Enumerating k-arc-connected orientations

Sarah Blind[*]   Kolja Knauer[†‡]   Petru Valicov[†§]

July 28, 2020

Abstract

We study the problem of enumerating the k-arc-connected orientations of a graph G, i.e., generating each exactly once. A first algorithm using submodular flow optimization is easy to state, but intricate to implement. In a second approach we present a simple algorithm with $O(knm^2)$ time delay and amortized time $O(m^2)$, which improves over the analysis of the submodular flow algorithm. As ingredients, we obtain enumeration algorithms for the α-orientations of a graph G in $O(m^2)$ time delay and for the outdegree sequences attained by k-arc-connected orientations of G in $O(knm^2)$ time delay.

1 Introduction

In an enumeration problem one receives a typically small input and wants to output every element of a typically large resulting set exactly once[1]. Since the runtime of an enumeration algorithm depends on the output, usually one measures how much it exceeds the size of the output in terms of the size of the input. More precisely, one wants to control the *delay*, i.e., the maximum time between two consecutive outputs (including the time before the first and after the last output) in terms of the input or at least the average over these, called the *amortized time*. Clearly, the delay is an upper bound for the amortized time.

An instance that illustrates this challenge perfectly is given an undirected (not necessarily simple) graph G, enumerate all its orientations with a given property. Many types of orientations have been studied with respect to their enumeration complexity, see [7] for an overview. Here we give just some examples, before describing the set-up of this paper.

The set of all orientations of $G = (V, E)$ can be identified with the set of vectors $\{0, 1\}^m$ where $m = |E|$. Thus, enumerating all orientations of G using a Gray code can be done with constant delay, see [14]. It gets more interesting when enumerating *acyclic* orientations, i.e. those that have no directed cycles. In [40] an algorithm for enumerating all acyclic orientations with $O(n^3)$ time delay but linear amortized time $O(n)$ was given, where $n = |V|$. In [3] the delay was reduced to $O(mn)$ with an increase in amortized time to $O(m + n)$. Another improvement was obtained in [8], where an algorithm of delay and amortized time $O(m)$ is given.

Strongly connected orientations are those such that for any two vertices $u, v \in V$ there is a directed path from u to v. In [9] an enumeration algorithm of strongly connected orientations with $O(m)$ time delay is given. The main objective of the present paper is to enumerate a parameterized version of strong orientations, generalizing the above. Namely, we enumerate the *k-arc-connected orientations* of G, i.e., those where at least k arcs have to be removed in order to destroy the strong connectivity[2]. Note that an orientation is strongly connected if and only if it is 1-connected. See Figure 1 for examples.

---

[*]Université de Lorraine, LGIPM, F-57000 Metz, France

Email: `sarah.blind@univ-lorraine.fr`

[†]Departament de Matemàtiques i Informàtica, Universitat de Barcelona (UB), Barcelona, Spain

[‡]Aix-Marseille Univ, Université de Toulon, CNRS, LIS, Marseille, France

Email: `{kolja.knauer,petru.valicov}@lis-lab.fr`

[§]LIRMM, CNRS, Université de Montpellier, Montpellier, France.

[1]We use the term *enumeration* instead of the sometimes used terms *generation* or *listing*.

[2]When there is no ambiguity, we will abbreviate k-arc-connected by saying *k-connected*.

Figure 1: Three strongly connected orientations. Only those in the middle and on the right side are 2-arc-connected.

The concept of k-connectivity is a classic in the theory of directed graphs. Let us review some of the fundamental results of the area. A theorem of Nash-Williams [34] asserts that G admits a k-connected orientation if and only if G is $2k$-edge connected. A result of Lovász [32] yields an easy algorithm to find a k-connected orientation of G if there is one, that runs in $O(n^6)$. More involved techniques have been developed to improve this runtime, see [21, 31]. Important contributions to the theory of k-connected orientations come from Frank. In particular, he showed that k-connected orientations are an instance of the theory of submodular flows, see [18]. Finding a submodular flow (and hence a k-connected orientation) can be done in polynomial time. There is a considerable amount of literature proposing different algorithmic solutions for this task on simple graphs, multigraphs, and mixed graphs, see e.g. [18, 20, 22, 28]. The existence of these polynomial algorithms is the central fact behind a first enumeration algorithm that we propose here. Another fundamental result of Frank is that any two k-connected orientations of G can be transformed into each other by reversing directed paths and directed cycles [19]. Disassembling this result lies at the core of our second algorithm.

All our algorithms are *backtrack searches*, i.e., we build a tree whose vertices are partial solutions and the leaves are full solutions. We explore the tree on the fly by a depth-first search in order to reach all the leaves. The delay of the algorithm depends on the depth of the tree and the time spent at each node.

The first algorithm can be described as follows. Using a min-cost k-connected orientation algorithm due to [28] and following ideas of [2] we design an algorithm that decides in $O(n^3k^3 + kn^2m)$ if a given mixed graph $G = (V, E \cup A)$ can be extended to an orientation that is k-connected. This yields an enumeration algorithm for k-connected orientations with $O(m(n^3k^3 + kn^2m))$ time delay. The idea is to simply start with $G = (V, E)$ as the root vertex of the backtrack search tree and at each node create two children by picking an edge, fixing it to be oriented in one way or the other, and verifying if a k-connected extension exists. Since this tree has depth in $O(m)$ the previously stated runtime follows. Note that any other algorithm finding a k-connected extension of a partial orientation could be used in this approach, see for example [18, 20, 22]. See Section 3 for a detailed description.

The main part of the paper presents an alternative enumeration algorithm. Its main advantages are its simplicity and an improvement of the runtime. Moreover, it consists of two parts enumerating objects of independent interest. The idea is based on Frank's result [19] that any two k-connected orientations can be transformed into each other by reversals of directed cycles and directed paths. Our algorithm splits this into two parts: one reversing cycles, the other paths.

