On Approximate Min-Max Theorems for Graph Connectivity Problems

by

Lap Chi Lau

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Computer Science
University of Toronto

Abstract

On Approximate Min-Max Theorems for

Graph Connectivity Problems

Lap Chi Lau

Doctor of Philosophy

Graduate Department of Computer Science

University of Toronto

2006

Given an undirected graph G and a subset of vertices $S \subseteq V(G)$, we call the vertices in S the terminal vertices and the vertices in $V(G) - S$ the Steiner vertices. In this thesis, we study two problems whose goals are to achieve high "connectivity" among the terminal vertices.

The first problem is the Steiner Tree Packing problem, where a Steiner tree is a tree that connects the terminal vertices (Steiner vertices are optional). The goal of this problem is to find a largest collection of edge-disjoint Steiner trees.

The second problem is the Steiner Rooted-Orientation problem. In this problem, there is a root vertex r among the terminal vertices. The goal is to find an orientation of all the edges in G so that the Steiner rooted-connectivity is maximized in the resulting directed graph D. Here, the Steiner rooted-connectivity is defined to be the maximum k so that the root vertex has k arc-disjoint paths to each terminal vertex in D.

Both problems are generalizations of two classical graph theoretical problems: the edge-disjoint s, t-paths problem and the edge-disjoint spanning trees problem. Polynomial time algorithms and exact min-max relations are known for the classical problems. However, both problems that we study are NP-complete, and thus exact min-max relations are not expected. In the following, we say S is l-edge-connected in G if we need to remove at least l edges in order to disconnect two vertices in S. Clearly, the maximum

l for which S is l-edge-connected in G is an upper bound on the optimal value for both problems that we study (i.e. the number of edge-disjoint Steiner trees, and the Steiner rooted-connectivity in an orientation).

The main result of the STEINER TREE PACKING problem is the following approximate min-max relation:

- If S is $24k$-edge-connected in G, then there are k edge-disjoint Steiner trees.

This answers Kriesell's conjecture affirmatively up to a constant multiple. We also generalize the above result to the STEINER FOREST PACKING problem. These results will appear in Chapter 3.

The main result of the STEINER ROOTED-ORIENTATION problem is the following approximate min-max relation:

- If S is $2k$-hyperedge-connected in a hypergraph H, then there is a Steiner rooted k-hyperarc-connected orientation of H.

Here, an orientation of a hyperedge e is to designate one vertex in e as the tail vertex and other vertices as the head vertices. The above result is best possible in terms of the connectivity bound. We have also considered the element-connectivity version of this problem, and proved a similar result. These results will appear in Chapter 4.

The proofs of the approximate min-max relations are constructive, and they imply the first polynomial time constant factor approximation algorithms for both problems. The proofs are based on a new technique of graph decomposition to reduce the problems into simpler instances (e.g. bipartite graphs). Then, powerful tools from combinatorial optimization (e.g. submodular flows, matroid union, edge splitting-off) can be applied to solve these NP-complete problems approximately in these simpler instances.

We shall start this thesis by describing the relations of the problems that we study to the network multicasting problem, which is the starting point of this work.

Acknowledgements

First and foremost, I would like to extend my deepest gratitude to my advisor Professor Michael Molloy, for his support, encouragement and trust throughout the years. His insightful comments have improved my papers, talks, and this thesis tremendously, and his advice has always helped me to make the right decision. I am also greatly indebted to Professor Derek Corneil, for nourishing me to be a better researcher, and always being there for help. I am also very grateful to Professor Avner Magen, for his friendship and countless hours of discussions, and to Professor Allan Borodin and Professor Joseph Cheriyan for invaluable suggestions. It was a real privilege for me to study in the Theory Group of the Department of Computer Science at the University of Toronto.

During my graduate studies, I had the good fortune to receive a Microsoft Fellowship and I enjoyed three wonderful summers working in the Theory Group of Microsoft Research in Redmond. I would like to offer my most sincere thanks to my mentors Dr. Kamal Jain and Professor László Lovász, for introducing me to fascinating research areas and sharing their insights and expertise, which have greatly broadened my knowledge and horizon. I am also very grateful to Christian Borgs, Jennifer Chayes and Microsoft Research for their hospitality.

Another group which has influenced me immensely is the Egreváry Research Group in Combinatorial Optimization (EGRES) in Budapest, in which I spent the winter of 2004. This is a truly exciting place for researchers in combinatorial optimization. During my visit, I learned some powerful tools and met many excellent young researchers. The group's passion of pursuing deep understanding in the subject has inspired me immensely. I am indebted to Professor András Frank for sharing his experiences and unique understanding of the area, and for his hospitality and advice. Special thanks go to my coauthor, Tamás Király, for the very enjoyable collaboration, and for his help on many administrative issues which made my stay in Budapest possible.

Finally, thanks to my dear parents for always believing in me, to my beloved wife Pui Ming and our wonderful daughter Ching Lam for making it all worthwhile. I dedicate this thesis to them.

To my parents, my wife Pui Ming, and my daughter Ching Lam.

Contents

Chapter 1

Overview

Sending data through a network is a task that is indispensable in our modern lives. The question of how to send data efficiently, despite the considerable amount of work, remains a central and challenging question in different areas of research from information theory to computer networking. The work of this thesis is motivated by a scenario known as the *network multicasting problem*, where a sender must transmit all its data to a set of receivers. This happens, for example, when a source tries to send a movie to a set of receivers over the Internet. Our objective is to maximize the transmission rate of the slowest receiver (or in other words, to minimize the completion time of the slowest receiver), subject to the capacity constraints in the network. We call the maximum achievable rate the *multicasting capacity*.

In this thesis we shall take a graph theoretical approach to the network multicasting problem, for which a network is modeled as a graph where a network node is represented by a vertex and a network link is represented by an edge. This chapter is intended to be a high-level overview of the thesis, and aims at presenting the motivation and the contribution of this work. Formal definitions and technical work will appear in subsequent chapters.

In the coming paragraphs we describe how previous research on the network multicas-

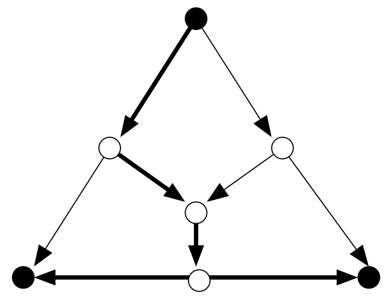

Figure 1.1: In this figure the top black vertex is the sender, and the bottom black vertices are the receivers. The bold lines form a directed Steiner tree, which can be used to transmit one unit of data from the sender to each receiver simultaneously.□

ting problem motivates the work of this thesis. First we describe the traditional setting of the network multicasting problem, where the direction of data movement along each edge is fixed. In a standard model where data can only be received, duplicated, and forwarded, a *directed Steiner tree* (also known as a *directed multicast tree* in the networking literature) is used to transmit one unit of data. See Figure 1.1 for an illustration. For the ease of visualizing the idea, we make the simplifying assumption that each edge has capacity one. That is, each edge can be used by at most one tree. Therefore, to maximize the transmission rate, one needs to find the maximum number of *edge-disjoint* directed Steiner trees. This is known as the Directed Steiner Tree Packing problem in the literature. For example, in Figure 1.2 (a), there are two edge-disjoint directed Steiner trees, which can be used to transmit two units of data from the sender to both receivers simultaneously.

A question comes up naturally: In this standard model, can we characterize which graphs have multicasting capacity at least k? This is equivalent to the question: Can we characterize which graphs have k edge-disjoint directed Steiner trees? In Figure 1.2 (b), by removing the two outgoing edges from the sender in Figure 1.2 (a), there are no more directed paths from the sender to the receivers. Since each edge can be used by at most

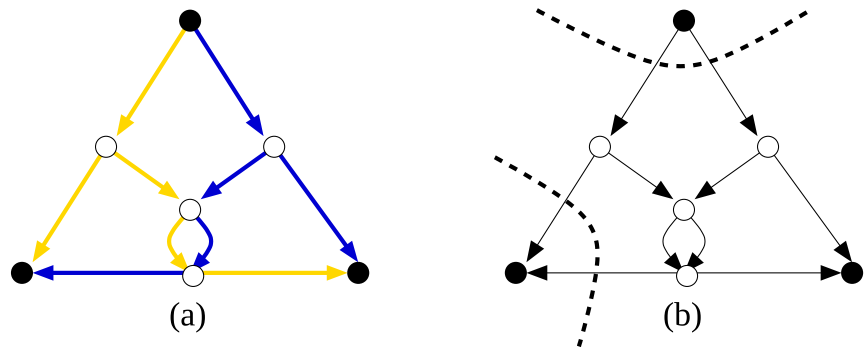

Figure 1.2: (a) The two edge-disjoint directed Steiner trees are of different colours. (b) The dotted lines indicate the edges in two of the bottlenecks (there are other bottlenecks). The two outgoing edges from the sender is a bottleneck. The two incoming edges to the left receiver is also a bottleneck.□

one tree, this implies that the multicasting capacity is at most 2, which certifies that the solution in Figure 1.2 (a) is optimal. The two outgoing edges from the sender can be thought of as a *bottleneck* of the graph. To generalize this observation, we define the bottleneck as a set of edges whose removal disconnects the sender from some receiver (i.e. after removing the bottleneck from the graph, there is no directed path from the sender to some receivers). Clearly, the capacity of a smallest bottleneck is an upper bound on the multicasting capacity. Is the multicasting capacity always equal to the capacity of a smallest bottleneck? In general, however, this is not true. For example the graph in Figure 1.1 has at most one edge-disjoint directed Steiner tree, but the size of a smallest bottleneck is 2. Furthermore, there are graphs for which this ratio is unbounded [1].

The simple but powerful idea of *network coding* overcomes the inefficiency of the standard model. In the traditional setting, data can only be duplicated and forwarded; that is, the data of an outgoing edge of a vertex must be a copy of the data of some incoming edge of the same vertex. In the network coding setting, data can also be *encoded and decoded*; that is, the data of an outgoing edge of a vertex can be an arbitrary

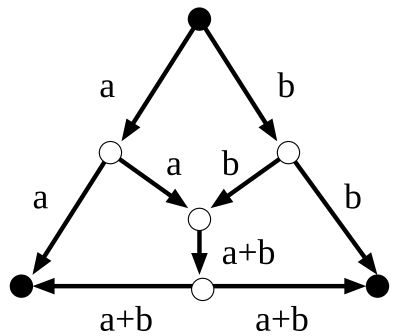

Figure 1.3: The transmission scheme is shown. In the middle vertex, the data on its outgoing edge is the addition (modulo 2) of the data on its incoming edges. This is the only vertex that does encoding in this example. In the left receiver, by adding (modulo 2) the two incoming data, the original data can be obtained; similarly for the right receiver. They are the only vertices that do decoding in this example.□

function of the data of the incoming edges of the same vertex. With network coding, the multicasting capacity of the example in Figure 1.1 is two (see Figure 1.3), which is optimal as it is equal to the capacity of a smallest bottleneck. A seminal result by Ahlswede, Cai, Li and Yeung [2] in 2000 proves that in every directed graph:

> *With network coding, the multicast capacity is equal to the capacity of a smallest bottleneck.*

This along with the aforementioned graphs from [1] shows, in particular, that the *coding advantage* can be unbounded. Here, coding advantage is defined to be the ratio of the multicasting capacity with network coding over the multicasting capacity without network coding (i.e. directed Steiner tree packing). Furthermore, the optimal transmission scheme with network coding can be computed in polynomial time. In fact, some relatively simple transmission scheme, namely linear network coding [65], suffices for the network multicasting problem. These results have generated much interest, and network coding has become a very active research area (see e.g. [2, 65, 83, 16, 63, 64, 68, 69, 70]).

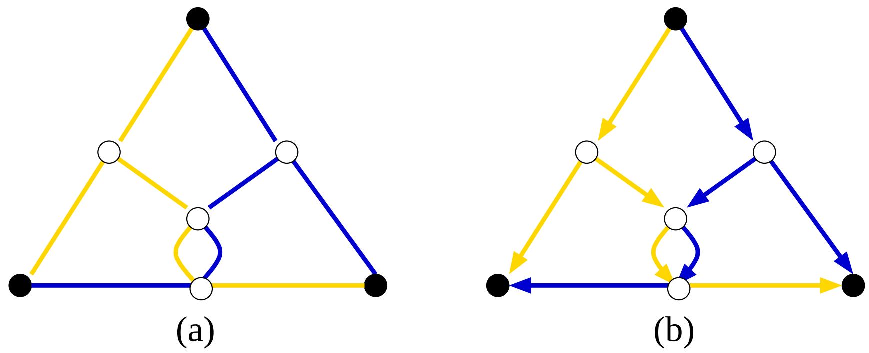

Figure 1.4: (a) The two edge-disjoint undirected Steiner trees are of different colours. (b) Having two edge-disjoint undirected Steiner trees, it is clear how to send the data from the sender to the receivers; simply move the data away from the sender in each tree.□

Our research focuses on the *undirected* version of the network multicasting problem, where data can be moved in *either direction* along an edge. There are practical networks which are undirected, for example the wireless networks. Motivated by the results on directed graphs, we are interested in the role of network coding on undirected graphs. Of particular interest is the coding advantage of the network multicasting problem in undirected graphs. Prior to our work, there were experimental results suggesting that the coding advantage in this scenario is marginal [63]. We study this problem from a theoretical point of view, and ask the following questions:

1. How can we compute the multicasting capacity if network coding is not used?

2. How can we compute the multicasting capacity if network coding is used?

3. How large can the coding advantage be?

The first problem is formulated as the UNDIRECTED STEINER TREE PACKING problem (STEINER TREE PACKING for short), where our objective is to find the maximum number of edge-disjoint undirected Steiner trees of a given graph. See Figure 1.4 for an illustration. To apply network coding, one first needs to know the directions of data

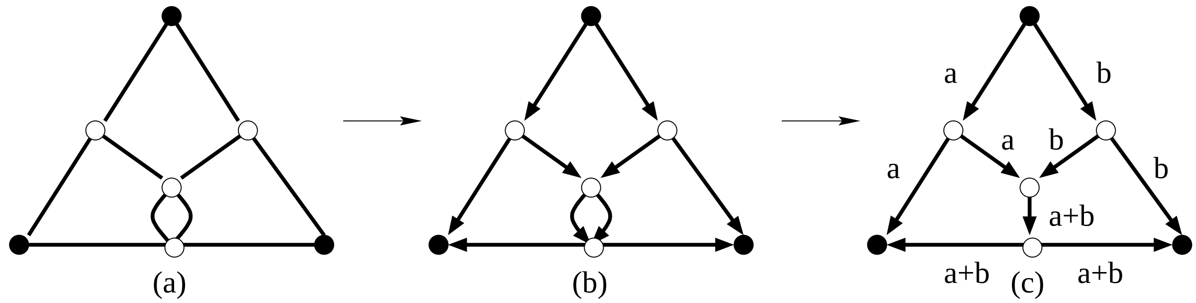

Figure 1.5: Given an undirected graph as in (a). The objective of the STEINER ROOTED ORIENTATION problem is to find an orientation of the edges as in (b), which maximizes the capacity of a smallest bottleneck in the resulting directed graph. Then, by the theorem of Ahlswede, Cai, Li, Yeung, the multicast capacity with network coding is equal to the capacity of a smallest bottleneck as shown in (c).□

movement. We formulate the second problem as the STEINER ROOTED-ORIENTATION problem, where our objective is to assign a direction to each edge of the undirected graph so as to maximize the multicasting capacity with network coding. By the above theorem of Ahlswede, Cai, Li, and Yeung, this is equivalent to the problem of assigning a direction to each edge of the undirected graphs so as to maximize the size of a smallest bottleneck in the resulting directed graph. See Figure 1.5 for an example.

Both problems are generalizations of well-studied problems in graph theory. In fact, there is an outstanding conjecture by Kriesell [57, 58] on the STEINER TREE PACKING problem that, if true, would imply that the coding advantage in the undirected network multicasting problem is at most two.

It turns out that both problems are NP-complete (the first problem was previously known to be NP-complete and we show the NP-completeness of the second problem in this thesis), which means that there might be no efficient exact algorithms for the problems. On the other hand, we are able to find efficient approximation algorithms for both problems, which give solutions close to the optimal solutions. In particular, we prove that, in both cases, the multicasting capacity is at least a constant fraction of the capacity of a smallest bottleneck. This helps us to answer the third question:

> *The coding advantage of the network multicasting problem in undirected graphs is at most a constant.*

This contrasts with the results in directed graphs, and provides a theoretical answer to the experimental observations. The main result on the Steiner Tree Packing problem also answers Kriesell's conjecture affirmatively up to a constant factor. This result and its generalization will be presented in Chapter 3.

Interestingly, the new technique developed to tackle the Steiner Tree Packing problem can also be applied to the Steiner Rooted Orientation problem, where we give generalizations of some fundamental results in graph orientations and also provide simpler proofs of some well-known theorems. The generalization has some implication to the network multicasting problem as well. The results on the Steiner Rooted Orientation problem and its generalization will be presented in Chapter 4. Formal definitions and background materials are presented in Chapter 2.

Chapter 2

The Basics

2.1 Definitions and Notation

The aim of this section is to introduce the necessary definitions and notation for this thesis. An index of the terminology and a list of notation are attached at the end of this thesis. Readers who are familiar with the topic could skip this section and only look back (via the index) if needed.

2.1.1 Undirected and Directed Graphs

All graphs are finite in this thesis.

An *undirected graph* (or a *graph*) $G = (V, E)$ consists of a set $V = V(G)$ of elements called *vertices* (or *nodes* or *points*) and a family $E = E(G)$ of unordered pairs of vertices called *edges*. We call $V(G)$ the *vertex set* of G and $E(G)$ the *edge set* of G. A vertex $v \in V(G)$ is *incident* with an edge $e \in E(G)$ if $v \in e$; then e is an edge *at* v. The two vertices incident with an edge are its *endvertices* (or *endpoints*). If $xy \in E(G)$, we say that the vertices x and y are *adjacent* (or *neighbours*). *Parallel edges* (multiple pairs with the same end-vertices) are allowed but not loops; sometimes we say G is a *multigraph* to illustrate this point. The *neighbourhood* $N_G(v)$ of a vertex v in G is the set of vertices adjacent to v. The following is an important notation:

$$\delta_G(X) := \{uv \in E(G) \mid |X \cap \{u, v\}| = 1\}.$$

In words, $\delta_G(X)$ consists of the edges with one endpoint in X and the other endpoint in $V(G)-X$. More generally, $\delta_G(X, Y)$ denotes the set of edges with one endpoint in X and the other endpoint in Y, i.e., $\delta_G(X) = \delta_G(X, V(G)-X)$. The *degree* $d_G(X)$ of a set X in G is defined to be $|\delta_G(X)|$; the *degree* $d_G(v)$ of a vertex v is just a shorthand of $d_G(\{v\})$. The *degree* $d_G(X, Y)$ between two sets is $|\delta_G(X, Y)|$. Finally, a graph is *Eulerian* if every vertex is of even degree.

A *directed graph* (or a *digraph*) $D = (V, A)$ consists of a set $V = V(D)$ of elements called vertices (or nodes or points) and a family $A = A(D)$ of ordered pairs of vertices called *arcs*. We call $V(D)$ the vertex set of D and $A(D)$ the *arc set* of D. For an arc (u, v) (or uv for short) the first vertex u is its *tail* and the second vertex v is its *head*; sometimes we may write $\overrightarrow{uv}$ for uv to emphasize the direction. We also say the arc uv *leaves* u and *enters* v. The head and the tail of an arc are its *endvertices* (or *endpoints*); we say the end-vertices are *adjacent* (or *neighbours*). Sometimes, we say D is a *directed multigraph* to emphasize that there are multiple arcs with the same end-vertices. The following notation will be used frequently:

$$\delta_D^{in}(X) := \{uv \in A(D) \mid u \in V(D) - X, v \in X\}; \qquad \delta_D^{out}(X) := \delta_D^{in}(V(D) - X).$$

In words, $\delta_D^{in}(X)$ consists of the arcs that enter X and $\delta_D^{out}(X)$ consists of the arcs that leave X. $\delta_D(X, Y)$ consists of the arcs with one endpoint in X and the other endpoint in Y. The *indegree* $d_D^{in}(X)$ of a set X in D is defined to be $|\delta_D^{in}(X)|$; similarly the *outdegree* $d_D^{out}(X)$ is defined to be $|\delta_D^{out}(X)|$. The *indegree* $d_D^{in}(v)$ and the *outdegree* $d_D^{out}(v)$ of a vertex are defined to be $|\delta_D^{in}(\{v\})|$ and $|\delta_D^{out}(\{v\})|$ respectively. The *degree* $d_D(X, Y)$ between two sets is defined to be $|\delta_D(X, Y)|$. Finally, a digraph D is an *Eulerian digraph* if $d_D^{in}(v) = d_D^{out}(v)$ for every $v \in V(D)$.

2.1.2 Hypergraphs and Directed Hypergraphs

A *hypergraph* $H = (V, \mathcal{E})$ consists of a set $V = V(H)$ of elements called vertices (or nodes or points) and a family $\mathcal{E} = \mathcal{E}(H)$ of subsets of $V(H)$ called *hyperedges*. We call $V(H)$ the vertex set of H and $\mathcal{E}$ the *hyperedge set* of H. Usually we denote a hyperedge with a lowercase letter (e.g. e, f) just like in the graph case, but sometimes we denote it with an uppercase letter (e.g. Z) to clarify that it is a subset of vertices. The *rank* of H is the cardinality of the largest hyperedge of H. Generalizing the notation of graphs, we define:

$$\delta_H(X) := \{Z \in \mathcal{E}(H) \mid 0 < |Z \cap X| < |Z|\}.$$

In other words, a hyperedge $e \in \delta_H(X)$ if e contains some vertex in X and contains some vertex in $V(H) - X$. The *degree* $d_H(X)$ of a set X in H is $|\delta_H(X)|$; the *degree* $d_H(v)$ of a vertex is $|\delta_H(\{v\})|$.

There are at least two natural ways to define *directed hypergraphs*. We first define the model used in [33] which we refer as *directed in-hypergraph* (*in-hypergraphs* for short). An in-hypergraph $\overrightarrow{H} = (V, \overrightarrow{\mathcal{E}})$ consists of a set $V = V(\overrightarrow{H})$ of elements called vertices (or nodes or points) and a family $\overrightarrow{\mathcal{E}} = \overrightarrow{\mathcal{E}}(\overrightarrow{H})$ of subsets of $V(\overrightarrow{H})$ called *in-hyperarcs*. An in-hyperarc is a subset $Z \subseteq V$ with a designated *head* vertex $v \in Z$, and it is denoted by Z^v. The vertices of $Z - v$ are called the *tail* vertices of Z^v. We call $\overrightarrow{\mathcal{E}}$ the *in-hyperarc set* of $\overrightarrow{H}$. An in-hyperarc Z^v *enters* a set X if $v \in X$ and $Z - X \neq \emptyset$; an in-hyperarc Z^v *leaves* a set X if it enters $V(\overrightarrow{H}) - X$. We denote $\delta^{in}_{\overrightarrow{H}}(X)$ and $\delta^{out}_{\overrightarrow{H}}(X)$ the set of in-hyperarcs that enter X and leave X respectively. The *indegree* $d^{in}_{\overrightarrow{H}}(X)$ of a set X is $|\delta^{in}_{\overrightarrow{H}}(X)|$ and the *outdegree* $d^{out}_{\overrightarrow{H}}(X)$ is $|\delta^{out}_{\overrightarrow{H}}(X)|$.

We refer to the second model, which is called *star hypergraphs* in [6], as *directed out-hypergraphs* (*out-hypergraphs* for short). All the definitions are defined similarly as for in-hypergraphs, except that we have *out-hyperarcs* instead of in-hyperarcs. An out-hyperarc is a subset $Z \subseteq V$ with a designated *tail* vertex $v \in Z$, and it is denoted by Z^v. The vertices of $Z - v$ are called the *head* of Z^v. An out-hyperarc Z^v *enters* a set X if

$v \notin X$ and $Z \cap X \neq \emptyset$; an out-hyperarc Z^v *leaves* a set X if it enters $V(\overrightarrow{H}) - Z$. The definitions of $\delta^{in}_{\overrightarrow{H}}(X)$, $\delta^{out}_{\overrightarrow{H}}(X)$, $d^{in}_{\overrightarrow{H}}(X)$ and $d^{out}_{\overrightarrow{H}}(X)$ are the same as for in-hypergraphs.

A *bipartite graph* $B = (X, Y; E)$ is a graph with vertex set $X \cup Y$ and edge set E, for which every edge in E has one endpoint in X and the other endpoint in Y. For $B = (X, Y; E)$, X and Y are called the *partite sets* of B. For a hypergraph $H = (V, \mathcal{E})$, the *bipartite representation* $B = (V, \mathcal{E}; E)$ of H is a bipartite graph with partite sets V and $\mathcal{E}$, and a vertex $v \in V(H)$ is adjacent to a vertex $Z \in \mathcal{E}(H)$ in B if $v \in Z$ in H. For $B = (V, \mathcal{E}; E)$, we call V the *vertex partite set* and $\mathcal{E}$ the *hyperedge partite set*. For a directed hypergraph $\overrightarrow{H} = (V, \overrightarrow{\mathcal{E}})$, the *bipartite representation* $B = (V, \overrightarrow{\mathcal{E}}; A)$ of H is a directed bipartite graph with the following arc set: for $u \in V(H)$ and $Z \in \overrightarrow{\mathcal{E}}(H)$, $uZ \in A(B)$ if and only if u is a tail of Z, and $Zu \in A(B)$ if and only if u is a head of Z. For $B = (V, \overrightarrow{\mathcal{E}}; A)$, we call V the vertex partite set and $\overrightarrow{\mathcal{E}}$ the *hyperarc partite set*. So, for example, in the bipartite representation of an in-hypergraph (out-hypergraph), every vertex in the hyperarc partite set has outdegree (indegree) exactly 1.

2.1.3 Deletions and Contractions

The following terminology is defined for hypergraphs, when it specializes to graphs we omit the word "hyper". Let $H = (V, \mathcal{E})$ and $H' = (V', \mathcal{E}')$ be two hypergraphs. If $V' \subseteq V$ and $\mathcal{E}' \subseteq \mathcal{E}$, then we say H' is a *subhypergraph* of H (or H is a *superhypergraph* of H'), written $H' \subseteq H$. Less formally, we say H *contains* H'. Given a subset $X \subseteq V(H)$, a hyperedge $Z \in E(H)$ is *induced* in X if $Z \subseteq X$. The number of hyperedges induced by X is denoted by $i_H(X)$. If $H' \subseteq H$ and H' contains all the hyperedges induced by V', then H' is an *induced subhypergraph* of H; we say that V' *induces* or *spans* H' in H, and write $H' := H[V']$. $H' \subseteq H$ is a *spanning subhypergraph* of H if V' spans all of H, i.e., $V' = V$.

If U is any set of vertices, we write $H - U$ for $H[V - U]$. In words, $H - U$ is obtained from H by *deleting* all the vertices in U and all the hyperedges that intersect U. If

$U = \{v\}$, we simply write $H - v$ instead of $H - \{v\}$. For a subset $\mathcal{F} \subseteq \mathcal{E}$, we write $H - \mathcal{F} := (V, \mathcal{E} - \mathcal{F})$; as above $H - \{e\}$ is abbreviated to $H - e$.

Let $X \subseteq V$ be a subset of vertices. By H/X we denote the hypergraph obtained from H by *contracting* X into a single vertex x, and "keeping" all the hyperedges in $\delta_H(X)$ and removing all the hyperedges induced in X. Formally, H/X is a hypergraph $(V', \mathcal{E}')$ with vertex set $V' := \{(V - X) \cup \{x\}\}$ and hyperedge set

$$E' := \{Z \mid Z \in \mathcal{E} \text{ and } Z \subseteq V - X\} \cup \{(Z - X) \cup \{x\} \mid Z \in \mathcal{E} \text{ and } X \text{ separates } Z\},$$

where X *separates* Z means $Z \cap X \neq \emptyset$ and $Z \cap (V(H) - X) \neq \emptyset$.

We need to clarify the definition of contraction in directed hypergraphs; all other definitions in this subsection apply equally well to directed hypergraphs, (i.e., one just needs to substitute "edge" by "arc"). Let $\overrightarrow{H} = (V, \overrightarrow{\mathcal{E}})$ be a out-hypergraph (for in-hypergraphs this is completely analogous) and $X \subseteq V$ be a subset of vertices. Formally, $\overrightarrow{H}/X$ is a hypergraph $(V', \overrightarrow{\mathcal{E}}')$ with vertex set $V' := \{(V - X) \cup \{x\}\}$ and hyperarc set

$$\overrightarrow{\mathcal{E}}' := \{Z \mid Z \in \overrightarrow{\mathcal{E}} \text{ and } Z \subseteq V - X\} \cup \{(Z - X) \cup \{x\} \mid Z \in \overrightarrow{\mathcal{E}} \text{ and } X \text{ separates } Z\}.$$

If a hyperarc Z with its tail in $\overrightarrow{H}$ is in X, then its tail in $\overrightarrow{H}'$ is x; all other hyperarcs have their tails in $\overrightarrow{H}'$ the same as their tails in $\overrightarrow{H}$. Intuitively, $\overrightarrow{H}/X$ denotes the out-hypergraph obtained from $\overrightarrow{H}$ by *contracting* X into a single vertex x, and "keeping" all the hyperarcs in $\delta^{in}_{\overrightarrow{H}}(X)$ and $\delta^{out}_{\overrightarrow{H}}(X)$ with the directions unchanged and removing all the hyperarcs induced in X.

2.1.4 Orientations of Graphs and Hypergraphs

The notion of *graph orientations* (or *hypergraph orientations*) provides a link between graphs and digraphs (or between hypergraphs and directed hypergraphs). The *underlying graph* of a digraph $D = (V, A)$ is obtained by replacing each arc $uv \in A$ by an edge uv, i.e. ignoring the directions. An *orientation* of a graph is a directed graph $D = (V, A)$ whose underlying graph is G; intuitively, an orientation of a graph is obtained by assigning a direction to each edge. We switch to hypergraphs in the following. The *underlying hypergraph* of a directed hypergraph $\overrightarrow{H} = (V, \overrightarrow{\mathcal{E}})$ is obtained by ignoring the distinction of head(s) and tail(s) in a hyperarc. An *orientation* of a hypergraph is a directed in-hypergraph (out-hypergraph) $\overrightarrow{H} = (V, \overrightarrow{\mathcal{E}})$ by designating the head vertex (the tail vertex) for each hyperedge.

2.1.5 Paths and Cycles

A *path of a graph* G is a sequence of distinct vertices $\{v_0, v_1, \ldots, v_k\}$ so that $v_i v_{i+1} \in E(G)$ for all $0 \leq i < k$. If $P = v_0, \ldots, v_k$ is a path, then $C := P + v_k v_0$ is a *cycle* if $v_k v_0$ is an edge of G. A *path of a digraph* D is a sequence of distinct vertices $\{v_0, v_1, \ldots, v_k\}$ so that $v_i v_{i+1} \in A(D)$ for all $0 \leq i < k$. If $P = v_0, \ldots, v_k$ is a path, then $C := P + v_k v_0$ is a *directed cycle* if $v_k v_0$ is an arc of D. A *path of a hypergraph* H is an alternating sequence of distinct vertices and hyperedges $\{v_0, e_0, v_1, e_1, \ldots, e_{k-1}, v_k\}$ so that $v_i, v_{i+1} \in e_i$ for all $0 \leq i < k$. A *path of a directed hypergraph* $\overrightarrow{H}$ is an alternating sequence of distinct vertices and hyperarcs $\{v_0, a_0, v_1, a_1, \ldots, a_{k-1}, v_k\}$ so that v_i is a tail of a_i and v_{i+1} is a head of a_i for all $0 \leq i < k$. Alternatively, a path in a (directed) hypergraph H is just a path in the bipartite representation B of H between two vertices in the vertex partite set of B. See Figure 2.1 for an illustration. In all the above sequences, v_0 and v_k are *linked* and are called the *ends* of the paths. In hypergraphs, we say it is a v_0, v_k-path to specify the ends; in directed hypergraphs, in addition, we often say it is a path from v_0 to v_k to emphasize the direction. Finally, the number of edges on a path is its *length*, in the above cases the paths are of length k.

