Enumerating Minimal Dicuts and Strongly Connected Subgraphs and Related Geometric Problems *

E. Boros† K. Elbassioni† V. Gurvich† L. Khachiyan‡

Abstract

We consider the problems of enumerating all minimal strongly connected subgraphs and all minimal dicuts of a given 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-complete 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 $G = (V, E)$ be a *strongly connected* digraph on V with arc set $E \subseteq V \times V$. A minimal strongly connected subgraph of G is a minimal subset of arcs $X \subseteq E$ such that the digraph (V, X) is strongly connected. A *minimal directed cut*, or a *dicut* in G 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 edges whose removal leaves no directed path from s to t. The analogous notions for undirected graphs correspond respectively to spanning trees, minimal cuts, and s, t-cuts.

These notions play an important role in network reliability, where edges or arcs represent communication or transportation links, which may work or fail

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

†RUTCOR, Rutgers University, 640 Bartholomew Road, Piscataway NJ 08854-8003; ({boros,elbassio, gurvich}@rutcor.rutgers.edu).

‡Department of Computer Science, Rutgers University, 110 Frelinghuysen Road, Piscataway, New Jersey 08854-8004; (leonid@cs.rutgers.edu.)

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, 7, 16].

It is quite easy to see that the number of spanning trees, cuts, s, t-paths, s, t-cuts, etc. is, in general, exponential in the size of the graph. For this reason the 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 digraph $G = (V, E)$, let us denote by $\mathcal{F}_{sc}$ the family of minimal strongly connected subgraphs (V, X) of G. Then the *dual* family $\mathcal{F}_{sc}^d$ is the set of minimal dicuts of G:

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

where $X \subseteq E$ is a *transversal of* $\mathcal{F}_{sc}$ if and only if $X \cap Y \neq \emptyset$ for all $Y \in \mathcal{F}_{sc}$.

In this paper, we shall consider the problem of listing incrementally all the elements of the families $\mathcal{F}_{sc}$ and $\mathcal{F}_{sc}^d$. Enumeration algorithms for listing minimal families $\mathcal{F}_\pi$ of subgraphs satisfying a number of monotone properties π 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., [9, 13, 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. 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 in case of many other monotone families, [6, 15]. In this paper we focus on a further case relevant to reliability theory, when this symmetry is broken. Specifically, we show the problem of incrementally generating all minimal dicuts of a given strongly connected directed graph is NP-hard.

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

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 Sections 3 and 4 respectively. In the next section, we discuss some geometric generalizations of these problems related to the enumeration of all minimal subsets of a given set of points $\mathcal{A} \subseteq \mathbb{R}^n$ 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 illustrate that the former problem is NP-hard, while the latter is at least as hard as the well-known hypergraph transversal problem, and discuss the relation of these two problems to the still open problem of vertex enumeration.

2 Some related geometric problems

Let $\mathcal{A} \subseteq \mathbb{R}^n$ be a given subset of vectors in $\mathbb{R}^n$. Fix a point $z \in \mathbb{R}^n$, say $z = \mathbf{0}$, and consider the following four geometric objects: A *simplex* is a minimal subset $X \subseteq \mathcal{A}$ of vectors containing z in its convex hull: $z \in \operatorname{conv}(X)$. An *anti-simplex* is a maximal subset $X \subseteq \mathcal{A}$ of vectors not containing z in its convex hull: $z \notin \operatorname{conv}(X)$. A *body* is a minimal (full-dimensional) subset $X \subseteq \mathcal{A}$ of vectors containing z in the interior of its convex hull: $z \in \operatorname{int}\operatorname{conv}(X)$. An *anti-body* is a maximal subset $X \subseteq \mathcal{A}$ of vectors not containing z in the interior of its convex hull: $z \notin \operatorname{int}\operatorname{conv}(X)$. Equivalently, a simplex (body) is a minimal collection of vectors not contained in an *open* (*closed*) half-space through z. An anti-simplex (anti-body) is a maximal collection of vectors contained in an open (closed) half space. It is known that $|X| \leq n+1$ for any simplex $X \subseteq \mathcal{A}$ and that $n+1 \leq |X| \leq 2n$ 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 the complementary families $\mathcal{S}(\mathcal{A})^{dc}$ and $\mathcal{B}(\mathcal{A})^{dc}$ of their duals (where for a family $\mathcal{F} \subseteq 2^{\mathcal{A}}$, the complementary family $\mathcal{F}^c = \{\mathcal{A} \setminus X \mid X \in \mathcal{F}\}$) 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:

$$Ax \geq \mathbf{e}, \quad x \in \mathbb{R}^n \tag{1}$$

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:

$$Ax \geq \mathbf{0}, \quad x \neq \mathbf{0}. \tag{2}$$

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 (1) 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 vectors not containing any simplex. Thus, the anti-simplices of $\mathcal{A}$ correspond to the complements of the *minimal feedback arc sets* (i.e. minimal sets of arcs whose removal breaks every directed circuit in G).