By reversing directed cycles, we obtain an algorithm that enumerates all orientations of G with a prescribed outdegree sequence also known as α-*orientations*. See for instance Figure 1, where the orientations on the middle and right have th same outdegrees and can be transformed into one another by revering a directed triangle. These orientations are of current interest with respect to computational properties (see e.g. [1]) and model many combinatorial objects such as domino and lozenge tilings of a plane region [39, 43], spanning trees of a planar graph [23], perfect matchings (and d-factors) of a bipartite graph [15, 30, 37], Schnyder woods of a planar triangulation [5], Eulerian orientations

of a graph [15], and contact representations of planar graphs with homothetic triangles, rectangles, and d-gons [16, 17, 26]. The set of α-orientations of a planar graph can be endowed with a natural distributive lattice structure [15] and therefore, in this case the enumeration can be done in linear amortized time [27]. It is a famous question whether the enumeration of the elements of a distributive lattice can be done in constant amortized time [38]. This is open even when restricted to distributive lattices coming from the α-orientations of a planar graph. Our enumeration algorithm for general α-orientations runs in $O(m^2)$ time delay and is explained in Section 4.

Reversing directed paths, yields an enumeration algorithm for all outdegree sequences that are attained by k-connected orientations of G. This algorithm runs in $O(m^2kn)$ time delay and is explained in Section 5.

In Section 6 we combine these algorithms, leading to an enumeration algorithm for k-connected orientations with time delay $O(m^2kn)$. The split into two algorithms leads to a structured traversal of the solution space which allows to prove an amortized runtime in $O(m^2)$.

We close the paper in Section 7 with some open questions.

2   Preliminaries

In this section we introduce some digraph basics. All graphs we consider in this paper are loopless multigraphs and consequently their orientations also may have multiple parallel or anti-parallel arcs. We will also consider *mixed multigraphs*. These are of the form $G = (V, E \cup A)$, where E is a multiset of undirected edges and A is a multiset of directed arcs. Analogously to undirected graphs, an *orientation* of a mixed graph consists in fixing a direction for each of its undirected edges.

The digraph obtained from $D = (V, A)$ by reversing a set of arcs $B \subseteq A$ is denoted by D^B. If $A = B$ we write D^- instead of D^B. Given an arc $a = (u, v)$, we denote by $a^- = (v, u)$ the reversed arc. Similarly, the reversed arc set of B is denoted by $B^- = \{a^- \mid a \in B\}$. Given a digraph $D = (V, A)$ and a subset $X \subseteq V$ we will denote its *outdegree* $\delta_D^+(X) = |\{a = (u, v) \in A \mid u \in X \not\ni v\}|$. In the case where $X = \{v\}$ consists of a single element we just write $\delta_D^+(v)$.

Note that by the definition of k-connectivity we have that D is k-connected if and only if $\delta_D^+(X) \geq k$ for every $\emptyset \neq X \subsetneq V$. There is a directed version of Menger's Theorem that provides another characterization of k-connectivity, see [33]. For its statement, we denote by $\lambda(u, v)$ the maximum number of pairwise arc-disjoint directed paths from u to v.

Theorem 1 (Local Menger Theorem for digraphs)**.** *Let $D = (V, A)$ be a digraph and $u, v \in V$. We have $\lambda(u, v) = \min\{\delta_D^+(X) \mid u \in X \not\ni v\}$.*

Thus, as a consequence of Theorem 1, D is k-connected if and only $\lambda(u, v) \geq k$ for all $u, v \in V$.

3   A first enumeration algorithm

We present a more detailed description of the first enumeration algorithm mentioned in the introduction. In [2] the following result is shown based on a min-cost k-connected orientation algorithm due to [28]:

Theorem 2 ([2, 28])**.** *It can be decided in time $O(k^3n^3 + kn^2m)$, whether a mixed graph admits a k-connected orientation.*

By Theorem 2, one can easily design a backtrack search algorithm for enumerating all k-connected orientations with polynomial delay (see Algorithm 1). At each step it takes a new edge and checks for both its orientations if there exists a k-connected orientation of the graph respecting the so far fixed orientation.

Algorithm 1: enumeration of k-connected orientations via min-cost orientations

---

Input: A graph $G = (V, E)$ and an integer $k \in \mathbb{N}$

Output: The k-connected orientations of G

1 begin

2   Fix any linear ordering on E;

3   Enumerate$(G = (V, E, \emptyset))$;

4 end

5 Function Enumerate$(G = (V, E, F))$:

6   **if** $E \neq \emptyset$ **then**

7   Take the smallest $e = \{u, v\} \in E$;

8   **if** $G' := (V, E \setminus \{e\}, F \cup \{(u, v)\})$ *admits a k-connected orientation* **then**

9   Enumerate(G');

10   **if** $G' := (V, E \setminus \{e\}, F \cup \{(v, u)\})$ *admits a k-connected orientation* **then**

11   Enumerate(G');

12   **else**

13   Output G;

14 end

This algorithm takes as input an undirected graph $G = (V, E)$. Clearly, all k-connected orientations are produced and since each node built by this algorithm gives rise to disjoint branches it does not repeat the same solution twice. The depth of the binary execution tree is m and at each node we check the orientability of a mixed graph which is solvable in $O(k^3 n^3 + k n^2 m)$ by Theorem 2. We conclude

Proposition 3. *Let G be a graph and $k \in \mathbb{N}$. Algorithm 1 enumerates all k-connected orientations of G with delay and amortized time in $O(m(k^3 n^3 + k n^2 m))$.*

A downside of Algorithm 1 is that Theorem 2 is based on a submodular flow algorithm and its implementation is intricate. Moreover, the literature on submodular flows and the available runtime analyses in different settings such as simple and multigraphs are hard to extract. A more detailed discussion of these approaches is discussed in [4].

In the following we will give two simple algorithms of independent interest that combined yield an alternative solution for the enumeration of k-connected orientations. The resulting algorithm improves over the delay and amortized runtime of Algorithm 1.

