

Vector Connectivity in Graphs*

Endre Boros

RUTCOR, Rutgers University, New Jersey, USA
640 Bartholomew Rd, Piscataway NJ 08854-8003, USA
{boros}@rutcor.rutgers.edu

Pinar Heggernes, Pim van 't Hof

Department of Informatics, University of Bergen, Norway
{pinar.heggernes,pim.vanthof}@ii.uib.no

Martin Milanič

University of Primorska, UP IAM, Muzejski trg 2, SI6000 Koper, Slovenia
University of Primorska, UP FAMNIT, Glagoljaška 8, SI6000 Koper, Slovenia
martin.milanic@upr.si

March 3, 2013

Abstract

Motivated by challenges related to domination, connectivity, and information propagation in social and other networks, we initiate the study of the VECTOR CONNECTIVITY problem. This problem takes as input a graph G and an integer k_v for every vertex v of G , and the objective is to find a vertex subset S of minimum cardinality such that every vertex v either belongs to S , or is connected to at least k_v vertices of S by disjoint paths. If we require each path to be of length exactly 1, we get the well-known VECTOR DOMINATION problem, which is a generalization of the famous DOMINATING SET problem and several of its variants. Consequently, our problem becomes NP-hard if an upper bound on the length of the disjoint paths is also supplied as input. Due to the hardness of these domination variants even on restricted graph classes, like split graphs, VECTOR CONNECTIVITY seems to be a natural problem to study for drawing the boundaries of tractability for this type of problems. We show that VECTOR CONNECTIVITY can actually be solved in polynomial time on split graphs, in addition to cographs and trees. We also show that the problem can be approximated in polynomial time within a factor of $\ln n + 2$ on all graphs.

Keywords: vector connectivity, approximation algorithm, polynomial time algorithm, split graph, cograph, tree

*A short version of this paper was accepted for presentation at the conference TAMC 2013 [13].

1 Introduction and Motivation

Connectivity between parts of a graph via disjoint paths is one of the best studied subjects in graph theory and graph algorithms, where NETWORK FLOW and DISJOINT PATHS and many of their variants are among the most well-known problems. In this paper, we introduce, motivate, and study a natural network problem, which we call VECTOR CONNECTIVITY. Given a graph $G = (V, E)$ and a vector \mathbf{k} indexed by the vertices of G , such that $\mathbf{k} = (k_v : v \in V)$ and k_v is between 0 and the degree of v for each vertex $v \in V$, the task of VECTOR CONNECTIVITY is to find a set $S \subseteq V$ of minimum cardinality that satisfies the following: every vertex v of G is either in S or is connected to at least k_v vertices of S via paths that pairwise intersect in no other vertex than v .

In VECTOR CONNECTIVITY there is no restriction on the lengths of the involved disjoint paths. If each path is restricted to be of length exactly 1, we get the well-known VECTOR DOMINATION problem; this problem was introduced by Harant et al. [10] as a generalization of the classical problems DOMINATING SET and VERTEX COVER. The DOMINATING SET problem and its variants have been studied extensively, as they naturally appear in a wide variety of theoretical and practical applications. This has led to a vast amount of papers and several books on domination, e.g., [11, 12]. DOMINATING SET and hence VECTOR DOMINATION are also among the toughest NP-hard problems as they remain NP-hard on various classes of graphs, such as planar graphs of maximum degree 3, bipartite graphs, and most interesting for our study: split graphs [6, 12]. The popularity and the difficulty of these domination problems, the connection between VECTOR DOMINATION and VECTOR CONNECTIVITY, and the question whether allowing paths of unbounded length rather than direct edges or bounded-length paths can result in tractability, are among the motivations for studying the VECTOR CONNECTIVITY problem.

Chlebík and Chlebíková [1] showed that DOMINATING SET, and consequently VECTOR DOMINATION, cannot be approximated in polynomial time within a factor of $(1 - \epsilon) \ln n$ for any constant $\epsilon > 0$ unless $\text{NP} \subseteq \text{DTIME}(n^{O(\log \log n)})$, even when restricted to the class of bipartite graphs or split graphs. On the positive side, Cicalese et al. [3] presented a greedy algorithm for VECTOR DOMINATION with approximation factor $\ln(2\Delta) + 1$, where Δ denotes the maximum degree of the input graph. Moreover, they showed that the problem can be solved in polynomial time on trees and cographs. If one asks for disjoint paths of bounded length rather than direct edges, it is not known in general whether the problem can be approximated within a factor of $O(\log n)$. This gives another motivation to study the unbounded-length paths case, which is exactly the VECTOR CONNECTIVITY problem.

In this paper, we show that VECTOR CONNECTIVITY can be approximated within a factor of $\ln n + 2$ in polynomial time on general graphs, which we find interesting due to the known and unknown approximation results mentioned above. Furthermore, we show that VECTOR CONNECTIVITY can be solved in polynomial time on split graphs, cographs, and trees. We find in particular the tractability result on split graphs surprising, as it is in contrast with the aforementioned NP-hardness and inapproximability results for the DOMINATING SET problem on split graphs. Furthermore, these intractability results imply that if paths are required to be of length at most an input bound then the problem remains NP-hard on split graphs. However, split graphs do not have any induced paths of length 4 or more. Hence our positive result on split graphs implies that the bounded-length path version of VECTOR CONNECTIVITY, which is a generalization of VECTOR DOMINATION, is solvable in polynomial time on split graphs if

the bound is at least 3.

Note that the classes of split graphs, cographs, and trees are all subclasses of perfect graphs, but they are not contained in each other. They form some of the most studied graph classes on which many algorithms have been given, and they play the main role in several books, e.g., in the monograph on perfect graphs by Golumbic [7], and in the monograph by Mahadev and Peled [15] on threshold graphs, which form a subclass of both split graphs and cographs.

Before we proceed to the technical part presenting and proving our results, we end this section by mentioning another motivation, which comes from information propagation in social networks. One famous problem of this type is TARGET SET SELECTION (see, e.g., [2, 14, 16]), where every vertex v has a threshold t_v such that v gets activated if at least t_v of its neighbors are activated, and the task is to select a minimum cardinality vertex subset that results in the activation of all vertices eventually. The practical application behind this problem is the desire by manufacturers to give away their products to a selected small group of people, based on the scenario that every potential customer will decide to buy the product if he or she has enough friends who possess the product. Another possible scenario can be that every potential customer will decide to buy the product only if he or she has enough independent ways to learn about the product. VECTOR CONNECTIVITY fits into this scenario if we assume that information spreads freely along the paths of the network.

2 Definitions and Notation

Unless otherwise stated, we work with undirected simple graphs $G = (V, E)$, where V is the set of vertices, E is the set of edges, and $|V|$ is denoted by n . We use standard graph terminology. In particular, the degree of a vertex v in G is denoted by $d_G(v)$, the maximum degree of a vertex in G is denoted by $\Delta(G)$, and $V(G)$ refers to the vertex set of G . For a given rooted tree T , we write T_v to denote the subtree rooted at vertex v , including vertex v .

