

A Linear Programming Based Approach to the Steiner Tree Problem with a Fixed Number of Terminals

Matias Siebert¹, Shabbir Ahmed¹, and George Nemhauser¹

¹H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA

Abstract

We present a set of integer programs (IPs) for the Steiner tree problem with the property that the best solution obtained by solving all, provides an optimal Steiner tree. Each IP is polynomial in the size of the underlying graph and our main result is that the linear programming (LP) relaxation of each IP is integral so that it can be solved as a linear program. However, the number of IPs grows exponentially with the number of terminals in the Steiner tree. As a consequence, we are able to solve the Steiner tree problem by solving a polynomial number of LPs, when the number of terminals is fixed.

1 Introduction

Given an undirected graph $G = (V, E)$, where V is the set of nodes, E is the set of edges, $c_e \geq 0$ is the cost of each edge $e \in E$, and $R \subseteq V$ is a set of *terminal nodes*, a Steiner tree is a subgraph of G , which is a tree such that all terminal nodes are connected. The Steiner tree problem is to find such a tree that minimizes the sum of the edge costs over all edges in the tree. The tree may contain nodes in $V \setminus R$, which are called *Steiner nodes*. The Steiner tree problem is NP-Hard [17], in fact, it is even NP-Hard to find approximate solutions whose cost is within a factor $\frac{96}{95}$ of the cost of an optimal solution [3, 8].

An important line of research has been to study integer linear programming formulations (IP) for this problem. In [15] the authors provide a catalog of Steiner tree formulations, and show the equivalence of some of these formulations. In [14] the author studies the vertex-weighted version of the undirected Steiner tree problem, and presents a complete description of the polytope when the graph is *series-parallel*. By projecting this formulation, some facet-defining inequalities for the Steiner tree polytope are obtained. In [20], the authors compare different formulations of the Steiner tree problem in terms of the strength of their linear relaxations (LP), and propose a new polynomial size formulation with a stronger LP relaxation than the ones analyzed in their paper.

In every Steiner tree, we can pick an arbitrary node $r \in R$ and find a direction of the edges, such that we can construct a unique arborescence rooted at r whose leaf nodes correspond to nodes in $R \setminus \{r\}$. Note that an r -arborescence is a directed tree where every node in the arborescence is reachable from the root node r . Consequently, the Steiner tree problem can be modeled as an r -arborescence whose leaves are nodes in $R \setminus \{r\}$ and where all terminal nodes are reached from r . A well studied formulation that uses this fact, is the *Bidirected Cut* formulation [11, 25]. This is a compact and simple formulation, whose integrality gap is known to be at most 2, but is believed to be close to one. The best known lower bound on the integrality gap for this formulation is $\frac{36}{31} \approx 1.16$ [6]. An upper bound on the integrality gap for this formulation has been found for special cases. For *quasi-bipartite* graphs, the integrality gap is upper bounded by $\frac{3}{2}$ [22], and for *claw-free* graphs it is upper bounded by $\ln(4)$ [12].

Another well studied formulation, is the so called *Hypergraphic* formulation [19, 21, 24]. This formulation has one variable for each *component* of the graph, where a *component* is a tree whose leaves correspond to

terminal nodes. The number of variables and constraints for this formulation are proportional to the number of *components*, which is exponential in the number of terminal nodes. On the positive side, we have that the *Hypergraphic* formulation is stronger than the *Bidirected Cut* formulation [21]. The integrality gap of the *Hypergraphic* formulation is lower bounded by $\frac{8}{7} \approx 1.14$ [19] and upper bounded by $\ln(4)$ [16]. Moreover, in [12] the authors show that in *claw-free* graphs both formulations are equivalent. On the down side, solving the *Hypergraphic* formulation is strongly NP-Hard [16]. To address this problem, researchers have proposed to solve it only considering components with at most k leaves, where k is a fixed value. This considerably reduces the number of constraints and variables, and the solution of this restricted formulation is proven to be a factor of ρ_k of an optimum solution, where $\rho_k \leq 1 + \frac{1}{\lceil \log_2(k) \rceil}$ [5]. A complete review of the *Hypergraphic* formulation and its variants can be found in [7].

Finally, in [10] the authors propose an algorithm that finds an optimal solution for the Steiner tree problem in $\mathcal{O}(n^3 + n^2 2^b + n 3^b)$, where $n = |V|$, $m = |A|$, and $b + 1 = |R|$. The result of this paper suggests that, for a fixed number of terminals, there may exist a polynomial size LP formulation of the Steiner tree problem, which is the motivation for the results presented in this paper.

Our main contribution is a set of independent IPs with the property that the best solution obtained by solving all, provides an optimal Steiner tree. Each IP is polynomial in the size of the underlying graph and our main result is that the LP relaxation of each IP is integral. Furthermore, the set of LPs can be solved in parallel. The drawback, is that the number of LPs grows exponentially with the number of terminal nodes so the complete algorithm is polynomial only for a fixed number of terminals.

The remainder of this paper is organized as follows. Section 2 analyzes the structure of Steiner trees, and also introduces definitions used in the paper. Section 3 presents the proposed model, the proof of integrality, and an analysis of the size of the problem. In Section 4 we present computational experiments obtained by solving the Steiner tree problem using our proposed formulation. Section 5 gives conclusions.

2 Solution Structure

As mentioned before, several formulations use the fact that every Steiner tree contains an r -arborescence, some of which use a flow-based formulation, and others a cut-based formulation. Both of those common approaches work with the directed graph version D of G . In our approach, we exploit two main properties that every solution must have.

1. We can always pick an arbitrary node $r \in R$ as our root node, and find the min-cost r -arborescence that spans $R \setminus \{r\}$. Moreover, this r -arborescence can be modeled as a Min-cost Multi-Commodity Network Flow, where we have one commodity for every node in $R \setminus \{r\}$, and we want to send 1 unit of flow from r to the corresponding node $i \in R \setminus \{r\}$.
2. Since the solution must correspond to an r -arborescence, there must be exactly one path from the root node r to each one of the nodes in $R \setminus \{r\}$, which means that if any subset S of the corresponding commodities of nodes in $R \setminus \{r\}$ share a path from r to node $i \in V$, and then split in node i , then they will never meet again. For example, if $S = \{k_1, k_2, k_3\}$ and in node i the set S splits into $S_1 = \{k_1, k_2\}$ and $S_2 = \{k_3\}$, then k_1 and k_2 will never share an arc again with k_3 .

