Steiner connectivity problems in hypergraphs

Florian Hörsch, Zoltán Szigeti

November 4, 2022

Abstract

We say that a tree T is an S-Steiner tree if $S \subseteq V(T)$ and a hypergraph is an S-Steiner hypertree if it can be trimmed to an S-Steiner tree. We prove that it is NP-hard to decide, given a hypergraph $\mathcal{H}$ and some $S \subseteq V(\mathcal{H})$, whether there is a subhypergraph of $\mathcal{H}$ which is an S-Steiner hypertree. As corollaries, we give two negative results for two Steiner orientation problems in hypergraphs. Firstly, we show that it is NP-hard to decide, given a hypergraph $\mathcal{H}$, some $r \in V(\mathcal{H})$ and some $S \subseteq V(\mathcal{H})$, whether this hypergraph has an orientation in which every vertex of S is reachable from r. Secondly, we show that it is NP-hard to decide, given a hypergraph $\mathcal{H}$ and some $S \subseteq V(\mathcal{H})$, whether this hypergraph has an orientation in which any two vertices in S are mutually reachable from each other. This answers a longstanding open question of the Egerváry Research group. On the positive side, we show that the problem of finding a Steiner hypertree and the first orientation problem can be solved in polynomial time if the number of terminals $|S|$ is fixed.

1 Introduction

This article is concerned with Steiner tree problems in hypergraphs and Steiner connectivity orientation problems in hypergraphs. Any undefined notation can be found in Section 2.

The first part of the article deals with finding Steiner hypertrees in hypergraphs. There exists a rich literature on Steiner tree problems in graphs. For example, the problem of finding a Steiner tree minimizing a given weight function on the edges has been studied to a significant depth. It is well-known to be NP-hard [8], several approximation results are known ([10],[1]) and the problem is known to be fixed parameter tractable when parameterized by the number of terminals [3]. Another branch of research is concerned with the problem of packing Steiner trees. In particular, a famous conjecture of Kriesell [11] remains open but several partial results are known ([6],[16]). The corresponding algorithmic problem has been proven to be NP-hard by Kaski [9]. On the other hand, we can trivially decide in polynomial time whether a given graph G contains a single S-Steiner tree for some given $S \subseteq V(G)$. We here show that this situation drastically changes when considering hypergraphs. When $S = V(\mathcal{H})$ for a hypergraph $\mathcal{H}$, the problem can be solved in polynomial time using the concept of hypergraphic matroids which was introduced by Loréa [12] and exploited by Frank, Király and Kriesell [6]. Dealing with Steiner hypertrees in hypergraphs, we formally consider the following problem:

Steiner Hypertree (SHT):

Input: A hypergraph $\mathcal{H}$, a set $S \subseteq V(\mathcal{H})$.

Question: Does $\mathcal{H}$ contain an S-Steiner hypertree?

On the negative side, we show the following result:

Theorem 1. *SHT is NP-hard.*

On the positive side, we are able to show that the problem can be solved in polynomial time if the number of terminals is fixed.

Theorem 2. *There is a function $f : \mathbb{Z}_{\geq 0} \rightarrow \mathbb{Z}_{\geq 0}$ and an algorithm that solves SHT and runs in $O(f(|S|)n^{|S|}m^2)$.*

In the second part of this article, we apply these results to orientation problems in hypergraphs. We first deal with rooted connectivity. Formally, we consider the following problem:

Steiner Rooted Connected Orientation of Hypergraphs (SRCOH):

Input: A hypergraph $\mathcal{H}$, a vertex $r \in V(\mathcal{H})$, a set $S \subseteq V(\mathcal{H})$.

Question: Is there an orientation $\vec{\mathcal{H}}$ of $\mathcal{H}$ that is (r, S)-Steiner rooted connected?

It turns out that SHT and SRCOH are closely related. In particular, SRCOH can be solved in polynomial time when restricted to graphs. On the other hand, using the above mentioned relation, we prove the following result showing that such an algorithm is unlikely to exist for general hypergraphs.

Theorem 3. *SRCOH is NP-hard.*

Again exploiting this relation, we can conclude the following result from Theorem 2.

Theorem 4. *There is a function $f : \mathbb{Z}_{\geq 0} \rightarrow \mathbb{Z}_{\geq 0}$ and an algorithm that solves SRCOH and runs in $O(f(|S|)n^{|S|+1}m^2)$.*

We also deal with a more symmetric connectivity problem in orientations of hypergraphs. For graphs, a fundamental result of Nash-Williams [13] states that for any positive integer k, a graph has a k-arc-connected orientation if and only if it is $2k$-edge-connected. Actually, Nash-Williams proved an even stronger result giving a complete characterization of the cases when a graph G has an orientation $\vec{G}$ satisfying $\lambda_{\vec{G}}(u, v) \geq r(u, v)$ for some arbitrary symmetric requirement function $r : V(G) \times V(G) \rightarrow \mathbb{Z}_{\geq 0}$. The case of global dyperedge-connectivity in hypergraphs has been dealt with by Frank, Király and Király [5]. The Egerváry Research group [4] raised the question whether these approaches can be combined in order to find orientations satisfying local symmetric dyperedge-connectivity requirements of hypergraphs. We answer this question to the negative even for the very special case when the requirement function r evaluates to 1 when both arguments belong to a fixed set of vertices and 0 otherwise. Formally, we consider the following problem:

Steiner Strongly Connected Orientation of Hypergraphs (SSCOH):

Input: A hypergraph $\mathcal{H}$, a set $S \subseteq V(\mathcal{H})$.

