Algorithmica (2008) 50: 159–172
DOI 10.1007/s00453-007-9074-x

On Enumerating Minimal Dicuts and Strongly Connected Subgraphs

**Leonid Khachiyan · Endre Boros ·
Khaled Elbassioni · Vladimir Gurvich**

Received: 19 April 2006 / Accepted: 19 June 2006 / Published online: 27 October 2007

Abstract We consider the problems of enumerating all minimal strongly connected subgraphs and all minimal dicuts of a given strongly connected directed graph $G = (V, E)$. We show that the first of these problems can be solved in incremental polynomial time, while the second problem is NP-hard: given a collection of minimal dicuts for G, it is NP-hard to tell whether it can be extended. The latter result implies, in particular, that for a given set of points $\mathcal{A} \subseteq \mathbb{R}^n$, it is NP-hard to generate all maximal subsets of $\mathcal{A}$ contained in a closed half-space through the origin. We also discuss the enumeration of all minimal subsets of $\mathcal{A}$ whose convex hull contains the origin as an interior point, and show that this problem includes as a special case the well-known hypergraph transversal problem.

1 Introduction

Let V be a finite set of vertices and $G = (V, E)$ a *connected* undirected graph on V with edge set $E \subseteq V \times V$. A minimal *cut* in G is a minimal set of edges

This research was supported by the National Science Foundation (Grant IIS-0118635). The third and fourth authors are also grateful for the partial support by DIMACS, the National Science Foundation's Center for Discrete Mathematics and Theoretical Computer Science.

Our friend and co-author, Leonid Khachiyan tragically passed away on April 29, 2005.

E. Boros · V. Gurvich
RUTCOR, Rutgers University, 640 Bartholomew Road, 08854-8003, Piscataway, NJ USA

E. Boros
e-mail: boros@rutcor.rutgers.edu

V. Gurvich
e-mail: gurvich@rutcor.rutgers.edu

K. Elbassioni (✉)
Max-Planck-Institut für Informatik, Saarbrücken, Germany
e-mail: elbassio@mpi-sb.mpg.de

the removal of which breaks G into two components. Given two specified vertices $s, t \in V$, a minimal (s,t)-*cut* in G is a minimal set of edges whose removal disconnects s and t. Similar terms are also defined for directed graphs (digraphs). Given a *strongly connected* digraph $G = (V, E)$ with arc set $E \subseteq V \times V$, a *minimal directed cut*, or a *dicut* is a minimal subset of arcs the removal of which leaves a non-strongly connected digraph. Given two specified vertices $s, t \in V$, a minimal (s,t)-*dicut* is a minimal subset of arcs whose removal leaves no directed path from s to t.

These notions play an important role in network reliability, where edges or arcs represent communication or transportation links, which may work or fail independently, and where the main problem is to determine the probability that the network is working, based on the individual edge/arc failure probabilities. It turns out that such network reliability computations require in the general case the list of minimal cuts, dicuts, (s,t)-cuts, etc., depending on the type of connectivity the network is ought to maintain (i.e., all-terminal, two-terminal, strong, etc.), see e.g., [1, 3, 6, 16].

It is easy to see that the number of spanning trees, cuts, (s,t)-paths, (s,t)-cuts, etc. may, in general, be exponential in the size of the graph. For this reason efficiency of their generation is measured in both the input and output sizes, e.g., we shall talk about the complexity of generation "per cut".

Given a (strongly) connected (di)graph $G = (V, E)$, we shall consider the problem of listing all minimal subgraphs of G, i.e. the family $\mathcal{F}_\pi \subseteq 2^E$ of all minimal subsets of E, satisfying a given monotone property $\pi : E \mapsto \{0, 1\}$. For instance, if $\pi(X)$ is the property that the subgraph with edge set $X \subseteq E$ is connected, then $\mathcal{F}_\pi$ is the family of spanning trees of G. Note that, with each family of subgraphs $\mathcal{F}_\pi$ satisfying a monotone property π, we can associate the dual family

$$\mathcal{F}_\pi^d = \{X \subseteq E \ : \ X \text{ is a minimal transversal of } \mathcal{F}_\pi\},$$

where $X \subseteq E$ is a *transversal of* $\mathcal{F}_\pi$ if and only if $X \cap Y \neq \emptyset$ for all $Y \in \mathcal{F}_\pi$. Let us also introduce the *complementary* family $\mathcal{F}_\pi^c = \{E \setminus X \mid X \in \mathcal{F}_\pi\}$ whose elements are complementary to the elements of $\mathcal{F}_\pi$. Thus if $\pi(X)$ is the property that $G' = (V, X)$ is connected, then $\mathcal{F}_\pi^d$ is the family of all minimal cuts of $G = (V, E)$, and $\mathcal{F}_\pi^{dc}$ is the family of maximal non-connected subgraphs of G.

Enumeration algorithms for listing subgraphs satisfying a number of monotone properties are well known. For instance, it is known [18] that the problems of listing all minimal cuts or all spanning trees of an undirected graph $G = (V, E)$ can be solved with delay $O(|E|)$ per generated cut or spanning tree. It is also known (see e.g., [8, 12, 17]) that all minimal (s,t)-cuts or (s,t)-paths, can be listed with delay $O(|E|)$ per cut or path, both in the directed and undirected cases. Furthermore, if $\pi(X)$ is the property that the subgraph (V, X) of a directed graph $G = (V, E)$ contains a directed cycle, then $\mathcal{F}_\pi$ is the family of minimal directed circuits of G, while $\mathcal{F}_\pi^d$ consists of all minimal feedback arc sets of G (i.e. minimal sets of arcs whose removal breaks every directed circuit in G). Both of these families can be generated with polynomial delay per output element, see e.g. [19].

It is quite remarkable that in all these cases both the family $\mathcal{F}_\pi$ and its dual $\mathcal{F}_\pi^d$ can be generated efficiently, unlike for many other monotone families, [5, 15]. In this paper we focus on a case, relevant to reliability theory, when this symmetry is broken.

Given a strongly connected digraph $G = (V, E)$, let $sc(X)$ be the property that $X \subseteq E$ is strongly connected on V. Then $\mathcal{F}_{sc}$ is the family of minimal strongly connected subgraphs (V, X) of G, and $\mathcal{F}_{sc}^d$ is the family of minimal dicuts of G. In this note, we show that the problem of incrementally generating all minimal dicuts is NP-hard.