2.1.6 Edge-Connectivity of Graphs and Hypergraphs

The following terminology is defined for hypergraphs; when we specialize to graphs we omit the word "hyper". A hypergraph H is *connected* if every pair of its vertices are linked

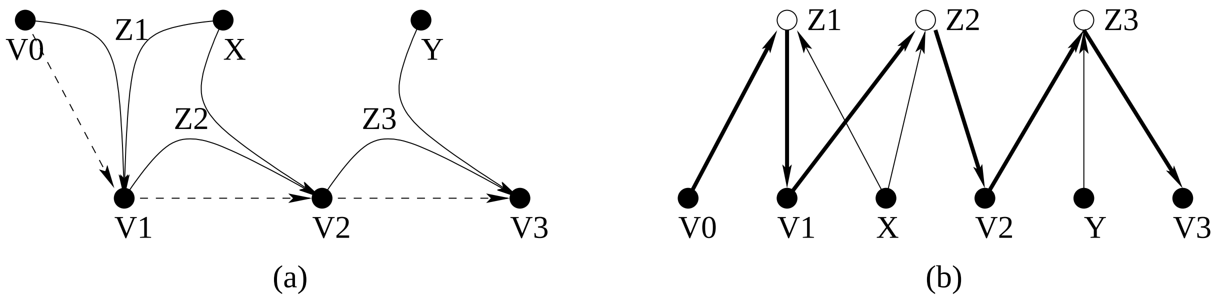

Figure 2.1: (a) A path in an in-hypergraph from v_0 to v_3. The dotted lines represent the movement of the path. (b) A path in the bipartite representation of the in-hypergraph of (a) from v_0 to v_3.

by a path. A maximal connected subgraph of H is called a *component*. We say v is a *cut vertex* if H is connected but $H-v$ is not connected. There are two natural ways to define the "hyperedge-connectivity" of a hypergraph. Intuitively, the first definition captures how "robust" a hypergraph is. A hypergraph H is *k-hyperedge-connected* if $H-F$ is connected for every set of hyperedges F with $|F|<k$. The second definition captures how many "connections" a hypergraph has. A hypergraph H is *k-hyperedge-connected* if any two of its vertices can be linked by k *hyperedge-disjoint paths*, i.e., k-paths that do not share a hyperedge. The largest integer for which H is k-hyperedge-connected is the *hyperedge-connectivity* of H. An extension of Menger's theorem to hypergraphs states that these two definitions are actually equivalent. We shall discuss Menger's theorem (and its many variants) in some depth later in Section 2.2.1 and Section 2.2.2.

As discussed in the overview, there are situations where we are only interested in a subset of vertices $S \subseteq V(H)$. We call the vertices in S the *terminal vertices* and the vertices in $V(H)-S$ the *Steiner vertices*. Given $S \subseteq V(H)$, a subset of vertices X is a *S-separating set* (or X *separates* S) if $0<|X\cap S|<|S|$. A set of hyperedges $F \subseteq E(H)$ is a *S-hyperedge-cut* (or a S-cut) if $F=\delta_H(X)$ for some S-separating set X. Notice that an S-cut of a graph is a formal definition of what we meant by a *bottleneck* in the overview. A subset $S \subseteq V(H)$ is *k-hyperedge-connected* in H if every S-cut has at least k hyperedges. For example, a hypergraph H is k-hyperedge-connected if every $V(H)$-cut has at least k hyperedges. Equivalently, as we shall see, S is k-hyperedge-connected in H if every pair of its vertices can be linked by k hyperedge-disjoint paths. The largest integer for which S is k-hyperedge-connected in H is the *hyperedge-connectivity of* S in H (or S*-hyperedge-connectivity* of H).

2.1.7 Arc-Connectivity of Digraphs and Directed Hypergraphs

In this subsection we do not distinguish between in-hypergraphs and out-hypergraphs, since the definitions apply to both cases. The following terminology is defined for directed hypergraphs; when we specialize to digraphs we omit the word "hyper".

A directed hypergraph $\overrightarrow{H}$ is *strongly connected* if there is a directed path from s to t and a directed path from t to s for any $s, t \in V(\overrightarrow{H})$. Similarly, there are two natural ways to define "hyperarc-connectivity" in a directed hypergraph. $\overrightarrow{H}$ is *strongly k-hyperarc-connected* if $\overrightarrow{H} - F$ is strongly connected for every set of hyperarcs F with $|F| < k$. Equivalently, as we shall see, a directed hypergraph $\overrightarrow{H}$ is *strongly k-hyperarc-connected* if any two of its vertices can be linked by k *hyperarc-disjoint paths* - k paths that do not share a hyperarc. The largest integer for which $\overrightarrow{H}$ is strongly k-hyperarc-connected is the *hyperarc-connectivity* of $\overrightarrow{H}$.

A set of hyperarcs $F \subseteq E(\overrightarrow{H})$ is a S*-hyperarc-cut* (or a S-cut) if $F = \delta^{in}_{\overrightarrow{H}}(X)$ (or $F = \delta^{out}_{\overrightarrow{H}}(X)$) for some S-separating set X. A subset $S \subseteq V(\overrightarrow{H})$ is *strongly k-hyperarc-connected* in $\overrightarrow{H}$ if every S-cut has at least k hyperarcs. Equivalently, as we shall see, S is strongly k-hyperarc-connected in $\overrightarrow{H}$ if any two of its vertices can be linked by k hyperarc-disjoint paths. The largest integer for which S is strongly k-hyperarc-connected in $\overrightarrow{H}$ is the *hyperarc-connectivity of* S in $\overrightarrow{H}$ (or S*-hyperarc-connectivity* of $\overrightarrow{H}$).

2.1.8 Local-Connectivity, Rooted-Connectivity and Partition-Connectivity

The following terminology is defined for hypergraphs; when we specialize to graphs we omit the word "hyper". We use the notation $\lambda_H(s,t)$ for the maximum number of hyperedge-disjoint paths between s and t in H. These values are called the *local hyperedge-connectivity* between s and t. Given a subset $S \subseteq V$ and a specified *root vertex* r, the *rooted-hyperedge-connectivity* of S in H is defined to be $\min_{v \in S} \lambda_H(r,v)$. When $S = V$, we will simply say rooted-hyperedge-connectivity of H.

Given a hypergraph $H = (V, \mathcal{E})$ and a partition $\mathcal{P} = \{P_1, \dots, P_t\}$ of V, $e_H(\mathcal{P})$ denotes the number of hyperedges which are not contained in a P_i. We also say those hyperedges are the *crossing hyperedges* of $\mathcal{P}$. A hypergraph H is called *k-partition-connected* if $e_H(\mathcal{P}) \geq k(|\mathcal{P}| - 1)$ for every partition $\mathcal{P}$ of V. Equivalently, one has to delete at least kt hyperedges to dismantle H into $t+1$ components for every t. Partition-connectivity clearly implies edge-connectivity; contrary to the graph case, however, a 1-connected hypergraph needs not be 1-partition-connected. For example, a hypergraph with a single hyperedge $e = V$ is 1-connected but not 1-partition-connected. In particular, a k-partition-connected hypergraph requires at least $k(|V(H)| - 1)$ hyperedges.

2.1.9 Trees and Arborescences

An *acyclic* graph, one not containing any cycle, is called a *forest*. A connected forest is called a *tree*. Given a graph G and a subset of vertices $S \subseteq V(G)$, a subgraph $T \subseteq G$ is called a *S-Steiner tree* (or a *S-tree*) if T is a tree and $S \subseteq V(T)$. A graph G has an S-tree if and only if S is connected in G.

Given a vertex r called the *root vertex*, a *r-arborescence* is a digraph T with a tree as its underlying graph and a path from the root to every other vertex. Given a digraph D, the root vertex r, and a subset of vertices $S \subseteq V(G)$, a sub-digraph $T \subseteq D$ is called a

(r, S)-*arborescence* if T is an r-arborescence and $S \subseteq V(T)$. If the root vertex r is clear from the context, we will not mention r in the above definitions.

2.1.10 Submodular and Supermodular Set Functions

Let V be a finite ground set. Two sets X and Y are called *co-disjoint* if $X \cup Y = V$; that is $V - X$ and $V - Y$ are disjoint. 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 not co-disjoint.

A *family* of sets $\mathcal{F}$ is a collection of (not necessarily distinct) subsets of V. $\mathcal{F}$ is a *laminar family* if it contains no intersecting members; $\mathcal{F}$ is a *cross-free family* if it contains no crossing members. For a function $m : V \to \mathbb{R}$ we use the notation $m(X) := \sum(m(x) : x \in X)$.

Let V be a finite ground set and $f : 2^V \to \mathbb{R}$ be a real valued function defined on the subsets of V. The set-function f is called *fully submodular* (or *submodular*) if the following inequality holds for any two subsets X and Y of V:

$$f(X) + f(Y) \geq f(X \cup Y) + f(X \cap Y). \tag{2.1}$$

The set function f is called *fully supermodular* (or *supermodular*) if for any two subsets X and Y of V:

$$f(X) + f(Y) \leq f(X \cup Y) + f(X \cap Y). \tag{2.2}$$

f is called *modular* if it is both submodular and supermodular; that is, $f(X) = \sum_{x \in X} f(x)$. There is an alternative way to characterize submodularity:

Proposition 2.1.1 *A set function $f : 2^V \to \mathbb{Z} \cup \{\infty\}$ is submodular if and only if*

$$f(X + v) - f(X) \geq f(Y + v) - f(Y)$$

for all $X \subseteq Y \subseteq V$ and $v \in V - Y$.

Let G be a graph and D be a digraph. The following proposition implies that the functions $d_G(.)$ and $d_D^{in}(.)$ (and hence $d_D^{out}(.)$) are submodular; it can be verified easily by checking that every edge has the same contribution to both sides.

Proposition 2.1.2 *For $X, Y \subseteq V$,*

$$\begin{aligned} d_G(X) + d_G(Y) &= d_G(X \cap Y) + d_G(X \cup Y) + 2d_G(X, Y), \\ d_G(X) + d_G(Y) &= d_G(X - Y) + d_G(Y - X) + 2d_G(X \cap Y, V - X \cup Y), \\ d_D^{in}(X) + d_D^{in}(Y) &= d_D^{in}(X \cap Y) + d_D^{in}(X \cup Y) + d_D(X, Y). \end{aligned}$$

The following is an example of a supermodular function. Recall that $i_G(X)$ denotes the number of edges of G induced in $X \subseteq V(G)$.

Proposition 2.1.3 *For $X, Y \subseteq V$,*

$$i_G(X) + i_G(Y) = i_G(X \cup Y) + i_G(X \cap Y) - d_G(X, Y).$$

Finally, we say a function f is *intersecting submodular* if (2.1) holds for any two intersecting sets; f is *crossing submodular* if (2.1) holds for any two crossing sets. Intersecting and crossing supermodular functions are defined in the same way. These functions will be very useful in graph connectivity problems.

2.2 Background

The aim of this section is to provide a comprehensive background on the subjects related to our results, and the goal is to make the results in this thesis self-contained. Part of the materials follow the presentation of [6, 15, 27, 31, 53].

2.2.1 A Proof of Menger's Theorem

Menger [73] proved that the two notions of edge-connectivity (stated in Section 2.1.6) are in fact equivalent. This result is one of the cornerstones of graph theory.

Here we give a short proof of the directed version, which can then be used to derive the other versions by standard reductions. This proof underlies the idea behind the main results of this thesis, which we shall discuss at the end of this subsection. In the following, a set X is a $\overline{s}t$ set if $s \notin X$ and $t \in X$.

Theorem 2.2.1 (Menger [73]) *Let $D = (V, A)$ be a digraph, and $s, t \in V$ be distinct vertices. There are k arc-disjoint paths from s to t if and only if*

$$d^{in}(X) \geq k \text{ for every } \overline{s}t \text{ set } X \subseteq V. \tag{2.3}$$

Proof. To set up this problem as a special case of the Steiner Tree Packing problem, we say $\{s, t\}$ are the terminal vertices and all the other vertices are Steiner vertices. Suppose, by way of contradiction, that the statement is false. Let $\mathcal{D}$ be a minimal counterexample so that it has the minimum number of arcs among all the counterexamples. We shall prove that $\mathcal{D}$ does not exist, and hence the theorem follows.

First we show that there is no arc between two Steiner vertices in $\mathcal{D}$. Suppose $a = uv$ is such an arc. If $\mathcal{D} - a$ satisfies (2.3), then $\mathcal{D} - a$ has k arc-disjoint paths from s to t by the choice of $\mathcal{D}$. This clearly contradicts the assumption that $\mathcal{D}$ is a counterexample. So we assume that $\mathcal{D} - a$ does not satisfy (2.3). Then there exists an $\overline{s}t$ set X for which uv enters X and $d^{in}(X) = k$. So, $\{s, u\} \subseteq V(G) - X$ and $\{t, v\} \subseteq X$. In particular, this implies that $|V(G) - X| \geq 2$ and $|X| \geq 2$ (see Figure 2.2 (a)). Now, we contract X of $\mathcal{D}$ into a single vertex v_1 to form D_1 and contract $V(G) - X$ of $\mathcal{D}$ into a single vertex v_2 to form D_2 (see Figure 2.2 (b)). Since s, t satisfies (2.3) in $\mathcal{D}$, it follows that s, v_1 and v_2, t satisfy (2.3) in D_1 and D_2 respectively. Notice that both D_1 and D_2 have fewer arcs than $\mathcal{D}$. So, by the choice of $\mathcal{D}$, there are k arc-disjoint paths $\{P_1^1, \ldots, P_k^1\}$ from s to v_1 in D_1 and k arc-disjoint paths $\{P_1^2, \ldots, P_k^2\}$ from v_2 to t in D_2 (see Figure 2.2 (c)). Since v_1 has indegree k, each path P_i^1 uses exactly one arc of $\delta^{in}_{D_1}(v_1)$. Similarly, since v_2 has outdegree k, each path P_j^2 uses exactly one arc of $\delta^{out}_{D_2}(v_2)$. By renaming if necessary, we can assume that P_i^1 and P_i^2 use the same arc of $\delta^{in}_{\mathcal{D}}(X)$. Therefore, by identifying

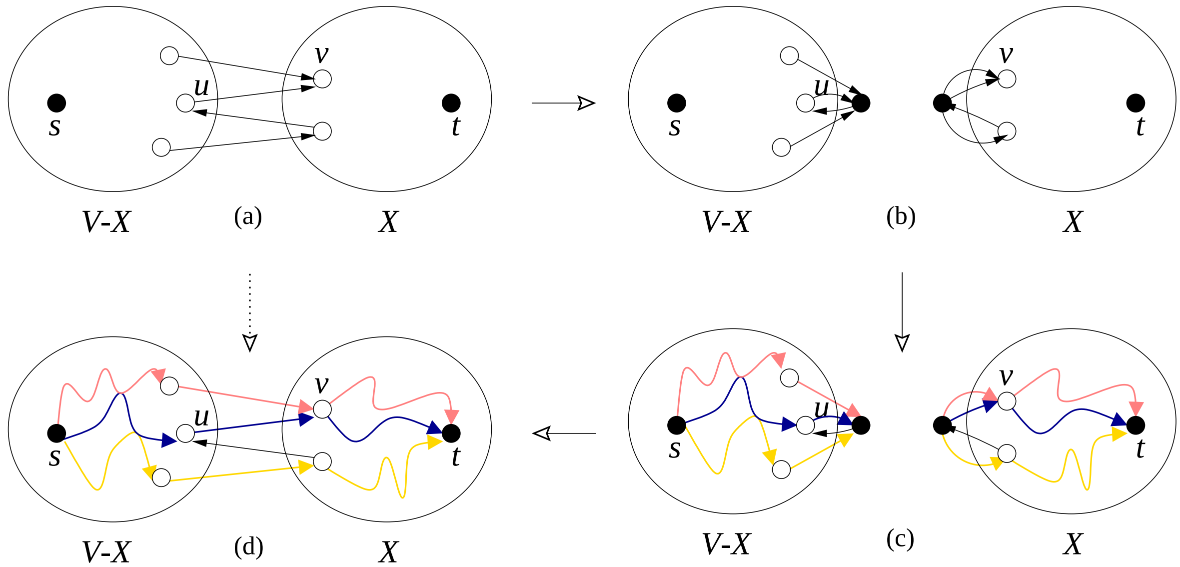

Figure 2.2: An illustration of the proof of Theorem 2.2.1.□

the corresponding arcs in $\mathcal{D}$, $\{P_1^1 \cup P_1^2, \ldots, P_k^1 \cup P_k^2\}$ are k arc-disjoint paths in $\mathcal{D}$ (see Figure 2.2 (d)). This contradicts the assumption that $\mathcal{D}$ is a counterexample. So, there is no arc between two Steiner vertices in $\mathcal{D}$.

Now the structure of $\mathcal{D}$ has become very restrictive. Notice that we can assume s has no incoming arcs and t has no outgoing arcs; otherwise we can just remove them and (2.3) would not be violated. So, each Steiner vertex v has only incoming arcs from s and outgoing arcs to t. By replacing the two arcs sv and vt of $\mathcal{D}$ by st and calling the resulting digraph D', it is easy to verify that (2.3) is satisfied on D'. By the choice of $\mathcal{D}$, there are k arc-disjoint paths between s and t in D'. Clearly, by taking the same k paths and replacing st by sv and vt if necessary, there are k arc-disjoint paths between s and t in $\mathcal{D}$. This contradicts the assumption that $\mathcal{D}$ is a counterexample. Therefore, we can further assume that there are no Steiner vertices in $\mathcal{D}$. So, $\mathcal{D}$ is only a digraph with two vertices. Clearly, to satisfy (2.3), there must be k parallel arcs from s to t. These are the k arc-disjoint paths from s to t in $\mathcal{D}$, which shows that $\mathcal{D}$ is not a counterexample. Hence $\mathcal{D}$ does not exist and this completes the proof. ∎

As remarked earlier, the above proof actually consists of almost all the main ingredients (in a simple form) of the proofs of the main results in this thesis. We identify them now to put the future proofs into context. In the first step we decompose a graph into two smaller graphs by contracting vertices; this will be called the *cut decomposition* operation. Then we combine the solutions (i.e. arc-disjoint paths in the above proof) of the smaller graphs to give a solution of the original graph. Notice that the solutions of the smaller graphs have intersections (i.e. the arcs in $\delta_D^{in}(X)$) and the solution in D_1 defines a partial solution on D_2 (i.e. which arc of $\delta_D^{out}(v_2)$ belongs to which path). To give a solution of the original graph, one needs to argue that any such partial solution on D_2 can be *extended* to a solution of D_2; we will say this is the *extension property*. The cut decomposition operation along with the extension property allow us to assume very restricted structures on the graph (e.g. there is no edge between two Steiner vertices) which greatly simplify the analysis. In the above proof, the extension property comes naturally (i.e. one just needs to rename the paths of D_2). However, in future problems, one needs to define and prove some appropriate extension property to make the above step works and this is usually the most difficult step in the proof. Finally, we remark that the operation of replacing two edges (or arcs) sv and vt by st is called the *edge splitting-off* operation. This turns out to be a very important operation.

2.2.2 Different Versions of Menger's theorem

From the directed version of Menger's theorem, one can derive all the following results. They are very similar, but we include all of them for the ease of further references. The following proposition follows immediately from Theorem 2.2.1.

Proposition 2.2.2 *For a digraph $D = (V, A)$, a subset $S \subseteq V$, and a positive integer k, the following are equivalent:*

1. *There are k arc-disjoint paths from any vertex of S to any other vertex of S.*

2. $d_D^{in}(X) \geq k$ *for every* S*-separating set* X.

3. S *remains strongly connected in* D *upon removal of* $k-1$ *arcs.*

Menger's theorem for undirected graphs

Theorem 2.2.3 (Menger [73]) *Let* $G = (V, E)$ *be a graph, and* $s, t \in V$ *be distinct vertices. There are* k *edge-disjoint paths between* s *and* t *if and only if*

$$d(X) \geq k \text{ for every } \overline{s}t \text{ set } X \subseteq V. \tag{2.4}$$

Proposition 2.2.4 *For a graph* $G = (V, E)$*, a subset* $S \subseteq V$*, and a positive integer* k*, the following are equivalent:*

1. *There are* k *edge-disjoint paths between any two vertices of* S.

2. $d_G(X) \geq k$ *for every* S*-separating set* X.

3. S *remains connected in* G *upon removal of* $k-1$ *edges.*

Menger's theorem for hypergraphs

Theorem 2.2.5 *Let* $H = (V, \mathcal{E})$ *be a hypergraph, and* $s, t \in V$ *be distinct vertices. There are* k *hyperedge-disjoint paths between* s *and* t *if and only if*

$$d_H(X) \geq k \text{ for every } \overline{s}t \text{ set } X \subseteq V. \tag{2.5}$$

Proposition 2.2.6 *For a hypergraph* $H = (V, \mathcal{E})$*, a subset* $S \subseteq V$*, and a positive integer* k*, the following are equivalent:*

1. *There are* k *hyperedge-disjoint paths between any two vertices of* S.

2. $d_H(X) \geq k$ *for every* S*-separating set* X.

3. S *remains connected in* H *upon removal of* $k-1$ *hyperedges.*

Menger's theorem for directed hypergraphs

The following results hold for both in-hypergraphs and out-hypergraphs.

Theorem 2.2.7 *Let $\overrightarrow{H} = (V, \overrightarrow{\mathcal{E}})$ be a directed hypergraph, and $s, t \in V$ be distinct vertices. There are k hyperarc-disjoint paths between s and t if and only if*

$$d_{\overrightarrow{H}}(X) \geq k \text{ for every } \overline{s}t \text{ set } X \subseteq V. \tag{2.6}$$

Proposition 2.2.8 *For a directed hypergraph $\overrightarrow{H} = (V, \overrightarrow{\mathcal{E}})$, a subset $S \subseteq V$, and a positive integer k, the following are equivalent:*

1. *There are k hyperarc-disjoint paths from any vertex of S to any other vertex of S.*
2. *$d^{in}_{\overrightarrow{H}}(X) \geq k$ for every S-separating set X.*
3. *S remains strongly connected in $\overrightarrow{H}$ upon removal of $k-1$ hyperarcs.*

2.2.3 Edge Splitting-Off Preserving Edge-Connectivity

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. Notice that parallel edges and loops may arise. However, any loop created will be removed. The resulting graph will be denoted by G^{ef}. In the proof of Theorem 2.2.1, we have already used the splitting-off technique (although in a very simple manner).

When a splitting-off operation is performed, the local edge-connectivity never increases. The content of the splitting-off theorems 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-connectivities between vertices distinct from v.

These theorems prove to be extremely powerful in attacking edge-connectivity problems; we shall see a couple of examples later. The first splitting-off theorem was proved by Lovász in 1974 (see [67]). The proof below is due to Frank [27]; we sketch it here for completeness.

Theorem 2.2.9 *Suppose that in an undirected graph $G = (V, E)$*

$$d(X) \geq K = 2k \text{ for every } \emptyset \neq X \subset V - s \tag{2.7}$$

where $s \in V$ is a given vertex of even degree. Then for every edge $f = st$ there is an edge $e = su$ so that $\{e, f\}$ can be split off without violating (2.7).

Proof. (see [67, 27]) Call a set $\emptyset \neq X \subset V - s$ *dangerous* if $d(X) \leq K + 1$. Splitting-off a pair of edges $\{e, f\}$ is said to be *suitable* if it does not destroy (2.7). Clearly, splitting-off $\{e, f\}$ is suitable if and only if there is no dangerous set X with $u, t \in X$ and $s \notin X$.

Lemma 2.2.10 *The union of two dangerous $\overline{s}t$ sets is dangerous.*

Proof. Let X and Y be two dangerous $\overline{s}t$ sets. If $X \subseteq Y$ or $Y \subseteq X$, then we have nothing to prove. So we assume $X - Y \neq \emptyset$ and $Y - X \neq \emptyset$. Since $t \in X \cap Y$ and $s \notin X \cup Y$, it follows that $d(X \cap Y, V - (X \cup Y)) \geq 1$. By Proposition 2.1.2, we have $(K+1) + (K+1) \geq d(X) + d(Y) = d(X - Y) + d(Y - X) + 2d(X \cap Y, V - (X \cup Y)) \geq K + K + 2$. So, $d(X) = K + 1, d(Y) = K + 1, d(X - Y) = K, d(Y - X) = K$, and $d(X \cap Y, V - (X \cup Y)) = 1$. From this $d(X, V - (X \cup Y)) = d(Y, V - (X \cup Y))$ follows; otherwise, say the left hand side is smaller, then a simple but tedious counting argument shows that $d(X - Y) < d(Y - X) = K$, contradicting (2.7). Therefore, $d(X \cup Y) = 2d(X, V - (X \cup Y)) + 1$, which is an odd number. Since s is a vertex of even degree, this implies that $X \cup Y \subset V - s$.

Suppose, by way of contradiction, that $X \cup Y$ is not dangerous. Then $d(X \cup Y) \geq K + 2$. In fact, we must have $d(X \cup Y) \geq K + 3$ since $d(X \cup Y)$ is an odd number. Now, by Proposition 2.1.2, $(K+1)+(K+1) = d(X)+d(Y) \geq d(X \cap Y)+d(X \cup Y) \geq K+(K+3)$ and this contradiction proves the claim. ∎

From Lemma 2.2.10 it follows that the union M of all dangerous $\overline{s}t$-sets is dangerous; in particular, $M \subset V - s$. Now there must be an edge $e = su$ with $u \notin M$ since otherwise $d(V - (M + s)) = d(M + s) = d(M) - d(s) \leq d(M) - 2 \leq K - 1$ contradicting (2.7). By the choice of M, the splitting-off of $\{e, f\}$ is splittable. ∎

Mader [71], answering an earlier conjecture of Lovász, proved the following powerful generalization of Lovasz's result.

Theorem 2.2.11 (MADER'S SPLITTING-OFF LEMMA [71]) *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 $x, y \in V - s$.*

Mader's proof, which does not use submodularity explicitly, is quite complicated. Frank [26] extended the idea of Lovász's proof to give a considerably simpler proof of Mader's theorem. In fact, the proof of Theorem 2.2.9 is a prototype for proofs of many splitting-off theorems.

Notice that in the above theorems, if s is a vertex of even degree, then we can apply a suitable splitting-off operation at s repeatedly until s is of degree 0. This is called a *complete splitting-off* at s.

2.2.4 Graph Orientations Achieving High Arc-Connectivity

Graph orientations provide a link between graphs and digraphs. There is a huge literature on results concerning graph orientations satisfying certain properties. In this subsection we only focus on graph orientations achieving high (local-)arc-connectivity.

The underlying graph of any strongly k-arc-connected digraph is $2k$-edge-connected. Is every $2k$-edge-connected graph the underlying graph of some strongly k-arc-connected digraph? The special case when $k = 1$ was proved by Robbins [81] in 1939. The general case is proved by Nash-Williams [74] in 1960. We present a simple proof using the Lovász splitting-off lemma (Theorem 2.2.9).

Theorem 2.2.12 (Nash-Williams Weak Orientation Theorem) [74] *The edges of an undirected graph G can be oriented so that the resulting directed graph is strongly k-arc-connected if and only if G is $2k$-edge-connected.*

Proof. We reproduce the proof from [27]. The necessary condition is trivial. We prove the sufficient condition by induction on the number of edges of G, which is $2k$-edge-connected by assumption. If $G-e$ is $2k$-edge-connected, then $G-e$ has a strongly k-arc-connected orientation and so does G. So assume that for every $e \in E(G)$, $G-e$ is not $2k$-edge-connected. By Menger's theorem (Proposition 2.2.4), there exists a set X for which $d_{G-e}(X) = 2k-1$, and thus $d_G(X) = 2k$.

First we claim that G has a vertex of degree $2k$. Suppose not. We consider a minimal set X for which $d_G(X) = 2k$. Since every vertex is of degree greater than $2k$, X contains at least two vertices and at least one edge. Pick an arbitrary edge $e = uv$ induced in X. Since $G-e$ is not $2k$-edge-connected, there exists a nontrivial $u\overline{v}$ set Y with $d(Y) = 2k$. Notice that $X \cup Y \neq V(G)$, for otherwise $V(G) - Y$ contradicts the minimality of X. Now, by Proposition 2.1.2, $2k + 2k = d(X) + d(Y) \geq d(X \cup Y) + d(X \cap Y) \geq 2k + 2k$. So $d(X \cap Y) = 2k$ and hence $X \cap Y$ contradicts the minimality of X. Therefore, G has a vertex of degree $2k$.

Let s be a vertex of degree $2k$. By Theorem 2.2.9, there is a complete splitting-off at s so that the resulting graph G' is $2k$-edge-connected. By induction there is a strongly k-arc-connected orientation D' of G'. Now, suppose $uv \in E(D')$ is an edge obtained from splitting-off $su, sv \in E(D)$ and uv is oriented as $\overrightarrow{uv}$ in D', then we orient $su, sv \in E(D)$ as $\overrightarrow{us}, \overrightarrow{sv}$ (see Figure 2.3 for an illustration). All other edges of D which are not adjacent to s are oriented as in D'. We claim that D is a strongly k-arc-connected orientation. By Menger's theorem (Proposition 2.2.2), it suffices to check $\delta^{in}_D(X) \geq k$ for

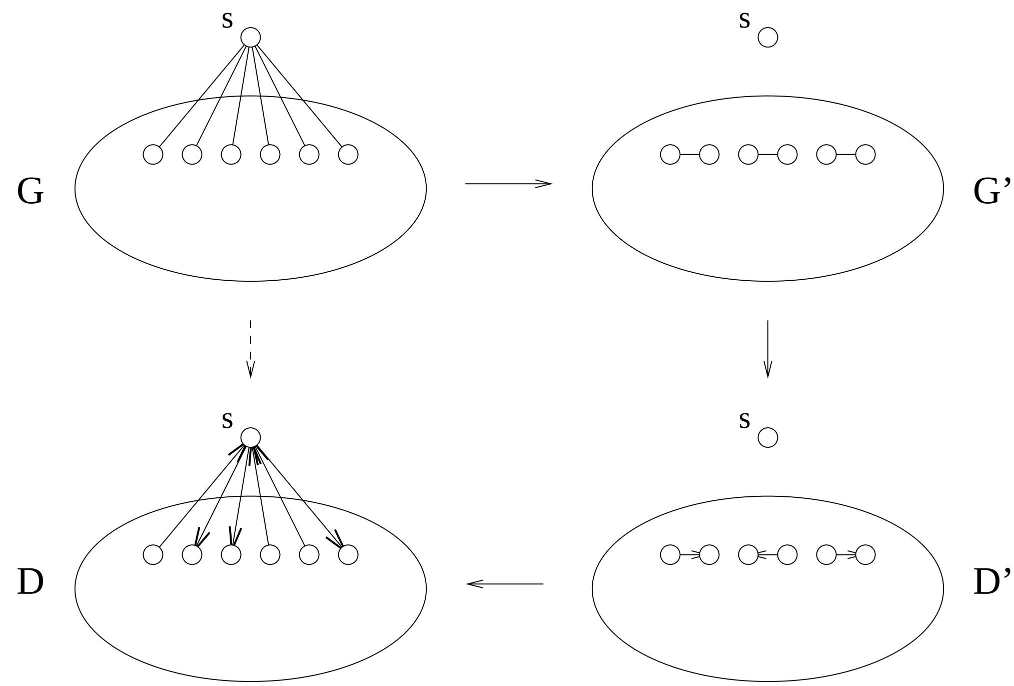

Figure 2.3: An illustration of the proof of Theorem 2.2.12.□

every $\emptyset \neq X \subset V$. If $s \notin X$, then $\delta_D^{in}(X) = \delta_{D'}^{in}(X) \geq k$ as required. If $X = \{s\}$, then $\delta_D^{in}(X) = \delta_D^{in}(s) = k$. Finally, if $\{s\} \subset X$, then $\delta_D^{in}(X) = \delta_{D'}^{in}(X - s) \geq k$ as required. So D is a strongly k-arc-connected orientation; this completes the proof. ∎

In fact, Nash-Williams has proved a much stronger theorem which achieves optimal local-arc-connectivity for all pair of vertices.

Theorem 2.2.13 (Nash-Williams Strong Orientation Theorem) [74] *Every undirected graph $G = (V, E)$ has an orientation D so that*

$$\lambda_D(x, y) = \lfloor \lambda_G(x, y)/2 \rfloor \text{ for all } x, y \in V. \tag{2.8}$$

In addition, the orientation can be chosen such that the difference between the indegree and the outdegree of each vertex is at most 1.

Nash-Williams calls an orientation satisfying (2.8) *well-balanced*. Nash-Williams' proof uses a sophisticated inductive argument. The starting idea of Nash-Williams' approach is an observation that the theorem is trivial for Eulerian graphs. Indeed, an Eulerian graph always has an orientation so that the resulting digraph is Eulerian, and this orientation satisfies (2.8). For a graph that is not Eulerian, Nash-Williams' idea is to augment it to an Eulerian graph by adding a perfect matching M on the vertices with odd degree, find an Eulerian orientation of the resulting graph $G + M$, and finally leave out the edges of M. Naturally, the resulting orientation of the original graph can be expected to satisfy (2.8) only if the auxiliary matching fulfills certain requirements. Nash-Williams' proves that such a matching, which is called a *feasible odd-vertex pairing*, always exists.

Lovász' splitting-off lemma immediately implies Nash-Williams' weak orientation theorem. One may wonder if Mader's splitting-off lemma would also imply immediately Nash-Williams' strong orientation theorem. Mader [71] was indeed able to prove the strong orientation theorem relying on his splitting-off lemma. The proof, however, can hardly be considered simpler than the original one.

Frank [27] gave a more illuminating proof by combining the ideas from Nash-Williams' and Mader's proofs as well as from his short proof of Mader's splitting-off lemma. He leaves it as a challenge to find a really simple proof and an ultimate answer to Nash-Williams hopes cited below.