Question: Is there an orientation $\vec{\mathcal{H}}$ of $\mathcal{H}$ that is strongly connected in S?

As a rather simple consequence of Theorem 3, we are able to prove the following:

Theorem 5. *SSCOH is NP-hard.*

In Section 2, we provide some more formal definitions and preliminary results we need for our proofs. In Section 3, we give the reductions proving Theorems 1,3 and 5 and we prove Theorems 2 and 4.

2 Preliminaries

In this section, notation and some auxiliary results are collected. In Section 2.1, we give the necessary definitions and in Section 2.2, we give the preliminary results.

2.1 Definitions

A *hypergraph* $\mathcal{H}$ consists of a vertex set $V(\mathcal{H})$ and a hyperedge set $\mathcal{E}(\mathcal{H})$ where each $e \in \mathcal{E}(\mathcal{H})$ is a subset of $V(\mathcal{H})$ of size at least 2. Throughout the article, we use n and m for the number of vertices and hyperedges of $\mathcal{H}$, respectively. If each hyperdege is of size exactly 2, we call $\mathcal{H}$ a graph. We say that a graph G is a *trimming* of $\mathcal{H}$ if G is obtained from $\mathcal{H}$ by replacing every $e \in \mathcal{E}(\mathcal{H})$ by an edge containing two distinct vertices of e. For some graph G and $X \subseteq V(G)$, we use $d_G(X)$ for the number of edges in G that have exactly one endvertex in X. For a single vertex $v \in V(G)$, we use $d_G(v)$ instead of $d_G(\{v\})$ and call this value the degree of v. For a non-negative integer k, a graph G is called *k-edge-connected* if $d_G(X) \geq k$ for every nonempty $X \subsetneq V(G)$. A *tree* is an edge-minimal 1-edge-connected graph. A *uv-path* is a tree P in which $d_P(u) = d_P(v) = 1$ and $d_P(w) = 2$ for all $w \in V(P) - \{u, v\}$ hold. A *path* is a uv-path for some u and v. Given a *terminal set* $S \subseteq V(T)$, we say that T is an *S-Steiner tree*. Two paths P_1, P_2 are called *internally vertex-disjoint* if $d_{P_1}(v) + d_{P_2}(v) \leq 2$ for all $v \in V(P_1) \cup V(P_2)$. An S-Steiner tree is called *small* if $|V(T)| \leq 2|S| - 2$. A *subdivision* of a graph G_1 in a graph G_2 is a mapping $\phi : V(G_1) \to V(G_2)$ together with a collection of paths $\mathcal{P} = \{P_e : e \in E(G_1)\}$ such that for every $e = uv \in E(G_1)$, P_e is a uv-path and the paths in $\mathcal{P}$ are pairwise internally vertex-disjoint. For a graph G and a vertex v of G of degree 2, we mean by *splitting off* v the graph obtained from G by deleting the vertex v and by adding a new edge between u and w, where uv, vw are the edges of G containing v.

For a hypergraph $\mathcal{H}$, we denote by $G(\mathcal{H})$ the *incidence graph* of $\mathcal{H}$, i.e. the graph which is obtained from $\mathcal{H}$ by replacing every $e \in \mathcal{E}(\mathcal{H})$ by a new vertex z_e and edges vz_e for all $v \in e$. Given a terminal set S, a small S-Steiner tree T and a hypergraph $\mathcal{H}$ with $S \subseteq V(\mathcal{H})$, a subdivision $(\phi, \mathcal{P})$ of T in $G(\mathcal{H})$ is called *special* if $\phi(s) = s$ for all $s \in S$ and $\phi(v) \in V(\mathcal{H})$ for all $v \in V(T)$. An *S-Steiner hypertree* is a hypergraph that can be trimmed to an S-Steiner tree.

A *dypergraph* $\mathcal{D}$ consists of a vertex set $V(\mathcal{D})$ and a dyperedge set $\mathcal{A}(\mathcal{D})$ where each $a \in \mathcal{A}(\mathcal{D})$ is a tuple $(tail(a), head(a))$ where $head(a)$ is a vertex in $V(\mathcal{D})$ and $tail(a)$ is a nonempty subset of $V(\mathcal{D}) - head(a)$.