Now, let us consider bodies and anti-bodies of $\mathcal{A}$ (associated with the graph G). Fix a vertex $v \in V$ and consider the system of inequalities (2) 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 (2) (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 is 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 vectors $\mathcal{A} \subseteq \mathbb{R}^n$, and a partial list $\mathcal{X} \subseteq \mathcal{B}(\mathcal{A})^{dc}$ of anti-bodies of S, 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 (2), and a partial list of maximal feasible subsystems of (2), 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* [10], 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 $A \subseteq \mathbb{R}^n$, 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 $\mathcal{H} = \{H_1, \ldots, H_n\} \subseteq 2^{\{1,\ldots,m\}}$, 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 $j = 1, \ldots, n$, let $\mathbf{e}_j$ be the jth unit vector, containing 1 in position j and 0 elsewhere. For $i = 1, \ldots, m$, 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 $\mathcal{A} = \{-\mathbf{e}_1, \ldots, -\mathbf{e}_m\} \bigcup \{v^1, \ldots, v^m\}$. Let $X \in \mathcal{B}(\mathcal{A})$ be a body. If, for some $j \in \{1, \ldots, n\}$, $-\mathbf{e}_j \notin X$, 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 vectors $-\mathbf{e}_1, \ldots, -\mathbf{e}_n$. Now it is easy to see that the set $X' = X \cap \{v^1, \ldots, v^m\}$ is a minimal subset of vectors for which there exists, for each $j = 1, \ldots, n$, a vector $v \in X'$ 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 vectors $\{v^i \ : \ i \in X\} \cup \{-\mathbf{e}_1, \ldots, -\mathbf{e}_n\}$ forms a body. □

It should be mentioned that the best currently known algorithm for the hypergraph transversal problem runs in incremental *quasi-polynomial* time (see [11]). 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, and the transversal problem of enumerating anti-simplices, currently remains open. For the special case of vectors $A \subseteq \mathbb{R}^n$ 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 vectors in general position, the number of anti-bodies of $\mathcal{A}$ is bounded by a polynomial in the number of bodies and n:

$$|\mathcal{B}(\mathcal{A})^{dc}| \leq (n + 1)|\mathcal{B}(\mathcal{A})|, \tag{3}$$

see [8]. A polynomial inequality, similar to (3), for vectors 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 (see, e.g., [6]). This follows from the fact that under the assumption that (3) holds, the problem of incrementally generating bodies reduces in polynomial time to the hypergraph transversal problem.

However, a polynomial bound similar to (3) 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)-path of length 2 each. 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 $(t, v'), (v', v), (v, v''), (v'', s)$. Then 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 antibodies $|\mathcal{B}(\mathcal{A})^{dc}|$ exceeds 2^k.

Let us finally mention that, although the status of the problem of enumerating all maximal feasible subsystems of (1) 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*

$$Ax \geq b, \quad x \in \mathbb{R}^n \tag{4}$$

has no solution $x \geq \mathbf{0}$. Let $\mathcal{F}$ be the set of maximal subsystems of (4) 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 $b = (0, 0, \ldots, 0, 1)$, and each entry in A is either, $-1, 1$, or 0.

The proof of this theorem is omitted due to lack of space, but can be found in [5].

3 Proof of Theorem 1

Let us state first the following easy but 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 $S, T, R \subseteq V$ such that*

(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, and*

(D3) *if $(a, b) \in E$ and $a \in S$ ($b \in T$) then $b \notin R$ ($a \notin R$), and*

(D4) *$X = \{(a, b) \in E \ : \ a \in S, \ b \in T\}$ is the set of all arcs from S to T.*

To prove the theorem, we 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 $\{x_1, \overline{x}_1, \ldots, x_n, \overline{x}_n\}$. 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.

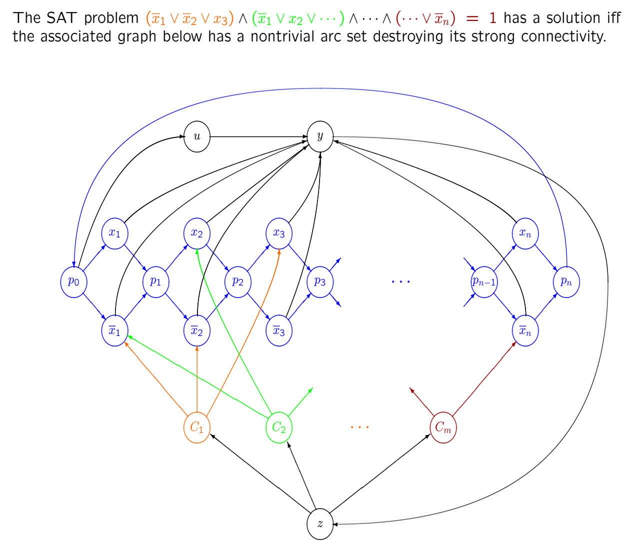

Figure 1: An example for the NP-hard reduction.