Theorem 1 *Given a strongly connected directed graph $G = (V, E)$ and a partial list $\mathcal{X} \subseteq \mathcal{F}_{sc}^d$ of minimal dicuts of G, it is NP-hard to determine if the given list is complete, i.e. if $\mathcal{F}_{sc}^d = \mathcal{X}$.*

We show also that on the contrary, listing minimal strongly connected subgraphs can be done efficiently.

Theorem 2 *Given a strongly connected directed graph $G = (V, E)$, all minimal strongly connected subgraphs of G can be listed in incremental polynomial time.*

We prove Theorems 1 and 2 in Sects. 2 and 3 respectively. We close with some geometric generalizations of these problems in Sect. 4. Specifically, for a given set of points $\mathcal{A} \subseteq \mathbb{R}^n$, we consider the problems of enumerating all minimal subsets of $\mathcal{A}$ whose convex hull contains the origin as an interior point, and all maximal subsets of $\mathcal{A}$ whose convex hull does not contain the origin as an interior point. We will show that the former is at least as hard as the well-known hypergraph transversal problem, while the latter turns out to be NP-hard by Theorem 1. Finally, we will discuss the relation of these problems to the problem of vertex enumeration, i.e., to the problem of generating all vertices of a polyhedron given as the set of feasible solutions to a system of linear inequalities. This generation problem is known to be NP-hard (see [14]), however it is still open for bounded polyhedra.

2 Proof of Theorem 1

Let us state first the following easy but very useful characterization of minimal dicuts in a directed graph. For a directed graph $G = (V, E)$ and a subset $S \subseteq V$ of its vertices let us denote by $G[S]$ the subgraph of G induced by the vertex set S.

Lemma 1 *Given a strongly connected digraph $G = (V, E)$, an arc set $X \subseteq E$ is a minimal dicut if and only if there exist vertex sets $S, T, R \subseteq V$ satisfying the following four properties.*

(D1) *$S \neq \emptyset$, $T \neq \emptyset$, S, T, R are pairwise disjoint, $S \cup T \cup R = V$.*
(D2) *$G[S]$ and $G[T]$ are both strongly connected.*
(D3) *No arc of G is connecting S to R, or R to T.*
(D4) *$X = \{(a, b) \in E \ : \ a \in S,\ b \in T\}$ is the set of all arcs from S to T.*

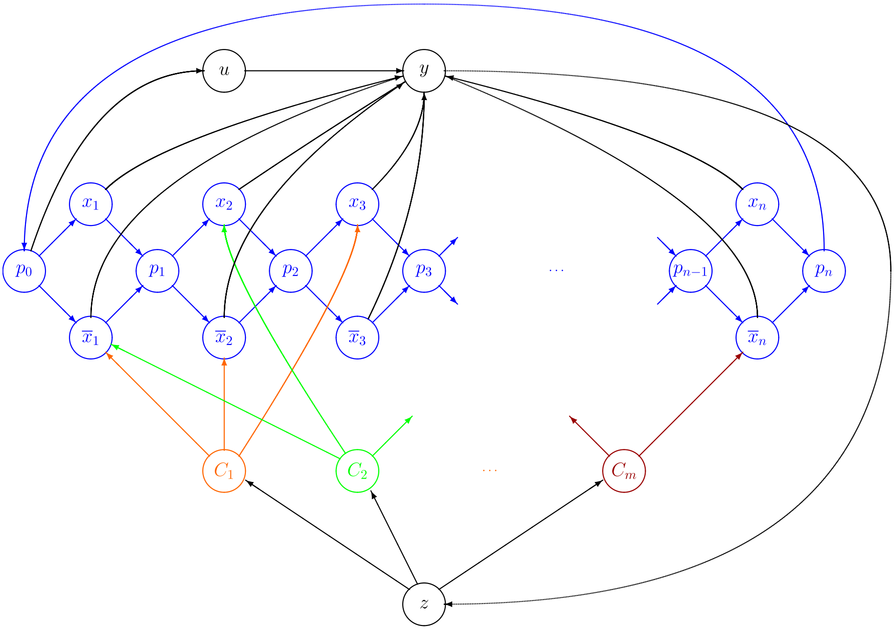

Fig. 1 An example for the NP-hard reduction. The SAT problem $(\overline{x}_1 \vee \overline{x}_2 \vee x_3) \wedge (\overline{x}_1 \vee x_2 \vee \cdots) \wedge \cdots \wedge (\cdots \vee \overline{x}_n) = 1$ has a solution iff the associated graph has a nontrivial arc set destroying its strong connectivity

Proof The statement follows from the fact that the component digraph of the graph obtained from G by deleting a minimal dicut X must have exactly one minimal and one maximal strongly connected components T and S, respectively. □

To prove the theorem, we use a polynomial transformation from the satisfiability problem. Let $\Phi = C_1 \wedge \cdots \wedge C_m$ be a conjunctive normal form of m clauses and $2n$ literals $\{x_1, \overline{x}_1, \ldots, x_n, \overline{x}_n\}$. In what follows, it is assumed without loss of generality that for each $i = 1, \ldots, n$, both x_i and $\overline{x}_i$ occur in Φ. We construct a strongly connected digraph $G = (V, E)$ with $|V| = m + 3n + 4$ vertices and $|E| = \sum_{i=1}^{m} |C_i| + m + 6n + 4$ arcs. See Fig. 1 for an example.

The Vertices The vertex set of G is defined as follows. There are m vertices $C_1, \ldots, C_m$ corresponding to the m clauses, $2n$ vertices $x_1, \overline{x}_1, \ldots, x_n, \overline{x}_n$ corresponding to the $2n$ literals, $n + 1$ vertices $p_0, p_1, \ldots, p_n$, and finally 3 other vertices z, u, y.

The Arcs There is an arc (z, C_j) from vertex z to every clause vertex C_j for $j = 1, \ldots, m$, an arc (ℓ, y) from each literal $\ell \in \{x_1, \overline{x}_1, \ldots, x_n, \overline{x}_n\}$ to vertex y, and an arc (C, ℓ) for each clause C and literal ℓ appearing in it. For $i = 1, \ldots, n$, we also have arcs (p_{i-1}, x_i), $(p_{i-1}, \overline{x}_i)$, (x_i, p_i), and $(\overline{x}_i, p_i)$. Finally, we add the arcs (p_n, p_0), (p_0, u), (u, y), and (y, z).

