



Approximate min–max theorems for Steiner rooted-orientations of graphs and hypergraphs [☆]

Tamás Király ^{a,1}, Lap Chi Lau ^b

^a *MTA-ELTE Egerváry Research Group, Department of Operations Research, Eötvös Loránd University, Hungary*

^b *Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong*

Received 24 November 2006

Available online 20 March 2008

Abstract

Given an undirected hypergraph and a subset of vertices $S \subseteq V$ with a specified root vertex $r \in S$, the STEINER ROOTED-ORIENTATION problem is to find an orientation of all the hyperedges so that in the resulting directed hypergraph the “connectivity” from the root r to the vertices in S is maximized. This is motivated by a multicasting problem in undirected networks as well as a generalization of some classical problems in graph theory. The main results of this paper are the following approximate min–max relations:

- Given an undirected hypergraph H , if S is $2k$ -hyperedge-connected in H , then H has a Steiner rooted k -hyperarc-connected orientation.
- Given an undirected graph G , if S is $2k$ -element-connected in G , then G has a Steiner rooted k -element-connected orientation.

Both results are tight in terms of the connectivity bounds. These also give polynomial time constant factor approximation algorithms for both problems. The proofs are based on submodular techniques, and a graph decomposition technique used in the STEINER TREE PACKING problem. Some complementary hardness results are presented at the end.

© 2008 Elsevier Inc. All rights reserved.

Keywords: Hypergraph; Steiner tree; Orientation

[☆] A preliminary version appeared in the proceedings of FOCS 2006.

E-mail addresses: tkiraly@cs.elte.hu (T. Király), chi@cse.cuhk.edu.hk (L.C. Lau).

¹ Research supported by OTKA K60802, ADONET MCR TN 504438, and OMF B-01608/2006.

1. Introduction

Let $H = (V, \mathcal{E})$ be an undirected hypergraph. An *orientation* of H is obtained by assigning a direction to each hyperedge in H . In our setting, a *hyperarc* (a directed hyperedge) is a hyperedge with a designated *tail vertex* and other vertices as *head vertices*. This model has been used in network multicasting [5,24], which is the main motivation for our research. Given a set $S \subseteq V$ of *terminal vertices* (the vertices in $V - S$ are called the *Steiner vertices*) and a *root vertex* $r \in S$, we say a directed hypergraph is *Steiner rooted k -hyperarc-connected* if there are k hyperarc-disjoint paths from the root vertex r to each terminal vertex in S . Here, a *path* in a directed hypergraph is an alternating sequence of distinct vertices and hyperarcs $\{v_0, a_0, v_1, a_1, \dots, a_{k-1}, v_k\}$ so that v_i is the tail of a_i and v_{i+1} is a head of a_i for all $0 \leq i < k$. The STEINER ROOTED-ORIENTATION problem is to find an orientation of H so that the resulting directed hypergraph is Steiner rooted k -hyperarc-connected, and our objective is to maximize k .

When the STEINER ROOTED-ORIENTATION problem specializes to graphs, it is a common generalization of some classical problems in graph theory. When there are only two terminals ($S = \{r, v\}$), it is the edge-disjoint paths problem solved by Menger [26]. When all vertices in the graph are terminals ($S = V$), it can be shown to be equivalent to the edge-disjoint spanning trees problem solved by Tutte [29] and Nash-Williams [28]. An alternative common generalization of the above problems is the STEINER TREE PACKING problem studied in [16,20,21]. Notice that if a graph G has k edge-disjoint *Steiner trees* (i.e. trees that connect the terminal vertices S), then G has a Steiner rooted k arc-connected orientation. The converse, however, is not true. As we shall see, significantly sharper approximate min–max relations and also approximation ratio can be achieved for the STEINER ROOTED-ORIENTATION problem, especially when we consider hyperarc-connectivity and element-connectivity.

Given a hypergraph H , we say S is *k -hyperedge-connected in H* if there are k hyperedge-disjoint paths between every pair of vertices in S . It is not difficult to see that for a hypergraph H to have a Steiner rooted k -hyperarc-connected orientation, S must be at least k -hyperedge-connected in H . The main focus of this paper is to determine the smallest constant c so that the following holds: If S is ck -hyperedge-connected in H , then H has a Steiner rooted k -hyperarc-connected orientation.

1.1. Previous work

Graph orientations is a well-studied subject in the literature, and there are many ways to look at such questions (see [2]). Here we focus on graph orientations achieving high connectivity. In the following $\lambda(x, y)$ denotes the maximum number of edge-disjoint paths from x to y , which is called the *local-edge-connectivity* from x to y . Nash-Williams [27] proved the following deep generalization of Robbins' theorem which achieves optimal local-arc-connectivity for all pairs of vertices:

Every undirected graph G has an orientation D so that $\lambda_D(x, y) \geq \lfloor \lambda_G(x, y)/2 \rfloor$ for all $x, y \in V$.

Nash-Williams' original proof is quite complicated, and until now this is the only known orientation result achieving high *local-arc-connectivity*. Subsequently, Frank, in a series of works [8–11], developed a general framework to solve graph orientation problems achieving high *global-arc-connectivity* by using the *submodular flow* problem. Recently, this framework has

been generalized to solve hypergraph orientation problems achieving high global-hyperarc-connectivity [13].

1.2. Results

The main result of this paper is the following approximate min–max theorem on hypergraphs.

Theorem 1.1. *Suppose H is an undirected hypergraph, S is a subset of terminal vertices with a specified root vertex $r \in S$. Then H has a Steiner rooted k -hyperarc-connected orientation if S is $2k$ -hyperedge-connected in H .*

Theorem 1.1 is best possible in terms of the connectivity bound. This is shown by any $2k$ -regular $2k$ -edge-connected non-complete graph G by setting $S = V(G)$ (e.g. a $2k$ -dimensional hypercube).

The proof of Theorem 1.1 is constructive, and implies a polynomial time constant factor approximation algorithm for the problem. When the above theorem specializes to graphs, this gives a new and simpler algorithm (without using Nash-Williams' orientation theorem) to find a Steiner rooted k -arc-connected orientation in a graph when S is $2k$ -edge-connected in G . On the other hand, we prove that finding an orientation which maximizes the Steiner rooted-arc-connectivity in a graph is NP-complete (Theorem 6.1).

Following the notation on approximation algorithms on graph connectivity problems, by an *element* we mean either an edge or a Steiner vertex. For graph connectivity problems, element-connectivity is regarded as of intermediate difficulty between vertex-connectivity and edge-connectivity (see [7,17]). A directed graph is *Steiner rooted k -element-connected* if there are k element-disjoint directed paths from r to each terminal vertex in S . We prove the following approximate min–max theorem on element-connectivity, which is tight in terms of the connectivity bound. We also prove the NP-completeness of the problem of deciding if there is a Steiner rooted k -element-connected orientation (Theorem 6.4).

Theorem 1.2. *Suppose G is an undirected graph, S is a subset of terminal vertices with a specified root vertex $r \in S$. Then G has a Steiner rooted k -element-connected orientation if S is $2k$ -element-connected in G .*

1.3. Techniques

A key ingredient in the proof of Theorem 1.1 is the use of an “extension property” (see [21,22]) to help decompose a general hypergraph into hypergraphs with substantially simpler structures. Then, in those simpler hypergraphs, we apply submodularity in an effective way to solve the problem (and also prove the extension property). An important building block of our approach is the following class of polynomial time solvable graph orientation problems, which we call the DEGREE-SPECIFIED STEINER ROOTED-ORIENTATION problem.

Theorem 1.3. *Suppose G is an undirected graph, S is a subset of terminal vertices with a specified root vertex $r \in S$, and m is an in-degree specification on the Steiner vertices (i.e. $m : (V(G) - S) \rightarrow \mathbb{Z}^+$). Then deciding whether G has a Steiner rooted k -arc-connected orientation with the specified indegrees can be solved in polynomial time.*

Perhaps Theorem 1.3 does not seem to be very useful at first sight, but it turns out to be surprisingly powerful in some situations when we have a rough idea on what the indegrees of Steiner vertices should be like. To prove Theorem 1.3, we shall reduce this problem to membership in a base polyhedron from which we can also derive a sufficient and necessary condition for the existence of a degree-specified Steiner rooted k -arc-connected orientation. This provides us with a crucial tool in establishing the approximate min–max relations.

2. The basics

Let $H = (V, \mathcal{E})$ be an undirected hypergraph. Given $X \subseteq V$, we say a hyperedge e enters X if $0 < |e \cap X| < |e|$. The rank of H is the cardinality of the largest hyperedge of H . We define $\delta_H(X)$ to be the set of hyperedges that enter X , and $d_H(X) := |\delta_H(X)|$. We also define $i(X)$ to be the number of induced hyperedges in X . In a directed hypergraph $\vec{H} = (V, \vec{\mathcal{E}})$, a hyperarc a enters a set X if the tail of a is not in X and some head of a is in X . We define $\delta_H^{\text{in}}(X)$ to be the set of hyperarcs that enter X , and $d_H^{\text{in}}(X) := |\delta_H^{\text{in}}(X)|$. Similarly, a hyperarc a leaves a set X if a enters $V - X$. We define $\delta_H^{\text{out}}(X)$ to be the set of hyperarcs that leave X , and $d_H^{\text{out}}(X) := |\delta_H^{\text{out}}(X)|$.