4 Orientations with prescribed outdegree sequence

Let $G = (V, E)$ be a graph and $\alpha : V \to \mathbb{N}$. We say that an orientation D of G is an α-*orientation* if $\delta_D^+(v) = \alpha(v)$ for all $v \in V$. We will denote by $\mathcal{O}_\alpha(G)$ the set of all α-orientations of G. In the present section we discuss the enumeration of the set $\mathcal{O}_\alpha(G)$. Before presenting the algorithm let us give a folklore result about α-orientations:

Lemma 4. *Let G be a graph and D and D' be two orientations of G. We have $\delta_D^+ = \delta_{D'}^+$ if and only if orientation D' can be obtained from D by reversing a set of arc-disjoint directed cycles.*

Proof. Reversing the direction of a directed cycle does not change the outdegree function. Therefore, if D' is obtained from D by reversing a set of arc-disjoint directed cycles, then $\delta_D^+ = \delta_{D'}^+$.

Let conversely $D = (V, A)$ and $D' = (V, A')$ be orientations of G such that $\delta_D^+ = \delta_{D'}^+$. Consider the arcs $A \setminus A'$ of D whose direction differ in D'. Observe that in the directed subgraph $D \setminus A'$ formed by these arcs, the indegree and the outdegree coincide at each node. Thus, $D \setminus A'$ is Eulerian and can be decomposed into an arc-disjoint union of directed cycles. This concludes the proof.   □   □

Our enumeration algorithm for $\mathcal{O}_\alpha(G)$ takes a digraph $D = (V, A) \in \mathcal{O}_\alpha(G)$ and a set of *fixed* arcs $F \subseteq A$ as input. Initially D is an arbitrary α-orientation and $F = \emptyset$. The algorithm recursively

constructs all orientations of $\mathcal{O}_\alpha(G)$ such that at each call all possible α-orientations contain the current F. At each recursive call, an arc $a \notin F$ will be fixed with the two possible cases to consider: The branch with $a = (u, v)$ is known to be non-empty since D is a possible extension, thus we recurse with D and $F \cup a$. For the branch with $a^- = (v, u)$ we check if there is a directed path P in $D \setminus F$ from v to u. If this is the case, then the directed cycle (P, a) is reversed and we recurse with the new α-orientation $D^{P \cup \{a\}}$ and $F \cup \{a^-\}$.

The set of fixed arcs is extended along an arbitrary linear ordering of the arcs. Whenever all arcs are fixed we output the current orientation. We give a presentation of our algorithm in pseudo-code (see Algorithm 2).

Algorithm 2: Backtrack search for α-*orientations*

    Input: A graph $G = (V, E)$ and $\alpha : V \to \mathbb{N}$
    Output: All elements of $\mathcal{O}_\alpha(G)$

1 **begin**
2   **if** *there exists* $D = (V, A) \in \mathcal{O}_\alpha(G)$ **then**
3   Fix an arbitrary linear order on A;
4   EnOPODS(D, $\emptyset$);
5 **end**

6 **Function** EnOPODS(D, F):
7   **if** $F \neq A$ **then**
8   Take the smallest $a = (u, v) \in A \setminus F$;
9   EnOPODS(D, $F \cup \{a\}$);
10   **if** $D \setminus F$ *has a directed path P from v to u* **then**
11   EnOPODS($D^{P \cup \{a\}}$, $F \cup \{a^-\}$);
12   **else**
13   Output D;
14 **end**

Algorithm 2 takes an undirected graph $G = (V, E)$ as input, but after finding one initial α-orientation D, it just consists of recursive calls of the function EnOPODS(D, F).

Lemma 5. *Let $D \in \mathcal{O}_\alpha(G)$ and $F \subseteq A$. The function EnOPODS(D, F) generates each orientation in $\mathcal{O}_\alpha(G)$ that coincides with D on F exactly once and runs with $O(m^2)$ time delay.*

Proof. We prove that each of the claimed α-orientations is generated exactly once by induction on $|A \setminus F|$. If $|A \setminus F| = 0$, then EnOPODS(D, F) = EnOPODS(D, A) = $\{D\}$ and we are done. Let now $|A \setminus F| > 0$ and $a = (u, v) \in A \setminus F$. By induction hypothesis EnOPODS $(D, F \cup \{a\})$ generates each α-orientation that coincides with D on $F \cup \{a\}$ exactly once and EnOPODS($D^{P \cup \{a\}}, F \cup \{a^-\}$) generates each α-orientation that coincides with $D^{P \cup \{a\}}$ on $F \cup \{a^-\}$ exactly once. Clearly, both sets are disjoint since they differ with respect to the orientation of a. Thus, no repetitions are produced. Since (P, a) is a directed cycle by Lemma 4 we have $D^{P \cup \{a\}} \in \mathcal{O}_\alpha(G)$. Since $P \cap F = \emptyset$ we have that $D^{P \cup \{a\}}$ coincides with D on F.

To see that we do not miss any orientation, we prove that if there is no directed path P from u to v in $D \setminus F$, then there exists no α-orientation fixing F and reversing a. By contraposition, suppose that $D' = (V, A')$ is an α-orientation that coincides with D on F but differs on a. Then by Lemma 4, there is a set of arc-disjoint directed cycles in D' whose union is $A' \setminus A$. Since both digraphs coincide on F, these cycles are disjoint from F. Since both digraphs differ on a, one of the directed cycles C contains a^- in D'. Thus, the path $P = (C \setminus \{a^-\})^-$ is a directed path in $D \setminus F$ from v to u.

For the complexity, note that in each recursion step the algorithm checks the presence of a directed path, which can be done by a single BFS from the source vertex u towards the target v. In general the complexity of a BFS algorithm is $O(m + n)$, however in our case the BFS tree will be constructed only on the strongly connected component of D containing u and thus the complexity is in $O(m)$. The depth of our recursion tree is bounded by m. Thus, the total time delay is bounded by $O(m^2)$. $\square$ $\square$

Observe that Algorithm 2 has to use a separate method for finding a first element $D \in \mathcal{O}_\alpha(G)$. It is well-known that this problem can be reduced to a flow-problem, see e.g. [15]. Therefore, this preprocessing step can be done in $O(mn)$ time, see [36]. With Lemma 5 we obtain:

Theorem 6. *Let G be a graph and $\alpha : V \to \mathbb{N}$. Algorithm 2 enumerates $\mathcal{O}_\alpha(G)$ with time delay in $O(m^2)$.*

5 Outdegree sequences

In this section we will present an algorithm to enumerate the possible outdegree sequences among the k-connected orientations of a graph G.