Using these two properties, our formulation is based on two type of decisions, (*i*) when we split the subsets of commodities, and (*ii*) into which subsets of commodities. Note that all the commodities have the same source node, so they all start together. At some point, the set of all commodities will split into different subsets of commodities, and those subsets will also split into other subsets, and so on, until we have each commodity by itself. Our idea is to include a variable that tells us where a subset splits, and the partition that it uses. For instance, say that at some point we have the set $S = \{k_1, k_2, k_3\}$ of commodities that share

a path. Eventually, S is split, so the model has to decide where the split is going to happen. In addition, the model has to decide the way that S is split. Note that there are 4 possible partitions of S .

- (i) $\{k_1, k_2, k_3\} \rightarrow \{\{k_1, k_2\}, \{k_3\}\}$
- (ii) $\{k_1, k_2, k_3\} \rightarrow \{\{k_1, k_3\}, \{k_2\}\}$
- (iii) $\{k_1, k_2, k_3\} \rightarrow \{\{k_2, k_3\}, \{k_1\}\}$
- (iv) $\{k_1, k_2, k_3\} \rightarrow \{\{k_1\}, \{k_2\}, \{k_3\}\}$

Two important observations are, first, we only consider proper partitions (i.e. a set is not a partition of itself), and second, the next partition that we can use will depend on the way that we partitioned S . If we take, for instance, partition (i) then we will have to eventually partition the set $\{k_1, k_2\}$ into $\{k_1\}, \{k_2\}$, but if we take partition (iv) then we do not have to perform any more partitions.

2.1 Notation

In this section, we define the notation used in the rest of the paper. In Section 2.2, we provide an example to clarify the concepts introduced in this section.

We define $D = (V, A)$ as a directed graph such that, for every edge $\{i, j\}$ in E we have two arcs (i, j) and (j, i) in A , where V and E are the set of nodes and edges of the original undirected graph G , and where $c_a = c_e$ for all $a \in A$ such that both of its end nodes correspond to the end nodes of $e \in E$. Let $r \in R$ be the arbitrarily selected root node, then we define a set K of commodities, with one commodity for each of the nodes in $R \setminus \{r\}$. Let s_k and t_k be the source and sink nodes of commodity k , where $s_k = r$ for all $k \in K$, and t_k is the corresponding node in $R \setminus \{r\}$. Let S be the set of all subsets of K with at least one element, i.e., $|S| = 2^{|K|} - 1$, and let P be the set of all proper partitions of the sets of S that have cardinality of at least 2.

The underlying r -arborescence of any Steiner tree uses only a subset of the sets in S (the set of all subsets of K), which contains the set of all commodities, and all the singletons. This collection of sets of S forms a laminar family. A family \mathcal{C} of sets is called *laminar* if for every $A, B \in \mathcal{C}$ we have $A \subseteq B$ or $B \subseteq A$ or $A \cap B = \emptyset$. Let \mathcal{L}_b be the set of all possible laminar families when $|K| = b$. For laminar family $l \in \mathcal{L}_b$, let $S(l)$ be the set of subsets of K that appear in laminar family l , and let $P(l)$ be the set of partitions used to split the subsets in $S(l)$.

For any $s' \in S(l)$ such that $|s'| \leq |K| - 1$, we say that \hat{s} is its *parent* set, if s' is one of the sets obtained when we partition set \hat{s} according to laminar family l . Note that there is only one parent set for every $s' \in S(l)$, $|s'| \leq |K| - 1$. Similarly, for any set $\hat{s} \in S(l)$ such that $|\hat{s}| \geq 2$, we say that s' is a *child* set of \hat{s} , if we obtain s' when we partition \hat{s} . Note that for every set $\hat{s} \in S(l)$, $|\hat{s}| \geq 2$ we have at least 2 child sets. Let $S_l(p)$ be the set of all child nodes of partition p in laminar family l , and let $P_l(s)$ be the partition p that splits s in laminar family l , which is defined only for sets s , such that $|s| \geq 2$. Finally, we say that a Steiner tree T in G follows an (r, l) *structure*, if the arborescence rooted in r contained in T belongs to laminar family $l \in \mathcal{L}_b$, where $b = |R \setminus \{r\}|$.

2.2 Example

Suppose we have an instance with 4 terminal nodes. Let r be the root node, and let $K = \{k_1, k_2, k_3\}$. Then,


$$S = \left\{ \underbrace{\{k_1\}}_{s_1}; \underbrace{\{k_2\}}_{s_2}; \underbrace{\{k_3\}}_{s_3}; \underbrace{\{k_1, k_2\}}_{s_4}; \underbrace{\{k_1, k_3\}}_{s_5}; \underbrace{\{k_2, k_3\}}_{s_6}; \underbrace{\{k_1, k_2, k_3\}}_{s_7} \right\}$$

and


$$P = \left\{ \underbrace{\{(k_1, k_2), (k_3)\}}_{p_1}; \underbrace{\{(k_1, k_3), (k_2)\}}_{p_2}; \underbrace{\{(k_2, k_3), (k_1)\}}_{p_3}; \underbrace{\{(k_1), (k_2), (k_3)\}}_{p_4}; \right. \\ \left. \underbrace{\{(k_1), (k_2)\}}_{p_5}; \underbrace{\{(k_1), (k_3)\}}_{p_6}; \underbrace{\{(k_2), (k_3)\}}_{p_7} \right\} \\ = \left\{ \underbrace{\{s_4, s_3\}}_{p_1}; \underbrace{\{s_5, s_2\}}_{p_2}; \underbrace{\{s_6, s_1\}}_{p_3}; \underbrace{\{s_1, s_2, s_3\}}_{p_4}; \underbrace{\{s_1, s_2\}}_{p_5}; \underbrace{\{s_1, s_3\}}_{p_6}; \underbrace{\{s_2, s_3\}}_{p_7} \right\}$$

In this case p_1, p_2, p_3 and p_4 are the partitions of the set whose cardinality is 3, which is the set $\{k_1, k_2, k_3\}$. The partitions p_5, p_6 and p_7 are the partitions of sets of cardinality 2, which are 3 different sets $\{k_1, k_2\}, \{k_1, k_3\}$ and $\{k_2, k_3\}$.


Figure (1) shows the tree representation of all the possible laminar families for the case $K = \{k_1, k_2, k_3\}$.




(a) Tree representation of laminar family $l_1 \in \mathcal{L}_3$.



(b) Tree representation of laminar family $l_2 \in \mathcal{L}_3$.



(c) Tree representation of laminar family $l_3 \in \mathcal{L}_3$.



(d) Tree representation of laminar family $l_4 \in \mathcal{L}_3$.