Let V be a finite ground set. Two subsets X and Y are intersecting if $X - Y, Y - X, X \cap Y$ are all non-empty. X and Y are crossing if they are intersecting and $X \cup Y \neq V$. For a function $m : V \rightarrow \mathbb{R}$ we use the notation $m(X) := \sum(m(v) : v \in X)$. Let $f : 2^V \rightarrow \mathbb{R}$ be a function defined on the subsets of V . The set-function f is called (intersecting, crossing) submodular if the following inequality holds for any two (intersecting, crossing) subsets X and Y of V :

$$f(X) + f(Y) \geq f(X \cup Y) + f(X \cap Y). \tag{1}$$

The set function f is called (intersecting, crossing) supermodular if the reverse inequality of (1) holds for any two (intersecting, crossing) subsets X and Y of V .

2.1. Base polyhedra and graph orientations

The following two results show the connection between orientation problems and the theory of base polyhedra.

Lemma 2.1. (See [14].) *Let $G = (V, E)$ be an undirected graph, $x : V \rightarrow \mathbb{Z}^+$ an in-degree specification, and $h : 2^V \rightarrow \mathbb{Z}^+$ a non-negative function. Then G has an orientation D such that $d_D^{\text{in}}(X) \geq h(X)$ for all $X \subseteq V$ and $d_D^{\text{in}}(v) = x(v)$ for every $v \in V$ if and only if $x(V) = |E|$ and*

$$x(X) \geq i(X) + h(X) \quad \text{for every } X \subseteq V.$$

Notice that the set function $i(X)$ is supermodular, so if $h(X)$ is intersecting supermodular then so is $i(X) + h(X)$.

Theorem 2.2. (See [23].) *Let $h : 2^V \rightarrow \mathbb{Z}^+$ be a non-negative intersecting supermodular set function, and let l be a non-negative integer. The polyhedron*

$$\mathcal{B} := \{x \in \mathbb{R}^V : x(X) \geq h(X) \text{ for } X \subseteq V, x(V) = l\}$$

is non-empty if and only if the following conditions hold:

1. $h(\emptyset) = 0$,
2. $\sum_{X \in \mathcal{F}} h(X) \leq l$ for every partition \mathcal{F} of V .

If \mathcal{B} is non-empty, then it is a base polyhedron, so its vertices are integral.

2.2. Mader’s splitting-off theorem

Let G be an undirected graph. *Splitting-off* a pair of edges $e = uv$, $f = vw$ means that we replace e and f by a new edge uw (parallel edges may arise). The resulting graph will be denoted by G^{ef} . When a splitting-off operation is performed, the local edge-connectivity never increases. The content of the splitting-off theorem is that under certain conditions there is an appropriate pair of edges $\{e = uv, f = vw\}$ whose splitting-off preserves all local or global edge-connectivity between vertices distinct from v . The following theorem by Mader [25] proves to be very useful in graph connectivity problems.

Theorem 2.3. *Let $G = (V, E)$ be a connected undirected graph in which $0 < d_G(s) \neq 3$ and there is no cut-edge incident with s . Then there exists a pair of edges $e = su$, $f = st$ so that $\lambda_G(x, y) = \lambda_{G^{ef}}(x, y)$ holds for every $x, y \in V - s$.*

3. Degree-specified Steiner orientations

In this section we consider the DEGREE-SPECIFIED STEINER ORIENTATION problem,² which will be the basic tool for proving the main theorems. Note that we shall only consider this problem in graphs. Given a graph $G = (V, E)$, a terminal set $S \subseteq V(G)$ and a connectivity requirement function $h : 2^S \rightarrow \mathbb{Z}$, we say the connectivity requirement function $h^* : 2^V \rightarrow \mathbb{Z}$ is the *Steiner extension* of h if $h^*(X) = h(X \cap S)$ for every $X \subseteq V$. Suppose G, S, h are given as above, and an in-degree specification $m(v)$ for each Steiner vertex is given. The goal of the DEGREE-SPECIFIED STEINER ORIENTATION problem is to find an orientation D of G that covers the Steiner extension h^* of h , with an additional requirement that $d_D^{\text{in}}(v) = m(v)$ for every $v \in V(G) - S$.

This problem is a generalization of the hypergraph orientation problem studied in [3,13,18]. Given a hypergraph $H = (V, \mathcal{E})$, we construct the bipartite representation B of H for which the terminal vertices correspond to $V(H)$ and the Steiner vertices correspond to $\mathcal{E}(H)$. Now, by specifying the indegree of each Steiner vertex to be exactly 1, an orientation of B with the specified indegrees corresponds to a hypergraph orientation of H .

We show that the DEGREE-SPECIFIED STEINER ORIENTATION problem can be solved in polynomial time if h is a non-negative intersecting supermodular set function. Notice that h^* is not an intersecting submodular function in general, and therefore Theorem 2.2 cannot be directly applied. Nonetheless, we can reformulate the problem so that we can use Theorem 2.2.

Since the indegrees of the vertices in $V - S$ are fixed, we have to determine the indegrees of the vertices in S . By Lemma 2.1, a vector $x : S \rightarrow \mathbb{Z}^+$ with $x(S) = |E| - m(V - S)$ is the vector of indegrees of a degree-specified Steiner orientation if and only if $x(X) + m(Z) \geq h^*(X \cup Z) +$

² The study of this problem is suggested by Frank (personal communication).

$i(X \cup Z) = h(X) + i(X \cup Z)$ for every $X \subseteq S$ and $Z \subseteq V - S$. Let us define the following set function on S :

$$h'(X) := h(X) + \max_{Z \subseteq V-S} (i(X \cup Z) - m(Z)) \quad \text{for } X \subseteq S.$$

It follows that there is a degree-specified Steiner orientation such that x is the vector of indegrees of the vertices of S if and only if $x(X) \geq h'(X)$ for every $X \subseteq S$ and $x(S) = |E| - m(V - S)$.

Lemma 3.1. *The set function h' is intersecting supermodular if h is intersecting supermodular.*

Proof. Let $X_1 \subseteq S$ and $X_2 \subseteq S$ be two intersecting sets. There are sets $Z_1 \subseteq V - S$ and $Z_2 \subseteq V - S$ such that $h'(X_1) = h(X_1) + i(X_1 \cup Z_1) - m(Z_1)$ and $h'(X_2) = h(X_2) + i(X_2 \cup Z_2) - m(Z_2)$. By the properties of the set functions involved, we have the following inequalities:

- $h(X_1) + h(X_2) \leq h(X_1 \cap X_2) + h(X_1 \cup X_2)$.
- $i(X_1 \cup Z_1) + i(X_2 \cup Z_2) \leq i((X_1 \cap X_2) \cup (Z_1 \cap Z_2)) + i((X_1 \cup X_2) \cup (Z_1 \cup Z_2))$.
- $m(Z_1) + m(Z_2) = m(Z_1 \cap Z_2) + m(Z_1 \cup Z_2)$.

Thus

$$\begin{aligned} h'(X_1) + h'(X_2) &= h(X_1) + h(X_2) + i(X_1 \cup Z_1) + i(X_2 \cup Z_2) - m(Z_1) - m(Z_2) \\ &\leq h(X_1 \cap X_2) + i((X_1 \cap X_2) \cup (Z_1 \cap Z_2)) - m(Z_1 \cap Z_2) \\ &\quad + h(X_1 \cup X_2) + i((X_1 \cup X_2) \cup (Z_1 \cup Z_2)) - m(Z_1 \cup Z_2) \\ &\leq h'(X_1 \cap X_2) + h'(X_1 \cup X_2). \quad \square \end{aligned}$$

Let us consider the following polyhedron:

$$\mathcal{B} := \{x \in \mathbb{R}^S : x(X) \geq h'(X) \text{ for every } X \subseteq S, x(S) = |E| - m(V - S)\}.$$

The integer vectors of this polyhedron correspond to indegree vectors of degree-specified Steiner orientations. By Theorem 2.2, \mathcal{B} is non-empty if and only if the following two conditions hold:

1. $h'(\emptyset) = 0$,
2. $\sum_{X \in \mathcal{F}} h'(X) \leq |E| - m(V - S)$ for every partition \mathcal{F} of S .