An easy consequence of Lemma 4 is the following that can also be found in [19].

Lemma 7. *If $D, D' \in \mathcal{O}_\alpha(G)$, then D is k-connected if and only if D' is k-connected.*

Proof. Let $D, D' \in \mathcal{O}_\alpha(G)$ and D be k-connected. Thus, for all $\emptyset \neq X \subsetneq V$ we have $\delta_D^+(X) \geq k$. By Lemma 4 since $D' \in \mathcal{O}_\alpha(G)$ it can be obtained from D by reversing a set of disjoint directed cycles. But reversing a directed cycle in a digraph does not change the outdegree of the subsets of vertices of G. Therefore, after reversing the set of directed cycles to obtain D', we have $\delta_{D'}^+(X) \geq k$ for all $\emptyset \neq X \subsetneq V$ and D' is k-connected. $\square$ $\square$

Given $G = (V, E)$, Lemma 7 allows to define a function $\alpha : V \to \mathbb{N}$ to be *k-connected* if there is some k-connected $D \in \mathcal{O}_\alpha(G)$. In this case we call α a k-connected outdegree sequence. Having in mind that we want to enumerate all k-connected orientations of G and already are able to enumerate $\mathcal{O}_\alpha(G)$ for any given α, we are left with enumerating the k-connected outdegree sequences. Here in order to change the outdegree sequence of D we will reverse a directed path P_{uv} from u to v and thus increase $\delta_D^+(u)$ by one, decrease $\delta_D^+(v)$ by one, and leave the remaining outdegrees unchanged. More generally we have:

Observation 1. *Let D be an orientation of graph G and $X \subseteq V(G)$. If D' is obtained from D by reversing a path from a vertex u to a vertex v, then we have*

1. $\delta_{D'}^+(X) = \delta_D^+(X)$ if $u, v \in X$ or $u, v \notin X$

2. $\delta_{D'}^+(X) = \delta_D^+(X) + 1$ if $u \notin X$ and $v \in X$

3. $\delta_{D'}^+(X) = \delta_D^+(X) - 1$ if $u \in X$ and $v \notin X$

In order to maintain connectivity, we have to reverse paths without decreasing the number of arc-disjoint directed paths between pairs of vertices too much.

Lemma 8. *Let P_{uv} be a directed path from vertex u to vertex v in D. For all vertices u', v', we have $\lambda_{DP_{uv}}(u', v') \geq \min(\lambda_D(u, v) - 1, \lambda_D(u', v'))$. Furthermore, we have $\lambda_{DP_{uv}}(u, v) = \lambda_D(u, v) - 1$.*

Proof. With Menger's Theorem (Theorem 1) and Observation 1, Lemma 8 can be easily proved:

$$
\begin{aligned}
\lambda_{DP_{uv}}(u', v') \quad &= \min\{\delta_{DP_{uv}}^+(X) \mid X \subset V, v' \notin X \ni u'\} \quad &\text{(Theorem 1)}\\
&= \min(\{\delta_D^+(X) - 1 \mid X \subseteq V, v, v' \notin X \ni u', u\} \quad &\text{(Observation 1)}\\
&\cup \{\delta_D^+(X) \mid X \subseteq V, v' \notin X \ni u', u, v \text{ or } v', u, v \notin X \ni u'\}\\
&\cup \{\delta_D^+(X) + 1 \mid X \subseteq V, u, v' \notin X \ni u', v\})\\
&\geq \min\{\lambda_D(u, v) - 1, \lambda_D(u', v')\} \quad &\text{(Theorem 1)}
\end{aligned}
$$

Note that in the case $u' = u$ and $v' = v$ the second and third part of the union in the above equation are empty. We thus get:

$$
\begin{aligned}
\lambda_{DP_{uv}}(u, v) \quad &= \min\{\delta_{DP_{uv}}^+(X) \mid X \subseteq V, u' \in X, v' \notin X\} \quad &\text{(Theorem 1)}\\
&= \min\{\delta_{DP}^+(X) - 1 \mid X \subseteq V, u' \in X, v' \notin X\} \quad &\text{(Observation 1)}\\
&= \lambda_D(u, v) - 1 \quad &\text{(Theorem 1)}
\end{aligned}
$$

$\square$ $\square$

In a k-connected digraph D we call a directed path P *flippable* if D^P is k-connected. Lemma 8 implies that a path from u to v is flippable, if and only if all of them are which is equivalent to $\lambda_D(u,v) > k$. In this case we call the pair (u,v) *flippable*. The following gives an algorithm, that is basically equivalent to the Edmonds-Karp algorithm for maximum flows [12, 13]:

Lemma 9. *Let D be k-connected, it can be decided in time $O(km)$ if (u,v) is flippable.*

Proof. By Lemma 8 a simple algorithm consists in finding a directed path P_{uv} in D, reverse it and iterate in $D^{P_{uv}}$. We have $\lambda_D(u,v) > k$ if and only if this procedure can be applied $k+1$ times. Each execution is a BFS, which yields the claimed runtime. $\qquad\square$ $\qquad\qquad\square$

The following is a slight refinement of a result of Frank [19] and constitutes the last ingredient for our algorithm:

Lemma 10. *Let $G = (V,E)$ be a graph and D, D' be two k-connected orientations of G. For every $v \in V$ with $\delta_D^+(v) < \delta_{D'}^+(v)$, there exists $u \in V$ such that $\delta_D^+(u) > \delta_{D'}^+(u)$ and (u,v) is flippable in D.*

Proof. Let $v \in V$ be with $\delta_D^+(v) < \delta_{D'}^+(v)$. We will prove that there is a $u \in V$ such that $\delta_D^+(u) > \delta_{D'}^+(u)$ and $\delta_D^+(X) > k$ whenever $u \in X$ and $v \notin X$. By Menger's Theorem (Theorem 1) this implies that (u,v) is flippable. Since $\delta_D^+(v) < \delta_{D'}^+(v)$ and because $\sum\limits_{w \in V} \delta_D^+(w) = |E| = \sum\limits_{w \in V} \delta_{D'}^+(w)$, there exists at least one vertex u such that $\delta_D^+(u) > \delta_{D'}^+(u)$. By contradiction suppose that every vertex u with $\delta_D^+(u) > \delta_{D'}^+(u)$, is contained in a set X with $\delta_D^+(X) = k$ and $v \notin X$. Among all these sets for all such u, we consider those that are inclusion maximal and collect them in $\mathcal{X}$.