> *The comparatively complicated nature of the foregoing proof, ... as contrasted with the comparatively simple and natural character of Theorems 1 and 2 might suggest that conceivably the most simple, natural, and insightful proof of those theorems has not yet been found... I have sometimes wondered whether there might be a way of using matroids, or something like matroids, to give a better and more illuminating proof of our two theorems.*
>
> *They do not seem particularly closely related to much other existing work in graph theory, ... these theorems seem to have a somewhat natural character*

which would suggest that there must ultimately be a place for them in the overall structure of graph theory.

As we shall see in the next subsection, as far as the weak orientation theorem is concerned, Nash-Williams' anticipation was correct. There are simple proofs for this result using submodular functions - structures with "somewhat matroid-like" features.

2.2.5 Submodular Flows and Graph Orientations

In this subsection we shall introduce a powerful generalization of flows due to Edmonds and Giles [18] - submodular flows. Many important theorems in graph theory and combinatorial optimization are special cases of this theory. We shall only focus on its applications to (hyper)graph orientations.

Let $D = (V, A)$ be a digraph, $\mathcal{F}$ be a crossing family of subsets of V, and $b : \mathcal{F} \to \mathbb{Z}$ be a crossing submodular function. Given such $D, \mathcal{F}, b$, a *submodular flow* is a function $x : A \to \mathbb{R}$ satisfying:

$$x^{in}(U) - x^{out}(U) \leq b(U) \text{ for each } U \in \mathcal{F}, \tag{2.9}$$

where $x^{in}(U)$ is a shorthand for $x(\delta^{in}(U))$ and $x^{out}(U)$ is a shorthand for $x(\delta^{out}(U))$. Equivalently, given a crossing supermodular function $p : \mathcal{F} \to \mathbb{Z}$, a function $x : A \to \mathbb{R}$ satisfying

$$x^{in}(U) - x^{out}(U) \geq p(U) \text{ for each } U \in \mathcal{F} \tag{2.10}$$

is a submodular flow.

Given two functions $f : A \to \mathbb{Z} \cup \{-\infty\}$ and $g : A \to \mathbb{Z} \cup \{\infty\}$, a submodular flow is *feasible* with respect to f, g if $f(a) \leq x(a) \leq g(a)$ holds for all $a \in A$. The set of feasible submodular flows (with respect to given $D, \mathcal{F}, b, f, g$) has very nice properties which makes submodular flows a very powerful tool in combinatorial optimization.

Theorem 2.2.14 (THE EDMONDS-GILES THEOREM) [18] *Let $D = (V, A)$ be a directed multigraph. Let $\mathcal{F}$ be a crossing family of subsets of V such that $\emptyset, V \in \mathcal{F}$, let $b : \mathcal{F} \to \mathbb{Z} \cup \{\infty\}$ be crossing submodular on $\mathcal{F}$ with $b(\emptyset) = b(V) = 0$, and let $f \leq g$ be functions on A such that $f : A \to \mathbb{Z} \cup \{-\infty\}$ and $g : A \to \mathbb{Z} \cup \{\infty\}$. The system of linear inequalities*

$$\{f \leq x \leq g \text{ and } x^{in}(U) - x^{out}(U) \leq b(U) \qquad \text{for all } U \in \mathcal{F}\} \tag{2.11}$$

has an integer optimal solution (provided it has a solution) for any objective function

$$\min\{\sum_{a \in A(D)} c(a) \cdot x(a) \ : \ x \text{ satisfies (2.11)}\},$$

where $c : A \to \mathbb{R}$ is a (cost) function on A.

Since linear programs can be solved in polynomial time, the above theorem implies that any discrete optimization problem that can be modeled as a submodular flow problem can be solved in polynomial time. For some applications, we also need to characterize when a feasible flow exists. The following theorem, characterizing when a feasible submodular flow exists with respect to functions f, g and b is due to Frank [24]:

Theorem 2.2.15 (FEASIBILITY THEOREM FOR FULLY SUBMODULAR FLOWS) [24] *Let $D = (V, A)$ be a directed multigraph, let $f \leq g$ be modular functions on A such that $f : A \to \mathbb{Z} \cup \{-\infty\}$ and $g : A \to \mathbb{Z} \cup \{\infty\}$ and let b be a fully submodular function on 2^V. There exists an integer valued feasible submodular flow if and only if*

$$f^{in}(U) - g^{out}(U) \leq b(U) \qquad \text{for all } U \subseteq V.$$

In particular there exists a feasible integer valued submodular flow if and only if there exists any feasible submodular flow.

Frank [24] also proved the following feasibility theorem for intersecting submodular flows, which we will use later.

Theorem 2.2.16 (FEASIBILITY THEOREM FOR INTERSECTING SUBMODULAR FLOWS) [24] *Let $D = (V, A)$ be a directed multigraph and let $f \leq g$ be modular functions on A such that $f : A \to \mathbb{Z} \cup \{-\infty\}$ and $g : A \to \mathbb{Z} \cup \{\infty\}$. Let $\mathcal{F}$ be an intersecting family of subsets of V such that $\emptyset, V \in \mathcal{F}$ and let b be an intersecting submodular function on $\mathcal{F}$. Then there exists a feasible submodular flow with respect to f, g and b if and only if*

$$f^{in}(\bigcup_{i=1}^{t} X_i) - g^{out}(\bigcup_{i=1}^{t} X_i) \leq \sum_{i=1}^{t} b(X_i) \tag{2.12}$$

holds whenever $X_1, \ldots, X_t$ are disjoint members of $\mathcal{F}$. Furthermore, if f, g, b are all integer valued functions, then there exists a feasible integer valued submodular flow with respect to f, g and b.

For completeness, we also present the feasibility theorem for crossing supermodular function; note that this will not be used later.

Theorem 2.2.17 (FEASIBILITY THEOREM FOR CROSSING SUBMODULAR FLOWS) [24] *Let $D = (V, A)$ be a directed multigraph and let $f \leq g$ be modular functions on A such that $f : A \to \mathbb{Z} \cup \{-\infty\}$ and $g : A \to \mathbb{Z} \cup \{\infty\}$. Let $\mathcal{F}$ be a crossing family of subsets of V such that $\emptyset, V \in \mathcal{F}$ and let b be a crossing submodular function on $\mathcal{F}$. Then there exists a feasible submodular flow with respect to f, g and b if and only if*

$$f^{in}(\bigcup_{i=1}^{t} X_i) - g^{out}(\bigcup_{i=1}^{t} X_i) \leq \sum_{i=1}^{t} b(X_{ij})$$

holds for every subpartition $X_1, \ldots, X_t$ of V such that every X_i is the intersection of co-disjoint sets $X_{i1}, \ldots, X_{il}$. Furthermore, if f, g, b are all integer valued functions, then there exists a feasible integer valued submodular flow with respect to f, g and b.

Now we show how to use submodular flows as a tool to tackle graph orientation problems.

First we relate graph orientations achieving high arc-connectivity to graph orientations that "cover" a given supermodular function. A graph G is said to *cover* a set function p if $d_G(X) \geq p(X)$ for every $X \subseteq V$. Similarly, a digraph D covers p if $\delta_D^{in}(X) \geq p(X)$ for every $X \subseteq V$. By setting $p(X) := k$ for $\emptyset \neq X \subset V$, $p(\emptyset) := 0$ and $p(V) := 0$, by Menger's theorem (Proposition 2.2.2), an orientation D of G covers p if and only if D is a strongly k-arc-connected orientation of G. Notice the above p is crossing supermodular but not intersecting supermodular. Choose any vertex r to be the root. By setting $p(X) := k$ for $X \subseteq V$ with $r \notin X$ and $p(X) := 0$ otherwise, by Menger's theorem (Theorem 2.2.1), an orientation D of G covers p if and only if D is a rooted k-arc-connected orientation of G. Notice that this p is intersecting supermodular but not fully supermodular; for example, inequality (2.2) does not hold for two disjoint sets $X_1, X_2 \subset V$.

Here we reduce the problem of finding an orientation covering an intersecting supermodular function h to a submodular flow problem, which will have applications in Chapter 4. The approach taken is due to Frank [27]. First, we start with an arbitrary orientation D of G as a reference orientation. Clearly, G has an orientation covering h if and only if it is possible to reorient some arcs of D so as to get an orientation covering h. Suppose we interpret the function $x : A \to \{0, 1\}$ as follows: $x(a) = 1$ if we reorient a in D and $x(a) = 0$ means that we leave the orientation of a as it is in D. Then G has an orientation covering h if and only if we can choose x so that the following holds:

$$\delta_D^{in}(U) - x^{in}(U) + x^{out}(U) \geq h(U) \text{ for all } U \subset V.$$

This is equivalent to:

$$x^{in}(U) - x^{out}(U) \leq \delta_D^{in}(U) - h(U) = b(U) \text{ for all } U \subset V.$$

Since $\delta_D^{in}(U)$ is fully submodular and $h(U)$ is intersecting supermodular, the function b is intersecting submodular. Note this is exactly the formulation of a submodular flow problem. Now Theorem 2.2.16 implies that the existence of an orientation if and only if

$$f^{in}(\bigcup_{i=1}^{t} X_i) - g^{out}(\bigcup_{i=1}^{t} X_i) \leq \sum_{i=1}^{t} b(X_i)$$

holds whenever $X_1, \ldots, X_t$ are disjoint members of $\mathcal{F}$ (i.e. $X_i \cap X_j = \emptyset$ for $i \neq j$). In this orientation problem, $f \equiv 0$ and $g \equiv 1$. By the definition of $b := \delta_D^{in} - h$, the above equation is equivalent to

$$-\delta_D^{out}(\bigcup_{i=1}^{t} X_i) \leq \sum_{i=1}^{t}(\delta_D^{in}(X_i) - h(X_i)).$$

Rearranging the terms, we have

$$\sum_{i=1}^{t} h(X_i) \leq \sum_{i=1}^{t} \delta_D^{in}(X_i) + \delta_D^{out}(\bigcup_{i=1}^{t} X_i).$$

Notice that since $X_1, \ldots, X_t$ are disjoint, the right hand side just counts the number of undirected edges in G which enter some X_i (that is, edges with precisely one endpoint in some X_i). Therefore, we have the following theorem which is important to the STEINER ROOTED ORIENTATION problem.

Theorem 2.2.18 *Let $G = (V, E)$ be an undirected graph. Let $h : 2^V \rightarrow Z \cup \{-\infty\}$ be an intersecting supermodular function with $h(\emptyset) = h(V) = 0$. Then there exists an orientation D of G satisfying*

$$\delta_D^{in}(X) \geq h(X) \text{ for all } X \subset V$$

if and only if

$$e_{\mathcal{P}} \geq \sum_{i=1}^{t} h(X_i)$$

holds for every subpartition $\mathcal{P} = \{X_1, X_2, \ldots, X_t\}$ of V. Here $e_{\mathcal{P}}$ counts the number of edges which enter some member of $\mathcal{P}$.

From this model many results on graph orientations can be derived. For example, the Nash-Williams' weak orientation theorem can be obtained immediately as follows: Given a $2k$-edge-connected undirected graph G. First find an arbitrary orientation D of G. Then G has a strongly k-arc-connected orientation if and only if we can choose x so that

$$d_D^{in}(U) + x^{out}(U) - x^{in}(U) \geq k \qquad \text{for all } \emptyset \neq U \subset V.$$

Now one just needs to verify that $x \equiv \frac{1}{2}$ is a feasible submodular flow if G is $2k$-edge-connected. Then by Theorem 2.2.14 there is an integral solution as well, and this gives us a strongly k-arc-connected orientation. Furthermore, we can also solve the weighted version where the two possible orientations of an edge may have different costs and the goal is to find the cheapest strongly k-arc-connected orientation of the graph; this again follows from the Edmonds-Giles theorem. Finally, we remark that using submodular flows is the only known method to solve the above weighted problem, as well as other more advanced orientation problems (e.g. orientations of mixed graphs, degree-constrainted orientations, etc). Also, we remark that there seems to be no easy way of formulating orientation problems concerning local-arc-connectivities (e.g. Nash-Williams strong orientation theorem) as submodular flow problems.

2.2.6 Disjoint Trees and Disjoint Arborescences

One approach to capture high edge-connection is to require the graph to be not dismantleable into smaller parts by leaving out only a few edges. Another possible approach is to require the graph or digraph to contain several edge-disjoint simple connected constituents. Menger showed that these two approaches are equivalent when the simple connected constituents are paths (see Section 2.2.2). Paths, however, only capture the local-edge-connectivity between two vertices.

To capture the global edge-connectivity of a graph, one naturally comes up with the notion of edge-disjoint spanning trees. When does a graph have k edge-disjoint spanning trees? If it does, then the graph is clearly k-edge-connected. On the other hand, it is not clear whether any edge-connectivity will imply the existence of k edge-disjoint spanning trees. Tutte [85] and Nash-Williams [75] show that this is closely related to partition-connectivity (see Section 2.1.8 for definitions).

Theorem 2.2.19 *(Tutte [85]; Nash-Williams [75]) A multigraph contains k edge-disjoint spanning trees if and only if it is k-partition-connected.*

This implies the following surprising corollary.

Corollary 2.2.20 *Every $2k$-edge-connected multigraph G has k edge-disjoint spanning trees.*

Proof. Consider a partition $\mathcal{P} = \{P_1, \ldots, P_t\}$ of $V(G)$. Since G is $2k$-edge-connected, $e_G(\mathcal{P}) = \frac{1}{2}\sum_{i=1}^{t} d_G(P_i) \geq \frac{1}{2}\sum_{i=1}^{t} 2k = kt$. This implies that G is k-partition-connected, and the corollary follows from Theorem 2.2.19. ∎

To generalize the Tutte-Nash-Williams theorem to hypergraphs, it is natural to replace spanning trees by connected spanning sub-hypergraphs. In hypergraphs, however, being k-partition-connected is a sufficient but not a necessary condition to have k hyperedge-disjoint connected spanning sub-hypergraphs. In fact, while being k-partition-connected is in P, having k hyperedge-disjoint connected spanning sub-hypergraphs is NP-complete.

Theorem 2.2.21 *[34] The problem of deciding whether a hypergraph $H = (V, \mathcal{E})$ can be decomposed into k connected spanning sub-hypergraphs is NP-complete for every integer $k \geq 2$.*

A less intuitive generalization, which allows the necessary condition to hold, is to replace spanning trees by partition-connected spanning sub-hypergraphs. Note that a spanning tree is a partition-connected spanning subgraph. Frank, Király and Kriesell [34] proved the following theorem on decomposing a hypergraph into partition-connected sub-hypergraphs.

Theorem 2.2.22 *[34] A hypergraph H can be decomposed into k hyperedge-disjoint partition-connected sub-hypergraphs if and only if H is k-partition-connected.*

They also derive the following corollary which is a generalization of Corollary 2.2.20. Recall that a hypergraph is of rank r if every hyperedge is of size at most r.

Corollary 2.2.23 *[34] A rk-edge-connected hypergraph H of rank r can be decomposed into k partition-connected sub-hypergraphs, and hence into k connected spanning subhypergraphs.*

As we shall see in Chapter 3, these results on hypergraphs are the basis for the STEINER TREE PACKING problem. The proof of Theorem 2.2.22 is based on matroid theory, which we have not discussed. Instead, we shall outline a proof using results on graph orientations. First we present the following powerful theorem by Edmonds [17]. The short and elegant proof is due to Lovász [66] using submodularity.

Theorem 2.2.24 *[17] A directed graph has k edge-disjoint spanning r-arborescences if and only if the following cut condition:*

$$d^{in}(X) \geq k \tag{2.13}$$

holds for every set $X \subseteq V - r$.

Proof. We reproduce the proof from [66]. The necessity is easy. We prove the sufficiency by induction on k. The base case $k = 0$ is trivial. Let F be a set of arcs such that (i) F is an r-arborescence, and (ii) $d^{in}_{D-F}(X) \geq k - 1$ for every set $X \subseteq V - r$. Given an F that satisfies (i) and (ii) and $V(F) \subset V(D)$ ($F = \emptyset$ at the beginning), we shall show that we can always add an arc a so that $F + a$ still satisfies (i) and (ii). Therefore, eventually F will cover all the vertices and the theorem follows by induction.

Call a set $X \subseteq V - r$ *tight* if $d^{in}_{D-F}(X) = k-1$, notice that any tight set must intersect $V(F)$ by (2.13). Consider a minimal tight set S not contained in $V(F)$; if no such set exists, then $F + a$ satisfies (i) and (ii) for any arc $a = uv$ with $u \in V(F)$ and $v \notin V(F)$. There must be an arc $a = uv$ with $u \in V(F) \cap S$ and $v \in S - V(F)$; for otherwise, $d^{in}_D(S - V(F)) = d^{in}_{D-F}(S - V(F)) \leq d^{in}_{D-F}(S) = k - 1$ which contradicts (2.13). The first equality holds since no arc in F enters $S - V(F)$, the second inequality holds since there is no arc $a = uv$ with $u \in V(F) \cap S$ and $v \in S - V(F)$, and the third equality holds because S is tight.

The next step is the crucial use of submodularity:

Proposition 2.2.25 *The intersection of two intersecting tight sets is a tight set.*

Proof. Let X and Y be two intersecting tight sets. By Proposition 2.1.2, we have $(k-1)+(k-1) = d^{in}_{D-F}(X)+d^{in}_{D-F}(Y) \geq d^{in}_{D-F}(X\cap Y)+d^{in}_{D-F}(X\cup Y) \geq (k-1)+(k-1)$. So we must have equality throughout, and this implies that $X \cap Y$ is tight. ∎

Now we claim that $F + a$ satisfies (i) and (ii). The validity of (i) is trivial. Suppose (ii) does not hold; then there is a tight set Y such that $u \notin Y$ and $v \in Y$. By Proposition 2.2.25, $S \cap Y$ is again a tight set, which contradicts the minimality of S since $u \in S$ but $u \notin Y$. Therefore, (ii) must also hold and this proves the theorem. ∎

To prove the Tutte-Nash-Williams theorem (Theorem 2.2.19), we use the following orientation theorem by Frank [28]. Notice that Frank's theorem is implied by the Tutte-Nash-Williams theorem, but Frank gave a short and direct proof of it.

Theorem 2.2.26 *[28] Given a graph $G = (V, E)$ and a vertex $r \in V$, G has an orientation for which $d^{in}(X) \geq k$ for every $X \subseteq V - r$ if and only if G is k-partition-connected. In other words, G has a rooted k-arc-connected orientation if and only if G is k-partition-connected.*

Proof. We reproduce the proof from [28]. Necessity is clear. To see the sufficiency, extend G by a minimum number of edges rv ($v \in V(G)$) to form G' so that G' has a rooted k-arc-connected orientation D'. If this minimum is zero, then we are done; so assume that it is positive. We can assume that $\delta^{in}_{D'}(r) = 0$. Call a set *critical* if $\delta^{in}_{D'}(X) = k$. The following is straightforward.

Proposition 2.2.27 *The intersection and the union of two critical sets with non-empty intersection are critical.*

Let $e = rt$ be a new arc in the given orientation and let T be the set of vertices reachable from t along a path.

Proposition 2.2.28 *If Z is critical and $Z \cap T \neq \emptyset$, then $Z \subseteq T$.*

Proof. Assume $Z \not\subseteq T$. For $Y := V - T$ we have $k = \delta^{in}(Y) + \delta^{in}(Z) = \delta^{in}(Y \cap Z) + \delta^{in}(Y \cup Z) + d(Y, Z) \geq k + 0 + d(Y, Z) \geq k$, where $d(Y, Z)$ denotes the number of arcs connecting $Y - Z$ and $Z - Y$ (in either direction). From this we get $\delta^{in}(Y \cup Z) = 0$ and $d(Y, Z) = 0$. The first inequality implies that $t \in Z$ (by the definition of T and by the assumption that $T \cap Z \neq \emptyset$), while the second one implies that $t \notin Z$ (because of edge st); this contradiction proves the claim. ∎

Suppose there is a vertex $v \in T$ which is not contained in any critical set. Let P be a directed path from t to v. Reorient the edges of P and discard the edge e. The new orientation is still a rooted k-arc-connected orientation, a contradiction to the minimality of the number of new rv edges.

So we can assume every vertex in T belongs to a critical set. Let $V_1, \ldots, V_{t-1}$ denote the maximal critical sets in T. By Proposition 2.2.27 and Proposition 2.2.28, these are disjoint sets and form a partition of T. Let $V_t := V - T$ and $\mathcal{P} := \{V_1, \ldots, V_t\}$. Since $\delta^{in}_{D'}(V_t) = 0$, we have $k(t-1) = \sum_{i=1}^{t-1} \delta^{in}_{D'}(V_i) = \sum_{i=1}^{t} \delta^{in}_{D'}(V_i) = e_{G'}(\mathcal{P}) > e_G(\mathcal{P})$, where $e_G(\mathcal{P})$ denotes the number of cross edges of $\mathcal{P}$ in G. However, this contradicts the fact that G is k-partition-connected. ∎

Now we are ready to prove the Tutte-Nash-Williams theorem. The necessity is clear. We prove the sufficiency here. Choose an arbitrary vertex r as the root. Since G is k-partition-connected, by Theorem 2.2.26, G has a rooted k-arc-connected orientation D. By Theorem 2.2.24, D can be decomposed into k arc-disjoint r-arborescences. The underlying graph of each of these r-arborescences is a spanning tree in G, this proves the theorem.

To prove the result by Frank, Király and Kriesell (Theorem 2.2.22), one can take a similar approach. Choose an arbitrary vertex r as the root. If H is k-partition-connected, then H can be oriented as a rooted k-hyperarc-connected in-hypergraph $\overrightarrow{H}$, where each hyperarc has a designated head and all other vertices are tails. This step is shown in [33]; one can also derive this from Theorem 2.2.18. Given $\overrightarrow{H}$, by submodularity, one can "shrink" each hyperarc into an ordinary arc without destroying rooted k-arc-connectivity. Here, by shrinking we mean deleting all but one tail of each hyperarc. (The proof of this fact is very similar to the proof of Lemma 4.3.2.) After shrinking every hyperarc into an ordinary arc, we can apply Edmonds' theorem (Theorem 2.2.24) to construct k arc-disjoint r-arborescences. By un-shrinking the hyperarc and ignoring the orientations, we have k hyperedge-disjoint partition-connected spanning sub-hypergraphs, as required.

Chapter 3

Steiner Forest Packing

The results in this chapter are based on [60, 61].

3.1 Introduction

A fundamental result of Menger, proved in 1927, states that for any two vertices $a, b \in V(G)$ the maximum number of edge-disjoint a, b-paths is equal to the minimum size of an a, b-edge-cut [73]. Since then, many *min-max relations* of this type have been being discovered (see [84]), and they are some of the most powerful and beautiful results in combinatorics (e.g. max-flow min-cut, max-matching min-odd-set-cover, etc.). Furthermore, some of the most fundamental polynomial time (exact) algorithms have been designed around such relations.

Like min-max relations in the development of exact algorithms, *approximate min-max relations* are vital in the development of approximation algorithms. For example, a seminal work of Leighton and Rao [62] (which shows that for any n-node multicommodity flow problem with uniform demands, the max-flow for the problem is within an $O(\log n)$ factor of the upper bound implied by the min-cut) leads to approximation algorithms for many different problems.

In this chapter, we present an approximate min-max relation for a generalization of the edge-disjoint a, b-paths problem, namely the STEINER FOREST PACKING problem. Given an undirected multigraph G and a set $\mathcal{S} := \{S_1, \ldots, S_t\}$ of disjoint subsets of vertices of G, a *Steiner $\mathcal{S}$-subgraph F* (or simply an $\mathcal{S}$-subgraph) is a subgraph of G such that each S_i is connected in F for $1 \leq i \leq t$. An acyclic Steiner $\mathcal{S}$-subgraph is called a *Steiner $\mathcal{S}$-forest* (or simply an $\mathcal{S}$-forest). We call each S_i a *terminal group*. Observe that we can assume that each terminal group has at least two vertices. The STEINER FOREST PACKING problem is to find a largest collection of edge-disjoint $\mathcal{S}$-forests. Note that an $\mathcal{S}$-subgraph contains an $\mathcal{S}$-forest as a subgraph. Therefore, to show a graph has k edge-disjoint $\mathcal{S}$-forests, it suffices to show that a graph has k edge-disjoint $\mathcal{S}$-subgraphs. An important special case of the STEINER FOREST PACKING problem is when there is only one terminal group in $\mathcal{S}$, i.e. $\mathcal{S} = \{S\}$, then the problem is known as the STEINER TREE PACKING problem. We call an $\mathcal{S}$-subgraph in such a case an S-subgraph. A Steiner S-tree (or simply an S-tree) is just a minimal S-subgraph.

This STEINER TREE PACKING problem and its generalization (where different specified subsets of vertices have to be connected by edge-disjoint trees) have attracted considerable attention from researchers in different areas. It has applications in routing problems arising in VLSI circuit design [56, 72, 80, 37, 38, 39, 40, 41, 86, 51], where an effective way of sharing different signals amongst cells in a circuit can be achieved by the use of edge-disjoint Steiner trees. It also has a variety of computer network applications such as multicasting [76, 9, 10, 4, 87, 36], video-conferencing [45] and network information flow [83, 63], where simultaneous communications can be facilitated by using edge-disjoint Steiner trees.

When $S = V(G)$, the STEINER TREE PACKING problem is known as the SPANNING TREE PACKING problem. Recall from Chapter 2 (Theorem 2.2.19) that Tutte [85] and Nash-Williams [75] independently proved that a graph has k edge-disjoint spanning trees if and only if $E_G(\mathcal{P}) \geq k(|\mathcal{P}| - 1)$ for every partition $\mathcal{P}$ of $V(G)$ into nonempty classes, where $E_G(\mathcal{P})$ denotes the number of edges connecting distinct classes of $\mathcal{P}$. As a corollary

(Corollary 2.2.20) of the Tutte-Nash-Williams result, every $2k$-edge-connected graph has k edge-disjoint spanning trees. Karger [52] exploited this approximate min-max relation to give the best known algorithm (near linear time) to compute a minimum cut of a graph.

The Steiner Tree Packing problem, however, is NP-complete (a proof will be given in Section 3.8). Therefore, under the assumption that NP $\neq$ co-NP, a min-max relation like the Tutte-Nash-Williams theorem does not exist. Nonetheless, Kriesell [57, 58] conjectures that the approximate min-max corollary of the Tutte-Nash-Williams theorem does generalize to the Steiner Tree Packing problem. In the following, we say a set of vertices S is k-edge-connected in G if for every pair of vertices $a, b \in S$ there are k edge-disjoint paths between a and b in G.

Kriesell's conjecture: [57, 58]

If S is $2k$-edge-connected in G, then G has k edge-disjoint S-trees.

The conjecture is best possible for every k as shown by any $2k$-regular $2k$-edge-connected graph G which is not a complete graph by setting $S = V(G)$ (e.g. a $2k$-dimensional hypercube). To see this, any $2k$-regular graph has kn edges but a spanning tree requires $n - 1$ edges. So there are just not enough edges to have more than k edge-disjoint spanning trees when $n \geq k + 2$, i.e. when G is not a complete graph.

3.1.1 Previous Work

Prior to this work, Kriesell's conjecture was open despite several attempts. It was not known to be true even when $2k$ is replaced by any $o(n) \cdot k$ (not even when $k = 2$ [49]). Similarly, not even a polynomial time $o(n)$ approximation algorithm was known for the Steiner Tree Packing problem. That is, in simple graphs, no known polynomial time approximation algorithm has an asymptotic performance better than the naive algorithm of simply finding one spanning tree.

In the special case where every Steiner vertex has an even degree, Kriesell [58] proves that his conjecture is true. By replacing each edge by two parallel edges, we have the following interesting corollary of this result: if S is $2k$-edge-connected in G, there is a collection of $2k$ S-trees such that every edge is used by at most 2 such S-trees. So, we have a 2-approximation algorithm if we allow *half integral solutions*. Also, the special case where there are no edges between Steiner vertices (i.e. the Steiner vertices induce an independent set) is considered by Kriesell [57] and Frank, Király and Kriesell [34]. In particular, it is proven in [57] that if G has no edges between Steiner vertices and S is $(k+1)k$-edge-connected in G, then G has k edge-disjoint S-trees. This result is improved in [34] by replacing $(k+1)k$ with $3k$; it is based on a generalization of the Tutte-Nash-Williams theorem to hypergraphs using matroid theory. Recently, Kriesell [59] proves that if S is $(l+2)k$-edge-connected in G where l is the maximum size of a *bridge* (see [59] for its definition), then G has k edge-disjoint S-trees; this result is a common generalization of the Tutte-Nash-Williams theorem (when $l = 0$) and the case where there is no edge between Steiner vertices (when $l = 1$).

For the general case, Petingi and Rodriguez [78] prove that if S is $(2(\frac{3}{2})^{|V(G)-S|} \cdot k)$-edge-connected in G, then G has k edge-disjoint S-trees. Kriesell [58], by using the result for the case that every Steiner vertex has an even degree, improves this by weakening the connectivity requirement to $2|V(G) - S| + 2k$. Jain, Mahdian and Salavatipour [49], by using a *shortcutting* procedure, prove that if S is $(|S|/4 + o(|S|))k$-edge-connected in G, then G has k edge-disjoint S-trees; this improves an exponential connectivity bound in terms of $|S|$ obtained earlier by Kriesell [58]. In both papers [58, 49], an optimal bound of $\lceil \frac{4}{3}k \rceil$ on the connectivity requirement is obtained for the case $|S| = 3$.

Jain, Mahdian, Salavatipour also study a natural linear programming relaxation of the STEINER TREE PACKING problem. The FRACTIONAL STEINER TREE PACKING problem is formulated [49] by the following linear program. In the following $\mathcal{T}$ denotes the collection of all S-trees in a graph G, and c_e is the given *capacity* of the edge e.

$$
\begin{aligned}
\text{maximize} \quad & \textstyle\sum_{T \in \mathcal{T}} x_T \\
\text{subject to} \quad & \forall e \in E : \textstyle\sum_{T \in \mathcal{T}} x_T \leq c_e \\
& \forall T \in \mathcal{T} : x_T \geq 0
\end{aligned}
\tag{3.1}
$$

By using the Ellipsoid algorithm on the dual of the above linear program, Jain, Mahdian and Salavatipour [49] show that there is a polynomial time α-approximation algorithm for the FRACTIONAL STEINER TREE PACKING problem if and only if there is a polynomial time α-approximation algorithm for the MINIMUM STEINER TREE problem. The MINIMUM STEINER TREE problem is to find a minimum weight S-tree for a given weighted graph. Robins and Zelikovsky [82] give a 1.55 approximation algorithm, and Bern and Plassmann [3] show that it is APX-hard (no polynomial time approximation scheme unless P=NP). Therefore, by using the results of the MINIMUM STEINER TREE problem, the FRACTIONAL STEINER TREE PACKING problem is APX-hard but can be approximated within a factor of 1.55 to the optimal solution [49]. As a corollary, the (integral) STEINER TREE PACKING problem is shown to be APX-hard [49].

Besides designing approximation algorithms, effort has been put in to designing faster exact algorithms by integer programming approaches [72, 80, 37, 38, 39, 40, 41, 86, 51] as well as designing practical heuristic methods [76, 9, 10, 4, 87, 36, 45, 83].

All the previous work mentioned above is related to the STEINER TREE PACKING problem. For the STEINER FOREST PACKING problem, Chekuri and Shepherd [12] gave a 2-approximation algorithm when the input graph is Eulerian. Specifically, they show that given an Eulerian graph, if each S_i is $2k$-edge-connected in G, then G has k edge-disjoint $\mathcal{S}$-forests. A related problem is the MINIMUM STEINER FOREST problem, where the goal is to find a minimum cost $\mathcal{S}$-forest F in G. Goemans and Williamson [42] gave a primal-dual 2-approximation algorithm for the MINIMUM STEINER FOREST problem. Chekuri and Shepherd show that their result, when used with linear programming, implies an alternative 2-approximation algorithm for the MINIMUM STEINER FOREST problem.

3.1.2 Our Results

The main result of this chapter is the following sufficient condition for the existence of k edge-disjoint $\mathcal{S}$-forests.

Theorem 3.1.1 *If each S_i is $30k$-edge-connected in G, then G has k edge-disjoint $\mathcal{S}$-forests.*

In the rest of this chapter, we let $Q := 30$ for notational convenience. In the special case of STEINER TREE PACKING problem, a slightly better bound can be obtained. This answers Kriesell's conjecture affirmatively up to a constant factor.

Theorem 3.1.2 *If S is $24k$-edge-connected in G, then G has k edge-disjoint S-trees.*

The proof of Theorem 3.1.1 has three main ingredients: (i) a new technique to help decompose a general graph into graphs with special structures, which we introduce in Section 3.3.3 (and we foreshadowed in Section 2.2.1), (ii) the edge splitting-off lemma by Mader [71], and (iii) a result by Frank, Király and Kriesell [34] on packing hypertrees. The proof is constructive so that if each S_i is $30k$-edge-connected in G, then a collection of k edge-disjoint $\mathcal{S}$-forests can be constructed in polynomial time. For a graph to have k edge-disjoint $\mathcal{S}$-forests, each terminal group must be at least k edge-connected in G. Therefore, Theorem 3.1.1 implies the first polynomial time constant factor approximation algorithm for the STEINER FOREST PACKING and the STEINER TREE PACKING problem (see Section 3.9 for algorithmic aspects of Theorem 3.1.1).