Figure 1: Tree representation of all possible laminar families for the case $K = \{k_1, k_2, k_3\}$

In this case, we have,

$$\begin{aligned}
P(l_1) = \{p_1, p_5\} & \quad S(l_1) = \{\{k_1, k_2, k_3\}; \{k_1, k_2\}; \{k_1\}; \{k_2\}; \{k_3\}\} & = \{s_7, s_4, s_1, s_2, s_3\} \\
P(l_2) = \{p_2, p_6\} & \quad S(l_2) = \{\{k_1, k_2, k_3\}; \{k_1, k_3\}; \{k_1\}; \{k_2\}; \{k_3\}\} & = \{s_7, s_5, s_1, s_2, s_3\} \\
P(l_3) = \{p_3, p_7\} & \quad S(l_3) = \{\{k_1, k_2, k_3\}; \{k_2, k_3\}; \{k_1\}; \{k_2\}; \{k_3\}\} & = \{s_7, s_6, s_1, s_2, s_3\} \\
P(l_4) = \{p_4\} & \quad S(l_4) = \{\{k_1, k_2, k_3\}; \{k_1\}; \{k_2\}; \{k_3\}\} & = \{s_7, s_1, s_2, s_3\}
\end{aligned}$$

Consider the laminar family $l_1 \in \mathcal{L}_3$ shown in Figure (1a). We have that $\{k_1, k_2, k_3\}$ is the *parent* set of $\{k_1, k_2\}$ and $\{k_3\}$, and that $\{k_1, k_2\}$ is the *parent* set of $\{k_1\}$ and $\{k_2\}$. On the other hand, sets $\{k_1\}$ and $\{k_2\}$ are the *child* sets of $\{k_1, k_2\}$, and, set $\{k_1, k_2\}$ and $\{k_3\}$ are *child* sets of $\{k_1, k_2, k_3\}$. Finally, we have

$$\begin{aligned}
S_{l_1}(p_1) &= \{\{k_1, k_2\}; \{k_3\}\} = \{s_4, s_3\} \\
S_{l_1}(p_5) &= \{\{k_1\}; \{k_2\}\} = \{s_1, s_2\} \\
P_{l_1}(\{k_1, k_2, k_3\}) &= P_{l_1}(s_7) = \{p_1\} \\
P_{l_1}(\{k_1, k_2\}) &= P_{l_1}(s_4) = \{p_5\}
\end{aligned}$$

3 IP Model

We propose an IP model \mathcal{Z}_l that finds an optimal solution for a given laminar family $l \in \mathcal{L}_b$, where $b = |K|$. Since we are considering a fixed laminar family l , we are given $S(l)$, the set of subsets of K that we use, and the partitions that are performed, as well as the *parent* and *child* sets of each set in $S(l)$. Therefore, the main purpose of our model is to choose the node where each partition is performed. Once we know where each partition is performed, it is easy to determine the arcs that minimize cost, since the problem reduces to finding a shortest path between the nodes where each subset $s \in S(l)$ begins and ends.

For notational simplicity, we do not index the variables of \mathcal{Z}_l by l . Our proposed formulation \mathcal{Z}_l is,

$$\begin{aligned}
\mathcal{Z}_l : \quad & \min \sum_{a \in A} c_a \left(\sum_{s \in S(l)} f_a^s \right) \\
& \sum_{a \in \delta^+(i)} f_a^s - \sum_{a \in \delta^-(i)} f_a^s = \hat{y}_i^s - \bar{y}_i^s \quad \forall i \in V, s \in S(l) \quad (1) \\
& w_i^p = \bar{y}_i^s \quad \forall i \in V, s \in S(l) : |s| \geq 2, p \in P_l(s) \quad (2) \\
& w_i^p = \hat{y}_i^s \quad \forall i \in V, p \in P(l), s \in S_l(p) \quad (3) \\
& \sum_{i \in V} w_i^p = 1 \quad \forall p \in P(l) \quad (4) \\
& \hat{y}_r^K = 1 \quad (5) \\
& \hat{y}_i^K = 0 \quad \forall i \in V \setminus \{r\} \quad (6) \\
& \bar{y}_{t_k}^k = 1 \quad \forall k \in K \quad (7) \\
& \bar{y}_i^k = 0 \quad \forall k \in K, i \in V \setminus \{d_k\} \quad (8) \\
& f \in \{0, 1\}^{|A| \times |S(l)|} \quad (9) \\
& (\hat{y}, \bar{y}) \in \{0, 1\}^{|V| \times |S(l)|} \quad (10) \\
& w \in \{0, 1\}^{|V| \times |P(l)|} \quad (11)
\end{aligned}$$

The variable f_a^s is 1 if we send flow in arc a for subset s , and 0 otherwise. This means, that all of the commodities that compose subset s use arc a . The variable \hat{y}_i^s is 1 if commodities in set s *start* sharing a

path in node i , and 0 otherwise. The variable \bar{y}_i^s is 1 if commodities in set s end sharing a path in node i , and 0 otherwise. Finally, w_i^p is 1 if a partition p is performed in node i , and 0 otherwise. Constraint (1) is the flow conservation constraints for all the subsets that belong to laminar family l , which relates variables f , \hat{y} and \bar{y} . Constraint (2) states that if partition p is performed in node i , then the subset s that is partitioned has to stop sharing in node i . Constraint (3) states that if partition p occurs in node i , then the *child* sets of partition p start to share in node i . Constraint (4) imposes that the partitions can occur only in one node. Constraints (5) and (6) say that the set of all commodities start sharing in the root node. Finally, constraints (7) and (8) state that every commodity has to send its flow to its sink node.

The LP relaxation of \mathcal{Z}_l , which we refer to as $LP(\mathcal{Z}_l)$, is defined by the same objective function and set of linear equalities, but relaxing the integrality condition of the variables. This means that we replace (9), (10) and (11) with

$$f \in [0, 1]^{|A| \times |S(l)|} \quad (12)$$

$$(\hat{y}, \bar{y}) \in [0, 1]^{|V| \times |S(l)|} \quad (13)$$

$$w \in [0, 1]^{|V| \times |P(l)|} \quad (14)$$

For all $l \in \mathcal{L}_b$, let Q_l^{IP} and Q_l^{LP} be the feasible regions of \mathcal{Z}_l and $LP(\mathcal{Z}_l)$ respectively. For all $e \in E$, let $A(e) = \{a \in A : a \text{ is a directed arc of } e \in E\}$, and let $\chi \in \{0, 1\}^{|E|}$. χ represents a solution to the Steiner tree problem in G , where $\chi_e = 1$ if edge e belongs to the solution.

