Computing elementary flux modes involving a set of target reactions

Laszlo David [1,2,3,*], Alexander Bockmayr [1,2]

November 2013

Abstract

Elementary flux mode (EM) analysis is an important tool in the constraint-based analysis of genome-scale metabolic networks. Due to the combinatorial complexity of these networks, as well as the advances in the level of detail to which they can be reconstructed, an exhaustive enumeration of all EMs is often not practical. Therefore, in recent years interest has shifted towards searching EMs with specific properties.

We present a novel method that allows computing EMs containing a given set of target reactions. This generalizes previous algorithms where the set of target reactions consists of a single reaction. In the one-reaction case, our method compares favorably to the previous approaches. In addition, we present several applications of our algorithm for computing EMs containing two target reactions in genome-scale metabolic networks.

A software tool to demonstrate the algorithms described in this paper is available at `https://sourceforge.net/projects/caefm`.

Contact: Laszlo.David@fu-berlin.de, Alexander.Bockmayr@fu-berlin.de

1 Introduction

Elementary (flux) modes (EMs) (Schuster and Hilgetag, 1994; Schuster *et al.*, 1999, 2000, 2002; Papin *et al.*, 2004) are an important concept for the structural analysis of metabolic networks, with many practical applications (see e.g. (Zanghellini *et al.*, 2013) and references therein). As a consequence, the development of methods for the computation of EMs has become an active research area over the past years (Gagneur and Klamt, 2004; von Kamp and Schuster, 2006; Terzer and Stelling, 2008; Terzer, 2009; de Figueiredo *et al.*, 2009; Rezola *et al.*, 2011, 2013). The computational complexity of enumerating all EMs is not known (Acuña *et al.*, 2009). However, there exist several algorithms and software packages for an exhaustive enumeration in a given metabolic network (Gagneur and Klamt, 2004; Terzer and Stelling, 2008; von Kamp and Schuster, 2006; Terzer, 2009). While these methods work very well for small networks, due to the possibly exponential number of EMs, they may fail for medium or large genome-scale networks. With the ever increasing size of genome-scale metabolic network reconstructions, EM analysis nowadays can often be used only under additional assumptions (e.g., modifying the system boundary of the network, blocking a large number of uptake reactions etc.). These extra

*to whom correspondence should be addressed

[1] Department of Mathematics and Computer Science, Freie Universität Berlin, Arnimallee 6, 14195 Berlin, Germany

[2] DFG-Research Center Matheon, Berlin, Germany

[3] Berlin Mathematical School (BMS), Berlin, Germany

assumptions may have the bad side effect of changing the structure of the network, sometimes introducing artificial pathways (Marashi *et al.*, 2012).

One way to deal with genome-scale networks is to define a subset of interest $\mathcal{R}_I$ of the full reaction set $\mathcal{R}$, without altering the network topology. Kaleta *et al.* (2009) look for sets of reactions in $\mathcal{R}_I$, called flux patterns, which indicate the existence of an EM having those reactions in its support. They enumerate a basis for this set, the elementary flux patterns. Urbanczik and Wagner (2005) and Marashi *et al.* (2012) project the steady-state flux cone onto the subspace defined by $\mathcal{R}_I$, and enumerate the partial EMs for this subspace.

Given the difficulty of computing and analyzing the full set of all EMs, recent research has focused on finding a special subset of EMs (de Figueiredo *et al.*, 2009; Rezola *et al.*, 2011, 2013). de Figueiredo *et al.* (2009) describe a mixed-integer programming method to enumerate the k shortest EMs ($k \geq 1$). Their method has been extended to find shortest EMs involving one reaction of choice and also to enumerate a minimal generating set of EMs (Rezola *et al.*, 2011). In the last paper, the authors note that their method cannot be applied to find elementary modes involving two predefined reactions.

The problem of finding an EM involving two or more given reactions was also considered by Acuña *et al.* (2009). They give a more general formulation, in that they not only want their EM to involve a certain set of target reactions $T \subset \mathcal{R}$ (with $|T| = t > 0$), but also to avoid another set of reactions $F \subset \mathcal{R}$ (with $F \cap T = \emptyset$). In other words, the goal is to find an EM e, with $e_i \neq 0$ for $i \in T$ and $e_i = 0$ for $i \in F$. The authors show that defining a set of reactions F to be avoided does not add to the difficulty of the problem, and in fact reactions belonging to F can simply be removed from the network. Acuña *et al.* (2010) study the complexity and prove that the decision problem, whether an EM involving two or more target reactions exists, is not solvable in polynomial time, unless $P = NP$.

To the best of our knowledge, no algorithm for this problem has been published so far. Here, we develop a mixed-integer programming approach for the more general problem of computing k elementary modes ($k \geq 1$) involving a given set of target reactions T, for $|T| \geq 2$. Computational experiments show that the method can be applied even to large genome-scale networks.

2 Preliminaries

We consider a metabolic network N with a set of metabolites $\mathcal{M}$ and a set of reactions $\mathcal{R}$. Formally, the network N is defined as $N = (\mathcal{M}, \mathcal{R}, S, Irr)$, with a stoichiometric matrix $S \in \mathbb{R}^{\mathcal{M}\times\mathcal{R}}$, and a set of irreversible reactions $Irr \subseteq \mathcal{R}$. For $I \subseteq \mathcal{M}$ and $J \subseteq \mathcal{R}$, we denote by S_{IJ} the submatrix of S defined by the rows from I and the columns from J. Any subset of reactions $\mathcal{R}' \subseteq \mathcal{R}$ defines a metabolic *subnetwork* $N' = (\mathcal{M}, \mathcal{R}', S', Irr')$, with $S' = S_{N'} = S_{\mathcal{M}\mathcal{R}'} \in \mathbb{R}^{\mathcal{M}\times\mathcal{R}'}$ and $Irr' = Irr \cap \mathcal{R}'$.

The stoichiometric constraints $Sv = 0$ (mass balance) and the thermodynamic constraints $v_{Irr} \geq 0$ define the *steady-state flux cone* $C = \{v \in \mathbb{R}^{\mathcal{R}} \mid Sv = 0, v_{Irr} \geq 0\}$. Any vector $v \in C\backslash\{0\}$ is called a *flux mode*. The *support* of v is the set $supp(v) = \{j \in \mathcal{R} \mid v_j \neq 0\}$ of active reactions in v. A reaction that is not active in any flux mode $v \in C$ is called *blocked*. A flux mode $v \in C \setminus \{0\}$ with minimal support (with respect to $\subseteq$) is called an *elementary mode* (EM). The goal of this paper is to study the following *problem* ($\mathscr{P}$):

> Given a metabolic network $N = (\mathcal{M}, \mathcal{R}, S, Irr)$, a set of target reactions $T \subseteq \mathcal{R}$ and $k \geq 1$, compute a set E of EMs in N, $|E| = k$, such that $supp(e) \supseteq T$, for all $e \in E$.