If \mathcal{B} is non-empty, then it is a base polyhedron, so its vertices are integral. As we have seen, such a vertex is the indegree vector of a degree-specified Steiner orientation. Thus the non-emptiness of \mathcal{B} is equivalent to the existence of a degree-specified orientation. Since a vertex of a base polyhedron given by an intersecting supermodular set function can be found in polynomial time, we obtained the following results:

Theorem 3.2. *Let $G = (V, E)$ be an undirected graph with a terminal set $S \subseteq V$. Let $h : 2^S \rightarrow \mathbb{Z}^+$ be a non-negative intersecting supermodular set function and $m : (V - S) \rightarrow \mathbb{Z}^+$ be an indegree specification. Then G has an orientation covering the Steiner extension h^* of h with the specified indegrees if and only if $i(Z) \leq m(Z)$ for every $Z \subseteq V - S$ and for every partition \mathcal{F} of S*

$$\sum_{X \in \mathcal{F}} \left(h(X) + \max_{Z \subseteq V-S} (i(X \cup Z) - m(Z)) \right) \leq |E| - m(V - S).$$

Theorem 3.3. *If h is non-negative and intersecting supermodular, then the DEGREE-SPECIFIED STEINER ORIENTATION problem can be solved in polynomial time.*

3.1. Steiner rooted-orientations of graphs

In the following we focus on the STEINER ROOTED-ORIENTATION problem. First we derive Theorem 1.3 as a corollary of Theorem 3.2. In contrast with Theorem 3.3, the STEINER ROOTED-ORIENTATION problem is NP-complete (Theorem 6.1). That said, in general, finding an in-degree specification for the Steiner vertices to maximize the Steiner rooted-edge-connectivity is hard.

Proof of Theorem 1.3. Let S be the set of terminal vertices and $r \in S$ be the root vertex. Set $h(X) := k$ for every $X \subseteq S$ with $r \notin X$, and $h(X) := 0$ otherwise. Then h is an intersecting supermodular function on S . By Menger’s theorem, an orientation is Steiner rooted k -arc-connected if and only if it covers the Steiner extension of h . Thus, by Theorem 3.2, the problem of finding a Steiner rooted-orientation with the specified indegrees can be solved in polynomial time. \square

The following theorem can be derived from Theorem 3.2 (by “hardwiring” the indegrees of the Steiner vertices to be 1), which will be used to prove Theorem 1.1 for rank 3 hypergraphs. This is also implicit in [3], we omit the proof here.

Theorem 3.4. *Let $G = (V, E)$ be an undirected graph with terminal set $S \subseteq V(G)$. If every Steiner vertex is of degree at most 3 and there is no edge between two Steiner vertices in G , then G has a Steiner rooted k -edge-connected orientation if and only if*

$$\delta(\mathcal{P}) \geq k(t - 1)$$

holds for every partition $\mathcal{P} = (V_1, \dots, V_t)$ of $V(G)$ such that each V_i contains a terminal vertex, where $\delta(\mathcal{P})$ denotes the number of edges with one endpoint in V_i and the other endpoint in V_j for $i \neq j$. In fact, there exists such an orientation with every Steiner vertex of indegree 1.

4. Proof of Theorem 1.1

In this section, we present the proof of the main result of this paper (Theorem 1.1). We shall consider a minimal counterexample \mathcal{H} of Theorem 4.2 with the minimum number of edges and then the minimum number of vertices. Note that Theorem 4.2 is a stronger version of Theorem 1.1 with an “extension property” introduced (Definition 4.1). The extension property allows us to apply a graph decomposition procedure to simplify the structures of \mathcal{H} significantly (Corollaries 4.5, 4.6). With these structures, we can construct a bipartite graph representation B of \mathcal{H} . Then, the DEGREE-SPECIFIED STEINER ROOTED-ORIENTATION problem can be applied in the bipartite graph B to establish a tight approximate min–max relation (Theorem 4.10).

We need some notation to state the extension property. A hyperarc a is in $\delta^{\text{in}}(X; \bar{Y})$ if a enters X and $a \cap Y = \emptyset$. If Y is an emptyset, then $\delta^{\text{in}}(X; \bar{Y})$ is the same as $\delta^{\text{in}}(X)$. We use $d^{\text{in}}(X; \bar{Y})$ to denote $|\delta^{\text{in}}(X; \bar{Y})|$. A hyperarc a is in $\vec{E}(X, Y; \bar{Z})$ if a leaves X , enters Y and $a \cap Z = \emptyset$. If Z is an emptyset, we denote $\vec{E}(X, Y; \bar{Z})$ by $\vec{E}(X, Y)$. We use $\vec{d}(X, Y; \bar{Z})$ to denote $|\vec{E}(X, Y; \bar{Z})|$, and $\vec{d}(X, Y)$ to denote $|\vec{E}(X, Y)|$. The following extension property is at the heart of our approach.

Definition 4.1. Given $H = (V, \mathcal{E})$, $S \subseteq V$ and a vertex $s \in S$, a Steiner rooted-orientation D of H extends s if:

- (i) $d_D^{\text{in}}(s) = d_H(s)$;
- (ii) $d_D^{\text{in}}(Y; \bar{s}) \geq \vec{d}_D(Y, s)$ for every $Y \subseteq V$ for which $Y \cap S = \emptyset$.

As mentioned previously, we shall prove the following stronger theorem which immediately implies Theorem 1.1.

Theorem 4.2. Suppose H is an undirected hypergraph, S is a subset of terminal vertices with a specified root vertex $r \in S$. Then H has a Steiner rooted k -hyperarc-connected orientation if S is $2k$ -hyperedge-connected in H . In fact, given any vertex $s \in S$ of degree $2k$, H has a Steiner rooted k -hyperarc-connected orientation that extends s . We call the special vertex s the sink of H .

The next lemma shows that the choice of the root vertex does not matter. The proof idea is that we can reverse the directions of the arcs in the r, v -paths.

Lemma 4.3. Suppose there exists a Steiner rooted k -hyperarc-connected orientation that extends s with r as the root. Then there exists a Steiner rooted k -hyperarc-connected orientation that extends s with v as the root for every $v \in S - s$.

Proof. Let D be a Steiner rooted k -hyperarc-connected orientation that extends s with r as the root. Let $v \neq r$ be another terminal vertex which is not the special sink s . By assumption, there are k hyperarc-disjoint paths $\{\vec{P}_1, \dots, \vec{P}_k\}$ between r and v . Now, let D' be an orientation with the same orientation as D except the orientations of all the hyperarcs in $P_1 \cup \dots \cup P_k$ are reversed. To be more precise, let $\vec{P}_i = \{v_l, a_0, v_1, a_1, \dots, a_{l-1}, v_l\}$ where a_i has v_i as the tail and v_{i+1} as a head, then $\overleftarrow{P}_i = \{v_l, \overleftarrow{a}_{l-1}, \dots, \overleftarrow{a}_0, v_0\}$ where \overleftarrow{a}_i has v_{i+1} as the tail and v_i as a head. For a directed path $\vec{P} = \{v_0, a_0, v_1, a_1, \dots, a_{l-1}, v_l\}$, we say a hyperarc a_i enters a subset of vertices X if $v_i \notin X$ and $v_{i+1} \in X$; and a_i in \vec{P} leaves X if $v_i \in X$ and $v_{i+1} \notin X$. We claim that D' is a Steiner rooted k -hyperarc-connected orientation that extends s with v as the root. First we check that $d_{D'}^{\text{in}}(X) \geq k$ for every $X \subseteq V(H)$ which satisfies $v \notin X$ and $X \cap S \neq \emptyset$. If $r \in X$, then $\{\overleftarrow{P}_1, \dots, \overleftarrow{P}_k\}$ are k hyperarc-disjoint paths from v to r in D' , where \overleftarrow{P}_i denotes the reverse path of \vec{P}_i . Hence $d_{D'}^{\text{in}}(X) \geq k$ for such X . So we assume $r \notin X$. As D is a Steiner rooted k -hyperarc-connected orientation, we have $d_D^{\text{in}}(X) \geq k$. Recall that D and D' differ only on the orientations of the paths in $\{P_1, \dots, P_k\}$. Notice that each path \vec{P}_i has both endpoints outside of X , and thus \vec{P}_i enters X the same number of times as it leaves X . Therefore, by reorienting \vec{P}_i to \overleftarrow{P}_i for all i , we have $d_{D'}^{\text{in}}(X) = d_D^{\text{in}}(X) \geq k$ for those X which contains a terminal but contains neither v nor r . This confirms that D' is a Steiner rooted k -hyperarc-connected orientation with v as the root. To finish the proof, we need to check that D' extends s as defined in Definition 4.1. Since s is a sink in D , by reorienting paths which do not start and end in s , s is still a sink in D' . So the first condition in Definition 4.1 is satisfied. For a subset $Y \subseteq V(H)$ with $Y \cap S = \emptyset$, \vec{P}_i enters Y and leaves Y the same number of times. Let a_1 be a hyperarc that enters Y and a_2 be a hyperarc that leaves Y in D . Suppose we reverse a_1 and a_2 in D' . We have four cases to consider.

- $s \in a_1$ and $s \in a_2$. Then $d_{D'}^{\text{in}}(Y; \bar{s}) = d_D^{\text{in}}(Y; \bar{s}) \geq \vec{d}_D(Y, s) = \vec{d}_{D'}(Y, s)$.
- $s \in a_1$ and $s \notin a_2$. Then $d_{D'}^{\text{in}}(Y; \bar{s}) = d_D^{\text{in}}(Y; \bar{s}) + 1 \geq \vec{d}_D(Y, s) + 1 = \vec{d}_{D'}(Y, s)$.




Fig. 1. An illustration of the proof of Lemma 4.4.