Let us first prove that the members of $\mathcal{X}$ are pairwise disjoint. Indeed, consider two sets X, X' with $\delta_D^+(X) = \delta_D^+(X') = k$ and $v \notin X \cup X'$ with $X \cap X' \neq \emptyset$. Since $X \cup X'$, $X \cap X'$ are both not empty, $X \cup X' \neq V$ and D is k-connected, we know that $\delta_D^+(X \cup X') \geq k$ and $\delta_D^+(X \cap X') \geq k$. Therefore, we have

$$k + k = \delta_D^+(X) + \delta_D^+(X') \geq \delta_D^+(X \cup X') + \delta_D^+(X \cap X') \geq k + k$$

This gives $\delta_D^+(X \cup X') = k$, i.e., X or X' is not maximal.

Now, let us count the number c of edges of G not contained in any subgraph of G induced by $X \in \mathcal{X}$ and denote $Y = V \setminus \bigcup_{X \in \mathcal{X}} X$. Observe that since the elements of $\mathcal{X}$ are disjoint, in any orientation of G the number c is just the sum of outdegrees of $X \in \mathcal{X}$ plus the sum of outdegrees of vertices of Y. Thus, if we furthermore denote $t = |\mathcal{X}|$ we can do this counting with respect to D and D' and obtain the following contradiction

$$c = \sum_{X \in \mathcal{X}} \delta_D^+(X) + \sum_{w \in Y} \delta_D^+(w) = kt + \sum_{w \in Y} \delta_D^+(w) < \sum_{X \in \mathcal{X}} \delta_{D'}^+(X) + \sum_{w \in Y} \delta_{D'}^+(w) = c$$

More precisely, since D' is k-connected we have that $kt \leq \sum\limits_{X \in \mathcal{X}} \delta_{D'}^+(X)$. Recall that we supposed that if w is such that $\delta_D^+(w) > \delta_{D'}^+(w)$, then $w \notin Y$. Therefore, we have $\delta_D^+(w) \leq \delta_{D'}^+(w)$ for all $w \in Y$. Adding this to the fact that $\delta_D^+(v) < \delta_{D'}^+(v)$ and $v \in Y$, we obtain the strict inequality claimed above. This concludes the proof. $\qquad\square$ $\qquad\qquad\square$

Note that a complete analogue of Lemma 10 holds for the case $\delta_D^+(v) > \delta_{D'}^+(v)$:

Lemma 11. *Let $G = (V,E)$ be a graph and D, D' be k-connected orientations of G. For every $v \in V$ with $\delta_D^+(v) > \delta_{D'}^+(v)$, there exists $u \in V$ such that $\delta_D^+(u) < \delta_{D'}^+(u)$ and (v,u) is flippable in D.*

Proof. Observe that an equivalent statement of the lemma is that for every vertex v such that $\delta_D^-(v) < \delta_{D'}^-(v)$, there exists a vertex u such that $\delta_D^-(u) > \delta_{D'}^-(u)$ and (v,u) is flippable in D. By fully reorienting all the arcs of D and D', we obtain graphs D^- and D'^- and get a further equivalent statement: for every vertex v such that $\delta_{D^-}^+(v) < \delta_{D'^-}^+(v)$, there exists a vertex u such that $\delta_{D^-}^+(u) > \delta_{D'^-}^+(u)$ and (u,v) is flippable in D^-. This is precisely the statement of Lemma 10, so we are done. $\qquad\square$ $\qquad\square$

Now we are ready de describe the general enumeration algorithm for k-connected outdegree sequences:

Algorithm 3: Enumeration of k-connected outdegree sequences

Input: A graph $G = (V, E)$, an integer k
Output: All k-connected outdegree sequences of G

1 begin
2   **if** *there exists a k-connected orientation D of G* **then**
3   Fix an arbitrary linear order on V;
4   EnODS$(D, \emptyset)$;
5 end

6 Function EnODS(D, F):
7   **if** $F \neq V$ **then**
8   take the smallest $v \in V \setminus F$;
9   Reverse$^-(D, F, v)$;
10   Reverse$^+(D, F, v)$;
11   EnODS$(D, F \cup \{v\})$;
12   **else**
13   Output δ_D^+;
14 end

15 Function Reverse$^-(D, F, v)$:
16   **if** *there exists $u \in V \setminus F$ such that (v, u) is flippable* **then**
17   Take a directed path P_{vu} from v to u ;
18   Reverse$^-(D^{P_{vu}}, F, v)$;
19   EnODS$(D^{P_{vu}}, F \cup \{v\})$;
20 end

21 Function Reverse$^+(D, F, v)$:
22   **if** *there exists $u \in V \setminus F$ such that (u, v) is flippable* **then**
23   Take a directed path P_{uv} from u to v ;
24   Reverse$^+(D^{P_{uv}}, F, v)$;
25   EnODS$(D^{P_{uv}}, F \cup \{v\})$;
26 end

Lemma 12. *Let D be a k-connected orientation of $G = (V, E)$ and $F \subseteq V$. The function EnODS(D,F) enumerates the k-connected outdegree sequences coinciding with D on F with time delay in $O(knm^2)$.*

Proof. We show the first part of the lemma by induction on $|V \setminus F|$. If $|V \setminus F| = 0$, then EnODS(D,F) outputs δ_D^+ and the claim holds. Consider now the case $|V \setminus F| > 0$ and let $v \in V \setminus F$ be the next vertex. By induction EnODS$(D,F \cup \{v\})$ generates every k-connected outdegree sequence coinciding with D on $F \cup \{v\}$ exactly once. We have to show that Reverse$^+(D,F,v)$ (resp. Reverse$^-(D,F,v)$) enumerates all k-connected outdegree sequences coinciding with D on F and having outdegree of v larger (resp. smaller) than $\delta_D^+(v)$. Also, we have to show that each of these outdegree sequences will be generated exactly once. Note that this implies that globally each solution is produced exactly once.