Trivial Dicuts Let us call a minimal dicut *trivial* if it is the set of arcs leaving a single vertex $v \in V$, or the set of arcs entering a single vertex $v \in V$. Clearly, not all sets of arcs leaving or entering a vertex are minimal dicuts. However, the number of such minimal dicuts does not exceed twice the number of vertices, which is polynomial in n and m.

Non-Trivial Minimal Dicuts Let us now show that any non-trivial dicut yields a satisfying assignment for Φ and conversely, any satisfying assignment for Φ gives a non-trivial minimal dicut for G. This will prove Theorem 1.

Let $\sigma = (\ell_1, \ell_2, \dots, \ell_n)$ be the set of literals assigned the value $True$ in a satisfying truth assignment for Φ. We define a minimal dicut X of G corresponding to σ. For this, we use the characterization of Lemma 1, i.e. give the corresponding sets S, T, and R. Let

$$\begin{aligned} T &= \{p_0, \overline{\ell}_1, p_1, \overline{\ell}_2, \dots, p_{n-1}, \overline{\ell}_n, p_n\}, \\ S &= \{z, C_1, \dots, C_m, \ell_1, \dots, \ell_n, y\}, \\ R &= \{u\}. \end{aligned} \tag{1}$$

Then it easy to verify that the assignments (1) satisfy the conditions (D1–D3) of Lemma 1, and define a non-trivial minimal dicut.

To see the converse direction, let us consider a non-trivial minimal dicut $X \subseteq E$. We use Lemma 1 again to present a corresponding satisfying truth assignment for Φ. Since X is a minimal dicut, there exist sets S, T, and R satisfying conditions (D1–D4) of Lemma 1. We present the case analysis in the following steps:

- (S0) None of the arcs (p_n, p_0), (p_0, u), (u, y), (y, z), and (z, C_j) for $j = 1, \dots, m$ can belong to X, since each of these arcs alone form a trivial minimal dicut.
- (S1) We must have $|S| \geq 2$ and $|T| \geq 2$ by the non-triviality of X.
- (S2) We must have $y \in S$, since otherwise all vertices from which y is reachable in $E \setminus X$ must belong to $R \cup T$ by (D3), implying by (S0) that $\{p_0, u, y\} \subseteq R \cup T$. Thus, $S \subseteq V \setminus \{p_0, u, y\}$ would follow, implying $|S| = 1$ in contradiction with (S1), since the vertex set $V \setminus \{p_0, u, y\}$ induces an acyclic subgraph of G.
- (S3) Then, $\{y, z, C_1, \dots, C_m\} \subseteq S$ is implied by (S2), (D3) and (S0).
- (S4) We must have $p_0 \in T$, since otherwise all vertices reachable from p_0 in $E \setminus X$ must belong to $S \cup R$ by (D3), implying by (S0) that $\{p_0, u, y\} \subseteq S \cup R$. Thus $T \subseteq V \setminus \{p_0, u, y\}$ would follow, implying $|T| = 1$ in contradiction with (S1), since the vertex set $V \setminus \{p_0, u, y\}$ induces an acyclic subgraph of G.
- (S5) Then, $\{p_0, p_n\} \subseteq T$ is implied by (S4), (D3) and (S0).
- (S6) The strong connectivity of $G[T]$ and (S5) then imply that there exist literals $\ell_i \in \{x_i, \overline{x}_i\}$ for $i = 1, \dots, n$ such that $\{p_0, \ell_1, p_1, \ell_2, \dots, p_{n-1}, \ell_n, p_n\} \subseteq T$.
- (S7) The strong connectivity of $G[S]$ and (S3) then imply that there exist literals $\overline{\ell}_{i_j} \in S$ which belong to C_j for all $j = 1, \dots, m$. Furthermore, we are guaranteed by (S6) that $\ell_{i_j} \neq \overline{\ell}_{i_k}$ for all $k, j = 1, \dots, m$.

Now it is easy to see by (S7) and by the construction of the graph G that assigning $\overline{\ell}_{i_j} \leftarrow True$, for all $j = 1, \dots, m$, yields a satisfying truth assignment for Φ.

Let us remark before closing this section that analogously to Lemma 1, characterizations of minimal cuts, (s,t)-cuts, and (s,t)-dicuts are also well-known, and can be used for efficient generation. For completeness, we include here these analogous statements:

Lemma 2 *Given a connected graph $G=(V,E)$, an edge set $X \subseteq E$ is a minimal cut if and only if there exist vertex sets $S, T \subseteq V$ such that*

(U1) $S \neq \emptyset$, $T \neq \emptyset$, $S \cap T = \emptyset$, $S \cup T = V$, *and*
(U2) $G - X = G[S] \cup G[T]$, *and both* $G[S]$ *and* $G[T]$ *are connected.*

Furthermore, given two vertices $s,t \in V$, an edge set $X \subseteq E$ is a minimal (s,t)-cut if and only if there exist vertex sets $S, T \subseteq V$ satisfying (U1), (U2) *and*

(U3) $s \in S$, $t \in T$.

Lemma 3 *Given a strongly connected digraph $G=(V,E)$ and two specified vertices $s,t \in V$, an arc set $X \subseteq E$ is a minimal (s,t)-dicut if and only if there exist disjoint vertex sets $S, T \subseteq V$ satisfying* (D4) *of Lemma* 1, *and the following properties*:

(D5) $s \in S$, $t \in T$,
(D6) *for any $v \in S$, there is a directed path from s to v in $G[S]$, and*
(D7) *for any $v \in T$, there is a directed path from v to t in $G[T]$.*

Proofs of Lemmas 2 and 3 can be found, e.g., in [17].