- $s \notin a_1$ and $s \in a_2$. Then $d_D^{\text{in}}(Y; \bar{s}) = d_D^{\text{in}}(Y; \bar{s}) - 1 \geq \vec{d}_D(Y, s) - 1 = \vec{d}_{D'}(Y, s)$.
- $s \notin a_1$ and $s \notin a_2$. Then $d_{D'}^{\text{in}}(Y; \bar{s}) = d_D^{\text{in}}(Y; \bar{s}) \geq \vec{d}_D(Y, s) = \vec{d}_{D'}(Y, s)$.

Since we have $d_D^{\text{in}}(Y; \bar{s}) \geq \vec{d}_D(Y, s)$ to start with, by reorienting \vec{P}_i to \vec{P}_i , we still have $d_{D'}^{\text{in}}(Y; \bar{s}) \geq \vec{d}_{D'}(Y, s)$. Hence the second condition in Definition 4.1 is also satisfied. Therefore, D' is a Steiner rooted k -hyperarc-connected orientation that extends s . This proves the lemma. \square

In the following we say a set X is *tight* if $d_{\mathcal{H}}(X) = 2k$ and $X \cap S$ and $S - X$ are non-empty; X is *nontrivial* if $|X| \geq 2$ and $|V(\mathcal{H}) - X| \geq 2$. The following is the key lemma where we use the graph decomposition technique (see Fig. 1 for an illustration).

Lemma 4.4. *There is no nontrivial tight set in \mathcal{H} .*

Proof. Suppose there exists a nontrivial tight set U , i.e. $d_{\mathcal{H}}(U) = 2k$, $|U| \geq 2$ and $|V(\mathcal{H}) - U| \geq 2$. By Lemma 4.3 we may assume that $r \in U$ and $s \notin U$. Contract $V(\mathcal{H}) - U$ of \mathcal{H} to a single vertex v_1 and call the resulting hypergraph H_1 (notice this may create parallel hyperedges); similarly, contract U of \mathcal{H} to a single vertex v_2 and call the resulting hypergraph H_2 . We assume $s \in H_2$. See Fig. 1(b) for an illustration. So, $V(H_1) = U \cup \{v_1\}$, $V(H_2) = (V(\mathcal{H}) - U) \cup \{v_2\}$ and there is an one-to-one correspondence between the hyperedges in $\delta_{H_1}(v_1)$ and the hyperedges in $\delta_{H_2}(v_2)$. To be more precise, for a hyperedge $e \in \delta_{\mathcal{H}}(U)$, it decomposes into $e_1 = (e \cap V(H_1)) \cup \{v_1\}$ in H_1 and $e_2 = (e \cap V(H_2)) \cup \{v_2\}$ in H_2 and we refer them as the corresponding hyperedges of e in H_1 and H_2 , respectively.

Since U is nontrivial, both H_1 and H_2 are smaller than \mathcal{H} . We set $S_1 := (S \cap V(H_1)) \cup v_1$ and $S_2 := (S \cap V(H_2)) \cup v_2$, and set the sink of H_1 to be v_1 and the sink of H_2 to be s . Clearly, S_1 is $2k$ -hyperedge-connected in H_1 and S_2 is $2k$ -hyperedge-connected in H_2 . By the minimality of \mathcal{H} , H_2 has a Steiner rooted k -hyperarc-connected orientation D_2 that extends s . By Lemma 4.3, we can choose the root of D_2 to be v_2 . Similarly, by the minimality of \mathcal{H} , H_1 has a Steiner rooted k -hyperarc-connected orientation D_1 that extends v_1 . Let the root of D_1 be r . See Fig. 1(c) for an illustration.

We shall prove that the concatenation D of the two orientations D_1, D_2 gives a Steiner rooted k -hyperarc-connected orientation of \mathcal{H} that extends s . Notice for a hyperedge e in $\delta_{\mathcal{H}}(U)$, its corresponding hyperedge e_1 in H_1 is oriented with v_1 as a head (by the extension property of D_1), and its corresponding hyperedge e_2 in H_2 is oriented so that v_2 is the tail (as v_2 is the root of D_2). So, in D , the orientation of e is well defined and has its tail in H_1 . See Fig. 1(d) for an illustration. Now we show that D is a Steiner rooted k -hyperarc-connected orientation. By Menger’s theorem, it suffices to show that $d_D^{\text{in}}(X) \geq k$ for any $X \subseteq V(\mathcal{H})$ for which $r \notin X$ and $X \cap S \neq \emptyset$.

Suppose $X \cap S_1 \neq \emptyset$. Then $d_{D_1}^{\text{in}}(X - V(H_2)) \geq k$ by the orientation D_1 of H_1 . Since v_1 is the sink of G_1 , there is no hyperarc going from $V(H_2)$ to $V(H_1)$ in D . Hence we have $d_D^{\text{in}}(X) \geq d_{D_1}^{\text{in}}(X - V(H_2)) \geq k$.

Suppose $X \cap S_1 = \emptyset$. Let $X_1 = X \cap H_1$ and $X_2 = X \cap H_2$. The case that $X_1 = \emptyset$ follows from the properties of D_2 . So we assume both X_1 and X_2 are non-empty. We have the following inequality:

$$d_D^{\text{in}}(X) \geq d_{D_1}^{\text{in}}(X_1; \bar{v}_1) + d_{D_2}^{\text{in}}(X_2) - \vec{d}_D(X_1, X_2). \tag{2}$$

Note that $\vec{d}_{D_1}(X_1, v_1) \geq \vec{d}_D(X_1, X_2)$. So, by property (ii) of Definition 4.1, $d_{D_1}^{\text{in}}(X_1; \bar{v}_1) \geq \vec{d}_{D_1}(X_1, v_1) \geq \vec{d}_D(X_1, X_2)$. Hence $d_D^{\text{in}}(X) \geq d_{D_2}^{\text{in}}(X_2) \geq k$, where the second inequality is by the properties of D_2 .

This implies that D is a Steiner rooted k -hyperarc-connected orientation of \mathcal{H} . To finish the proof, we need to check that D extends s . The first property of Definition 4.1 follows immediately from our construction. It remains to check that property (ii) of Definition 4.1 still holds in D . Consider a subset $Y \subset V(\mathcal{H})$ with $Y \cap S = \emptyset$. Let $Y_1 = Y \cap H_1$ and $Y_2 = Y \cap H_2$. The following inequality is important:

$$d_D^{\text{in}}(Y; \bar{s}) \geq d_{D_1}^{\text{in}}(Y_1; \bar{v}_1) + d_{D_2}^{\text{in}}(Y_2; \bar{s}) - \vec{d}_D(Y_1, Y_2; \bar{s}). \tag{3}$$

By property (ii) of the extension property of D_1 , we have $d_{D_1}^{\text{in}}(Y_1; \bar{v}_1) \geq \vec{d}_{D_1}(Y_1, v_1) \geq \vec{d}_D(Y_1, Y_2; \bar{s}) + \vec{d}_D(Y_1, s)$. Therefore, $d_D^{\text{in}}(Y; \bar{s}) \geq \vec{d}_D(Y_1, s) + d_{D_2}^{\text{in}}(Y_2; \bar{s})$. By property (ii) of the extension property of D_2 , we have $d_{D_2}^{\text{in}}(Y_2; \bar{s}) \geq \vec{d}_{D_2}(Y_2, s)$. Hence, by (3), $d_D^{\text{in}}(Y; \bar{s}) \geq \vec{d}_D(Y_1, s) + \vec{d}_{D_2}(Y_2, s) = \vec{d}_D(Y_1, s) + \vec{d}_D(Y_2, s) = \vec{d}_D(Y, s)$, as required. This shows that D extends s , which contradicts that \mathcal{H} is a counterexample. \square

The following are two important properties obtained from Lemma 4.4.

Corollary 4.5. *Each hyperedge of \mathcal{H} of size at least 3 contains only terminal vertices.*

Proof. Suppose e is a hyperedge of \mathcal{H} of size at least 3 and $t \in e$ is a Steiner vertex. Let H' be a hypergraph with the same vertex and edge set as \mathcal{H} except we replace e by $e' := e - t$. If H' is $2k$ -hyperedge-connected, then by the choice of \mathcal{H} , H' has a Steiner rooted k -hyperedge-connected orientation, hence \mathcal{H} also has one; a contradiction. Therefore, there exists a set X which separates two terminals with $d_{\mathcal{H}}(X) = 2k$ and $d_{H'}(X) < 2k$. So $e \in \delta_{\mathcal{H}}(X)$. Suppose $t \in X$. Since X contains a terminal, $|X| \geq 2$. Also, $e - t$ must be contained in $V(\mathcal{H}) - X$; otherwise $d_{\mathcal{H}}(X) = d_{H'}(X)$. Hence $|V(\mathcal{H}) - X| \geq |e - t| \geq 2$. Therefore, X is a nontrivial tight set, which contradicts Lemma 4.4. \square

Corollary 4.6. *There is no edge between two Steiner vertices in \mathcal{H} .*




Fig. 2. The bipartite representation B of \mathcal{H} .