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 $\sigma = (\ell_1, \ell_2, \ldots, \ell_n)$ 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:

$$T = \{p_0, \overline{\ell}_1, p_1, \overline{\ell}_2, \ldots, p_{n-1}, \overline{\ell}_n, p_n\}, \ \ S = \{z, C_1, \ldots, C_m, \ell_1, \ldots, \ell_n, y\}, \ \ R = \{u\}. \tag{5}$$

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 $j = 1, \ldots, m$ 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, $\{y, z, C_1, \ldots, C_m\} \subseteq S$ is implied by (S2), (D3) and (S0) (First, $z \notin R$ follows directly from (D3). Second $z \notin T$ follows from (S0) and (S2). Then (S0) implies that $C_j \in S$ for $j = 1, \ldots, m$).

(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 $i = 1, \ldots, n$ such that $\{p_0, \ell_1, p_1, \ell_2, \ldots, p_{n-1}, \ell_n, p_n\} \subseteq T$.

(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 $j = 1, \ldots, m$. Furthermore, we are guaranteed by (S6) that $\ell_{i_j} \neq \overline{\ell}_{i_k}$ for all $k, j = 1, \ldots, m$.

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 $j = 1, \ldots, m$, yields a satisfying truth assignment for Φ. □

4 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 this step that we have fixed some order on the arcs of G).

Theorem 2 readily follows from the following lemma.

Lemma 2

(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 smaller difference $|S \setminus S'|$. This shows 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 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 $B(b) \cup B(z^1) \cup \ldots \cup B(z^{i-1})$ 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). □

As mentioned earlier, by performing a traversal 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 form Lemma 2 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 conclude 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 by the construction described in the end of Section 2. 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)-path for G[1]. Now the claim follows by using the amplification technique of [14], 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 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

[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 [12]

$\Phi(x_1, \overline{x}_1, \ldots, x_n, \overline{x}_n)$ 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 $2n+2s$ 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.

References

[1] U. Abel and R. Bicker, Determination of all cutsets between a node pair in an undirected graph, *IEEE Transactions on Reliability* **31** (1986) pp. 167-171.

[2] D. Avis, K. Fukuda, A Pivoting algorithm for convex hulls and vertex enumeration of arrangements and polyhedra, *Symp. Comput. Geom.* 1991, pp. 98-104.

[3] V.K. Bansal, K.B. Misra and M.P. Jain, Minimal pathset and minimal cutset using search technique, *Microelectr. Reliability* **22** (1982) pp. 1067-1075.

[4] B. Bollobas, Graph Theory: An introductory course, Springer Verlag, 1979.

[5] E. Boros, K. Elbassioni, V. Gurvich and L. Khachiyan, On enumerating minimal dicuts and strongly connected subgraphs, DIMACS Technical Report 2003-35, Rutgers University, **http://dimacs.rutgers.edu/TechnicalReports/2003.html**.

[6] E. Boros, V. Gurvich, L. Khachiyan and K. Makino, Dual-bounded generating problems: partial and multiple transversals of a hypergraph, *SIAM J. Comput.*, **30** (2001), pp. 2036–2050.

[7] C. J. Colburn, The Combinatorics of network reliability. Oxford Univ. Press, 1987.

[8] R. Collado, A. Kelmans and D. Krasner, On convex polytopes in the plane “containing” and “avoiding” zero, DIMACS Technical Report 2002-33, Rutgers Univ. **http://dimacs.rutgers.edu/TechnicalReports/2002.html**.

[9] N.D. Curet, J. DeVinney, and M.E. Gaston, An efficient network flow code for finding all minimum cost s-t cutsets, *Comp. and Oper. Res.* **29** (2002), pp. 205-219.

[10] T. Eiter and G. Gottlob, Identifying the minimal transversals of a hypergraph and related problems, *SIAM J. Comput.*, 24 (1995), pp. 1278-1304.

[11] M. L. Fredman and L. Khachiyan (1996), On the complexity of dualization of monotone disjunctive normal forms, *J. Algorithms*, **21**, pp. 618–628.

[12] M. R. Garey and D. S. Johnson, Computers and intractability: A guide to the theory of NP-completeness, W.H. Freeman and Co., 1979.

[13] D. Gusfield and D. Naor, Extracting maximum information about sets of minimum cuts, *Algorithmica* **10** (1993) pp. 64-89.

[14] M.R. Jerrum, L.G. Valiant and V.V. Vazirani, Random generation of combinatorial structures from a uniform distribution, *Theor. Comp. Sc.*, **43** (1986) pp. 169–188.

[15] E. Lawler, J. K. Lenstra and A. H. G. Rinnooy Kan, Generating all maximal independent sets: NP-hardness and polynomial-time algorithms, *SIAM Journal on Computing*, 9 (1980) pp. 558-565.

[16] J. S. Provan and M. O. Ball, Computing network reliability in time polynomial in the number of cuts, *Operations Research* **32** (1984) pp. 516-526.

[17] J. S. Provan and D. R. Shier, A paradigm for listing (s, t) cuts in graphs, *Algorithmica* **15**, (1996), pp. 351–372.

[18] R. C. Read and R. E. Tarjan, Bounds on backtrack algorithms for listing cycles, paths, and spanning trees, *Networks*, 5 (1975) pp. 237-252.

[19] B. Schwikowski, E. Speckenmeyer, On enumerating all minimal solutions of feedback problems, *Discrete Applied Mathematics*, **117** (2002), pp. 253–265.