3 Enumerating Minimal Strongly Connected Subgraphs

Let $G=(V,E)$ be a given strongly connected digraph and $\mathcal{F}_{sc} \subseteq 2^E$ the family of all minimal strongly connected subgraphs of G. We generate all elements of $\mathcal{F}_{sc}$ by performing a traversal (for instance, breadth-first-search) of a directed "supergraph" $\mathcal{G} = (\mathcal{F}_{sc}, \mathcal{E})$ on vertex set $\mathcal{F}_{sc}$. Let $S \in \mathcal{F}_{sc}$ be a "vertex" of $\mathcal{G}$, then we define the neighborhood $\mathcal{E}^+(S) \subseteq \mathcal{F}_{sc}$ of the immediate successors of S in $\mathcal{G}$ to consist of all minimal strongly connected subgraphs T of G which can be obtained from S by the following process:

1. Let $e=(a,b) \in S$ be an arc of G such that the graph $(V, E \setminus e)$ is strongly connected. Delete e from S.
2. Add a minimal set W of arcs from $E \setminus S$ to restore the strong connectivity of $(S \setminus e) \cup W$, i.e. the reachability of b from a.
3. Lexicographically delete some arcs Y from $S \setminus e$ to guarantee the minimality of $T = (S \setminus (Y \cup e)) \cup W$.

(We assume in Step 3 that we have fixed some order on the arcs of G.)

To illustrate, consider the strongly connected subgraph S on three vertices a, b, c with arcs $\{(a,b),(b,a),(b,c),(c,b)\}$. Let $e=(a,b)$ and $W=\{(a,c)\}$. Then $(S \setminus e) \cup W$ is strongly connected and can be made minimal by deleting arc (b,c).

Theorem 2 readily follows from the following lemma.

Lemma 4

(i) *The supergraph $\mathcal{G}$ is strongly connected.*
(ii) *For each vertex $S \in \mathcal{F}_{sc}$, the neighborhood of S can be generated with polynomial delay.*

Proof (i) Let $S, S' \in \mathcal{F}_{sc}$ be two distinct vertices of $\mathcal{G}$. To show that $\mathcal{G}$ contains an (S, S')-path, consider an arbitrary arc $e = (a, b) \in S \setminus S'$. Since $S \setminus \{e\} \cup S'$ is strongly connected, we can find a minimal set of arcs $W \subseteq S' \setminus S$ such that b is reachable from a in $S \setminus \{e\} \cup W$. Lexicographically minimizing the set of arcs $S \setminus \{e\} \cup W$ over $S \setminus e$, we obtain an element S'' in the neighborhood of S with a difference $|S'' \setminus S'|$ smaller than $|S \setminus S'|$. This implies that $\mathcal{G}$ is strongly connected and has diameter linear in n.
(ii) Let us start with the observation that for any two distinct minimal sets W and W' in Step 2, the neighbors resulting after Step 3 are distinct. Therefore, it suffices to show that all minimal arc sets W in Step 2 can be generated with polynomial delay.

For convenience, let us color the arcs in $S \setminus \{e\}$ black, and color the remaining arcs in $E \setminus S$ white. So we have to enumerate all minimal subsets W of white arcs such that b is reachable from a in $G(W) = (V, (S \setminus \{e\}) \cup W)$. Let us call such subsets of white arcs *minimal white* (a, b)*-paths*. The computation of these paths can be done by using a backtracking algorithm that performs depth first search on the following recursion tree $\mathbf{T}$ (see [18] for general background on backtracking algorithms). Each node (z, W_1, W_2) of the tree is identified with a vertex $z \in V$ and two disjoint subsets of white arcs W_1 and W_2, such that W_1 is a minimal white (z, b)-path that can be extended to a minimal white (a, b)-path by adding some arcs from $W \setminus (W_1 \cup W_2)$. The root of the tree is $(b, \emptyset, \emptyset)$ and the leaves of the tree are those nodes (z, W_1, W_2) for which $a \in B(z)$, where $B(z)$ consists of all vertices $v \in V$ such that z can be reached from v in $S \setminus \{e\}$, that is by using only black arcs. As we shall see, the set W_1 for each leaf of $\mathbf{T}$ is a minimal white (a, b)-path, and each minimal white (a, b)-path will appear exactly once on a leaf of $\mathbf{T}$.

We now define the children of an internal node (z, W_1, W_2), where $a \notin B(z)$. Let Z be the set of all white arcs from $W \setminus (W_1 \cup W_2)$ which enter $B(z)$. Pick an arc $e \in Z$. If the tail y of e is reachable from a in $G(W \setminus (W_1 \cup W_2 \cup Z))$, then $(y, W_1 \cup \{e\}, W_2 \cup (Z \setminus \{e\}))$ is a child of the node (z, W_1, W_2) in $\mathbf{T}$. It is easy to see that $W_1 \cup \{e\}$ is indeed a minimal white (y, b)-path: Let $e_1, e_2, \ldots, e_k = e$ be the set of arcs added to $W_1 \cup \{e\}$, in that order, and let $(z^1, W_1^1, W_2^1), (z^2, W_1^2, W_2^2), \ldots, (z^k, W_1^k, W_2^k)$ be the set of nodes from the root of tree $\mathbf{T}$ to node $(y, W_1 \cup \{e\}, W_2 \cup (Z \setminus \{e\}))$. Then the set $W_1^k \setminus \{e_1\}$ does not contain any white path since it contains no arcs entering $B(b)$. More generally, for $i = 1, \ldots, k$, the set $W_1^k \setminus \{e_i\}$ does not contain any white arc entering $B(b) \cup B(z^1) \cup \cdots \cup B(z^{i-1})$ and hence it contains no white (y, b)-path. Note also that this construction guarantees that in addition to the minimality of the white (y, b)-path $W_1^k = W_1 \cup \{e\}$, it can be extended to a minimal white (a, b) path by adding some arcs from $(W \setminus (W_1 \cup W_2 \cup Z))$. Similar arguments also show that distinct leaves of $\mathbf{T}$ yield distinct minimal white (a, b)-paths, and that all such distinct paths appear as distinct leaves in $\mathbf{T}$.

Note that the depth of the backtracking tree is at most $|V|$, and that the time spent at each node is polynomial in $|V|$ and $|E|$. This proves (ii). $\square$

As mentioned earlier, by performing a transversal on the nodes of the supergraph $\mathcal{G}$, we can generate the elements of $\mathcal{F}_{sc}$ in incremental polynomial time. However, we cannot deduce from Lemma 4 that the set $\mathcal{F}_{sc}$ can be generated with polynomial delay since the size of the neighborhood of a given vertex $S \in \mathcal{F}_{sc}$ may be exponentially large.