Let us prove this for Reverse$^+(D,F,v)$. So let D' be a k-connected orientation of G such that $\delta_D^+(F) \equiv \delta_{D'}^+(F)$ and $\delta_D^+(v) < \delta_{D'}^+(v)$. By Lemma 10 there exists a vertex $u \in V \setminus F$ such that (u, v) is flippable, i.e., for any path P_{uv} the orientation $D^{P_{uv}}$ is k-connected, its outdegree sequence coincides with D on F and $\delta_D^+(v) + 1 = \delta_{D^{P_{uv}}}^+(v)$.

We proceed by induction on $\delta_{D'}^+(v) - \delta_D^+(v)$ to show that $\delta^+(D')$ is enumerated exactly once. So for the base case $\delta_{D'}^+(v) - \delta_D^+(v) = 1$ we have $\delta_{D'}^+(v) = \delta_{D^{P_{uv}}}^+(v)$ and by induction $\delta_{D'}^+$ will be enumerated exactly once by the next call of EnODS$(D^{P_{uv}}, F \cup \{v\})$ and not at all by Reverse$^+(D^{P_{uv}}, F, v)$ since

the latter outputs degree sequences with $\alpha(v) > \delta^+_{D^{P_{uv}}}(v)$. Suppose now that $\delta^+_{D'}(v) - \delta^+_D(v) > 1$. We have $\delta^+_{D'}(v) - \delta^+_{D^{P_{uv}}}(v) < \delta^+_{D'}(v) - \delta^+_D(v)$, so by induction hypothesis $\texttt{Reverse}^+(D^{P_{uv}}, F, v)$ enumerates the outdegree sequence of $\delta^+_{D'}(v)$ exactly once.

The analogue proof works for $\texttt{Reverse}^-$ using Lemma 11.

For the analysis of complexity note that in each call of $\texttt{Reverse}^+$ or $\texttt{Reverse}^-$ for at most n times it has to be checked if a pair (u, v) is flippable. The latter can be done in time $O(km)$ by Lemma 9, finding a directed path from u to v is done in $O(m)$. So a call costs $O(knm)$.

Finally, the depth of the recursion tree is in $O(m)$. To see this compare the $\delta^+_{D'}$ of a leaf orientation with the δ^+_D of orientation D at the root. Between any two calls of $\texttt{EnODS}$, there will be a sequence of at most $\deg(v)$ calls of $\texttt{Reverse}^+$ or $\texttt{Reverse}^-$. This way δ^+_D will be approached to $\delta^+_{D'}$ coordinate by coordinate, where previous coordinates are not affected by modifications on latter coordinates. Thus, there are at most $\sum_{v \in V} \deg(v) = 2m$ calls and we get an overall time delay of $O(knm^2)$. □ □

Note that Algorithm 3 has to use a separate method for finding a first k-connected orientation D of G. This preprocessing step can be done in $O(k^3 n^3 + kn^2 m)$ [28]. On the other hand, recall that in a k-connected orientation we have $kn \leq m$. Therefore, together with Lemma 12 we obtain:

Theorem 13. *Let G be a graph and $k \in \mathbb{N}$. Algorithm 3 enumerates all k-connected outdegree sequences of G in $O(knm^2)$ time delay.*

6 k-connected orientations

Putting the above together we obtain an algorithm to enumerate k-connected orientations.

Algorithm 4: Simple enumeration of k-connected orientations

Input: A graph $G = (V, E)$, an integer k
Output: All k-connected orientations of G

1 begin
2   **if** *there exists a k-connected orientation D of G* **then**
3   Fix an arbitrary linear order on V;
4 $\texttt{EnODS'}(D, \emptyset)$;
5 end

6 Function $\texttt{EnODS'}(D, F)$:
7   **if** $F \neq V$ **then**
8   take the smallest $v \in V \setminus F$;
9 $\texttt{Reverse}^-(D, F, v)$;
10 $\texttt{Reverse}^+(D, F, v)$;
11 $\texttt{EnODS'}(D, F \cup \{v\})$;
12   **else**
13 $\texttt{EnOPODS}(D, \emptyset)$;
14 end

We need the following easy result for analyzing the amortized complexity:

Lemma 14. *Let $G = (V, E)$ be a graph and α be a k-connected out-degree sequence, then $|\mathcal{O}_\alpha(G)| \geq (k-1)n + 2$.*

Proof. Let $D \in \mathcal{O}_\alpha(G)$. We will show that D contains at least $(k-1)n + 1$ directed cycles. Since for each directed cycle C, the orientation D^C is a different element of $\mathcal{O}_\alpha(G)$, we obtain the result.