The CAPACITATED STEINER FOREST PACKING problem is a generalization of the STEINER FOREST PACKING problem where each edge e has an integer capacity c_e which bounds the number of forests that can use e (the STEINER FOREST PACKING problem is the special case where $c_e = 1$ for all $e \in E(G)$). We extend the algorithm for the STEINER FOREST PACKING problem to give a polynomial time constant factor approximation algorithm for the CAPACITATED STEINER FOREST PACKING problem (see Section 3.10).

3.2 Overview of the Main Proof

In this section, we present the ideas leading to the proof of Theorem 3.1.1.

Steiner Tree Packing in Quasi-Bipartite Graphs

To understand our approach, it is illuminating to start with the ground work. In [34], Frank, Király and Kriesell consider a hypergraph generalization of the SPANNING TREE PACKING problem. Recall that a hypergraph H is k-partition-connected if $E_H(\mathcal{P}) \geq k(|\mathcal{P}| - 1)$ holds for every partition $\mathcal{P}$ of $V(H)$ into non-empty classes, where $E_H(\mathcal{P})$ denotes the number of hyperedges intersecting at least two classes. The main theorem in [34] is the following.

Theorem 3.2.1 *A hypergraph H is k-partition-connected if and only if H can be decomposed into k sub-hypergraphs each of which is 1-partition-connected.*

The proof is based on the observation that the hyperforests (see [34] for the definition of a hyperforest) of a hypergraph form the family of independent sets of a matroid and thus Edmonds' matroid partition theorem can be applied. We have also outlined an alternative proof in Section 2.2.6 based on an orientation result of Frank and an arborescence packing result of Edmonds.

Consider an instance of the STEINER TREE PACKING problem where G has no edge between Steiner vertices; such a G is referred as a *quasi-bipartite graph* in the literature. We construct a hypergraph H with vertex set S, where S is the set of terminal vertices of G. For every Steiner vertex v in G, we add a hyperedge in H consisting of the neighbours $N_G(v)$ of v in G (notice that $N_G(v) \subseteq S$ because G is quasi-bipartite). Also, for every edge $uv \in E(G)$, $u, v \in S$, we add an edge uv in H. As we shall see in Section 3.5, we can assume every Steiner vertex is of degree 3 in G by using Mader's splitting-off lemma. So the hyperedges in H are of size at most 3. If we have k edge-disjoint connected sub-hypergraphs of H, then we can easily construct k edge-disjoint S-trees in G; we just need to replace a hyperedge in H by the corresponding Steiner vertex and the edges incident to it. In fact, the converse is also true since every Steiner vertex is of degree 3, so edge-disjoint S-trees use disjoint Steiner vertices, and hence they correspond to edge-disjoint connected sub-hypergraphs of H. Also, using the fact that each Steiner vertex is of degree 3, it can be proved that if S is $3k$-edge-connected in G, then the hypergraph H constructed is k-partition-connected. Therefore, by applying Theorem 3.2.1, the following result on the STEINER TREE PACKING problem is obtained as a corollary.

Theorem 3.2.2 *[34] If G has no edge between Steiner vertices and S is $3k$-edge-connected in G, then G has k edge-disjoint S-trees.*

In [60], Theorem 3.2.2 is used to prove Theorem 3.1.2 with a slightly bigger constant 26. It was pointed out by Oleg Pikhurko [79] that using Theorem 3.2.1 instead would yield an improvement.

Steiner Tree Packing in General Graphs

Given an instance of the STEINER TREE PACKING problem, our method is to reduce the general case to the seemingly restrictive case of Theorem 3.2.2. The key observation is that Theorem 3.1.2 holds with a rich combinatorial property, which we call *the extension property*. The extension property roughly (formally defined in Section 3.3.3) says that for any edge-partition of the edges incident to a "small" degree vertex, the edge-partition can be extended to edge-disjoint S-trees such that each class in the edge-partition is contained in one S-tree.

The proof is basically divided into two steps. Given a graph G with l edges between Steiner vertices, we search for a minimum S-cut in G with a Steiner edge, and decompose G through the cut, resulting in two graphs G_1 and G_2 with a total of at most $l-1$ edges between Steiner vertices. A cut decomposition lemma (Lemma 3.4.4) shows that if Theorem 3.1.2 holds in both G_1 and G_2 with the extension property, then we can always "piece together" the solutions in G_1 and G_2 so that Theorem 3.1.2 also holds in G with the extension property. Therefore, by applying the cut decomposition step recursively, we reduce an instance with l edges between Steiner vertices to at most $l+1$ instances with no edges between Steiner vertices. By the cut decomposition lemma, if all those $l+1$ graphs (without an edge between Steiner vertices) satisfy Theorem 3.1.2 with the extension property, then G satisfies Theorem 3.1.2 (with the extension property) by "piecing" their solutions together. This key step removes the difficulty of having edges between Steiner vertices, and gives new insight into the core of the problem. It should be mentioned that the Steiner Tree Packing problem remains APX-hard when there are no edges between Steiner vertices.

The second step, of course, is to prove that Theorem 3.1.2 does indeed hold with the extension property when there are no edges between Steiner vertices. By using Mader's splitting lemma (see Section 3.5), we can assume that every Steiner vertex is of degree 3. With a sufficiently high connectivity assumption ($24k$ in Theorem 3.1.2), we can use Theorem 3.2.2 to show that the extension property holds for any graph with no edges between Steiner vertices and with every Steiner vertex of degree 3. This step is more involved, and will appear in Section 3.7. Intuitively, it says that if S is highly edge-connected in G, then any edge-partition of the edges incident to a "small" degree vertex can be extended to edge disjoint S-trees.

Steiner Forest Packing in Eulerian Graphs

In [12], Chekuri and Shepherd consider the Steiner Forest Packing problem when G is Eulerian. By assuming G is Eulerian, as we shall see in Section 3.5, we can assume that G has no Steiner vertices, i.e. $V(G) = S_1 \cup \ldots \cup S_t$. Then they find a "core" C which contains at least one group and show that $G[C]$ has k edge-disjoint spanning trees $T_1, \ldots, T_k$, by using Tutte [85] and Nash-Williams' [75] result on spanning tree packing. Now they contract C in G and obtain a new graph G^*. Note that C contains a group and thus has at least two vertices, so G^* has fewer vertices than G. By induction, G^* has k edge-disjoint Steiner forests $F_1, F_2, \ldots, F_k$. Now, as each tree is *spanning* in C, $F_1 \cup T_1, \ldots, F_k \cup T_k$ are the desired k edge-disjoint Steiner forests in G. The base case in their proof is the Tutte and Nash-Williams result on spanning tree packing.

Steiner Forest Packing in General Graphs

If G is non-Eulerian, we are unable to assume that $V(G) = S_1 \cup \ldots \cup S_t$, and the situation is much more complicated. Here is the main difficulty: Even if we assume the existence of a "core" C of G so that $G[C]$ has k edge-disjoint Steiner trees $T_1, \ldots, T_k$ connecting the terminal groups inside C, and also the existence of k edge-disjoint Steiner forests $F_1, \ldots, F_k$ of G^* as constructed above, we cannot guarantee that a terminal group which is connected in F_i in G^* is still connected in $F_i \cup T_i$ in G since T_i does not necessarily span C.

Here the extension property introduced earlier comes into the picture. Roughly, we show that there are k edge-disjoint Steiner trees in $G[C]$ that "extend" $F_1, \ldots, F_k$ so that $F_1 \cup T_1, \ldots, F_k \cup T_k$ are actually k edge-disjoint Steiner forests. Unlike the situation in the Steiner Tree Packing problem, however, we also need to prove structural properties on $F_1, \ldots, F_k$ in order for them to be extended (not every $F_1, \ldots, F_k$ can be extended). This requires us to revise and generalize the extension property for Steiner Tree Packing. In particular we need to add an additional requirement to the extension theorem, which causes the constant in Theorem 3.1.1 (i.e. approximation ratio) to be slightly bigger than that in Theorem 3.1.2.

3.3 The Setup

3.3.1 Notation and Definitions

We repeat some definitions for the sake of having them all available in one subsection. We have an undirected multigraph G and a set $\mathcal{S} := \{S_1, \ldots, S_t\}$ of disjoint subsets of vertices of G. Let $S^* := S_1 \cup S_2 \cup \cdots \cup S_t$. Each vertex in S^* is called a *terminal vertex*, and each vertex in $V(G) - S^*$ is called a *Steiner vertex*. We call each S_i a *terminal group*, note that we can assume that each group is of size at least 2. A subgraph H of G is a *$\mathcal{S}$-subgraph* if each S_i is connected in H for $1 \leq i \leq t$; a subgraph H is a *double $\mathcal{S}$-subgraph* of G if H is a $\mathcal{S}$-subgraph of G and every vertex in S^* is of degree at least 2 in H. Likewise, given $S \subseteq V(G)$, we say a subgraph H of G is a *S-subgraph* if H is connected and $S \subseteq V(H)$. And a subgraph H is a *double S-subgraph* of G if H is a S-subgraph of G and every vertex in S is of degree at least 2 in H. Furthermore, we say a set of vertices S is k-edge-connected in G if for every pair of vertices $a, b \in S$ there are k edge-disjoint paths between a and b in G. We say a subgraph H *spans* a subset of vertices U if $U \subseteq V(H)$.

For a subset of vertices X, $\delta_G(X)$ denotes the set of edges in G with one endpoint in X and the other endpoint in $V(G) - X$. Notice that if we remove the edges in $\delta_G(X)$ from G, then X is disconnected from $V(G) - X$. We also call a set of edges Y an *edge-cut* if $G - Y$ is disconnected. A subset of vertices X is a *group separating set* if

1. $S^* \cap X \neq \emptyset$ and $S^* \cap (V(G) - X) \neq \emptyset$;

2. for each S_i, either $S_i \subseteq X$ or $S_i \subseteq V(G) - X$.

The following is an important notion mentioned in the previous section: A *core* C is a group separating set with $d_G(C) \leq Qk$ and $|C|$ minimal. Let $R \subseteq V(G) - S^*$ be a specified subset of Steiner vertices. A subset of vertices X is an *R-isolating set* if

1. $R \cap X \neq \emptyset$;

2. $S^* \cap X = \emptyset$.

Given a vertex v, we denote by $\delta(v)$ the set of edges with an endpoint in v. $\mathcal{P}_k(v) := \{\delta_1(v), \ldots, \delta_k(v)\}$ is an edge-subpartition if

1. $\delta_1(v) \cup \delta_2(v) \cup \ldots \cup \delta_k(v) \subseteq \delta(v)$;

2. $\delta_i(v) \cap \delta_j(v) = \emptyset$ for $i \neq j$.

Also, $\mathcal{P}_k(v)$ is a *balanced edge-partition* of $\delta(v)$ if

1. $\delta_1(v) \cup \delta_2(v) \cup \ldots \cup \delta_k(v) = \delta(v)$;

2. $\delta_i(v) \cap \delta_j(v) = \emptyset$ for $i \neq j$;

3. $|\delta_i(v)| \geq 2$ for $1 \leq i \leq k$.

Furthermore, $\mathcal{P}_k(v)$ is a *balanced edge-subpartition* if it can be extended to a balanced edge-partition. More formally, $\mathcal{P}_k(v)$ is a balanced edge-subpartition of $\delta(v)$ if there exists a balanced edge-partition $\mathcal{P}'_k(v) := \{\delta'_1(v), \ldots, \delta'_k(v)\}$ of $\delta(v)$ so that $\delta_i(v) \subseteq \delta'_i(v)$ for $1 \leq i \leq k$. Note that $\delta_i(v)$ of a balanced edge-subpartition could be an empty set. Equivalently, $\mathcal{P}_k(v)$ is a balanced edge-subpartition of $\delta(v)$ if the number of edges of $\delta(v)$ not appearing in $\mathcal{P}_k(v)$ is at least the number of $\delta_i(v)$ with $|\delta_i(v)| = 1$ plus two times the number of $\delta_i(v)$ with $|\delta_i(v)| = 0$.

The open neighbourhood of a vertex v in G is denoted by $N_G(v)$. We use $N_{\delta_i}(u)$ to denote the set of neighbours joined to u by $\delta_i(u)$. Given k edge-disjoint subgraphs $\{H_1, \ldots, H_k\}$ of G, a vertex v is *balanced* with respect to $\{H_1, \ldots, H_k\}$ (or just balanced if $\{H_1, \ldots, H_k\}$ is clear from the context) if $\mathcal{P}_k(v) := \{H_1 \cap \delta(v), \ldots, H_k \cap \delta(v)\}$ is a balanced edge-subpartition. The following are two situations where a vertex v is balanced with respect to k edge-disjoint subgraphs $\{H_1, \ldots, H_k\}$; notice that the proofs are trivial, we list them here for future reference.

Proposition 3.3.1 *Let v be a vertex and $\{H_1, \ldots, H_k\}$ be k edge-disjoint subgraphs. Then v is balanced with respect to $\{H_1, \ldots, H_k\}$ if:*

1. *there are $2k$ edges in $\delta(v)$ not used in any of $\{H_1, \ldots, H_k\}$;*

2. *v is of degree at least 2 in each H_i.*

Finally, given G, $\mathcal{S} := \{S_1, \ldots, S_t\}$ and $R \subseteq V(G) - S^*$, $\{H_1, \ldots, H_k\}$ are k edge-disjoint $\mathcal{S}$-subgraphs that *balance* $S^* \cup R$ if every vertex in $S^* \cup R$ is balanced with respect to $\{H_1, \ldots, H_k\}$.

3.3.2 Main Theorem

Now we are ready to state the main theorem in this chapter. This theorem is stronger than Theorem 3.1.1; the stronger statement will help us to reduce the STEINER FOREST PACKING problem to a strengthened version of the STEINER TREE PACKING problem in Section 3.6.2. Recall that in this chapter $Q = 30$.

Theorem 3.3.2 (THE MAIN THEOREM)
Given G, $\mathcal{S} := \{S_1, \ldots, S_t\}$, and $R \subseteq V(G) - S^*$. If each S_i is Qk-edge-connected in G and each vertex in R is of degree at least Qk, then there are k edge-disjoint $\mathcal{S}$-subgraphs that balance $S^* \cup R$.

Note that Theorem 3.3.2 trivially implies Theorem 3.1.1, because every $\mathcal{S}$-subgraph contains an $\mathcal{S}$-forest. The vertices in R are not needed for the STEINER TREE PACKING problem, but are very important in the proof of the STEINER FOREST PACKING problem. When we apply the decomposition procedure as mentioned in the outline, vertices in R serve as intermediate vertices to combine the solutions (i.e. edge-disjoint subgraphs) in the smaller graphs. Requiring R to be balanced allows the solutions to be extended, and thus makes the approach mentioned in the outline work. In the remainder of this section, we develop the techniques to reduce this main theorem to a strengthened version of the STEINER TREE PACKING problem.

3.3.3 The Extension Theorem

The extension property defined below is crucial in applying a divide-and-conquer strategy to decompose the original problem instance to smaller instances with nice structures. Note that the extension property is defined only for the STEINER TREE PACKING problem, so we shall use S instead of S^* to denote the set of all terminal vertices.

Definition 3.3.3 (THE EXTENSION PROPERTY)
Given G, and $S, R \subseteq V(G)$ with $S \cap R = \emptyset$, and an edge-subpartition $\mathcal{P}_k(v) := \{\delta_1(v), \ldots, \delta_k(v)\}$ of a vertex v, k edge-disjoint S-subgraphs $\{H_1, \ldots, H_k\}$ extend *$\mathcal{P}_k(v)$ if:*

1. *$\delta_i(v) \subseteq E(H_i)$ for each $1 \leq i \leq k$;*

2. *$H_i - v$ is a $(S \cup N_{\delta_i}(v) - v)$-subgraph for each $1 \leq i \leq k$;*

In this case, we also say $\mathcal{P}_k(v)$ is extendible. *If every vertex in $S \cup R - v$ is balanced with respect to $\{H_1, \ldots, H_k\}$ as well, then we say $\mathcal{P}_k(v)$ is* balanced-extendible.

Observe that the above definition makes sense even if $S = R = \emptyset$. Intuitively, the first property of the extension property is to make sure that when we use the cut decomposition operation (to be defined in Section 3.4) on G to obtain G_1 and G_2 and apply the induction hypothesis (Theorem 3.3.4), we can choose the k edge-disjoint subgraphs of G_1 and G_2 to be "compatible", so that they combine to define k edge-disjoint subgraphs in G. The second property of the extension property is to make sure that the set of terminals S is connected in each of the k edge-disjoint subgraphs in G; this will be proved formally in Lemma 3.4.2 using the concept of "natural extension" to be defined in Section 3.4.1.

The following extension theorem, which is at the heart of Theorem 3.1.2, gives sufficient conditions for an edge-subpartition to be balanced-extendible.

Theorem 3.3.4 (THE EXTENSION THEOREM)
Given G, and $S, R \subseteq V(G)$ with $S \cap R = \emptyset$. If S is Qk-edge-connected in G and every

vertex in R is of degree at least Qk, then there are k edge-disjoint S-subgraphs that balance $S \cup R$.

Furthermore, given a vertex v, an edge-subpartition $\mathcal{P}_k(v)$ is balanced-extendible if either of the following happens:

1. *$v \in S$, v is of degree Qk and $\mathcal{P}_k(v)$ is a balanced edge-subpartition.*

2. *$N_G(v) \subseteq S \cup R$, v is of degree at most Qk, $v \notin S \cup R$, and there is no R-isolating set Y with $v \in Y$ and $d_G(Y) \leq Qk$.*

Roughly speaking, part (2) of the extension theorem is to deal with situations where we could not guarantee that $\mathcal{P}_k(v)$ is a balanced edge-subpartition. We remark that when $\mathcal{P}_k(v)$ is not a balanced edge-subpartition, it may not even be extendible, and so we need the extra conditions guaranteed by part (2).

3.4 Technique - Cut Decomposition

Given a multigraph G and a subset of vertices $Y \subset V(G)$, the cut decomposition operation constructs two multigraphs G_1 and G_2 from G as follows. G_1 is obtained from G by contracting $V(G) - Y$ to a single vertex v_1, and keeping all edges from Y to v_1 (this may produce multiple edges). Similarly, G_2 is obtained from G by contracting Y to a single vertex v_2, and keeping all edges from $V(G) - Y$ to v_2. So, $V(G_1) = Y \cup \{v_1\}, \delta_G(Y) \subseteq E(G_1)$ and $V(G_2) = (V(G) - Y) \cup \{v_2\}, \delta_G(Y) \subseteq E(G_2)$ (see Figure 3.1 for an illustration). Notice that for each edge $e \in \delta_G(Y)$, e appears in both G_1 and G_2 (i.e. e in G_1 is incident with v_1 where e in G_2 is incident with v_2). So, given an edge e incident with v_1 in G_1, we refer to the same edge in G_2 incident with v_2 as the *corresponding edge* of e in G_2, and vice versa. The cut decomposition operation will be used several times later. The following are two basic properties of G_1 and G_2:

Proposition 3.4.1 (PROPERTIES OF G_1 AND G_2)

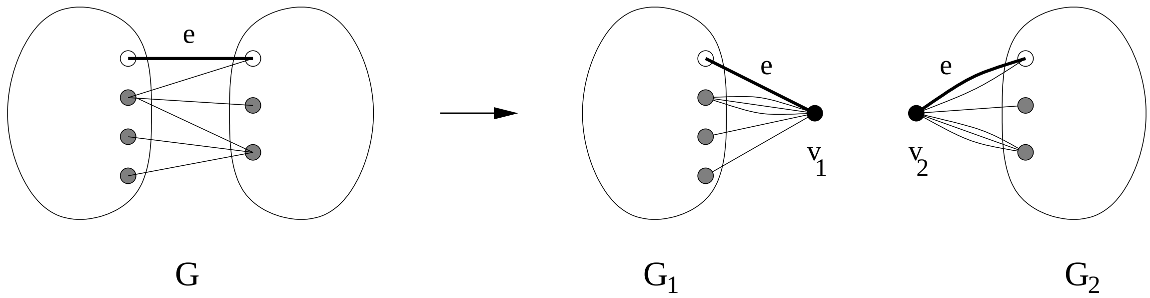

Figure 3.1: The construction of G_1 and G_2 from G.

1. *For each pair of vertices u, v in G_i, the maximum number of edge-disjoint paths between u, v in G_i is at least the maximum number of edge-disjoint paths between u, v in G. In particular, if a set S is Qk-edge-connected in G, then $S \cap V(G_1)$ and $S \cap V(G_2)$ are Qk-edge-connected in G_1 and G_2 respectively.*

2. *The degree of each vertex v in $V(G_i) - \{v_i\}$ is equal to the degree of v in G.*

3.4.1 Natural Extensions

In the following, we show how to use the extension property to combine edge-disjoint subgraphs in two graphs obtained from the cut decomposition operation. Let G_1 and G_2 be two graphs obtained from the cut decomposition operation of G. Suppose $\{H_1^2, \ldots, H_k^2\}$ are k edge-disjoint subgraphs in G_2. Let $\mathcal{P}_k(v_2) := \{H_1^2 \cap \delta(v_2), \ldots, H_k^2 \cap \delta(v_2)\}$ be an edge-subpartition of $\delta(v_2)$; so $\mathcal{P}_k(v_2)$ is an *edge-subpartition induced by* $\{H_1^2, \ldots, H_k^2\}$. By the construction of the cut decomposition operation, there is a one-to-one correspondence between the edges in $\delta_{G_2}(v_2)$ and $\delta_{G_1}(v_1)$. Hence, $\mathcal{P}_k(v_2)$ naturally defines an edge-subpartition $\mathcal{P}_k(v_1) := \{\delta_1(v_1), \ldots, \delta_k(v_1)\}$ of $\delta_{G_1}(v_1)$, where $\delta_i(v_1)$ is defined to be the corresponding edges of $H_i^2 \cap \delta(v_2)$.

Now, suppose $\{H_1^1, \ldots, H_k^1\}$ are k edge-disjoint subgraphs in G_1 that extend $\mathcal{P}_k(v_1)$; note that $H_i^1 - v$ might have no edges if $|N_{\delta_i}(v_1)| \leq 1$. Intuitively, each H_i^1 in G_1 simulates the role of v_2 in G_2 to connect H_i^2. Let H_i^1 be a subgraph of G_1 and H_i^2 be a subgraph of

G_2, we define $H_i^1 \uplus H_i^2$ to be a subgraph in G whose edge set is the union of the edge set of H_i^1 and the edge set of H_i^2. In the following lemma, we prove that if a, b are connected in H_i^2, then a, b are connected in $H_i^1 \uplus H_i^2$. Because of this lemma and the intuition, we call any such $\{H_1^1, \ldots, H_k^1\}$ that extends $\mathcal{P}_k(v_1)$ defined above in G_1 a *natural extension* of $\{H_1^2, \ldots, H_k^2\}$. Or we just say $\{H_1^1, \ldots, H_k^1\}$ in G_1 *naturally extends* $\{H_1^2, \ldots, H_k^2\}$ in G_2. This construction will be used several times later.

Lemma 3.4.2 *Let G_1 and G_2 be two graphs obtained from the cut decomposition operation of G. Suppose $\{H_1^2, \ldots, H_k^2\}$ are k edge-disjoint subgraphs in G_2, and $\{H_1^1, \ldots, H_k^1\}$ in G_1 is a natural extension of $\{H_1^2, \ldots, H_k^2\}$. Then $\{H_1^1 \uplus H_1^2, \ldots, H_k^1 \uplus H_k^2\}$ are k edge-disjoint subgraphs in G. Furthermore, if $a, b \in V(G_1) - v_1$ are connected in H_i^1, or if $a, b \in V(G_2) - v_2$ are connected in H_i^2, then a, b are connected in $H_i^1 \uplus H_i^2$ of G.*

Proof. For each i, let $H_i := H_i^1 \uplus H_i^2$. By assumption, H_i^1 and H_j^1 are edge-disjoint for $i \neq j$ in G_1, and H_i^2 and H_j^2 are edge-disjoint for $i \neq j$ in G_2. Since H_i^1 and H_i^2 use exactly the same edges in the edge-cut, H_i and H_j are edge-disjoint for $i \neq j$ in G.

Suppose $a, b \in V(G_1) - v_1$ are connected in H_i^1. Since $\{H_1^1, \ldots, H_k^1\}$ extend $\mathcal{P}_k(v_1)$, $H_i^1 - v_1$ is connected (by property (ii) of Definition 3.3.3) and so a, b are connected in $H_i^1 - v_1$. So, there is a path from a to b in $H_i^1 - v_1$ and thus in H_i. Therefore, a, b are connected in H_i.

Suppose $a, b \in V(G_2) - v_2$ are connected in H_i^2. If a and b are connected in H_i^2 without using v_2, then there is a path from a to b in $H_i^2 - v_2$ and thus in H_i. Hence they are connected in H_i. So we consider the case that they are connected in H_i^2 using v_2 (see Figure 3.2 for an illustration). Let e_1 and e_2 be the edges incident to v_2 in a path that connects a and b in G_2. Let e_1' and e_2' be the corresponding edges in G_1. By our construction, $e_1', e_2' \in H_i^1$. Let u_1 and u_2 be the endpoints of e_1' and e_2' in G_1, note that u_1 and u_2 need not be distinct as in Figure 3.2. Since $\{H_1^1, \ldots, H_k^1\}$ extend $\mathcal{P}_k(v_1)$, $H_i^1 - v_1$ is connected and so there is a path in $H_i^1 - v_1$ between u_1 and u_2 in G_1. By combining

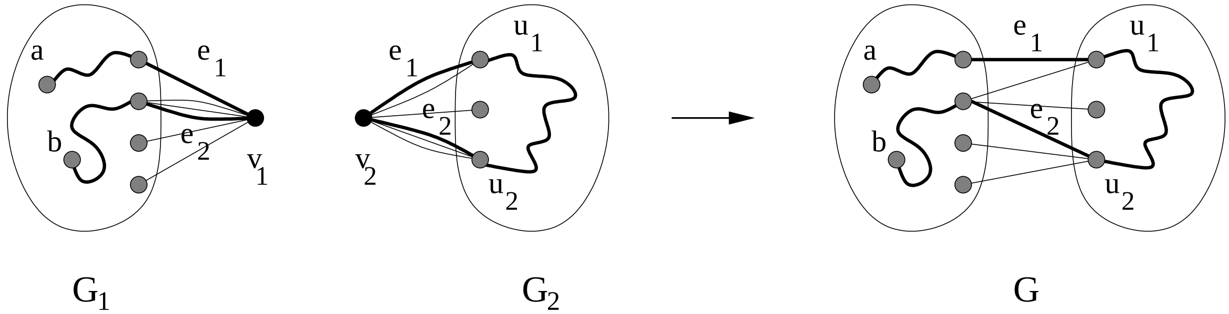

Figure 3.2: If $a, b \in C_2$ is connected in H_i^2 by a path through v_2 in G_2, they are still connected in H_i through G_1.

the edges in the a, v_2-path in H_i^2, the edges in the u_1, u_2-path in $H_i^1 - v_1$ and the edges in the v_2, b-path in H_i^2, we get a path from a to b in H_i. As a result, a and b are connected in H_i. ∎

3.4.2 An Application for Steiner Forest Packing

In this subsection we state the first application which will be crucial in Section 3.6. Given G, $\mathcal{S} := \{S_1, \ldots, S_t\}$, and $R \subseteq V(G) - S^*$. Suppose Y is a group separating set. Let G_1 and G_2 be two graphs obtained from the cut decomposition operation of G so that $V(G_1) = Y \cup \{v_1\}$ and $V(G_2) = (V(G) - Y) \cup \{v_2\}$. Let $\mathcal{S}_1, \mathcal{S}_2$ be the set of terminal groups that are contained in G_1, G_2 respectively. Also, let S_i^* be the union of the vertices in the terminal groups in $\mathcal{S}_i$. Furthermore, let $R_1 := R \cap V(G_1)$ and $R_2 := R \cap V(G_2)$.

Lemma 3.4.3 *Let G_1 and G_2 be defined above. Suppose $\{H_1^2, \ldots, H_k^2\}$ are k edge-disjoint $\mathcal{S}_2$-subgraphs in G_2 that balance $S_2^* \cup R_2$. Suppose further that $\{H_1^1, \ldots, H_k^1\}$ are k edge-disjoint S_1^*-subgraphs in G_1 that balance $S_1^* \cup R_1$ and naturally extend $\{H_1^2, \ldots, H_k^2\}$. Then $\{H_1^1 \uplus H_1^2, \ldots, H_k^1 \uplus H_k^2\}$ are k edge-disjoint $\mathcal{S}$-subgraphs of G that balance $S^* \cup R$.*

Proof. For each i, let $H_i := H_i^1 \uplus H_i^2$. By Lemma 3.4.2, $\{H_1, \ldots, H_k\}$ are k edge-disjoint subgraphs of G. Consider a terminal group S. First assume S is in G_1. Since

S_1^* is connected in H_i^1 for each $1 \leq i \leq k$ and $S_1^* \subseteq V(G_1) - v_1$, Lemma 3.4.2 implies that S_1^* and thus S is connected in H_i. Now assume S is in G_2. Since S is connected in H_i^2 for each i and $S \subseteq V(G_2) - v_2$, by Lemma 3.4.2, S is connected in H_i for each i. Therefore, each $H_i = H_i^1 \uplus H_i^2$ is a $\mathcal{S}$-subgraph. Finally, since every vertex in $S_1^* \cup R_1$ is balanced with respect to $\{H_1^1, \ldots, H_k^1\}$ and every vertex in $S_2^* \cup R_2$ is balanced with respect to $\{H_1^2, \ldots, H_k^2\}$, it follows that every vertex in $S^* \cup R$ is balanced with respect to $\{H_1^1 \uplus H_1^2, \ldots, H_k^1 \uplus H_k^2\}$. This completes the proof. ∎

3.4.3 An Application for Steiner Tree Packing

The second application, which will be useful in Section 3.6, deals with the extension property. Consider G, $S, R \subseteq V(G)$ with $S \cap R = \emptyset$. Suppose Y is a S-separating set. Let G_1 and G_2 be two graphs obtained from the cut decomposition operation on Y. Let $v \in G_2 - v_2$ and $\mathcal{P}_k(v)$ be an edge-subpartition of v. Let $S_1 := S \cap Y, S_2 := S \cap (V(G) - Y), R_1 := R \cap Y$, and $R_2 := R \cap (V(G) - Y)$.

As we mentioned after the main theorem (Theorem 3.3.2), R is not needed for the STEINER TREE PACKING problem. However, in the proof of the main theorem (Section 3.6), we shall reduce the STEINER FOREST PACKING problem to a stronger version of the STEINER TREE PACKING problem, and there R is essential.

In the following lemma, we need $(S_i \cup v_i)$-subgraphs instead of just S_i-subgraphs; this is to make sure that S is connected in the resulting subgraphs $\{H_1^1 \uplus H_1^2, \ldots, H_k^1 \uplus H_k^2\}$. It is instructive to compare this lemma with the previous lemma (Lemma 3.4.3), where v_1 and v_2 are not involved. This is because in the previous lemma each terminal group is contained in either G_1 or G_2, while the terminal group S in the current lemma is split over G_1 and G_2.

Lemma 3.4.4 *Let G_1 and G_2 be defined as above. Suppose $\{H_1^2, \ldots, H_k^2\}$ are k edge-disjoint $(S_2 \cup v_2)$-subgraphs in G_2 that balance $S_2 \cup R_2$ and extend $\mathcal{P}_k(v)$. Suppose $\{H_1^1, \ldots, H_k^1\}$ are k edge-disjoint $(S_1 \cup v_1)$-subgraphs in G_1 that balance $S_1 \cup R_1$ and* *naturally extend* $\{H_1^2, \ldots, H_k^2\}$*. Then* $\{H_1^1 \uplus H_1^2, \ldots, H_k^1 \uplus H_k^2\}$ *are* k *edge-disjoint* S*-subgraphs in* G *that balance* $S \cup R$ *and extend* $\mathcal{P}_k(v)$*.*

Proof. For each i, let $H_i := H_i^1 \uplus H_i^2$. By Lemma 3.4.2, $\{H_1, \ldots, H_k\}$ are k edge-disjoint subgraphs. Let $\mathcal{P}_k(v) := \{\delta_1(v), \ldots, \delta_k(v)\}$. We first verify that H_i spans $S \cup N_{\delta_i}(v)$. Consider $w \in S \cup N_{\delta_i}(v)$. Suppose $w \in G_2$. Since $\{H_1^2, \ldots, H_k^2\}$ are $(S_2 \cup v_2)$-subgraphs that extend $\mathcal{P}_k(v)$, we have $w \in H_i^2$ and thus $w \in H_i$. Suppose $w \in S_1$. Since $\{H_1^1, \ldots, H_k^1\}$ are $(S_1 \cup v_1)$-subgraphs, we have $w \in H_i^1$ and thus $w \in H_i$. The only case left is when $w \in N_{\delta_i}(v) \cap V(G_1)$. Then $vv_2 \in H_i^2$ and $v_1 w \in H_i^1$. Let $\mathcal{P}_k(v_1) = \{\delta_1(v_1), \ldots, \delta_k(v_1)\}$. Notice that $v_1 w \in \delta_i(v_1)$. Since $\{H_1^1, \ldots, H_k^1\}$ are S_1-subgraphs that extend $\mathcal{P}_k(v_1)$, we have $w \in H_i^1$ and thus $w \in H_i$. Hence H_i spans $S \cup N_{\delta_i}(v)$.