Flux coupling analysis (Burgard *et al.*, 2004) defines four binary relations for pairs of unblocked reactions $i, j \in \mathcal{R}$. We say that i is *directionally coupled* to j and write $i \to j$, if $v_i \neq 0$ implies $v_j \neq 0$, for all $v \in C$. The reactions i and j are *partially coupled* if both $i \to j$ and

$j \to i$ hold. Two partially coupled reactions i and j are *fully coupled* if there is $\lambda \neq 0$ such that $v_i = \lambda v_j$, for all $v \in C$. If neither $i \to j$ nor $j \to i$, then the reactions i and j are uncoupled. Following (Marashi and Bockmayr, 2011), we call two uncoupled reactions i and j *mutually exclusive* if they are never active together in the same EM, otherwise they are called *sometimes coupled*.

For the rest of this paper, we assume that all reactions are irreversible, i.e., $Irr = \mathcal{R}$, and that none of the reactions is blocked. These assumptions do not limit the applicability of our methods. Reversible reactions can be split into a forward and backward reaction. In general, this operation may induce a number of artificial EMs in the form of two-cycles. However, these two-cycles do not make problem ($\mathscr{P}$) more difficult to solve. In fact, the only case where such two-cycles satisfy the conditions of problem ($\mathscr{P}$) is for $|T| = 1$. In this case, there is exactly one artificial EM that needs to be filtered out from the final solution set. For $|T| > 1$ no pairs of split reactions will be part of the same EM. All blocked reactions can be identified in polynomial time, by solving a linear number of linear programs. Afterwards, they can be removed from the network without altering the underlying flux cone.

3 Methods

In the following, we divide the general problem ($\mathscr{P}$) into three subproblems and discuss them individually: the one-reaction case, the two-reaction case and the general t-reaction case (where $t = |T| > 2$). The underlying details vary in each case, but there is a general concept followed by all three methods. In every case, we aim to incrementally find an alternating sequence N^1, e^1, N^2, e^2, ..., N^k, e^k, of subnetworks and EMs of the input network N with the following properties for all $i \in \{1, ..., k\}$:

1. The target reactions r_1, ..., r_t are part of every subnetwork N^i.
2. In every subnetwork N^i, reaction r_1 is directionally coupled to the reactions r_2, ..., r_t.
3. No subnetwork N^i has as flux mode any of the EMs e^l for $l < i$.
4. e^i is an EM in N^i involving r_1.

Clearly, the main difficulty in finding such a sequence of subnetworks is in imposing condition (2). In turn, once a new subnetwork $N^i = (\mathcal{M}, \mathcal{R}^i, S^i)$ has been found, a corresponding EM e^i can be computed by solving the linear program LP(N^i) (Acuña *et al.*, 2009).

$$
\begin{array}{rlrcll}
\text{LP}(N^i)\text{: } & \min & 0 & & & \\
& \text{s.t.} & S^i v & = & 0, & \\
& & v_{r_1} & \geq & 1, & \\
& & v_j & \geq & 0, & \forall j \in \mathcal{R}^i.
\end{array}
$$

Using a simplex-based method, we can compute a basic feasible solution e^i of LP(N^i), which corresponds to a vertex of the truncated flux cone of N^i and thus defines an EM in N^i. Due to the conservation property of EMs (see e.g. Lemma 1 in (Marashi and Bockmayr, 2011)), e^i is also an EM in N. The set $E^i := \{e^1, ..., e^i\}$, $i \geq 1$, contains EMs of the original network N that involve every target reaction. Terminating the search after k EMs have been found provides a solution to problem $\mathscr{P}$.

3.1 The one-reaction problem

If the set of target reactions consists of only one reaction, condition (2) is trivially satisfied for every subnetwork satisfying condition (1). Assuming N^1, e^1, N^2, e^2, ... N^i, e^i have already been computed, we can determine a new subnetwork N^{i+1} by solving the following mixed-integer linear program (MILP1).

$$\begin{array}{rrcll} \text{MILP1}(E^i): \min & 0 & & & \\ \text{s.t.} & Sv & = & 0, & \\ & v_{r_1} & \geq & 1, & \\ a_l \leq & v_l & \leq & Ma_l, & \forall l \in \mathcal{R}, \\ \sum_{l \in supp(e^u)} & a_l & \leq & |supp(e^u)| - 1, & \forall u \in [1..i], \\ & v_l & \geq & 0, & \forall l \in \mathcal{R}, \\ & a_l & \in & \{0,1\}, & \forall l \in \mathcal{R}. \end{array}$$

The constraints in (MILP1) are the same as for the computation of the so-called shortest EMs in (de Figueiredo *et al.*, 2009). There are two groups of variables: v represents the steady-state flux values of the reactions, while the 0-1 vector a models the support of v (i.e., $v_l > 0 \Leftrightarrow a_l = 1$). The variables v_l and a_l are linked by the 3rd constraint, using a suitably large constant $M > 0$. An important difference to (de Figueiredo *et al.*, 2009) is the objective function (see also Sect. 4.2). We do not try to find the smallest set of reactions satisfying the constraints of (MILP1). Instead computing any feasible solution is sufficient. This turns out to be enough to derive a subnetwork that satisfies conditions (1-3). Indeed, given a feasible solution (v', a') of (MILP1), define N^{i+1} by $\mathcal{R}^{i+1} := \{l \in \mathcal{R} \mid a'_l = 1\}$. This subnetwork N^{i+1} clearly satisfies conditions (1) and (2), while the 4th constraint in (MILP1), the so-called *no-good constraints*, guarantees condition (3). By solving $\text{LP}(N^{i+1})$, we obtain an EM e^{i+1}.

Table 1 summarizes the method for the one-reaction case. The two conditional exit points of the algorithm are Step 2 and 6. If the algorithm terminates at Step 2, Prop. 1 assures all EMs will be found. In contrast, if the exit occurs at Step 6, we enumerate k EMs.

Step	Action
0.	Initialize $i := 1, E := \emptyset$.
1.	Try to find a feasible solution (v', a') of MILP1(E).
2.	If MILP1(E) is infeasible, then STOP.
3.	Otherwise, use (v', a') to derive subnetwork N^i.
4.	Find a basic feasible solution e^i of $\text{LP}(N^i)$.
5.	Let $E := E \cup \{e^i\}$ and $i := i + 1$.
6.	If $i > k$ then STOP.
7.	Go to Step 1.

Table 1: Algorithm 1 for the one-reaction case.