Definition 1. Let $l \in \mathcal{L}_b$, and let $x = (f, w, \hat{y}, \bar{y})$ be a feasible solution for Q_l^{IP} . We define ϕ to be a mapping of x to the original problem space, where

$$\phi(x) = \chi \in \{0, 1\}^{|E|}$$

given by

$$\chi_e = \min \left\{ 1, \sum_{s \in S(l)} \sum_{a \in A(e)} f_a^s \right\} \quad \text{for all } e \in E.$$

3.1 Structure of feasible solutions of Q_l^{IP}

We assume that Q_l^{IP} is feasible, and let $x = (f, w, \hat{y}, \bar{y})$ be an arbitrary feasible solution for Q_l^{IP} . A feasible solution can have two components, it must have an acyclic component, denoted by x_{nc} , and it may also have a cycle component, denoted by x_c . Then, $x = x_{nc} + x_c$. The acyclic component is always present and is a feasible solution to the problem, while the cycle component may not be present, i.e., we may have $x_c = 0$. The cycle component is not feasible on its own, since it is only composed of cycles, and therefore it only fulfills the flow constraints (1). We analyze each component separately.

First, consider the acyclic component, $x_{nc} = (f_{nc}, w_{nc}, \hat{y}_{nc}, \bar{y}_{nc})$. Since this component has to be feasible, then by constraints (2)-(8), \hat{y}_{nc}^s and \bar{y}_{nc}^s have exactly one non-zero component for all $s \in S(l)$, which we call i_s and j_s respectively. There are two cases, either $i_s \neq j_s$, or $i_s = j_s$. In the first case, because of constraint (1), there must be a path for s from i_s to j_s . In the second case, since x_{nc} has no cycles, $f_{nc}^s = 0$, or equivalently the set s has no path. On the other hand, constraints (2)-(4) ensure that the path of every parent set ends in the same node where the path of its child sets start. By constraints (5) and (6) we have that the set of all commodities must start in r , and by constraints (7) and (8), each singleton k must end its path in t_k . Therefore, in every feasible solution, there is a path from r to t_k for all k in K , which corresponds to the union of the paths of all sets containing k . Consequently, all Steiner trees with (r, l) structure have a corresponding acyclic feasible solution to \mathcal{Z}_l , but not all acyclic solutions of \mathcal{Z}_l maps to an (r, l) structured

Steiner tree. We will discuss the case in which a feasible acyclic solution is not an (r, l) structured Steiner tree in Proposition 1.

Now, consider the cycle component. Let $x_c = (f_c, w_c, \hat{y}_c, \bar{y}_c) \neq 0$, i.e, there is at least one set $\bar{s} \in S(l)$ which has at least one cycle, and for simplicity of the argument, suppose that \bar{s} has only one cycle. The only way that this can happen is when the cycle uses arcs that are not in the path between $i_{\bar{s}}$ and $j_{\bar{s}}$, the path used by \bar{s} in the non-cycle component, otherwise it would be an infeasible solution because it would violate constraints (1) and (9). In the cycle component, w_c, \hat{y}_c and \bar{y}_c are always 0, and if $x_c \neq 0$ then $f_c \neq 0$. Consequently, we have $w = w_{nc}, \hat{y} = \hat{y}_{nc}, \bar{y} = \bar{y}_{nc}$, and $f = f_{nc} + f_c$. If a solution does not contain cycles, then $f = f_{nc}$.

Proposition 1. *For all $l \in \mathcal{L}_b$, let x be a feasible solution of \mathcal{Q}_l that does not contain any cycles, then $\phi(x)$ is either an (r, l) structured Steiner tree, or there exists a Steiner tree, with possible different (r, l) structure, contained in the support of $\phi(x)$.*

Proof. As was pointed out in Section 3.1, in the non-cycle component of every feasible solution x of \mathcal{Q}_l , there is a path from r to t_k for every commodity $k \in K$. This is a property of every Steiner tree with an (r, l) structure. But, there are also other cases where a feasible solution x of \mathcal{Q}_l does not map to an (r, l) structured Steiner tree, as shown in Figure 2. We observe in Figure 2a, a feasible solution x of \mathcal{Q}_{l_1} whose mapping $\phi(x)$ is not an (r, l_1) structured Steiner tree¹. Note that commodities k_2 and k_3 share arcs, but not using set $\{k_2, k_3\}$ because this set is not in l_1 .




Figure 2: Feasible solution x of \mathcal{Q}_{l_1} , whose mapping $\phi(x)$ is not an (r, l_1) structured Steiner tree, and (r, l_3) structured Steiner tree contained in support of $\phi(x)$.

Now, suppose $x \in \mathcal{Q}_l$ is acyclic and does not map to an (r, l) structured Steiner tree. The mapping $\phi(x)$ tells us which edges of the original undirected graph G are used by x . We can construct a subgraph defined by those edges, and we know there must be a Steiner tree in it, because in solution x there is at least one path from r to t_k for all $k \in K$. Then, we can take any Steiner tree within the constructed subgraph, which will have a unique (r, l^*) structure, with l^* not necessarily equal to l . We can do this by taking any spanning tree in the constructed subgraph. For instance, Figure 2b shows an (r, l_3) structured Steiner tree within the support of $\phi(x)$ shown in Figure 2a. □ □

3.2 Structure of feasible solutions of \mathcal{Q}_l^{LP}

The structure of the feasible solutions of \mathcal{Q}_l^{LP} , is similar to the solutions of \mathcal{Q}_l^{IP} , because they also have a non-cycle, and a cycle component. For simplicity, we use the same notation as in Section 3.1. We assume that \mathcal{Q}_l^{LP} is feasible, and let $x = (f, w, \hat{y}, \bar{y})$ be an arbitrary solution of \mathcal{Q}_l^{LP} . Let $x = x_{nc} + x_c$, where x_{nc} is the non-cycle component of x , and x_c is the cycle component of x .

¹See Figure 1 for description of laminar families l_1 and l_3 .

The analysis for this case is similar to the analysis in Section 3.1, but since the variables now can be fractional, we may have solutions where the non-cycle and the cycle component use the same arc a for the same set s , but the sum of acyclic and cyclic components has to be at most 1 for all arcs and for all sets. Moreover, in the cycle component, we may have that two different cycles, for the same s , use the same arc a . Thus, the cycle part will be a collection of weighted cycles for each $s \in S(l)$, such that for every arc $a \in A$, the sum of acyclic and cyclic flow in a for s is at most 1.

