Egerváry Research Group
on Combinatorial Optimization

Technical reportS

TR-2007-07. Published by the Egerváry Research Group, Pázmány P. sétány 1/C, H–1117, Budapest, Hungary. Web site: www.cs.elte.hu/egres. ISSN 1587–4451.

An algorithm to increase the node-connectivity of a digraph by one

András Frank and László A. Végh

June 2007

An algorithm to increase the node-connectivity of a digraph by one*

András Frank and László A. Végh**

Abstract

We develop a combinatorial polynomial-time algorithm to make a $(k-1)$-connected digraph k-connected by adding a minimum number of new edges. In [7] a min-max theorem was proved (in a more general form) for the minimum number of new edges whose addition makes a $(k-1)$-connected directed graph k-connected. In this paper we describe a new, constructive proof that gives rise to a combinatorial polynomial-time algorithm.

1 Introduction

A directed graph $D = (V, H)$ is called **k-edge-connected** if the number of edges entering X, called the **in-degree** of X, is at least k for every non-empty proper subset X of nodes. A 1-edge-connected digraph is called **strongly connected.** D is called **k-node-connected**, in short, **k-connected** if it has at least $k+1$ nodes and discarding any subset of less than k nodes results in a strongly connected digraph. It is well-known, by versions of Menger's theorem, that D is k-edge-connected (respectively, k-node-connected) if and only if there are k edge-disjoint (openly disjoint) directed paths from each node to every other (and has at least $k+1$ nodes in the k-node-connected case.)

The directed edge-connectivity (node-connectivity) augmentation problem consists of finding a minimum number of edges whose addition to a given digraph results in a k-edge-connected (k-node-connected) digraph. In this paper we consider the special problem of augmenting connectivity by one, that is, we augment the connectivity of a digraph which is already $(k-1)$-edge connected ($(k-1)$-node-connected). Note that for $k = 1$ edge- and node-connectivity coincide. A combinatorial polynomial-time algorithm was developed for the corresponding augmentation problem by K.P. Eswaran and R.E. Tarjan [2] in 1976.

The edge-connectivity augmentation problem was solved in [5] where both a min-max theorem and a combinatorial polynomial-time algorithm were given. As far as

*Research supported by the Hungarian National Foundation for Scientific Research Grant, OTKA T037547, K60802 and by European MCRTN ADONET, Grant Number 504438

**Department of Operations Research, Eötvös University, Pázmány P. s. 1/c, Budapest, Hungary, H-1117, MTA-ELTE Egerváry Research Group (EGRES). e-mail: {frank,veghal}@cs.elte.hu

the node-connectivity augmentation problem is concerned, a min-max theorem was proved in [7]. Here we state it only for the problem of augmenting node-connectivity by one. In a digraph $D = (V, H)$, we call an ordered pair (X, Y) of disjoint non-empty subsets of V a **one-way** pair if there is no edge of D from X to Y. The first member X is called the **tail** of the pair while the second member Y is the **head.** D is clearly $(k-1)$-connected if and only if $|V-(X\cup Y)| \geq k-1$ holds for every one-way pair. We say that in a $(k-1)$-connected digraph a one-way pair is **tight** if $|V-(X\cup Y)| = k-1$. Two pairs are **independent** if their tails are disjoint or their heads are disjoint.

Theorem 1.1 ([7]). *The minimum number of directed edges whose addition to a $(k-1)$-connected digraph $D = (V, H)$ with $|V| \geq k+1$ results in a k-connected digraph is equal to the maximum number of pairwise independent tight one-way pairs.*

The proof in [5] for the corresponding edge-connectivity augmentation theorem relied on the edge splitting-off technique and was thus algorithmic. The proof of Theorem 1.1 in [7] used the uncrossing technique and hence it could not provide a polynomial algorithm. Instead, the theorem itself was used to justify the polynomiality of an algorithm for computing the minimum. That algorithm however relied heavily on the ellipsoid method. The dual optimum could also be computed by a method of T. Fleiner [3] using the min-max theorem, the ellipsoid method, and a clever uncrossing procedure.

It remained an important open problem to find a purely combinatorial algorithm for node-connectivity augmentation. In [8] the first author and T. Jordán exhibited a combinatorial polynomial-time algorithm to make a strongly connected digraph 2-connected, and showed that a similar approach may be used for any fixed k. That is, the running time of the algorithm is polynomial in the size of the digraph but exponential in k. Recently, A. Benczúr and the second author [10] have given a combinatorial algorithm for the general case polynomial also in k.

The present approach is a combinatorial algorithm for augmenting connectivity by one: the special case when the starting digraph is $(k-1)$-connected. The advantage of our approach is that it is much simpler than [10], although has slightly worse running time bounds.

The motivation of our algorithm is a previous algorithm of [6] for Győri's theorem. The general result of [7] contains not only various connectivity augmentation problems but it also implies a deep min-max theorem of E. Győri [9] on the minimum number of generators of family of subpaths of a directed path. (Győri's result found a beautiful application in combinatorial geometry concerning the minimum number of rectangles covering a vertically convex rectilinear polygon in the plane.)

In this paper, instead of considering connectivity augmentation directly, we investigate an equivalent problem. Let $G = (A, B; E)$ be a bipartite graph with color classes A and B. It is well known by Hall's theorem that there exists a matching covering A if and only if $|X| \leq |\Gamma(X)|$ holds for every $X \subseteq A$. G is called **elementary bipartite** if $|A| = |B|$ and one has the stronger property $|X|+1 \leq |\Gamma(X)|$ for every $\emptyset \neq X \subset A$. This is equivalent to the property that removing any edge of G the remaining graph still contains a complete matching. As a generalization, we say that the bipartite graph $G = (A, B; E)$ is **k-elementary** for $k \geq 0$ if $|X| + k \leq |\Gamma(X)|$ or $\Gamma(X) = B$

for every $\emptyset \neq X \subseteq A$. (Note that $|A| = |B|$ is not assumed.) A problem analogous to connectivity augmentation is as follows. Given a $(k-1)$-elementary bipartite graph, add a minimum number of edges to get a k-elementary bipartite graph.