Proposition 1. *For any EM $e \notin E$ and sufficiently large $M > 0$, there is a flux mode v' in N such that $(v', supp(e))$ is a feasible solution for MILP1(E).*

A formal proof of Prop. 1 is omitted here. It can be easily obtained by suitably scaling the vector e.

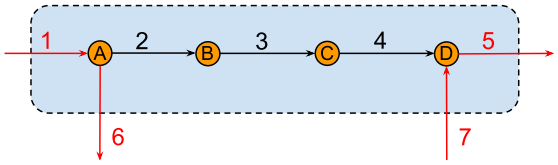

Figure 1: Example network. Thick arrows represent the optimal solution of (MILP2)

We note that it is possible to initialize $N^1 := N$, thus avoiding the need to solve the very first mixed-integer program. The computationally hard part is Step 1 of the algorithm, while the rest of the steps can be computed in polynomial time.

3.2 The two-reaction problem

A natural idea to find a shortest EM involving a pair of reactions $\{r_1, r_2\} \subseteq \mathcal{R}$ would be to extend the previous method by forcing both r_1 and r_2 to be active, while minimizing the total number of active reactions:

$$
\begin{array}{rlrlll}
\text{(MILP2)}: \min & \sum_{l\in\mathcal{R}} a_l & & & & \\
\text{s.t.} & Sv & = & 0, & & \\
& v_{r_1} & \geq & 1, & & \\
& v_{r_2} & \geq & 1, & & \\
a_l \leq & v_l & \leq & Ma_l, & \forall l \in \mathcal{R}, & \\
& v_l & \geq & 0, & \forall l \in \mathcal{R}, & \\
& a_l & \in & \{0,1\}, & \forall l \in \mathcal{R}. &
\end{array}
$$

(MILP2) indeed finds a shortest flux mode containing r_1 and r_2. However, it may fail to produce an EM. As illustrated by Fig. 1, if we apply (MILP2) for reactions 1 and 5, a non-elementary flux mode will be found, involving the reactions 1,6,7,5. However, an optimal solution (v^*, a^*) of (MILP2) still has interesting properties, which will turn out to be useful for refining our method. These are described in the following propositions.

Proposition 2. *Let $v^* = \sum_{i=1}^{s} \lambda_i e^i$, with $\lambda_i > 0$, be any decomposition of an optimal solution (v^*, a^*) of (MILP2) into s EMs in N. Then for all $i \in [1..s]$, $e^i_{r_1} > 0$ or $e^i_{r_2} > 0$. There exists $i \in [1..s]$ with $e^i_{r_1} > 0$ and $e^i_{r_2} > 0$, if and only if $s = 1$.*

Proof. Suppose there exists $i \in [1..s]$ such that $e^i_{r_1} = 0$ and $e^i_{r_2} = 0$. Let $p := \min\{v^*_j / e^i_j \mid j \in supp(e^i)\}$ and $v' := v^* - pe^i$. Then there exist $\lambda > 0$ and $a' \in \{0,1\}^{\mathcal{R}}$ such that $(\lambda v', a')$ is a feasible solution of (MILP2), with $\sum_{l\in\mathcal{R}} a'_l < \sum_{l\in\mathcal{R}} a^*_l$, in contradiction to the optimality of v^*.

From $s = 1$, we get $v^* = \lambda_1 e^1, \lambda_1 > 0$, and therefore $e^i_{r_1} > 0$ and $e^i_{r_2} > 0$. Conversely, suppose there exists $i \in [1..s]$ such that $e^i_{r_1} > 0$ and $e^i_{r_2} > 0$. Up to scaling, e^i is a feasible solution of (MILP2) with $supp(e^i) \subseteq supp(v^*)$. From the optimality of v^*, we get $supp(e^i) = supp(v^*)$, which implies $s = 1$. □

Prop. 2 shows that any EM participating in a decomposition of v^* must contain at least one of the two target reactions. The only case when an EM contains both r_1 and r_2 is when the optimal solution v^* itself is an EM.

The result can be formulated in a stronger form, extending it to the whole subnetwork $N^* = (\mathcal{M}, \mathcal{R}^*, S_{\mathcal{M}\mathcal{R}^*})$ defined by v^* (with $\mathcal{R}^* = supp(v^*)$). Every EM e in this subnetwork will satisfy $e_{r_1} > 0$ or $e_{r_2} > 0$. If v^* is itself an EM, then also the subnetwork N^* will have only one EM, namely v^*. The following Corollary summarizes our previous results.

Corollary 3. *In the subnetwork N^* defined by v^*, the reactions r_1 and r_2 are either fully coupled or mutually exclusive.*

Our next proposition shows that any flux mode using the reactions r_1 and r_2 can be scaled by a positive factor so that it becomes a feasible solution of (MILP2).

Proposition 4. *Let $v \in C$ with $v_{r_1} > 0$ and $v_{r_2} > 0$. Then for sufficiently large $M > 0$, there exists a feasible solution (v', a') of (MILP2), such that $v' = \lambda v$, for some $\lambda > 0$.*

Proof. Let $\lambda := 1/\min\{v_i \mid i \in supp(v)\}$ and $M := \max\{v_i/v_j \mid i, j \in supp(v)\}$. Now consider $v' := \lambda v$ and let $a'_l := 1$, for all $l \in supp(v)$, and $a'_l := 0$, for all $l \in \mathcal{R} \setminus supp(v)$. For all $l \in supp(v)$, $v'_l = v_l/\min\{v_i \mid i \in supp(v)\} \geq 1$. Thus, $v'_{r_1} \geq 1$ and $v'_{r_2} \geq 1$. By definition of a', for all $l \in supp(v)$, $a'_l = 1 \leq v'_l = v_l/\min\{v_i \mid i \in supp(v)\} \leq \max\{v_i/v_j \mid i, j \in supp(v)\} = M = Ma'_l$. Since $v' \in C$ also holds, we conclude that (v', a') is a feasible solution of (MILP2). □

In general, when decomposing an arbitrary steady-state flux vector into EMs, the number of participating EMs ranges from using only one to many. Moreover, a decomposition does not necessarily have to be unique. It turns out that in our special case, a decomposition is much more constrained. In fact, no decomposition can contain more than two EMs. As direct corollary of this result, we will also get the uniqueness of the decomposition.