Given a graph $G = (V, E)$, a set $S \subseteq V$ and a vertex $v \in V \setminus S$, a v - S fan of order k is a collection of k paths P_1, \dots, P_k such that (1) every P_i is a path connecting v to a vertex of S , and (2) the paths are pairwise vertex-disjoint except at v , i.e., for all $1 \leq i < j \leq k$, it holds that $V(P_i) \cap V(P_j) = \{v\}$. Given an integer-valued vector $\mathbf{k} = (k_v : v \in V)$ with $k_v \in \{0, 1, \dots, d_G(v)\}$ for every $v \in V$, a *vector connectivity set* for (G, \mathbf{k}) is a set $S \subseteq V$ such that there exists a v - S fan of order k_v for every $v \in V \setminus S$. We say that k_v is the *requirement* of vertex v . The minimum size of a vector connectivity set for (G, \mathbf{k}) is denoted by $\kappa(G, \mathbf{k})$.

The VECTOR CONNECTIVITY problem is the problem of finding a vector connectivity set of minimum size, and can be formally stated as follows:

VECTOR CONNECTIVITY

Input: A graph $G = (V, E)$ and a vector $\mathbf{k} = (k_v : v \in V) \in \mathbb{Z}_+^V$ with $k_v \in \{0, 1, \dots, d_G(v)\}$ for all $v \in V$.

Task: Find a vector connectivity set for (G, \mathbf{k}) of size $\kappa(G, \mathbf{k})$.

For every $v \in V$ and every set $S \subseteq V \setminus \{v\}$, we say that v is k -connected to S if there is a v - S fan of order k in G . Hence, given an instance (G, \mathbf{k}) of VECTOR CONNECTIVITY, a set $S \subseteq V$ is a vector connectivity set for (G, \mathbf{k}) if and only if every $v \in V \setminus S$ is k_v -connected to S .

For a subset $A \subseteq V$, we write $\mathbf{k}|_A$ to denote the sub-vector of \mathbf{k} indexed by elements of A , and $\mathbf{1}_A$ denotes the all-one vector indexed by elements of A . We let $\sigma(v, A)$ to denote the maximum order of a v - A fan in G . In other words, $\sigma(v, A) = \max\{s \mid v \text{ is } s\text{-connected to } A\}$. Let $B \subseteq A$, let \mathcal{A} be a v - A fan and let \mathcal{B} be a v - B fan. We say that \mathcal{A} *contains* \mathcal{B} if the collection of paths in \mathcal{B} is a subcollection of the paths in \mathcal{A} .

A set of vertices in a graph is a *clique* if they are all pairwise adjacent, and it is an *independent set* if no two of them are adjacent. A graph is a *split graph* if its vertex set can be partitioned into a clique C and an independent set I , where (C, I) is called a *split partition* of G . Split graphs can be recognized and a split partition can be computed in linear time [9].

For two vertex-disjoint graphs G_1 and G_2 , $G_1 \oplus G_2$ denotes the *disjoint union* of G_1 and G_2 , i.e., $G_1 \oplus G_2 = (V(G_1) \cup V(G_2), E(G_1) \cup E(G_2))$, and $G_1 \otimes G_2$ denotes the *join* of G_1 and G_2 , i.e., the graph obtained by adding to $G_1 \oplus G_2$ all edges of the form $\{uv \mid u \in V(G_1), v \in V(G_2)\}$. The class of *cographs* is defined recursively through the following operations: a single vertex is a cograph; if G_1 and G_2 are vertex-disjoint cographs, then $G_1 \oplus G_2$ is a cograph; if G_1 and G_2 are vertex-disjoint cographs, then $G_1 \otimes G_2$ is a cograph. Cographs, split graphs, and trees are not related to each other inclusion-wise.

A well-known characterization of cographs is via cotrees. A *cotree* T of a cograph G is a rooted tree with two types of interior nodes, \oplus -nodes and \otimes -nodes, that has the following property: there is a bijection between the vertices of G and the leaves of T such that two vertices u and v are adjacent in G if and only if the lowest common ancestor of the leaves u and v in T is a \otimes -node. In particular, every node t of T corresponds to an induced subgraph of G , which is the disjoint union or the join of the subgraphs of G corresponding to the children of t . A graph is a cograph if and only if it has a cotree [4]. Cographs can be recognized and a cotree can be generated in linear time [5, 8]. For our purposes, it is convenient to use the binary version of a cotree, which is commonly used for algorithms on cographs: the recursive definition of cographs implies that we can assume the cotree to be binary. We will call this a *nice* cotree. Clearly, given a cotree of a cograph, a nice cotree can be obtained in linear time.

3 A Polynomial-Time Approximation Algorithm

In this section, we show that VECTOR CONNECTIVITY can be approximated in polynomial time by a factor of $\ln n + 2$ on all graphs. We will achieve this by showing that VECTOR CONNECTIVITY can be recast as a particular case of the well-known MINIMUM SUBMODULAR COVER problem, which will allow us to apply a classical approximation result due to Wolsey [22].

First, we recall some definitions and results about submodular functions, hypergraphs and matroids that we will use in our proofs (see, e.g., [20]). Given a finite set U , a function $g : 2^U \rightarrow \mathbb{Z}_+$ is *submodular* if for every $X, Y \subseteq U$ with $X \subseteq Y$ and every $x \in U \setminus Y$, we have that $g(Y \cup \{x\}) - g(Y) \leq g(X \cup \{x\}) - g(X)$. An instance of the (unweighted) MINIMUM SUBMODULAR COVER problem consists of a set U and an integer-valued, non-decreasing, submodular function $g : 2^U \rightarrow \mathbb{Z}_+$. The objective is to pick a set $S \subseteq U$ of minimum cardinality such that $g(S) = g(U)$.

A *hypergraph* is a pair $H = (U, \mathcal{E})$ where U is a finite set of *vertices* and \mathcal{E} is a set of subsets of U , called *hyperedges*. A *matroid* is a hypergraph $M = (U, \mathcal{F})$ such that \mathcal{F} is nonempty and closed under taking subsets, and its elements, called *independent sets*, satisfy the following

“exchange property”: for every two independent sets A and B such that $|A| < |B|$, there exists an element of B whose addition to A results in a larger independent set. (This is just one of the many equivalent ways to define matroids, see, e.g., [17, 21].) Given a matroid $M = (U, \mathcal{F})$, the *rank function* of M is the function $r_M : 2^U \rightarrow \mathbb{Z}_+$ that assigns to every subset S of U the maximum size of an independent set contained in S . The following property of rank functions of matroids is well known (see, e.g., [20]).

Lemma 1. *For every matroid M , its rank function r_M is submodular.*

A *gammoid* is a hypergraph $\Gamma = (U, \mathcal{E})$ derived from a triple (D, S, T) where $D = (V, A)$ is a digraph and $S, T \subseteq V$, such that $U = S$ and a subset S' of S forms a hyperedge if and only if there exist $|S'|$ vertex-disjoint directed paths in D connecting S' to a subset of T .