Connectivity augmentation by one can be reduced to this problem. For a digraph $D = (V, H)$ construct the bipartite graph $G = (A, B; E)$ in the following way. With each $v \in V$ associate vertices $v' \in A$ and $v'' \in B$. Each edge $uv \in H$ defines an edge $u'v'' \in E$, and each vertex $v \in V$ defines an edge $v'v'' \in E$. D is clearly k-connected if and only if G is k-elementary. Thus augmenting connectivity of a $(k-1)$-connected digraph by one can be reduced to augmenting a $(k-1)$-elementary bipartite graph to k-elementary. It is not difficult to show that a reduction is possible in the other direction as well.

For a bipartite graph G, let $\tau(G)$ denote the minimum number of edges whose addition to G results in a k-elementary bipartite graph. Let us call such a set of edges an **augmenting edge set**. For a $(k-1)$-elementary bipartite graph G, the set $\emptyset \neq X \subseteq A$ with $\Gamma(X) \neq B$ and $|\Gamma(X)| = |X| + k - 1$ is called **tight**. Two tight sets X and Y are **independent** if $X \cap Y = \emptyset$ or $\Gamma(X \cup Y) = B$. Let $\nu(G)$ denote the number of pairwise independent tight sets. The following min-max formula is a direct consequence of the min-max formula of [7].

Theorem 1.2. *Let $G = (A, B; E)$ be a $(k-1)$-elementary bipartite graph. Then $\nu(G) = \tau(G)$.*

The main purpose of this paper is to give an algorithm which computes an optimal augmentation. Assume we are given a subroutine for determining the optimum value $\nu(G)$ for arbitrary $(k-1)$-elementary bipartite graph G. Making use of Theorem 1.2, one can easily construct an optimal augmenting edge set the following way. First compute $\nu(G)$, and let J be the set of edges not in E between A and B. In each step choose an edge $e \in J$, compute $\nu(G+e)$, and remove e from J. If $\nu(G+e) = \nu(G) - 1$, then add the edge e to G, otherwise keep the same G. Note that Theorem 1.2 ensures the existence of an edge e with $\nu(G+e) = \nu(G) - 1$.

In this paper we develop a subroutine for determining $\nu(G)$. Furthermore, we also present an other algorithm, which uses this subroutine only once, and finds an optimal augmenting set directly. The method is based on a new, algorithmic proof of Theorem 1.2.

We conclude the section by listing some definitions and notation. For sets X and Y, $X \subset Y$ means that X is a proper subset of Y. Let $G = (A, B; E)$ be a $(k-1)$-elementary bipartite graph. Let $\mathcal{T}$ denote the set of tight sets, that is:

$$\mathcal{T} := \{X : \emptyset \neq X \subseteq A, \Gamma(X) \neq B, |X| + k - 1 = |\Gamma(X)|\}.$$

Two tight sets $X, Y \in \mathcal{T}$ are **independent** if $X \cap Y = \emptyset$ or $\Gamma(X \cup Y) = B$. Two non-independent sets are called **dependent**. If $X \subseteq Y$ or $Y \subseteq X$, then we call X and Y **comparable**. X and Y are **crossing** if they are dependent but not comparable. We also use the terms X **crosses** Y or Y **crosses** X in this case.

A set $\mathcal{F} \subseteq \mathcal{T}$ is called **crossing** if $X \cup Y, X \cap Y \in \mathcal{F}$ for any two crossing sets $X, Y \in \mathcal{F}$. Making use of the submodularity of $|\Gamma(X)|$, the following inequalities show that $\mathcal{T}$ itself is crossing.

$$|X \cup Y| + k - 1 + |X \cap Y| + k - 1 \leq |\Gamma(X \cup Y)| + |\Gamma(X \cap Y)| \leq \\ \leq |\Gamma(X)| + |\Gamma(Y)| = |X| + k - 1 + |Y| + k - 1. \qquad (1)$$

We have equality throughout, thus $X \cup Y, X \cap Y \in \mathcal{T}$ follows. If a crossing system $\mathcal{F} \subseteq \mathcal{T}$ contains no two crossing members $\mathcal{F}$ is called **cross-free.**

For a set $K \in \mathcal{F}$ let $\mathcal{F} \div K$ denote the set of sets in $\mathcal{F}$ not crossing K. Similarly, for a subset $\mathcal{K} \subseteq \mathcal{F}$ let $\mathcal{F} \div \mathcal{K}$ denote the set of sets in $\mathcal{F}$ crossing no element of $\mathcal{K}$. Let us call a cross-free subset $\mathcal{F} \subseteq \mathcal{T}$ **complete** if $\mathcal{T} \div \mathcal{F} = \mathcal{F}$, which means that $\mathcal{F}$ is a maximal cross-free subset of $\mathcal{T}$.

A directed edge $e = uv$ **augments** the tight set $X \in \mathcal{T}$ if $u \in X$, $v \in B - \Gamma(X)$. We say that a set F of edges **augments** $\mathcal{F}$ or that F is an **augmenting edge set** of $\mathcal{F}$ if for every member of $\mathcal{F}$ there is an edge in F augmenting it. Let $\tau(\mathcal{F})$ denote the minimum number of augmenting edges, and $\nu(\mathcal{F})$ the maximum number of pairwise independent elements of $\mathcal{F}$.

The following theorem is a slight generalization of Theorem 1.2 and is also an easy consequence of the min-max formula of [7].

Theorem 1.3 ([7]). *For a crossing system $\mathcal{F} \subseteq \mathcal{T}$ the minimum number $\tau(\mathcal{F})$ of edges augmenting $\mathcal{F}$ equals the maximum number $\nu(\mathcal{F})$ of pairwise independent members of $\mathcal{F}$.*

Note that we need at least $\nu(\mathcal{F})$ edges since two pairs are independent if and only if they cannot be augmented by the same edge.

The rest of the paper is organized as follows. In Section 2 we give the description of the Dual Oracle, a subroutine which determines $\nu(\mathcal{T})$. In Section 2.2 we analyze the oracle and the first algorithm which relies on this oracle. In Section 3, we give a new proof for Theorem 1.3, and sketch a second algorithm. For this algorithm, we present only the main ideas, and omit the technical details which can be done similarly as for the first algorithm.

