

MCS²: minimal coordinated supports for fast enumeration of minimal cut sets in metabolic networks

Reza Miraskarshahi¹, Hooman Zabeti¹, Tamon Stephen² and Leonid Chindelevitch^{1,*}

¹School of Computing Science and ²Department of Mathematics, Simon Fraser University, Burnaby, BC V5A 1S6, Canada

*To whom correspondence should be addressed.

Abstract

Motivation: Constraint-based modeling of metabolic networks helps researchers gain insight into the metabolic processes of many organisms, both prokaryotic and eukaryotic. Minimal cut sets (MCSs) are minimal sets of reactions whose inhibition blocks a target reaction in a metabolic network. Most approaches for finding the MCSs in constrained-based models require, either as an intermediate step or as a byproduct of the calculation, the computation of the set of elementary flux modes (EFMs), a convex basis for the valid flux vectors in the network. Recently, Ballerstein *et al.* proposed a method for computing the MCSs of a network without first computing its EFMs, by creating a dual network whose EFMs are a superset of the MCSs of the original network. However, their dual network is always larger than the original network and depends on the target reaction. Here we propose the construction of a different dual network, which is typically smaller than the original network and is independent of the target reaction, for the same purpose. We prove the correctness of our approach, minimal coordinated support (MCS²), and describe how it can be modified to compute the few smallest MCSs for a given target reaction.

Results: We compare MCS² to the method of Ballerstein *et al.* and two other existing methods. We show that MCS² succeeds in calculating the full set of MCSs in many models where other approaches cannot finish within a reasonable amount of time. Thus, in addition to its theoretical novelty, our approach provides a practical advantage over existing methods.

Availability and implementation: MCS² is freely available at <https://github.com/RezaMash/MCS> under the GNU 3.0 license.

Contact: leonid@sfu.ca

Supplementary information: [Supplementary data](#) are available at *Bioinformatics* online.

1 Introduction

Constraint-based modeling of metabolic networks has been a major subfield of systems biology thanks to its ability to identify key qualitative characteristics of networks for analyzing and extracting useful information (Bordbar *et al.*, 2014; Lewis *et al.*, 2012; Price *et al.*, 2004). A metabolic network is a collection of chemical reactions which comprise the metabolic activities (i.e. the biochemical transformation of molecules into other molecules for the purpose of maintenance and growth) of a specific organism. One important application of metabolic network analysis is to find interventions that can block a reaction of interest, typically referred to as the *target reaction*, with applications in drug target identification (Harder *et al.*,

2016; Hartman *et al.*, 2014; Imielinski and Belta, 2008; Trinh *et al.*, 2006; von Kamp and Klamt, 2017) and metabolic engineering (Mahadevan *et al.*, 2015). When this is achieved by disabling one or more other reactions, the disabled reactions are called a *cut set*. A cut set is called ‘minimal’ if no proper subset of it can disable the target reaction. The concept of minimal cut sets (MCS) was introduced by Klamt and Gilles (2004) and its applications are examined in detail in Klamt (2006).

At the moment, the main approach used for enumerating the MCSs for a target reaction is to compute the elementary flux modes (EFMs) containing the target and then use a dualization procedure to produce the MCSs (Gainer-Dewar and Vera-Licona, 2017).

Here, *flux modes* are possible distributions of fluxes through the reactions, and those can be modeled as hyperedges on the vertex set of possible reactions. *EFMs* are flux modes which are support-minimal, and it is known that any flux mode can be written as a non-negative linear combination of EFMs. Given the full set of EFMs, MCS can be obtained through the dualization of the hypergraph they define (Haus et al., 2008; Klamt and Gilles, 2004). Two approaches to do this are Berge’s algorithm (Berge, 1984) and Fredman and Khachiyan’s dualization procedure (Fredman and Khachiyan, 1996). However, both suffer from poor worst-case complexity and produce mixed results in practice. A comparatively new approach (Ballerstein et al., 2012) produces the MCSs without first computing the EFMs. It works by generating a dual network, which is larger than the original network and depends on the target reaction, and then computing a subset of the EFMs of that network with a specific property, which guarantees that they are precisely the MCSs in the original network. We call this the *target-specific dual network method*. In this article, we develop a new method, minimal coordinated support (MCS²), which also generates a dual network (either explicitly or implicitly), but in a way that is independent of the target reaction from the original network, then computes the MCSs from those EFMs of the dual network that satisfy a certain property, which also guarantees that they are precisely the MCSs for the target reaction in the original network. MCS² is based on a generalization of some theoretical results by the last author (Chindelevitch, 2014).

We implement MCS² and find it to be effective on most instances we test it on. We compare it to three alternate methods for enumerating all the MCSs for a target set. The first two methods are to compute the EFMs, and then dualize them with either Berge’s algorithm, or an optimized implementation of Fredman–Khachiyan dualization, respectively. For Berge’s algorithm, we used the implementation in CellNetAnalyzer (Klamt et al., 2007), containing the enhancements described in Eiter et al. (2008) and Haus et al. (2008). For Fredman–Khachiyan dualization, we used the recent implementation of Sedaghat et al. (2018).

The *target-specific dual network method* (Ballerstein et al., 2012) first creates a dual network based on the given stoichiometric matrix and the given target reaction. It then proceeds to compute the EFMs of that dual network. Following some post-processing, the supports of these EFMs are reduced to give the required MCSs; because the MCSs correspond to only a subset of the vectors produced, this post-processing includes removing any supersets. Like MCS², the target-specific dual network method reports all the MCSs without first computing the EFMs or requiring them as an input. The authors of Ballerstein et al. (2012) did not provide a publicly available implementation of their method, so we did so ourselves, including all the enhancements mentioned in their [Supplementary Materials](#). Most of these enhancements have improved the performance of the target-specific dual network method, for instance by reducing the size of the intermediate results.

For the majority of the models we investigated, we find that MCS² is more efficient than these other methods, in terms of both running time and memory use. On the negative side, we show that our approach does not allow the enumeration of all MCSs through a given target reaction in incremental polynomial time, something that therefore remains a major open problem in the field.