For some $X \subseteq V(\mathcal{D})$, we say that a dyperedge $a \in \mathcal{A}(\mathcal{D})$ *enters* X if $head(a) \in X$ and $tail(a) - X \neq \emptyset$. We denote by $\delta^-_{\mathcal{D}}(X)$ the set of dyperedges in $\mathcal{A}(\mathcal{D})$ that enter X. We use $d^-_{\mathcal{D}}(X)$ for $|\delta^-_{\mathcal{D}}(X)|$. For $u, v \in V(\mathcal{D})$, we use $\lambda_{\mathcal{D}}(u, v)$ for $\min\{d^-_{\mathcal{D}}(X) : X \subseteq V(\mathcal{D}), v \in X, u \in V(\mathcal{D}) - X\}$. For some $S \subseteq V(\mathcal{D})$, we say that $\mathcal{D}$ is *strongly connected in S* if $\lambda_{\mathcal{D}}(u, v) \geq 1$ for every pair (u, v) in S. We say that $\mathcal{D}$ is *strongly connected* if $\mathcal{D}$ is strongly connected in $V(\mathcal{D})$. For some $u, v \in V(\mathcal{D})$, we say that v is *reachable* from u if $\lambda_{\mathcal{D}}(u, v) \geq 1$. If for some $r \in V(\mathcal{D})$ and $S \subseteq V(\mathcal{D})$, every $v \in S$ is reachable from r, we say that $\mathcal{D}$ is *(r, S)-Steiner rooted connected*. If a dypergraph $\vec{\mathcal{H}}$ is obtained from a hypergraph $\mathcal{H}$ by choosing a head for each hyperedge, we say that $\vec{\mathcal{H}}$ is an *orientation* of $\mathcal{H}$. A dypergraph in which the tail of each dyperedge is of size 1 is called a *digraph*. The dyperedges of a digraph are called *arcs*. For some dypergraph $\mathcal{D}$, we let $D(\mathcal{D})$ denote the digraph in which every dyperedge $a \in \mathcal{A}(\mathcal{D})$ is replaced by a vertex z_a, an arc vz_a for all $v \in tail(a)$ and an arc $z_a head(a)$. The *underlying hypergraph* of $\mathcal{D}$ is the hypergraph on the same vertex set and that contains the hyperedge $tail(a) \cup head(a)$ for all $a \in \mathcal{A}(\mathcal{D})$. If $\mathcal{D}$ is a digraph, we speak of the *underlying graph*. For a non-negative integer k, a digraph D is called *k-arc-connected* if $d^-_D(X) \geq k$ for every nonempty $X \subsetneq V(D)$. We say that a digraph D is a *directed trimming* of $\mathcal{D}$ if D is obtained from $\mathcal{D}$ by replacing every $a \in \mathcal{A}(\mathcal{D})$ by an arc whose head is the head of a and whose tail is a vertex in $tail(a)$. An r-arborescence is a digraph B with $r \in V(B)$ and which is arc-minimal with the property that every vertex in $V(B)$ is reachable from r. A *directed uv-path* is a u-arborescence in which v is the only vertex that is not the tail of any arc. For some $S \subseteq V(B)$, we speak of an *(r, S)-Steiner arborescence*. A *circuit* is a strongly connected digraph satisfying $|A(D)| = |V(D)|$. A *circuit* is a strongly connected digraph D satisfying $|A(D)| = |V(D)|$.

2.2 Preliminaries

For the first reduction in Section 3, we consider the following well-known problem. For a binary variable x, the *literals* over x are x and $\bar{x}$, the negation of x.

3SAT

Input: A set of binary variables X, a set of clauses $\mathcal{C}$ each of which contains 3 literals over X.

Question: Is there an assignment $\phi : X \to \{TRUE, FALSE\}$ such that every clause of $\mathcal{C}$ contains at least one true literal?

For the first reduction, we need the following well-known result, see [8].

Theorem 6. *3SAT is NP-hard.*

For the second reduction, we need the following result that can be found in [6].

Proposition 1. *Let $\mathcal{D}$ be a dypergraph, $r \in V(\mathcal{D})$ and $S \subseteq V(\mathcal{D})$. Then $\mathcal{D}$ contains a subdypergraph that can be transformed into an (r, S)-Steiner arborescence by a directed trimming if and only if all vertices in S are reachable from r in $\mathcal{D}$.*

For the last reduction, we need the following result.

Proposition 2. *Let $\mathcal{D}$ be a dypergraph. Then for any pair (u, v) of vertices in $V(\mathcal{D})$, we have $\lambda_{D(\mathcal{D})}(u, v) = \lambda_{\mathcal{D}}(u, v)$.*

Proof. First let $X \subseteq V(D(\mathcal{D}))$ with $v \in X, u \in V(D(\mathcal{D})) - X$ and $d^-_{D(\mathcal{D})}(X) = \lambda_{D(\mathcal{D})}(u, v)$. Let $X' = X \cap V(\mathcal{D})$. Then for every $a \in \delta^-_{\mathcal{D}}(X')$, we have $head(a) \in X$ and $tail(a) - X \neq \emptyset$, so either the arc $z_a head(a)$ enters X in $D(\mathcal{D})$ or the arc wz_a enters X in $D(\mathcal{D})$ for some $w \in tail(a)$. Hence $\lambda_{\mathcal{D}}(u, v) \leq d^-_{\mathcal{D}}(X') \leq d^-_{D(\mathcal{D})}(X) = \lambda_{D(\mathcal{D})}(u, v)$.

Now let $X \subseteq V(\mathcal{D})$ with $v \in X, u \in V(\mathcal{D}) - X$ and $d^-_{\mathcal{D}}(X) = \lambda_{\mathcal{D}}(u, v)$. Let $X' \subseteq V(D(\mathcal{D}))$ be the set that contains X and the vertex z_a for all $a \in A(\mathcal{D})$ for which $tail(a) \subseteq X$ holds. Now every arc entering X' in $D(\mathcal{D})$ is of the form $z_a head(a)$ such that a enters X in $\mathcal{D}$. Hence $\lambda_{D(\mathcal{D})}(u, v) \leq d^-_{D(\mathcal{D})}(X') \leq d^-_{\mathcal{D}}(X) = \lambda_{\mathcal{D}}(u, v)$. □