2 The Dual Oracle

The following theorem is the essence of the Dual Oracle.

Theorem 2.1. *For a complete cross-free system $\mathcal{K} \subseteq \mathcal{T}$ the maximum number of pairwise independent sets is equal in $\mathcal{K}$ and $\mathcal{T}$, that is, $\nu(\mathcal{K}) = \nu(\mathcal{T})$.*

Clearly, $\nu(\mathcal{K}) \leq \nu(\mathcal{T})$ for every $\mathcal{K} \subseteq \mathcal{T}$. The advantage of a cross-free system is that we can easily determine the maximum number of pairwise independent sets. This is due to the fact that whenever it contains two dependent sets they are comparable. Thus considering the partially ordered set $(\mathcal{K}, \subseteq)$ an antichain consists of pairwise independent sets. A maximum antichain in a poset can be easily found by an algorithm based on Dilworth's theorem stating the equality of the size of a minimum chain cover and a maximum antichain. In order to prove Theorem 2.1, we need some elementary propositions.

Claim 2.2. *If $X, Y \in \mathcal{T}$ are dependent, then* $\Gamma(X) \cap \Gamma(Y) = \Gamma(X \cap Y)$

Proof. This follows since the second inequality in (1) holds with equality. □

Claim 2.3. *If $Y \in \mathcal{T}$, $X \subseteq A$ and $\Gamma(X) \subseteq \Gamma(Y)$, then $X \subseteq Y$.*

Proof. If X is not a subset of Y, then $|X \cup Y| > |Y|$. On the other hand, $|\Gamma(Y)| = |Y| + k - 1$ and $\Gamma(Y) = \Gamma(X \cup Y)$, thus $|\Gamma(X \cup Y)| < |X \cup Y| + k - 1$ contradicting the fact that G is $(k-1)$-elementary. □

Lemma 2.4. *For a crossing family $\mathcal{F}$ of and for any $K \in \mathcal{T}$, the subfamily $\mathcal{F} \div K$ is crossing.*

Proof. Let $\mathcal{F}' = \mathcal{F} \div K$ and let X and Y be two crossing members of $\mathcal{F}'$. We have to prove that neither $X \cup Y$ nor $X \cap Y$ crosses K.

First assume that K is comparable with both X and Y. It is not possible that $X \subseteq K \subseteq Y$ or $Y \subseteq K \subseteq X$ as X and Y are not comparable. Therefore either $K \subseteq X, Y$ or $K \supseteq X, Y$. In the first case K is contained in both $X \cup Y$ and $X \cap Y$, in the second case it contains both of them.

Second, assume that K is independent from both X and Y. If $K \cap X = \emptyset$ and $K \cap Y = \emptyset$, then K is disjoint from both $X \cup Y$ and $X \cap Y$. If $K \cap X = \emptyset$ and $\Gamma(K \cup Y) = B$ then $K \cap (X \cap Y) = \emptyset$ and $\Gamma(K \cup (X \cup Y)) = B$, thus $X \cap Y, X \cup Y \in \mathcal{F}'$. Finally, if $\Gamma(K \cup X) = \Gamma(K \cup Y) = B$ then $\Gamma(K \cup (X \cup Y)) = B$ shows $X \cup Y \in \mathcal{F}'$, and $\Gamma(K \cup (X \cap Y)) = B$ also follows by Claim 2.2.

In the third case K is independent from one of X and Y, say from X, and comparable with the other, Y. If $K \cap X = \emptyset$ then $K \subseteq Y$ follows since $X \cap Y \neq \emptyset$. This gives $K \cap (X \cap Y) = \emptyset$, $K \subseteq X \cup Y$, giving $X \cap Y, X \cup Y \in \mathcal{F}'$. Finally, if $\Gamma(K \cup X) = B$, then $\Gamma(X \cup Y) \neq B$ implies $Y \subseteq K$, giving $\Gamma(K \cup (X \cup Y)) = B$, $X \cap Y \subseteq K$. □

Lemma 2.5. *(i) Suppose K and L are dependent, $K \cap L$ and M are also dependent, but L and M are independent for some $K, L, M \in \mathcal{T}$. Then $(\Gamma(L) - \Gamma(K)) - \Gamma(M) \neq \emptyset$ and $K - L \subseteq M$.*

(ii) Let K and L be dependent, $K \cup L$ and M also dependent, but L and M independent for some $K, L, M \in \mathcal{T}$. Then $(\Gamma(L) - \Gamma(K)) \cap \Gamma(M) = \emptyset$ and $M \cap (K - L) \neq \emptyset$.

Proof. (i) Observe $K \cap L \cap M \neq \emptyset$, thus the independency of L and M implies $\Gamma(L \cup M) = B$. By the dependency of $K \cap L$ and M, $\Gamma(M \cup (K \cap L)) \neq B$ giving the first part of the claim using Claim 2.2. For the second part, consider $\Gamma((K \cap L) \cup M) = \Gamma((K \cup M) \cap (L \cup M)) = \Gamma(K \cup M) \cap \Gamma(L \cup M) = \Gamma(K \cup L)$. By Claim 2.3, we have $K \cup L \subseteq (K \cap L) \cup M$, implying $K - L \subseteq M$.

(ii) As $K \cup L$ and M are dependent, L and M can be independent only if $M \cap L = \emptyset$. Since $M \cap (K \cup L) \neq \emptyset$, we have $M \cap (K - L) \neq \emptyset$. Claim 2.2 gives $\Gamma(K \cup L) \cap \Gamma(M) = \Gamma((K \cup L) \cap M) = \Gamma(K \cap M) \subseteq \Gamma(K)$, showing the first part of the claim. □

Now we are ready to prove Theorem 2.1. The proof is based on the following lemma:

Lemma 2.6. *For a crossing system $\mathcal{F}$ and a pair $K \in \mathcal{F}$ we have $\nu(\mathcal{F}) = \nu(\mathcal{F} \div K)$.*

First we show how Theorem 2.1 follows from this lemma. Let $\mathcal{K} = \{K_1, \dots, K_\ell\}$. First apply this lemma for $\mathcal{T}$ and K_1, then in the ith step for $\mathcal{T} \div \{K_1, \dots K_{i-1}\}$ and K_i. Note that $\mathcal{T} \div \{K_1, \dots K_{i-1}\}$ is a crossing system by applying inductively Lemma 2.4. Thus we have $\nu(\mathcal{T}) = \nu(\mathcal{T} \div K_1) = \dots = \nu(\mathcal{T} \div \mathcal{K})$, hence the claim follows by $\mathcal{T} \div \mathcal{K} = \mathcal{K}$.