Now, we show that $H_i - v$ is a $(S \cup N_{\delta_i}(v) - v)$-subgraph of G. Note that $\{H_1^1, \ldots, H_k^1\}$ are k edge-disjoint connected subgraphs that naturally extend $\{H_1^2, \ldots, H_k^2\}$, and hence naturally extend $\{H_1^2 - v, \ldots, H_k^2 - v\}$. Observe also that since $\{H_1^2, \ldots, H_k^2\}$ extend $\mathcal{P}_k(v)$, $\{H_1^2 - v, \ldots, H_k^2 - v\}$ are connected. For any $a, b \in V(H_i) - v$, we consider the following three cases. If $a, b \in V(H_i^2) - v_2 - v$, since a, b are connected in $H_i^2 - v$, Lemma 3.4.2 implies that a, b are connected in $H_i - v$. If $a, b \in V(H_i^1) - v_1$, since a, b are connected in H_i^1, Lemma 3.4.2 implies that a, b are connected in $H_i - v$. Now we consider the case that $a \in V(H_i^1) - v_1$ and $b \in V(H_i^2) - v_2 - v$. Since $H_i^2 - v$ is connected and spans v_2, there is a path from b to v_2 in $H_i^2 - v$. Since H_i^1 is connected and spans v_1, there is a path from v_1 to a in G_1. Therefore, there is an a, b-path in $H_i - v$ by combining the edges in the b, v_2-path and the edges in the v_1, a-path. So, each $H_i - v$ is a $(S \cup N_{\delta_i}(v) - v)$-subgraph in G. Therefore, $\{H_1, \ldots, H_k\}$ are k edge-disjoint S-subgraphs in G that extend $\mathcal{P}_k(v)$.

Finally, since $\{H_1^1, \ldots, H_k^1\}$ balance $S_1 \cup R_1$ and $\{H_1^2, \ldots, H_k^2\}$ balance $S_2 \cup R_2$, this implies that $\{H_1, \ldots, H_k\}$ balance $S \cup R$ as well. This completes the proof. ∎

3.5 Tool - Mader's Splitting-off Lemma

A basic tool in many edge-connectivity problems is Mader's splitting-off lemma (see Section 2.2.3 for more details). Let G be a graph, $e_1 = xy$, $e_2 = xz$ be two edges, $y \neq z$. The operation of obtaining G' from G by deleting e_1 and e_2 and then adding exactly one new edge between y and z (multiple edges between y and z may be produced) is said to be splitting-off at x. This splitting-off operation at x is said to be suitable, if the number of edge-disjoint a, b-paths in G' is at least the number of edge-disjoint a, b-paths in G for every pair $a, b \in V(G) - x$. Mader's splitting-off lemma provides a sufficient condition for the existence of a suitable splitting at a certain vertex x:

Lemma 3.5.1 (MADER'S SPLITTING-OFF LEMMA) *[71] Let x be a vertex of a graph G. Suppose that x is not a cut vertex and that x is incident with at least 4 edges and adjacent to at least 2 vertices. Then there exists a suitable splitting-off operation at x in G.*

We remark that Lemma 3.5.1 is equivalent to Theorem 2.2.11, but the form in the above lemma is more convenient for our purpose.

Let $\mathcal{G}$ be a counterexample of Theorem 3.3.2 or Theorem 3.3.4 with the minimum number of edges and then the minimum number of vertices. The following lemmas restrict the structures of $\mathcal{G}$ by applying Mader's splitting-off lemma.

Lemma 3.5.2 *There is no cut vertex in $\mathcal{G}$.*

Proof. The first case is when $\mathcal{G}$ is a minimal counterexample to Theorem 3.3.2. Suppose w is a cut vertex in $\mathcal{G}$. Let $\{C_1, \ldots, C_m\}$ be the connected components of $\mathcal{G} - w$ where $m \geq 2$. We construct $G_j = \mathcal{G}[C_j \cup \{w\}]$ for $1 \leq j \leq m$. We say S_i is *isolated* if $S_i \subseteq V(G_j)$ for some j; otherwise we say S_i is *separated*. For each group S_i, let S_i^j be $S_i \cap V(G_j)$. For each separated group S_i, we add w to S_i^j for $1 \leq j \leq m$.

We claim that each S_i^j is Qk-edge-connected in G_j for $1 \leq j \leq m$. Notice that it is possible that $S_i^j = \emptyset$ for an isolated S_i, or $S_i^j = \{w\}$ for a separated S_i, we will consider these degenerate sets to be Qk-edge-connected. So, when S_i is isolated, each S_i^j is Qk-edge-connected since S_i is Qk-edge-connected in $\mathcal{G}$. Now we consider the case when S_i is separated. For each $a \in S_i^j$ such that $a \neq w$, since S_i is separated, there is a vertex $b \neq w$ that is in S_i^l for some $l \neq j$. By the assumption that S_i is Qk-edge-connected in $\mathcal{G}$, there are Qk edge-disjoint paths between a and b in $\mathcal{G}$. Since w is a cut vertex, all the paths from a to b must pass through w and thus there are Qk edge-disjoint paths between a and w in G_j. As a is an arbitrary vertex of $S_i^j - w$, this proves our claim that S_i^j is Qk-edge-connected in G_j for each j.

Note that each G_j has fewer edges than $\mathcal{G}$. So, by the minimality of $\mathcal{G}$, Theorem 3.3.2 holds in G_j for $1 \leq j \leq m$. That is, there are k edge-disjoint Steiner forests $\{F_1^j, \ldots, F_k^j\}$ in G_j so that S_i^j is connected in each such forest for each terminal group i. In particular, a forest F_l^j spans S_i^j for all i. Now, by setting $F_l = F_l^1 \cup F_l^2 \cup \cdots \cup F_l^m$, F_l is a subgraph that spans $S_i = \cup S_i^j$ for all i in $\mathcal{G}$. For an isolated S_i, S_i is obviously connected in F_l. For a separated S_i, S_i is connected in F_i because each $S_i^j = (S_i \cap V(G_j)) \cup w$ is connected in F_i^j. So, each F_l is a $\mathcal{S}$-subgraph in $\mathcal{G}$. And thus $\mathcal{G}$ has k edge-disjoint $\mathcal{S}$-forests, a contradiction. Therefore, by the minimality of $\mathcal{G}$, $\mathcal{G}$ has no cut vertex.

In the case when $\mathcal{G}$ is a minimal counterexample to Theorem 3.3.4, the arguments are very similar. We need to consider two cases which correspond to whether the cut vertex is the vertex to be extended. Both cases can be handled in exactly the same manner, only that we have to check the extension property also. But it is totally straightforward and we omit the details. ∎

Lemma 3.5.3 *Every vertex in $V(\mathcal{G}) - S^* - R$ is incident with exactly three edges and adjacent to exactly three vertices.*

Proof. Let $\mathcal{G}$ be a minimal counterexample of Theorem 3.3.2, and w be a vertex in $V(\mathcal{G}) - S^* - R$. Suppose w is adjacent to only one vertex u. Since each S_i is Qk-edge-connected in $\mathcal{G}$, each S_i is still Qk-edge-connected in $\mathcal{G} - w$. By the minimality of $\mathcal{G}$, Theorem 3.3.2 holds in $\mathcal{G} - w$ and hence in $\mathcal{G}$, a contradiction. So we can assume that w is adjacent to at least two vertices.

Suppose w is incident with only two edges, by the previous argument, w is adjacent to two vertices $\{y, z\}$. Since each S_i is Qk-edge-connected in $\mathcal{G}$ and $w \notin S^* \cup R$, each S_i is Qk-edge-connected in $\mathcal{G} - w + yz$ which has one fewer edge than $\mathcal{G}$. By the minimality of $\mathcal{G}$, Theorem 3.3.2 holds in $\mathcal{G} - w + yz$. For any k edge-disjoint $\mathcal{S}$-subgraphs $\{H_1, \ldots, H_k\}$ of $\mathcal{G} - w + yz$ that balance $S^* \cup R$, if yz is in H_i, we can construct H_i' from H_i by replacing yz with $\{wy, wz\}$. So $\mathcal{G}$ also has k edge-disjoint $\mathcal{S}$-subgraphs that balance $S^* \cup R$, a contradiction. Hence we can further assume that w is incident with more than two edges.

Suppose w is incident with at least four edges. By the previous argument, w is adjacent to at least two vertices. And by Lemma 3.5.2, w is not a cut vertex. Therefore, by Lemma 3.5.1, there exists a suitable splitting-off of $\mathcal{G}$ at w, say the resulting graph is G^*. Since each S_i is Qk-edge-connected in $\mathcal{G}$ and the splitting-off operation is suitable, each S_i is Qk-edge-connected in G^* which has one fewer edge than $\mathcal{G}$. By the minimality of $\mathcal{G}$, Theorem 3.3.2 holds in G^*. By a similar argument as in the previous paragraph, it follows that Theorem 3.3.2 also holds in $\mathcal{G}$; a contradiction. Therefore, the only possibility left is when w is incident with exactly three edges.

Suppose w is incident with three edges but adjacent to only two vertices $\{y, z\}$ so that there are two edges e_1, e_2 between w and y. Since $w \notin S^* \cup R$ and w is incident with exactly three edges and adjacent only to $\{y, z\}$, e_1 and e_2 cannot be in two edge-disjoint paths connecting two terminal vertices. Note also that any path that uses e_1 can use e_2 instead. Since each S_i is Qk-edge-connected in $\mathcal{G}$, it follows that S is Qk-edge-connected in $\mathcal{G} - e_1$ which has one fewer edge than $\mathcal{G}$. By the minimality of $\mathcal{G}$, Theorem 3.1.1 holds in $\mathcal{G} - e_1$. Hence Theorem 3.1.1 also holds in $\mathcal{G}$, a contradiction. As a result, every vertex $w \in V(G) - S^* - R$ of $\mathcal{G}$ must be incident with exactly 3 edges and adjacent to exactly

3 vertices; this completes the proof when $\mathcal{G}$ is a counterexample of Theorem 3.3.2.

Now consider the case when $\mathcal{G}$ is a counterexample of Theorem 3.3.4. All the cases are handled in exactly the same manner as above, only that we have to check the extension property also. The only case that needs to be mentioned is when w is adjacent to the vertex v to be extended, and a suitable splitting-off operation is applied by replacing wv and wu with uv for some $u \neq v$. In this case the edges incident to v have been changed, but there is a one-to-one correspondence between the new edges and the old edges. And the extension property holds in the new graph if and only if the extension property holds in the original graph. The remaining details are omitted. ∎

3.6 Proof of the Main Theorem

The strategy we will use is similar to the strategy used in packing Steiner forests in Eulerian graphs as outlined in Section 3.2 - the goal is to reduce the STEINER FOREST PACKING problem to the STEINER TREE PACKING problem. First, we shall find a core C and prove some structural results about it. Recall that a core C is a group separating set with $d_G(C) \leq Qk$ and $|C|$ minimal. We apply the cut decomposition operation on $\mathcal{G}$ and C to obtain two graphs G_1 and G_2, where $V(G_1) = C \cup \{v_1\}$. As usual, we set $R_1 := R \cap V(G_1)$ and $S_1^* := S^* \cap V(G_1)$ (recall that $S^* := S_1 \cup \ldots \cup S_t$). Then we shall show that the main theorem (Theorem 3.3.2) follows from induction and the extension theorem (Theorem 3.3.4).

3.6.1 Group Separating Cut and Core

Let C be a core of $\mathcal{G}$, where $\mathcal{G}$ is a minimal counterexample of Theorem 3.3.2. The following lemma provides structural properties of a core of $\mathcal{G}$. This will be used to reduce the main theorem to the extension theorem.

Lemma 3.6.1 *Let G_1 be defined as above. Then S_1^* is Qk-edge-connected in G_1, and v_1 is of degree at most Qk. Furthermore, at least one of the following must be true:*

1. *$S_1^* \cup \{v_1\}$ is Qk-edge-connected in G_1.*

2. *$N_{G_1}(v_1) \subseteq S_1^* \cup R_1$, and there is no R_1-isolating set Y with $v_1 \in Y$ and $d_{G_1}(Y) \leq Qk$.*

Proof. First we prove that S_1^* is Qk-edge-connected in G_1; note that S_1^* maybe the union of several terminal groups. Suppose not, then there exists a set $Y \subseteq V(G_1)$ such that $Y \cap S_1^* \neq \emptyset$, $(V(G_1) - Y) \cap S_1^* \neq \emptyset$, and $d_{G_1}(Y) < Qk$. Without loss of generality, we can assume that $v_1 \notin Y$. Since $d_{G_1}(Y) < Qk$, each terminal group in G_1 is either contained in Y or disjoint from Y; otherwise this contradicts the assumption that each group is Qk-edge-connected in $\mathcal{G}$ and thus in G_1. So, Y is a group separating set in $\mathcal{G}$ with $d_{\mathcal{G}}(Y) < Qk$ and $Y \subset V(G_1) - v_1 = C$ (as $(V(G_1) - Y) \cap S_1^* \neq \emptyset$). This contradicts the fact that C is a core, and so S_1^* is Qk-edge-connected in G_1. The fact that v_1 is of degree at most Qk follows from our construction and the fact that $d_{G_1}(Y) \leq Qk$.

Next, we prove that if $N_{G_1}(v_1) \not\subseteq S_1^* \cup R_1$, then $S_1^* \cup \{v_1\}$ is Qk-edge-connected in G_1. First, we show that v_1 must be of degree Qk. Suppose, by way of contradiction, that v_1 is a vertex of degree less than Qk. Let $w \in N_{G_1}(v_1)$ be a vertex in $V(G_1) - S_1^* - R_1$; recall that this vertex exists by our assumption. Recall also that by our construction of G_1, $V(G_1) = C \cup \{v_1\}$. By Lemma 3.5.3, w is of degree 3. Also $v_1 w \in E(G_1)$ and v_1 is of degree less than Qk. So we have $d_{G_1}(C - w) = d_{\mathcal{G}}(C - w) \leq Qk$. Since $C - w$ is also a group separating set, this contradicts the fact that C is a core. So, v_1 is of degree exactly Qk.

Suppose, by way of contradiction, that $S_1^* \cup \{v_1\}$ is not Qk-edge-connected in G_1. Then there exists $Y \subseteq V(G_1)$ such that $v_1 \notin Y$, $Y \cap S_1^* \neq \emptyset$ and $d_{G_1}(Y) < Qk$. Since we have proved that S_1^* is Qk-edge-connected in G_1, we must have $S_1^* \subseteq Y$, so Y is also a group separating set. Also, since v_1 is of degree Qk but $d_{G_1}(Y) < Qk$, we have $|V(G_1) - Y| \geq 2$ and hence $Y \subset C$ (recall that $V(G_1) = C + \{v_1\}$). Thus Y contradicts the fact that C

is a core. Therefore, if $N_{G_1}(v_1) \nsubseteq S_1^* \cup R_1$, then $S_1^* \cup \{v_1\}$ is Qk-edge-connected in G_1.

Finally we prove that there is no R_1-isolating set Y with $v_1 \in Y$ and $d_{G_1}(Y) \leq Qk$. Suppose, by way of contradiction, that such a Y exists. Since $v_1 \in Y$, $V(G_1) - Y$ is a group separating set in G. Also, $|Y| \geq 2$ as $v_1 \in Y$ and Y contains at least one vertex in R_1. So $V(G_1) - Y$ is a group separating set with $V(G_1) - Y \subset C$. This contradicts the fact that C is a core. Therefore, there is no R_1-isolating set Y with $v_1 \in Y$ and $d_{G_1}(Y) \leq Qk$. We remark that this statement always holds, not just for graphs in Case 2. ∎

3.6.2 The Reduction

Now, assuming Theorem 3.3.4, we are ready to prove Theorem 3.3.2 which we restate below. This is where we reduce the Steiner Forest Packing problem to a strengthened version of the Steiner Tree Packing problem (i.e., the Steiner Tree Packing problem with the additional requirements as stated in Theorem 3.3.4). As usual, we will assume $\mathcal{G}$ is a minimal counterexample of Theorem 3.3.2 and prove that it cannot exist. Recall the statement of Theorem 3.3.2:

> (The Main Theorem) Given G, $\mathcal{S} := \{S_1, \ldots, S_t\}$, and $R \subseteq V(G) - S^*$. If each S_i is Qk-edge-connected in G and each vertex in R is of degree at least Qk, then there are k edge-disjoint $\mathcal{S}$-subgraphs that balance $S^* \cup R$.

Proof. First suppose that S^* is Qk-edge-connected in $\mathcal{G}$. Then applying Theorem 3.3.4 with $S := S^*$ implies the theorem in this case, a contradiction. Note that in this case we do not even need the extension property.

So assume S^* is not Qk-edge-connected in $\mathcal{G}$. Then there exists a pair of vertices $a, b \in S^*$ with $\lambda(a, b) < Qk$. By Menger's theorem (Theorem 2.2.3), there exists a set X with $a \in X$, $b \notin X$ and $d(X) < Qk$. Since a terminal group is Qk-edge-connected, a terminal group must be either contained in X or disjoint from X. Therefore X is a group separating set with $d(X) < Qk$, and hence a core C exists. We apply the cut decomposition operation on $\mathcal{G}$ and C to obtain two graphs G_1 and G_2, where $V(G_1) = C \cup \{v_1\}$. Let $S_1^* := S^* \cap V(G_1), S_2^* := S^* \cap V(G_2), R_1 := R \cap V(G_1), R_2 := R \cap V(G_2)$, and $\mathcal{S}_1$ and $\mathcal{S}_2$ be the groups contained in S_1^* and S_2^* respectively. By the first property of the cut decomposition operation (Proposition 3.4.1), each terminal group in $\mathcal{S}_1$ and $\mathcal{S}_2$ is Qk-edge-connected in G_1 and G_2 respectively. Also, by the second property of the cut decomposition operation (Proposition 3.4.1), vertices in R_1 and R_2 are of degree at least Qk in G_1 and G_2 respectively.

Based on Lemma 3.6.1, there are only two possibilities for the structure of G_1. The first case is that $S_1^* \cup \{v_1\}$ is Qk-edge-connected in G_1, let $S_1' := S_1^* \cup \{v_1\}$ and $R_2' := R_2 \cup \{v_2\}$. Notice that v_1, v_2 are of degree Qk since $S_1^* \cup \{v_1\}$ is Qk-edge-connected in G_1, and so every vertex in R_2' has degree at least Qk in G_2. Note that both G_1 and G_2 are smaller than $\mathcal{G}$ because they both contain at least one terminal group which has at least two terminal vertices. By the minimality of $\mathcal{G}$, there are k edge-disjoint $\mathcal{S}_2$-subgraphs $\{H_1^2, \ldots, H_k^2\}$ in G_2 that balance $S_2^* \cup R_2'$. Since $v_2 \in R_2'$, v_2 is balanced with respect to $\{H_1^2, \ldots, H_k^2\}$. Recall that v_1 is of degree Qk. So, by applying Theorem 3.3.4 (1) on G_1 (with $S := S_1'$, $R := R_1$ and $v := v_1$), there are k edge-disjoint S_1'-subgraphs $\{H_1^1, \ldots, H_k^1\}$ that balance $S_1' \cup R_1$ and naturally extend $\{H_1^2, \ldots, H_k^2\}$. Therefore, by Lemma 3.4.3, we obtain k edge-disjoint $\mathcal{S}$-subgraphs in $\mathcal{G}$ that balance $S^* \cup R$. This contradicts the fact that $\mathcal{G}$ is a counterexample of Theorem 3.3.2 and so the first case of Lemma 3.6.1 cannot happen.

The second case is that $N_{G_1}(v_1) \subseteq S_1^* \cup R_1$, and there is no R_1-isolating set Y with $v \in Y$, and $d_{G_1}(Y) \leq Qk$. The following arguments are similar to the previous case. By the minimality of $\mathcal{G}$, there are k edge-disjoint $\mathcal{S}_2$-subgraphs $\{H_1^2, \ldots, H_k^2\}$ in G_2 that balance $S_2^* \cup R_2$. (Notice that R_2 is used here instead of R_2' because v_2 may have degree strictly less than Qk, and so $\{H_1^2, \ldots, H_k^2\}$ need not balance v_2.) Now, by applying Theorem 3.3.4 (2) on G_1 (with $S := S_1^*$, $R := R_1$ and $v := v_1$), we have k edge-disjoint S_1^*-

subgraphs $\{H_1^1, \ldots, H_k^1\}$ in G_1 that balance $S_1^* \cup R_1$ and naturally extend $\{H_1^2, \ldots, H_k^2\}$. Therefore, by Lemma 3.4.3, we obtain k edge-disjoint $\mathcal{S}$-subgraphs $\{H_1, \ldots, H_k\}$ in $\mathcal{G}$ that balance $S^* \cup R$. So, $\mathcal{G}$ is not a counterexample and this completes the proof. ∎

3.7 The Extension Theorem

In this section, we will prove Theorem 3.3.4 by showing that a minimal counterexample $\mathcal{G}$ of Theorem 3.3.4 does not exist. Recall the statement of Theorem 3.3.4:

> (THE EXTENSION THEOREM)
> Given G, and $S, R \subseteq V(G)$ with $S \cap R = \emptyset$. If S is Qk-edge-connected in G and every vertex in R is of degree at least Qk, then there are k edge-disjoint S-subgraphs that balance $S \cup R$. Furthermore, given a vertex v, an edge-subpartition $\mathcal{P}_k(v)$ is balanced-extendible if either of the following happens:
>
> 1. $v \in S$, v is of degree Qk and $\mathcal{P}_k(v)$ is a balanced edge-subpartition.
> 2. $N_G(v) \subseteq S \cup R$, v is of degree at most Qk, and there is no R-isolating set Y with $v \in Y$ and $d_G(Y) \leq Qk$.

Proof Outline

The proof consists of a series of technical lemmas, which might not be easy to follow. This outline intends to give a higher level structure of the proof. We remark that some parts of the outline might not be precise, but they should be informative enough to give a brief understanding of the approach.

As mentioned earlier, Frank, Király and Kriesell proved a hypertree packing theorem (Theorem 3.2.1), which can be applied to construct edge-disjoint Steiner trees in graphs with no edges between two Steiner vertices. This will be the key tool in our proof. In fact, the reason for introducing the extension property is to allow us to prove that a minimal counterexample $\mathcal{G}$ has no edges between two Steiner vertices, which allows us to apply Theorem 3.2.1.

The first step of the proof (Lemma 3.7.1) is to combine the cut decomposition technique and the extension property to prove that $\mathcal{G}$ has no edges between two Steiner vertices. The next important step (Lemma 3.7.2) is to prove that $S \cup R$ is highly edge-connected in $\mathcal{G}$. This allows us, in some cases, to find edge-disjoint $(S \cup R)$-subgraphs (instead of just S-subgraphs) so as to show that the resulting subgraphs balance R. With Lemma 3.7.1 and Lemma 3.7.2, we can then handle the second case of the extension theorem.

The first case of the extension theorem, however, doesn't come as quickly. Let v be the vertex to be extended. It is not so easy to deal with the edge-subpartition $\mathcal{P}_k(v)$ of $\delta(v)$ directly, particularly when $N(v)$ contains Steiner vertices. So instead we consider $G' := \mathcal{G} - v - W$ where W is the set of Steiner vertices adjacent to v. First we show in Lemma 3.7.3 that if $S \cup R - v$ is $6k$-edge-connected in G', then we can apply Theorem 3.2.2 to construct the desired S-subgraphs in the extension theorem. Let Z be the edge cutset of G' whose removal separates $S \cup R - v$, so we have $|Z| < 6k$. We show in Lemma 3.7.4 that $G' - Z$ has exactly two connected components C_1 and C_2.

Then we introduce the concepts of diverging paths and common paths, which will be used to establish the edge-connectivity of the terminal vertices in C_1 and C_2. We remark that the proof at this point is divided into two cases; the first case is when both components contain some vertices in S and the second case is when one component contains no vertices in S. The details therein are quite different, but the general strategy is similar. Let $S_1 := S \cap V(C_1)$, $R_1 := R \cap V(C_1)$, $S_2 := S \cap V(C_2)$ and $R_2 := R \cap V(C_2)$. In Lemma 3.7.7 (respectively Lemma 3.7.14), we use the diverging paths to show that $S_i \cup R_i$ are highly edge-connected in C_i for $i \in \{1, 2\}$. This allows us to use Theorem 3.2.1 to construct edge-disjoint $(S_i \cup R_i)$-subgraphs in C_i. Finally, in Lemma 3.7.10 (respectively Lemma 3.7.15), we use the edge-disjoint $(S_i \cup R_i)$-subgraphs in each component and the fact that $\mathcal{P}_k(v)$ is a balanced edge-subpartition to construct edge-disjoint S-subgraphs in $\mathcal{G}$ that satisfy the requirements of the extension theorem. This shows that $\mathcal{G}$ is not a counterexample, and we are done.

This finishes the proof outline, and we now start with the first step of the proof.

There Is No Edge Between Two Vertices in $V(\mathcal{G}) - S - R$

Lemma 3.7.1 *There is no edge in $\mathcal{G}$ with both endpoints in $V(\mathcal{G}) - S - R$.*

Proof. Suppose, by way of contradiction, that e is such an edge. If S is Qk-edge-connected in $\mathcal{G} - e$, then by the minimality of $\mathcal{G}$, $\mathcal{G} - e$ satisfies Theorem 3.3.4 and hence so does $\mathcal{G}$. So we assume that S is not Qk-edge-connected in $\mathcal{G} - e$. That is, there exists $Y \subseteq V(\mathcal{G})$ and $e \in \delta_{\mathcal{G}}(Y)$ so that $Y \cap S \neq \emptyset$, $(V(\mathcal{G}) - Y) \cap S \neq \emptyset$ and $d_{\mathcal{G}}(Y) = Qk$. Without loss of generality we assume $v \notin Y$, where v is the vertex to be extended. We apply the cut decomposition operation on $\mathcal{G}$ and Y to obtain two graphs G_1 and G_2, so that $V(G_1) = Y \cup \{v_1\}$ and $v \in G_2$. Set $S_1 := S \cap V(G_1)$, $R_1 := R \cap V(G_1)$, $S_2 := S \cap V(G_2)$ and $R_2 := R \cap V(G_2)$.

By the first property of the cut decomposition operation (Proposition 3.4.1), S_1 and S_2 are Qk-edge-connected in G_1 and G_2 respectively. Let $a \in S_1$ and $b \in S_2$. Since S is Qk-edge-connected in $\mathcal{G}$, there are Qk edge-disjoint paths from a to b in $\mathcal{G}$. All such paths must pass through $\delta_{\mathcal{G}}(Y)$, and hence in G_1 and G_2 there are Qk edge-disjoint paths from a to v_1 and from b to v_2 respectively. Therefore, $S_1 \cup v_1$ and $S_2 \cup v_2$ are Qk-edge-connected in G_1 and G_2 respectively. Also, by the second property of the cut decomposition operation (Proposition 3.4.1), vertices in R_1 and R_2 are of degree at least Qk in G_1 and G_2 respectively. Notice that since both G_1 and G_2 contain a vertex in $V(\mathcal{G}) - S - R$ (an endpoint of e) and a vertex in S, both G_1 and G_2 have fewer vertices than $\mathcal{G}$. Hence, by the minimality of $\mathcal{G}$, G_1 and G_2 both satisfy Theorem 3.3.4.

Let $\mathcal{P}_k(v)$ be the edge-subpartition of $\delta(v)$ to be extended. By the minimality of $\mathcal{G}$, there are k edge-disjoint $(S_2 \cup v_2)$-subgraphs $\{H_1^2, \ldots, H_k^2\}$ in G_2 that balance $S_2 \cup R_2 \cup v_2$

and extend $\mathcal{P}_k(v)$. Notice that since the edge-cut $\delta_{\mathcal{G}}(Y)$ is of size exactly Qk, v_1 is of degree exactly Qk. Also, since v_2 is balanced with respect to $\{H_1^2, \ldots, H_k^2\}$, the edge-subpartition induced on v_2 by $\{H_1^2, \ldots, H_k^2\}$ is a balanced edge-subpartition. Therefore, by substituting $S := S_1 \cup v_1$, v_1 satisfies all the requirements of the first case of the extension theorem. So, by the minimality of $\mathcal{G}$, there are k edge-disjoint $(S_1 \cup v_1)$-subgraphs $\{H_1^1, \ldots, H_k^1\}$ in G_1 that balance $S_1 \cup R_1$ and naturally extend $\{H_1^2, \ldots, H_k^2\}$. Now, by applying Lemma 3.4.4, there are k edge-disjoint S-subgraphs in $\mathcal{G}$ that balance $S \cup R$ and extend $\mathcal{P}_k(v)$. This, however, contradicts the assumption that $\mathcal{G}$ is a counterexample to Theorem 3.3.4. Therefore, there is no edge between two vertices in $V(\mathcal{G}) - S - R$. ∎

R-isolating set

Now we shall prove that $S \cup R$ is highly edge-connected in $\mathcal{G}$. This is an important step towards the proof.

Lemma 3.7.2 *$S \cup R$ is $(Q-2)k$-edge-connected in $\mathcal{G}$.*

Proof. Suppose, by way of contradiction, that $S \cup R$ is not $(Q-2)k$-edge-connected in $\mathcal{G}$. Consider an R-isolating set Y with $d_{\mathcal{G}}(Y)$ minimum. Since S is Qk-edge-connected but $S \cup R$ is not $(Q-2)k$-edge-connected, we have $d_{\mathcal{G}}(Y) < (Q-2)k$. Apply the cut decomposition operation on $\mathcal{G}$ and Y to obtain two graphs G_1 and G_2 where $V(G_1) = Y + v_1$. Let $R_1 := R \cap V(G_1)$ and $R_2 := R \cap V(G_2)$. Also, since Y is a R-isolating set, $S \subseteq V(G_2)$. Let v be the vertex to be extended in $\mathcal{G}$. So, we assume that v satisfies one of the requirements of the extension theorem. If v satisfies the first requirement of the extension theorem, then $v \in S$, and so it must be in G_2 because $S \subseteq V(G_2)$. If v satisfies the second condition of the extension theorem, then v also must be in G_2 because there is no R-isolating set with $v \in Y$ and $d_{\mathcal{G}}(Y) \leq Qk$. So, it follows that v must be in G_2. By the properties of the cut decomposition operation, S is Qk-edge-connected in G_2 and every vertex of R_2 is of degree at least Qk in G_2. Note that $R_1 \neq \emptyset$, hence $|Y| \geq 2$ as each vertex in R is of degree at least Qk while $d_{\mathcal{G}}(Y) \leq (Q-2)k < Qk$. As $|Y| \geq 2$, G_2 is smaller than $\mathcal{G}$. By the minimality of $\mathcal{G}$, G_2 has k edge-disjoint S-subgraphs $\{H_1^2, \ldots, H_k^2\}$ that balance $S \cup R_2$ and extend $\mathcal{P}_k(v)$.

Let $l := d_{\mathcal{G}}(Y)$, then $l < (Q-2)k$ by assumption. We claim that each vertex $r \in R_1$ has l edge-disjoint paths to v_1 in G_1. Suppose not, then there exists $Y' \subseteq V(G_1)$ so that $r \in Y'$, $v_1 \notin Y'$ and $d_{G_1}(Y') < l$. Hence, Y' is a R-isolating set in $\mathcal{G}$ with $d_{\mathcal{G}}(Y') < l$, but this contradicts the minimality of $d_{\mathcal{G}}(Y)$. Therefore, each vertex in R_1 has l edge-disjoint paths to v_1. Choose a vertex $r^* \in R_1$ so that the total length of the l edge-disjoint paths $\{P_1, \ldots, P_l\}$ from v_1 to r^* is minimized. This, for example, can be computed by using minimum cost flow. Let $P := \{P_1, \ldots, P_l\}$ and $H := P_1 \cup P_2 \cup \ldots \cup P_l$. We claim that r is of degree at most l in H for each $r \in R_1$. Suppose not, let $r \in R_1$ be of degree greater than l in H. Then $r \neq r^*$, and hence each (v_1, r^*)-path in P either has two edges incident to r or has no edge incident to r. So r is of even degree in H. We assume r is of degree $l+1$ in H, and thus we assume that $l+1$ is an even number. (We remark that the case when r is of degree greater than $l+1$ is easier.) So there are $(l+1)/2$ paths from v_1 to r^* that pass through r, say $\{P_1, \ldots, P_{(l+1)/2}\}$. For $1 \leq i \leq (l+1)/2$, let the subpath of P_i from v_1 to r be P_i', and the subpath of P_i from r^* to r be P_i''. Then we claim that $H - P''_{(l+1)/2}$ contains l edge-disjoint paths from v_1 to r; indeed, $\{P_1', \ldots, P'_{(l+1)/2}, P_{(l+1)/2+1} \cup P_1'', \ldots, P_l \cup P''_{(l-1)/2}\}$ are l edge-disjoint paths from v_1 to r in $H - P''_{(l+1)/2}$. This contradicts the choice of r^* and thus r is of degree at most $l < (Q-2)k$ for each $r \in R_1$ in H. Since every vertex $r \in R_1$ is of degree at least Qk in G_1, there are at least $2k$ edges in $\delta(r)$ not used in $H = P_1 \cup P_2 \cup \ldots \cup P_l$, and so each $r \in R_1$ is balanced with respect to $\{P_1, \ldots, P_l\}$ by Proposition 3.3.1 (1).