Proposition 3. *Let $\mathcal{H}$ be a hypergraph and let $e \in \mathcal{E}(\mathcal{H})$. Further, let $\vec{\mathcal{H}}_0$ be an orientation of $\mathcal{H}$ and let $\vec{e}_0$ be the orientation of e in $\vec{\mathcal{H}}_0$. Suppose that there is some $x \in tail(\vec{e}_0)$ such that $\lambda_{\vec{\mathcal{H}}_0}(head(\vec{e}_0), x) \geq 1$. Then there is an orientation $\vec{\mathcal{H}}_1$ of $\mathcal{H}$ such that $\lambda_{\vec{\mathcal{H}}_1}(u, v) = \lambda_{\vec{\mathcal{H}}_0}(u, v)$ for every pair (u, v) in $V(\mathcal{H})$ and $head(\vec{e}_1) = x$ where $\vec{e}_1$ is the orientation of e in $\vec{\mathcal{H}}_1$.*

Proof. We obtain by Proposition 2 that $D(\vec{\mathcal{H}}_0)$ contains a directed path from $head(\vec{e}_0)$ to x. As this directed path contains none of the arcs $xz_{\vec{e}_0}$ and $z_{\vec{e}_0} head(\vec{e}_0)$, we obtain that $D(\vec{\mathcal{H}}_0)$ contains a circuit containing the arcs $xz_{\vec{e}_0}$ and $z_{\vec{e}_0} head(\vec{e}_0)$. Let D_1 be the digraph obtained from $D(\vec{\mathcal{H}}_0)$ by reversing all the arcs of this cycle. Note that there is an orientation $\vec{\mathcal{H}}_1$ of $\mathcal{H}$ such that $D(\vec{\mathcal{H}}_1) = D_1$. Observe that $head(\vec{e}_1) = x$. Further, for all $u, v \in V(\mathcal{H})$, by Proposition 2, we have $\lambda_{\vec{\mathcal{H}}_1}(u, v) = \lambda_{D(\vec{\mathcal{H}}_1)}(u, v) = \lambda_{D(\vec{\mathcal{H}}_0)}(u, v) = \lambda_{\vec{\mathcal{H}}_0}(u, v)$. □

For the proof of Theorem 2, we need the following result due to Kawarabayashi, Kobayashi and Reed [7] which improves upon an earlier result of Robertson and Seymour [15].

Lemma 1. *Let G be a graph and (u_1, v_1), $\ldots$, (u_k, v_k) pairs of vertices in $V(G)$. Then there exist a function $f : \mathbb{Z}_{\geq 0} \to \mathbb{Z}_{\geq 0}$ and an algorithm that runs in $O(f(k)n^2)$ and decides whether there is a set of internally vertex-disjoint paths $P_1, \ldots, P_k$ such that P_i is a $u_i v_i$-path for $i = 1, \ldots, k$.*

We further need the following well-known property of trees.

Proposition 4. *Let T be a tree. Then the number of vertices in $v \in V(T)$ with $d_T(v) = 1$ is at least two more than the number of vertices in $V(T)$ with $d_T(v) \geq 3$.*

We further require the following classic theorem due to Cayley [2].

Theorem 7. *The number of distinct labelled trees on a ground set of n vertices is n^{n-2}.*

3 Proofs

In this section, we give the proofs of the main theorems of this article.

3.1 Steiner hypertrees

This section is dedicated to proving the results on finding Steiner hypertrees in a given hypergraph. We first prove the negative result when the size of terminal is not fixed in Section 3.1.1 and then prove the positive result for a fixed number of terminals in Section 3.1.2.

3.1.1 The proof of Theorem 1

Proof. (of Theorem 1) We prove this by a reduction from 3SAT. Let $(X, \mathcal{C})$ be an instance of 3SAT. We now create an instance $(\mathcal{H}, S)$ of SHT. For every $x \in X$, we let $V(\mathcal{H})$ contain 2 vertices w_x and $w_{\bar{x}}$. Next, for every $C \in \mathcal{C}$, we let $V(\mathcal{H})$ contain a vertex z_C. Further, we let $V(\mathcal{H})$ contain one more vertex a. Let $W = \bigcup_{x \in X}\{w_x, w_{\bar{x}}\}$ and $Z = \{z_C : C \in \mathcal{C}\}$. For every $x \in X$, we let $\mathcal{E}(\mathcal{H})$ contain a hyperedge $e_x = \{a, w_x, w_{\bar{x}}\}$. Next, for every $C \in \mathcal{C}$, we let $\mathcal{E}(\mathcal{H})$ contain a hyperedge $e_C = \{\{w_\ell : \ell \in C\} \cup z_C\}$. Finally, we set $S = Z \cup \{a\}$. This finishes the description of $(\mathcal{H}, S)$.

An illustration can be found in Figure 1.

We now prove that $(\mathcal{H}, S)$ is a positive instance of SHT if and only if $(X, \mathcal{C})$ is a positive instance of 3SAT.

First suppose that $(X, \mathcal{C})$ is a positive instance of 3SAT, so there is an assignment $\phi : X \to \{True, False\}$ that satisfies every clause of $\mathcal{C}$.