Proof of Lemma 2.6. Trivially, $\nu(\mathcal{F} \div K) \le \nu(\mathcal{F})$. Consider a maximum independent subset $\mathcal{L}$ of $\mathcal{F}$ which has the most common members with $\mathcal{F} \div K$. Suppose by contradiction that $\mathcal{L} \cap (\mathcal{F} \div K) < \nu(\mathcal{F})$ and choose an element $L \in \mathcal{L} - (\mathcal{F} \div K)$. By definition L crosses K. We claim that either $\mathcal{L} \setminus \{L\} \cup \{L \cap K\}$ or $\mathcal{L} \setminus \{L\} \cup \{L \cup K\}$ is independent. This leads to contradiction, since the new system intersects $\mathcal{F} \div K$ in a strictly larger subset than $\mathcal{L}$ does.

Suppose that neither $\mathcal{L} \setminus \{L\} \cup \{L \cap K\}$ nor $\mathcal{L} \setminus \{L\} \cup \{L \cup K\}$ is independent. Then there is an element $M \in \mathcal{L}$ crossing $L \cap K$, and an other element $M' \in \mathcal{L}$ crossing $L \cup K$. If $M = M'$, then M is clearly dependent with L, a contradiction.

Assume now $M \ne M'$. The conditions of Lemma 2.5(i) hold for K, L and M, and the conditions of (ii) hold for for K, L and M'. We claim that M and M' are dependent. It follows as $\Gamma(L) - \Gamma(K)$ contains an element in $B - \Gamma(M \cup M')$, and $K - L$ contains an element in $M \cap M'$. □

2.1 Constructing a complete cross-free subset

A straightforward approach to construct a complete cross-free subset of $\mathcal{T}$ would be to select sets greedily, that is, as long as possible choose sets which do not cross the the previously selected ones. The difficulty arises from the fact that it is not clear how to decide whether a given cross-free system is complete or not. (Note that the size of $\mathcal{T}$ may be exponentially large.) To overcome this difficulty we work with a special kind of cross-free systems. Let us call a cross-free subset $\mathcal{H}$ **down-closed** if it fulfills the following property:

$$Z \text{ crosses some element of } \mathcal{H} \text{ whenever } K \in \mathcal{H},\ Z \subset K,\ Z \in \mathcal{T} - \mathcal{H}. \tag{2}$$

This means that if $\mathcal{H}$ has an element containing Z, then Z cannot be added to $\mathcal{H}$. Given a down-closed system, the following claim provides a straightforward way to decide whether it is complete.

Claim 2.7. *A down-closed system is complete if and only if it contains all the maximal members of $\mathcal{T}$.*

Proof. On the one hand, any complete cross-free system should contain all the maximal sets of $\mathcal{T}$ since a maximal set cannot cross any other set. On the other hand, suppose by contradiction that a down-closed system $\mathcal{H}$ contains all the maximal members, but it is not complete. Choose a $Z \notin \mathcal{H}$ with $\mathcal{H} \cup \{Z\}$ cross-free. There is a maximal element $K \in \mathcal{T}$ with $Z \subseteq K$. By our assumption $K \in \mathcal{H}$, contradicting the definition of the down-closed system. □

Assume we are given a down-closed system $\mathcal{H}$ which is not complete. In the following, we investigate how a set $K \in \mathcal{T} - \mathcal{H}$ can be found with the property that $\mathcal{H} \cup \{K\}$ is down-closed as well.

As $\mathcal{H}$ is not complete, we can find a maximal element M with $M \in \mathcal{T} - \mathcal{H}$. Let

$$\mathcal{L}_1 := \{K \in \mathcal{H} : \ K \subseteq M\}; \ \mathcal{L}_2 := \{K \in \mathcal{H} : K \not\subseteq M\} \tag{3}$$

We say that a set Z **fits** the pair $(\mathcal{H}, M)$ if (a) $Z \in \mathcal{T} - \mathcal{H}, Z \subseteq M$; (b) Z is independent of all members in $\mathcal{L}_2$ and (c) either $K \subset Z$ or $K \cap Z = \emptyset$ for every $K \in \mathcal{L}_1$.

Lemma 2.8. *If Z is a minimal member of $\mathcal{T} - \mathcal{H}$ fitting $(\mathcal{H}, M)$, then $\mathcal{H} \cup \{Z\}$ is down-closed.*

This is a straightforward consequence of the following claim.

Claim 2.9. *Let $Z \in \mathcal{T} - \mathcal{H}$, $Z \subseteq M$. The following two properties are equivalent: (i) Z fits $(\mathcal{H}, M)$; (ii) $\mathcal{H} \cup \{Z\}$ is cross-free.*

Proof. (i)$\Rightarrow$(ii) is straightforward. For the other direction we have to verify (b) and (c) of the above definition. By (2), either $K \subset Z$ or Z and K are independent for arbitrary $K \in \mathcal{H}$. By contradiction to (b), suppose $K \subset Z$ for some $K \in \mathcal{L}_2$. In this case $K \subset Z \subseteq M$, contradicting the definition of $\mathcal{L}_2$. For (c) we need $K \cap Z = \emptyset$ if K and Z are independent for some $K \in \mathcal{L}_1$. This follows by $K, Z \subseteq M$, thus $\Gamma(K \cup Z) \subseteq \Gamma(M) \subset B$. □

Observe that M itself fits $(\mathcal{H}, M)$ ensuring the existence of a Z satisfying the conditions of Lemma 2.8. So $K = Z$ is an appropriate selection. Such a Z can be found using bipartite matching theory. The description of this subroutine is quite technical and rather standard, therefore it is moved to an Appendix.