Given the challenges in enumerating all MCSs (in part due to their large number, which can be exponential in the size of the network), some recent work (von Kamp and Klamt, 2014; Vieira et al., 2016) uses mixed-integer linear programming (MILP) formulations to enumerate a subset of the MCSs, in increasing order of size. In

some practical applications, quickly obtaining a few MCSs of minimum size may be more desirable than enumerating the complete set. We therefore adapt the MCS² approach to use MILP formulations to address this task. We also implement this method, which we call MCS²-MILP, and compare it to MCSEnumerator, the target-specific dual network approach adapted to MILP (von Kamp and Klamt, 2014). The comparison shows that MCS²-MILP performs at least as well as MCSEnumerator using a state-of-the-art MILP solver (IBM).

We conclude that MCS² is a promising approach for the computation of MCSs in metabolic networks, and expect it to be a beneficial addition to the analysis tools available for metabolic network models.

We now introduce the terminology we will be using throughout this article. When we speak of a metabolic network, it is understood that we are talking about a model in the constraint-based modeling formalism.

Definition 1 (Stoichiometric matrix). The *stoichiometric matrix* S is an $m \times n$ matrix with each row representing a metabolite (indexed from 1 to m) and each column, a reaction (indexed from 1 to n). The entry S_{ij} indicates how many units of metabolite i are produced (if $S_{ij} > 0$) or consumed (if $S_{ij} < 0$) by reaction j . A vector v is *feasible* with respect to S if it is in the *null space* of S , i.e. if it satisfies $Sv = 0$.

Definition 2 (Reaction irreversibility). The set \mathcal{I} of *irreversible* reactions is a subset of the set of reactions constrained to have only non-negative fluxes. Its complement \mathcal{I}^c , the set of *reversible* reactions, is allowed to have fluxes of any sign. A vector v respects the reaction irreversibility constraints if $v_i \geq 0 \forall i \in \mathcal{I}$, also written as $v_{\mathcal{I}} \geq 0$.

Definition 3 (Metabolic network). A metabolic network \mathcal{M} is a pair (S, \mathcal{I}) , where $S \in \mathbb{Q}^{m \times n}$ is a stoichiometric matrix and $\mathcal{I} \subseteq [n]$ is the set of irreversible reactions. A vector v is a *flux mode* if it is feasible with respect to S and respects the irreversibility constraints, i.e. $Sv = 0$ and $v_{\mathcal{I}} \geq 0$. The set of all such vectors is called the network’s *flux cone*.

Definition 4 (Reconfigured network). Let (S, \mathcal{I}) be a metabolic network. We can *reconfigure* this network by replacing S with $S' = [S | -S_{\mathcal{I}^c}]$ which $S_{\mathcal{I}^c}$ contains the columns of S that are not in \mathcal{I} , and then consider all reactions as irreversible. This is equivalent to splitting each reversible reaction in the network into its forward reaction and reverse reaction.

Definition 5 (Null space matrix and network). Let S be a matrix. A *null space matrix* of S is a matrix whose rows form a basis of the null space of S . The *null space network* of a metabolic network with stoichiometric matrix S is the fully reversible metabolic network (i.e. with $\mathcal{I} = \phi$) whose stoichiometric matrix is a null space matrix of S .

Definition 6 (Positive and negative support). Let v be a vector. The *positive support* of v , $\mathcal{R}_+(v)$, is the set of positions i where v_i is positive: $\mathcal{R}_+(v) := \{i | v_i > 0\}$. The *negative support* of v , $\mathcal{R}_-(v)$, is the set of positions i where v_i is negative: $\mathcal{R}_-(v) := \{i | v_i < 0\}$. Their union $\mathcal{R}(v)$ is the *support* of v : $\mathcal{R}(v) := \mathcal{R}_+(v) \cup \mathcal{R}_-(v)$.

Definition 7 (Coordinated support). Let v be a vector of size n and let $A \subseteq \{1, 2, \dots, n\}$ be a set of positions. The *A-coordinated support* of v , $\mathcal{R}_A(v)$, is the union of its negative support on the positions in A and its support everywhere else: $\mathcal{R}_A(v) := (\mathcal{R}_-(v) \cap A) \cup (\mathcal{R}(v) \cap A^c)$, where A^c is the complement of the set A with respect to $\{1, 2, \dots, n\}$.

Definition 8 (EFM). Let $\mathcal{M} = (S, \mathcal{I})$ be a metabolic network, and let v be a flux mode of \mathcal{M} . It is an *EFM* if its support is minimal among all the

flux modes of \mathcal{M} , i.e. for any flux mode w which $\mathcal{R}(w) \subsetneq \mathcal{R}(v) \Rightarrow w = 0$ (Gagneur and Klamt, 2004; Schuster *et al.*, 2000).

Definition 9 (MCS). Let $\mathcal{M} = (S, \mathcal{I})$ be a metabolic network, and let t be a reaction. C is a *cut set* for t if $Sv = 0, v_{\mathcal{I}} \geq 0, v_C = 0 \Rightarrow v_t = 0$. C is a MCS if it is inclusion-minimal: for any set $D, D \subsetneq C$ implies that $\exists v$ s.t. $Sv = 0, v_{\mathcal{I}} \geq 0, v_D = 0, v_t \neq 0$ (Klamt and Gilles, 2004).

Definition 10 (Canonical form of a network). Let $\mathcal{M} = (S, \mathcal{I})$ be a metabolic network. We say that \mathcal{M} is in *canonical form* if it satisfies:

1. No blocked reactions: for every reaction i , there exists a flux vector v with $v_i = 1$
2. Proper directedness: for every reaction $i \in \mathcal{I}^C$, there exists a flux vector w with $w_i = -1$
3. No enzyme subsets: no reaction pair $i \neq j$ satisfies $v_i = \kappa v_j$ with $\kappa \in \mathbb{Q}$ for all flux vectors v
4. No redundant constraints: S has full row rank

A metabolic network can be reduced to an equivalent one in canonical form (a.k.a. compressed form) in time polynomial in m and n (Chindelevitch, 2014).