Now, we show that the k edge-disjoint S-subgraphs $\{H_1^2, \ldots, H_k^2\}$ in G_2 can be extended to k edge-disjoint S-subgraphs in $\mathcal{G}$, by using the edge-disjoint paths from v_1 to r^* in G_1 constructed in the preceding paragraph. For each H_i^2 in G_2, let $E_i(v_2) := E(H_i^2) \cap \delta(v_2)$ and $E_i(v_1)$ be the corresponding edge set in G_1. Let H_i^1 be the (r^*, v_1)-paths in P of G_1 containing the corresponding edges in $E_i(v_1)$. Then $\{H_1^1, \ldots, H_k^1\}$ balance R_1 as argued in the previous paragraph, and naturally extend $\{H_1^2, \ldots, H_k^2\}$ since paths in each H_i^1 share a common vertex r^*. By Lemma 3.4.4, $\{H_1^1 \uplus H_1^2, \ldots, H_k^1 \uplus H_k^2\}$ are k edge-disjoint subgraphs of $\mathcal{G}$ that balance $S \cup R$ and extend $\mathcal{P}_k(v)$. This contradicts the assumption that $\mathcal{G}$ is a counterexample. Hence, $S \cup R$ is $(Q-2)$-edge-connected in $\mathcal{G}$. ∎

Elimination of the Second Case

Suppose v in $\mathcal{G}$ satisfies the requirements of the second case in the extension theorem. By Lemma 3.7.2, $S \cup R$ is $(Q-2)k$-edge-connected in $\mathcal{G}$. Consider $\mathcal{G} - v$. Since v is of degree at most Qk, we claim that $S \cup R$ is at least $(Q/2-2)k$-edge-connected in $\mathcal{G} - v$. To see this, consider any two vertices $a, b \in S \cup R$. There are at least $(Q-2)k$ edge-disjoint paths between a and b in $\mathcal{G}$. Among those paths, at most $Qk/2$ paths pass through v since v is of degree at most Qk. So, there are at least $(Q-2)k - Qk/2 = (Q/2-2)k$ paths between a and b in $\mathcal{G} - v$ and the claim follows. Since $Q = 30$, $S \cup R$ is at least $6k$-edge-connected in $\mathcal{G} - v$.

By Lemma 3.7.1, $\mathcal{G}$ has no edge between two vertices in $V(\mathcal{G}) - S - R$ and hence neither does $\mathcal{G} - v$. So, by Theorem 3.2.2, there are $2k$ edge-disjoint $(S \cup R)$-subgraphs $\{H_1', \ldots, H_{2k}'\}$ in $\mathcal{G} - v$. By setting $H_i := H_{2i-1}' \cup H_{2i}'$, we obtain k edge-disjoint double $(S \cup R)$-subgraphs in $\mathcal{G} - v$. Let $\mathcal{P}_k(v) = \{\delta_1(v), \ldots, \delta_k(v)\}$. Since $N_{\mathcal{G}}(v) \subseteq S \cup R$, it follows that $\{H_1 \cup \delta_1(v), \ldots, H_k \cup \delta_k(v)\}$ are k edge-disjoint double $(S \cup R)$-subgraphs in $\mathcal{G}$ that extend $\mathcal{P}_k(v)$. Clearly, $\{H_1 \cup \delta_1(v), \ldots, H_k \cup \delta_k(v)\}$ balance $S \cup R$ as well. Therefore, $\mathcal{G}$ is not a counterexample and so v does not satisfy the requirements of the second case in the extension theorem.

Henceforth, we assume v satisfies the requirements of the first case in the extension theorem.

When $|S| \leq 2$

If $|S| = 1$, then the extension theorem holds trivially. Here we shall show the extension theorem for the case when $|S| = 2$. Let $S = \{v, u\}$ where v is the vertex to be extended. Since S is Qk-edge-connected in $\mathcal{G}$, Menger's theorem (Theorem 2.2.3) implies that there are Qk edge-disjoint paths $\{P_1, \ldots, P_{Qk}\}$ between v and u in $\mathcal{G}$. Without loss of generality, we assume that $\mathcal{P}_k(v) = \{\delta_1(v), \ldots, \delta_k(v)\}$ is a balanced edge-partition (instead of a balanced edge-subpartition). Since v is of degree Qk, each edge in $\delta(v)$ is used exactly once. Let $P_{i_1}, \ldots, P_{i_t}$ be the paths that intersect with $\delta_i(v)$; so $t \geq 2$. We construct H_i^j as follows. Set $H_i^1 := P_{i_1}$ and $H_i^2 := H_i^1 \cup P_{i_2}$. For $j \geq 3$, let P'_{i_j} be the subpath of P_{i_j} which connects v to H_i^{j-1}. That is, P'_{i_j} is an "ear" of H_i^{j-1} which has its two endvertices in H_i^{j-1} but not the other vertices. Then we set $H_i^j := H_i^{j-1} \cup P'_{i_j}$ for $3 \leq j \leq t$, and $H_i := H_i^t$.

Now we verify that $\{H_1, \ldots, H_t\}$ are k edge-disjoint S-subgraphs which extend $\mathcal{P}_k(v)$ and balance $S \cup R$. It is clear that H_i is an S-subgraph, since $P_{i_1} \subseteq H_i$. Also, by our construction, $H_i - v$ spans $N_{\delta_i}(v)$; this is because each edge in $\delta_i(v)$ is used in exactly one path from $\{P_{i_1}, \ldots, P_{i_t}\}$. This implies that $H_i - v$ is an $(S - v)$-subgraph that spans $N_{\delta_i}(v)$. Hence $\{H_1, \ldots, H_k\}$ extend $\mathcal{P}_k(v)$.

It remains to check that $\{H_1, \ldots, H_k\}$ balance $S \cup R$. Since $P_{i_1} \cup P_{i_2} \subseteq H_i$, u and v are of degree at least 2 in each H_i, and thus are balanced by Proposition 3.3.1 (2). For each vertex $w \in R$, we shall argue that there are at least $2k$ edges which are not used in any of $\{H_1, \ldots, H_k\}$, and thus w is balanced by Proposition 3.3.1 (1). To see this, since $w \neq u$ and $w \neq v$, w is either of degree 0 or degree 2 in each path P_j. Assume that w appears in l paths from $P_1, \ldots, P_{Qk}$, say $P_1, \ldots, P_l$. Recall that $w \in R$ and thus w is of degree at least Qk. If $2l \leq (Q-2)k$, then we are done. So suppose $2l > (Q-2)k$, which implies $l - 2k \geq 2k$ so long as $Q \geq 10$. By our construction of H_i where we iteratively add "ears", for each i, at most 2 paths in $P_{i_1}, \ldots, P_{i_t}$ use two edges incident to w. So there are at most $2k$ paths among $P_1, \ldots, P_l$ from which both edges incident to w are used in $\{H_1, \ldots, H_k\}$. For the remaining paths ($l - 2k$ of them), we can save one edge incident to w by the adding "ears" process. So at least $l - 2k \geq 2k$ edges incident to w from $P_1, \ldots, P_l$ are not used in $\{H_1, \ldots, H_k\}$; and so we can apply Proposition 3.3.1 (1). Therefore, $\{H_1, \ldots, H_k\}$ are k edge-disjoint S-subgraphs that extend $\mathcal{P}_k(v)$ and balance $S \cup R$.

Henceforth, we assume that $|S| \geq 3$.

Construction and Properties of G'

Let W be the set of neighbours of v in $V(\mathcal{G}) - S - R$ and B be the set of neighbours of v in $S \cup R$. By Lemma 3.5.3, each $w_i \in W$ is incident with exactly three edges and adjacent to exactly three vertices, so we let $N_{\mathcal{G}}(w_i) := \{v, x_i, y_i\}$ and call $\{x_i, y_i\}$ a *couple*. By Lemma 3.7.1, x_i and y_i are in $S \cup R$. For each $b_i \in B$, we denote by $c(b_i)$ the number of multiple edges between v and b_i.

Set $G' := \mathcal{G} - v - W$. Let Z be a minimum $(S \cup R - v)$-edge-cutset of G' and $\{C_1, \ldots, C_l\}$ be the connected components of $G' - Z$. We let $S_i := S \cap V(C_i)$, $R_i := R \cap V(C_i)$ and $B_i := B \cap V(C_i)$. Also, $c(B_i)$ denotes the sum of the $c(b)$ for $b \in B_i$ and X_i denotes the collection of couples with both vertices in C_i. By the minimality of Z, each edge e in Z connects two vertices in different components, and we call it a *crossing edge*. Similarly, a couple $\{x_i, y_i\}$ is a *crossing couple* if x_i and y_i are in different components, and we denote the collection of crossing couples by X_C.

Lemma 3.7.3 *$(S \cup R - v)$ is at most $(6k - 1)$-edge-connected in G'.*

Proof. Since $|S| \geq 3$, we have $|S \cup R - v| \geq 2$. If $(S \cup R - v)$ is $6k$-edge-connected in G', then by Theorem 3.2.2, there are $2k$ edge-disjoint $(S \cup R - v)$-subgraphs $\{H''_1, \ldots, H''_{2k}\}$ in G'. Notice that since the union of two edge-disjoint $(S \cup R - v)$-subgraphs is a double $(S \cup R - v)$-subgraph, by setting $H'_i := H''_{2i-1} \cup H''_{2i}$, $\{H'_1, \ldots, H'_k\}$ are k edge-disjoint double $(S \cup R - v)$-subgraphs of G'. Now, let $H_i := H'_i \cup \{vb_j | vb_j \in \delta_i(v)\} \cup \{vw_j, w_jx_j | vw_j \in$

$\delta_i(v)\}$. Notice that $H_i - v$ is connected since H_i' is a $(S \cup R - v)$-subgraph and b_j, x_j are terminal vertices. So, by our construction, $\delta_i(v) \subseteq H_i$, and $H_i - v$ is a $(S \cup R \cup N_{\delta_i}(v) - v)$-subgraph. Also, since $\mathcal{P}_k(v)$ is a balanced edge-subpartition, $\{H_1, \ldots, H_k\}$ balance $S \cup R$ as well. Therefore, $\{H_1, \ldots, H_k\}$ are k edge-disjoint S-subgraphs of $\mathcal{G}$ that balance $S \cup R$ and extend $\mathcal{P}_k(v)$. This contradicts with the fact that $\mathcal{G}$ is a counterexample. So $(S \cup R - v)$ must be at most $(6k-1)$-edge-connected in G'. ∎

Lemma 3.7.4 *$G' - Z$ has 2 connected components.*

Proof. To show that $G' - Z$ has 2 connected components, we just need to show that G' has at most 2 connected components, then the rest follows from the minimality of Z. Notice that from our construction of G' from $\mathcal{G}$, for each vertex $u \in V(G') - S - R$, we have $N_{G'}(u) = N_{\mathcal{G}}(u)$ by Lemma 3.7.1. Since $\mathcal{G}$ is connected, no component in G' contains only vertices in $V(\mathcal{G}) - S - R$; otherwise by Lemma 3.7.1 it is an isolated vertex and is also isolated in $\mathcal{G}$, which implies that $\mathcal{G}$ is disconnected too. Therefore, it remains to show that there are at most two components in G' that contain vertices in $S \cup R$.

Suppose there are three connected components containing vertices in $S \cup R$. Let $u_1, u_2, u_3 \in S \cup R$ be vertices in C_1, C_2, C_3 respectively. Since u_1, u_2, u_3 each have at least $(Q-2)k$ edge-disjoint paths to v in $\mathcal{G}$ and v is of degree Qk, there exists a vertex $w \in N_{\mathcal{G}}(v)$ such that u_1, u_2, u_3 all have a path to w in $\mathcal{G}$. If $w \in S \cup R$, then clearly u_1, u_2, u_3 are still connected in G', a contradiction. If $w \in V(\mathcal{G}) - S - R$, then w is of degree 3 by Lemma 3.5.3. Let the two neighbours of w in G' be x and y. Then there must exist a pair of vertices, say u_1 and u_2, that both have a path to the same neighbour of w, say x, in G'. So there is a path from C_1 to C_2, a contradiction. ∎

Diverging paths and common paths

Consider a vertex $u \neq v$ where $u \in S$. Since $v \in S$ and S is Qk-edge-connected in $\mathcal{G}$, by Menger's theorem, there are Qk edge-disjoint paths, denoted by $P(u) :=$

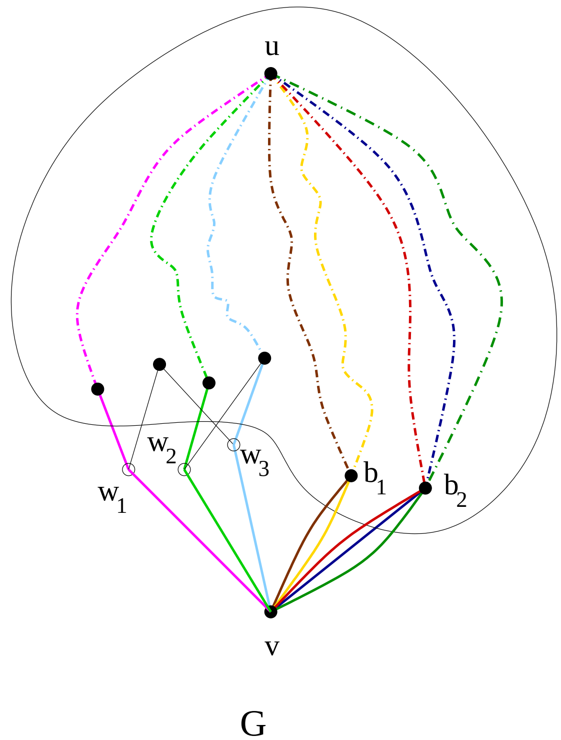

Figure 3.3: The paths in dotted lines are paths in $P'(u)$.

$\{P_1(u), \ldots, P_{Qk}(u)\}$, from u to v. Note that since v is of degree exactly Qk, each path in $P(u)$ uses exactly one edge in $\delta(v)$. Furthermore, since w_i is of degree 3 by Lemma 3.5.3, each w_i is used by exactly one path in $P(u)$. Similarly, for $u \neq v$ where $u \in R$, there are $(Q-2)k$ edge-disjoint paths from u to v denoted by $P(u) := \{P_1(u), \ldots, P_{(Q-2)k}(u)\}$. By choosing a set of $(Q-2)k$ paths which minimize the total length (i.e. $\sum_{i=1}^{(Q-2)k} |P_i(u)|$), we can assume that each path in $P(u)$ uses exactly one edge in $\delta(v)$ and at most one vertex in W. Consider $P_i(u)$ induced in G', denoted by $P'_i(u)$. Let $P'(u) := \{P'_1(u), \ldots, P'_{Qk}(u)\}$ for $u \in S$, and similarly $P'(u) := \{P'_1(u), \ldots, P'_{(Q-2)k}(u)\}$ for $u \in R$. Notice that paths in $P'(u)$ are edge-disjoint in G', and we call them the *diverging paths* from u. See Figure 3.3 for an illustration.

We plan to use the diverging paths from a and b for any two vertices $a, b \in S \cup R$ in the same component of $G' - Z$ to establish the edge-connectivity of $S \cup R$ in each component of $G' - Z$. We say v_1 and v_2 have λ *common paths* if there are λ edge-disjoint paths starting from v_1, λ edge-disjoint paths starting from v_2, and an one-to-one mapping of the paths from v_1 to the paths from v_2 so that each pair of paths in the mapping ends in the same vertex.

The following lemma gives a lower bound on the number of edge-disjoint paths between two vertices based on the number of their common paths; note that this lemma holds for any graph G.

Lemma 3.7.5 *If v_1 and v_2 have $2\lambda + 1$ common paths in G, then there exist $\lambda + 1$ edge-disjoint paths from v_1 to v_2 in G.*

Proof. Suppose not. By Menger's theorem (Theorem 2.2.3), there is an edge-cutset T of size at most λ that disconnects v_1 and v_2 in G. Since $|T| \leq \lambda$, at least $\lambda + 1$ paths starting from v_1 remain in $G - T$; and the same holds for v_2. So, v_1 and v_2 still have at least $(\lambda + 1) + (\lambda + 1) - (2\lambda + 1) = 1$ common path in $G - T$. This implies that v_1 and v_2 are connected in $G - T$, a contradiction. ∎

When Both Components of $G' - Z$ Contain Some Vertices in S

In this subsection, we consider the case that both components contain some vertices in S. The lemmas in this section all share this assumption.

Lemma 3.7.6 *If both components of $G' - Z$ contain some vertices in S, then there are at least $Qk - 2|Z|$ crossing couples, that is, $|X_C| \geq Qk - 2|Z|$.*

Proof. Let $u_1 \in S$ be in C_1. In G', from the diverging paths in $P'(u_1)$, u_1 has at least $c(B_2) + |X_2|$ edge-disjoint paths to C_2. Since Z is an edge-cut in G', it follows that $c(B_2) + |X_2| \leq |Z|$. Similarly, by considering a vertex $u_2 \in S$ in C_2, we have $c(B_1) + |X_1| \leq |Z|$. By Lemma 3.7.4, there are only two components in $G' - Z$. So $Qk = |X_C| + |X_1| + |X_2| + c(B_1) + c(B_2) \leq |X_C| + 2|Z|$, and we have $|X_C| \geq Qk - 2|Z|$. ∎

With Lemma 3.7.6, we can show that $S_i \cup R_i$ is highly edge-connected in C_i for $i = \{1, 2\}$.

Lemma 3.7.7 *If both components of $G' - Z$ contain some vertices in S, then for $i \in \{1, 2\}$, $S_i \cup R_i$ is $(Q/2 - 8)k$-edge-connected in C_i of $G' - Z$ and S_i is $(Q/2 - 6)k$-edge-connected in C_i of $G' - Z$.*

Proof. Consider any two vertices $a, b \in S_1 \cup R_1$ in C_1. If $a \in S_1$, then $P'(a)$ has $|X_C|$ edge-disjoint paths to the crossing couples, one path for each crossing couple. If $a \in R_1$, then $P'(a)$ has at least $|X_C| - k_a$ edge-disjoint paths to the crossing couples, where k_a denotes the number of crossing couples with no path to a. In the following argument, we say $P'(a)$ has at least $|X_C| - k_a$ paths to different crossing couples, with the understanding that $k_a = 0$ for $a \in S$. Assume that, among those $|X_C| - k_a$ paths, z_a paths use edges in Z. Then, in $G' - Z$, a has at least $|X_C| - k_a - z_a$ edge-disjoint paths such that each starts from a and ends in a different crossing couple. We define k_b and z_b similarly. By the same argument, in $G' - Z$, b has at least $|X_C| - k_b - z_b$ edge-disjoint paths such that each starts from b and ends in a different crossing couple. Therefore, in $G' - Z$, a and b have at least $(|X_C| - k_a - z_a) + (|X_C| - k_b - z_b) - |X_C| = |X_C| - k_a - k_b - z_a - z_b$ pairs of edge-disjoint paths such that each pair of paths ends in the same crossing couple. Since a, b are in the same component, each such pair ends in the same member of a crossing couple. So, a and b have at least $|X_C| - k_a - k_b - z_a - z_b$ common paths in C_i.

Also, in G', for $a \in S_1$, $P'(a)$ has $c(B_2)$ edge-disjoint paths to B_2 and $|X_2|$ edge-disjoint paths to the couples in X_2. So, for $a \in S_1$, $P'(a)$ has $c(B_2) + |X_2|$ edge-disjoint paths to C_2 in G'. Similarly, for $a \in R_1$, $P'(a)$ has $c(B_2) + |X_2| - m_a$ edge-disjoint paths to C_2 in this way, where m_2 denotes the number of members of B_2 with no path to a. Notice that since $P'(a)$ has at least $(Q - 2)k$ paths for $a \in R_i$, $k_a + m_a \leq 2k$; if $a \in S_1$, $k_a + m_a = 0$. As mentioned in the previous paragraph, $P'(a)$ also has z_a edge-disjoint paths to crossing couples that use edges in Z. These $c(B_2) + |X_2| - m_a + z_a$ paths are edge-disjoint, because each of these paths has a different destination. Since Z is an edge-cut, Z has at least one edge in each such path. So, in $G'[C_1]$, a has at least $c(B_2)+|X_2|-m_a+z_a$ edge-disjoint paths such that each starts from a and ends in an endpoint of a different crossing edge in Z. Similarly, $P'(b)$ has $c(B_2)+|X_2|-m_b+z_b$ edge-disjoint paths such that each starts from b and ends in an endpoint of a different crossing edge in Z. Therefore, in $G'[C_1]$, a and b have at least $(c(B_2)+|X_2|-m_a+z_a)+(c(B_2)+|X_2|-m_b+z_b)-|Z| = 2c(B_2)+2|X_2|-m_a-m_b+z_a+z_b-|Z|$ pairs of paths such that each pair of paths ends in the same endpoint of a crossing edge. These pairs of paths are edge-disjoint from those paths mentioned in the previous paragraph, because they all have different destinations. So, a and b have at least $2c(B_2)+2|X_2|-m_a-m_b+z_a+z_b-|Z|$ more common paths in C_1.

As a result, by the previous two paragraphs, a and b have at least $(|X_C|-k_a-k_b-z_a-z_b)+(2c(B_2)+2|X_2|-m_a-m_b+z_a+z_b-|Z|) = 2c(B_2)+2|X_2|+|X_C|-(k_a+m_a)-(k_b+m_b)-|Z|$ common paths in C_1. Since $k_a+m_a \leq 2k$ and $k_b+m_b \leq 2k$, a and b have at least $2c(B_2)+2|X_2|+|X_C|-4k-|Z|$ common paths in C_1. Recall from the proof of Lemma 3.7.6 that $Qk = |X_C|+|X_1|+|X_2|+c(B_1)+c(B_2)$ and $c(B_1)+|X_1| \leq |Z|$, so $|X_C|+|X_2|+c(B_2)+|Z| \geq Qk$ and hence $|X_C|+|X_2|+c(B_2) \geq Qk-|Z|$. So a and b have at least $Qk+c(B_2)+|X_2|-4k-2|Z| \geq Qk-4k-2|Z| > (Q-16)k$ common paths in C_1; the last inequality is because Lemma 3.7.3 implies that $|Z| < 6k$. Furthermore, we observe that if a and b are both in S_1, then by the same calculation a and b have at least $(Q-12)k$ common paths in C_1; this is because $k_a+m_a = k_b+m_b = 0$.

Remark 3.7.8 *If we use the inequality $|X_C|+|X_2|+c(B_2) \geq (Q-2)k-|Z|$ instead of the inequality $|X_C|+|X_2|+c(B_2) \geq Qk-|Z|$ used above, then we conclude that a and b have at least $(Q-18)k$ common paths in C_1. This will be used in Lemma 3.7.14 which is a counterpart of the current lemma.*

Now, by Lemma 3.7.5, there are at least $(Q/2-8)k$ edge-disjoint a,b-paths in C_1 for $a,b \in S_1 \cup R_1$. Furthermore, for $a,b \in S_1$, there are at least $(Q/2-6)k$ edge-disjoint

a, b-paths in C_1. So, $S_1 \cup R_1$ is $(Q/2-8)$-edge-connected in C_1 and S_1 is $(Q/2-6)k$-edge-connected in C_1. By the same argument, we can show that $S_2 \cup R_2$ is $(Q/2-8)k$-edge-connected in C_2 and S_2 is $(Q/2-6)k$-edge-connected in C_2. ∎

The following is a technical lemma which shows that we can construct edge-disjoint S-subgraphs even if we delete a limited number of edges.

Lemma 3.7.9 *Let G be a graph and S be a subset of $V(G)$. Suppose S is $3k$-edge-connected in G, every vertex in $V(G)-S$ is of degree 3, and there is no edge between vertices in $V(G)-S$. Let T be a set of edges with $|T| \leq k$. Then $G-T$ has k edge-disjoint S-subgraphs.*

Proof. This is slightly stronger than Theorem 3.2.2, and we shall use the same approach used in [34] to prove it. Recall that a hypergraph H is k-partition-connected if $E_H(\mathcal{P}) \geq k(|\mathcal{P}|-1)$ holds for every partition $\mathcal{P}$ of $V(H)$ into non-empty classes, where $E_H(\mathcal{P})$ denotes the number of hyperedges intersecting at least two classes. From G, we construct a hypergraph H with vertex set S. For every vertex $v \in V(G)-S$ in G, we add a hyperedge of size 3 in H consisting of the neighbours $N_G(v)$ of v in G, notice that $N_G(v) \subseteq S$ by assumption and hence is well-defined in H. Also, for every edge $uv \in E(G)$ with $u, v \in S$, we add a hyperedge uv in H. For each edge $e \in T$ of $E(G)$, we remove the hyperedge in H that contains e. Let the resulting hypergraph be H'. We first claim that a connected sub-hypergraph of H' corresponds to an S-subgraph of $G-T$. To see this, we just need to replace a hyperedge in H' by the corresponding Steiner vertex and the edges incident to it; then it is an S-subgraph of G. The fact that it is also an S-subgraph of $G-T$ follows from our construction that the hyperedges containing edges in T are removed. Therefore, if we can prove that H' has k edge-disjoint connected sub-hypergraphs, then we can easily construct k edge-disjoint S-subgraphs in $G-T$. We now prove this claim using Theorem 3.2.1.

Consider a partition $P := \{P_1, \ldots, P_l\}$ of the vertex set S where $l \geq 2$. We argue that $d_H(P_i) \geq 3k$ for each i; in other words, there are at least $3k$ hyperedges in H "crossing" P_i. Suppose, by way of contradiction, that $d_H(P_i) < 3k$. Consider a hyperedge e of size 3 in $\delta_H(P_i)$, e either contains one vertex in P_i or one vertex in $S - P_i$. By our construction, e corresponds to a vertex w of degree 3 in $V(G) - S$. So we can remove just one (graph) edge incident to w in G to decrease $d_H(P_i)$ by 1 in the corresponding hypergraph. Clearly, if $e \in \delta_H(P_i)$ is of size 2, then we can remove just one edge in G to decrease $d_H(P_i)$ by 1. Since $d_H(P_i) < 3k$, we can remove less than $3k$ (graph) edges in G to disconnect the corresponding hypergraph. This implies that, by removing less than $3k$ edges in G, S is disconnected in G. But this contradicts the fact that S is $3k$-edge-connected in G, hence $d_H(P_i) \geq 3k$.

Since each vertex in $V(G) - S$ is of degree 3, each hyperedge in H can intersect at most 3 classes of P. Therefore, in H, there are at least $3kl/3 = kl$ crossing hyperedges for P. Constructing H' from H removes at most k hyperedges. So, in H', there are still at least $(l-1)k$ crossing hyperedges. Since P is an arbitrary partition, H' is k-partition-connected. By Theorem 3.2.1, we have k edge-disjoint connected sub-hypergraphs in H'. Hence we have k edge-disjoint S-trees in $G - T$. This proves the lemma. ∎

Now we can finally prove the extension theorem for the case when both components contain some vertices in S. Note that the assumption that $\mathcal{P}_k(v)$ is a balanced edge-subpartition is crucial in this lemma.

Lemma 3.7.10 *If both components of $G' - Z$ contain some vertices in S, then $\mathcal{G}$ has k edge-disjoint S-subgraphs $\{H_1, \ldots, H_k\}$ that balance $S \cup R$ and extend $\mathcal{P}_k(v)$.*

Proof. The strategy is to combine k edge-disjoint subgraphs in each component to construct k edge-disjoint subgraphs in $\mathcal{G}$ with the desired property. By the previous lemma, each $S_i \cup R_i$ is $(Q/2 - 8)k$-edge-connected in C_i. Since $Q = 30$, $S_i \cup R_i$ is at least $6k$-edge-connected in C_i. We remark that in the special case of STEINER TREE

PACKING, since $R_i = \emptyset$ and S_i is $(Q/2-6)k$-edge-connected in C_i, we just need $Q = 24$ to guarantee that S is $6k$-edge-connected in C_i.

Now, we pick arbitrarily $\min\{k, |Z|\}$ edges in Z and call them the *connecting edges*; they will be used to connect the subgraphs in each component. For each connecting edge e with an endpoint $w \in V(\mathcal{G}) - S - R$ in C_i, we remove one edge e' in C_i which is incident with w and we call e' a *reserve edge* of e. By Lemma 3.7.1, the other endpoint of e' must be a vertex in $S \cup R$, so for each connecting edge e the path e, e' is from terminal vertex to terminal vertex. Since w is a Steiner vertex, w is of degree 3 by Lemma 3.5.3. We must have $\delta(w) \cap Z = \{e\}$; otherwise this contradicts the minimality of Z. This has two implications. The first implication is that each connecting edge will be assigned a different reserve edge. The second implication is that, as we will show, after the removal of the reserve edges, $V(C_i)$ is still connected and we denote it by C_i'. Clearly, $S_i \cup R_i$ is still connected in C_i' because $S_i \cup R_i$ is $6k$-edge-connected in C_i and we remove at most k reserve edges. For any Steiner vertex w for which $\delta(w) \cap Z \neq \emptyset$, we remove exactly two edges incident to w, and since $|\delta(w) \cap Z| = 1$ as argued above, w is still connected to a vertex in $S_i \cup R_i$ through the remaining edge. Hence $V(C_i) = V(C_i')$ is still connected.

Since we remove at most k reserved edges from C_i and $S_i \cup R_i$ is $6k$-edge-connected in C_i, by using Lemma 3.7.9 (with $S := S_i \cup R_i$), there are $2k$ edge-disjoint $(S_i \cup R_i)$-subgraphs in C_i'. By grouping two $(S_i \cup R_i)$-subgraphs into one, we have k edge-disjoint double $(S_i \cup R_i)$-subgraphs $\{H_1^i, \ldots, H_k^i\}$ in C_i' for $i \in \{1, 2\}$. We remark that if $|S_i \cup R_i| = 1$, then the subgraphs might have no edges. We shall handle this case later.

Now we set $H_j := H_j^1 \cup H_j^2 \cup \{vb_i | vb_i \in \delta_j(v)\} \cup \{vw_i, w_i x_i, w_i y_i | vw_i \in \delta_j(v)\}$ for $1 \leq j \leq k$. Notice that by construction $\delta_j(v) \subseteq E(H_j)$ and $H_j - v$ spans $N_{\delta_j}(v)$ for $1 \leq j \leq k$. Each $H_j - v$ is almost an $(S \cup R - v)$-subgraph except that the two subgraphs in the two components might not be connected to each other. Suppose there is a crossing couple $\{x_i, y_i\}$ such that $vw_i \in \delta_j(v)$, then $H_j - v$ is connected and thus is an $(S \cup R - v)$-subgraph of $\mathcal{G}$ so that $\delta_j(v) \subseteq E(H_j)$ and $H_j - v$ extends $\delta_j(v)$. So, if they are all connected by crossing couples, then $\{H_1, \ldots, H_k\}$ are k edge-disjoint S-subgraphs that extend $\mathcal{P}_k(v)$. Let $\{vw_1, \ldots, vw_{|X_C|}\}$ be the set of edges such that the corresponding couples are crossing. By Lemma 3.7.6, $|X_C| \geq Qk - 2|Z|$. Since $\mathcal{P}_k(v)$ is a balanced edge-subpartition, we can assume that $|E_i(v)| \geq 2$ for $1 \leq i \leq k$. So, there are at most $\min\{k, |Z|\}$ classes of $\mathcal{P}_k(v)$ with no edges in $\{vw_1, \ldots, vw_{Qk-2|Z|}\}$. Hence there are at most $\min\{k, |Z|\}$ H_j's, say $\{H_1, \ldots, H_{\min\{k,|Z|\}}\}$, that are not connected by the crossing couples. Now, by adding each connecting edge and its reserve edge (if necessary) to a different H_j that has not yet been connected by a crossing couple, all H_j are connected since we have $|Z|$ connecting edges. Therefore, $\{H_1, \ldots, H_k\}$ are k edge-disjoint $(S \cup R)$-subgraphs of $\mathcal{G}$ that extend $\mathcal{P}_k(v)$. (Note that the assumption that $\mathcal{P}_k(v)$ is a balanced edge-subpartition is crucial here.)