Proposition 5. *Let (v^*, a^*) be an optimal solution of (MILP2) and let $v^* = \sum_{i=1}^{s} \lambda_i e^i$, with $\lambda_i > 0$, for all $i \in [1..s]$, be a decomposition of v^* into s EMs with pairwise different support. Then $s \leq 2$.*

Proof. Assume $s > 2$. From Prop. 2 it follows that there exist $i, j \in [1..s]$ such that e^i and e^j both contain r_1 or both contain r_2. Without loss of generality, we assume that e^i and e^j both contain r_1. Then, since $s \neq 1$, both e^i and e^j do not contain r_2. Let $\mathcal{R}_{diff} := supp(e^j) \setminus supp(e^i)$. Since e^i is an EM, we have $\mathcal{R}_{diff} \neq \emptyset$. Let $p := \min\{v^*_r/e^j_r \mid r \in \mathcal{R}_{diff}\}$. The vector $v' := v^* - pe^j$ satisfies $supp(v') \subsetneq supp(v^*)$. However, it need not be feasible for (MILP2) because it may violate some constraint $v_r \geq 0$. Let $V := \{r \in \mathcal{R} \mid v'_r < 0\} \subseteq supp(e^j) \cap supp(e^i)$. If $V \neq \emptyset$, define $q := \min\{v'_r/e^i_r \mid r \in V\} < 0$, otherwise $q := 0$. Then $v'' := v' + (1-q)e^i \geq 0, v''_{r_1} > 0, v''_{r_2} > 0$, and we still have $supp(v'') \subsetneq supp(v^*)$. By Prop. 4, this implies the existence of a feasible solution for (MILP2) with a smaller objective function value than v^*, which is a contradiction. □

Corollary 6. *Any decomposition of v^* into EMs is unique.*

Proof. The result is trivial if v^* is itself an EM. Thus we only have to consider the case $s = 2$. Assume by contradiction that $v^* = \lambda_1 e^1 + \lambda_1 e^2 = \mu_1 e^3 + \mu_2 e^4$, where at least three of e^1, e^2, e^3, e^4 have pairwise different support. Clearly, then $v^* = (\lambda_1 e^1 + \lambda_1 e^2 + \mu_1 e^3 + \mu_2 e^4)/2$, which contradicts the result of Prop. 5. □

Of special interest is Cor. 3, which asserts that in the subnetwork N^* defined by v^*, r_1 and r_2 are either fully coupled or mutually exclusive. Thus, to make the optimal solution of (MILP2) an EM, it is enough to add constraints that exclude the second case. These additional constraints must disallow the existence of an EM in N^* involving exactly one of r_1 and r_2. This can be achieved by requiring that r_1 should be directionally coupled to r_2 in N^* (or alternatively that r_2 is directionally coupled to r_1). To formulate this mathematically, we delete r_2 from N^* and require that r_1 is blocked in the resulting subnetwork N'. More formally, if $S' := S_{*,supp(v^*)\setminus\{r_2\}}$ is the stoichiometric matrix of N', then the following system should be infeasible:

$$\begin{aligned} S'z &= 0, \\ z_{r_1} &= 1, \\ z &\geq 0. \end{aligned}$$

By applying Farkas' Lemma (see e.g. (Schrijver, 1998)), this infeasibility requirement can be turned into a feasibility condition in the dual space. Let $y \in \mathbb{R}^m$, $x \in \mathbb{R}$. Then the following system in y and x should be feasible:

$$\begin{aligned} (S')^T y + u^{r_1} x &\geq 0, \\ x &\leq -1. \end{aligned}$$

Here u^r is the r-th unit vector (with an entry 1 for component r, and 0 otherwise) and $.^T$ denotes transposition of a matrix.

This formulation inherently uses S', the stoichiometric matrix of the new subnetwork. Naturally, information about it is not derivable independently. Thus the first constraint set needs to be reformulated to dynamically adjust itself according to the current solution (v, a) of (MILP2). This leads to the following constraints:

$$\begin{aligned} S^T y + u^{r_1} x &\geq -M(\mathbf{1} - a + u^{r_2}), \\ x &\leq -1, \end{aligned} \qquad \mathrm{DirC}(r_1, r_2)$$

where $\mathbf{1}$ denote a vector all whose components are 1.

By using a large enough constant M, the first constraint becomes trivially satisfiable for inactive reactions ($a_l = 0$) and for r_2, where the right-hand side simplifies to $-M$. In contrast, for active reactions ($a_l = 1$) different from r_2, the right-hand side sums up to 0, thus effectively activating the constraint. The inequalities $\mathrm{DirC}(r_1, r_2)$ are called *directional coupling constraints for r_1 implying r_2*. Extending (MILP2) with $\mathrm{DirC}(r_1, r_2)$, allows computing a shortest EM through r_1 and r_2. The following Prop. 7 summarizes our construction. It guarantees that by adding the constraints $\mathrm{DirC}(r_1, r_2)$, any feasible solution of (MILP2) defines a subnetwork in which condition (2) is satisfied.

Proposition 7. *Let (v', a') be any feasible solution of (MILP2) augmented with the directional coupling constraints DirC(r_1, r_2). Then in the subnetwork N' defined by supp(v'), r_1 is directionally coupled to r_2.*

By iteratively adding no-good constraints corresponding to already found EMs, we are able to enumerate any number of EMs in an increasing order of length. The resulting mixed-integer linear program can be expected to work for smaller-scale network models, but due to the difficulty of proving optimality in mixed-integer linear programs, the algorithm will most likely turn impractical for larger models. The reason for the bottle-neck is clearly that in every iteration we aim to find the shortest EM not yet discovered. Similar to the one-reaction case, we next trade the optimality condition on the length for an easier to solve program. The final form of our method is given in the mixed-integer linear program (MILP3). Tab. 2 summarizes the algorithm in the two-reaction case.

$$
\begin{array}{rrcll}
\text{MILP3}(E^i): \min & 0 & & & \\
\text{s.t.} & Sv & = & 0 & \\
& v_{r_1} & \geq & 1, & \\
& v_{r_2} & \geq & 1, & \\
a_l \leq & v_l & \leq & M_0\, a_l, & \forall l \in \mathcal{R}, \\
& S^T y + u^{r_1} x & \geq & M_1(a - \mathbf{1} - u^{r_2}), & \\
& -x & \geq & 1, & \\
& \sum_{l \in supp(e^u)} a_l & \leq & |supp(e^u)| - 1, & \forall u \in [1..i], \\
& v_l & \geq & 0, & \forall l \in \mathcal{R}, \\
& a_l & \in & \{0,1\}, & \forall l \in \mathcal{R}, \\
& x, y_m & \in & \mathbb{R}, & \forall m \in \mathcal{M}.
\end{array}
$$

Step	Action
0.	Initialize $i := 1, E := \emptyset$.
1.	Try to find a feasible solution (v', a') of MILP3(E).
2.	If MILP3(E) is infeasible, then STOP.
3.	From (v', a') derive subnetwork N^i.
4.	Find a basic feasible solution e^i of LP(N^i).
5.	Let $E := E \cup \{e^i\}$ and $i := i + 1$.
6.	If $i > k$ then STOP.
7.	Go to Step 1.