2 The MCS² method

Let S_i be the i th row of the stoichiometric matrix S . Then S_{ir} represents the amount of metabolite i consumed or produced by reaction r (in these cases, $S_{ir} < 0$ and $S_{ir} > 0$, respectively). Assume that reaction r produces metabolite i if it has a positive flux. Then, in a steady state where no reaction consuming metabolite i is active, reaction r must be inactive in the forward direction. If reaction r happens to be reversible, it must be consuming metabolite i , and its flux must be negative. This shows that reaction r is blocked in the forward direction if we disable every reaction that can consume metabolite i , i.e. every irreversible reaction with a negative value in row i and every reversible reaction with a non-zero value in row i . Disabling a reaction can be done by removing the corresponding column from the stoichiometric matrix or adding the constraint that this reaction's flux is zero to the problem. The set of such reactions is then a cut set for the forward direction of reaction r . Every row gives us some, not necessarily minimal, cut set in this manner.

We can apply the same reasoning to linear combinations of the metabolites. Consider a new *virtual* metabolite x , which represents a linear combination of rows S_i and S_j corresponding to metabolites i and j respectively, say for example $u_x = 2S_i - S_j$.

Since the fluxes producing and consuming each metabolite are balanced in any admissible vector, so are the fluxes of their linear combinations, so the virtual metabolite x must also be balanced. If we pick a reaction with a positive value in u_x , it produces a virtual metabolite x when it has a positive flux. It will therefore be blocked if we cut all irreversible reactions with negative values in u_x and all reversible reactions with non-zero values in u_x . Thus, we can obtain cut sets from the vector u_x , which is a member of the row space of S , as we did with S_i and S_j . A proposal for finding cut sets by analyzing the row space of the stoichiometric matrix was introduced in the Ph.D. thesis of Chindelevitch (Chindelevitch, 2014). The intuition we described shows how vectors in the row space can generate cut sets. However, the lemmas proven in Chindelevitch (2014) work only for the fully reversible networks (i.e. with $\mathcal{I} = \phi$) or fully reversible networks (i.e. with $\mathcal{I} = [n]$). We generalize them here to networks with both irreversible reactions and reversible reactions.

2.1 Enumerating the full set of MCSs

In the MCS² method, the dual network is the null space network of the original network, i.e. a fully reversible network whose stoichiometric matrix is the null space matrix of the original stoichiometric matrix. The EFMs in this dual network map to MCSs of the original network, though, as we will see, the mapping can be many to one. The dual network has the same number of reactions, n , as the original one, but it typically has fewer metabolites; if the original network has m metabolites and the stoichiometric matrix is full rank, the dual has $n - m$ metabolites.

Lemma 1 (MCSs for an irreversible reaction). Let $\mathcal{M} = (S, \mathcal{I})$ be a metabolic network. Let $t \in \mathcal{I}$ be an irreversible target reaction. Then C is a cut set for t if and only if there exists a vector u in the row space of S , $u \in \text{Row}(S)$, such that $u_t = 1$ and $\mathcal{R}_{\mathcal{I}}(u) \subseteq C$.

Proof. This lemma is an extension of Lemma 3 of Chindelevitch (2014). We observe that C being a cut set for irreversible reaction t is equivalent to:

$$S_{-C}v = 0 \text{ and } \forall i \in \mathcal{I} - C \quad v_i \geq 0 \Rightarrow v_t = 0 \quad (1)$$

Based on Farkas' Lemma and the irreversibility of t , we only need to find a constraint that implies $v_t \leq 0$. Thus, there exists a y such that:

$$y^T S_{-C} = e_t + \sum_{i \in \mathcal{I} - \{C, t\}} \alpha_i e_i, \text{ s.t. } \forall i \in \mathcal{I} - \{C \cup t\} \alpha_i \geq 0 \quad (2)$$

Here, e_i is a vector with a 1 in the i th position and 0 elsewhere. Thus:

$$y^T S = u, u_t = 1, \forall i \in \mathcal{I} - C \quad u_i \geq 0, \forall j \in \mathcal{I}^C - C \quad u_j = 0 \quad (3)$$

Therefore $\mathcal{R}_{-}(u) \cap \mathcal{I} \subseteq C \cap \mathcal{I}$ and $\mathcal{R}(u) \cap \mathcal{I}^C \subseteq C \cap \mathcal{I}^C$, and so $\mathcal{R}_{\mathcal{I}}(u) \subseteq C$.

For the other direction, suppose that u satisfies Equation (3) for some C . Then the union of the indices i in \mathcal{I} for which $u_i < 0$ and the indices i in \mathcal{I}^C for which $u_i \neq 0$ is a subset of C , i.e. $\mathcal{R}_{\mathcal{I}}(u) \subseteq C$. Then equality (2) holds and, by Farkas' lemma, so does condition (1). \square

Lemma 2 (MCSs for one direction of a reversible reaction). Let $\mathcal{M} = (S, \mathcal{I})$ be a metabolic network. Let t be a reversible target reaction. Then C is a cut set for the forward (reverse) direction of t if and only if there exists a vector $u \in \text{Row}(S)$ such that $\mathcal{R}_{\mathcal{I}}(u) \subseteq C$ and $u_t = 1$ ($u_t = -1$), respectively.

Proof. If we assume that t is irreversible for a moment, the first part already follows from the previous lemma. For the second part, replace t with $-t$ in S to create S' . Reaction t is blocked in the forward direction in S' if and only if reaction t is blocked in the reverse direction in S , and there is a bijection between the vectors in $\text{Row}(S)$ and those in $\text{Row}(S')$ via the mapping that multiplies the t th coordinate by -1 . \square

With these lemmas, Algorithm 1 can be used to find the MCS for a set of target reactions $T = \{t_1, t_2, \dots, t_k\}$ in an arbitrary metabolic network $\mathcal{M} = (S, \mathcal{I})$, where T has separate elements for the opposite directions of a reversible reaction. We call this method MCS² because it computes MCSs as the MCS² of the EFMs in the dual network.