Remark 3.7.11 *If $|Z| \geq k$, then the above proof would succeed without using the assumption that $\mathcal{P}_k(v)$ is a balanced edge-subpartition. This is because there are enough connecting edges to make all H_j connected, and so we don't need to use the crossing couples for that purpose.*

To finish off the proof, we need to make sure that $\{H_1, \ldots, H_k\}$ balance $S \cup R$. We do so by making sure that $\{H_1, \ldots, H_k\}$ are double $(S \cup R)$-subgraphs; thus each $v \in S \cup R$ has degree at least 2 in each H_i, and so v is balanced by Proposition 3.3.1 (2). Suppose $|S_1 \cup R_1| \geq 2$, then $\{H_1^1, \ldots, H_k^1\}$ are double $(S_1 \cup R_1)$-subgraphs and hence $\{H_1, \ldots, H_k\}$ balance $S_1 \cup R_1$. The same argument can be applied to $S_2 \cup R_2$. The subtle case is when $|S_1 \cup R_1| = 1$, say $S_1 = \{x\}$ and $R_1 = \emptyset$, where each H_i^1 might have no edges. Note that x is in every crossing couple in this case. Let $\{H_1, \ldots, H_l\}$ be the S-subgraphs in which x is a degree 1 vertex. Suppose $\{\{x, y_1\}, \{x, y_2\}, \ldots, \{x, y_c\}\}$ are crossing couples such that $\{\{vw_1, xw_1, y_1w_1\}, \ldots, \{vw_c, xw_c, y_cw_c\}\} \subseteq E(H_j)$ and $\{vw_1, vw_2, \ldots, vw_c\} \subseteq \delta_j(v)$. If $c > 2$, then we can delete $\{xw_3, \ldots, xw_c\}$ from H_j without affecting the properties of H_j that are required by the extension theorem (as it is already connected and balances v, x).

We repeat this procedure until there are at least l edges, say $\{xw_1, \ldots, xw_l\}$, that are not used in any H_j. Then we can add each such edge to a different S-subgraph in $\{H_1, \ldots, H_l\}$ to make x to be of degree at least 2 in each of $\{H_1, \ldots, H_k\}$. We do the same "switching" procedure if $|S_2| = 1$. Since there are at least $Qk - 2|Z| > (Q-12)k = 18k$ crossing couples and there are only 2 components in $G' - Z$, there are more than enough edges for this "switching" procedure. Finally, $\{H_1, \ldots, H_k\}$ are k edge-disjoint S-subgraphs of $\mathcal{G}$ that balance $S \cup R$ and extend $\mathcal{P}_k(v)$. ∎

In the special case of STEINER TREE PACKING, since $R = \emptyset$, each component must contain some vertices in S and hence the results in this subsection suffice. Therefore, the minimal counterexample $\mathcal{G}$ does not exist and hence the extension theorem (Theorem 3.3.4) follows. As we remarked in the proof of Lemma 3.7.10, we just need $Q := 24$ for STEINER TREE PACKING. So, we have Theorem 3.1.2 as claimed.

When One Component Contains No Vertices in S

Without loss of generality, we assume that C_1 contains vertices in S and C_2 contains only vertices in R. Lemma 3.7.12 and Lemma 3.7.14 in the following are counterparts of Lemma 3.7.6 and Lemma 3.7.7. However, it should be pointed out that we only have a weaker bound on the number of crossing couples in Lemma 3.7.12, and hence the strategy in Lemma 3.7.10 will not always work. In Lemma 3.7.15, which is a counterpart of Lemma 3.7.10, we use a different strategy to handle the remaining case.

Lemma 3.7.12 *If C_2 contains only vertices in R, then there are at least $(Q-2)k - 2|Z|$ crossing couples, that is, $|X_C| \geq (Q-2)k - 2|Z|$.*

Proof. Let $u_1 \in S$ be in C_1. In G', u_1 has at least $c(B_2) + |X_2|$ edge-disjoint paths in $P'(u_1)$ to C_2. Since Z is an edge-cut in G', it follows that $c(B_2) + |X_2| \leq |Z|$. Let $u_2 \in R$ be in C_2. In G', we have at least $c(B_1) + |X_1| - 2k \leq |Z|$ edge-disjoint paths in $P'(u_2)$ to C_1. By Lemma 3.7.4, there are only two components in $G' - Z$. So $Qk =$

$|X_C|+|X_1|+|X_2|+c(B_1)+c(B_2) \leq |X_C|+2|Z|+2k$, and we have $|X_C| \geq (Q-2)k-2|Z|$. ∎

Remark 3.7.13 *In the proof of Lemma 3.7.12, we have $Qk = |X_C|+|X_2|+c(B_2)+|X_1|+c(B_1)$ and $c(B_1)+|X_1|-2k \leq |Z|$. This implies that $|X_C|+|X_2|+c(B_2) \geq (Q-2)k-|Z|$, which will be useful in Lemma 3.7.14.*

The following lemma is a counterpart of Lemma 3.7.7 which shows that $S_1 \cup R_1$ is highly edge-connected in C_1, but with a slightly weaker lower bound.

Lemma 3.7.14 *If C_2 contains only vertices in R, then $S_1 \cup R_1$ is at least $(Q/2-9)k$-edge-connected in C_1 of $G' - Z$.*

Proof. Consider any two vertices $a, b \in S_1 \cup R_1$ in C_1. The proof is almost identical to the proof of Lemma 3.7.7, except that we only have a weaker lower bound on $|X_C|$. By Remark 3.7.13, we have $|X_C| + |X_2| + c(B_2) \geq (Q-2)k - |Z|$. Then, by Remark 3.7.8, we conclude that a and b have at least $(Q-18)k$ common paths in C_1. This implies that there are at least $(Q/2-9)k$ edge-disjoint a, b-paths in C_1, and hence $S_1 \cup R_1$ is $(Q/2-9)k$-edge-connected in C_1. ∎

Now we are ready to prove the extension theorem for the case when one component contains no vertices in S.

Lemma 3.7.15 *If C_2 contains only vertices in R, then $\mathcal{G}$ has k edge-disjoint S-subgraphs that balance $S \cup R$ and extend $\mathcal{P}_k(v)$.*

Proof. If $|Z| \geq k$, we have the desired result by Remark 3.7.11. Henceforth we assume that $|Z| < k$, in which case we may not have enough connecting edges since the bound on crossing couples is weaker. Here we use another strategy to handle this situation. Let $E' := \{e_1, \ldots, e_{|X_2|+c(B_2)}\}$ be the set of edges incident to v in $\mathcal{G}$ so that either (i) the other endpoint of e_i is in B_2 or (ii) the other endpoint of e_i has both of its neighbours in C_2. Intuitively, E' is the set of edges which we would like to be extended properly.

Let $u_1 \in S$. From $P(u_1)$ (the diverging paths) in $\mathcal{G}$, there are $|X_2| + c(B_2)$ edge-disjoint paths from u_1 to v such that each uses exactly one edge in E'. From these paths, in $\mathcal{G} - E(C_1)$, there are $|X_2| + c(B_2)$ edge-disjoint paths $P := \{P_1, \ldots, P_{|X_2|+c(B_2)}\}$ with the following property: each P_i starts from v and ends in some vertex of C_1, and $e_i \in P_i$. Since each vertex $w \in V(C_1) - S - R_1$ is of degree 3, by the minimality of $|Z|$, it has at most one neighbour in C_2. Therefore, each $w \in V(C_1) - S - R_1$ can be in at most one path $P_i \in P$. Now, for each such w, we remove one edge e' in C_1 which is incident with w and set $P'_i := P_i \cup \{e'\}$; for every P_i not containing such a w, we set $P'_i := P_i$. We call those edges removed the reserve edges. Notice that by Lemma 3.7.1, the other endpoint of e' must be in $S \cup R_1$. So $P' := \{P'_1, \ldots, P'_{|X_2|+c(B_2)}\}$ are edge-disjoint paths with the following property: each P'_i starts from v and ends in some vertex $u_i \in S \cup R_1$, and $e_i \in P'_i$. These paths will be used to extend the edge-subpartition induced on E'.

In constructing P' from P, we remove at most $|X_2| + c(B_2) \leq |Z| \leq k$ reserve edges from C_1; the first inequality appeared in the proof of Lemma 3.7.6. Following the same argument as in Lemma 3.7.10, $V(C_1)$ remains connected after the removal of the reserve edges. Let the resulting component be C'_1. By Lemma 3.7.14, $S \cup R_1 - v$ is $(Q/2-9)k$-edge-connected in C_1. As $Q = 30$, $S \cup R_1 - v$ is $6k$-edge-connected in C_1. Since we remove at most k edges from C_1 and $S \cup R_1 - v$ is $6k$-edge-connected in C_1, by using Lemma 3.7.9 (with $S := S \cup R_1 - v$), there are $2k$ edge-disjoint $(S \cup R_1 - v)$-subgraphs in C'_1. By grouping two $(S \cup R_1 - v)$-subgraphs into one, we have k edge-disjoint double $(S \cup R_1 - v)$-subgraphs $\{H^1_1, \ldots, H^1_k\}$ in C'_1. Now, we set $H_j := H^1_j \cup \{vb_i | vb_i \in \delta_j(v)\} \cup \{vw_i, w_ix_i, w_iy_i | vw_i \in \delta_j(v)\} \cup \{P'_i | e_i \in \delta_j(v)\}$.

We claim that $\{H_1, \ldots, H_k\}$ are k edge-disjoint S-subgraphs that balance $S \cup R$ and extend $\mathcal{P}_k(v)$. By our construction, $\delta_j(v) \subseteq H_j$ and $H_j - v$ is connected and spans $N_{\delta_j}(v)$. So, $\{H_1, \ldots, H_k\}$ are k edge-disjoint S-subgraphs that extend $\mathcal{P}_k(v)$. Every vertex in R_1

is of degree at least 2 in each H_i^1 and thus H_i, and so it is balanced by Proposition 3.3.1 (2). Also, by our construction of $\{H_1, \dots, H_k\}$, every vertex in R_2 is used by at most $|X_2| + c(B_2) \leq |Z| < 6k$ paths from P' (the first inequality is stated in the proof of Lemma 3.7.12 and the second inequality follows from Lemma 3.7.3), and thus has degree at most $12k$ in H_j. Since $Q = 30$, every vertex in R_2 has at least $Qk - 12k > 2k$ edges not used in $\{H_1, \dots, H_k\}$, and hence $\{H_1, \dots, H_k\}$ balance R_2 by Proposition 3.3.1 (1); we remark that $\{H_1, \dots, H_k\}$ need not span R_2 as they are only required to be S-subgraphs. This completes the proof of the lemma. ∎

Putting Lemma 3.7.10 and Lemma 3.7.15 together shows that a minimal counterexample $\mathcal{G}$ of Theorem 3.3.4 does not exist, and this completes the proof of Theorem 3.3.4. ■

We remark that using the current method, it seems that we need $Q \geq 4$ because of the double subgraphs condition. Also, it does not seem straightforward to improve the bound on Q significantly by a refinement of the current techniques.

Acknowledgment

The idea of using Lemma 3.7.9 in the proof of Lemma 3.7.10 is from Oleg Pikhurko in Carnegie Mellon University, who generously allows his idea to be put into this thesis. This improves the previous constants of [60] and [61] from 26 and 32 to 24 and 30 respectively.

3.8 NP-completeness

As we have seen, the special case of Steiner Tree Packing when there is no edge between two Steiner vertices and every Steiner vertex is of degree 3 is a major building block of our main result. It is natural to ask whether we can prove a stronger result in this special case (for example beat the bound of Theorem 3.2.2). Notice that the bound of Theorem 3.2.2 is obtained from an exact min-max relation by Theorem 3.2.1. So, one may ask whether an exact min-max relation can be proved for this rather special case. In this subsection, however, we shall show that this special case remains NP-complete. Hence obtaining such a nice characterization is unlikely.

Theorem 3.8.1 *The* Steiner Tree Packing *problem, when restricted to a graph with no edge between two Steiner vertices and with every Steiner vertex of degree 3, is NP-complete.*

Proof. In fact, we shall prove that even packing 3 edge-disjoint Steiner trees in this special case is NP-complete. We show that 3-edge-colourability of cubic graphs can be reduced to this problem. 3-edge-colourability of cubic graphs is shown to be NP-complete by Holyer [44] (see also Schrijver [84], chapter 28, page 468-469).

Here is the construction. Given a cubic graph G, we subdivide each edge of $E(G)$ to form H. That is, for each edge $e = uv$, we add a new vertex x_e and replace the edge e by two edges ux_e and $x_e v$. Also, we add a root vertex r to H and connect it to all the vertices in H corresponding to edges in G. The vertices in the union of $V(G)$ and r are terminal vertices, and all the vertices corresponding to edges in G are Steiner vertices. By our construction, in H, there is no edge between two Steiner vertices and every Steiner vertex is of degree 3. Now, we argue that G is 3-edge-colourable if and only if H has 3 edge-disjoint Steiner trees. Hence the latter problem is NP-complete.

Suppose G is 3-edge-colourable. Let E_1 be the edge set which has colour red. By setting $T_1 := \{ux_e, vx_e, rx_e |$ for $uv = x_e \in E_1\}$, we claim that T_1 is a Steiner tree in H. Since G is a cubic graph, each vertex is adjacent to three edges and they must be of different colours in a valid 3-edge-colouring. So E_1 is a perfect matching in G, and hence T_1 spans every terminal vertex. Also, T_1 is connected because every terminal vertex has a path to r. By constructing T_2 and T_3 similarly using E_2 and E_3, we obtain three Steiner trees in H. Since each edge in G is assigned at most one colour, the three trees we constructed used disjoint Steiner vertices and hence are edge-disjoint. This proves one direction.

Suppose H has three disjoint Steiner trees T_1, T_2, T_3. Let us colour the edge set of each tree with a different colour. For a terminal vertex $v \neq r$, v is of degree 3 and hence the edges incident to it must be of different colours; so in particular every edge incident to v must be coloured and v must be a leaf in each of the three trees. For each Steiner vertex x_e, we claim that all of its three incident edges must be in the same tree. Let say x_e be adjacent to u, v, r in H. By the previous argument, ux_e must be coloured, say red. Since the red colour subgraph is connected and u, v are leaves, rx_e must also be coloured red. Also, vx_e must be coloured. If vx_e is of another colour, say blue, then the blue colour subgraph is not connected since v is a leaf. Therefore, vx_e must be of colour red also. So, for each Steiner vertex, all of its incident edges must be in the same tree. Now, for each Steiner tree T_i, let E_i be the set of edges in G corresponding to the Steiner vertices used by T_i in H. By our previous argument, E_1, E_2 and E_3 must be edge-disjoint in G and their union is equal to $E(G)$. Also, each terminal vertex (except the root which is not in G) in H is of degree 1 in each Steiner tree, so each vertex in G is of degree 1 in each E_i. Hence $\{E_1, E_2, E_3\}$ is a valid 3-edge-colouring of G. This proves the other direction, and we are done. ∎

3.9 Algorithmic Aspects

In this section, we present a polynomial time algorithm to find the edge-disjoint $\mathcal{S}$-forests guaranteed by Theorem 3.1.1. Here is an overview:

The proof of Theorem 3.3.2 (which implies Theorem 3.1.1) considers a minimum counterexample and shows that it does not exist. This proof structure can be rephrased as a recursive algorithm. In the proof we showed that a minimum counterexample must have various special structures (e.g. no edges between Steiner vertices); otherwise we argued that it is not minimal. As detailed below, these steps correspond to a reduction procedure which takes an instance and reduces it into instances with the special structures, which we call the base instances. Then we shall point out how to solve a base instance by referencing to the proofs, and highlight the major steps. The algorithm will be a pure recursive algorithm, i.e. no dynamic programming technique is required. To show that it is a polynomial time algorithm, we will argue that the number of recursive calls is polynomially bounded, and each call can be carried out in polynomial time.

The starting point of the proof is to reduce the main theorem (Theorem 3.3.2) of the STEINER FOREST PACKING problem to the extension theorem (Theorem 3.3.4) for the STEINER TREE PACKING problem. The relevant part of the proof is in Section 3.6. In terms of algorithmic procedure, we reduce an instance of the STEINER FOREST PACKING problem to instances where we need to find edge-disjoint Steiner trees which extend a specified edge-subpartition of a specified vertex. Now we describe how the reduction can be done. Let $S_1, \ldots, S_t$ be the terminal groups and $S^* := S_1 \cup \ldots \cup S_t$. Notice that edge-connectivities can be tested by algorithms for computing minimum cuts, which is well-known to be solvable in polynomial time. If S^* is Qk-edge-connected, then we can treat S^* as one terminal group and the problem reduces easily to the STEINER TREE PACKING problem, and we are done. Otherwise, we need to identify a core C, and apply the cut decomposition operation to obtain two graphs where $V(G_1) = C \cup \{v_1\}$ and $V(G_2) = (V(G) - C) \cup \{v_2\}$. Note that a group separating set $X \subseteq V(G)$ can be found in polynomial time by a minimum cut computation. If there is no group separating set in $G[X]$, then X is a core. So a core can be obtained by repeatedly finding a group separating set until all terminal groups in the resulting graph are Qk-edge-connected.

We then recursively solve the STEINER FOREST PACKING problem on G_2, and the edge-disjoint Steiner forests obtained in G_2 induce an edge-subpartition $\mathcal{P}_k(v_2)$ on v_2. This naturally defines an edge-subpartition $\mathcal{P}_k(v_1)$ on v_1. Then we use the extension theorem (which we shall explain below) to find edge-disjoint Steiner trees of G_1 that extend $\mathcal{P}_k(v_1)$. The proof of Theorem 3.3.2 asserts that the combination of the forests in G_2 and the trees in G_1 gives edge-disjoint Steiner forests in G. Notice that whenever we invoke a recursive call, the number of terminal groups in G_2 is less than the number of terminal groups in G. So the number of Steiner tree instances we reduce to (i.e. the number of recursive calls) is at most the number of terminal groups. Hence the procedure of reducing the STEINER FOREST PACKING problem to the extension theorem of the STEINER TREE PACKING problem is polynomial time solvable, assuming that the extension theorem can be solved in polynomial time.

Before we go into the details of the extension theorem, we mention briefly the use of the edge splitting-off technique. By Mader's result (Lemma 3.5.1), for a given Steiner vertex, there exists at least a pair of edges incident to it whose splitting-off preserves the edge-connectivity of the terminal vertices. This was used in the proof to assume that every Steiner vertex is of degree 3 in a minimal counterexample. In terms of algorithmic procedure, this allows us to reduce an input instance to an instance where every Steiner vertex is of degree 3. To implement the edge splitting-off technique, a straightforward procedure is to find a Steiner vertex of degree greater than 3, try every pair of edges incident to it until a suitable splitting-off is found, which can be verified by minimum cut computations. Clearly, this can be done in polynomial time. Notice that every time we split-off a pair of edges, the resulting graph has one fewer edge than the original graph. So we can reduce to an instance with every Steiner vertex of degree 3 in a polynomial number of steps.

Now we discuss the algorithm for the extension theorem of the STEINER TREE PACKING problem. The first key step is to transform the input graph G into instances with no edge between Steiner vertices (Lemma 3.7.1). Suppose there is an edge e between two Steiner vertices. If S is Qk-edge-connected in $G - e$, then we can simply remove e from G. Otherwise, the edge e must be in an edge-cutset T of size Qk whose removal disconnects two terminal vertices. We can identify this edge-cutset by a minimum cut computation as follows: If S is not Qk-edge-connected in $G - e$, then $C \cup \{e\}$ is the desired edge-cutset where C is a minimum S-edge-cutset of $G - e$. Let v be the vertex to be extended. Let C_1 and C_2 be the two components of $G - T$ where $v \in C_2$. Then we apply the cut decomposition operation to obtain two graphs, where $V(G_1) = C_1 + v_1$ and $V(G_2) = C_2 + v_2$. We recursively find edge-disjoint Steiner trees in G_2 which extend v; those trees define an edge-subpartition of v_1. Then we recursively find edge-disjoint Steiner trees in G_1 which extend v_1. The proof of Lemma 3.7.1 guarantees that the combination of the trees in these two graphs give the desired edge-disjoint Steiner trees for the original graph G. Suppose G has l edges between Steiner vertices. Then the total number of edges between Steiner vertices in G_1 and G_2 is at most $l - 1$. This is because the edge e has become an edge incident to v_1 in G_1 and an edge incident to v_2 in G_2, where v_1 and v_2 are terminal vertices in G_1 and G_2 respectively. The recursion of this procedure stops when there are no edges between Steiner vertices. Therefore, starting with a graph with l edges between Steiner vertices, the above recursion will yield at most $l + 1$ graphs with no edges between Steiner vertices. In particular, this implies that the number of recursive calls is polynomially bounded.

The next step is to transform the input graph G so that $S \cup R$ is $(Q-2)k$-edge-connected in G (Lemma 3.7.2). The strategy used is similar to the previous steps, and here we sketch the procedure and highlight the difference. If $S \cup R$ is not $(Q-2)k$-edge-connected in G, then we find an R-isolating set X in G using a minimum cut computation. We apply the cut decomposition operation on X to obtain two graphs, where $G_1 := X + v_1$ and $G_2 := (V(G) - X) + v_2$. We recursively find edge-disjoint Steiner trees on G_2, which has fewer edges than G. Here, instead of finding edge-disjoint Steiner trees on G_1, we simply find edge-disjoint paths between v_1 and a carefully chosen vertex $r^* \in R \cap X$. The vertex r^*, as shown in Lemma 3.7.2, can be chosen to be a vertex so that the total length of the edge-disjoint paths to v_1 is minimized, which can be computed in polynomial time by a minimum cost flow computation. Then, Lemma 3.7.2 asserts that the combination of the Steiner trees of G_2 and the paths of G_1 gives edge-disjoint Steiner trees of G with the desired properties.

A key tool is a theorem by Frank, Király and Kriesell (Theorem 3.2.1) to find edge-disjoint Steiner trees in graphs with no edges between Steiner vertices. This theorem can be implemented as a polynomial time algorithm. Some details can be found in the end of Section 2.2.6 and we will not repeat them here. Using this algorithm, one can easily see that the procedure, which corresponds to "Elimination of the Second Case" in the proof, can be implemented in polynomial time.

Finally, we have described all the reduction steps (i.e. recursions), and we arrive at the base instances for which no more recursion will be invoked. The relevant part of the proof starts with the title "Construction and Properties of G'" in Section 3.7. The first step is to construct $G' := G - v - W$ where W is the Steiner vertex adjacent to v. If $S \cup R - v$ is $6k$-edge-connected in G', then using the algorithm for Theorem 3.2.1, we can construct edge-disjoint Steiner trees with the desired properties (see Lemma 3.7.3). Otherwise, we find an edge-cutset Z with $|Z| < 6k$ by a minimum cut computation. Here the proof is divided into two cases; we only describe the first case since the second case is analogous. Let the components of $G' - Z$ be C_1 and C_2, and let $S_1 := C_1 \cap S$, $R_1 := C_1 \cap R$, $S_2 := C_2 \cap S$, and $R_2 := C_2 \cap R$. Lemma 3.7.7 asserts that $S_i \cup R_i$ is $6k$-edge-connected in C_i for $i \in \{1, 2\}$. This allows us to apply the algorithm for Theorem 3.2.1 to construct edge-disjoint Steiner trees in C_1 and C_2 respectively. Then, following the proof of Lemma 3.7.10 which involves only elementary operations, we can combine the Steiner trees in C_1, the Steiner trees in C_2, the edges incident to v, and the edge cutset Z to construct edge-disjoint Steiner trees in G with the desired properties. This shows that a base instance can be solved in polynomial time, and concludes the description of the algorithm for the STEINER FOREST PACKING problem.

3.10 Capacitated Version

In this section, we consider the Capacitated Steiner Forest Packing problem. This is a generalization of the Steiner Forest Packing problem, where every edge e has an integer capacity c_e which bounds the number of forests that can use e. The Steiner Forest Packing problem is the special case when $c_e = 1$ for every edge e.

A naive strategy to solve the capacitated version is to replace each edge e of G by c_e multiple edges and apply the algorithm for Theorem 3.1.1 on the resulting graph G'. However, this only gives a pseudo-polynomial time approximation algorithm for the Capacitated Steiner Forest Packing problem to G, as the running time depends upon the edge capacities which may not be polynomially related to the number of vertices of G.

The strategy we shall use is to first solve the Fractional Steiner Forest Packing problem, which is a relaxation of the Capacitated Steiner Forest Packing problem where each forest is allowed to have fractional value. We remark that even this fractional relaxation is NP-complete, as was proved in [49]. A 2-approximation algorithm for the Fractional Steiner Forest Packing problem, however, can be obtained by using exactly the same approach in [49] for the Fractional Steiner Tree Packing problem. In [49], an approximate optimal fractional solution for the Fractional Steiner Tree Packing problem is computed by using the ellipsoid algorithm for solving linear programs, with the approximation algorithm for the Minimum Steiner Tree problem as a subroutine (more precisely as a separation oracle). Similarly, an approximate optimal fractional solution for the Fractional Steiner Forest Packing problem can be computed by the ellipsoid method; we just need to replace the separation oracle by an approximation algorithm for the Minimum Steiner Forest problem (e.g. [42]). We remark that the factor 2 comes from the approximation ratio of the Minimum Steiner Forest problem; but as we shall see, this number is not important as long as it is not bigger than $Q = 30$. For more details, we refer the reader to [49].

Once we have obtained a 2-approximate fractional solution of the problem, we can obtain an integral solution by rounding down the fractional solution. Not surprisingly, this would not always be a Q-approximate integral solution for the problem. However, we will prove that for those instances on which this procedure fails, we can use the naive procedure mentioned earlier to solve the problem. That is, the size of the graph obtained by the naive procedure in those instances is at most a polynomial times the size of the original graph. This combines to give an approximation algorithm for the CAPACITATED STEINER FOREST PACKING problem, which will be proved formally in the following theorem.

Theorem 3.10.1 *There is a polynomial time algorithm for the* CAPACITATED STEINER FOREST PACKING *that constructs an integral solution of value at least* $\lfloor \frac{\tau}{Q} \rfloor$*, where* τ *is the value of an optimal integral solution.*

Proof. Given an instance of the CAPACITATED STEINER TREE PACKING problem, let τ^*, τ be the value of an optimal fractional, integral solution, respectively. We first use the approximation algorithm for the FRACTIONAL STEINER FOREST PACKING problem [49] to obtain a fractional solution of value β such that $2\beta \geq \tau^*$. One feature of that algorithm is that there are at most a polynomial number of forests in the fractional solution with $x_T > 0$, say $\{x_1, \ldots, x_{p(n)}\}$. Suppose $\sum_{i=1}^{p(n)} \lfloor x_i \rfloor \geq \frac{2}{Q} \sum_{i=1}^{p(n)} x_i$, then $\sum_{i=1}^{p(n)} \lfloor x_i \rfloor \geq \frac{2}{Q} \sum_{i=1}^{p(n)} x_i = \frac{2}{Q}\beta \geq \frac{1}{Q}\tau^* \geq \frac{1}{Q}\tau$. So, $\{\lfloor x_1 \rfloor, \ldots, \lfloor x_{p(n)} \rfloor\}$ is an integral solution which is at least $\frac{\tau}{Q}$, and we are done.

Otherwise, $\sum_{i=1}^{p(n)} x_i > \frac{Q}{2} \sum_{i=1}^{p(n)} \lfloor x_i \rfloor$. Then, $(\frac{Q}{2} - 1) \sum_{i=1}^{p(n)} \lfloor x_i \rfloor < \sum_{i=1}^{p(n)} (x_i - \lfloor x_i \rfloor) \leq p(n)$, which implies $\sum_{i=1}^{p(n)} \lfloor x_i \rfloor < \frac{2}{Q-2} p(n)$. So, $\beta = \sum_{i=1}^{p(n)} x_i = \sum_{i=1}^{p(n)} \lfloor x_i \rfloor + \sum_{i=1}^{p(n)} (x_i - \lfloor x_i \rfloor) < \frac{2}{Q-2} p(n) + p(n) = \frac{Q}{Q-2} p(n)$. Therefore, $\tau^* < \frac{2Q}{Q-2} p(n)$. Note that in any solution, at most a value of τ^* capacity is used in an edge. So, if $c_e > \tau^*$, the excess capacity $c_e - \tau^*$ will never be used. Now, to find an integral solution, we replace every edge e

of G by $\min\{c_e, \lfloor \tau^* \rfloor\}$ multiple edges and call the resulting graph G'. Notice that the total number of edges in G' is bounded by a polynomial of n and the value of an optimal solution in G' is the same as in G. So, we can apply the algorithm for Theorem 3.1.1 to obtain $\lfloor \frac{\tau}{Q} \rfloor$ edge-disjoint $\mathcal{S}$-forests of G' in polynomial time, which correspond to an integral solution of G with value at least $\lfloor \frac{\tau}{Q} \rfloor$. Therefore, in either case, the integral solution constructed is at least $\lfloor \frac{\tau}{Q} \rfloor$. ∎

3.11 Steiner Network Packing

The following is a general problem that captures the STEINER FOREST PACKING problem. Given an undirected multigraph G and a connectivity requirement r_{uv} for each pair of vertices $u, v \in V(G)$, find a largest collection of edge-disjoint subgraphs of G such that in each subgraph there are r_{uv} edge-disjoint paths from u to v for all $u, v \in V(G)$. Since edge-connectivity is transitive, Theorem 3.3.2 is equivalent to the following (see also [12]).

Theorem 3.11.1 *Given an undirected multigraph G and a connectivity requirement $r_{uv} \in \{0, 1\}$ for $u, v \in V(G)$. If there are $Qk \cdot r_{uv}$ edge-disjoint paths for all $u, v \in V(G)$, then there are k edge-disjoint forests such that in each forest there are r_{uv} paths between u, v for all $u, v \in V(G)$.*

Proof. Given an instance of the STEINER FOREST PACKING problem, by setting $r_{uv} = 1$ for all $u, v \in S_i$ for all i, the result of Theorem 3.1.1 follows. Now we prove the other direction. Consider an instance of the STEINER NETWORK PACKING problem with $r_{uv} \in \{0, 1\}$ for all $u, v \in V(G)$. We observe that if $r_{uv} = 1$ and $r_{vw} = 1$, then in the resulting Steiner network, we must also have a path between u and w. Hence we can set $r_{uw} = 1$ as well. By setting each equivalence class of size at least 2 to a terminal group, it is an instance of the STEINER FOREST PACKING problem. So the statement follows from Theorem 3.1.1. ∎

I conjecture that Theorem 3.11.1 can be generalized to arbitrary non-negative integer connectivity requirements:

Conjecture 3.11.2 *Given an undirected multigraph G and a connectivity requirement r_{uv} for each pair of vertices $u, v \in V(G)$. There exists a universal constant c so that the following holds. If there are $ck \cdot r_{u,v}$ edge-disjoint paths for all $u, v \in V(G)$, then there are k edge-disjoint subgraphs $H_1, \ldots, H_k$ in G such that in each subgraph there are r_{uv} edge-disjoint paths between u and v for all $u, v \in V(G)$.*

It would be interesting to first verify Conjecture 1 for Eulerian graphs; notice that $c = 2$ has not been ruled out. This may also yield insights into the generalized Steiner network problem, where the goal is to find a minimum cost subgraph which satisfies all the connectivity requirements, i.e. to find a minimum cost subgraph which has $r_{u,v}$ paths for all $u, v \in V$. For example, the problem of finding a minimum Steiner tree in a graph is a special case of the generalized Steiner network problem. The only known constant factor (factor 2) approximation algorithm for the generalized Steiner network problem is due to Jain's iterative rounding technique [47], and an algorithm of a more combinatorial nature is left as an open problem. The Steiner network packing problem is closely related to the generalized Steiner network problem. Intuitively, if one has an algorithm to find many edge-disjoint Steiner networks, then there should be some Steiner network in the collection which uses few edges (and hence low costs). This intuition can actually be realized using linear programming techniques. Very roughly, we could use the packing algorithm to decompose the linear program solution into (fractionally) disjoint Steiner networks, and then find the cheapest Steiner network among them. In fact, if Conjecture 3.11.2 is true for any constant c in Eulerian graphs, then this would imply a more combinatorial c-approximation algorithm for the generalized Steiner network problem.

Chapter 4

Steiner Orientations

The results in this chapter are based on joint work with Tamás Király [54].

4.1 Introduction

Let $H = (V, E)$ be an undirected hypergraph. Recall that an orientation of H is obtained by assigning a direction to each hyperedge in H. In this chapter, we use the out-hypergraph model defined in Section 2.1.2, where a hyperarc (a directed hyperedge) is a hyperedge with a designated tail vertex and other vertices as head vertices. 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. Recall that a path in a directed hypergraph is an alternating sequence of distinct vertices and hyperarcs $\{v_0, a_0, v_1, a_1, \ldots, a_{k-1}, v_k\}$ 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 terminal vertices (i.e. $S = \{r, v\}$), it is the edge-disjoint paths problem solved by Menger [73]. To see this, given k edge-disjoint paths between r, v in G, one can easily construct a Steiner rooted k-arc-connected orientation D of G by orienting each r, v-path away from r (the remaining edges can be oriented arbitrarily). For the other direction, suppose we have a Steiner rooted k-arc-connected orientation D of G, then there are k arc-disjoint paths from r to v in D. These paths are k edge-disjoint r, v-paths in G by ignoring the directions. When all vertices in the graph are terminals (i.e. $S = V(G)$), the STEINER ROOTED-ORIENTATION problem can be shown to be equivalent to the edge-disjoint spanning trees problem solved by Tutte [85] and independently later by Nash-Williams [75]. To see this, given k edge-disjoint spanning trees in G, one can easily construct a Steiner rooted k-arc-connected orientation D of G by orienting each spanning tree as an r-arborescence (the remaining edges can be oriented arbitrarily). For the other direction, suppose we have a Steiner rooted k-arc-connected orientation D of G, then by Edmonds' theorem (Theorem 2.2.24) there are k arc-disjoint r-arborescences in D. These r-arborescences are k edge-disjoint spanning trees in G by ignoring the directions.