Proof. This follows from a similar argument as in Corollary 4.5. Let e be an edge which connects two Steiner vertices. If $\mathcal{H} - e$ is $2k$ -hyperedge-connected, then by the choice of \mathcal{H} , $\mathcal{H} - e$ has a Steiner rooted k -hyperarc-connected orientation, hence \mathcal{H} also has one; a contradiction. Otherwise, there exists a set X which separates two terminals with $d_{\mathcal{H}}(X) = 2k$ and $d_{\mathcal{H}-e}(X) < 2k$. So $e \in \delta_{\mathcal{H}}(X)$. Since X contains a terminal vertex and an endpoint of e which is a Steiner vertex, $|X| \geq 2$. Similarly, $|V(\mathcal{H}) - X| \geq 2$. Hence X is a nontrivial tight set, which contradicts Lemma 4.4. \square

4.1. The bipartite representation of \mathcal{H}

Using Corollaries 4.5 and 4.6, we shall construct a bipartite graph from \mathcal{H} , which allows us to apply the results on the DEGREE-SPECIFIED STEINER ROOTED-ORIENTATION problem to \mathcal{H} . Let S be the set of terminal vertices in \mathcal{H} . Let \mathcal{E}' be the set of hyperedges in \mathcal{H} which do not contain a Steiner vertex, i.e. a hyperedge e is in \mathcal{E}' if $e \cap (V(\mathcal{H}) - S) = \emptyset$. We construct a bipartite graph $B = (S, (V(\mathcal{H}) - S) \cup \mathcal{E}'; E)$ from the hypergraph \mathcal{H} as follows. Every vertex v in \mathcal{H} corresponds to a vertex v in B , and also every hyperedge $e \in \mathcal{E}'$ corresponds to a vertex v_e in B . By Corollary 4.5, hyperedges which intersect $V(\mathcal{H}) - S$ are graph edges (i.e. hyperedges of size 2); we add these edges to $E(B)$. For every hyperedge $e \in \mathcal{E}'$, we add $v_e w$ to $E(B)$ if and only if $w \in e$ in \mathcal{H} . Let the set of terminal vertices in B be S (the same set of terminal vertices in \mathcal{H}); all other vertices are non-terminal vertices in B . By Corollaries 4.5 and 4.6, there is no edge between two non-terminal vertices in B . Hence B is a bipartite graph. To distinguish the non-terminal vertices corresponding to Steiner vertices in \mathcal{H} and the non-terminal vertices corresponding to hyperedges in \mathcal{E}' , we call the former the Steiner vertices and the latter the hyperedge vertices. See Fig. 2 for an illustration.

4.2. Rank 3 hypergraphs

To better illustrate the idea of the proof, we first prove Theorem 4.2 for the case of rank 3 hypergraphs. This motivates the proof for general hypergraphs, which is considerably more complicated. It is not needed for the main result. Readers may choose to skip it.

Lemma 4.7. \mathcal{H} is not a rank 3 hypergraph.

Proof. Since \mathcal{H} is of rank 3, all hyperedge vertices in B are of degree at most 3. The use of the rank 3 assumption is the following simple observation, which allows us to relate the hyperedge-connectivity of \mathcal{H} to edge-connectivity of B .

Proposition 4.8. *S is 2k-hyperedge-connected in \mathcal{H} if and only if S is 2k-edge-connected in B.*

Proof. Consider $a, b \in S$. If there are $2k$ hyperedge-disjoint paths from a to b in \mathcal{H} , then clearly there are $2k$ edge-disjoint paths from a to b in B . Suppose there are $2k$ edge-disjoint paths from a to b in B . Since each hyperedge vertex $z \in \mathcal{E}'$ is of degree at most 3, no two edge-disjoint paths in B share a hyperedge vertex. Hence there are $2k$ hyperedge-disjoint paths from a to b in \mathcal{H} . \square

We remark that Proposition 4.8 does not hold for hypergraphs of rank greater than 3. With Proposition 4.8, we can apply Mader’s splitting off theorem to prove the following:

Lemma 4.9. *Steiner vertices of \mathcal{H} are of degree at most 3.*

Proof. If a Steiner vertex v is not of degree 3 in \mathcal{H} , then it is not of degree 3 in B . So we can apply Mader’s splitting-off theorem (Theorem 2.3) to find a suitable splitting at v in B . Let $e_1 = s_1v$ and $e_2 = vs_2$ be the pair of edges that we split-off, and $e = s_1s_2$ be the new edge. By Corollary 4.6, s_1 and s_2 are terminal vertices. We add a new Steiner vertex v_e to $V(B)$ and replace the edge s_1s_2 by two new edges $v_e s_1$ and $v_e s_2$. Since B is bipartite, the resulting graph, denoted by B' , is bipartite. Notice that B' corresponds to a hypergraph H' with $V(H') = V(\mathcal{H})$ and $E(H') = E(\mathcal{H}) - \{e_1, e_2\} + \{e\}$. S remains k -edge-connected in B' , so by Proposition 4.8, S is k -hyperedge-connected in H' . By the minimality of \mathcal{H} , there is a Steiner rooted k -hyperarc-connected orientation of H' . Suppose s_1s_2 in H' is oriented as $\overrightarrow{s_1s_2}$ in H' , then we orient vs_1 and vs_2 as $\overrightarrow{s_1v}$ and $\overrightarrow{vs_2}$ in \mathcal{H} . All other hyperedges in \mathcal{H} have the same orientations as the corresponding hyperedges in H' . It is easy to see that this orientation is a Steiner rooted k -hyperarc-connected orientation of \mathcal{H} , and also the extension property holds, a contradiction. \square

Now we are ready to finish the proof of Lemma 4.7. Construct $B' = B - s$, where we remove all edges in B which are incident with s . We shall use Theorem 3.4 to prove that there is a Steiner rooted k -arc-connected orientation of B' . Since S is $2k$ -edge-connected in B , for any partition $\mathcal{P} = \{P_1, \dots, P_t\}$ of $V(B')$ such that each P_i contains a terminal vertex, we have $\sum_{i=1}^t d_{B'}(P_i) = \sum_{i=1}^t d_B(P_i) - d_B(s) \geq 2kt - 2k = 2k(t - 1)$. So there are at least $k(t - 1)$ edges crossing \mathcal{P} in B' .

By Theorem 3.4, there is a Steiner rooted k -edge-connected orientation D' of B' with the additional property that each Steiner vertex has indegree exactly 1. By orienting the edges in $\delta_{B'}(s)$ to have s as the head, we obtain an orientation D of B . Note that each Steiner vertex still has indegree exactly 1, and so D corresponds to a hypergraph orientation of \mathcal{H} . Also, by this construction, property (i) of Definition 4.1 is satisfied.

Consider an arbitrary Y for which $Y \cap S = \emptyset$. Since every vertex y in Y is of degree at most 3 by Lemma 4.9, y can have at most one outgoing arc to s ; otherwise $d_{\mathcal{H}}(\{s, y\}) < 2k$ which contradicts our connectivity assumption since $d_{\mathcal{H}}(s) = 2k$ (recall that $d_{\mathcal{H}}(s) = 2k$ as s is the sink). Since Y induces an independent set by Corollary 4.6 and each vertex in Y has indegree exactly 1, each $y \in Y$ has an incoming arc from outside Y . Notice that those incoming arcs are of size 2 by Corollary 4.5, So we have $d_D^{\text{in}}(Y; \bar{s}) \geq \vec{d}(Y, s)$. This implies that D satisfies property (ii) of Definition 4.1 as well.

Finally we verify that D is a Steiner rooted k -hyperedge-connected orientation. Consider a subset $X \subseteq V(\mathcal{H})$ which contains a terminal but not the root. If X contains a terminal other

than s , then clearly $d_D^{\text{in}}(X) \geq k$ by the orientation on $\mathcal{H} - s$. So suppose $X \cap S = s$. As argued above, since each Steiner vertex v is of degree 3, v has at most one outgoing arc to s . As each Steiner vertex is of indegree 1 and there is no edge between two Steiner vertices, we have $d_D^{\text{in}}(X) \geq d_D^{\text{in}}(s) = 2k$ as s is the sink. This shows that D is a Steiner rooted k -hyperarc-connected orientation that extends s , which contradicts the assumption that \mathcal{H} is a counterexample. \square

4.3. Applying degree-specified Steiner orientation

For the proof of Theorem 4.2 for the case of rank 3 hypergraphs, a crucial step is to apply Mader’s splitting-off lemma to the bipartite representation B of \mathcal{H} to obtain Lemma 4.9. In general hypergraphs, however, a suitable splitting at a Steiner vertex which preserves the edge-connectivity of S in B might not preserve the hyperedge-connectivity of S in \mathcal{H} . And there is no analogous edge splitting-off result which preserves hyperedge-connectivity.