Lemma 2 ([18, 19]). *Every gammoid is a matroid.*

For any instance $(G = (V, E), \mathbf{k})$ of VECTOR CONNECTIVITY, we define a function $f : 2^V \rightarrow \mathbb{Z}_+$ as follows:

$$f(X) = \sum_{v \in V} f_v(X), \text{ where } X \subseteq V, \text{ and} \tag{1}$$

$$f_v(X) = \begin{cases} \min\{\sigma(v, X), k_v\} & \text{if } v \notin X; \\ k_v & \text{if } v \in X. \end{cases}$$

Observe that a set $S \subseteq V$ satisfies $f(S) = f(V)$ if and only if S is a vector connectivity set for (G, \mathbf{k}) . Consequently, Lemma 3 below immediately implies that VECTOR CONNECTIVITY is a special case of MINIMUM SUBMODULAR COVER.

Lemma 3. *Let $(G = (V, E), \mathbf{k})$ be an instance of VECTOR CONNECTIVITY. Then the function $f : 2^V \rightarrow \mathbb{Z}_+$, given by (1), satisfies the following properties:*

- (i) $f(\emptyset) = 0$;
- (ii) f is integer-valued, i.e., $f(X) \in \mathbb{Z}_+$ for every $X \subseteq V$;
- (iii) f is non-decreasing, i.e., $f(X) \leq f(Y)$ whenever $X \subseteq Y \subseteq V$;
- (iv) f is submodular.

Proof. It is easy to verify that properties (i)–(iii) hold. In order to show that f is submodular, it suffices to show that all the functions $f_v(\cdot)$ are submodular, that is, that for all $X \subseteq Y \subseteq V$ and for all $w \in V \setminus Y$,

$$f_v(Y \cup \{w\}) - f_v(Y) \leq f_v(X \cup \{w\}) - f_v(X). \tag{2}$$

Suppose first that $f_v(Y) = k_v$. Then $f_v(Y \cup \{w\}) = k_v$ and the left-hand side of inequality (2) is equal to 0. Hence inequality (2) holds since f_v is non-decreasing.

Now suppose that $f_v(Y) < k_v$, which implies that $f_v(Y) = \sigma(v, Y)$. If $f_v(X \cup \{w\}) = k_v$, then $f_v(Y \cup \{w\}) = k_v$ and inequality (2) holds due to the fact that f_v is non-decreasing. In what follows, we assume that $f_v(X \cup \{w\}) < k_v$, which implies that $f_v(X \cup \{w\}) = \sigma(v, X \cup \{w\})$, and also that $v \neq w$. Since $\sigma(v, Y \cup \{w\}) \leq \sigma(v, Y) + 1 \leq k_v$, we have, by the definition of $f_v(\cdot)$

and using the fact that $v \notin Y \cup \{w\}$, the equality $f_v(Y \cup \{w\}) = \sigma(v, Y \cup \{w\})$. Since f_v is non-decreasing, $f_v(Y) < k_v$ implies that $f_v(X) < k_v$, and hence $f_v(X) = \sigma(v, X)$. Inequality (2) then simplifies to

$$\sigma(v, Y \cup \{w\}) - \sigma(v, Y) \leq \sigma(v, X \cup \{w\}) - \sigma(v, X). \quad (3)$$

Hence, in order to prove Lemma 3, it suffices to show that inequality (3) holds for any fixed vertex $v \in V$, i.e., that the function $g_v : 2^{V \setminus \{v\}} \rightarrow \mathbb{Z}_+$, defined by $g_v(W) = \sigma(v, W)$ for all $W \subseteq V \setminus \{v\}$, is submodular. Consider the gammoid Γ derived from the triple $(D, V \setminus \{v\}, N_G(v))$ where D is the digraph obtained from G by replacing each edge with a pair of oppositely directed arcs. By Lemma 2, Γ is a matroid. It follows directly from the definition that function g_v is equal to the rank function r_Γ of Γ . Therefore, by Lemma 1, the function g_v is submodular, which completes the proof of Lemma 3. \square

Theorem 1. VECTOR CONNECTIVITY can be approximated within a factor of $\ln n + 2$ in polynomial time.

Proof. Let $(G = (V, E), \mathbf{k})$ be an instance of VECTOR CONNECTIVITY with $|V| = n$. From the definition of the function f , given by (1), it follows that a set $S \subseteq V$ satisfies $f(S) = f(V)$ if and only if S is a vector connectivity set for (G, \mathbf{k}) . Hence, an optimal solution to the VECTOR CONNECTIVITY problem is provided by a minimum size subset $S \subseteq V$ such that $f(S) = f(V)$, i.e., by an optimal solution for MINIMUM SUBMODULAR COVER. An approximation to such a set S can be found in the following way.

Let \mathbb{A} denote the natural greedy strategy which starts with $S = \emptyset$ and iteratively adds to S the element $v \in V \setminus S$ such that $f(S \cup \{v\}) - f(S)$ is maximum, until $f(S) = f(V)$ is achieved. The maximum order of a v - S fan can be computed in polynomial time using an easy reduction to the well-known MAXIMUM FLOW problem, and thus the function f is polynomially computable. Therefore, the greedy strategy can be implemented in polynomial time. Moreover, Wolsey [22] proved that if f satisfies the four properties listed in Lemma 3, then algorithm \mathbb{A} is an $H(\tau)$ -approximation algorithm for MINIMUM SUBMODULAR COVER, and consequently for VECTOR CONNECTIVITY, where $H(j) = \sum_{i=1}^j \frac{1}{i}$ denotes the j -th harmonic number, and $\tau = \max_{y \in V} f(\{y\}) - f(\emptyset)$. For every $y \in V$, we have

$$f(\{y\}) = \sum_{v \in V \setminus \{y\}} f_v(\{y\}) + f_y(\{y\}) \leq n - 1 + k_y \leq n + \Delta(G).$$

Since $f(\emptyset) = 0$, this implies $\tau \leq n + \Delta(G)$. Hence, algorithm \mathbb{A} is an $H(n + \Delta(G))$ -approximation algorithm for VECTOR CONNECTIVITY. Since $H(n) \leq \ln n + 1$ for $n \geq 1$, we can further bound the approximation ratio ρ of \mathbb{A} from above as follows:

$$\rho \leq H(n + \Delta(G)) \leq \ln(n + \Delta(G)) + 1 \leq \ln(2n) + 1 = \ln n + \ln 2 + 1 \leq \ln n + 2,$$

yielding the desired result. \square

4 A Polynomial-Time Algorithm for Split Graphs

Recall that the VECTOR DOMINATION problem on split graphs is both NP-hard and hard to approximate within a factor of $(1 - \epsilon) \ln n$ for any constant $\epsilon > 0$. In this section, we give a

polynomial-time algorithm to solve the VECTOR CONNECTIVITY problem on split graphs. Our algorithm is based on the following lemma.