The *cycle space* of D is the set of vectors in $\mathbb{R}^m$ that are linear combinations of *signed incidence vectors* of cycles of D. Given a cycle C with a direction of traversal its signed incidence vector has an entry for every arc of D, which is 1 if a is traversed forward, -1 if it is traversed backward and 0 if $a \notin C$. It is well known, that the cycle space of a strongly connected D can be generated by linear combinations of the vectors associated to its directed cycles, see e.g. [24]. Moreover, it can be found in

most books on (algebraic) graph theory that the dimension of the cycle space of a weakly connected digraph is $m - n + 1$, see e.g. [25, 29]. In particular, D has at least $m - n + 1$ directed cycles. Since D is k-connected it has at least kn edges. Therefore D contains at least $(k-1)n+1$ directed cycles. □ □

Theorem 15. *Let G be a graph and $k \in \mathbb{N}$. Algorithm 4 enumerates all k-connected orientations of G with $O(knm^2)$ time delay. If $k \geq 2$ the amortized time is in $O(m^2)$.*

Proof. The correctness and the delay follow directly from Theorem 13 and Theorem 6. Let us compute the amortized time complexity as an average over the delays. Let s be the number of solutions, i.e., the total number of k-connected orientations and t be the number of k-connected outdegree sequences of G. Since by Lemma 14 for every k-connected outdegree sequence α there are at least $(k-1)n+2$ orientations, we have that $t \leq \frac{s}{(k-1)n+2}$. Thus there exist constants c and c', such that the overall runtime of our algorithm is bounded by $cknm^2 t + c'm^2 s \leq cknm^2 \frac{s}{(k-1)n} + c'm^2 s = O(m^2)s$, where for the last equality we use $k \geq 2$. Hence the amortized complexity is in $O(m^2)$. □ □

7 Discussion

We have given a simple algorithm for enumerating the k-arc-connected orientations of a graph with a low amortized time complexity. While polynomial delay algorithms were accessible through the theory of submodular flows before, our result is a contribution to the general quest in enumeration to lower polynomial delays and amortized times as much as possible. Other instances of this quest are the different approaches to enumerate acyclic orientations mentioned in the introduction [3, 8, 40] and the major open problem of constant amortized time enumeration of the elements of a distributive lattice [38]. And by [15] our enumeration of α-orientations can be seen as related to the latter problem.

The weakness of our enumeration algorithm for k-arc-connected orientations is that finding the initial solution, i.e., finding a k-connected orientation of a graph, is algorithmically the most complicated part. The best implementations using splitting off techniques are rather simple and lead to computation time of roughly $O(n^5)$ [21, 31], it would be interesting to find simpler methods.

Our original motivation was the hunt for counterexamples to a conjecture of Neumann-Lara [35] stating that the vertex set of every oriented planar graph (that is having no directed cycles of length 2) can be vertex-partitioned into two subsets each inducing an acyclic digraph. It is not hard to see that a minimal counterexample to this conjecture has to be a 2-arc-connected (planar) digraph. However, with a little more work one can actually prove that a minimal counterexample has to be 2-*vertex-connected*, i.e., at least two vertices have to be removed in order to destroy strong connectivity. For illustration, the left and middle orientations in Figure 1 have vertex-connectivity 1 while the one on the right has vertex-connectivity 2. Indeed, using SageMath [41] one can compute that among the 830918 strong orientations of that graph, only 3842 are 2-arc-connected, and 3734 of those are 2-vertex-connected. This also gives an idea, how inefficient it would be to generate all strong orientations and filter with respect to 2-arc-connectivity.

Deciding if a graph admits a k-vertex-connected orientation can be done in polynomial time if and only if $k \leq 2$ (unless P=NP), see [11, 42]. Thus, this enumeration problem would be viable only for $k = 2$. As shown by the middle and right orientation in Figure 1 having the same α does not guarantee the same vertex-connectivity. Thus, our approach does not carry through in this setting. In [2] it is posed as open problem whether it can be decided in polynomial time, if a mixed graph can be oriented such that it is 2-vertex-connected. A positive answer would allow polynomial time enumeration as in Algorithm 1. Note that concerning our initial motivation it would be interesting to answer this question even restricted to planar graphs.

A dual analogue of k-arc-connectivity could be called k-*acyclicity*, where at least k arcs of D have to be contracted in order to destroy its acyclicity. As mentioned in the introduction, acyclic orientations can be enumerated with polynomial delay. On the other hand, it is easy to see that a graph G admits a 2-acyclic orientation if and only if G is the cover graph of a poset. The corresponding recognition problem is NP-complete [6] and the proof can be extended to see that testing whether G admits a k-acyclic orientation is NP-complete for any $k \geq 2$. Using [10, Theorem 13] one can show that

enumerating k-acyclic orientations (as well as k-vertex-connected orientations for $k > 2$) are DelNP-hard under D-reductions and IncNP-hard under I-reduction. Since with an NP-oracle one can decide if a given partial orientation extends to one of these types, a backtrack search algorithm like Algorithm 1 yields that the above problems can be solved with delay a polynomial number of calls of an NP-oracle. Thus these problems are in the class DelNP (and a fortiori IncNP), and therefore are indeed complete in both settings.

A way of relaxing acyclicity is to consider orientations of *digirth* at least d, i.e., all directed cycles are of length at least d. We believe that in this setting polynomial delay enumeration should be possible, maybe even when combined with an arc-connectivity requirement.

References

[1] Oswin Aichholzer, Jean Cardinal, Tony Huynh, Kolja Knauer, Torsten Mütze, Raphael Steiner, and Birgit Vogtenhuber. Flip distances between graph orientations. In Ignasi Sau and Dimitrios M. Thilikos, editors, *Graph-Theoretic Concepts in Computer Science*, pages 120–134, Cham, 2019. Springer International Publishing.

[2] Joergen Bang-Jensen, Jing Huang, and Xuding Zhu. Completing orientations of partially oriented graphs. *Journal of Graph Theory*, 87(3):285–304, 2018.

[3] Valmir C. Barbosa and Jayme L. Szwarcfiter. Generating all the acyclic orientations of an undirected graph. *Information Processing Letters*, 72(1):71 – 74, 1999.

[4] Sarah Blind. *Output-sensitive algorithms for enumeration problems in graphs*. PhD thesis, Université de Lorraine, 2019.

[5] Enno Brehm. 3-orientations and Schnyder-3-tree-decompositions, 2000. Diploma Thesis, FU Berlin.

[6] Graham Brightwell. On the complexity of diagram testing. *Order*, 10(4):297–303, Dec 1993.

[7] Alessio Conte. *Enumeration Algorithms for Real-World Networks: Efficiency and Beyond*. PhD thesis, Università di Pisa, 2018.

