  is the stoichiometric matrix that denotes  $m \times n$ 
  ▷  $rev$  is a vector containing the reaction reversibilities
2:    $K_{i,j} \leftarrow$  kernel matrix of  $N$ 
3:   for  $i1 \in K$  do
4:     for  $i2_{i+1} \in K$  do
5:       factor  $\leftarrow i1 / i2_{i+1}$ 
6:       if factor  $\neq$  None then
7:          $n_{i1} \leftarrow n_{i1} + n_{i2} /$  factor
8:         if  $rev_{i2} =$  False then
9:            $rev_{i1} \leftarrow$  False
10:        end if
11:        remove  $n_{i2}$  from  $N$ 
12:        remove  $rev_{i2}$  from  $rev$ 
13:        remove  $i2$  from  $K$ 
14:      end if
15:    end for
16:  end for
17:  for  $m \in N$  do
18:    if  $|m| = 0$  then
19:      remove  $m$  from  $N$ 
20:    end if
21:  end for
22:  return  $\Delta N, \Delta rev$ 
23: end function

```

The last compression step (Algorithm 4) is called *echelon* compression, as the redundancies it targets can be identified by analyzing the reduced row echelon form of the stoichiometric matrix N . These redundancies are caused by metabolites that are in a so-called conservative relation with each other, meaning that there is a linear dependency between them [22]. Such linear dependencies show in the transposed reduced echelon form as a column that breaks the diagonal pattern of ones. In Figure S1 a the metabolites K and J are linearly dependent on each other and thereby build a conservative relation. K can only be produced if J is produced. However, J as well can only be produced while K is being produced. Therefore the reactions $R14$ and $R15$ are strongly linked thus any change in the concentration of the metabolite J leads to the same change in the concentration of the metabolite K . Hence, one of the two metabolite, in this case K , can be removed from the stoichiometric matrix N without losing any necessary information.

If no more redundancies can be found during any compression step, the flux cone (9) is reconfigured by splitting all reversible reactions into two – a forward and a backward reaction. Thereby the flux cone is transformed for calculations with `mplrs`. Afterwards, the compressed input files for `mplrs` and `efmtool`, as well as a log file and an info file, concerning the performed compressions and later needed for decompressions, are created. Now, the computations using either `mplrs` or `efmtool` can be started.

post-processing

When calculating EFM/Vs from a loss-free compressed network the number of resulting EFM/Vs is the same as if the calculations were done with the original network. However, the EFM/Vs are compressed and are still containing the merged reactions. Additionally, `mplrs`' output files also contain the previously split reversible reactions which need to be lumped together again. Thus, post-processing and decompression are needed to get the *full* set of EFM/Vs.

The bottom part of Figure 3 shows the main steps during post-processing. First, the info file, containing all information on the previously applied compressions, and the `mplrs`, respectively, `efmtool`'s output files are read. Since `mplrs`' output besides the calculated modes also includes the split reactions, as well as additional information, it requires additional steps before decompressions can be started. Therefore users need to specify whether `mplrs`' or `efmtool`'s output is to be decompressed. After parsing all input files a reverse decompression stack is built and each EFM, respectively EFV, is being decompressed one after the other in reverse order of the previously applied compressions until all modes are decompressed and written to a specified output. The resulting decompressed EFMs now contain all reactions as the original, uncompressed model and are ready for further analysis.

Algorithm 4 Pseudo-code snippet: *echelon*

```

1: function FINDS AND REMOVES REDUNDANCIES DUE TO CONSERVATIVE RELATIONS( $N_{m,n}$ )
  ▷ finds conservative relations and removes redundant metabolites from  $N$ 
  ▷  $N$  is the stoichiometric matrix that denotes  $m \times n$ 
2:    $E_{i,j} \leftarrow$  reduced row echelon form of  $N$ 
3:    $k \leftarrow 0$ 
4:   for  $i \in E$  do
5:     while  $i_k \neq 1$  do
6:        $v \leftarrow v$  append  $k$ 
7:        $k \leftarrow k + 1$ 
8:     end while
9:      $k \leftarrow k + 1$ 
10:  end for
11:  remove all  $m \in v$  from  $N$ 
12:  return  $\Delta N$ 
13: end function

```