Table 2: Algorithm 2 for the two-reaction case.

3.3 The general t-reaction case

Although this problem seems to be much harder at first sight, it turns out that the previous results provide all the ingredients necessary to tackle this general case. We propose two strategies for building a mixed-integer linear program that can be used in a similar fashion to (MILP3).

In the *cascade* method, we extend (MILP2) with DirC(r_1, r_2), DirC(r_2, r_3), ..., DirC(r_{t-1}, r_t). Based on Prop. 7 and the transitivity of directional coupling (Larhlimi *et al.*, 2012), in any feasible solution of this new MILP, reaction r_1 will imply reactions r_2, ..., r_t, thus satisfying condition (2). Similarly, in the *hub* method, DirC(r_1, r_2), DirC(r_1, r_3), ..., DirC(r_1, r_t) are added to (MILP2).

Alternative coupling strategies can be thought of. Indeed, by constructing any spanning tree on the vertices r_1, r_2,... r_t, with r_1 being the root of the tree, and taking the union of the directionality constraints corresponding to each edge, we create the conditions for a subnetwork where every desired reaction is directionally coupled to r_1. At this time, it is unclear whether there is a practical advantage in choosing one of these methods compared to the others. In all cases, the total number of constraints (and variables) is the same and grows linearly with t.

3.4 Flux uncoupling

The computational complexity of the problem to decide for a pair of uncoupled reactions whether they are sometimes coupled or mutually exclusive (Marashi and Bockmayr, 2011) is NP-complete (Acuña *et al.*, 2010). However, this problem can be seen as a special case of Sect. 3.2. Indeed,

rather than enumerating some (or all) EMs involving these two reactions, we are merely interested in the existence of any. For this purpose, it is enough to execute Step 1 of Algorithm 2. If MILP3($\emptyset$) is feasible (resp. infeasible), we can conclude that our two reactions are sometimes coupled (resp. mutually exclusive).

While this works for any pair of reactions, the above solution might not be optimal if our goal is to find all uncoupling relations. For a given pair of reactions, rather than stopping after Step 1 of Algorithm 2, one could continue executing Steps 2-7 and potentially compute an EM, if one exists. This EM can then be used to decide not only about the uncoupling relation for the current pair of reactions, but for other pairs as well. More specifically, once an EM e has been computed for two reactions i and j, this EM can also be used to show that other pairs of reactions in $supp(e)$ are sometimes coupled. This may greatly reduce the number of pairs for which solving a MILP is necessary. Moreover, by performing FCA (David *et al.*, 2011; Larhlimi *et al.*, 2012) as a heuristic presolving step, many uncoupling pairs can be deduced without having to solve a MILP.

3.5 Choosing big-M values

One issue that we have not addressed so far is the importance of choosing right values for M_0 and M_1. From the theoretical side, we are assured about the existence of M_0 and M_1 values for which the algorithm behaves as intended. However, when implementing these algorithms, it becomes crucial to choose correct values. On one hand, we are inclined to select large constants to guarantee the correctness of the solutions. If we select constants that are not large enough, we risk cutting off feasible solutions corresponding to EMs we may interested in. On the other hand, the larger these values are, the less numerically stable the MILPs become. In the following, we analyze how setting M_0 and M_1 affects the output of the algorithms.

We note that every feasible solution (v, a) of (MILP3) satisfies $v_i \in [1, M_0]$ for $a_i = 1$. Thus, M_0 represents the largest ratio v_i/v_j that can occur for non-zero fluxes v_i and v_j.

M_1 appears only in the directional coupling constraints. For any flux vector that contains r_1 but not r_2, the directional coupling constraints will be infeasible, regardless of the choice of M_1. Thus the only bad case occurs when an EM containing both r_1 and r_2 that should be feasible is rendered infeasible by the directional coupling constraints. This happens if M_1 is chosen not large enough, such that some of the constraints that should be trivially satisfied become unsatisfiable. Let v' be an EM for which this case occurs, i.e., $v'_{r_1} \geq 0$, $v'_{r_2} \geq 0$, with infeasible $DirC(r_1, r_2)$ constraints. Let $a \in \{0,1\}^n$ be a vector corresponding to the support of v'. We investigate the following LP-relaxation of (MILP3).

$$\begin{array}{lrcl} \min & 0 & & \\ & S^T y + u^{r_1} x & \geq & M_1(a - \mathbf{1} - u^{r_2}), \\ & -x & \geq & 1. \end{array}$$

Its corresponding dual-LP reads

$$\begin{array}{lrcl} \max & w - M_1 \sum_{i \in \{i \mid a_i = 0\} \cup \{r_2\}} v_i & & \\ & Sv & = & 0, \\ & v_{r_1} - w & = & 0, \\ & v, w & \geq & 0. \end{array}$$

Note that the null vector 0 is always dual feasible. Since in our assumption the primal is infeasible, the dual must be unbounded. We conclude that the primal is infeasible because $w = v_{r_1} > M_1 \sum_{i \in \{i \mid a_i = 0\} \cup \{r_2\}} v_i$. In order to avoid this problem, we recommend over-approximating M_1 with $M_1 \geq v_{r_1}/v_i, i \in \mathcal{R} \setminus \{r_1\}$.

4 Results and Discussion

The above described methods have been implemented in Matlab. The MILP solver of choice was Gurobi 5.0. In the following, we present several use-case scenarios that have been performed on different real-world networks. These tests were aimed at validating the correctness of the methods, and also to motivate their existence, by applying them on networks where an exhaustive EM enumeration would fail. Furthermore, we ran benchmarking tests to measure the running time of the algorithms and other statistical properties. All computations were performed using a single Intel T2600 (2.16 GHz) processor on a 32-bit Windows 7 system, with a maximum memory of 640MB allocated to Matlab.

Checking the correctness of the resulting flux vectors is easy. A simple rank test as in (Kamp and Schuster, 2006) can prove the EM property, while checking for non-zero entries in the target reactions assures that we are indeed using them. Due to the nature of the algorithms, the EM property never gets violated. Indeed, because the last step of every method presented involves solving an LP to optimality using a Simplex based method, the result will be a vertex of the truncated flux cone.

4.1 Use-case scenario 1

de Figueiredo *et al.* (2008) compare graph-based pathway enumeration with EM analysis in the context of discovering pathways producing G6P (KEGG entry C00668) from AcCoA (KEGG entry C00024). The underlying network was based on the human Krebs-cycle with two possible configurations, one of which was able to display the required phenotype, while the other was not. Answering this type of questions is a perfect use-case scenario for our method. Indeed, we ran our algorithms, trying to find EMs containing the required reactions, and in both configurations we were able to replicate the answers presented by the authors.