The MCS² method computes the null space network of the original network. Then, by applying coordinated support to the EFMs in which each target reaction is active in turn, it generates the cut sets related for this target. All the MCSs are among these cut sets, and are obtained by pruning. The null space network is fully reversible, and since the null space of the null space is the original space, the dual network of the dual network is equivalent to the original network all of whose reactions have become reversible. An example




Fig. 1. Example of a metabolic network with its associated dual network created by the nullspace matrix. Some of its FMs involving target reaction 1 are shown; their \mathcal{I} -coordinated supports result in cut sets for it in the original network

$$\left[\begin{array}{c|c} S & -S \\ \hline I & I \end{array} \right]$$

Fig. 2. This $(m+n) \times 2n$ matrix is the nullspace of the reconfigured nullspace of stoichiometry matrix S . The double description method begins on this space and finds extreme rays with length $2n$

is shown in Figure 1, where the dual network appears alongside the original one. As shown in the example, flux modes with an active target reaction in the dual network map to cut sets for the target reaction in the original network.

Flux modes finders such as FluxModeCalculator (van Klinken and Willems van Dijk, 2016) reconfigure the network before applying the double description method. The double description method is an algorithmic approach for finding the extreme rays of a pointed cone described by linear constraints. The reconfigured network is $N' = [N; -N]$, where N is the null space network of S . Figure 2 shows the null space matrix of N' , which is the starting point of double description method. The double description method begins by using elementary row operations to put the matrix in the form suggested in Wagner (2004) which contains an identity matrix of size $m+n$. When the double description method finishes, it outputs the extreme rays describing a cone in $2n$ -dimensional space (Terzer and Stelling, 2008). These extreme rays are the non-zero vectors in the flux cone with minimal support. On the other hand, \mathcal{I} -coordinated support does not count non-zero values in some dimensions, namely, those that correspond to positive values in irreversible reactions. If we ignore these dimensions, we project the cone into a lower-dimensional subspace. While the image of a pointed cone remains a pointed cone, the extreme rays of the new cone are those in the original flux cone with minimal support in the remaining dimensions. Figure 3 shows why all the MCS² are among the minimal supports, while arguing that there may be some redundant results among them as well. The next theorem formalizes this:

Theorem 1 (Correctness of the method). Algorithm 1 returns precisely the set of MCS of the network \mathcal{M} for a single target reaction.




Fig. 3. Each extreme ray of the projected cone is an image of an extreme ray in the original cone, while some extreme rays of the original cone do not project to extreme rays. It is also possible that two or more extreme rays in the original cone project onto the same one. Our desired projections lie in the plane where the value in the target position is 1

Proof. Let t be the target reaction. We prove the inclusion in both directions. First, let $C \in \mathcal{C}'$ be one of the sets returned by the method above. Then C is a cut set for t in the reconfigured network, by Lemma 1 and by construction. Indeed, \mathcal{F} contains flux modes of N involving t , which are precisely the vectors in the row space of S involving t , and \mathcal{C} (as well as \mathcal{C}') contains the \mathcal{I} -coordinated supports of these vectors.

Now, let C be a MCS for t in \mathcal{M} . We will show that $C \in \mathcal{C}'$. By Lemma 1, there exists a vector $u \in \text{Row}(S)$ such that $u_t = 1$ and $C = \mathcal{R}_{\mathcal{I}}(u)$. Since $u \in \text{Row}(S) \iff u \in \text{Null}(N)$, u is a conical combination of the EFMs of N . Note that the results of Müller and Regensburger (2016) imply that since the space to which u belongs is linear (i.e. it does not need to satisfy any non-negativity constraints), this conical combination can be chosen to be conformal, meaning that there are no cancellations involved in any component. Let such a conformal conical combination be given by

$$\forall 1 \leq i \leq k \quad u = \alpha_1 f_1 + \dots + \alpha_k f_k \quad \text{where } \alpha_i > 0 \quad (4)$$

Since all the coefficients are strictly positive in (4), we deduce that

$$\mathcal{R}_{\mathcal{I}}(u) = \mathcal{R}_{\mathcal{I}}(f_1) \cup \dots \cup \mathcal{R}_{\mathcal{I}}(f_k).$$

Indeed, each $j \in \mathcal{R}_{\mathcal{I}}(u) \cap \mathcal{I}$ must have a negative component in at least one of the f_i , as otherwise the j th component of the right-hand side of (4) will be non-negative, which gives the \subseteq direction, and the fact that the combination is conformal gives the \supseteq direction, as otherwise there would be a cancellation.

In particular, we deduce that $\mathcal{R}_{\mathcal{I}}(f_i) \subseteq \mathcal{R}_{\mathcal{I}}(u)$ for each $1 \leq i \leq k$. In this case, the minimality of C implies that either $\mathcal{R}_{\mathcal{I}}(f_i) = \mathcal{R}_{\mathcal{I}}(u)$ or f_i has a 0 in position t , for each $1 \leq i \leq k$. But since u has a 1 in position t , there must be at least one f_i in the first category, so that $\mathcal{R}_{\mathcal{I}}(f_i) = \mathcal{R}_{\mathcal{I}}(u) = C$ and therefore, $C \in \mathcal{C}'$. Once again, by the minimality of C we conclude that $C \in \mathcal{C}'$ since it cannot be a superset of the \mathcal{I} -coordinated support of another vector in \mathcal{C} , concluding the proof. \square

2.1.1 Limitations

Our method is limited to blocking one direction of a given reaction. In practice, blocking one direction of a given reaction is the typical objective (Burgard et al., 2001). To block multiple reactions, it is possible to compute the MCSs of every target reaction, take unions

space and are optional. These advantages are shared with the target-specific dual network approach.

However, there are additional advantages that MCS² has over this method. First, the null space matrix is typically smaller than the original matrix, especially if the original matrix is nearly full rank, while the target-specific dual network has a matrix that is always larger. This difference in input size can lead to substantial resource savings during EFM computation. Second, and perhaps most importantly, the dual network is independent of the target reaction in our method, while it is not with the target-specific approach. This means that we can calculate these EFMs once and use them for any given target reaction to be blocked.