It suffices to prove that $\mathcal{H}$ can be trimmed to a $(Z \cup a)$-Steiner tree. For every $x \in X$, we trim e_x to an edge e'_x where $e'_x = aw_x$ if $\phi(x) = TRUE$ and $e'_x = aw_{\bar{x}}$ if $\phi(x) = FALSE$. Now consider some $C \in \mathcal{C}$. As ϕ is satisfying, we can choose some $x \in X$ such that either $x \in C$ and $\phi(x) = TRUE$ or $\bar{x} \in C$ and

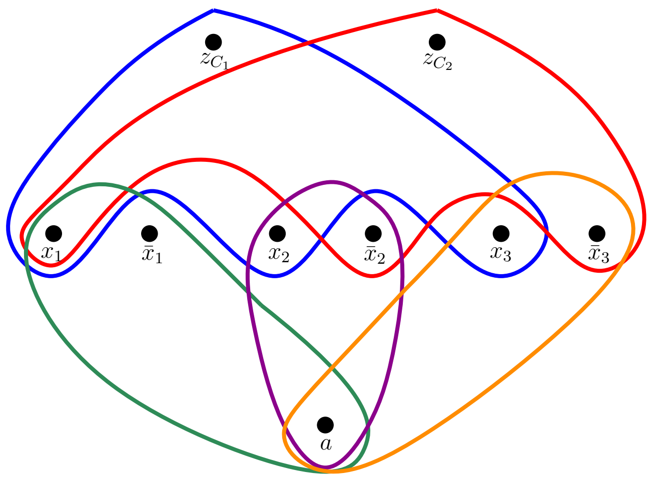

Figure 1: An example of the construction of $(\mathcal{H}, S)$ for a 3SAT formula with $X = \{x_1, x_2, x_3\}$ and $\mathcal{C} = \{C_1 = \{x_1, x_2, x_3\}, C_2 = \{x_1, \bar{x}_2, \bar{x}_3\}\}$.

$\phi(x) = FALSE$ hold. We trim e_C to an edge e'_C where $e'_C = w_x z_C$ in the former case and $e'_C = w_{\bar{x}} z_C$ in the latter case. Now let T be the graph that contains all vertices contained in e'_x for some $x \in X$ and all the vertices contained in e'_C for some $C \in \mathcal{C}$ and whose edge set is $\{e'_x : x \in X\} \cup \{e'_C : C \in \mathcal{C}\}$. Clearly, we have $Z \cup a \subseteq V(T)$. It also follows directly from the construction that T is a tree. Hence T is a $(Z \cup a)$-Steiner tree and so $\mathcal{H}$ is a $(Z \cup a)$-Steiner hypertree.

Now suppose that $(\mathcal{H}, Z \cup a)$ is a positive instance of SHT, so there is a subhypergraph $\mathcal{T}$ of $\mathcal{H}$ which can be trimmed to a $(Z \cup a)$-Steiner hypertree T. We now define a truth assignment $\phi : X \to \{TRUE, FALSE\}$ in the following way: if the hyperedge e_x is contained in $\mathcal{E}(\mathcal{T})$ and is trimmed to the edge aw_x in T, we set $\phi(x) = TRUE$. Otherwise, we set $\phi(x) = FALSE$.

In order to prove that ϕ is satisfying, first observe that as T is a $(Z \cup a)$-Steiner tree and as z_C is contained in only one hyperedge of $\mathcal{E}(\mathcal{H})$, we obtain that $d_T(z_C) = 1$ for all $C \in \mathcal{C}$. Now fix some $C^* \in \mathcal{C}$. Let b be the unique vertex such that the edge e_{C^*} is trimmed to bz_{C^*} in T. As T is a $(Z \cup a)$-Steiner tree, we obtain that T contains a $z_{C^*}a$-path P. As $d_T(z_C) = 1$ for all $C \in \mathcal{C}$, it follows $V(P) \cap Z = z_{C^*}$. By construction, this yields that $P = abz_{C^*}$. It follows that the unique hyperedge in $\mathcal{E}(\mathcal{H})$ containing a and b is trimmed to ab in T. If $b = w_x$ for some $x \in X$, we obtain that $x \in C$ and $\phi(x) = TRUE$ by definition of ϕ. If $b = w_{\bar{x}}$ for some $x \in X$, we obtain that $\bar{X} \in C$ and $\phi(x) = FALSE$ by definition of ϕ. In either case, C^* is satisfied by ϕ. As C^* was chosen arbitrarily, ϕ is a satisfying assignment for $(X, \mathcal{C})$.

As the size of $\mathcal{H}$ is clearly polynomial in the size of $(X, \mathcal{C})$ and by Proposition 6, the statement follows. $\square$

3.1.2 Polynomial algorithm for a fixed number of terminals

This section is dedicated to proving Theorem 2. We first show that in order to do so, it suffices to consider a related problem in the incidence graph of the given hypergraph.

Lemma 2. *Let $\mathcal{H}$ be a hypergraph and $S \subseteq V(\mathcal{H})$. Then $\mathcal{H}$ contains an S-Steiner hypertree if and only if $G(\mathcal{H})$ contains an S-Steiner tree T that satisfies $d_T(z_e) = 2$ for all $e \in \mathcal{E}(\mathcal{H})$ with $z_e \in V(T)$.*

Proof. First suppose that $\mathcal{H}$ contains an S-Steiner hypertree $\mathcal{T}$ that can be trimmed to an S-Steiner tree T. Let T' be obtained from T by subdividing every edge $\tilde{e} \in E(T)$, creating the vertex z_e, where $e \in \mathcal{E}(\mathcal{H})$ is the edge from which $\tilde{e}$ is obtained by trimming. Observe that T' is a subgraph of $G(\mathcal{H})$. By construction, we have $S \subseteq V(T) \subseteq V(T')$ and $d_{T'}(z_e) = 2$ for all $e \in \mathcal{E}(\mathcal{H})$ with $z_e \in V(T)$. Finally, as T is a tree, so is T'.