Our key observation is that, if we were able to apply Mader’s lemma as in the proof of Lemma 4.7, then every Steiner vertex would end up with indegree $\lfloor d(v)/2 \rfloor$ in the resulting orientation of B . So, we apply the DEGREE-SPECIFIED STEINER ROOTED-ORIENTATION problem by “hardwiring” $m(v) = \lfloor d(v)/2 \rfloor$ for the Steiner vertices. Also, we “hardwire” the indegree of the sink to be $2k$ for the extension property. (In the example of Fig. 2, the indegrees of the Steiner vertices are specified to be 3, 2, 1 from left to right; the sink becomes a non-terminal vertex with specified indegree $2k$.) Quite surprisingly, such an orientation always exists when S is $2k$ -hyperedge connected in \mathcal{H} . The following theorem is the final (and most technical) step to the proof of Theorem 4.2, which shows that a minimal counterexample of Theorem 4.2 does not exist.

Theorem 4.10. *Suppose that S is $2k$ -hyperedge-connected in $H = (V, \mathcal{E})$, there is no edge between two Steiner vertices, and no hyperedge of size at least 3 contains a Steiner vertex. Let $s_0 \in S$ be a vertex of degree $2k$. Then H has a Steiner rooted k -hyperarc-connected orientation that extends s_0 .*

Proof. We will use the theorem on the DEGREE-SPECIFIED STEINER ROOTED-ORIENTATION problem of graphs (Theorem 3.2). To get an instance of that problem, we consider the bipartite representation $B = (V', E')$ of H that was defined in Section 4.1 (i.e. we replace each hyperedge in \mathcal{E}' by a hyperedge vertex). Let the set of terminals in B be $S' := S - s_0$. The indegree specification $m' : V' - S' \rightarrow \mathbb{Z}^+$ is defined by

$$m'(v) := \begin{cases} \lfloor d_H(v)/2 \rfloor & \text{if } v \text{ is a Steiner vertex,} \\ 1 & \text{if } v \text{ is a hyperedge vertex,} \\ 2k & \text{if } v = s_0 \text{ is the sink.} \end{cases}$$

By Theorem 3.2, this graph has a Steiner rooted k -arc-connected orientation with the specified indegrees if and only if the following conditions hold:

$$i(Z) \leq m'(Z) \quad \text{for every } Z \subseteq V' - S', \tag{4}$$

$$\sum_{X \in \mathcal{F}} \left(h(X) + \max_{Y \subseteq V' - S'} (i(X \cup Y) - m'(Y)) \right) \leq |E'| - m'(V' - S') \tag{5}$$

for every partition \mathcal{F} of S' , where $h : S' \rightarrow \mathbb{Z}^+$ is defined by

$$h(X) := \begin{cases} k & \text{if } \emptyset \neq X \subseteq S' - r, \\ 0 & \text{otherwise.} \end{cases}$$

It is easy to see that condition (4) is always satisfied, since the only edges spanned by $V' - S'$ are those incident to s_0 , and $d_B(s_0) = 2k = m'(s_0)$.

Proposition 4.11. *Condition (5) is satisfied if*

$$\sum_{e \in \mathcal{E}} (|\{X \in \mathcal{F} : e \cap X \neq \emptyset\}| - 1) + \sum_{v \notin \bigcup \mathcal{F} + s_0} \left\lceil \frac{d_H(v)}{2} \right\rceil \geq k(|\mathcal{F}| - 1) \tag{6}$$

for every subpartition \mathcal{F} of V for which $S \cap X \neq \emptyset$ for every $X \in \mathcal{F}$, and $S \cap (\bigcup \mathcal{F}) = S - s_0$.

Proof. Suppose that there is a partition \mathcal{F} of S' where (5) does not hold. By the definition of h ,

$$k(|\mathcal{F}| - 1) + \sum_{X \in \mathcal{F}} \max_{Y \subseteq V' - S'} (i(X \cup Y) - m'(Y)) > |E'| - m'(V' - S'). \tag{7}$$

For a given $X \in \mathcal{F}$ we can determine the set Y where the maximum is attained. We can assume that s_0 is not in Y , since its inclusion would increase $m'(Y)$ by $2k$, and $i(X \cup Y)$ can increase by at most $2k$.

We can assume that Y contains all the hyperedge vertices corresponding to hyperedges that are not disjoint from X . The inclusion of such a vertex increases $m'(Y)$ by 1, and increases $i(X \cup Y)$ by at least 1. By a similar argument, we may assume that Y does not contain hyperedge vertices corresponding to hyperedges that are disjoint from X , since the inclusion of such a vertex would not increase $i(X \cup Y)$.

Finally, if we take into account the above observations, the inclusion in Y of a Steiner vertex v increases $m'(Y)$ by $\lfloor d_H(v)/2 \rfloor$, and increases $i(X \cup Y)$ by $|\{e \in \mathcal{E} : v \in e \subseteq X + v\}|$. Therefore we may assume that a Steiner vertex v is included in Y if and only if $|\{e \in \mathcal{E} : v \in e \subseteq X + v\}| > \lfloor d_H(v)/2 \rfloor$.

For a given X , we determined a set $Y \subseteq V' - S'$ where the maximum in (7) is attained. Let $X^* := X \cup (Y \cap (V - S))$. If $X_1 \subseteq S - s_0$ and $X_2 \subseteq S - s_0$ are disjoint sets, then X_1^* and X_2^* are also disjoint, since a node in $V - S$ cannot have more than half of its neighbors in both X_1 and X_2 . So if \mathcal{F} is a partition of $S - s_0$, then $\mathcal{F}^* := \{X^* : X \in \mathcal{F}\}$ is a subpartition of V for which $S \cap X^* \neq \emptyset$ for every $X \in \mathcal{F}^*$, and $S \cap (\bigcup \mathcal{F}^*) = S - s_0$.

Since (7) holds for \mathcal{F} , the following holds for \mathcal{F}^* :

$$k(|\mathcal{F}^*| - 1) = k(|\mathcal{F}| - 1) > |E'| - m'(V' - S') - \sum_{X \in \mathcal{F}} \max_{Y \subseteq V' - S'} (i(X \cup Y) - m'(Y)).$$

Here

$$\begin{aligned} |E'| &= |\mathcal{E}| + \sum_{e \in \mathcal{E}'} (|e| - 1), \\ m'(V' - S') &= |\mathcal{E}'| + 2k + \sum_{v \in V - S} \left\lfloor \frac{d_H(v)}{2} \right\rfloor, \\ \max_{Y \subseteq V' - S'} (i(X \cup Y) - m'(Y)) &= \sum_{e \in \mathcal{E}} \max\{0, |e \cap X^*| - 1\} - \sum_{v \in X^* \cap (V - S)} \left\lfloor \frac{d_H(v)}{2} \right\rfloor. \end{aligned}$$

Using these identities, and the fact that $d_H(s_0) = 2k$, we get the following inequalities:

$$\begin{aligned}
 & k(|\mathcal{F}^*| - 1) \\
 & > |\mathcal{E}| + \sum_{e \in \mathcal{E}'} (|e| - 2) - \sum_{e \in \mathcal{E}} \sum_{X^* \in \mathcal{F}^*} \max\{0, |e \cap X^*| - 1\} - 2k - \sum_{v \notin \bigcup \mathcal{F}^* + s_0} \left\lfloor \frac{d_H(v)}{2} \right\rfloor \\
 & = \sum_{e \in \mathcal{E}} \left(|e| - 1 - \sum_{X^* \in \mathcal{F}^*} \max\{0, |e \cap X^*| - 1\} \right) - 2k - \sum_{v \notin \bigcup \mathcal{F}^* + s_0} \left\lfloor \frac{d_H(v)}{2} \right\rfloor \\
 & = \sum_{e \in \mathcal{E}} \left(|e \cap (V - \bigcup \mathcal{F}^*)| + |\{X^* \in \mathcal{F}^*: e \cap X^* \neq \emptyset\}| - 1 \right) - 2k \\
 & \quad - \sum_{v \notin \bigcup \mathcal{F}^* + s_0} \left\lfloor \frac{d_H(v)}{2} \right\rfloor \\
 & = \sum_{e \in \mathcal{E}} \left(|e \cap (V - (\bigcup \mathcal{F}^* + s_0))| + |\{X^* \in \mathcal{F}^*: e \cap X^* \neq \emptyset\}| - 1 \right) \\
 & \quad - \sum_{v \notin \bigcup \mathcal{F}^* + s_0} \left\lfloor \frac{d_H(v)}{2} \right\rfloor \\
 & = \sum_{e \in \mathcal{E}} (|\{X^* \in \mathcal{F}^*: e \cap X^* \neq \emptyset\}| - 1) + \sum_{v \notin \bigcup \mathcal{F}^* + s_0} \left\lfloor \frac{d_H(v)}{2} \right\rfloor.
 \end{aligned}$$

But this means that property (6) does not hold for the subpartition \mathcal{F}^* . \square

Notice that Proposition 4.11 is formulated in terms of the original hypergraph H . We will prove that the bipartite representation B of H has the desired degree-specified orientation by showing that the conditions in Proposition 4.11 are satisfied if S is $2k$ -hyperedge-connected in H .