2.2 Description of the Dual Oracle

Given the above subroutine for constructing a complete down-closed system, we have the following oracle to determine the value $\nu(G)$ in a $(k-1)$-elementary bipartite graph $G = (A, B; E)$: we construct a complete down-closed system, then we apply Dilworth's theorem. (It is well-known that computing a maximum antichain and a minimum chain-decomposition of a partially ordered set can be reduced to a maximum matching computation in a bipartite graph.) The size of the maximum antichain will give the value $\nu(G)$.

A trivial upper bound for the size of the optimal augmenting edge set and by Theorem 1.2 also for the number of pairwise independent sets is $|A||B|$. A better bound can be given following the argument of Theorem 4.5 in [7]. It can be proved that there is an optimal augmenting set which is a matching, hence the maximum independent system is of cardinality at most $|A|$. A chain can also have at most $|A|$ elements, thus the cardinality of a complete cross-free system is at most $s = |A|^2$.

As shown in the Appendix, if s is an upper bound for the size of a complete down-closed system, then it can be constructed in $O(s|A|^3|B|) = O(|A|^5|B|)$ running time.

Finding a maximum antichain in a poset of size $O(s)$ can be reduced to finding a maximum matching in a bipartite graph on $O(s)$ vertices and $O(s^2)$ edges. Using the Hopcroft-Karp algorithm this can be done in $O(s^{2.5})$ running time. This gives $O(|A|^5)$ for $s = |A|^2$, so the total running time of the Dual Oracle is $O(|A|^5|B|)$.

As we have already indicated in the Introduction, the Dual Oracle may be used to compute the optimal augmentation. For this, we need to call the Dual Oracle at most $|A||B|$ times, thus the total complexity is $O(|A|^6|B|^2)$. For connectivity augmentation by one, this gives $O(n^8)$, where n is the number of vertices of the graph. (For comparison, the running time of the algorithm in [10] is $O(n^7)$ for the same problem.) However, the correctness of the present approach does rely on Theorem 1.2. In the next section we use a more direct approach for finding the optimal augmentation.

3 Algorithmic Proof of Theorem 1.3

In this section we give a proof of Theorem 1.3 and sketch another algorithm, which uses the Dual Oracle only once. After a complete down-closed system $\mathcal{K}$ is determined, an augmenting set of $\mathcal{K}$ can be transformed to an augmenting set of the entire $\mathcal{T}$. This will also give a new proof for Theorem 1.3. We begin with the definition of the elementary augmenting step.

Consider a crossing family $\mathcal{F}$, and let F' be a set of edges between A and B. We say that a pair (u, v) of nodes with $u \in A$, $v \in B$ is **bad** (with respect to $\mathcal{F}$ and F') if there is a member X of $\mathcal{F}$ with $u \in X$, $v \notin \Gamma(X)$, and X is not augmented by F'. Let $W(F') = W_{\mathcal{F}}(F')$ denote the set of bad pairs.

Consider an augmenting edge set F' of $\mathcal{F}' := \mathcal{F} \div K$. For two elements $f'_1 = x_1y_1, f'_2 = x_2y_2$ of F', define

$$f_1 := x_1y_2 \text{ and } f_2 := x_2y_1 \tag{4}$$

and let

$$F'' := F' - \{f'_1, f'_2\} \cup \{f_1, f_2\}. \tag{5}$$

We will say that F'' arises from F' by **flipping** $\{f'_1, f'_2\}$. A flipping is called **improving** if F'' augments a strictly larger subset of $\mathcal{F}$ than F' does. Note that this is equivalent to requiring that $W(F'') \subset W(F')$ and, since the total number of edges is $|A||B|$, we obtain that after at most $|A||B|$ improving flippings the resulting subset of edges must augment the whole $\mathcal{F}$. The following lemma, which is the heart of the proof of Theorem 1.3 and the algorithm, asserts the existence of an improving flipping.

Lemma 3.1. *Let $\mathcal{F} \subseteq \mathcal{T}$ be a crossing family. Let K be a member of $\mathcal{F}$ and F' an augmenting edge set of $\mathcal{F}' := \mathcal{F} \div K$. If F' does not augment $\mathcal{F}$, then there is an improving flipping.*

Proof. Let us choose two (not necessarily distinct) members X and Y of $\mathcal{F}$ that are not augmented by F' so that $X \subseteq Y$, X is minimal (in the sense that X' is augmented by F' for every $X' \in \mathcal{F}, X' \subset X$) while Y is maximal in an analogous sense.

Since F' does not augment X and Y, we have $X, Y \in \mathcal{F} - \mathcal{F}'$, that is, both X and Y cross K. Therefore $X \cap K \subset X$ and $Y \cup K \supset Y$. By the minimality of X, $X \cap K$ is augmented by F', that is, there is an edge $f_1' = x_1y_1$ in F' augmenting $X \cap K$. Since F' does not augment X, we must have $x_1 \in X \cap K$ and $y_1 \in \Gamma(X) - \Gamma(K)$. Analogously, there is an edge $f_2' = x_2y_2$ in F' augmenting $Y \cup K$ for which $x_2 \in K - Y, y_2 \in B - \Gamma(K \cup Y)$. Let f_1, f_2 and F'' be defined by (4) and (5).

We are going to show that flipping $\{f_1', f_2'\}$ is improving. Since X is augmented by F'' but not augmented by F', we have only to show that every member of $\mathcal{F}$ augmented by F' is augmented by F'', as well.

Suppose indirectly that there is a member M of $\mathcal{F}$ which is augmented by F' but not by F''. In particular, no element of $F' - \{f_1', f_2'\}$ augments M. It is not possible that both f_1' and f_2' augments M since then both f_1 and f_2 would augment M, that is, F'' would augment M. Therefore there is exactly one element in F' augmenting M and this only element is either f_1' or f_2'. Let us assume first that M is augmented by f_1'.