Lemma 4. *Let (G, \mathbf{k}) be an instance of VECTOR CONNECTIVITY, where G is a split graph. Let S be any set of vertices in G such that $k_u \geq k_v$ for every pair of vertices $u \in S$ and $v \in V(G) \setminus S$. Then there exists a v - S fan of order $\min\{k_v, |S|\}$ for every $v \in V(G) \setminus S$.*

Proof. Let (C, I) be a split partition of $G = (V, E)$, and for convenience let $S_I = S \cap I$ and $S_C = S \cap C$. We will call the vertices of $V \setminus S$ *free vertices*. Let v be a free vertex of G . We first show that every vertex $u \in S_I$ has at least $k_v - |S_C|$ free neighbors. To see this, let $u \in S_I$. It is obvious that u has at least $d_G(u) - |S_C|$ free neighbors. Since $u \in S$ and $v \in V \setminus S$, we have that $k_v \leq k_u$. This, together with the assumption that $k_v \leq d_G(v)$ for every $v \in V$, implies that u has at least $d_G(u) - |S_C| \geq k_u - |S_C| \geq k_v - |S_C|$ free neighbors.

Suppose that v is a vertex of C . Every vertex of S_C is a neighbor of v , and thus v is $\min\{k_v, |S_C|\}$ -connected to S_C . If $k_v \leq |S_C|$ then the lemma follows, so assume that $k_v > |S_C|$. Recall that every vertex $u \in S_I$ has at least $k_v - |S_C|$ free vertices in its neighborhood. Let $S' \subseteq S_I$ be any subset of S_I such that $|S'| = \min\{k_v - |S_C|, |S_I|\}$. Let G' be the bipartite subgraph of G obtained from the subgraph of G induced by $S' \cup (N_G(S') \setminus S_C)$ by deleting all edges of the form $\{xy \mid x, y \in N_G(S')\}$. Since $|S''| \leq |N_{G'}(S'')|$ for every subset $S'' \subseteq S'$, Hall's Theorem implies that there is a matching M in G' that saturates S' . Let Y be the set of endpoints of M that are not in S' . Then $Y \subseteq C$, and it is possible that $v \in Y$. Since both v and all the vertices of Y belong to the clique C , v can reach at least $|S'| = \min\{k_v - |S_C|, |S_I|\}$ vertices of S_I via disjoint paths that do not contain vertices of S_C , using the edges of M . Consequently, v is $\min\{k_v, |S|\}$ -connected to S , and the lemma follows.

Suppose now that v is a vertex of I . Since $k_v \leq d_G(v)$, v is $\min\{k_v, |S_C|\}$ -connected to S_C . Let \mathcal{P}_C be a v - S_C fan of order $\min\{k_v, |S_C|\}$ that is of smallest total path length. In particular, every path in \mathcal{P}_C is of length 1 or 2. If $k_v \leq |S_C|$ then the lemma follows, so assume that $k_v > |S_C|$. In this case, exactly $|S_C|$ neighbors of v are used by the paths in \mathcal{P}_C . However, v has at least $d_G(v) - |S_C| \geq k_v - |S_C|$ additional neighbors that are free vertices in C . Furthermore, we already proved that every $u \in S_I$ has at least $k_v - |S_C|$ free vertices in its neighborhood. Each such vertex is either a neighbor of v or a neighbor of a neighbor of v . Thus v can reach at least $\min\{k_v - |S_C|, |S_I|\}$ vertices of S_I via disjoint paths that intersect each other and the paths of \mathcal{P}_C only in vertex v . This shows that v is $\min\{k_v, |S|\}$ -connected to S , and the lemma follows. \square

Lemma 4 implies that we can sort the vertices of G by their \mathbf{k} -values in non-increasing order, and greedily pick vertices from the start of the sorted list to be in S until we have a vector connectivity set. This is formalized in the proof of the following theorem.

Theorem 2. VECTOR CONNECTIVITY can be solved in polynomial time on split graphs.

Proof. Let (G, \mathbf{k}) be an instance of VECTOR CONNECTIVITY, where $G = (V, E)$ is a split graph with split partition (C, I) . Given a subset $S \subseteq V$, we define $n(v) = k_v - \max\{t \mid v \text{ is } t\text{-connected to } S\}$ for every vertex $v \in V \setminus S$. Then, for every $v \in V \setminus S$, $n(v) \leq 0$ if and only if v is k_v -connected to S . In particular, if $n(v) > 0$ for some vertex $v \in V \setminus S$, then there exists a v - S fan of order $k_v - n(v)$ in G but no v - S fan of order at least $k_v - n(v) + 1$. Clearly, S is a vector connectivity set for (G, \mathbf{k}) if and only if $n(v) \leq 0$ for every $v \in V \setminus S$.

Lemma 4 suggests the following algorithm. We start by sorting the vertices of G as v_1, v_2, \dots, v_n so that $k_{v_i} \geq k_{v_j}$ for every $1 \leq i < j \leq n$. We set $S = \emptyset$ and $i = 0$. Then, as long as there are vertices $v \in V \setminus S$ such that $n(v) > 0$, we increase i by 1 and add vertex v_i to S . If $n(v) \leq 0$ for every vertex $v \in V \setminus S$, then the algorithm outputs the set S and terminates.

Let us prove that the algorithm is correct. Since the algorithm stops only when $n(v) \leq 0$ for every vertex $v \in V \setminus S$, the set S output by the algorithm is a vector connectivity set for (G, \mathbf{k}) . It remains to show that there is no vector connectivity set S' for (G, \mathbf{k}) such that $|S'| < |S|$. As a result of Lemma 4, every time we add a vertex u to S , $n(v)$ does not increase for any vertex $v \in V \setminus S$, and $n(v)$ decreases by exactly 1 for every vertex $v \in V \setminus S$ for which $n(v) > 0$ before u was added to S . This implies that when the algorithm terminates, there is a vertex $v \in V \setminus S$ such that $k_v = |S|$. Consequently, any vector connectivity set for (G, \mathbf{k}) must either contain at least $|S|$ vertices or contain v . Let S' be a vector connectivity set for (G, \mathbf{k}) and assume, for contradiction, that $|S'| < |S|$. Then S' contains v by the above arguments, and there exists a vertex $u \in S \setminus S'$. Since $k_u \geq k_v = |S|$ and u is k_u -connected to S' , we find that S' must contain at least $|S|$ vertices, contradicting the assumption that $|S'| < |S|$. This finishes the correctness proof of the algorithm. It is clear that the algorithm runs in polynomial time. \square

5 A Polynomial-Time Algorithm for Cographs

In this section we show that VECTOR CONNECTIVITY can be solved in polynomial time on cographs. We will in fact solve the following more general variant of VECTOR CONNECTIVITY. For a graph $G = (V, E)$, an integer-valued vector $\mathbf{k} = (k_v : v \in V)$, and an integer ℓ , we say that a set $S \subseteq V$ is a *vector connectivity set for (G, \mathbf{k}, ℓ)* if S is a vector connectivity set for (G, \mathbf{k}) such that $v \in S$ whenever $k_v \geq \ell$. Let us denote by $\kappa(G, \mathbf{k}, \ell)$ the minimum size of a vector connectivity set for (G, \mathbf{k}, ℓ) . Since $S = V$ is a vector connectivity set for (G, \mathbf{k}, ℓ) , the above parameter is well defined and satisfies $\kappa(G, \mathbf{k}, \ell) \leq |V|$. Clearly, the following relation holds, and hence solving the described variant indeed also solves VECTOR CONNECTIVITY.

Lemma 5. $\kappa(G, \mathbf{k}) = \kappa(G, \mathbf{k}, \max_{v \in V} k_v + 1)$.