2.2 Generating a partial set of MCSs via MILP

An alternative strategy for computing EFMs is via MILP, particularly when only a few small MCSs are required instead of a full enumeration (Rezola et al., 2013, 2015). Recall that EFMs are minimal-support vectors in the null space. Our method for finding small MCSs, which we call MCS²-MILP, similarly looks for vectors with MCS² in the row space, which is to say, EFMs with minimal coordinated support in the dual network.

Lemma 4 (MCSs of a target set of reactions in a fully irreversible metabolic network). Let S be the stoichiometric matrix of a fully irreversible metabolic network \mathcal{M} . Let T be a set of target reactions. Then C is a cut set for all the reactions in T if and only if there exist a vector $u \in \text{Row}(S)$ such that $T \subseteq \mathcal{R}_+(u)$ and $\mathcal{R}_-(u) = C$.

Proof. We need to show that every cut set for the target reactions arises from a vector in the row space with the described constraints, and every vector in the row space satisfying those constraints maps to a cut set.

Let C be a cut set for all reactions in T . Therefore, C is a cut set for each reaction in $T = \{t_1, t_2, \dots, t_k\}$ individually. Based on Lemma 1, there exist vectors $u_1, u_2, \dots, u_k \in \text{Row}(S)$ such that $t_i \in \mathcal{R}_+(u_i)$ and $\mathcal{R}_-(u_i) = C$ for $1 \leq i \leq k$. In other words, for all vectors u_i ($1 \leq i \leq k$) the only negative elements are the ones with indices belonging to C , and all other elements are non-negative, with a strictly positive value in the one with index t_i in the vector u_i , for $1 \leq i \leq k$. If we define the vector $u := u_1 + u_2 + \dots + u_k$, then $\mathcal{R}_-(u) = C$ and $T \subseteq \mathcal{R}_+(u)$, and u is clearly in $\text{Row}(S)$.

Now, let u be a vector in $\text{Row}(S)$ such that $T \subseteq \mathcal{R}_+(u)$ and $\mathcal{R}_-(u) = C$. Then $t_i \in \mathcal{R}_+(u)$ for all $1 \leq i \leq k$. Based on Lemma 1, $C = \mathcal{R}_-(u)$ is a cut set for the reaction t_i , for each $1 \leq i \leq k$. Therefore, C is a cut set for all the reactions in T , completing the proof. \square

Based on this Lemma we are able to find MCS for every set of target reactions without the restriction of only blocking one direction of a reaction. Since reversible reactions can be split into two reactions after re-configuration, we can block a reversible reaction in one direction or in both directions.

Let S' be the $m \times n'$ reconfigured matrix of the $m \times n$ stoichiometric matrix S with irreversible reactions \mathcal{I} . Since all the values in the stoichiometric matrix are proportions of consumed and produced metabolites, we can scale each row of S to have only integer entries without changing its structural properties.

We now describe how to encode the problem of finding the smallest MCS for a target set T as a MILP. Let $v \in \mathbb{Z}^m$ be a vector in the row space of the reconfigured matrix S' corresponding to the smallest MCS for target reaction set $T = \{t_1, t_2, \dots, t_k\}$. Then there exists a vector $y \in \mathbb{Z}^m$ s.t. $y^T S = v$. If we define $r^+, r^- \in \{0, 1\}^m$ as the indicator vectors of the positive and negative supports of v , respectively, we may force v_i to be non-positive if $r_i^+ = 0$, and force it to be non-negative if $r_i^- = 0$, by adding the following constraints using r_i^- and r_i^+ as indicator variables:

$$r_i^+ = 0 \Rightarrow v_i \leq 0, r_i^- = 0 \Rightarrow v_i \geq 0 \quad (5)$$

There must also be positive values in the target positions:

$$r_i^+ = 1 \quad \forall i \in T \quad (6)$$

These constraints also ensure that $v=0$ is not in our feasible space. To make v a vector in the row space of S' we need to add the y_j variables, namely, the entries of a vector y with size m . The constraint $y^T S' = v$ then ensures that v is an element of the row space of S' .

The objective function for finding the smallest MCS is:

$$\text{minimize} \sum_{i=1}^{n'} r_i^-, \quad (7)$$

since the cut set is precisely the negative support of v , i.e. r^- .

Suppose that we have found the smallest MCS $C \subseteq \{1, 2, \dots, n'\}$. To find the next smallest MCS, we need to exclude C and all its supersets from our feasible space. This is achieved by the following constraint:

$$\sum_{i \in C} r_i^- \leq |C| - 1 \quad (8)$$

We can keep excluding newly found MCSs and thus enumerate them in order of increasing size. As we stated above, in most scenarios we only wish to block an irreversible reaction or one direction of a reversible reaction. In those cases, we can avoid re-configuring the network to have a smaller stoichiometric matrix. Let t be the only target reaction. Instead of the constraints (6), we only need one constraint $r_t^+ = 1$ if we want to block it in forward direction, and we need the constraint $r_t^- = 1$ if we need to block it in the reverse direction. The objective function (7) and constraints (8) can be updated as follows to reflect the coordinated support instead of the negative support:

$$\text{minimize} \left(\sum_{\substack{i=1 \\ i \neq t}}^n r_i^- + \sum_{\substack{i \in \mathcal{I} \\ i \neq t}} r_i^+ \right) \quad \sum_{i \in C} r_i^- + \sum_{i \in C \cap \mathcal{I}} r_i^+ \leq |C| - 1$$

Unlike von Kamp and Klamt (2014), our problem formulation does not require any additional constraints, because they only reduce a part of the feasible space of our problem without affecting the optimum objective value. This concerns constraints such as $r_i^+ + r_i^- \leq 1$ and $r_i^+ \Rightarrow v_i > \epsilon$.