Now suppose that $G(\mathcal{H})$ contains an S-Steiner tree T' that satisfies $d_{T'}(z_e) = 2$ for all $e \in \mathcal{E}(\mathcal{H})$ with $z_e \in V(T')$. Let T be the graph with $V(T) = V(T') \cap V(\mathcal{H})$ and which contains an edge $\tilde{e} = uv$ for all $u, v \in V(T)$ for which there is some $e \in \mathcal{E}(\mathcal{H})$ with $uz_e, vz_e \in E(T')$. Observe that $\tilde{e}$ can be obtained from e by trimming and hence T can be obtained from a subhypergraph $\mathcal{T}$ of $\mathcal{H}$ by trimming. As T is a obtained from T' by contracting edges, we obtain that T is a tree and by construction, we have $S \subseteq V(T') \cap V(\mathcal{H}) = V(T)$. Hence T is an S-Steiner tree and so $\mathcal{T}$ is an S-Steiner hypertree. □

We next show that it suffices to deal with small S-Steiner trees instead of arbitrary ones which is important to limit the number of possible choices.

Lemma 3. *Let S be a set and $\mathcal{H}$ a hypergraph with $S \subseteq V(\mathcal{H})$. Then $G(\mathcal{H})$ contains an S-Steiner tree T that satisfies $d_T(z_e) = 2$ for all $e \in \mathcal{E}(\mathcal{H})$ with $z_e \in E(T)$ if and only if $G(\mathcal{H})$ contains a small S-Steiner tree as a special subdivision.*

Proof. First suppose that $G(\mathcal{H})$ contains a small S-Steiner tree T as a special subdivision. Let T' be the subgraph of $G(\mathcal{H})$ with $V(T') = \bigcup_{P \in \mathcal{P}} V(P)$ and $E(T') = \bigcup_{P \in \mathcal{P}} E(P)$. As the paths of $\mathcal{P}$ are pairwise internally vertex-disjoint, we obtain that T' can be obtained from T by subdividing edges several times. Hence T' is a tree. Further, we have $S \subseteq V(T) \subseteq V(T')$. Finally, as $V(T) \subseteq V(\mathcal{H})$ and the paths in $\mathcal{P}$ are pairwise internally vertex-disjoint, we obtain that $d_{T'}(z_e) = 2$ for all $e \in \mathcal{E}(\mathcal{H})$ with $z_e \in E(T')$. Hence T' is an S-Steiner tree that satisfies $d_{T'}(z_e) = 2$ for all $e \in \mathcal{E}(\mathcal{H})$ with $z_e \in E(T')$.

Now suppose that $G(\mathcal{H})$ contains an S-Steiner tree T' that satisfies $d_{T'}(z_e) = 2$ for all $e \in \mathcal{E}(\mathcal{H})$ with $z_e \in V(T')$. Choosing T' minimum, we may suppose that every vertex of degree 1 of T' is contained in S. Let T be obtained from T' by splitting off vertices of degree 2 which are not contained in S. Observe that, as $d_{T'}(z_e) = 2$ for all $e \in \mathcal{E}(\mathcal{H})$ with $z_e \in V(T')$, we have $V(T) \subseteq V(\mathcal{H})$. Further, by construction, we have $d_T(v) \geq 3$ for all $v \in V(T) - S$. By Proposition 4, we obtain $|V(T)| \leq 2|S| - 2$. Hence T is a small S-Steiner tree. In order to see that $G(\mathcal{H})$ contains T as a special subdivision, let ϕ be the identity map on $V(T)$. Further, for every $e = uv \in E(T)$, let P_e be the unique uv-path in T' and let $\mathcal{P} = \{P_e : e \in E(T)\}$. As T' is a tree, the P_e are pairwise internally vertex-disjoint and hence $(\phi, \mathcal{P})$ is a special subdivision of T in $G(\mathcal{H})$. □

We now show that a special subdivision of a fixed small S-Steiner tree can be found efficiently.

Lemma 4. *Let S be a set with $|S| = k$, T a small S-Steiner tree and $\mathcal{H}$ a hypergraph. Then there exist a function $f : \mathbb{Z}_{\geq 0} \to \mathbb{Z}_{\geq 0}$ and an algorithm that tests whether $G(\mathcal{H})$ contains T as a special subdivision and runs in $O(f(k)n^k m^2)$.*

Proof. When trying to find a special subdivision $(\phi, \mathcal{P})$ of T in $G(\mathcal{H})$, first observe that, as T is small, there are at most $n^{|V(T)-S|} \leq n^{k-2}$ possibilities to choose ϕ. We now fix some ϕ_0. In order to test whether there exists a special subdivision $(\phi, \mathcal{P})$ of T in $G(\mathcal{H})$ with $\phi = \phi_0$, it suffices to decide whether there exists a set of pairwise internally vertex-disjoint paths $\{P_e : e \in E(T)\}$ in $G(\mathcal{H})$ such that P_e is a uv-path for every $e = uv \in E(T)$. By Lemma 1, there is a function $f' : \mathbb{Z}_{\geq 0} \to \mathbb{Z}_{\geq 0}$ that tests this property and runs in $O(f'(|E(T)|)(n+m)^2)$. We obtain a total running time of $O(n^{k-2}(f'(|E(T)|)(n+m)^2)) = O(f'(2k)n^k m^2)$. □

We next prove that the number of small S-Steiner trees on a fixed ground set is bounded.