Claim 3.2. *Y and M are dependent.*

Proof. Suppose by contradiction that Y and M are independent. $K \cap Y$ and M are dependent as f_1' augments both. Thus we can apply Lemma 2.5(i) with $Y = L$, giving $K - Y \subseteq M$. This is contradiction since $x_2 \in K - Y$ and $x_2 \notin M$ as f_2 does not augment M. □

The assumption that M is not augmented by F'' gives $y_2 \in \Gamma(M) - \Gamma(Y)$, as otherwise f_1 would augment F, thus $Y \not\subseteq M$, implying $Y \cup M \supset Y$. By the maximality of Y, $Y \cup M$ is augmented by an element $f = xy$ of F' and f is different from f_1' and f_2' since $y_1, y_2 \in \Gamma(M \cup Y)$. As $x \in Y \cup M$, $y \in B - \Gamma(Y \cup M)$, f augments either M or Y. However, $f \in F'' \cap F'$ and hence f augments neither M nor Y, a contradiction.

The case when M is augmented only by f_2' also leads to contradiction by a similar argument using Lemma 2.5(ii). □

Proof of Theorem 1.3

$\nu \leq \tau$ follows since no two independent pairs can be augmented by the same edge. To see the other direction, we distinguish two cases.

Case 1. *$\mathcal{F}$ is cross-free.* By applying Dilworth's theorem to the partially ordered set $(\mathcal{F}, \subseteq)$, we obtain that there is a maximum subfamily $\mathcal{I}$ of $\mathcal{F}$ consisting of uncomparable members and that $\mathcal{F}$ can be decomposed into $\gamma := |\mathcal{I}|$ chains. Since $\mathcal{F}$ is assumed to be cross-free, the members of $\mathcal{I}$ are pairwise independent. Furthermore, it is easy to see that the chain-decomposition of $\mathcal{F}_s$ corresponds to a set F of γ edges augmenting $\mathcal{F}$. Hence we obtained the required covering F of $\mathcal{F}$ and independent subfamily $\mathcal{I}$ of $\mathcal{F}$ for which $|F| = |\mathcal{I}|$.

Case 2. *There is a member K of $\mathcal{F}$ crossing some other members of $\mathcal{F}$.* Let $\mathcal{F}' := \mathcal{F} \div K$. By Lemma 2.4, $\mathcal{F}'$ is a crossing family, so by induction, there is an independent subfamily $\mathcal{I}$ of $\mathcal{F}'$ and a covering F' of $\mathcal{F}'$ for which $|\mathcal{I}| = |F'|$. Choose F' in such a way that the number of bad pairs of nodes is minimum. By Lemma 3.1, this number is zero, that is, F' covers the whole $\mathcal{F}$. □

3.1 Description of the Algorithm

Our next goal is to transform the inductive proof above into an algorithm, that constructs an independent subset $\mathcal{I}$ of $\mathcal{T}$ and an augmenting edge set F of $\mathcal{T}$ so that $|\mathcal{I}| = |F|$. It consists of two phases.

In Phase 1 our algorithm uses the Dual Oracle. It determines a complete down-closed system $\mathcal{K} = \{K_1, \ldots, K_\ell\}$, and by Dilworth's theorem it finds a maximum antichain along with a minimum chain-decomposition. The chain-decomposition of $\mathcal{K}$ corresponds to a subset F' of edges augmenting $\mathcal{K}$ for which $|F'| = |\mathcal{I}|$. The antichain $\mathcal{I}$ will be output by the whole algorithm as a maximum cardinality independent subset of $\mathcal{T}$.

Phase 2 will terminate by outputting a covering of $\mathcal{T}$ of cardinality $|\mathcal{I}|$. Let $\mathcal{F}_0 = \mathcal{T}$ and $\mathcal{F}_j := \mathcal{T} \div \{K_1, \ldots, K_j\}$ for each $j = 1, \ldots, \ell$. From Phase 1, we have $\mathcal{F}_\ell = \mathcal{K}$ covered. By Lemma 3.1, when applied to $\mathcal{F}_{\ell-1}, \mathcal{F}_\ell, K_\ell$ in place of $\mathcal{F}, \mathcal{F}', K$, respectively, we can find an improving flipping and obtain a revised covering F'' of $\mathcal{F}_\ell$ which covers a strictly larger subset of $\mathcal{F}_{\ell-1}$ as F' does. Since the number of bad pairs is at most $|A||B|$ and an improving flipping reduces this number, after at most $|A||B|$ improving flippings the resulting covering of $\mathcal{F}_\ell$ will cover $\mathcal{F}_{\ell-1}$. Then we can iterate this step with $\mathcal{F}_{\ell-2}, \mathcal{F}_{\ell-1}, K_{\ell-1}$, …, $\mathcal{F}_0, \mathcal{F}_1, K_1$, and finally we get a cover F' of $\mathcal{T} = \mathcal{F}_0$. F' will be the output of the algorithm as a minimal edge set whose addition to G results in a k-elementary bipartite graph.

We have outlined the steps of the algorithm and proved its validity. Phase 1 can be preformed as described in Section 2. For the realization of Phase 2, we can use similar technics as in Section 2.2. However, we omit this analysis. Our reason for this is that the analysis is quite technical, and we cannot prove a better running time bound then $O(|A|^6|B|^2)$, which we had for the Dual Algorithm.

4 Concluding remarks

This approach can be extended to solve algorithmically other connectivity augmentation problems as well, for example, to increase the ST-edge-connectivity of a digraph by one. (A digraph is called k-edge-connected from S to T if there are k edge-disjoint paths from every node of S to every node of T). The main difficulty is that instead of (2) we have to maintain a more complicated property when selecting the elements K_j.

Finally, we remark that the method is suitable for solving a minimum cost version of the connectivity augmentation problem for node-induced costs. That is, the cost of a possible new directed edge uv is defined by $c^-(u) + c^+(v)$ where c^- and c^+ are two cost-functions on V. For general cost functions, even the special case of making the graph $(V, \emptyset)$ strongly connected is NP-complete as being a generalization of the TSP problem.

References

[1] A. Frank, An algorithm to increase the node-connectivity of a digraph, Proceedings of the 3rd Hungarian-Japanese Symposium on Discrete Mathematics and Its Applications, Tokyo, Japan (2003), 378-387.