Let $LP(\mathcal{Z}_l)(\lambda)$ denote the formulation $LP(\mathcal{Z}_l)$, when we replace 1 by $\lambda \in (0, 1]$ in constraints (4), (5), (7), (12), (13), and (14).

Lemma 1. *For every $l \in \mathcal{L}_b$ and $\lambda \in (0, 1]$, we have $LP(\mathcal{Z}_l)(\lambda) = \lambda [LP(\mathcal{Z}_l)(1)]$.*

Proof. For any $\lambda \in (0, 1]$, let $(f^\lambda, \hat{y}^\lambda, \bar{y}^\lambda, w^\lambda)$ denote an optimal solution for $LP(\mathcal{Z}_l)(\lambda)$. Then,

- $LP(\mathcal{Z}_l)(\lambda) \leq \lambda [LP(\mathcal{Z}_l)(1)]$. Since $(\lambda f^1, \lambda \hat{y}^1, \lambda \bar{y}^1, \lambda w^1)$ is a feasible solution to $LP(\mathcal{Z}_l)(\lambda)$, it holds that $LP(\mathcal{Z}_l)(\lambda) \leq \lambda [LP(\mathcal{Z}_l)(1)]$.
- $LP(\mathcal{Z}_l)(\lambda) \geq \lambda [LP(\mathcal{Z}_l)(1)]$. Since $(\frac{1}{\lambda} f^\lambda, \frac{1}{\lambda} \hat{y}^\lambda, \frac{1}{\lambda} \bar{y}^\lambda, \frac{1}{\lambda} w^\lambda)$ is a feasible solution to $LP(\mathcal{Z}_l)(1)$, it holds that $\frac{1}{\lambda} [LP(\mathcal{Z}_l)(\lambda)] \geq LP(\mathcal{Z}_l)(1)$

Then, $LP(\mathcal{Z}_l)(\lambda) = \lambda [LP(\mathcal{Z}_l)(1)]$ for all $\lambda \in (0, 1]$. □ □

Our main result is stated in the following theorem.

Theorem 1. *For all $l \in \mathcal{L}_b$, we have that $\text{Conv}(\mathcal{Q}_l^{IP}) = \mathcal{Q}_l^{LP}$.*

Proof. We will prove that for all cost vectors c , there exists an optimal solution to $LP(\mathcal{Z}_l)$ that is integral.

The following definitions will be used in the proof. For all $i \in V, s \in S(l)$, let $LP(\mathcal{Z}_l(i, s))$ be a subformulation of $LP(\mathcal{Z}_l)$, where the root is i (instead of r), the terminal nodes are t_k for all $k \in s$, and the splitting sequence is defined by the way s , and its subsets, are split in laminar family l . The cost vector of $LP(\mathcal{Z}_l(i, s))$ is the same as the one used in $LP(\mathcal{Z}_l)$, but only considering the components of $LP(\mathcal{Z}_l)$ that are present in $LP(\mathcal{Z}_l(i, s))$. Let $x^*(i, s)$ be an optimal solution to $LP(\mathcal{Z}_l(i, s))$.

Let $x^* = (f^*, \hat{y}^*, \bar{y}^*, w^*)$ be an optimal solution of $LP(\mathcal{Z}_l)$. Note that if w^* is integer, then \hat{y}^* and \bar{y}^* are integer because of constraints (4)-(8). This implies, that for all $s \in S(l)$, there is only one non-zero component in vectors \hat{y}^s and \bar{y}^s , which is equal to 1. Let i_s and j_s be those indexes, respectively. Now, for all $s \in S(l)$, constraints (1) and (9) correspond to the shortest i_s - j_s path polytope, which has only integer extreme points. Since each of these polytopes is independent of each other, we have that f^* is integer. Therefore, if w^* is integer, then x^* is integer.

Now, suppose that w^* is fractional. This implies, that for at least one set $s \in S(l)$, \bar{y}^s is fractional. Note that, we may have that \hat{y}^s is also fractional, but for at least one $s \in S(l)$, we claim that \hat{y}^s is integer, and \bar{y}^s is fractional. The justifications for this claim is the following. Suppose the partitions in $P(l)$ are ordered in decreasing order by the cardinality of the set they split, i.e., the first element of $P(l)$ splits K , and the last set of $P(l)$ splits a set of cardinality 2. Now, consider w^* , and let \hat{p} be the first element in $P(l)$ such that $w^{\hat{p}}$ is fractional. Let \hat{s} be the set that \hat{p} splits. Then, it is clear that $\hat{y}^{\hat{s}}$ is integer and $\bar{y}^{\hat{s}}$ is fractional.

Let \hat{s} be the largest set in $S(l)$ such that $\hat{y}^{\hat{s}}$ is integer, and $\bar{y}^{\hat{s}}$ is fractional, and for simplicity of the argument, suppose $\hat{s} = K$. Let $I(\hat{s})$ be the set of indexes, such that $\bar{y}_i^{\hat{s}} = \lambda_i > 0$ for all $i \in I(\hat{s})$, and $\sum_{i \in I(\hat{s})} \lambda_i = 1$. Let $\hat{p} \in P(l)$, be the partition that splits \hat{s} in laminar family l . By constraint (2) we have that $w_i^{\hat{p}} = \lambda_i$ for all $i \in I(\hat{s})$, as shown in Figure 3.




Figure 3: Solution illustration.

For all $i \in I(\hat{s})$ and $s \in S_l(\hat{p})^2$, by constraint (3), we have that $\hat{y}_i^s = \lambda_i$. Now, take a fixed $s \in S_l(\hat{p})$, and a fixed $i \in I(\hat{s})$. Because $\hat{y}_i^s = \lambda_i$, then we know that within solution x^* , we are sending λ_i units of flow from node i to $\{t_k\}_{k \in s}$. Since x^* is optimal, then by Lemma 1, that part of the solution has to correspond to $\lambda_i x^*(i, s)$, otherwise, we can improve solution x^* . On the other hand, within solution x^* we are sending λ_i units of flow from r to i , for all $i \in I(\hat{s})$. Let $x^*(r, i, \hat{s})$ be the lowest cost solution sending 1 unit of flow from r to $i \in I(\hat{s})$ for set \hat{s} , then, by Lemma 1, $\lambda_i x^*(r, i, \hat{s})$ corresponds to the lowest cost solution sending λ_i units of flow from r to $i \in I(\hat{s})$ for set \hat{s} , otherwise we can improve solution x^* . Figure 4 shows a representation of the last statement.




Figure 4: Solution illustration.