[8] Alessio Conte, Roberto Grossi, Andrea Marino, and Romeo Rizzi. Efficient enumeration of graph orientations with sources. *Discrete Appl. Math.*, 246:22–37, 2018.

[9] Alessio Conte, Roberto Grossi, Andrea Marino, Romeo Rizzi, and Luca Versari. *Directing Road Networks by Listing Strong Orientations*, pages 83–95. Springer International Publishing, 2016.

[10] Nadia Creignou, Markus Kröll, Reinhard Pichler, Sebastian Skritek, and Heribert Vollmer. A complexity theory for hard enumeration problems. *Discrete Applied Mathematics*, 268:191 – 209, 2019.

[11] Olivier Durand de Gevigney. On Frank's conjecture on k-connected orientations. *Journal of Combinatorial Theory, Series B*, 141:105 – 114, 2020.

[12] E. A. Dinits. Algorithm for solution of a problem of maximum flow in a network with power estimation. *Sov. Math., Dokl.*, 11:1277–1280, 1970.

[13] Jack Edmonds and Richard M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *J. Assoc. Comput. Mach.*, 19:248–264, 1972.

[14] Gideon Ehrlich. Loopless algorithms for generating permutations, combinations, and other combinatorial configurations. *J. ACM*, 20(3):500–513, 1973.

[15] Stefan Felsner. Lattice structures from planar graphs. *Electron. J. Combin.*, 11(1), 2004. Research Paper 15.

[16] Stefan Felsner. Rectangle and square representations of planar graphs. In *Thirty essays on geometric graph theory*, pages 213–248. Springer, New York, 2013.

[17] Stefan Felsner, Hendrik Schrezenmaier, and Raphael Steiner. Pentagon contact representations. *Electron. J. Combin.*, 25(3), 2018. Paper 3.39, 38.

[18] András Frank. An algorithm for submodular functions on graphs. *Ann. Discrete Math.*, 16:97–120, 1982.

[19] András Frank. A note on k-strongly connected orientations of an undirected graph. *Discrete Math.*, 39(1):103–104, 1982.

[20] Harold N. Gabow. A framework for cost-scaling algorithms for submodular flow problems. In *34th Annual Symposium on Foundations of Computer Science, Palo Alto, USA, November 1993*, pages 449–458, 1993.

[21] Harold N. Gabow. Efficient splitting off algorithms for graphs. In *Proceedings of the Twenty-sixth Annual ACM Symposium on Theory of Computing*, STOC '94, pages 696–705, New York, NY, USA, 1994. ACM.

[22] Harold N. Gabow. Centroids, representations, and submodular flows. *J. Algorithms*, 18(3):586 – 628, 1995.

[23] Patrick M. Gilmer and Richard A. Litherland. The duality conjecture in formal knot theory. *Osaka J. Math.*, 23(1):229–247, 1986.

[24] Petra M. Gleiss, Josef Leydold, and Peter F. Stadler. Circuit bases of strongly connected digraphs. *Discuss. Math., Graph Theory*, 23(2):241–260, 2003.

[25] Chris Godsil and Gordon Royle. *Algebraic graph theory.*, volume 207. New York, NY: Springer, 2001.

[26] Daniel Gonçalves, Benjamin Lévêque, and Alexandre Pinlou. Triangle contact representations and duality. *Discrete Comput. Geom.*, 48(1):239–254, 2012.

[27] Michel Habib, Raoul Medina, Lhouari Nourine, and George Steiner. Efficient algorithms on distributive lattices. *Discrete Applied Mathematics*, 110(2):169 – 187, 2001.

[28] Satoru Iwata and Yusuke Kobayashi. An algorithm for minimum cost arc-connectivity orientations. *Algorithmica*, 56(4):437–447, 2010.

[29] Ulrich Knauer and Kolja Knauer. *Algebraic graph theory. Morphisms, monoids and matrices. 2nd revised and extended edition.*, volume 41. Berlin: De Gruyter, 2nd revised and extended edition edition, 2019.

[30] Peter Che Bor Lam and Heping Zhang. A distributive lattice on the set of perfect matchings of a plane bipartite graph. *Order*, 20(1):13–29, 2003.

[31] Lap Chi Lau and Chun Kong Yung. Efficient edge splitting-off algorithms maintaining all-pairs edge-connectivities. *SIAM J. Comput.*, 42(3):1185–1200, 2013.

[32] László Lovász. *Combinatorial Problems and Exercises.* North-Holland, 1979.

[33] Karl Menger. Zur allgemeinen Kurventheorie. *Fundamenta Mathematicae*, 10(1):96–115, 1927.

[34] Crispin St. John Alvah Nash-Williams. On orientations, connectivity and odd-vertex-pairings in finite graphs. *Canadian Journal of Mathematics*, 12:555–567, 1960.

[35] Victor Neumann-Lara. Vertex colourings in digraphs, Some Problems. Technical report, Waterloo, Canada, 1985.

[36] James B. Orlin. Max flows in $O(nm)$ time, or better. In *Proceedings of the 45th annual ACM symposium on theory of computing, STOC '13. Palo Alto, CA, USA, June 1–4, 2013*, pages 765–774. 2013.

[37] James Propp. Lattice structure for orientations of graphs. ArXiv: math/0209005, September 2002.

[38] Gara Pruesse and Frank Ruskey. Gray codes from antimatroids. *Order*, 10(3):239–252, 1993.

[39] Eric Rémila. The lattice structure of the set of domino tilings of a polygon. *Theoret. Comput. Sci.*, 322(2):409–422, 2004.

[40] Matthew B. Squire. Generating the acyclic orientations of a graph. *J. Algorithms*, 26(2):275 – 290, 1998.

[41] The Sage Developers. SageMath, the Sage Mathematics Software System (Version 9.1).

[42] Carsten Thomassen. Strongly 2-connected orientations of graphs. *Journal of Combinatorial Theory, Series B*, 110:67 – 78, 2015.

[43] William P. Thurston. Conway's tiling groups. *Amer. Math. Monthly*, 97(8):757–773, 1990.