Let \mathcal{F} be a subpartition of V for which $S \cap X \neq \emptyset$ for every $X \in \mathcal{F}$, and $S \cap (\bigcup \mathcal{F}) = S - s_0$. Let \mathcal{E}_1 denote the set of hyperedges of H which enter exactly 1 member of \mathcal{F} , and let \mathcal{E}_2 denote the set of hyperedges of H which enter at least 2 members of \mathcal{F} . Let $d_{\mathcal{E}_1}(X) := d(X) \cap \mathcal{E}_1$ and $d_{\mathcal{E}_2}(X) := d(X) \cap \mathcal{E}_2$. Let $U := V - (\bigcup \mathcal{F} + s_0)$. Then the only hyperedges that are disjoint from every member of \mathcal{F} are the edges between U and s_0 , so

$$\begin{aligned}
 & \sum_{e \in \mathcal{E}} (|\{X \in \mathcal{F}: e \cap X \neq \emptyset\}| - 1) + \sum_{v \notin \bigcup \mathcal{F} + s_0} \left\lfloor \frac{d_H(v)}{2} \right\rfloor \\
 & \geq \sum_{X \in \mathcal{F}} \frac{d_{\mathcal{E}_2}(X)}{2} - d_H(U, s_0) + \sum_{v \in U} \frac{d_H(v)}{2} \\
 & = \sum_{X \in \mathcal{F}} \frac{d_{\mathcal{E}_2}(X)}{2} + \frac{d_H(U, S - s_0)}{2} - \frac{d_H(U, s_0)}{2}. \tag{8}
 \end{aligned}$$

Here

$$d_H(U, S - s_0) = \sum_{X \in \mathcal{F}} d_{\mathcal{E}_1}(X) - |\{e \in \mathcal{E}_1: e \cap U = \emptyset\}| = \sum_{X \in \mathcal{F}} d_{\mathcal{E}_1}(X) - d_H(V - U, s_0),$$

and so

$$d_H(U, S - s_0) - d_H(U, s_0) = \sum_{X \in \mathcal{F}} d_{\mathcal{E}_1}(X) - d_H(s_0) = \sum_{X \in \mathcal{F}} d_{\mathcal{E}_1}(X) - 2k.$$

Using this identity in inequality (8) we get that

$$\sum_{e \in \mathcal{E}} (|\{X \in \mathcal{F} : e \cap X \neq \emptyset\}| - 1) + \sum_{v \notin \bigcup \mathcal{F} + s_0} \left\lceil \frac{d_H(v)}{2} \right\rceil \geq \sum_{X \in \mathcal{F}} \left(\frac{d_{\mathcal{E}_2}(X) + d_{\mathcal{E}_1}(X)}{2} \right) - k \geq k(|\mathcal{F}| - 1),$$

where the last inequality holds because $d_{\mathcal{E}_2}(X) + d_{\mathcal{E}_1}(X) \geq 2k$ for every $X \in \mathcal{F}$ as S is $2k$ -hyperedge-connected in H .

We proved that the conditions of type (6) in Proposition 4.11 are satisfied. Therefore, we have the desired degree-specified orientation of the bipartite representation B of H . Since every hyperedge vertex has indegree 1 in B , this orientation corresponds to a Steiner rooted k -hyperarc-connected orientation of H . It remains to check that this orientation extends s_0 . The first property of the extension property (Definition 4.1) follows immediately from our construction, since the indegree of s_0 is $2k$. To check the second property of the extension property, we consider an arbitrary $Y \subset V(H)$ for which $Y \cap S = \emptyset$. Since s_0 is of degree $2k$ and S is $2k$ -hyperedge-connected in H , each vertex $v \in Y$ has at most $\lfloor d(v)/2 \rfloor$ edges to s_0 . Recall that the indegree of v in the orientation is $\lfloor d(v)/2 \rfloor$. Since there are no edges between two Steiner vertices, all the incoming arcs of v come from $V(H) - Y$. Notice that these incoming arcs are of size 2 by Corollary 4.5, and so do not intersect s_0 . Hence, $d^{\text{in}}(Y; \overline{s_0}) \geq \vec{d}(Y, s_0)$, as required. \square

Since a minimal counterexample \mathcal{H} must satisfy the condition of Theorem 4.10, Theorem 4.10 proves that \mathcal{H} does not exist. So Theorem 4.2 (and hence Theorem 1.1) is proven. We remark that in the proof of Theorem 4.10, the indegree specifications on the Steiner vertices have two uses. The major use is to apply Theorem 3.2 to establish the connectivity upper bound, which consists of the bulk of the proof. The other use is that it is crucial in proving the extension property (Definition 4.1).

5. Proof of Theorem 1.2

In this section we show another application of the DEGREE-SPECIFIED STEINER ORIENTATION problem. We consider the ELEMENT-DISJOINT STEINER ROOTED-ORIENTATION problem where our goal is to find an orientation D of G that maximizes the Steiner rooted-element-connectivity. The proof of Theorem 1.2 consists of two steps. The first step is to reduce the problem from general graphs to the graphs with no edges between Steiner vertices. This technique was used in [4,15] so we omit the proof here. The second step is to reduce the problem in this special instance into the DEGREE-SPECIFIED STEINER ROOTED-ORIENTATION problem. The idea is that if we specify the indegree of each Steiner vertex to be 1, then a Steiner rooted k -arc-connected orientation is a Steiner rooted k -element-connected orientation, since each Steiner vertex cannot be in two edge-disjoint paths. It turns out that such a degree-specified orientation always exists when S is $2k$ -element-connected in G .

We remark that the property that every Steiner vertex is of indegree 1 in the orientation will be used twice—once in Lemma 5.2 to establish the connectivity upper bound, and once in the following lemma for the reduction. In the following lemma conditions (1)–(3) have been proved in [4,15]: the construction involves the deletion or contraction of edges between Steiner vertices. It is easy to check that the property in (4) can be maintained during the inverse operations of deletion and contraction.

Lemma 5.1. (See also [4,15].) Given an undirected graph G and a set S of terminal vertices. Suppose S is k -element-connected in G . Then we can construct in polynomial time a graph G' with the following properties:

- (1) $S \subseteq V'$;
- (2) there is no edge between Steiner vertices in G' ;
- (3) S is k -element-connected in G' ;
- (4) if there is a Steiner rooted k' -element-connected orientation in G' with the indegrees of the Steiner vertices being 1, then there is a Steiner rooted k' -element-connected orientation in G .

The following lemma can be shown to be a special case of Theorem 4.10.

Lemma 5.2. Given an undirected graph $G = (V, E)$ and a set S of terminal vertices. If S is $2k$ -element-connected in G and there are no edges between vertices in $V(G) - S$, then G has a Steiner rooted k -element-connected orientation with the indegrees of the Steiner vertices being 1.

Proof. We construct a hypergraph H as follows. The vertex set of H is S . For each vertex $v \in V(G) - S$, we add a hyperedge $N_G(v)$ to H . Note that since there are no edges between two vertices in $V(G) - S$, $N_G(v) \subseteq S$ and so is well defined. Also, we keep all the edges in G between two vertices in S . Since S is $2k$ -element-connected in G , S is $2k$ -hyperedge-connected in H . Theorem 4.10 implies that H (which has no Steiner vertices) has a Steiner rooted k -hyperarc-connected orientation. This corresponds to a Steiner rooted k -element-connected orientation with the indegrees of the Steiner vertices being 1. \square

Theorem 1.2 follows immediately from Lemmas 5.2 and 5.1.

6. Hardness results

Nash-Williams' orientation theorem implies that the maximum k for which a graph has a Steiner strongly k -arc-connected orientation can be found in polynomial time. By the theorem, this is equivalent to finding the maximum k for which the graph is Steiner $2k$ -edge-connected, and this can be done using $O(n)$ flow computations. Moreover, the algorithmic proof of Nash-Williams' theorem provides an algorithm for finding such an orientation. Usually the rooted counterparts of graph connectivity problems are easier to solve. For example, finding a minimum cost k -arc-connected subgraph of a directed graph is NP-hard, while a minimum cost rooted k -arc-connected subgraph can be found in polynomial time [12]. It is a very rare phenomenon that the rooted version of a connectivity problem is more difficult than the non-rooted one. In this light, the following result is somewhat surprising.

Theorem 6.1. Given a graph G , a set of terminals S , and a root vertex $r \in S$, it is NP-complete to determine if G has a Steiner rooted k -arc-connected orientation.

Proof. First we introduce the NP-complete problem to be reduced to the STEINER ROOTED-ORIENTATION problem. Let $G = (V, E)$ be a graph, and $R : V \times V \rightarrow \mathbb{Z}^+$ a demand function for which $R(v, v) = 0$ for every $v \in V$. An R -orientation of G is an orientation where for every pair $u, v \in V$ there are at least $R(u, v)$ edge-disjoint paths from u to v .