It is important to mention that the network used in the previous study is quite small (20 metabolites and 26 reactions). EM enumeration tools work very well for models of this size. Thus, we performed the same analysis for a genome-scale reconstruction of the human metabolic model (Duarte *et al.*, 2007). A crucial difference, beside the size of the models, is that in the latter network, the two metabolites we were interested in are internal to the model. This is important for two reasons. First, the number of adjacent reactions to each metabolite can be more than one, making the selection of a pair of reactions involving G6P and AcCoA non-unique. Indeed, after eliminating the blocked reactions from the network, we still had 63 viable pairs of reactions.

Secondly, and more importantly, for two internal reactions r_1 and r_2, one cannot easily find out whether reaction r_1 is a predecessor of r_2 or vice-versa. In the case of boundary reactions, it is very natural to claim that in an EM, uptake reactions precede outgoing reactions. We could say that the products of the outgoing reactions are obtained from the substrates of the uptake reactions. A similar statement for internal reactions is not trivial. Deciding whether one internal reaction precedes another one is an unsolved problem. While our method can still identify EMs with internal reactions, the interpretation of these EMs has to be done with care.

For these reasons, two artificial transport reactions were added to the network, one importing AcCoA and the other secreting G6P. The biological interpretation of our method applied to this modified system would be the following. "Assuming we can inject AcCoA into the network and have a method to secret G6P. Can the underlying network convert AcCoA into G6P?". Allowing one hour of running time, we could identify 6 EMs with the desired property. The number of reactions taking part in these EMs ranged from 58 to 286.

In a similar way, a possible usage of the tool would be to enumerate EMs that produce biomass, allowing the system to grow, but at the same time also produce one or more by-

products (e.g. biofuels or toxic compounds). Such a set of EMs gives us insight into the growth-coupled production capabilities of a microorganism. Alternatively, these EMs could also be used as an input for other methods, such as constrained cut set computation (Hädicke and Klamt, 2011).

4.2 Use-case scenario 2

While the main novelty of the paper is Algorithm 2, Algorithm 1 provides an alternative method for computing EMs through a single given reaction. Compared to methods like (de Figueiredo *et al.*, 2009), our algorithm improves running time at the cost of computing not necessarily the shortest EMs. While shorter EMs are more likely to represent biological phenotypes that occur in practice, it is unclear if only the shortest ones should be of interest. Rezola *et al.* (2011) introduced the concept of *generating flux modes* (GFMs) and compared the 100 shortest EMs with the 100 shortest GFMs producing lysine in the *E. coli* iAF1260 model (Feist *et al.*, 2007). They argued that even though GMFs are longer and take more time to compute, they present a more varied description of the total solution set. We extended their experiment by also considering Algorithm 1. We enumerated 100 EMs producing lysine and compared their statistical properties with the other two methods. The results are summarized in Table 3.

Method	NoR	LI	AHD
shortest EM*	54	25-26	12.79
shortest GFMs*	132	25-37	16.21
Algorithm 1	272	25-57	26.082

Table 3: Comparison between shortest EMs, shortest GFMs and Algorithm 1. (NoR) - total number of reactions involved in any of the computed modes. (LI) - length interval of the modes. (AHD) - Average Hamming distance. '*' indicates numbers taken from (Rezola *et al.*, 2011).

It becomes clear from Table 3 that Algorithm 1 computes a broader range of pathways that can be used for lysine production. While it still finds some representatives of the shortest EMs, it also detects more advanced EMs. Fig. 2 presents the length distribution of the EMs obtained by Algorithm 1.

Even if we extend our interest beyond the set of shortest EMs, one might still want to avoid unreasonably long ones (i.e., those involving more than L reactions, for some $L \geq 1$). This can easily be achieved by adding the constraint $\sum_{l \in \mathcal{R}} a_l \leq L$ to the MILP formulation.

In our next experiment, we compared Algorithm 1 with the method of de Figueiredo *et al.* (2009) for computing shortest EMs. As input we used the metabolic network of *S. Cerevisiae iND750* (Duarte *et al.*, 2004) with ethanol (R_ETOH) being the target reaction, required to be active in all EMs. For various M values and choosing the flux variables to be either integer or continuous, 20 EMs were computed in each case. Table 4 summarizes the results.

Table 4 highlights the main differences in the two methods. Algorithm 1 finds EMs much quicker at the expense of individual EMs' length. As expected, in almost every case, the higher the constant M was chosen, the more expensive the computation becomes. de Figueiredo *et al.* (2009) remark that using continuous flux variables, the MILPs might be more time consuming to solve. However, when computing the shortest EMs in this particular network, we do not consistently observe this behavior. Indeed, while for $M = 10$, choosing the flux variables as continuous proved to be slightly more expensive, in all other cases continuous variables led to shorter running times.

In every iteration of Algorithm 1, the MILPs to be solved become more complex, since they

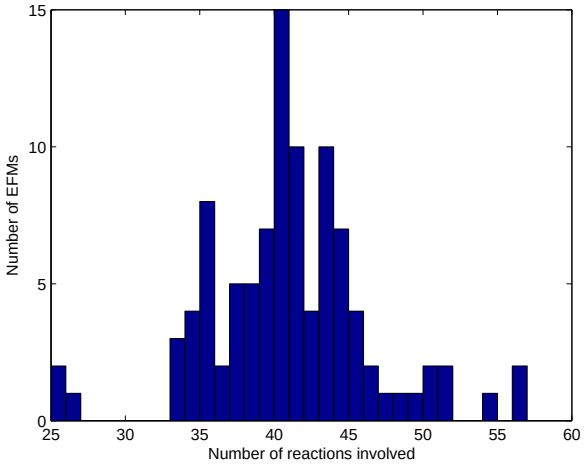

Figure 2: Histogram of the length of EMs found by Algorithm 1

Table 4: Comparing Algorithm 1 with the computation of shortest EMs (Running time in secs).