Lemma 5. *Let S be a set with $|S| = k$ for some integer $k \geq 2$ and X a set with $|X| = k - 2$ and $X \cap S = \emptyset$. Then there are at most $(2k-2)^{2k-3}$ labelled small Steiner trees T with $V(T) \subseteq S \cup X$.*

Proof. By Theorem 7, the lemma follows from the fact that the number of labelled trees on at most n vertices is at most $\sum_{\mu=1}^{n} \binom{n}{\mu} \mu^{\mu-2} < \sum_{\mu=1}^{n} n^{n-\mu} \cdot n^{\mu-2} = \sum_{\mu=1}^{n} n^{n-2} = n^{n-1}$. □

We are now ready to prove Theorem 2.

Proof. (of Theorem 2) Let $\mathcal{H}$ be a hypergraph and $S \subseteq V(\mathcal{H})$ with $|S| = k$. We need to decide in $O(f(k)n^k m^2)$ time whether $\mathcal{H}$ contains an S-Steiner hypertree. By Lemmas 2 and 3, it suffices to show that we can decide in $O(f(k)n^k m^2)$ time whether $G(\mathcal{H})$ contains a small S-Steiner tree as a special subdivision. In order to test whether this is the case, we can restrict ourselves to trees T that satisfy $V(T) \subseteq S \cup X$ for some fixed X with $|X| = k - 2$. By Lemma 5, there is a function $f_1 : \mathbb{Z}_{\geq 0} \to \mathbb{Z}_{\geq 0}$ such that there are at most $f_1(k)$ such trees. Next, by Lemma 4, there is a function $f_2 : \mathbb{Z}_{\geq 0} \to \mathbb{Z}_{\geq 0}$ such that, for each of these trees T, we can decide in $O(f_2(k)n^k m^2)$ time if $G(\mathcal{H})$ contains T as a special subdivision. We obtain a total running time of $O(f_1(k)f_2(k)n^k m^2)$. □

3.2 Steiner rooted-connected orientations

This section is dedicated to proving Theorems 3 and 4. The following is the key ingredient. Its graphic version is trivial and easily implies the hypergraphic one.

Lemma 6. *Let $\mathcal{H}$ be a hypergraph, $r \in V(\mathcal{H})$ and $S \subseteq V(\mathcal{H})$. Then $\mathcal{H}$ has a (r, S)-Steiner rooted connected orientation if and only if $\mathcal{H}$ contains a $(S \cup r)$-Steiner hypertree.*

Proof. First suppose that $\mathcal{H}$ contains a $(S \cup r)$-Steiner hypertree $\mathcal{T}$ that can be trimmed to an $(S \cup r)$-Steiner tree T. Then there is an orientation $\vec{T}$ of T that is an (r, S)-Steiner arborescence. Now consider the orientation $\vec{\mathcal{H}}$ in which for every hyperedge in $\mathcal{E}(\mathcal{T})$ we choose as its head the head of the corresponding arc in $\vec{T}$ and assign an arbitrary orientation to all remaining hyperedges. Let $\vec{\mathcal{T}}$ be the orientation of $\mathcal{T}$ obtained from $\vec{\mathcal{H}}$ by a restriction to the dyperedges whose corresponding hyperedges are contained in $\mathcal{E}(\mathcal{T})$. Then $\vec{\mathcal{T}}$ can be transformed into $\vec{T}$ by a directed trimming. As $\vec{\mathcal{T}}$ is a subdypergraph of $\vec{\mathcal{H}}$, we obtain by Proposition 1 that all vertices in S are reachable from r in $\vec{\mathcal{H}}$.

Now suppose that $\mathcal{H}$ has an (r, S)-Steiner rooted connected orientation $\vec{\mathcal{H}}$. By Proposition 1, $\vec{\mathcal{H}}$ contains a subdypergraph $\vec{\mathcal{T}}$ that can be transformed into an (r, S)-Steiner arborescence $\vec{T}$ by a directed trimming. Let $\mathcal{T}$ be the underlying hypergraph of $\vec{\mathcal{T}}$ and T be the underlying graph of $\vec{T}$. Then T is an $(S \cup r)$-Steiner tree, T can be obtained from $\mathcal{T}$ by trimming and $\mathcal{T}$ is a subhypergraph of $\mathcal{H}$. This finishes the proof. $\square$

Lemma 6 and Theorem 1 imply Theorem 3. Further, Lemma 6 and Theorem 2 imply Theorem 4.

3.3 Steiner strongly connected orientations

We here conclude Theorem 5 from Theorem 3.

Proof. (of Theorem 5) We prove this by a reduction from SRCOH. Let $(\mathcal{H}, r, S)$ be an instance of SRCOH. Let $\mathcal{H}'$ be obtained from $\mathcal{H}$ by adding the hyperedge $e^* = S \cup r$ and let $S' = S \cup r$. We will prove that $(\mathcal{H}', S')$ is a positive instance of SSCOH if and only if $(\mathcal{H}, r, S)$ is a positive instance of SRCOH.

First suppose that $(\mathcal{H}, r, S)$ is a positive instance of SRCOH, so there is an orientation $\vec{\mathcal{H}}$ of $\mathcal{H}$ in which all vertices of S are reachable from r. Let the orientation $\vec{\mathcal{H}'}$ of $\mathcal{H}'$ be obtained by choosing r as the head of $\vec{e^*}$ and giving all other hyperedges the orientation they have in $\vec{\mathcal{H}}$. As $\vec{\mathcal{H}}$ is a subdypergraph of $\vec{\mathcal{H}'}$, we obtain that all vertices in S are reachable from r in $\vec{\mathcal{H}'}$. Further, due to the orientation of e^*, r is also reachable from all vertices in S in $\vec{\mathcal{H}'}$. Hence $\vec{\mathcal{H}'}$ is strongly connected in S'.