Theorem 6.2. (See [13].) *The problem of finding an R -orientation of a graph is NP-complete, even if R has maximum value 3.*

In the following we show that the R -orientation problem can be reduced to the Steiner rooted orientation problem, thus the latter is NP-complete. Let $(G = (V, E), R)$ be an instance of the R -orientation problem. We define a graph $G' = (V', E')$ such that G is an induced subgraph of G' . In addition to the vertices of V , V' contains the root r , and vertices $b_{u,v}$ for every ordered pair $(u, v) \in V \times V$, $u \neq v$. In addition to the edges of E , E' contains the following 4 types of edges:

1. $\sum_{v \in V} R(u, v)$ edges from r to u for every $u \in V$,
2. $R(u, v)$ edges from v to $b_{u,v}$ for every pair u, v ,
3. $\sum_{w \in V-v} R(u, w)$ edges from u to $b_{u,v}$ for every pair u, v ,
4. $\sum_{w \in V} R(x, w)$ edges from x to $b_{u,v}$ for every triple x, u, v of distinct nodes.

Let

$$S := \{b_{u,v} : u, v \in V, u \neq v\},$$

$$k := \sum_{u,v \in V} R(u, v).$$

We set the vertices in S to be the terminal vertices, and all other vertices the Steiner vertices.

Lemma 6.3. *The graph G' has a Steiner rooted k -edge-connected orientation if and only if G has an R -orientation.*

Proof. Let D' be a Steiner rooted k -edge-connected orientation of G' . Since the degree of r is k in G' , each edge of type 1 must be oriented away from r . Since the degree of every node in S is k in G , each edge of types 2, 3, or 4 must be oriented towards S .

Let $(u, v) \in V \times V$ be a fixed pair. Since D' is a Steiner rooted k -edge-connected orientation, there are k edge-disjoint paths from r to $b_{u,v}$. Of these paths, $k - R(u, v)$ are necessarily composed of an edge of type 1 and an edge of type 3 or 4. The remaining $R(u, v)$ paths necessarily start with the edge ru , and end with the edge $vb_{u,v}$. Thus, in order to “complete” these paths, there must be $R(u, v)$ edge-disjoint paths from u to v in $D'[V]$. The above argument applied to all pairs $(u, v) \in V \times V$ shows that $D'[V]$ is an R -orientation of G . To prove the other direction of the claim, let D be an R -orientation of G . We define an orientation D' of G' by orienting the edges in E according to D , and orienting the other edges as described earlier in this proof. It is easy to see that the obtained digraph D' is a Steiner rooted k -edge-connected orientation of G' . \square

Since R has maximum value 3, the size of G' is polynomial in the size of G . Thus the construction is polynomial and this proves that the Steiner rooted orientation problem is NP-complete. \square

The question remains whether the Steiner rooted k -edge-connected orientation problem is polynomially solvable for fixed k . We do not even know whether it is solvable for $k = 2$ (for $k = 1$ it is easy).

For element-connectivity, can we show that the STEINER ROOTED-ORIENTATION problem is NP-complete. The proof is described in [19].

Theorem 6.4. *Given a graph G , a set of terminals S , and a root vertex $r \in S$, it is NP-complete to determine if G has a Steiner rooted k -element-connected orientation.*

One can consider minimum cost versions of the orientation problems discussed in this section. For each edge, the two different orientations have separate costs, and the cost of an orientation of the graph is the sum of the costs of the oriented edges. It turns out that in both the edge-disjoint and the element-disjoint cases the minimum cost problem is more difficult to approximate than the basic problem. Even for $k = 1$, when the edge-disjoint and element-disjoint problems coincide, we can obtain the following result:

Theorem 6.5. *The MINIMUM COST STEINER ROOTED-ORIENTATION problem (even for $k = 1$) is hard to approximate within a factor of $c \log(n)$ for some constant c unless $P = NP$.*

The proof, which is described in detail in [19], consists of the reducing the SET COVER problem (which is hard to approximate within a factor of $c \log(n)$ for some constant c [1,6]) to the MINIMUM COST STEINER ROOTED-ORIENTATION problem, such that the number of sets in the cover corresponds to the cost of the orientation.

7. Concluding remarks

The questions of generalizing Nash-Williams' theorem to hypergraphs and obtaining graph orientations achieving high vertex-connectivity remain wide open. We believe that substantially new ideas are required to solve these problems. The following problem seems to be a concrete intermediate problem which captures the main difficulty: If S is $2k$ -element-connected in an undirected graph G , is it true that G has a Steiner strongly k -element-connected orientation? We believe that settling it would be a major step towards the above questions.

Acknowledgments

We would like to thank András Frank and Attila Bernáth for useful suggestions. The second author would like to thank Zongpeng Li for motivating the problem and providing references, and Michael Molloy for valuable comments.

References

- [1] N. Alon, D. Moshkovitz, M. Safra, Algorithmic construction of sets for k -restrictions, ACM Trans. Algorithms 2 (2006) 153–177.
- [2] J. Bang-Jensen, G. Gutin, Digraphs: Theory, Algorithms and Applications, Springer-Verlag, London, 2000.
- [3] J. Bang-Jensen, S. Thomassé, Decompositions and orientations of hypergraphs, manuscript, 2001.
- [4] J. Cheriyan, M. Salavatipour, Packing element-disjoint Steiner trees, in: Proceedings of APPROX 2005, in: Lecture Notes in Comput. Sci., vol. 3625, 2005, pp. 52–61.
- [5] S. Deb, M. Effros, T. Ho, D. Karger, R. Koetter, D.S. Lun, M. Médard, N. Ratnakar, Network coding for wireless applications: A brief tutorial, in: Proc. International Workshop on Wireless Ad-hoc Networks (IWWAN), 2005.
- [6] U. Feige, M. Halldorsson, G. Kortsarz, A. Srinivasan, Approximating the domatic number, SIAM J. Comput. 32 (2002) 172–195.

- [7] L.K. Fleischer, K. Jain, D.P. Williamson, An iterative rounding 2-approximation algorithm for the element connectivity problem, in: Proceedings of the 42nd Annual Conference on the Foundation of Computer Science (FOCS), 2001, pp. 339–347.
- [8] A. Frank, On the orientations of graphs, *J. Combin. Theory Ser. B* 28 (1980) 251–261.
- [9] A. Frank, An algorithm for submodular functions on graphs, *Ann. Discrete Math.* (1982) 97–120.
- [10] A. Frank, Applications of submodular functions, in: Surveys in Combinatorics, Keele, 1993, in: London Math. Soc. Lecture Note Ser., vol. 187, Cambridge Univ. Press, Cambridge, 1993, pp. 85–136.
- [11] A. Frank, Orientations of graphs and submodular flows, *Congr. Numer.* 113 (1996) 111–142. Festschrift for C.St.J.A. Nash-Williams.
- [12] A. Frank, T. Király, Combined connectivity augmentation and orientation problems, *Discrete Appl. Math.* 131 (2003) 401–419.
- [13] A. Frank, T. Király, Z. Király, On the orientation of graphs and hypergraphs, *Discrete Appl. Math.* 131 (2003) 385–400.
- [14] S.L. Hakimi, On the degrees of the vertices of a directed graph, *J. Franklin Inst.* 279 (1965) 290–308.
- [15] H.R. Hind, O. Oellermann, Menger-type results for three or more vertices, *Congr. Numer.* 113 (1996) 179–204.
- [16] K. Jain, M. Mahdian, M.R. Salavatipour, Packing Steiner trees, in: Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2003, pp. 266–274.
- [17] K. Jain, I. Mándou, V.V. Vazirani, D.P. Williamson, A primal-dual schema based approximation algorithm for the element connectivity problem, in: Proceedings of the 10th Annual Symposium on Discrete Algorithms (SODA), 1999, pp. 484–489.
- [18] T. Király, Edge-connectivity of undirected and directed hypergraphs, PhD thesis, Eötvös Loránd University, 2003.
- [19] T. Király, L.C. Lau, Approximate min–max theorems for Steiner rooted-orientations of graphs and hypergraphs, EGRES Technical Report No. 2006-13.
- [20] M. Kriesell, Edge-disjoint trees containing some given vertices in a graph, *J. Combin. Theory Ser. B* 88 (2003) 53–63.
- [21] L.C. Lau, An approximate max-Steiner-tree-packing min-Steiner-cut theorem, in: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2004, pp. 61–70.
- [22] L.C. Lau, Packing Steiner Forests, in: Proceedings of the Eleventh Conference on Integer Programming and Combinatorial Optimization (IPCO), 2005, pp. 362–376.
- [23] L. Lovász, Submodular functions and convexity, in: *Mathematical Programming: The State of the Art*, Springer, 1983, pp. 235–257.
- [24] D.S. Lun, M. Médard, R. Koetter, Network coding for efficient wireless unicast, in: *IEEE International Zurich Seminar on Communications*, 2006.
- [25] W. Mader, A reduction method for edge-connectivity in graphs, *Ann. Discrete Math.* 3 (1978) 145–164.
- [26] K. Menger, Zur allgemeinen Kurventheorie, *Fund. Math.* 10 (1927) 95–115.
- [27] C.St.J.A. Nash-Williams, On orientations, connectivity and odd vertex pairings in finite graphs, *Canad. J. Math.* 12 (1960) 555–567.
- [28] C.St.J.A. Nash-Williams, Edge disjoint spanning trees of finite graphs, *J. London Math. Soc.* 36 (1961) 445–450.
- [29] W.T. Tutte, On the problem of decomposing a graph into n connected factors, *J. London Math. Soc.* 36 (1961) 221–230.