		Integer variables		Continuous variables	
	Method	Length	Time	Length	Time
M = 10	Shortest	6-10	1719s	6-10	2074s
	Algo. 1	6-15	16s	6-23	15s
M = 100	Shortest	6-10	8158s	6-10	3421s
	Algo. 1	6-21	21s	6-31	18s
M = 1000	Shortest	6-10	14362s	6-10	7780s
	Algo. 1	6-31	16s	6-50	29s

contain additional constraints. Over time, these constraints are expected to slow down the solving of (MILP2). Hence, the rate at which new EMs are computed is expected to decrease, the more EMs we enumerate. To study this effect, we enumerated 1000 EMs in the network of *S. Cerevisiae iND750* (Duarte *et al.*, 2004), with ethanol (R_ETOH) being the target reaction. The flux variables were chosen as continuous, while a value for M of 1000 was used. The total running time was 4350 seconds. For reference, in the same time frame we could compute only 16 EMs with the method presented by de Figueiredo *et al.* (2009). We measured the time (t_i) required for the computation of the i-th EM, for all $i \in \{1, ..., 1000\}$. These times displayed a wide range of variation. Thus, in order to display them we used the mean times $\bar{t}_i := 1/i \sum_{j=1}^{i} t_j$. Fig. 3 illustrates these results. In order to study how the lengths of consecutive EMs vary, we also measured in the previous experiment the length l_i of the i-th EM, for all $i \in \{1, ..., 1000\}$. Fig. 4 shows the evolution of the mean lengths $\bar{l}_i := 1/i \sum_{j=1}^{i} l_j$.

4.3 Statistical analysis and flux uncoupling

The main bottleneck of the method are cases when EMs of the required type do not exist. While finding an EM if it exists seemed to work well in practice, proving their non-existence (i.e., showing (MILP3) to be infeasible) is rather hard and time-consuming. Intuitively, one can think of the MILPs search-tree. Depending on how many feasible solutions there are, the solver might find one without traversing the whole tree. For proving the non-existence, the whole

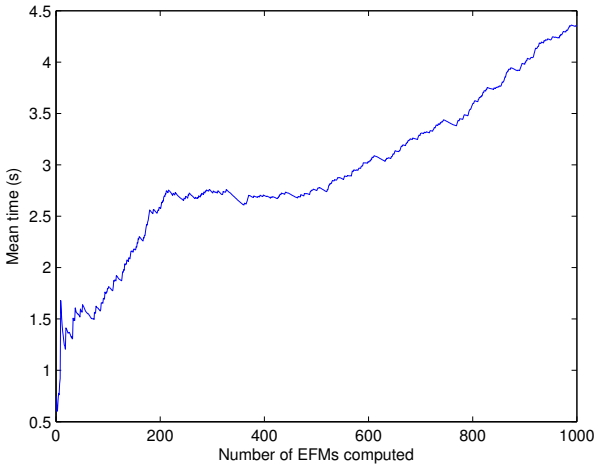

Figure 3: Mean time of the EMs computed.

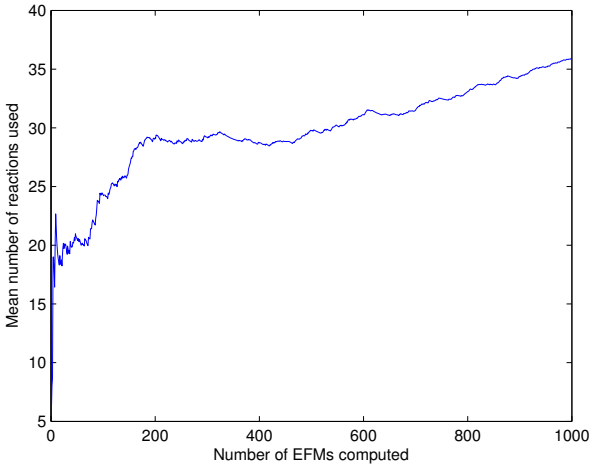

Figure 4: Mean length of the EMs computed.

search-tree must be traversed.

In the next experiment, we tried to compute an EM for every pair of reactions. A maximum time of 60 seconds was allocated for each pair. Reaching the timeout meant that we were unable to compute if this pair shared an EM or not. The test has been performed on two small to medium-sized real world networks, the central metabolism of *E. coli* (ECC) (Palsson, 2006) and the *H. pylori* (HP) (Thiele *et al.*, 2005) genome-scale metabolic network. For the constants M_0 and M_1, the value of 10000 was chosen, while the feasibility tolerance was set to 10^{-6}. The total running time in the case of ECC was approximately 3 hours, while the algorithm took close to 3 days in the case of HP.

Netw.	Reac.	Pairs	EM found	Mean length	No EM
ECC	90	8010	7691	24.36	176
HP	269	72092	66749	46.57	1862

Table 5: Summary of the correct computational results

Table 5 empirically sheds light on the nature of real-world networks. One can observe in these networks that most pairs of reactions share at least one EM, and the algorithm presented in this paper is able to find them.

Network	Timed out	Numerical error
ECC	78	65
HP	2078	1403

Table 6: Summary of the bad instances

From Table 6, it becomes clear that (although not many) there are cases where the algorithm does not produce a relevant result, either by not finishing before its time-out, or by producing an erroneous result. The latter cases are attributed to the incorrect choice of M_0 and M_1 and numerical imprecisions in the MILP solver.

Seeing the high probability for the existence of an EM containing two randomly chosen reactions, the question arises whether solving (MILP2) would suffice. If it correctly finds an EM in most cases, this would motivate using (MILP2) as a heuristic approach to the problem. Unless the two reactions are blocked, (MILP2) is always feasible, and the optimal solution is characterized by Cor. 3 which asserts that we either get an EM or a false positive. We performed the same experiment as before, and attempted to compute a flux vector for every pair of reactions. We decided about the EM property of the computed flux vectors by applying the rank test (Gagneur and Klamt, 2004, Lemma 2). For every instance, a maximum of 60 seconds was allowed. Table 7 summarizes these results.

Network	EM found	False positive	Timed out
ECC	5212	2686	202
HP	206	9213	62673

Table 7: The performance of (MILP2) on metabolic networks

It turns out that in the case of the small *E. Coli* network, (MILP2) performs reasonably well. In 67% of the cases where EMs exist, it correctly finds one. However, for the medium-scale *H. Pylori* network, solving (MILP2) to optimality almost never terminates in the allocated time.

Based on this empirical evidence we conclude that (MILP2) may not be a viable approach for medium- to large-scale metabolic networks.

5 Conclusion

We have presented novel methods to compute EMs involving any number of predefined target reactions. These algorithms can also be used to distinguish between mutually exclusive and sometimes coupled reactions. From the application on genome-scale metabolic networks, we conclude that the methods work as intended and are fast enough for practical use. They should become a valuable asset for constraint-based analysis of metabolic networks.

A prototype implementation in Matlab is available for download at `https://sourceforge.net/projects/caefm`.

Acknowledgement

We would like to thank Sayed-Amir Marashi for constructive discussions on the experiments. The proof of Prop. 5 is based on ideas from Yaron Goldstein. We would also like to thank Luis de Figueiredo for providing us a network model from (de Figueiredo *et al.*, 2008).

References

Acuña, V., Chierichetti, F., Lacroix, V., Marchetti-Spaccamela, A., Sagot, M., and Stougie, L. (2009). Modes and cuts in metabolic networks: Complexity and algorithms. *BioSystems*, **95**, 51–60.