Putting everything together, we have

$$x^* = \sum_{i \in I(\hat{s})} \lambda_i \left[x^*(r, i, \hat{s}) + \sum_{s \in S_l(\hat{p})} x^*(i, s) \right]$$

But, for all $i \in I(\hat{s})$, $x^*(r, i, \hat{s}) + \sum_{s \in S_l(\hat{p})} x^*(i, s)$ is a feasible solution. Then, x^* is a convex combination of $\{x^*(r, i, \hat{s}) + \sum_{s \in S_l(\hat{p})} x^*(i, s)\}_{i \in I(\hat{s})}$, and therefore, it is not an extreme point. Note that, $x^*(r, i, \hat{s})$ and $\sum_{s \in S_l(\hat{p})} x^*(i, s)$ may contain cycles. Furthermore, if $\hat{s} \neq K$, we can use the same argument, but we have to consider the solution of all the sets that contain \hat{s} as a subset. By definition of \hat{s} , in the solution of all those sets, \hat{y} and \bar{y} will be integers, and therefore, the entire solution of all those sets has to be integer. Thus, we can include that part of the solution in $x^*(r, i, \hat{s}) + \sum_{s \in S_l(\hat{p})} x^*(i, s)$ for all $i \in I(\hat{s})$, and use the same argument as before. Consequently, all extreme point solutions have \hat{y} , \bar{y} , and w integral, and therefore, all extreme points of \mathcal{Q}_l^{LP} are integral. \square \square

²Recall that $S_l(p)$ is the set of all child subsets of p in laminar family l .

As a corollary, we can solve the entire Steiner tree problem by optimizing each laminar family subproblem independently, and taking the solution with the lowest cost. For simplicity of notation, we define $\mathcal{Q}_l = \mathcal{Q}_l^{LP}$, which are the extreme points of \mathcal{Q}_l^{LP} .

Corollary 1. *Let $c \geq 0$ be the cost vector of a Steiner tree problem instance. Let $x^l = \text{Argmin}\{c^T x : x \in \mathcal{Q}_l\}$, and $v^l = c^T x^l$ for all $l \in \mathcal{L}_b$. Then, $\phi(x^{l^*})$ is a minimum cost Steiner tree, where $l^* = \text{Argmin}\{v^l : l \in \mathcal{L}_b\}$.*

Proof. Note that, since the cost is non-negative, then the optimal solution for each subproblem is acyclic. By Proposition 1, we know that $\phi(x^{l^*})$ is either an (r, l^*) structured Steiner tree, or $\phi(x^{l^*})$ contains a Steiner tree, with a different structure. But, since the cost vector is non-negative, then $\phi(x^{l^*})$ has to be a Steiner tree, otherwise, we can take any Steiner tree contained in the support of $\phi(x^{l^*})$, which will have a lower cost since it uses a subset of the edges used by $\phi(x^{l^*})$. Therefore, $\phi(x^{l^*})$ is a minimum cost Steiner tree. \square \square

Note that using a standard disjunctive programming argument, we can write an extended formulation for the entire problem based on the formulations for each laminar family [1, 9].

3.3 Problem size

The downside of our proposed formulation is the number of laminar families, which grows very fast as the number of terminals increase. In this section we provide an upper bound on the number of laminar families that we have to consider, and we show that this upper bound is tight.

In the following proposition, we use the notion of a *full binary tree*. Note that, a tree T is called a *full binary tree* if every node in T has either zero or two children.

Proposition 2. *It suffices only to consider laminar families in \mathcal{L}_b whose tree representation corresponds to a full binary tree.*

Proof. Let l_1 be a laminar family in \mathcal{L}_b and let $T(l_1)$ be the tree representation of l_1 , such that $T(l_1)$ is not a full binary tree. Therefore, there must be a node in $T(l_1)$ that has at least 3 child nodes, which implies that there must be a set in l_1 that is partitioned into at least 3 sets, let \hat{s} be such set. Suppose that the *child* sets of \hat{s} in l_1 are s_1, s_2, \dots, s_t with $t \geq 3$. For simplicity, assume $t = 3$. Now, let l_2 be the laminar family in \mathcal{L}_b such that $S(l_2) = S(l_1) \cup s'$ where $s' = s_1 \cup s_2$. This means that all partitions in l_1 and l_2 are the same, but in l_1 , \hat{s} is partitioned into s_1, s_2, s_3 , and in l_2 , \hat{s} is partitioned into s' and s_3 , and s' is partitioned into s_1 and s_2 . Now, let $x_1 = (f_1, w_1, \hat{y}_1, \bar{y}_1)$ be an optimal solution to \mathcal{Z}_{l_1} . We claim that we can construct a feasible solution to \mathcal{Z}_{l_2} using x_1 . Let i^* be the node where \hat{s} is split in x_1 , then we define $x_2 = (f_2, w_2, \hat{y}_2, \bar{y}_2) \in \mathcal{Q}_{l_2}$ as follows,

$$\begin{aligned} (f_2)_a^s &= \begin{cases} (f_1)_a^s & \forall a \in A, s \in S(l_1) \\ 0 & \forall a \in A, s = s' \end{cases} \\ (\hat{y}_2)_i^s &= \begin{cases} (\hat{y}_1)_i^s & \forall i \in V, s \in S(l_1) \\ 1 & i = i^*, s = s' \\ 0 & \forall i \in V \setminus \{i^*\}, s = s' \end{cases} \\ (\bar{y}_2)_i^s &= \begin{cases} (\bar{y}_1)_i^s & \forall i \in V, s \in S(l_1) \\ 1 & i = i^*, s = s' \\ 0 & \forall i \in V \setminus \{i^*\}, s = s' \end{cases} \end{aligned}$$

Note that the values of the w_2 vector will be determined by the values of \hat{y}_2 and \bar{y}_2 . Moreover, x_2 is a feasible solution to \mathcal{Z}_{l_2} , whose cost is the same as the cost of x_1 . Hence, $\mathcal{Z}_{l_2} \leq \mathcal{Z}_{l_1}$. The same argument holds for $t > 3$ as well. Consequently, we only need laminar families where each non-singleton set is partitioned into two proper subsets, which correspond to laminar families whose tree representation is a full binary tree. \square \square

Using Proposition 2, we can reduce the number of laminar families by considering only laminar families whose tree representations are full binary trees. For simplicity of notation, we denote by \mathcal{L}_b the reduced set of laminar families, when we have $b + 1$ terminals. We are interested in expressing $|\mathcal{L}_b|$ as a function of b . Note that, $|\mathcal{L}_b|$ is equivalent to the number of full binary labeled trees with b leaves. It is known that (see [4], and [23] Chapter 5.2.6).