In order to simplify the presentation of our algorithm, we assume in this section that in the input to the VECTOR CONNECTIVITY problem and its variant mentioned above, requirements k_v are allowed to be negative. If $k_v < 0$, no condition is imposed on vertex v , and it can be treated the same as if $k_v = 0$.

The first lemma below is an easy observation.

Lemma 6. Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two graphs such that $V_1 \cap V_2 = \emptyset$, and let $G = G_1 \oplus G_2$. Then it holds that

$$\kappa(G, \mathbf{k}, \ell) = \kappa(G_1, \mathbf{k}|_{V(G_1)}, \ell) + \kappa(G_2, \mathbf{k}|_{V(G_2)}, \ell).$$

Lemma 7. Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two graphs such that $V_1 \cap V_2 = \emptyset$, and let $G = G_1 \otimes G_2$. Let $n_1 = |V_1|$ and $n_2 = |V_2|$, and let $\mathcal{F} = \{0, 1, \dots, n_1\} \times \{0, 1, \dots, n_2\}$. For every integer ℓ , it holds that

$$\kappa(G, \mathbf{k}, \ell) = \min_{(i,j) \in \mathcal{F}} f(i, j)$$

with

$$f(i, j) = \max \left\{ \kappa \left(G_1, \mathbf{k}^{1ij}, \ell_1^{ij} \right), i \right\} + \max \left\{ \kappa \left(G_2, \mathbf{k}^{2ij}, \ell_2^{ij} \right), j \right\},$$

where

- $\mathbf{k}^{1ij} = \mathbf{k}|_{V_1} - \min\{i + j, n_2\} \cdot \mathbf{1}_{V_1}$,
- $\mathbf{k}^{2ij} = \mathbf{k}|_{V_2} - \min\{i + j, n_1\} \cdot \mathbf{1}_{V_2}$,
- $\ell_1^{ij} = \min\{\ell, i + j + 1\} - \min\{i + j, n_2\}$,
- $\ell_2^{ij} = \min\{\ell, i + j + 1\} - \min\{i + j, n_1\}$.

Proof. First, we show that $\min_{(i,j) \in \mathcal{F}} f(i,j) \leq \kappa(G, \mathbf{k}, \ell)$. Let $S \subseteq V$ be a minimum vector connectivity set for (G, \mathbf{k}, ℓ) , that is, $|S| = \kappa(G, \mathbf{k}, \ell)$, S is a vector connectivity set for (G, \mathbf{k}) , and $v \in S$ for all $v \in V$ such that $k_v \geq \ell$. Let $S_i = S \cap V_i$, for $i = 1, 2$, and let $I = |S_1|$ and $J = |S_2|$.

Claim 1. $\kappa(G_1, \mathbf{k}^{1IJ}, \ell_1^{IJ}) \leq I$

In order to prove Claim 1, it is enough to argue that S_1 is a vector connectivity set for $(G_1, \mathbf{k}^{1IJ}, \ell_1^{IJ})$. Equivalently, it suffices to prove that:

- (i) for every $v \in V_1 \setminus S_1$, there exists a v - S_1 fan in G_1 of order $k_v^{1IJ} = k_v - \min\{I + J, n_2\}$, and
- (ii) for every $v \in V_1$ such that

$$k_v^{1IJ} \geq \ell_1^{IJ} = \min\{\ell, I + J + 1\} - \min\{I + J, n_2\},$$

it holds that $v \in S_1$.

Let us first show that condition (i) follows from the fact that S is a vector connectivity set for (G, \mathbf{k}) . Indeed, for every $v \in V_1 \setminus S_1 = V_1 \setminus S$, there exists a v - S fan of order k_v in G . Let \mathcal{P} be a v - S fan of maximum order in G that minimizes the number of paths entirely contained in G_1 , and, subject to this condition, is of smallest total path length. Then \mathcal{P} contains all J one-edge paths of the form (v, x) where $x \in S_2$. Moreover, \mathcal{P} contains as many two-edge paths of the form (v, x, y) as possible, where $x \in V_2 \setminus S_2$ and $y \in S_1$; this number of paths is equal to $\min\{I, n_2 - J\}$. Every other path of \mathcal{P} is entirely contained in G_1 . Hence, the number of paths in \mathcal{P} entirely contained in G_1 is equal to $|\mathcal{P}| - J - \min\{I, n_2 - J\}$, which is at least $k_v - \min\{I + J, n_2\} = k_v^{1IJ}$. This shows that condition (i) holds.

To show (ii), suppose that $v \in V_1$ is a vertex with $k_v^{1IJ} \geq \ell_1^{IJ}$, that is,

$$k_v - \min\{I + J, n_2\} \geq \min\{\ell, I + J + 1\} - \min\{I + J, n_2\},$$

or, equivalently,

$$k_v \geq \min\{\ell, I + J + 1\}. \quad (4)$$

We want to show that $v \in S_1$. Suppose, for contradiction, that $v \in V_1 \setminus S_1$. Since S is a vector connectivity set for (G, \mathbf{k}, ℓ) and $v \notin S$, we have that $k_v \leq \ell - 1$ and $k_v \leq |S| = I + J$. Therefore, $k_v \leq \min\{\ell - 1, I + J\} = \min\{\ell, I + J + 1\} - 1$, in contradiction with (4). This shows that $v \in S_1$.

With an analogous argument, one can prove the following claim.

Claim 2. $\kappa(G_2, \mathbf{k}^{2IJ}, \ell_2^{IJ}) \leq J$

Claims 1 and 2 together imply that $f(I, J) = I + J$, and consequently $\min_{(i,j) \in \mathcal{F}} f(i, j) \leq f(I, J) = I + J = |S| = \kappa(G, \mathbf{k}, \ell)$.

Now we show that $\kappa(G, \mathbf{k}, \ell) \leq \min_{(i,j) \in \mathcal{F}} f(i, j)$. Let (I, J) be a pair from \mathcal{F} such that $f(I, J) = \min_{(i,j) \in \mathcal{F}} f(i, j)$. Write $\min_{(i,j) \in \mathcal{F}} f(i, j) = m_1 + m_2$, where

$$m_1 = \max \left\{ \kappa(G_1, \mathbf{k}^{1IJ}, \ell_1^{IJ}), I \right\}$$

and

$$m_2 = \max \left\{ \kappa(G_2, \mathbf{k}^{2IJ}, \ell_2^{IJ}), J \right\}.$$

Let S'_1 and S'_2 be minimum vector connectivity sets for $(G_1, \mathbf{k}^{1IJ}, \ell_1^{IJ})$ and $(G_2, \mathbf{k}^{2IJ}, \ell_2^{IJ})$, respectively. For $i = 1, 2$, let $S_i \subseteq V_i$ be a set of size m_i containing S'_i . Set $S = S_1 \cup S_2$. To show that $\kappa(G, \mathbf{k}, \ell) \leq \min_{(i,j) \in \mathcal{F}} f(i, j)$, it suffices to show that S is a vector connectivity set for (G, \mathbf{k}, ℓ) . Equivalently, it suffices to prove that:

- (I) for every $v \in V \setminus S$, there exists a v - S fan in G of order k_v , and
- (II) for every $v \in V$ such that $k_v \geq \ell$, it holds that $v \in S$.