Acuña, V., Marchetti-Spaccamela, A., Sagot, M., and Stougie, L. (2010). A note on the complexity of finding and enumerating elementary modes. *BioSystems*, **99**, 210–214.

Burgard, A. P., Nikolaev, E. V., Schilling, C. H., and Maranas, C. D. (2004). Flux coupling analysis of genome-scale metabolic network reconstructions. *Genome Research*, **14**, 301–312.

David, L., Marashi, S.-A., Larhlimi, A., Mieth, B., and Bockmayr, A. (2011). FFCA: a feasibility-based method for flux coupling analysis of metabolic networks. *BMC bioinformatics*, **12**(1), 236.

de Figueiredo, L. F., Schuster, S., Kaleta, C., and Fell, D. A. (2008). Can sugars be produced from fatty acids? a test case for pathway analysis tools. *Bioinformatics*, **24**(22), 2615–2621.

de Figueiredo, L. F., Podhorski, A., Rubio, A., Kaleta, C., Beasley, J. E., Schuster, S., and Planes, F. J. (2009). Computing the shortest elementary flux modes in genome-scale metabolic networks. *Bioinformatics*, **25**(23), 3158–3165.

Duarte, N. C., Herrgård, M. J., and Palsson, B. Ø. (2004). Reconstruction and validation of saccharomyces cerevisiae iND750, a fully compartmentalized genome-scale metabolic model. *Genome research*, **14**(7), 1298–1309.

Duarte, N. C., Becker, S. A., Jamshidi, N., Thiele, I., Mo, M. L., Vo, T. D., Srivas, R., and Palsson, B. . (2007). Global reconstruction of the human metabolic network based on genomic and bibliomic data. *Proceedings of the National Academy of Sciences*, **104**(6), 1777–1782.

Feist, A. M., Henry, C. S., Reed, J. L., Krummenacker, M., Joyce, A. R., Karp, P. D., Broadbelt, L. J., Hatzimanikatis, V., and Palsson, B. Ø. (2007). A genome-scale metabolic reconstruction for Escherichia coli K-12 MG1655 that accounts for 1260 ORFs and thermodynamic information. *Mol. Syst. Biol.*, **3**(121).

Gagneur, J. and Klamt, S. (2004). Computation of elementary modes: a unifying framework and the new binary approach. *BMC Bioinformatics*, **5**, 175.

Hädicke, O. and Klamt, S. (2011). Computing complex metabolic intervention strategies using constrained minimal cut sets. *Metabolic Engineering*, **13**(2), 204 – 213.

Kaleta, C., de Figueiredo, L. F., and Schuster, S. (2009). Can the whole be less than the sum of its parts? Pathway analysis in genome-scale metabolic networks using elementary flux patterns. *Genome Research*, **19**, 1872–1883.

Kamp, A. v. and Schuster, S. (2006). Metatool 5.0: fast and flexible elementary modes analysis. *Bioinformatics*, **22**(15), 1930–1931.

Larhlimi, A., David, L., Selbig, J., and Bockmayr, A. (2012). F2C2: a fast tool for the computation of flux coupling in genome-scale metabolic networks. *BMC bioinformatics*, **13**(1), 57.

Marashi, S.-A. and Bockmayr, A. (2011). Flux coupling analysis of metabolic networks is sensitive to missing reactions. *Biosystems*, **103**(1), 57 – 66.

Marashi, S.-A., David, L., and Bockmayr, A. (2012). Analysis of metabolic subnetworks by flux cone projection. *Algorithms for molecular biology : AMB*, **7**(1), 17.

Palsson, B. O. (2006). *Systems Biology: Properties of Reconstructed Networks*. Cambridge University Press, New York.

Papin, A. J., Stelling, J., Price, N. D., Klamt, S., Schuster, S., and Palsson, B. O. (2004). Comparison of network-based pathway analysis methods. *TRENDS in Biotechnology*, **22**(8), 400–405.

Rezola, A., de Figueiredo, L. F., Brock, M., Pey, J., Podhorski, A., Wittmann, C., Schuster, S., Bockmayr, A., and Planes, F. J. (2011). Exploring metabolic pathways in genome-scale networks via generating flux modes. *Bioinformatics*, **27**(4), 534–540.

Rezola, A., Pey, J., de Figueiredo, L. F., Podhorski, A., Schuster, S., Rubio, A., and Planes, F. J. (2013). Selection of human tissue-specific elementary flux modes using gene expression data. *Bioinformatics*, **29**(16), 2009–16.

Schrijver, A. (1998). *Theory of Linear and Integer Programming*. Wiley.

Schuster, S. and Hilgetag, C. (1994). On elementary flux modes in biochemical reaction systems at steady state. *Journal of Biological Systems*, **2**, 165–182.

Schuster, S., Dandekar, T., and Fell, A. (1999). Detection of elementary flux modes in biochemical networks : a promising tool for pathway analysis and metabolic engineering. *Trends Biotechnol.*, **17**, 53–60.

Schuster, S., Fell, D. A., and Dandekar, T. (2000). A general definition of metabolic pathways useful for systematic organization and analysis of complex metabolic networks. *Nature Biotechnology*, **18**, 326–332.

Schuster, S., Hilgetag, C., Woods, J. H., and Fell, P. A. (2002). Reaction routes in biochemical reaction systems: Algebraic properties , validated calculation procedure and example from nucleotide metabolism. *J Math Biol*, **45**(2), 153–181.

Terzer, M. (2009). *Large Scale Methods to Enumerate Extreme Rays and Elementary Modes*. Ph.D. thesis, Swiss Federal Institute of Technology, Zurich.

Terzer, M. and Stelling, J. (2008). Large-scale computation of elementary flux modes with bit pattern trees. *Bioinformatics*, **24**, 2229–2235.

Thiele, I., Vo, T. D., Price, N. D., and Palsson, B. O. (2005). Expanded metabolic reconstruction of *Helicobacter pylori* (iIT341 GSM/GPR): an in silico genome-scale characterization of single- and double-deletion mutants. *Journal of Bacteriology*, **187**, 5818–5830.

Urbanczik, R. and Wagner, C. (2005). Functional stoichiometric analysis of metabolic networks. *Bioinformatics*, **21**, 4176–4180.

von Kamp, A. and Schuster, S. (2006). Metatool 5.0: fast and flexible elementary modes analysis. *Bioinformatics*, **22**, 1930–1931.

Zanghellini, J., Ruckerbauer, D. E., Hanscho, M., and Jungreuthmayer, C. (2013). Elementary flux modes in a nutshell: Properties, calculation and applications. *Biotechnology Journal*, page n/a.