$$|\mathcal{L}_b| = (2b - 3)!! = \frac{(2(b - 1))!}{2^{b-1}(b - 1)!}$$

Note, however, that $|\mathcal{L}_b|$ is not always achieved, since there are graphs that do not yield some laminar families of \mathcal{L}_b , which is why $\frac{(2(b-1))!}{2^{b-1}(b-1)!}$ is an upper bound on the number of laminar families to consider. Now, recall Stirling's formula, which gives lower and upper bounds for $n!$. For any $n \in \mathbb{Z}^+$ we have

$$\sqrt{2\pi}n^{n+\frac{1}{2}}e^{-n} \leq n! \leq en^{n+\frac{1}{2}}e^{-n}$$

Then,

$$\begin{aligned} |\mathcal{L}_b| &= \frac{(2(b - 1))!}{2^{b-1}(b - 1)!} \\ &\leq \frac{e(2(b - 1))^{2(b-1)+\frac{1}{2}}e^{-2(b-1)}}{2^{b-1}\sqrt{2\pi}(b - 1)^{(b-1)+\frac{1}{2}}e^{-(b-1)}} \\ &= \frac{e2^{(b-1)}(b - 1)^{(b-1)}e^{-(b-1)}}{\sqrt{\pi}} \\ &= \frac{e}{\sqrt{\pi}} \left(\frac{2(b - 1)}{e} \right)^{(b-1)} \end{aligned}$$

We conclude that $|\mathcal{L}_b|$ grows as $\mathcal{O}\left(\left(\frac{2b}{e}\right)^b\right)$. Although this is a big number, it is small compared to the total number of laminar families. In fact, it can be shown that as b increases, the ratio between the number of laminar families corresponding to full binary trees and the total number of laminar families, goes to 0.

Now, we analyze the size of each subproblem $LP(\mathcal{Z}_l)$ for $l \in \mathcal{L}_b$. By Proposition 2, we use laminar families whose tree representation $T(l)$ is a full binary tree. An important property of full binary trees is that, the number of internal nodes is $L - 1$, where L is the number of leaves. In our case, there are b leaves, which implies $b - 1$ internal nodes, and $2b - 1$ nodes in the entire tree. This implies that $|S(l)| = 2b - 1$ for all $l \in \mathcal{L}_b$. Moreover, the number of non-singleton sets corresponds to the number of internal nodes in $T(l)$ which is $b - 1$, and also is the number of partitions in l . Finally, since we only consider l such that $T(l)$ is a full binary tree, the number of *child* sets of each partition is exactly 2. Using these facts, the number of variables and constraints are the same for all $l \in \mathcal{L}_b$, and the number of variables is $\mathcal{O}(nb + mb)$ and the number of constraints is $\mathcal{O}(nb + b)$, where $n = |V|$, $m = |A|$ and $b + 1 = |R|$. Therefore, $LP(\mathcal{Z}_l)$ is of polynomial size in the input data. Consequently, we have that each subproblem is polynomially solvable, but the number of subproblems grow as $\mathcal{O}\left(\left(\frac{2b}{e}\right)^b\right)$, which implies that the problem is fixed-parameter tractable with respect to the number of terminal nodes. This last remark is consistent with previous results [10, 13].

4 Computational Results

We provide computational results to evaluate the performance of our proposed formulation. We use instances from the SteinLib library [18]. Table 1 shows the number of laminar families as a function of the number of terminal nodes, and it shows how long it will take to solve the entire problem if each subproblem takes one second to execute, and they are solved sequentially.

Number of terminals	$ \mathcal{L}_b $	Running time
3	1	1 seconds
4	3	3 seconds
5	15	15 seconds
6	105	1.75 minutes
7	945	15.75 minutes
8	10,395	2.89 hours
9	135,135	37.54 hours
10	2,027,025	23.46 days
11	34,459,425	1.09 years

Table 1: Number of terminals, maximum number of laminar families, and total execution time if each subproblem takes 1 second to solve.

As we can see from Table 1, the expected time to run instances with 9 or more terminals will be very large, even if the execution time for each subproblem is 1 second. Thus, we have only run experiments on instances with at most 8 terminal nodes.

Our code is implemented in *Python* using *Gurobi* 7.5.1 as a solver. All tests were performed on a laptop with an *Intel Core i7* (3.3 GHz) processor, and 16 GB of memory, using the *macOS Sierra* operating system. As we previously discussed, we can solve each laminar family subproblem independently, and therefore, they can be solved in parallel, but we implemented our code such that we run each subproblem sequentially.

Table 2 summarizes the results. The first two columns correspond to the test set and instance name, the third column shows the number of nodes, the fourth column shows the number of arcs, the fifth and the sixth columns present the number of terminal nodes and total number of laminar families. Finally, the last two columns present the average execution time for each subproblem, and the total execution time for all the models, respectively. All the reported times are in seconds.

Test set	Instance	$ V $	$ A $	$ R $	$ \mathcal{L}_b $	Subproblem time	Entire problem time
LIN	<i>lin01</i>	53	160	4	3	0.003 s	0.01 s
LIN	<i>lin02</i>	55	164	6	105	0.005 s	0.48 s
LIN	<i>lin04</i>	157	532	6	105	0.021 s	2.18 s
LIN	<i>lin07</i>	307	1,052	6	105	0.060 s	6.26 s
I160	<i>i160-033</i>	160	640	7	945	0.016 s	15.43 s
I160	<i>i160-043</i>	160	5,088	7	945	0.151 s	142.63 s
I160	<i>i160-045</i>	160	5,088	7	945	0.164 s	154.74 s
LIN	<i>lin03</i>	57	168	8	10,395	0.007 s	71.26 s
PUC	<i>cc3-4p</i>	64	576	8	10,395	0.018 s	190.89 s
PUC	<i>cc3-4u</i>	64	576	8	10,395	0.019 s	192.65 s
I320	<i>i320-011</i>	320	3,690	8	10,395	0.197 s	1,806.00 s
I320	<i>i320-043</i>	320	20,416	8	10,395	0.910 s	9,468.00 s

Table 2: Instance details and execution times in seconds.

5 Conclusions and Future Work

We propose an LP based approach to the Steiner tree problem. The proposed approach consists of solving a set of independent IPs. The best solution among the set of IPs corresponds to an optimal Steiner tree.

Each IP is polynomial in the size of the underlying graph, and we prove that the LP relaxation of each IP is integral, so each IP can be solved as a linear program. The main issue is that the number of IPs to solve grows as $\mathcal{O}\left(\left(\frac{2b}{e}\right)^b\right)$ where $b + 1$ is the number of terminal nodes. Consequently, we are able to solve the Steiner tree problem by solving a polynomial number of LPs, when the number of terminals is fixed. This is consistent with previous results [10, 13].