We close this section with the following observation. It is well known that the number of spanning trees (i.e., minimal edge sets ensuring the connectivity of the graph) for an undirected graph G can be computed in polynomial time (see, e.g., [4]). In contrast to this result, given a strongly connected digraph G with m arcs, it is NP-hard to approximate the size of $\mathcal{F}_{sc}$ to within a factor of $2^{m^{1-\epsilon}}$, for any fixed $\epsilon > 0$. To see this, pick two vertices s and t in G and let $G\{s,t\}$, be the digraph obtained from G by adding, for each vertex $v \in V \setminus \{s,t\}$, two auxiliary vertices v', v'' and four auxiliary arcs (t,v'), (v',v), (v,v''), (v'',s). It is easy to see that any minimal strongly connected subgraph of $G\{s,t\}$ contains all the auxiliary arcs and some (s,t)-path in G. Hence there is a one-to-one correspondence between the set $\mathcal{F}_{sc}$ for $G\{s,t\}$ and the set of directed (s,t)-paths for G.[1] Now the claim follows by using the amplification technique of [13], which replaces each arc of G by $(2m)^{1/\epsilon}$ consecutive pairs of parallel paths of length 2. It follows then that any approximation of the number of (s,t)-paths in the resulting graph G' to within an accuracy of $2^{(m')^{1-\epsilon}}$, where m' is the number of arcs in G', can be used to compute the longest (s,t)-path in G, a problem that is known to be NP-hard.

A stronger inapproximability result for counting minimal dicuts is implied by the NP-hardness proof of Theorem 1: Unless P = NP, there is a constant $c > 0$, such that no polynomial-time algorithm can approximate the number of minimal dicuts of a given strongly connected directed graph G to within a factor of 2^{cm}, where m is the number of arcs of G. This can be seen, for instance, as follows. Let $\Phi(x_1, \overline{x}_1, \ldots, x_n, \overline{x}_n)$ be a CNF of k clauses on $2n$ literals. Replace Φ by

$$\Phi' = \Phi \wedge \bigwedge_{j=1}^{s} (y_j \vee z_j)(\overline{y}_j \vee \overline{z}_j),$$

where $y_1, \ldots, y_s$ and $z_1, \ldots, z_s$ are new variables. This way we obtain a new CNF Φ' of $k + 2s$ clauses on $2n + 4s$ literals such that Φ has a satisfying assignment if and only if Φ' has at least 2^s satisfying assignments. Let G be the digraph with $O(k+n+s)$ arcs constructed for Φ' in the proof of Theorem 1. Now the result follows from the fact that it is NP-hard to determine whether G has $O(k+n+s)$ or more than 2^s minimal dicuts.

4 Some Related Geometric Problems and Concluding Remarks

Let $\mathcal{A} \subseteq \mathbb{R}^n$ be a given subset of points in $\mathbb{R}^n$. Fix a point $z \in \mathbb{R}^n$, say $z = \mathbf{0}$, and consider the following four geometric objects:

[1] In particular, this shows that maximizing the number of arcs in a minimal strongly connected subgraph of G is NP-hard. Minimizing the number of arcs in such a subgraph is also known to be NP-hard [11].

A Simplex a minimal subset $X \subseteq \mathcal{A}$ of points containing z in its convex hull: $z \in \text{conv.hull}(X)$.

An Anti-Simplex a maximal subset $X \subseteq \mathcal{A}$ of points not containing z in its convex hull: $z \notin \text{conv.hull}(X)$.

A Body a minimal (full-dimensional) subset $X \subseteq \mathcal{A}$ of points containing z in the interior of its convex hull: $z \in \text{int conv.hull}(X)$.

An Anti-Body a maximal subset $X \subseteq \mathcal{A}$ of points not containing z in the interior of its convex hull: $z \notin \text{int conv.hull}(X)$.

Equivalently, a simplex (body) is a minimal collection of points not contained in an *open* (*closed*) half-space. An anti-simplex (anti-body) is a maximal collection of points contained in an open (closed) half-space. It is known that $|X| \le n+1$ for any simplex $X \subseteq \mathcal{A}$ and that $n+1 \le |X| \le 2n$ for any body $X \subseteq \mathcal{A}$.

For a given point set $\mathcal{A}$, denote respectively by $\mathcal{S}(\mathcal{A})$ and $\mathcal{B}(\mathcal{A})$ the families of simplices and bodies of $\mathcal{A}$ with respect to the origin $z = \mathbf{0}$. Then it is clear that $\mathcal{S}(\mathcal{A})^{dc}$ and $\mathcal{B}(\mathcal{A})^{dc}$ are respectively the families of anti-simplices and anti-bodies of $\mathcal{A}$.

Let $A \in \mathbb{R}^{m\times n}$, where $m = |\mathcal{A}|$, be the matrix whose rows are the points of $\mathcal{A}$. It follows from the above definitions that simplices and anti-simplices are in one-to-one correspondence respectively with the *minimal infeasible* and *maximal feasible* subsystems of the linear system of inequalities:

$$Ax \ge \mathbf{e}, \quad x \in \mathbb{R}^n \tag{2}$$

where $e \in \mathbb{R}^m$ is the m-dimensional vector of all ones. Similarly, it follows that bodies and anti-bodies correspond respectively to the minimal infeasible and maximal feasible subsystems of the system:

$$Ax \ge \mathbf{0}, \quad x \ne \mathbf{0}. \tag{3}$$

As a special case of the above problems, let $G = (V, E)$ be a directed graph, and let $\mathcal{A} \subseteq \{-1, 0, 1\}^V$ be the set of incidence vectors corresponding to the arc-set E, i.e. $\mathcal{A} = \{\chi(a, b) \ : \ (a, b) \in E\}$, where $\chi = \chi(a, b)$ is defined for an arc $(a, b) \in E$ by

$$\chi_v = \begin{cases} 1 & \text{if } v = a, \\ -1 & \text{if } v = b, \\ 0 & \text{otherwise.} \end{cases}$$

Denote by $A \in \mathbb{R}^{|E|\times|V|}$ the corresponding incidence matrix of G. Note that, for any subgraph G' of G, the corresponding subsystem of (2) defined by the arcs of G' is feasible if and only if G' is acyclic. Thus it follows that the simplices $\mathcal{S}(\mathcal{A})$ are in one-to-one correspondence with the *simple directed circuits* of G. By definition, an anti-simplex is a maximal subset of points not containing any simplex. Thus, the anti-simplices of $\mathcal{A}$ correspond to the complements of the *minimal feedback arc sets* (recall that feedback arc sets in a directed graph are sets of arcs whose removal breaks every directed circuit the graph).