Let $v \in V \setminus S$. Suppose that $v \in V_1 \setminus S_1$. Let \mathcal{P} be a v - S_1 fan in G_1 of order $k_v^{1IJ} = k_v - \min\{I + J, n_2\}$. Then, \mathcal{P} is also a v - S fan in G . Extend \mathcal{P} to a larger v - S fan in G by adding to it all m_2 one-edge paths of the form (v, x) where $x \in S_2$, and as many two-edge paths of the form (v, x, y) as possible, where $x \in V_2 \setminus S_2$ and $y \in S_1$; the number of these two-edge paths is equal to $\min\{n_2 - m_2, m_1 - k_v^{1IJ}\}$. The total number of paths in the so constructed v - S fan is equal to

$$\begin{aligned} |\mathcal{P}| + m_2 + \min\{n_2 - m_2, m_1 - k_v^{1IJ}\} &= \min\{k_v^{1IJ} + n_2, m_1 + m_2\} \\ &\geq \min\{k_v^{1IJ} + n_2, I + J\}. \end{aligned}$$

We would like to show that $\min\{k_v^{1IJ} + n_2, I + J\} \geq k_v$, or, equivalently, that

$$k_v^{1IJ} + n_2 \geq k_v \tag{5}$$

and

$$I + J \geq k_v. \tag{6}$$

Inequality (5) follows directly from the definition of k_v^{1IJ} . Since S_1 is a vector connectivity set for $(G_1, \mathbf{k}^{1IJ}, \ell_1^{IJ})$ and $v \in V_1 \setminus S_1$, it follows that

$$k_v^{1IJ} \leq \ell_1^{IJ} - 1,$$

that is,

$$k_v - \min\{I + J, n_2\} \leq \min\{\ell - 1, I + J\} - \min\{I + J, n_2\},$$

or, equivalently,

$$k_v \leq \min\{\ell - 1, I + J\}.$$

In particular, comparing k_v to the second term in the above minimum yields the desired inequality (6). This shows that for every $v \in V_1 \setminus S_1$, there exists a v - S fan in G of order k_v . We can show similarly that for every $v \in V_2 \setminus S_2$, there exists a v - S fan in G of order k_v . This shows (I).

To show (II), let $v \in V$ such that $k_v \geq \ell$. We need to show that $v \in S$. Without loss of generality, assume that $v \in V_1$. The condition $k_v \geq \ell$ implies that

$$k_v^{IJ} = k_v - \min\{I + J, n_2\} \geq \ell - \min\{I + J, n_2\} \geq \ell_1^{IJ},$$

which implies that $v \in S_1$ due to the fact that S_1 is a vector connectivity set for $(G_1, \mathbf{k}^{IJ}, \ell_1^{IJ})$. Hence, $v \in S$. This shows that condition (II) holds, and hence that $\kappa(G, \mathbf{k}, \ell) \leq \min_{(i,j) \in \mathcal{F}} f(i, j)$. We conclude that $\kappa(G, \mathbf{k}, \ell) = \min_{(i,j) \in \mathcal{F}} f(i, j)$. \square

Theorem 3. VECTOR CONNECTIVITY can be solved in polynomial time on cographs.

Proof. Consider the input $(G = (V, E), \mathbf{k})$ to the VECTOR CONNECTIVITY problem, where G is a cograph and $n = |V|$. By Lemma 5, computing the value of $\kappa(G, \mathbf{k})$ is equivalent to computing the value of $\kappa(G, \mathbf{k}, K)$ with $K = \max_{v \in V} k_v + 1$. We compute this value as follows. First, we compute a nice cotree T of G . We traverse T bottom up, processing a node only after all its children have been processed. When processing a node t of T , we compute all $O(n^2)$ values of

$$\kappa(H, \mathbf{k}|_{V(H)} - i \cdot \mathbf{1}_{V(H)}, \ell), \quad i \in \{0, 1, \dots, n\}, \quad \ell \in \{0, 1, \dots, K\},$$

where H is the induced subgraph of G corresponding to the subtree T_t . For every leaf of the cotree, corresponding to a single vertex v of G , each of the $O(n^2)$ values can be computed in $O(1)$ time as follows:

$$\kappa(\{\{v\}, \emptyset\}, k_v - i, \ell) = \begin{cases} 0 & \text{if } k_v - i \leq \min\{\ell - 1, 0\}, \\ 1 & \text{otherwise.} \end{cases}$$

Depending on whether an internal node t is a \oplus -node or a \otimes -node, we can use Lemma 6 or Lemma 7 to compute each of the $O(n^2)$ values of $\kappa(H, \mathbf{k}|_{V(H)} - i \cdot \mathbf{1}_{V(H)}, \ell)$ in time $O(n^2)$. Hence, each internal node of the modified cotree can be processed in time $O(n^4)$, yielding an overall time complexity of $O(n^5)$, since a cotree has $O(n)$ nodes.

A minimum vector connectivity set can also be computed in the stated time. In addition to the values of $\kappa(H, \mathbf{k}|_{V(H)} - i \cdot \mathbf{1}_{V(H)}, \ell)$ at each node of the cotree, we need to store also a minimum vector connectivity set achieving each of these values. These sets can be computed recursively as follows. For an internal node t with corresponding subgraph H , let H_1 and H_2 denote the subgraphs of G corresponding to the two children of t in T .

- If H corresponds to a leaf of T , then $V(H) = \{v\}$ for some $v \in V$, and a minimum vector connectivity set for $(H, k_v - i, \ell)$ is either empty or $\{v\}$, depending on whether $k_v - i \leq \min\{\ell - 1, 0\}$ or not.
- If $H = H_1 \oplus H_2$, then a minimum vector connectivity set for $(H, \mathbf{k}|_{V(H)} - i \cdot \mathbf{1}_{V(H)}, \ell)$ is given by the union of minimum vector connectivity sets for $(H_1, \mathbf{k}|_{V(H_1)} - i \cdot \mathbf{1}_{V(H)}, \ell)$ and $(H_2, \mathbf{k}|_{V(H_2)} - i \cdot \mathbf{1}_{V(H)}, \ell)$.

- If $H = H_1 \otimes H_2$, then a minimum vector connectivity set S for $(H, \mathbf{k}|_{V(H)} - i \cdot \mathbf{1}_{V(H)}, \ell)$ can be computed in $O(n^2)$ time: first compute a pair (I, J) minimizing the function f defined in Lemma 7 (with H, H_1, H_2 in place of G, G_1, G_2 , respectively), and then take the union of minimum vector connectivity sets S'_1 and S'_2 for $(H_1, \mathbf{k}^{1I}, \ell_1^{IJ})$ and $(H_2, \mathbf{k}^{2IJ}, \ell_2^{IJ})$, together with some extra vertices if necessary so that $|S \cap V(H_1)| \geq I$ and $|S \cap V(H_2)| \geq J$.