3 Implementation details

Except where noted, the implementations we discuss are in MATLAB. Each method that we consider requires an extreme ray computation, with the underlying cone varying. We used FluxModeCalculator's EFM generator (van Klinken and Willems van Dijk, 2016) for this purpose. Note that the optimized Berge algorithm implemented by CellNetAnalyzer (Klamt et al., 2007) uses the older EFM finder of CellNetAnalyzer by default. However, we observed that it is a slower implementation of an identical calculation, so we rewrote this part to use FluxModeCalculator in order to make a fair comparison. The MCS² method and the target-specific dual network method both need to remove redundant supersets from the obtained extreme rays, since the desired minimality is not

Table 1. Result of running the methods on the *hepatic polyamine and sulfur amino acid network* (Reyes-Palomares et al., 2012)

All times are in seconds	Optimized Berge	Improved MFK	Target-specific dual network	MCS ² dual network
Extreme ray computation	270.2	270.2	1191.9	79.8
Secondary process time	>18 000	>18 000	591.3	157.4
Total time	>18 000	>18 000	1783.2	237.2

Note: $m = 53$, $n = 73$; target reaction 1.

with respect to the full support of the vector, but only a specific part of it (the \mathcal{I} -coordinated support for the former, and support in the v -coordinates for the latter). We use an implementation in Java whose time complexity is $O(N^2)$ for a collection of N sets. All stoichiometric matrices are compressed by the Mongoose (Leonid et al., 2014) before processing, which converts them to a canonical form.

Since the null space is needed for the row space method, we calculate the null space basis matrix using Mongoose (Leonid et al., 2014). Since finding the MCSs in every method takes several seconds to several minutes, and the computation time of the null space basis matrix is less than a second in every case, we ignore this component of it. The reduced matrix given by Mongoose (for Berge and MFK), the target-specific dual matrix (for the target-specific dual method) and the null space basis matrix (for MCS²) get further compressed by FluxModeCalculator before processing. For the Berge algorithm, we used CellNetAnalyzer (Klamt et al., 2007). We also used an existing implementation of the improved modified Fredman–Khachiyan (MFK) algorithm (Sedaghat et al., 2018). However, we implemented the target-specific dual method from scratch using MATLAB and the source code of FluxModeCalculator. All the enhancements mentioned in the Supplementary Material of the original paper (Ballerstein et al., 2012) were implemented as well.

We used CPLEX (IBM) to solve the MILPs. The implementation was done via the Java API and has been implemented for single target reactions without network reconfiguration, and for multiple target reactions with network reconfiguration. Since the stoichiometric matrices needed to contain only integers, we used the integralize function of MONGOOSE (Leonid et al., 2014) to multiply each row by the smallest possible integer that makes all the values integer (which is the least common multiple of the denominators of its entries). We also tested the results of our MILP in small networks against other implementations to make sure that the results are consistent. The implementation of all the methods and the MILP version of our method are publicly available at <https://github.com/RezaMash/MCS> under the GNU 3.0 license. Some of them require the use of non-public modules available for academic use, such as CellNetAnalyzer (Klamt et al., 2007) and CPLEX (IBM).

4 Results

In this section, we summarize the performance of MCS² and MCS²-MILP in comparison to the other methods.

We ran the implementations on the first 5 models in our database in the GitHub repository. We compared the set of MCSs in these five small examples to confirm that all the implementations produced same results. For the other results presented in Tables 1–3, we again checked the number of MCSs reported by implementations if they finished, and the numbers matched in all cases. We then ran the methods on several models from the BioModels database (Li

Table 2. Result of running the methods on the *kinetic model of yeast network* (Stanford et al., 2013) with $m = 295$, $n = 285$; all the reactions were used as targets

All times are in seconds	Optimized Berge	Improved MFK	Target-specific Dual network	MCS ² Dual network
Extreme ray computation	86.0	86.0	>18 000	53.0
Secondary process time	>18 000 ^a	>18 000	—	13.6
Total time	>18 000 ^a	>18 000	>18 000	66.6

^aBerge computed the MCSs for the first 5 reactions before running out of time. The MFK and target-specific dual methods were not able to finish the computation of the MCSs even for the first reaction.

Table 3. Result of running the methods on *Fernandez2006 ModelB* (Fernandez et al., 2006) with $m = 75$, $n = 152$; target reaction 1

All times are in seconds	Optimized Berge	Improved MFK	Target-specific dual network	MCS ² dual network
Extreme ray computation	99.5	99.5	>18 000	>18 000
Secondary process	2.1	1445.1	—	—
Total time	101.6	1544.6	>18 000	>18 000

Note: This is an example where MCS² and the target-specific dual methods could not finish in time, while the Berge and MFK methods reported all 194 689 MCSs for the compressed network’s reaction 1 fairly quickly.

et al., 2010). There were a few models on which our method either was not able to finish in the given time (5 h) or took much longer to report the MCSs, while the optimized Berge was able to finish in time and beat our method (see Table 3 for an example). This is due to the large number of supersets generated in that example by MCS². However, MCS² always performed better than the target-specific dual network approach, despite all its suggested enhancements being implemented. In addition, as can be seen for the *hepatic polyamine and sulfur amino acid combined model* (Reyes-Palomares et al., 2012), the Berge and MFK methods could not finish in 5 h, but MCS² generated results in 4 min, and the dual method in 30 min.

The first task of every method is an extreme ray computation, which for Berge and MFK is the well-known EFM computation. Berge and MFK then proceed to generate the MCSs through dualization, while the secondary process of the target-specific dual network and MCS² approaches is removing the redundant cut sets. In the first two provided examples, the target is the forward direction of the first reaction. Table 2 shows the computation time for calculating the MCSs for all possible target reactions. In the *kinetic model of yeast metabolic network*, described in Stanford et al. (2013), our method’s advantage is clear—it was able to finish computing the MCSs for all the reactions in under 14 s. Note that the dimensions stated in the tables are the ones before compression is applied. The conclusion is there are models for which it was not feasible to enumerate the full set of MCSs for a given target reaction before our work, but it is feasible now with MCS².

We ran the MILP versions on larger networks alongside the MILP version of the target-specific dual approach, as described in von Kamp and Klamt (2014). This version is also a part of CellNetAnalyzer and is believed to be the state of the art for




Fig. 4. Time (in seconds) for computing each of the 40 smallest MCSs for reaction 10 (the first reaction which has at least 100 MCSs) of the Li2012 calcium-mediated synaptic plasticity model (Li et al., 2012)

extracting some of the smallest MCSs in increasing order of size. We were able to compute 100 MCSs for reaction 10 (the first reaction with at least 100 MCSs) in Li et al. (2012) model, which has 578 reactions after compression. The time required by both approaches on the first 40 MCSs is shown in Figure 4.