Now, let us consider bodies and anti-bodies of $\mathcal{A}$. Fix a vertex $v \in V$ and consider the system of inequalities (3) together with the equation $x_v = 0$ (or equivalently, remove the v-th column of A and the v-th component of x). Then it is easy to see that the subsystem of (3) (together with $x_v = 0$) defined by the arcs of a subgraph G' of G is infeasible if and only if G' is strongly connected. In particular, the family of bodies $\mathcal{B}(\mathcal{A})$ is in one-to-one correspondence with the family of minimal strongly connected subgraphs of G, and the family of anti-bodies $\mathcal{B}(\mathcal{A})^{dc}$ is in one-to-one correspondence with the (complementary) family of minimal dicuts of G.

Given a directed graph $G = (V, E)$, it is known that all simple circuits of G can be listed with polynomial delay (see, e.g., [18]). It is also known [19] that all minimal feedback arc sets for a directed graph G can be listed with polynomial delay. Theorem 2 states that we can also list, in incremental polynomial time, all minimal strongly connected subgraphs of G, while Theorem 1 states that such a result cannot hold for the family of minimal dicuts unless $P = NP$.

Thus, as a consequence of Theorem 1, we obtain the following negative result.

Corollary 1 *Given a set of points $\mathcal{A} \subseteq \mathbb{R}^n$, and a partial list $\mathcal{X} \subseteq \mathcal{B}(\mathcal{A})^{dc}$ of anti-bodies of $\mathcal{A}$, it is NP-hard to determine if the given list is complete, i.e. $\mathcal{X} = \mathcal{B}(\mathcal{A})^{dc}$. Equivalently, given an infeasible system* (3), *and a partial list of maximal feasible subsystems of* (3), *it is NP-hard to determine if the given partial list is complete.*

We now turn to the enumeration of all bodies for $\mathcal{A}$. In contrast to Theorem 2, the general case of the enumeration problem for $\mathcal{B}(\mathcal{A})$ turns out to be at least as hard as the well-known *hypergraph transversal problem* [9], which is not known to be solvable in incremental polynomial time.

Proposition 1 *The problem of incrementally enumerating bodies, for a given set of $m + n$ points $\mathcal{A} \subseteq \mathbb{R}^n$, includes as a special case the problem of enumerating all minimal transversals for a given hypergraph $\mathcal{H}$ with n hyperedges on m vertices.*

Proof Given a hypergraph $\mathcal{H} = \{H_1, \dots, H_n\} \subseteq 2^{\{1,\dots,m\}}$, we define a set of points $\mathcal{A} \subseteq \mathbb{R}^n$ such that the bodies of $\mathcal{A}$ are in one-to-one correspondence with the minimal transversals of $\mathcal{H}$. For $j = 1, \dots, n$, let $\mathbf{e}_j$ be the jth unit vector, containing 1 in position j and 0 elsewhere. For $i = 1, \dots, m$, let $v^i \in \{0, 1\}^n$ be the vector with components $v^i_j = 1$ if $i \in H_j$ and $v^i_j = 0$ if $i \notin H_j$. Now define

$$\mathcal{A} = \{-\mathbf{e}_1, \dots, -\mathbf{e}_n\} \cup \{v^1, \dots, v^m\}.$$

Let $X \in \mathcal{B}(\mathcal{A})$ be a body. If, for some $j \in \{1, \dots, n\}$, $-\mathbf{e}_j \notin X$, then $X \notin \mathcal{B}(\mathcal{A})$, because $X \subseteq \mathcal{A} \setminus \{-\mathbf{e}_j\} \subseteq \{x \in \mathbb{R}^n \mid x_j \geq 0\}$, and hence the convex hull of $\mathcal{A} \setminus \{-\mathbf{e}_j\}$ does not contain the origin as an interior point. We conclude therefore that X must contain the points $-\mathbf{e}_1, \dots, -\mathbf{e}_n$. Now it is easy to see that the set $X' = X \cap \{v^1, \dots, v^m\}$ is a minimal subset of points for which there exists, for each

$j = 1, \ldots, n$, a point $v \in X'$ with $v_j = 1$, i.e. X' is a minimal transversal of $\mathcal{H}$. Conversely, let X be a minimal transversal of $\mathcal{H}$. Then X is a minimal set with the property that $\sum_{i \in X} v^i = y$, for some vector $y > 0$, and consequently the set of points $\{v^i \ : \ i \in X\} \cup \{-\mathbf{e}_1, \ldots, -\mathbf{e}_n\}$ forms a body. $\square$

It should be mentioned that the best currently known algorithm for the hypergraph transversal problem runs in incremental *quasi-polynomial* time (see [10]). We also mention that the problem of generating simplices for a given set of points $\mathcal{A} \subseteq \mathbb{R}^n$ is equivalent with the well-known open problem of listing the vertices of a polytope given by its linear description:

Vertex Enumeration: *Given an $m \times n$ real matrix $A \in \mathbb{R}^{m \times n}$ and an n-dimensional vector $b \in \mathbb{R}^n$ such that the polyhedron $P = \{x \in \mathbb{R}^n \mid Ax = b, \ \ x \geq 0\}$ is bounded, enumerate all vertices of P.*

If the polyhedron P is bounded, i.e. if it is a polytope, then the vertices of P are in one-to-one correspondence with the simplices of the point set $\mathcal{A}$ whose elements are the columns of the augmented matrix $[A \mid -b]$. The complexity status of the vertex enumeration problem for polytopes and the transversal problem of enumerating anti-simplices, currently remain open. For the special case of points $A \subseteq \mathbb{R}^n$ in general position, we have $\mathcal{B}(\mathcal{A}) = \mathcal{S}(\mathcal{A})$, and consequently the problem of enumerating bodies of A turns into the problem of enumerating vertices of the polytope $\{x \in \mathbb{R}^n \mid Ax = 0, \mathbf{e}x = 1, x \geq \mathbf{0}\}$, each vertex of which is non-degenerate and has exactly $n + 1$ positive components. For such kinds of *simple* polytopes, there exist algorithms that generate all vertices with polynomial delay (see [2]). It is also worth mentioning that, as pointed out by Kelmans and Rubinov, for points in general position, the number of anti-bodies of $\mathcal{A}$ is bounded by a polynomial in the number of bodies and n:

$$|\mathcal{B}(\mathcal{A})^{dc}| \leq (n + 1)|\mathcal{B}(\mathcal{A})|, \tag{4}$$

see [7]. A polynomial inequality, similar to (4), for points not necessarily in general position, would imply that bodies, for any set of points $\mathcal{A} \subseteq \mathbb{R}^n$, could be generated in quasi-polynomial time. This follows from the fact that under the assumption that (4) holds, the problem of incrementally generating bodies reduces in polynomial time to the hypergraph transversal problem (see e.g., [5]).