Finally, let us remark that all the instances (H, \mathbf{k}', ℓ) for which $\kappa(H, \mathbf{k}', \ell)$ must be evaluated in order to compute the value of $\kappa(G, \mathbf{k}) = \kappa(G, \mathbf{k}, \max_{v \in V} k_v + 1)$ satisfy the property that for all $v \in V(H)$, either $k'_v \leq d_H(v)$ or $k'_v \geq \ell$. This can be proved by induction on the distance of a node t representing H from the root of the cotree T , and assures that the values of $\kappa(H, \mathbf{k}', \ell)$ are well defined for such instances. This completes the proof of Theorem 3. \square

6 A Polynomial-Time Algorithm for Trees

We have seen that VECTOR CONNECTIVITY is solvable in polynomial time on cographs and split graphs. These two graph classes do not contain graphs with long induced paths. In particular, cographs are equivalent to graphs that do not have induced paths of length 3 or more [5], and it is easy to observe that split graphs do not contain induced paths of length 4 or more. In this section, we give a polynomial-time algorithm to solve the VECTOR CONNECTIVITY problem on trees, a graph class that allows the existence of arbitrarily long induced paths.

Theorem 4. VECTOR CONNECTIVITY *can be solved in polynomial time on trees.*

Proof. Let (T, \mathbf{k}) be an instance of VECTOR CONNECTIVITY, where $T = (V, E)$ is a tree. We assume that T has at least two vertices and is rooted at an arbitrary vertex r . Since the requirements of the vertices do not change during the execution of the algorithm, we will simply speak of a vector connectivity set for T_v instead of a vector connectivity set for $(T_v, \mathbf{k}|_{V(T_v)})$, for every $v \in V$.

The idea of the algorithm is to construct a vector connectivity set for T of minimum size, starting from the leaves of T and processing a vertex only after all its children have been processed. At any step of the algorithm, let $S \subseteq V$ be the set of vertices that have thus far been chosen to belong to the solution. For any vertex v of T , we define $S_v = S \cap V(T_v)$. When processing a vertex v , the algorithm computes the values $f(v)$, $n(v)$ and $r(v)$, which are defined as follows. For every vertex $v \in V$, $f(v) = 1$ if the subtree T_v contains at least one vertex of S ; otherwise $f(v) = 0$. The value $r(v)$ denotes the number of children w of v for which $f(w) = 1$. Note that if $f(w) = 1$ for a child w of vertex v , then v is 1-connected but not 2-connected to S_w , regardless of how many vertices S_w contains. Furthermore, v is 1-connected to $S \setminus V(T_v)$ if S contains a vertex outside T_v . We let $n(v)$ denote whether or not a vertex in T_v “needs” an additional path to a vertex outside of T_v , indicated by 1 or 0, for every $v \in V$. More precisely, $n(v) = 0$ if for every vertex $w \in V(T_v)$, there is a w - S_v fan of order k_w in T_v , i.e., every vertex w of T_v , including v itself, is k_w -connected to S_v and hence also to S . On the other hand, $n(v) = 1$ if there is a vertex $w \in V(T_v)$ such that there is a w - S_v fan of order $k_w - 1$ but no w - S_v fan of order k_w in T_v .

We now describe the algorithm in detail. Initially, we set $S = \emptyset$. Let $v \in V$ be a leaf of T . We set $r(v) = 0$. If $k_v = 0$, then we set $f(v) = 0$ and $n(v) = 0$. If $k_v = 1$, then we set $f(v) = 0$ and $n(v) = 1$.

Next, let v be a vertex that is not a leaf and not the root, and assume that the children of v have all been processed. For every child w of v , if $n(w) = 1$ and v has a child $w' \neq w$ such that $f(w') = 1$, then we set $n(w) = 0$. We then compute $r(v)$ by adding up the f -values of all the children of v . If $k_v \leq r(v)$, then we set $n(v) = 1$ if v has a child w with $n(w) = 1$, and we set $n(v) = 0$ otherwise. If $k_v = r(v) + 1$, then we set $n(v) = 1$. If $k_v \geq r(v) + 2$, then we add v to S and set $n(v) = 0$. In each of the above cases, we set $f(v) = 1$ if $r(v) \geq 1$ or if v is added to S , and we set $f(v) = 0$ otherwise.

Finally, let v be the root of T . We set $n(v) = 0$. If $k_v \leq r(v)$, then we perform the following check: if v has a child w such that $n(w) = 1$ and $f(w') = 0$ for every other child $w' \neq w$ of v , then we add v to S . If $k_v \geq r(v) + 1$, then we add v to S . The algorithm outputs the set S and terminates as soon as the root has been processed.

In order to prove the correctness of the algorithm, we prove that, for every $v \in V$, the following three statements are true immediately after v is processed, where p denotes the parent of v in T .

- (i) If $n(v) = 0$, then S_v is a vector connectivity set for T_v .
- (ii) If $n(v) = 1$, then $S_v \cup \{p\}$ is a vector connectivity set for the subtree of T induced by $V(T_v) \cup \{p\}$;
- (iii) There is no vector connectivity set S' for T such that $|S' \cap V(T_v)| < |S_v|$.

Assume first that v is a leaf. Then $d_T(v) = 1$, so $k_v \in \{0, 1\}$. Note that the algorithm does not add v to S , regardless of whether $k_v = 0$ or $k_v = 1$. Hence $S_v = \emptyset$, and statement (iii) trivially holds in both cases. If $k_v = 0$, then we set $n(v) = 0$, and it is clear that all statements are satisfied in this case. If $k_v = 1$, then we set $n(v) = 1$. Since the set $\{p\}$ consisting of the parent of v is a vector connectivity set for the subtree of T induced by $V(T_v) \cup \{p\}$, statement (ii) holds.

Assume now that v is not a leaf and not the root. We again distinguish two cases, depending on whether or not v is added to S by our algorithm. Suppose v is added to S . Then $k_v \geq r(v) + 2$, and we set $n(v) = 0$ in this case. Since v is added to S , we have that every descendant w of v with $n(w) = 1$ becomes k_w -connected to S_v , so statement (i) holds. To see why statement (iii) holds in this case, let S' be any vector connectivity set for T . We know that $|S' \cap V(T_w)| \geq |S_w|$ for every child w of v , since statement (iii) holds for every child w . Hence, if $v \in S'$, then $|S' \cap V(T_v)| \geq |S_v|$ and statement (iii) holds. If $v \notin S'$, then the fact that $k_v \geq r(v) + 2$ implies that S' must contain at least one vertex of a subtree T_w for a child w of v with $f(w) = 0$. Hence we have that $|S' \cap V(T_v)| \geq |S_v|$ and statement (iii) is true also in this case.