MILP was used to find a subset of the MCSs (Song et al., 2017). The target-specific dual method has previously been used for this task, in a method called MCSEnumerator (von Kamp and Klamt, 2014; Vieira et al., 2016). As its authors state, not all the EFMs in the dual space result in valid MCSs, but by adding the appropriate constraints, one can remove the redundant results from the ILP's feasible space. To get a sense of how our approach, MCS²-MILP, performs compared to MCSEnumerator, we implemented the MILP described in von Kamp and Klamt (2014), currently part of CellNetAnalyzer. We ran our implementation of MCSEnumerator (von Kamp and Klamt, 2014) and MCS²-MILP on *E.coli iAF1260* for the sake of comparison, which showed a similar performance, as shown in Table 4. This table contains the result of running MILPs for each reaction as a target reaction once per iteration. In each iteration, we restricted the MILPs to not spend more than 1 min on finding MCSs. Table 5 shows the result of the same experiment repeated with models chosen from the BiGG database (King et al., 2016).

The results in Figure 4 and Table 5 suggest that while MCS²-MILP does not perform faster than the MCSEnumerator, it is slower either. On the other hand, MCS²-MILP is in the first version of itself. It has very simple structure compare to the MCSEnumerator and there is a possibility that it would improve. Even if not, MCS²-MILP is an alternative method which is easier to implement and understand and can help to have a better insight about what is exactly happening during the procedure of these methods.

5 Conclusions and future work

One key advantage of our method is that it does not depend on the target reaction to construct the dual network. The computations for one target reaction can therefore be reused for a different target reaction. Furthermore, it tends to operate on a smaller network than the original.

One limitation to our method is that it is primarily designed for single target reactions (rather than a target containing a set of reactions), while both are just as easily handled by the competitor methods. Although MCS² does not find the MCSs for a set of reactions directly, it can easily find the MCSs for each reaction individually, then prune any supersets from the union of these MCSs.

Table 4. Result of running MCS²-MILP and MCSEnumerator on the *E.coli iAF1260* network with 2382 reactions (981 reactions after compression)

Method used	Average number of MCSs	Average time for shortest MCS	Number of targets MILP failed on
MCS ² -MILP	12.74 MCSs	4.45 s	17
MCSEnumerator	12.07 MCSs	5.22 s	13

Table 5. Result of running MCS²-MILP and MCSEnumerator on the models from the BiGG database which initially have 2000–2600 reactions

Model ID	Average time for shortest MCS for MCSE numerator (s)	Average time for shortest MCS for MCS ² -MILP (s)	Reactions before (after) compression
iJO1366	4.66	3.98	2583 (1106)
iRC1080	7.12	7.19	2191 (1080)
STM_v1_0	1.82	1.83	2545 (1031)
iSbBS512_1146	14.10	19.03	2591 (1018)
iAF1260	5.22	4.45	2382 (981)
iSDY_1059	8.00	9.63	2539 (942)
iYL1228	1.88	2.11	2262 (805)

An alternate strategy for computing MCSs is via MILP, particularly when only a few short sets are required, rather than a complete enumeration (Rezola et al., 2013, 2015). We showed that MCS² can be easily adapted to this task via the MCS²-MILP method, which has shown performance not inferior to that of the state of the art. Another topic for further investigation can be the problem of finding a partial set of irreversible minimal cut sets (iMCSs) (Annika et al., 2018) with MILP. With proper additional constraints, we can find a partial set of iMCSs. There has been some work in enumerating iMCSs; the MILP version of this method is a suitable contender.

Another strategy is to alter the double description method to directly find rays with MCS² instead of minimal support, e.g. by ignoring some of dimensions of the reconfigured network. Here it is important to be careful about zero-cycle flux modes, which are flux modes that have fluxes in both direction of a split reversible reaction. These are not valid flux modes, but they do appear in the output of the double description method (Gagneur and Klamt, 2004) and they may cause the omission of some rays which contain them in their support.

As we mentioned, there are many models for which our method outperforms all other existing methods, while for some models, the best performance is obtained by the Berge algorithm. The challenge is to find out what features of these models are different, and then to decide ahead of time what method to choose for a given model.

Our method is based on novel insights, and may be refined further. Possible additional sources of improvement include identifying and removing unwanted supersets during the execution of the double description method and optimizing the process of superset removal during post-processing. We believe that our method opens the door to further ideas exploring this different kind of duality between EFMs and MCSs, and deeper insights into the structure of metabolic network models.

Funding

TS would like to acknowledge financial support from an NSERC Discovery Grant. LC would like to acknowledge financial support from a Sloan Foundation Fellowship and an NSERC Discovery Grant.

Conflict of Interest: none declared.