[2] K.P. Eswaran and R.E. Tarjan, Augmentation problems, SIAM J. Computing, Vol. 5, No. 4, December 1976, 653-665.

[3] T. Fleiner, Uncrossing a family of set-pairs, Combinatorica, Vol. 21, No. 1, (2001) pp. 145-150.

[4] L.R. Ford and D.R. Fulkerson, Flows in Networks, Princeton Univ. Press, Princeton NJ., 1962.

[5] A. Frank, Augmenting graphs to meet edge-connectivity requirements, SIAM J. Discrete Mathematics, (1992 February), Vol.5, No 1., pp.22-53.

[6] A. Frank, Finding minimum generators of path systems, J. Combinatorial Theory, Ser B., 75, (1999) pp. 237-244.

[7] A. Frank and T. Jordán, Minimal edge-coverings of pairs of sets, J. Combinatorial Theory, Vol. 65, No. 1 (1995, September) pp. 73-110.

[8] A. Frank and T. Jordán, Directed vertex-connectivity augmentation, in: Connectivity Augmentation of Networks: Structures and Algorithms, Mathematical Programming, (ed. A. Frank) Ser. B, Vol. 84, No. 3 (1999), pp. 537-553.

[9] E. Győri, A minimax theorem on intervals, J. Combinatorial Theory, Ser. B 37 (1984) 1-9.

[10] L.A. Végh, A.A. Benczúr, Primal-dual Approach for Directed Vertex Connectivity Augmentation and Generalizations, Proceedings of SODA 2005, pp. 186-194.

5 Appendix

In this Appendix we present how the subroutine for constructing a complete down-closed system can be implemented using bipartite matching theory.

Given the bipartite graph $G = (A, B; E)$ and a function $f : A \cup B \to \mathbb{N}$ we call the set $F \subseteq E$ an f**-factor** if $d_F(x) = f(x)$ for every $x \in A \cup B$ where $d_F(x)$ denotes the number of edges in F incident to x. Let $f(Z) = \sum_{x \in Z} f(x)$ for $Z \subseteq A \cup B$.

Claim 5.1. *Consider a bipartite graph $G = (A, B; E)$ and a function $f : A \cup B \to \mathbb{N}$ so that $f(A) = f(B)$ and $f(x) = 1$ or $f(y) = 1$ for every $xy \in E$. An f-factor exists if and only if for any $X \subseteq A$, $f(X) \le f(\Gamma(X))$ for every $X \subseteq A$.*

Proof. An easy consequence of Hall's theorem, replacing each $x \in A \cup B$ by $f(x)$ copies. Note that by the condition $f(x) = 1$ or $f(y) = 1$ for every $xy \in E$ no edge is used more than once. □

First we show how the maximal elements of $\mathcal{T}$ can be found. Let us consider elements $a \in A$ and $b \in B$ with $ab \notin E$. A set $X \in \mathcal{T}$ is called an ***ab*-set**, if ab covers X, that is, $a \in X$ and $b \notin \Gamma(X)$. For $ab \notin E$, consider the following f. Let $f(a) = f(b) = k+1$ and for $c \in A \cup B - a - b$, let $f(c) = 1$. For this f, an f-factor is called a ***k*-*ab*-factor**. If G is a $(k-1)$-elementary bipartite graph, then Claim 5.1 implies the existence of a $(k-1)$-ab-factor, denoted by F_{ab}.

Claim 5.2. *If there exists a k-ab-factor, then there is no ab-set.*

Proof. Assume X is an ab-set. As $X \in \mathcal{T}$, $|\Gamma(X)| = |X| + k - 1$. Since $a \in X$, $b \notin \Gamma(X)$, we have $f(X) = |X| + k$, $f(\Gamma(X)) = |X| + k - 1$, thus by Claim 5.1 no k-ab-factor exists. □

It is easy to see that any two ab-sets are dependent and the union and intersection of two ab-sets are ab-sets as well. Thus if the set of ab-sets is nonempty, then it contains a unique minimal and maximal element. Now we show how these can be found algorithmically. We say that the path $U = x_0y_0x_1y_1 \dots x_ty_t$ is an **alternating path** for $F \subseteq E$ from x_0 to y_t, if $x_i \in A$, $y_i \in B$, $x_iy_i \notin F$ for $i = 0, \dots, t$, and $y_ix_{i+1} \in F$ for $i = 0, \dots, t-1$. By the same conditions we also say that $x_0y_0x_1y_1 \dots x_t$ is an alternating path for F from x_0 to x_t.

Claim 5.3. *(a) If there exists an alternating path for F_{ab} between a and b, then there exists no ab-set. (b) Assume there is no alternating path for F_{ab} from a to b; let S denote the set of vertices c having an alternating path for F_{ab} from a to c, and let $X = A \cap S$. Then X is the unique minimal ab-set. (c) Assume no alternating path exists for F_{ab} from a to b; let S' denote the set of vertices c having an alternating path for F_{ab} from c to b, and let $Y = A - S'$. Then Y is the unique maximal ab-set.*

Proof. (a) Let U be an alternating path for F_{ab} from a to b. Then $F_{ab}\Delta U$ is a k-ab-factor so by Claim 5.2 no ab-set exists. (b) Let Z be an arbitrary ab-set. $\Gamma(Z)$ contains a unique y with $xy \in F_{ab}$ for every $x \in Z - a$. The number of $y \in B$ with $ay \in F_{ab}$ is exactly k, and all of them are contained in $\Gamma(Z)$. These are $|Z| + k - 1$ different elements of $\Gamma(Z)$, and since $Z \in \mathcal{T}$, $\Gamma(Z)$ has no other elements than these. This easily implies that Z contains every $x \in A$ for which there is an alternating path for F_{ab} from a to x, showing $X \subseteq Z$. It is left to prove that $X \in \mathcal{T}$. It is sufficient to show that there exists an $x \in X$ with $xy \in F_{ab}$ for every $y \in \Gamma(X)$. This follows from the definition of X, completing the proof of (b). The proof of (c) follows the same lines. □

At the initialization of the algorithm, we determine the sets F_{ab} by a single max-flow computation for every $a \in A$, $b \in B$, $ab \notin E$. By Claim 5.3 the maximal ab-sets can be found by a breadth-first search. The maximals among these will give the maximal elements of $\mathcal{T}$ (note that the maximal ab-set might be contained in some other $a'b'$-set). We will use the sets F_{ab} also in the later steps of the algorithm.