Now suppose that v is not added to S . Since statement (iii) holds for all children of v and $v \notin S$, statement (iii) is trivially true for v . Let us argue why statements (i) and (ii) hold. If $n(v) = 0$, then for every child w with $n(w) = 1$ just before v is processed there is another child $w' \neq w$ with $f(w') = 1$. Hence every child w with $n(w) = 1$ just before v is processed is k_w -connected to S , since it has a w - S_v fan of order $k_w - 1$ and one additional path from w via v to a vertex of $S \cap V(T_{w'})$ for some child $w' \neq w$ with $f(w') = 1$; the same holds for each descendant w'' of w with $n(w'') = 1$. Hence S_v is a vector connectivity set for T_v , and statement (i) holds in this case. If $n(v) = 1$ and $k_v \leq r(v)$, then at least one vertex $w \in V(T_v)$ is $(k_w - 1)$ -connected to S_v (in fact, to S_w), but needs a path from w to a vertex of S outside T_v to become k_w -connected to S . Clearly, adding the parent p of v to S_v ensures that all such

vertices w become k_w -connected to $S_v \cup \{p\}$. If $n(v) = 1$ and $k_v = r(v) + 1$, then every vertex w in T_v with $n(w) = 1$ that is $(k_w - 1)$ -connected to S_v becomes k_w -connected to $S_v \cup \{p\}$. In either case, statement (ii) holds.

Finally, assume that v is the root of T . Then we set $n(v) = 0$, so we need to show that statement (i) holds. This is clearly the case if v is added to S , since then every descendant w of v with $n(w) = 1$ becomes k_w -connected to S_v . If v is not added to S , then $k_v \leq r(v)$ and for every child w of v with $n(w) = 1$, there is another child $w' \neq w$ such that $f(w') = 1$ and consequently there is a path from w via v to a vertex of S outside T_w . Hence $S_v = S$ is a vector connectivity set for $T_v = T$, and statement (i) also holds in this case. If v was not added to S , then the validity of statement (iii) immediately follows from the fact that statement (iii) holds for every child of v . Suppose v was added to S , and let S' be any vector connectivity set for T . Since statement (iii) holds for any child w of v , we have that $|S' \cap V(T_w)| \geq |S_w|$. Just like in the case where v is not a leaf and not the root, we have that $|S' \cap V(T_v)| \geq |S_v|$ regardless of whether or not $v \in S'$. Hence statement (iii) is true also in this case.

Since these statements hold for the root of T , the set S constructed by the algorithm is a vector connectivity set for T of minimum size. The observation that all steps of the algorithm can be performed in polynomial time completes the proof of Theorem 4. \square

7 Concluding Remarks

In this paper, we initiated the study of the VECTOR CONNECTIVITY problem, which opens a research path with many interesting questions. The most prominent of these questions is of course the computational complexity of VECTOR CONNECTIVITY on general input graphs. Could it be that the problem is polynomial-time solvable on all graphs, or is its tractability heavily dependent on either the absence of long induced paths or on a tree-like structure of the input graph? On which other graph classes is VECTOR CONNECTIVITY solvable in polynomial time? Does VECTOR CONNECTIVITY admit a polynomial-time constant-factor approximation algorithm on general graphs?

Another interesting variant of the problem can be obtained by allowing the requirement k_v of each vertex v to be arbitrarily large, in which case a vertex v with $k_v > d_G(v)$ is forced to be in every vector connectivity set. Is it perhaps easier to prove this variant to be NP-hard in general? Note that the algorithms given in this paper, except the algorithm for split graphs, work in polynomial time also for this variant. Is this variant polynomial-time solvable on split graphs?

Acknowledgements

We are indebted to Ferdinando Cicalese and Ugo Vaccaro for initial discussions during which the VECTOR CONNECTIVITY problem was discovered and research on it started. We are also grateful to Nicola Apollonio for helpful discussions.

This work is supported by the Research Council of Norway (197548/F20) and by the Slovenian Research Agency (research program P1-0285 and research projects J1-4010, J1-4021, BI-US/12-13-029 and N1-0011: GReGAS, supported in part by the European Science Foundation).

References

- [1] Chlebík, M., Chlebíkova, J.: Approximation hardness of dominating set problems in bounded degree graphs. *Information and Computation* 206, 1264–1275 (2008).
- [2] Chen, N.: On the approximability of influence in social networks. *SIAM Journal on Discrete Mathematics*, 23, 1400–1415 (2009).
- [3] Cicalese, F., Milanič, M., Vaccaro, U.: On the approximability and exact algorithms for vector domination and related problems in graphs, *Discrete Applied Mathematics* (2012), doi:10.1016/j.dam.2012.10.007.
- [4] Corneil, D.G., Lerchs, H., Stewart Burlingham, L.: Complement reducible graphs. *Discrete Applied Mathematics* 3, 163–174 (1981)
- [5] Corneil, D.G., Perl, Y., Stewart, L.K.: A linear recognition algorithm for cographs. *SIAM J. Comput.* 14, 926–934 (1985)
- [6] Garey, M.R., Johnson, D.S.: *Computers and Intractability*. W.H. Freeman and Co., New York (1979)
- [7] Golumbic, M.C.: *Algorithmic Graph Theory and Perfect Graphs*. Second edition, *Annals of Discrete Mathematics* 57, Elsevier (2004)
- [8] Habib, M., Paul, C.: A simple linear time algorithm for cograph recognition. *Discrete Applied Mathematics* 145, 183–197 (2005)
- [9] Hammer, P.L., Simeone, B.: The splittance of a graph. *Combinatorica* 1, 275–284 (1981).
- [10] Harant, J., Prochnowski, A., Voigt, M.: On dominating sets and independent sets of graphs. *Combinatorics, Probability and Computing* 8, 547–553 (1999)
- [11] Haynes, T.W., Hedetniemi, S., Slater, P.: *Fundamentals of Domination in Graphs*. Marcel Dekker, New York (1998)
- [12] Haynes, T.W., Hedetniemi, S., Slater, P.: *Domination in Graphs: Advanced Topics*. Marcel Dekker, New York (1998)
- [13] Boros, E., Heggernes, P., van ’t Hof, P., Milanič, M.: *Vector Connectivity in Graphs*, To appear in *Proceedings of the 10th Annual Conference on Theory and Applications of Models of Computation (TAMC 2013)*, *Lecture Notes in Computer Science*, Springer, 2013.
- [14] Kempe, D., Kleinberg, J., Tardos, E.: Maximizing the spread of influence through a social network. In *Proc. 9th ACM KDD*, 137–146. ACM Press (2003)
- [15] Mahadev, N., Peled, U.: *Threshold graphs and related topics*. *Annals of Discrete Mathematics* 56, North Holland (1995)
- [16] Nichterlein, A., Niedermeier, R., Uhlmann, J., Weller, M.: On tractable cases of target set selection. *Social Network Analysis and Mining* (2012). Online available.
- [17] Oxley, J.: *Matroid Theory*. Oxford University Press, Oxford, 1992.

- [18] Perfect, H.: Applications of Menger's graph theorem. *Journal of Mathematical Analysis and Applications* 22 (1968) 96–111.
- [19] Pym, J.S.: A proof of the linkage theorem. *Journal of Mathematical Analysis and Applications* 27 (1969) 636–638.
- [20] Schrijver, A.: *Combinatorial Optimization. Polyhedra and Efficiency. Volumes A–C. Algorithms and Combinatorics*, 24. Springer-Verlag, Berlin, 2003.
- [21] Welsh, D.J.A.: *Matroid Theory*, L.M.S. Monographs, 8. Academic Press, London-New York, 1976.
- [22] Wolsey, L.A.: An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica* 2, 385—393 (1982)