References

- Annika, R. *et al.* (2018) Computing irreversible minimal cut sets in genome-scale metabolic networks via flux cone projection. *Bioinformatics*, **34**, 3011–3017.
- Ballerstein, K. *et al.* (2012) Minimal cut sets in a metabolic network are elementary modes in a dual network. *Bioinformatics*, **28**, 381–387.
- Berge, C. (1984) *Hypergraphs: Combinatorics of Finite Sets, Volume 45 of North-Holland Mathematical Library*. Elsevier, Toronto, Canada.
- Bordbar, A. *et al.* (2014) Constraint-based models predict metabolic and associated cellular functions. *Nat. Rev. Genet.*, **15**, 107.
- Burgard, A.P. *et al.* (2001) Minimal reaction sets for *Escherichia coli* metabolism under different growth requirements and uptake environments. *Biotechnol. Prog.*, **17**, 791–797.
- Chindelevitch, L. (2014) An exact arithmetic toolbox for a consistent and reproducible structural analysis of metabolic network models. *Nat Commun.*, **5**, 4893.
- Eiter, T. *et al.* (2008) Computational aspects of monotone dualization: a brief survey. *Discrete Appl. Math.*, **156**, 2035–2049.
- Fernandez, É. *et al.* (2006) DARPP-32 is a robust integrator of dopamine and glutamate signals. *PLoS Comput. Biol.*, **2**, e176.
- Fredman, M. and Khachiyan, L. (1996) On the complexity of dualization of monotone disjunctive normal forms. *J. Algorithm.*, **21**, 618–628.
- Gagneur, J. and Klamt, S. (2004) Computation of elementary modes: a unifying framework and the new binary approach. *BMC Bioinformatics*, **5**, 175.
- Gainer-Dewar, A. and Vera-Licona, P. (2017) The minimal hitting set generation problem: algorithms and computation. *SIAM J. Discrete Math.*, **31**, 63–100.
- Harder, B.-J. *et al.* (2016) Model-based metabolic engineering enables high yield itaconic acid production by *Escherichia coli*. *Metab. Eng.*, **38**, 29–37.
- Hartman, H. *et al.* (2014) Identification of potential drug targets in *Salmonella enterica* sv. Typhimurium using metabolic modelling and experimental validation. *Microbiology*, **160**, 1252–1266.
- Haus, U.-U. *et al.* (2008) Computing knock-out strategies in metabolic networks. *J. Comput. Biol.*, **15**, 259–268.
- IBM. CPLEX optimizer. www-01.ibm.com/software/integration/optimization/cplex-optimizer.
- Imieliński, M. and Belta, C. (2008) Exploiting the pathway structure of metabolism to reveal high-order epistasis. *BMC Syst. Biol.*, **2**, 40.
- King, Z. *et al.* (2016) BiGG Models: a platform for integrating, standardizing and sharing genome-scale models. *Nucleic Acids Res.*, **44**, D515–D522.
- Klamt, S. (2006) Generalized concept of minimal cut sets in biochemical networks. *Biosystems*, **83**, 233–247.
- Klamt, S. *et al.* (2007) Structural and functional analysis of cellular networks with CellNetAnalyzer. *BMC Syst. Biol.*, **1**, 2.
- Klamt, S. and Gilles, E. (2004) Minimal cut sets in biochemical reaction networks. *Bioinformatics*, **20**, 226–234.
- Leonid, C. *et al.* (2014) An exact arithmetic toolbox for a consistent and reproducible structural analysis of metabolic network models. *Nat. Commun.*, **5**, 4893.
- Lewis, N. *et al.* (2012) Constraining the metabolic genotype–phenotype relationship using a phylogeny of in silico methods. *Nat. Rev. Microbiol.*, **10**, 291.
- Li, C. *et al.* (2010) BioModels database: an enhanced, curated and annotated resource for published quantitative kinetic models. *BMC Syst. Biol.*, **4**, 92.
- Li, L. *et al.* (2012) Calcium input frequency, duration and amplitude differentially modulate the relative activation of calcineurin and CaMKII. *PLoS One*, **7**, e43810.
- Mahadevan, R. *et al.* (2015) Genome-scale strain designs based on regulatory minimal cut sets. *Bioinformatics*, **31**, 2844–2851.
- Müller, S. and Regensburger, G. (2016) Elementary vectors and conformal sums in polyhedral geometry and their relevance for metabolic pathway analysis. *Front. Genet.*, **7**, 90.
- Price, N. *et al.* (2004) Genome-scale models of microbial cells: evaluating the consequences of constraints. *Nat. Rev. Microbiol.*, **2**, 886.
- Reyes-Palomares, A. *et al.* (2012) A combined model of hepatic polyamine and sulfur amino acid metabolism to analyze S-adenosyl methionine availability. *Amino Acids*, **42**, 597–610.
- Rezola, A. *et al.* (2013) Selection of human tissue-specific elementary flux modes using gene expression data. *Bioinformatics*, **29**, 2009–2016.
- Rezola, A. *et al.* (2015) Advances in network-based metabolic pathway analysis and gene expression data integration. *Brief. Bioinformatics*, **16**, 265–279.
- Schuster, S. *et al.* (2000) A general definition of metabolic pathways useful for systematic organization and analysis of complex metabolic networks. *Nat. Biotechnol.*, **18**, 326.
- Sedaghat, N. *et al.* (2018) Speeding up dualization in the Fredman–Khachiyan algorithm B. In *Leibniz International Proceedings in Informatics, Caen, France*, Vol. 103. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- Song, H.-S. *et al.* (2017) Sequential computation of elementary modes and minimal cut sets in genome-scale metabolic networks using alternate integer linear programming. *Bioinformatics*, **33**, 2345–2353.
- Stanford, N. *et al.* (2013) Systematic construction of kinetic models from genome-scale metabolic networks. *PLoS One*, **8**, e79195.
- Terzer, M. and Stelling, J. (2008) Large-scale computation of elementary flux modes with bit pattern trees. *Bioinformatics*, **24**, 2229–2235.
- Trinh, C.T. *et al.* (2006) Design, construction and performance of the most efficient biomass producing *E. coli* bacterium. *Metab. Eng.*, **8**, 628–638.
- van Klinken, J.B. and Willems van Dijk, K. (2016) FluxModeCalculator: an efficient tool for large-scale flux mode computation. *Bioinformatics*, **32**, 1265–1266.
- Vieira, V. *et al.* (2016) Development of an integrated framework for minimal cut set enumeration in constraint-based models. In: *10th International Conference on Practical Applications of Computational Biology & Bioinformatics, Ávila, Spain*, pp. 193–201. Springer.
- von Kamp, A. and Klamt, S. (2014) Enumeration of smallest intervention strategies in genome-scale metabolic networks. *PLoS Comput. Biol.*, **10**, e1003378.
- von Kamp, A. and Klamt, S. (2017) Growth-coupled overproduction is feasible for almost all metabolites in five major production organisms. *Nat. Commun.*, **8**, 15956.
- Wagner, C. (2004) Nullspace approach to determine the elementary modes of chemical reaction systems. *The Journal of Physical Chemistry B*, **108**, 2425–2431.
- Zabeti, H. *et al.* (2018) A duality-based method for identifying elemental balance violations in metabolic network models. In: *Leibniz International Proceedings in Informatics, Caen, France*, Vol. 113. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.