Future research might be directed in trying to develop tools to deal with the large number of laminar families that need to be considered. One approach could be to use Column Generation to solve this problem. We can start with a small set of the laminar families, and then try to add new laminar families to the problem in a smart way. The main difficulty with this approach, is to develop an efficient pricing problem to determine which laminar families must be added to the problem.

Another line of research is to develop new approximation algorithms using our formulation. Since solving our formulation for a small set of laminar families is fast, then one may try to find laminar families which are good candidates to be optimal. For instance, one can look at the linear relaxation of the *Bidirected Cut* formulation, which is easy to solve, and then identify all the laminar families used by the fractional solution obtained. We can use those laminar families as candidates for an optimal solution.

In addition, we can use our formulation to improve solutions delivered by any approximation algorithm. This can be done, because every solution delivered by an approximation algorithm is a Steiner tree. We can identify the (r, l) structure of such a tree, and then find a minimum cost Steiner tree with (r, l) structure by solving $LP(\mathcal{Z}_l)$, which will take polynomial time since we are only solving one subproblem.

Another research path is to develop an algorithm to solve each subproblem efficiently. The structure of the optimal solution for each problem is well characterized. For each subproblem, the problem reduces to finding the splitting nodes, and then taking the union of shortest paths. Furthermore, we can use this algorithm to construct a local search algorithm. Suppose we have a solution for a given laminar family $l \in \mathcal{L}_b$. There are laminar families in \mathcal{L}_b that have a similar topology to that of l . Then, the idea is to develop an algorithm that searches for an improving laminar family in the neighborhood of l .

Finally, a different line of research is to study our proposed approach in special graphs. There are important results in the literature for special graphs, in particular there are results that provide an upper bound on the integrality gap for well-known formulations. In these type of graphs, we may find a bound on how bad the optimal solution of the subproblem is, for the “worst” structure compared with the optimal solution for the entire problem. Moreover, there are results that show that for special cases of planar graphs, there is an algorithm to solve the problem in polynomial time [2]. These results, may imply that for those type of graphs, the number of tree structures to consider is limited.

Acknowledgement

This research was partially supported by Office on Naval Research grants N00014-15-1-2078 and N00014-18-1-2075 to the Georgia Institute of Technology, and by the CONICYT (Chilean National Commission for Scientific and Technological Research) through the Doctoral Fellowship program “Becas Chile”, Grant No. 72160393

References

- [1] E. Balas. Disjunctive programming: Properties of the convex hull of feasible points. *Discrete Applied Mathematics*, 89(1-3):3–44, 1998.
- [2] M. Bern and D. Bienstock. Polynomially solvable special cases of the steiner problem in planar networks. *Annals of Operations Research*, 33(6):403–418, 1991.
- [3] M. Bern and P. Plassmann. The steiner problem with edge lengths 1 and 2. *Information Processing Letters*, 32(4):171–176, 1989.
- [4] L. J. Billera, S. P. Holmes, and K. Vogtmann. Geometry of the space of phylogenetic trees. *Advances in Applied Mathematics*, 27(4):733–767, 2001.
- [5] A. Borchers and D.-Z. Du. The k-steiner ratio in graphs. *SIAM Journal on Computing*, 26(3):857–869, 1997.
- [6] J. Byrka, F. Grandoni, T. Rothvoss, and L. Sanità. Steiner tree approximation via iterative randomized rounding. *Journal of the ACM (JACM)*, 60(1):6, 2013.
- [7] D. Chakrabarty, J. Könemann, and D. Pritchard. Hypergraphic lp relaxations for steiner trees. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 383–396. Springer, 2010.
- [8] M. Chlebík and J. Chlebíková. The steiner tree problem on graphs: Inapproximability results. *Theoretical Computer Science*, 406(3):207–214, 2008.
- [9] M. Conforti, G. Cornuéjols, and G. Zambelli. *Integer Programming*. Springer, 2014.
- [10] S. E. Dreyfus and R. A. Wagner. The steiner problem in graphs. *Networks*, 1(3):195–207, 1971.
- [11] J. Edmonds. Optimum branchings. *Journal of Research of the National Bureau of Standards*, 71:233–240, 1967.
- [12] A. E. Feldmann, J. Könemann, N. Olver, and L. Sanità. On the equivalence of the bidirected and hypergraphic relaxations for steiner tree. *Mathematical programming*, 160(1-2):379–406, 2016.
- [13] A. E. Feldmann and D. Marx. The complexity landscape of fixed-parameter directed steiner network problems. *arXiv preprint arXiv:1707.06808*, 2017.
- [14] M. X. Goemans. The steiner tree polytope and related polyhedra. *Mathematical programming*, 63(1-3):157–182, 1994.
- [15] M. X. Goemans and Y.-S. Myung. A catalog of steiner tree formulations. *Networks*, 23(1):19–28, 1993.
- [16] M. X. Goemans, N. Olver, T. Rothvoß, and R. Zenklusen. Matroids and integrality gaps for hypergraphic steiner tree relaxations. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 1161–1176. ACM, 2012.
- [17] R. M. Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.
- [18] T. Koch, A. Martin, and S. Voß. Steinlib: An updated library on steiner tree problems in graphs. In *Steiner trees in industry*, pages 285–325. Springer, 2001.
- [19] J. Könemann, D. Pritchard, and K. Tan. A partition-based relaxation for steiner trees. *Mathematical Programming*, 127(2):345–370, 2011.
- [20] T. Polzin and S. V. Daneshmand. A comparison of steiner tree relaxations. *Discrete Applied Mathematics*, 112(1-3):241–261, 2001.

- [21] T. Polzin and S. V. Daneshmand. On steiner trees and minimum spanning trees in hypergraphs. *Operations Research Letters*, 31(1):12–20, 2003.
- [22] S. Rajagopalan and V. V. Vazirani. On the bidirected cut relaxation for the metric steiner tree problem. In *SODA*, volume 99, pages 742–751, 1999.
- [23] R. P. Stanley. *Enumerative Combinatorics*, volume 2. Cambridge University Press, 1999.
- [24] D. M. Warme and J. S. Salowe. *Spanning trees in hypergraphs with applications to Steiner trees*. University of Virginia, 1998.
- [25] R. T. Wong. A dual ascent approach for steiner tree problems on a directed graph. *Mathematical programming*, 28(3):271–287, 1984.