To implement the basic step of the algorithm, consider a down-closed $\mathcal{H}$ which is not complete, a maximal element $M \in \mathcal{T} - \mathcal{H}$ and $\mathcal{L}_1$, $\mathcal{L}_2$ as defined by (3). Our task is to find a K fitting $(\mathcal{H}, M)$ and minimal subject to this property. Let $\mathcal{M}$ be the set of the maximal elements of $\mathcal{L}_1$.

Claim 5.4. *$\mathcal{M}$ consits of pairwise disjoint sets.*

Proof. Let $T_1, T_2 \in \mathcal{M}$. As they are maximal, they cannot be comparable, thus either $T_1 \cap T_2 = \emptyset$ or $\Gamma(T_1 \cup T_2) = B$. The latter is excluded since $T_1, T_2 \subset M$ implies $\Gamma(T_1 \cup T_2) \subseteq \Gamma(M) \subset B$. $\square$

Let us construct $G' = (A, B; E')$ from G as follows. The set E' contains E and some additional edges. For each $X \in \mathcal{L}_2$, let $xy \in E'$ for every $x \in X$, $y \in B - \Gamma(X)$. Furthermore, let $xy \in E'$ whenever $T \in \mathcal{M}$, $x \in T$ and $y \in \Gamma(T)$.

Claim 5.5. *Let $Z \in \mathcal{T} - \mathcal{H}$, $Z \subseteq M$. Z fits $(\mathcal{H}, M)$ if and only if Z is a tight set in G'.*

Proof. The tight sets of G' are those tight sets Z of G for which there is no edge in $E' - E$ augmenting Z, that is, no edge xy with $x \in Z$ and $y \in B - \Gamma(Z)$.

Z fits $(\mathcal{H}, M)$ if it is independent from all elements of $\mathcal{L}_2$, and for arbitrary $T \in \mathcal{M}$, either $T \cap Z = \emptyset$ or $T \subset Z$. If it satisfies these properties, no new edge in G' augments Z, thus Z is tight also in G'. For the other direction, if Z is dependent with some $X \in \mathcal{L}_2$, then there exists $x \in X \cap Z$, $y \in B - \Gamma(X \cup Z)$ with $xy \in E'$ augmenting Z. If for some $T \in \mathcal{M}$, T would cross X, then by Claim 2.3, $\Gamma(T) - \Gamma(Z) \neq \emptyset$, thus there exist $x \in T \cap X$, $y \in \Gamma(T) - \Gamma(Z)$ with $xy \in E'$ augmenting Z. $\square$

However, it is not enough to determine a minimum tight set of G', since the elements of $\mathcal{M}$ are among these, and we are looking for a $Z \in \mathcal{T} - \mathcal{H}$. To exclude the elements of $\mathcal{M}$, we add some further edges to G'. Let $Q \subseteq M$ be an arbitrary set. Let $Z(Q)$ denote the unique minimal X satisfying the following property:

$$X \in \mathcal{T},\ Q \subseteq X, \text{ and } X \text{ fits } (\mathcal{H}, M). \tag{6}$$

We will determine $Z(Q)$ for different Q sets in order to find K. $Z(Q)$ is well-defined since M itself satisfies (6); and if X and X' satisfy (6), then X and X' are dependent and it is easy to see that $X \cap X'$ also satisfies (6). The following claim gives an easy algorithm for finding $Z(Q)$ for a given Z.

Claim 5.6. *Fix some $a \in Q$, $b \in B - \Gamma(M)$. Let G'' denote the graph obtained from G' by adding all edges ay with $y \in \Gamma(Q)$. Let S denote the set of vertices c for which there exists an alternating path for F_{ab} from a to c, and let $X = A \cap S$. Then $Z(Q) = X$.*

Proof. As M is an ab-set in G'', applying Claim 5.3(a) for G'' instead of G, we get that G'' contains no alternating path for F_{ab} from a to b. By Claim 5.3(b), X is the unique minimal ab-set in G''. $\Gamma(X \cup Q) = \Gamma(X)$, thus by Claim 2.3, $Q \subseteq X$. By Claim 5.5, X is the unique minimal set satisfying (6), thus $Z(Q) = X$. $\square$

Let L denote the union of the elements of $\mathcal{M}$. First, we find a set Z_1 fitting $(\mathcal{H}, M)$ and $Z_1 - L \neq \emptyset$. Let us compute the set $Z(a)$ for any $a \in M - L$. By Claim 5.6, this can be done by a single breadth-first search. We get a good Z_1 by choosing a minimal element of the set $\{Z(a) : a \in M - L\}$.

Thus can be found by $M - L = O(|A|)$ breadth-first searches. Now either Z_1 is itself a minimal set fitting $(\mathcal{H}, M)$, or there exists a $Z_2 \subseteq L \cap Z_1$, also fitting $(\mathcal{H}, M)$. This is impossible if Z_1 contains only one element of $\mathcal{M}$, so in this case Z_1 is a minimal set fitting $(\mathcal{H}, M)$.

Assume now Z_1 contains at least two sets in $\mathcal{M}$. In order to obtain Z_2, let us compute $Z(T_i \cup T_j)$ for every two disjoint members $T_i, T_j \in \mathcal{M}$, $T_i, T_j \subset Z_1$. Choosing a minimal among these gives a minimal Z_2 fitting $(\mathcal{H}, M)$. This can be obtained by $O(|A|^2)$ breadth-first searches.

As Z_2 fits $(\mathcal{H}, M)$ and is minimal subject to this property, $K := Z_2$ is an appropriate choice.

5.1 Complexity

To find a complete down-closed system first we need $|A|^2$ Max Flow computations for computing the maximal members and the auxiliary graphs. The running time for determining a member of the complete down-closed system is dominated by $O(|A|^2)$ breadth first searches. Thus if s is an upper bound for the size of a skeleton then we can determine a complete down-closed system in $O(s|A|^3|B|)$ running time.