However, a polynomial bound similar to (4) does not hold in general as illustrated by the following example. Let $G = (V, E)$ be a directed graph on $k + 2$ vertices consisting of two special vertices (s, t), and k parallel directed (s, t)-paths of length 2 each. Then let $G\{s, t\}$ be the graph obtained by the construction described in the end of Sect. 3. It is not difficult to see that, for the set of incidence vectors $\mathcal{A} \subseteq \{-1, 0, 1\}^V$ corresponding to the arc-set of $G\{s, t\}$, the number of bodies of $\mathcal{A}$ is $|\mathcal{B}(\mathcal{A})| = k$, while the number of anti-bodies $|\mathcal{B}(\mathcal{A})^{dc}|$ exceeds 2^k.

Let us finally mention that, although the status of the problem of enumerating all maximal feasible subsystems of (2) is not known in general, the situation changes if we fix some set of inequalities, and ask for enumerating all its maximal extensions to a feasible subsystem. In fact, such a problem turns out to be NP-hard, even if we only fix non-negativity constraints.

Theorem 3 *Let $A \in \mathbb{R}^{m\times n}$ be an $m \times n$ matrix, $b \in \mathbb{R}^m$ be an m-dimensional vector, and assume that the system*

$$Ax \geq b, \quad x \in \mathbb{R}^n \tag{5}$$

has no solution $x \geq \mathbf{0}$. Let $\mathcal{F}$ be the set of maximal subsystems of (5) for which there exists a non-negative solution x. Then given a partial list $\mathcal{X} \subseteq \mathcal{F}$, it is NP-hard to determine if the list is complete, i.e. if $\mathcal{X} = \mathcal{F}$, even if b is a $0, 1$-vector, and entries in A are either, $-1, 1,$ or 0.

Proof We again use a polynomial transformation from the satisfiability problem. Let $\Phi = C_1 \wedge \ldots \wedge C_m$ be a conjunctive normal form of m clauses and $2n$ literals $\{x_1, \overline{x}_1, \ldots, x_n, \overline{x}_n\}$, and assume that for each $i = 1, \ldots, n$, both x_i and $\overline{x}_i$ occur in Φ. We construct a system on $m + 3n + 1$ variables and $\sum_{j=1}^m |C_j| + 3n + 1$ inequalities. For each clause C_j, $j = 1, \ldots, m$, we associate a variable y_j. For $i = 1, \ldots, n$, we associate a variable u_i with the positive literal x_i and a variable v_i with the negative literal $\overline{x}_i$. We also use other $n + 1$ variables, $z_1, \ldots, z_n$ and x. There are four types of constraints:

(C1) $-u_i \geq 0, \; -v_i \geq 0$, for $i = 1, \ldots, n$.
(C2) $x - z_i - u_i - v_i \geq 0$, for $i = 1, \ldots, n$.
(C3) $u_i - y_j - x \geq 0$, for every j and i such that x_i appears in C_j. Similarly, $v_i - y_j - x \geq 0$, for every j and i such that $\overline{x}_i$ appears in C_j.
(C4) $\sum_{j=1}^m y_j + \sum_{i=1}^n (u_i + v_i + z_i) + x \geq 1$.

It is easy to see that the system (C1)-(C4) has no non-negative solution, and furthermore that any of the following subsystems is maximal with the property that it has a non-negative solution:

(A) All constraints but (C4).
(B) All constraints but (C2) for some $i \in \{1, \ldots, n\}$.
(C) All constraints but (C3) for some $j \in \{1, \ldots, m\}$ and all i such that either x_i or $\overline{x}_i$ appears in C_j.

Let us call any subsystem having the form (A), (B), or (C), a *trivial* maximal feasible subsystem (MFS). The number of such subsystems is exactly $m + n + 1$. We now show that there exists a non-trivial MFS if and only if the formula Φ is satisfiable.

Let $\sigma = \{\ell_1, \ell_2, \ldots, \ell_n\}$ be the set of literals assigned the value $True$ in a satisfying truth assignment for Φ. We define a non-trivial MFS X corresponding to σ. Starting from an empty set X, we add (C4) and (C2) for all $i = 1, \ldots, n$ to X. For each $i = 1, \ldots, n$, if $\ell_i = x_i$, we add the inequality $-v_i \geq 0$ to X, otherwise we add the inequality $-u_i \geq 0$ to X. For $j = 1, \ldots, m$, we also add to X all inequalities of type (C3) corresponding to a pair (C_j, ℓ), of a clause C_j and any literal ℓ in σ such that ℓ satisfies C_j. Since σ is satisfying, for each clause C_j there exists at least one such literal ℓ_{i_j}, and this together with the fact that X includes constraints (C4), and (C2) for all i, implies indeed that X is non-trivial. So it remains to verify that X is an MFS for the system. To see this, note first that X has the following non-negative solution: $u_i = 1$ if $x_i \in \sigma$ and 0 otherwise, $v_i = 1$ if $\overline{x}_i \in \sigma$ and 0 otherwise, $x = 1$, $y_j = 0$ for all $j = 1, \ldots, m$, and $z_i = 0$ for all $i = 1, \ldots, n$. Second, let us see that if we add a new inequality to X, the resulting subsystem has no non-negative solution. If we add an inequality of type (C1) for some i, then we fix at value 0 some variable, say u_i, for which the corresponding literal ℓ appears in σ. Since, by our earlier assumption, ℓ appears in some clause C_j, the inequality $u_i - y_j - x \geq 0$ is already in X. But then the latter inequality implies with X that all variables must be 0, in contradiction to (C4). Suppose now that we add an inequality of type (C3) to X, say $u_i - y_j - x \geq 0$, corresponding to some clause C_j and some non-satisfying literal $x_i \notin \sigma$. Since $x_i \notin \sigma$, the inequality $-u_i \geq 0$ is already in X, implying again, with the added inequality that all variables must be 0, contradicting (C4) again. This shows that X is indeed an MFS.