Now suppose that $(\mathcal{H}', S')$ is a positive instance of SSCOH, so there is an orientation $\vec{\mathcal{H}'}$ of $\mathcal{H}'$ which is strongly connected in S'. By Proposition 3, we may suppose that r is the head of $\vec{e^*}$ in $\vec{\mathcal{H}'}$. Now let $\vec{\mathcal{H}}$ be the orientation of $\mathcal{H}$ which is obtained from $\vec{\mathcal{H}}'$ by deleting the dyperedge corresponding to e^*. Consider some $s \in S$. As $\vec{\mathcal{H}}'$ is strongly connected in S, there is a subdypergraph $\vec{\mathcal{T}}$ of $\vec{\mathcal{H}}'$ that can be trimmed to a directed rs-path. Clearly, this path does not contain an arc entering r and hence $\vec{\mathcal{T}}$ does not contain the dyperedge corresponding to e^*. It follows that $\vec{\mathcal{T}}$ is also a subdypergraph of $\vec{\mathcal{H}}$. As s was chosen arbitrarily, we obtain that all vertices in S are reachable from r in $\vec{\mathcal{H}}$.

As the size of $(\mathcal{H}', S')$ is clearly polynomial in the size of $(\mathcal{H}, r, S)$ and by Theorem 3, the statement follows.

4 Conclusion

We have shown hardness results for the problem of finding Steiner hypertrees in a given hypergraph and two Steiner orientations problems in hypergraphs. We further show that two of these problems become easy when we fix the number of terminals. For these two problems, one may ask whether they are fixed parameter tractable when parameterized by the number of terminals. More concretely, we pose the following problem:

Problem 1. *Is there an algorithm that solves SHT and runs in $O(f(|S|)n^{O(1)})$ for some computable function $f : \mathbb{Z}_{\geq 0} \to \mathbb{Z}_{\geq 0}$?*

Observe that due to Lemma 6, the analogous problem for SRCOH is equivalent to Problem 1.

For SSCOH even the question whether an analogous result to Theorems 2 and 4 exists remains open.

Problem 2. *Is there a polynomial time algorithm that solves SSCOH when $|S|$ is fixed?*

Finally, we could ask whether Theorem 2 can be generalized to finding packings of hypertrees.

Problem 3. *Can we decide in polynomial time whether a given hypergraph $\mathcal{H}$ contains a packing of q hyperedge-disjoint S-Steiner hypertrees when $|S|$ and q are fixed?*

□

References

[1] J. Byrka, F. Grandoni, T. Rothvoß, L. Sanitá, An improved LP-based approximation for Steiner tree, Proceedings of the 42nd ACM Symposium on Theory of Computing, pp. 583–592, 2010,

[2] A. Cayley, A theorem on trees, Quart. J. Pure Appl. Math. (23), 376–378, 1889,

[3] S. E. Dreyfus, R. A. Wagner, The Steiner problem in graphs, Networks. 1 (3), 195–207, 1971,

[4] Egres Open Problems, Problem: "*Well-balanced orientations of hypergraphs*", `http://lemon.cs.elte.hu/egres/open/Well-balanced_orientations_of_hypergraphs`,

[5] A. Frank, T. Király, Z. Király, On the orientation of graphs and hypergraphs, Discrete Appl. Math. 131 (2), 385–400, 2003,

[6] A. Frank, T. Király, M. Kriesell, On decomposing a hypergraph into k connected sub-hypergraphs, Discrete Applied Mathematics 131, 373–383, 2003,

[7] K. Kawarabayashi, Y. Kobayashi, B. A. Reed, The disjoint paths problem in quadratic time. J. Comb. Theo., Ser. B 102(2) (2012) 424–435,

[8] R. M. Karp, Reducibility Among Combinatorial Problems, In: R. E. Miller; J. W. Thatcher; J.D. Bohlinger, Complexity of Computer Computations, New York: Plenum, 85–103, 1972,

[9] P. Kaski, Packing Steiner trees with identical terminal sets. Inf. Process. Lett., 91(1):1–5, 2004,

[10] L. Kou, G. Markowsky, L. Berman, A fast algorithm for Steiner trees, Acta Informatica. 15 (2), 141–145, 1981,

[11] M. Kriesell, Edge disjoint trees containing some given vertices in a graph, Journal of Combinatorial Theory B 88, 53–65, 2003,

[12] M. Lorea, Hypergraphes et matroides, Cahiers Centre Etudes Rech. Oper. 17, 289–291, 1975,

[13] C. St. J. A. Nash–Williams, On orientations, connectivity, and odd vertex pairings in finite graphs, Canad. J. Math., 12, 555–567,1960,

[14] H. E. Robbins, A theorem on graphs with an application to a problem of traffic control, American Math. Monthly 46, 281–283, 1939,

[15] N. Robertson, P. D. Seymour, Graph minors. XIII. The disjoint paths problem, J. Comb. Theo., Ser. B, 63 (1) (1995) 65–110.

[16] D.B. West, H. Wu, Packing of Steiner trees and S-connectors in graphs, Journal of Combinatorial Theory B 102, 186–205, 2012.