Conversely, assume that X is a non-trivial MFS for the system (C1–C4). Then X must contain (C4), and (C2) for all $i = 1, \dots, n$, since otherwise, X is contained in some trivial MFS of type (A) or (B). Furthermore, the fact that X is not contained in any MFS of type (C) implies that, for each clause $j = 1, \dots, m$, there exists an index i_j such that X contains an inequality of the form $w_{i_j} - y_j - x \geq 0$, where w_{i_j} is either u_{i_j} or v_{i_j}. Now we define a satisfying truth assignment σ to be the set of literals corresponding to the set of variables $\{w_{i_j} \ : \ j = 1, \dots, m\}$. Let us first show that σ is indeed a truth assignment, i.e. no literal ℓ and its complement $\overline{\ell}$ appear together in σ. Assume on the contrary that for some i, both $x_i, \overline{x}_i \in \sigma$. Then there exist two distinct indices j and k for which the two inequalities $u_i - y_j - x \geq 0$ and $v_i - y_k - x \geq 0$ are both contained in X. But then adding these inequalities, we get

$$u_i + v_i - y_j - y_k - 2x \geq 0 \tag{6}$$

as a valid inequality for X. Since X also contains all inequalities of type (C2), we conclude by adding the ith such inequality to (6) that $z_i + y_j + y_k + x \leq 0$, from which we get that every variable must be 0, in contradiction to (C4). The fact that σ satisfies Φ immediately follows from the way we constructed σ. □

To conclude this section, let us consider an infeasible system of linear inequalities

$$Dx \geq f \quad \text{and} \quad D'x \geq f', \tag{7}$$

where the subsystem $D'x \geq f'$ is feasible. Let us observe that the problem of finding all maximal feasible subsystems of (7), which includes the given feasible subsystem $D'x \geq f'$, generalizes the generation of both anti-simplices and simplices. Clearly, the former problem can be written in the form (7) by considering (2) and all maximal extensions of an empty (feasible) subsystem. For the latter problem, note that the vertices of the polytope $\{x \in \mathbb{R}^n \mid Ax = b, \ x \geq 0\}$, where $b \neq 0$, are in one-to-one correspondence with the maximal feasible extensions of the subsystem $Ax = b, \ x \geq 0$ in the infeasible system $Ax = b, \ x \geq 0, \ x \leq 0$. Although the general problem, of generating maximal feasible extensions, is NP-hard as shown in Theorem 3, the special cases of generating simplices and anti-simplices remain open.

Acknowledgements We thank Marc Pfetsch for helpful discussions.

References

1. Abel, U., Bicker, R.: Determination of all cutsets between a node pair in an undirected graph. IEEE Trans. Reliab. **31**, 167–171 (1986)
2. Avis, D., Fukuda, K.: A pivoting algorithm for convex hulls and vertex enumeration of arrangements and polyhedra. In: Symposium on Computational Geometry, North Conway, NH, pp. 98–104 (1991)
3. Bansal, V.K., Misra, K.B., Jain, M.P.: Minimal pathset and minimal cutset using search technique. Microelectron. Reliab. **22**, 1067–1075 (1982)
4. Bollobas, B.: Graph Theory: An Introductory Course. Springer, Berlin (1979)
5. Boros, E., Gurvich, V., Khachiyan, L., Makino, K.: Dual-bounded generating problems: partial and multiple transversals of a hypergraph. SIAM J. Comput. **30**, 2036–2050 (2001)
6. Colburn, C.J.: The Combinatorics of Network Reliability. Oxford University Press, New York (1987)
7. Collado, R., Kelmans, A., Krasner, D.: On convex polytopes in the plane "containing" and "avoiding" zero. DIMACS Technical Report 2002-33, Rutgers University (2002)
8. Curet, N.D., DeVinney, J., Gaston, M.E.: An efficient network flow code for finding all minimum cost $s - t$ cutsets. Comput. Oper. Res. **29**, 205–219 (2002)
9. Eiter, T., Gottlob, G.: Identifying the minimal transversals of a hypergraph and related problems. SIAM J. Comput. **24**, 1278–1304 (1995)
10. Fredman, M.L., Khachiyan, L.: On the complexity of dualization of monotone disjunctive normal forms. J. Algorithms **21**, 618–628 (1996)
11. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-completeness. Freeman, New York (1979)
12. Gusfield, D., Naor, D.: Extracting maximum information about sets of minimum cuts. Algorithmica **10**, 64–89 (1993)
13. Jerrum, M.R., Valiant, L.G., Vazirani, V.V.: Random generation of combinatorial structures from a uniform distribution. Theor. Comput. Sci. **43**, 169–188 (1986)
14. Khachiyan, L., Boros, E., Borys, K., Elbassioni, K., Gurvich, V.: Generating all vertices of a polyhedron is hard. In: SODA '06: Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithm, pp. 758–765 (2006)
15. Lawler, E., Lenstra, J.K., Rinnooy Kan, A.H.G.: Generating all maximal independent sets: NP-hardness and polynomial-time algorithms. SIAM J. Comput. **9**, 558–565 (1980)
16. Provan, J.S., Ball, M.O.: Computing network reliability in time polynomial in the number of cuts. Oper. Res. **32**, 516–526 (1984)
17. Provan, J.S., Shier, D.R.: A paradigm for listing (s,t) cuts in graphs. Algorithmica **15**, 351–372 (1996)
18. Read, R.C., Tarjan, R.E.: Bounds on backtrack algorithms for listing cycles, paths, and spanning trees. Networks **5**, 237–252 (1975)
19. Schwikowski, B., Speckenmeyer, E.: On enumerating all minimal solutions of feedback problems. Discrete Appl. Math. **117**, 253–265 (2002)