The STEINER ROOTED-ORIENTATION problem is a common generalization of the above problems. An alternative common generalization of the above problems is the STEINER TREE PACKING problem studied in Chapter 3. It is instructive to compare these two problems on the setting of graphs. 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. For example, the underlying graph of Figure 1.1 has a Steiner rooted 2 arc-connected orientation, but not 2 edge-disjoint Steiner trees. 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. This has implications in the network multicasting problem, which will be discussed later.

Given a hypergraph H, recall that 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.

4.1.1 Previous Work

Graph orientations is a very well-studied subject in the literature, and there are many ways to look at such questions (see [5]). Here we focus on graph orientations achieving high connectivity (see Section 2.2.4 and Section 2.2.5 for more details). The starting point of this line of research is a theorem by Robbins [81] in 1939 which says that a graph G has a strongly-connected orientation if and only if G is 2-edge-connected. In 1960, Nash-Williams [74] generalized this theorem by proving that a graph G has a strongly k-arc-connected orientation if and only if G is $2k$-edge-connected; this theorem is often called the weak orientation theorem in the literature (Theorem 2.2.12). Recall that $\lambda(x, y)$ denotes the maximum number of edge-disjoint (or arc-disjoint) paths from x to y, which is called the local-edge-connectivity (or local-arc-connectivity) from x to y. Nash-Williams [74] also proved the following deep theorem (Theorem 2.2.13) which achieves optimal local-arc-connectivity for all pair of vertices:

(NASH-WILLIAMS' STRONG ORIENTATION THEOREM)

Every undirected graph G has an orientation D so that

$$\lambda_D(x, y) \geq \lfloor \lambda_G(x, y)/2 \rfloor \text{ for all } x, y \in V.$$

Nash-Williams' original proof uses a complicated inductive argument (see Section 2.2.4); until now this is the only known orientation result achieving high *local*-arc-connectivity.

Subsequently, Frank, in a series of works [22, 23, 25, 27, 30], developed a general framework to solve graph orientation problems achieving high *global*-arc-connectivity (see Section 2.2.5) by using the submodular flow problem introduced by Edmonds and Giles [18] (see Section 2.2.5). With this powerful tool, Frank greatly extended the range of orientation problems that can be solved concerning global-arc-connectivity. Some representative examples include finding a strongly k-arc-connected orientation with minimum weight [23], with in-degree constraints [22] and in mixed graphs [27]. Recently, this framework has been generalized to solve hypergraph orientation problems achieving high global-hyperarc-connectivity [33].

Extending graph orientation results to hyperarc-connectivity or to vertex-connectivity is more challenging. For the Steiner Rooted-Orientation problem, the only known result follows from Nash-Williams' strong orientation theorem: if S is $2k$-edge-connected in an undirected graph G, then G has a Steiner rooted k-arc-connected orientation. For hypergraphs, there is no known orientation result concerning Steiner rooted-hyperarc-connectivity. A closely related problem of characterizing hypergraphs that have an orientation which is Steiner strongly k-hyperarc-connected is posted as an open problem in [19] (and more generally an analog of Nash-Williams' strong orientation theorem in hypergraphs). For orientation results concerning vertex-connectivity, very little is known even for global rooted-vertex-connectivity (i.e. when there is no Steiner vertices). Frank [29] made a conjecture on a necessary and sufficient condition for the existence of a strongly k-vertex-connected orientation, which in particular would imply that a $2k$-vertex-connected graph has a strongly k-vertex-connected orientation (and hence a rooted k-vertex-connected orientation). The only positive result along this line is a sufficient condition due to Jordán [48] for the case $k = 2$: Every 18-vertex-connected graph has a strongly 2-vertex-connected orientation.

4.1.2 Results

The main result of this chapter is the following theorem on hypergraphs, which is tight in terms of the connectivity bound. This gives a positive answer to the rooted version of the open question in [19].

Theorem 4.1.1 *Suppose H is a hypergraph, S is a given subset of terminal vertices and $r \in S$ is the root vertex. Then H has a Steiner rooted k-hyperarc-connected orientation if S is $2k$-hyperedge-connected in H.*

We remark that no analogous result can be obtained for Steiner strongly k-hyperarc-connected orientations: for every constant C, there are hypergraphs which are Ck-hyperedge-connected but do not have a Steiner strongly k-hyperarc-connected orientation. This contrasts with the graph case where the Nash-Williams strong orientation theorem implies that: if S is $2k$-edge-connected in a graph G, then G has a strongly k-edge-connected orientation. We shall also give an alternative proof of this result in Section 4.6.

Theorem 4.1.1 is best possible in terms of the connectivity bound. For example, a cycle is 2-edge-connected but only has a rooted 1-arc-connected orientation. More generally, this is best possible for every k as shown by any $2k$-regular $2k$-edge-connected non-complete graph G by setting $S = V(G)$ (e.g. a $2k$-dimensional hypercube). To see this, any $2k$-regular graph has kn edges but a rooted $(k+1)$-arc-connected orientation requires at least $(k+1)(n-1)$ edges. So there are just not enough edges to have a rooted $(k+1)$-arc-connected orientation if G is not a complete graph (i.e. $n > k-1$).

The proof of Theorem 4.1.1 is constructive and so also implies the first 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's strong 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 4.7.1). This contrasts with the polynomial time solvable problem of finding an orientation which maximizes the Steiner strong arc-connectivity of a graph; this is implied by the Nash-Williams strong orientation theorem for which a polynomial time algorithm exists. It is a rare phenomenon that the rooted version of a connectivity problem is more difficult than the non-rooted one.

Following the standard notation on approximation algorithms for 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 [46, 21]). 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.

Theorem 4.1.2 *Suppose G is a graph, S is a given 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.*

This is optimal in terms of the connectivity bound as the tight examples for Theorem 4.1.1 show. The proof is constructive and so also implies the first polynomial time approximation algorithm for the problem. Finally, we also prove that this problem is NP-complete (Theorem 4.7.5).

4.1.3 Techniques

Since Nash-Williams strong orientation theorem, little progress has been made on the orientation problems concerning local-arc-connectivity, local-hyperarc-connectivity or vertex-connectivity. The difficulty is largely due to a lack of techniques to work with these more sophisticated connectivity notions. The main technical contribution of this chapter is a new method to use the submodular flow problem. A key ingredient in the proof of Theorem 4.1.1 is the use of an "extension property" (see Chapter 3) to help decompose a general hypergraph into hypergraphs with substantially simpler structures. Then, in those simpler hypergraphs, we apply the submodular flows technique in a very 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 4.1.3 *Suppose G is a graph, S is a given 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 in-degrees (so that the indegree of each Steiner vertex v is equal to $m(v)$) can be solved in polynomial time.*

Perhaps Theorem 4.1.3 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 4.1.3, we shall reduce this problem to a submodular flow problem from which we can also derive a necessary and sufficient condition for the existence of a Steiner rooted k-arc-connected orientation. This provides us with a crucial tool in establishing the connectivity bounds. As we shall see, Theorem 4.1.3 is a generalization of the results on the hypergraph global rooted-arc-orientation problems in [33, 53] and [6].

Interestingly, the proof of Theorem 4.1.2 is also based on the DEGREE-SPECIFIED STEINER ROOTED-ORIENTATION problem (Theorem 4.1.3) which was designed for edge-connectivity problems. For a similar step in the hypergraph orientation problem, we borrow a technique in [14] to obtain a graph with simpler structures.

4.1.4 The Network Multicasting Problem

The Steiner Rooted-Orientation problem is motivated by the multicasting problem in computer networks (see also Chapter 1), where the root vertex (the sender) must transmit all its data to the terminal vertices (the receivers) and the goal is to maximize the transmission rate that can be achieved simultaneously for all receivers. The connection is through a recent beautiful min-max theorem by Ahlswede et. al. [2] in 2000:

> (The Network Coding Theorem)
> Given a directed multigraph with unit capacity on each arc, if there are k arc-disjoint paths from the root vertex to each terminal vertex, then the root vertex can transmit k unit of data to all terminal vertices simultaneously.

They prove the theorem by introducing the innovative idea of network coding [2, 65], which has generated much interest in different areas from information theory to computer science [2, 65, 16, 63, 64, 68, 69, 70]. These studies focus on directed networks, for example the Internet, where the direction of data movement on each link is fixed a priori. On the other hand, there are practical networks which are *undirected*, i.e. data can be sent in either direction along a link. By using the above theorem, computing the maximum multicasting rate in undirected networks (with network coding supported) reduces to the Steiner Rooted-Orientation problem. This has been studied in the graph model [63, 64] and efficient approximation algorithms have been proposed. An important example of undirected networks is wireless networks (equipped with omni-directional antennas), for which many papers have studied the advantages of incorporating network coding (see [68] and the references therein). However, there are some aspects of wireless communications that are not captured by a graph model. One distinction is that wireless communications in such networks are inherently one-to-many instead of one-to-one. This motivates researchers to use a directed hypergraph model [16, 68, 69, 70] to study the multicasting problem (with network coding supported) in wireless networks. The directed hypergraph model used in [16, 68, 69, 70] coincides with the out-hypergraph model defined in Section 2.1.2. A simple reduction shows that the above network coding theorem by Ahlswede el. al. applies to directed hypergraphs as well. Therefore, computing the maximum multicasting rate in an undirected hypergraph (with network coding supported) reduces to the Steiner Rooted-Orientation problem of hypergraphs.

In the multicasting problem, the Steiner Tree Packing is used to transmit data when network coding is not supported. However, one cannot hope for analogous results of Theorem 4.1.1 or Theorem 4.1.2 for the corresponding Steiner Tree Packing problems. In fact, both the hyperedge-disjoint Steiner tree packing problem and the element-disjoint Steiner tree packing problem are shown to be NP-hard to approximate within a factor of $\Theta(\log n)$ [14]. (It was also shown in [7] that no constant connectivity bound implies the existence of two hyperedge-disjoint spanning sub-hypergraphs.) So, Theorem 4.1.1 indicates that multicasting with network coding in a hypergraph model could be much more efficient in terms of the throughput. To be more precise, if S is $2k$-hyperedge-connected in a hypergraph H, then Theorem 4.1.1 implies that we can transmit at least k units of data from the root to each receiver in S (assuming each edge can transmit one unit of data at a time) if network coding is supported. On the other hand, if network coding is not supported, then each unit of data needs to be transmitted through a hyperedge-disjoint subgraph which connects S. However, there are hypergraphs (e.g. the examples in [7]) for which S is $2k$-hyperedge-connected in H but H does not have more than $\Theta(k/\log n)$ hyperedge-disjoint subgraphs which connect S. Hence, in those hypergraphs, we can transmit at most $\Theta(k/\log n)$ unit of data from the root to each receiver in S (assuming each edge can transmit one unit of data at a time) if network coding is not supported.

4.2 Degree-Specified Steiner Orientations

In this section we consider the Degree-Specified Steiner Orientation problem, suggested by Frank (personal communication). Given a graph $G = (V, E)$, a terminal set $S \subseteq V(G)$ and a connectivity requirement function $h : 2^S \to \mathbb{Z}$, the connectivity requirement function $h^* : 2^V \to \mathbb{Z}$ 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 $d_D^{in}(v) = m(v)$ for every $v \in V(G) - S$.

Given a hypergraph $H = (V, \mathcal{E})$ and a connectivity requirement function $h : 2^V \to \mathbb{Z}$, the Hypergraph Orientation problem is to find an orientation $\overrightarrow{H}$ of H that covers h. Bang-Jensen and Thomassé [6] and Frank, Király and Király [33] studied the Hypergraph Orientation problem for out-hypergraphs and in-hypergraphs respectively. We now show that both are special cases of the Degree-Specified Steiner Orientation problem.

Given a hypergraph $H = (V, \mathcal{E})$, we consider the bipartite representation $B = (V, \mathcal{E}; E)$ of H. The vertices in the vertex partite set V are terminal vertices; while the vertices in the hyperedge partite set $\mathcal{E}$ are Steiner vertices. To define a connectivity requirement function on every subset of $V(B) = V \cup \mathcal{E}$, we use the Steiner extension h^* of h, i.e. $h^*(X)$ for $X \subseteq V(B)$ is defined to be $h(X \cap V)$. For out-hypergraphs, we specify the indegree of each Steiner vertex (which corresponds to a hyperedge) to be 1. Let D be an orientation of B with the specified indegrees. Since each Steiner vertex has indegree 1, D corresponds to an orientation $\overrightarrow{H}$ of H as follows. For each hyperedge vertex v in B, if uv is oriented as $\overrightarrow{uv}$ in D, then the hyperedge corresponds to v in H is oriented with u as the tail in $\overrightarrow{H}$. Also, since each Steiner vertex has indegree 1 in D, arc-disjoint paths in D corresponds to hyperarc-disjoint paths in $\overrightarrow{H}$. Therefore, there is a hypergraph orientation $\overrightarrow{H}$ of H that covers h if and only if there is an orientation D of B that covers the h^* with the specified indegrees. For example, there is a rooted k-hyperarc-connected orientation of H if and only if there is a Steiner rooted k-arc-connected orientation of B with the specified indegrees. The HYPERGRAPH ORIENTATION problem for in-hypergraphs can be reduced to the DEGREE-SPECIFIED STEINER ORIENTATION problem analogously.

4.2.1 Degree-Specified Orientations Covering Steiner Extensions of Intersecting Supermodular Set Functions

Given an undirected graph G and $S \subseteq V(G)$, an intersecting supermodular function $h : 2^S \to \mathbb{Z}$ and an in-degree specification $m : (V(G) - S) \to \mathbb{Z}^+$, we ask whether there exists an orientation D of G that covers the Steiner extension h^* of h with $d_D^{in}(v) = m(v)$ for each $v \in V(G) - S$. We reduce this problem to a submodular flow problem (see Section 2.2.5 for more details). In the following $d(X, Y)$ denotes the number of edges with one endpoint in X and the other endpoint in Y, and $E(Y)$ denotes the number of edges with both endpoints in Y. Also, we let $m(Y) := \sum_{v \in Y} m(v)$ for $Y \subseteq V(G) - S$.

Notice that h^* is not an intersecting supermodular function in general (e.g. inequality (2.2) does not hold for two sets X, Y with $X \cap Y \subseteq V(G) - S$), and therefore Theorem 2.2.18 cannot be directly applied. The difficulty is that we have no information about the indegrees of sets containing only Steiner vertices. To overcome this obstacle, we are motivated to add extra information on the Steiner vertices. It turns out that with the indegree specification m, we can define another connectivity requirement function $h' : 2^V \to \mathbb{Z} \cup \{-\infty\}$ which "encodes" the DEGREE-SPECIFIED STEINER ORIENTATION problem and such that h' is intersecting supermodular if h is intersecting supermodular.

$$h'(X) := \begin{cases} m(X) & \text{for } X = \{v\}, v \in V - S. \\ h(X) + \max_{Y \subseteq V(G)-S}\{d(X,Y) - m(Y) + E(Y)\} & \text{for } X \subseteq S, X \neq \emptyset. \\ 0 & \text{if } X = \emptyset \text{ or } X = V. \\ -\infty & \text{otherwise.} \end{cases}$$

The following lemma shows that the newly defined function h' is an intersecting supermodular function if h is intersecting supermodular.

Lemma 4.2.1 *If h is intersecting supermodular, then h' is an intersecting supermodular function.*

Proof. Let X_1 and X_2 be two intersecting sets. We need to show that $h'(X_1)+h'(X_2) \leq h'(X_1 \cap X_2) + h'(X_1 \cup X_2)$. If either $h'(X_1)$ or $h'(X_2)$ is $-\infty$, then the inequality holds trivially. Also, if $X_1 = \emptyset$ or $X_1 = V(G)$, then the inequality holds (as equality). So we just need to consider the case that both X_1 and X_2 are in S. Let X_1 and X_2 be two intersecting sets in S, and Y_1 and Y_2 be the sets in $V(G) - S$ which yield the maximum of $h'(X_1)$ and $h'(X_2)$ respectively. Then,

$$\begin{aligned}
& h'(X_1) + h'(X_2) \\
= \;& h(X_1) + h(X_2) - m(Y_1) - m(Y_2) + E(Y_1) + E(Y_2) + d(X_1, Y_1) + d(X_2, Y_2) \\
\leq \;& h(X_1 \cap X_2) + h(X_1 \cup X_2) - m(Y_1) - m(Y_2) + \\
& \quad E(Y_1) + E(Y_2) + d(X_1, Y_1) + d(X_2, Y_2) \\
= \;& h(X_1 \cap X_2) + h(X_1 \cup X_2) - m(Y_1 \cap Y_2) - m(Y_1 \cup Y_2) + \\
& \quad E(Y_1) + E(Y_2) + d(X_1, Y_1) + d(X_2, Y_2) \\
\leq \;& h(X_1 \cap X_2) + h(X_1 \cup X_2) - m(Y_1 \cap Y_2) - m(Y_1 \cup Y_2) + \\
& \quad E(Y_1 \cap Y_2) + E(Y_1 \cup Y_2) + d(X_1, Y_1) + d(X_2, Y_2) \\
\leq \;& h(X_1 \cap X_2) + h(X_1 \cup X_2) - m(Y_1 \cap Y_2) - m(Y_1 \cup Y_2) + \\
& \quad E(Y_1 \cap Y_2) + E(Y_1 \cup Y_2) + d(X_1 \cap X_2, Y_1 \cap Y_2) + d(X_1 \cup X_2, Y_1 \cup Y_2) \\
= \;& h(X_1 \cap X_2) - m(Y_1 \cap Y_2) + E(Y_1 \cap Y_2) + d(X_1 \cap X_2, Y_1 \cap Y_2) \\
& \quad h(X_1 \cup X_2) - m(Y_1 \cup Y_2) + E(Y_1 \cup Y_2) + d(X_1 \cup X_2, Y_1 \cup Y_2) \\
\leq \;& h(X_1 \cap X_2) + \max_{Y \subseteq V(G)-S} \{d(X_1 \cap X_2, Y) - m(Y) + E(Y)\} \\
& \quad h(X_1 \cup X_2) + \max_{Y \subseteq V(G)-S} \{d(X_1 \cup X_2, Y) - m(Y) + E(Y)\} \\
= \;& h'(X_1 \cap X_2) + h'(X_1 \cup X_2).
\end{aligned}$$

The inequalities are based on the fact that $h(.)$, $E(.)$ and $d(.,.)$ are intersecting supermodular function. ∎

The following lemma shows that h' "encodes" the DEGREE-SPECIFIED STEINER ORIENTATION problem.

Lemma 4.2.2 *Suppose $E(Y) \leq m(Y)$ for every $Y \subseteq V(G) - S$. G has an orientation covering h' if and only if G has an orientation covering the Steiner extension h^* of h with the indegree of each $v \in V(G) - S$ being exactly $m(v)$.*

Proof. Suppose there exists an orientation D of G that covers h'; thus $d_D^{in}(Z) \geq h'(Z)$ for every $Z \subseteq V(G)$. For each $v \in V(G)-S$, $h'(v) = m(v)$, and so $d_D^{in}(v) \geq h'(v) = m(v)$. If $d_D^{in}(v) = m(v)$ for all $v \in V(G)-S$, then for all $Y \subseteq V(G)-S$, $d_D^{in}(Y) = \sum_{v \in Y} d_D^{in}(v) - E(Y) = m(Y) - E(Y)$. Hence, for every $X \subseteq S$ and for every $Y \subseteq V(G) - S$, we have:

$$\begin{aligned}
& d_D^{in}(X \cup Y) \\
= \;& d_D^{in}(X) + d_D^{in}(Y) - d_G(X,Y) \\
\geq \;& h'(X) + d_D^{in}(Y) - d_G(X,Y) \\
= \;& h'(X) + m(Y) - E(Y) - d_G(X,Y) \\
\geq \;& h(X) + \max_{Y' \subseteq V(G)-S} \{d_G(X,Y') - m(Y') + E(Y')\} - (d_G(X,Y) - m(Y) + E(Y)) \\
\geq \;& h(X).
\end{aligned}$$

This implies that $d_D^{in}(Z) \geq h(Z \cap S)$ for every $Z \subseteq V$. Therefore, if $d_D^{in}(v) = m(v)$ for any $v \in V(G)-S$, then the orientation D that covers h' is an orientation that covers the Steiner extension h^* of h with the indegree of each $v \in V(G) - S$ being exactly $m(v)$.

Suppose there exists a vertex v with $d_D^{in}(v) > m(v)$. Let U be the set of vertices which can reach v by a directed path (i.e. a vertex u can reach v if there is a directed path from u to v). Assume, by way of contradiction, that $U \cap S = \emptyset$. Then, by the definition of U, we have $d_D^{in}(U) = 0$. Therefore, since $v \in U$, we have $E(U) = \sum_{u \in U} d_D^{in}(u) > m(U)$, which contradicts our hypothesis. So we must have $U \cap S \neq \emptyset$ and we pick a vertex

$w \in S$ which is closest to v. Let P_w be a shortest directed path from w to v. Now we reverse the direction of all the arcs in P_w and obtain a new orientation D'. By doing so, $d^{in}_{D'}(u) = d^{in}_{D}(u) \geq m(u) = h'(u)$ for each $u \in V(G) - S - v$; and for v, $d^{in}_{D'}(v) = d^{in}_{D}(v) - 1 \geq m(v) + 1 - 1 = h'(v)$. Hence $d^{in}_{D'}(u) \geq h'(u)$ for every $u \in V(G) - S$. Also, by reorienting P_w, we still have $d^{in}_{D'}(X) \geq d^{in}_{D}(X) \geq h'(X)$ for every $X \subseteq S$. By repeating the above procedure, we will eventually obtain an orientation that covers h' with each $v \in V(G) - S$ having indegree exactly $m(v)$. Hence this reduces to the case in the previous paragraph; this proves one direction.

Now we prove the easier direction. Suppose we have an orientation D that covers the Steiner extension h^* of h such that $d^{in}_{D}(v) = m(v)$ for each $v \in V(G) - S$. For each $X \subset S$, recalling that $h'(X) := h(X) + \max_{Y \subseteq V(G)-S}\{d(X,Y) - m(Y) + E(Y)\}$, we let $Y^* \subseteq V(G) - S$ be the set which yields the maximum for $h'(X)$. Notice that Y^* could be the emptyset, and so $h'(X) \geq h(X)$ for each $X \subset S$. Also, recall that since $d^{in}_{D}(v) = m(v)$ for each $v \in V(G) - S$, the arguments from the first part of this proof imply that $d^{in}_{D}(Y) = m(Y) - E(Y)$ for each $Y \subseteq V(G) - S$. Since D covers h^*, $d^{in}_{D}(X \cup Y^*) \geq h^*(X \cup Y^*) = h((X \cup Y^*) \cap S) = h(X)$. So,

$$
\begin{aligned}
& d^{in}_{D}(X) \\
= \; & d^{in}_{D}(X \cup Y^*) - d^{in}_{D}(Y^*) + d_G(X, Y^*) \\
= \; & d^{in}_{D}(X \cup Y^*) - m(Y^*) + E(Y^*) + d_G(X, Y^*) \\
\geq \; & h(X) - m(Y^*) + E(Y^*) + d_G(X, Y^*) \\
= \; & h(X) + \max_{Y \subseteq V(G)-S}\{d_G(X,Y) - m(Y) + E(Y)\} \\
= \; & h'(X).
\end{aligned}
$$

Therefore, D is an orientation that covers h'. This completes the proof. ∎

Now, combining the two lemmas and also Theorem 2.2.18, we have the following theorem.

Theorem 4.2.3 *The* DEGREE-SPECIFIED STEINER ORIENTATION *problem can be solved in polynomial time.*

Proof. Lemma 4.2.1 shows that h' is indeed an intersecting supermodular function. Suppose $E(Y) > m(Y)$ for some $Y \subseteq V(G) - S$. Then such an indegree-specified orientation does not exist, and we say "No". Otherwise, Lemma 4.2.2 shows that the DEGREE-SPECIFIED STEINER ORIENTATION problem is equivalent to finding an orientation that covers h'. Therefore, Theorem 2.2.18 implies the theorem. ∎

Now we derive Theorem 4.1.3 as a corollary of Theorem 4.2.3.

Proof of Theorem 4.1.3: 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 (recall that if $h(X) = k$ for every $\emptyset \neq X \subset S$ then h is only a crossing supermodular function but not intersecting supermodular). By Theorem 4.2.3, the problem of finding an orientation D that covers the Steiner extension h^* of h with the specified indegrees can be solved in polynomial time.

We shall show that this is equivalent to finding a Steiner rooted k-arc-connected orientation with the specified indegrees. Let $s \in S$ be a terminal vertex. An orientation D that covers h^* satisfies $d_D^{in}(X) \geq k$ for every $\overline{r}s$ set. So, by Menger's theorem (Theorem 2.2.1), there are k arc-disjoint paths from r to s in D. Notice this holds for an arbitrary terminal vertex, and thus D is a Steiner rooted k-arc-connected orientation. For the other direction, a Steiner rooted k-arc-connected orientation D of G has k arc-disjoint paths from the root to each terminal vertex. So $d_D^{in}(X) \geq k$ for every $X \subseteq V(G)$ which contains a terminal vertex but not the root. Hence D covers h^*. Therefore the two problems are equivalent, and thus the DEGREE-SPECIFIED STEINER ROOTED-ORIENTATION problem can be solved in polynomial time. ∎

The following min-max formula, which is derived from Theorem 2.2.18, will be very useful later in establishing connectivity bounds.

Theorem 4.2.4 *Let $G = (V, E)$ be an undirected graph with a terminal set $S \subseteq V(G)$. Let $h : 2^S \rightarrow \mathbb{Z}$ be an intersecting supermodular function and $m : (V(G) - S) \rightarrow \mathbb{Z}^+$ be an indegree specification. Let $e_{\mathcal{P}}$ be the number of edges which enter some X_i. Then G has an orientation covering the Steiner extension h^* of h with the specified indegrees if and only if*

$$e_{\mathcal{P}} \geq \sum_{i=1}^{t} h'(X_i)$$

holds for every subpartition $\mathcal{P} = \{X_1, X_2, \ldots, X_t\}$ of V.

4.3 Steiner Rooted Orientations of Graphs

In this section we study the STEINER ROOTED ORIENTATION problem on graphs. The DEGREE-SPECIFIED STEINER ROOTED ORIENTATION problem is shown to be polynomial time solvable in the previous section. The STEINER ROOTED ORIENTATION problem, however, is NP-complete as we shall see in Section 4.7. That is, in general, finding an in-degree specification for the Steiner vertices to maximize the Steiner rooted-edge-connectivity is hard. However, in some occasions when we have a good idea what the indegrees of Steiner vertices should be like, the DEGREE-SPECIFIED STEINER ROOTED ORIENTATION problem turns out to be a powerful way to tackle the STEINER ROOTED ORIENTATION problem.

As mentioned in the introduction, Nash-Williams' strong orientation theorem (Theorem 2.2.13) implies the following result: if S is $2k$-edge-connected in an undirected graph G, then G has a Steiner rooted k-arc-connected orientation. As a warm-up, we give the following well-known proof of the above result on Eulerian graphs.

Theorem 4.3.1 *Let $G = (V, E)$ be an undirected graph with terminal set $S \subseteq V(G)$. Suppose S is $2k$-edge-connected in G and every Steiner vertex is of even degree. Then G has a Steiner rooted k-arc-connected orientation.*

Proof. Let v be a Steiner vertex. By Lemma 3.5.1, there is a suitable splitting-off operation at v so that S is still $2k$-edge-connected in the resulting graph. After the splitting-off operation, v is still of even degree and so we can repeat this procedure until there is no edge incident to v. Applying this procedure for every Steiner vertex gives a new graph G' such that $V(G') = S$ and G' is $2k$-edge-connected. Also, given a rooted k-arc-connected orientation of G', it is easy to construct a rooted k-arc-connected orientation of G. Now, we apply the Tutte-Nash-Williams theorem (Theorem 2.2.19) to show that G' has a rooted k-edge-connected orientation. Consider any vertex partition $\mathcal{P} = (V_1, \ldots, V_t)$ of $V(G')$. Since G' is $2k$-edge-connected, we have $d_{G'}(V_i) \geq 2k$. So, $e_{\mathcal{P}} = \frac{1}{2}\sum_{i=1}^{t} d_{G'}(V_i) \geq kt$. And the Tutte-Nash-Williams theorem implies that G' has at least k edge-disjoint spanning trees. So, by making each spanning tree into an r-arborescence, G' has a rooted k-arc-connected orientation. This proves the theorem. ∎

In the above proof, if a Steiner vertex is of odd degree, then we may not be able to find a suitable splitting-off operation when it is reduced to degree 3. This difficulty motivates our next theorem, which gives a min-max formula for the problem when every Steiner vertex is of degree 3 and there is no edge between Steiner vertices. This is one example that the Degree-Specified Steiner Rooted-Orientation problem can be applied. The following lemma shows that we can "hardwire" the indegrees of the Steiner vertices to be 1 without loss of generality.

Lemma 4.3.2 *Let $G = (V, E)$ be an undirected graph with terminal set $S \subseteq V(G)$. Suppose that every Steiner vertex of G is of degree 3 and there is no edge between two Steiner vertices in G. If G has a Steiner rooted k-arc-connected orientation D, then there is a Steiner rooted k-arc-connected orientation D' with $d^{in}_{D'}(v) = 1$ for each Steiner vertex v.*

Proof. Consider a vertex $v \in V(G) - S$. If $d^{in}_{D}(v) = 0$ or $d^{in}_{D}(v) = 3$, then v plays no role in connecting the root to any terminal, so we can simply reorient one or two of its incident arcs without affecting the rooted-arc-connectivity.

The interesting case is when $d_D^{in}(v) = 2$. Let u, w be the two incoming neighbours of v, and x be the outgoing neighbour of v, and r be the root. We claim that either $D - uv$ or $D - wv$ is still a Steiner rooted k-arc-connected orientation. Suppose not. Since $D - uv$ is not a Steiner rooted k-arc-connected orientation but D is, there is a terminal vertex s so that there are k arc-disjoint paths from r to s in D but only $k - 1$ arc-disjoint paths from r to s in $D - uv$. So, by Menger's theorem (Theorem 2.2.1), there is a set X_1 with $r \notin X_1$, $X_1 \cap S \neq \emptyset$ (in particular $s \in X_1$), $d_D^{in}(X_1) = k$ and $uv \in \delta_D^{in}(X_1)$. Actually x must be in X_1; otherwise $X_1 - v$ is a set with $r \notin X_1 - v$, $(X_1 - v) \cap S \neq \emptyset$, and $d_D^{in}(X_1 - v) < k$, which contradicts the assumption that D is a Steiner rooted k-arc-connected orientation. By a similar reasoning, w must be in X_1; otherwise $X_1 - v$ will violate the assumption that D is a Steiner rooted k-arc-connected orientation. So we obtain a set X_1 with $r, u \notin X_1$, $w, x, v \in X_1$, and $d_D^{in}(X_1) = k$. By the same argument with the roles of u and w switched, we obtain a set X_2 with $r, w \notin X_2$, $u, x, v \in X_2$, $d_D^{in}(X_2) = k$. Now, by submodularity (Proposition 2.1.2), we have

$$k + k = d_D^{in}(X_1) + d_D^{in}(X_2) \geq d_D^{in}(X_1 \cup X_2) + d_D^{in}(X_1 \cap X_2) \geq k + k;$$

the last inequality holds because both $X_1 \cup X_2$ and $X_1 \cap X_2$ contain x but not r (recall that since there is no edge between two Steiner vertices, $x \in S$). So equalities hold throughout and we have $d_D^{in}(X_1 \cap X_2) = k$. Notice that $v, x \in X_1 \cap X_2$ and $u, w, r \notin X_1 \cap X_2$. Therefore, $X_1 \cap X_2 - v$ has the properties that $r \notin X_1 \cap X_2 - v$, $x \in X_1 \cap X_2 - v$ and $d_D^{in}(X_1 \cap X_2 - v) = k - 1$, which, as described above, contradicts the assumption that D is a Steiner rooted k-arc-connected orientation. So, either $D - uv$ or $D - wv$ is still a Steiner rooted k-arc-connected orientation. Then reorienting uv or wv will not affect the rooted k-arc-connectivity and indegree of v becomes 1. Hence we can assume every Steiner vertex has indegree precisely 1. ∎

Lemma 4.3.2 allows us to use the result for the DEGREE-SPECIFIED STEINER ROOTED ORIENTATION problem. We remark the following theorem is implicit in [6]. Note that the STEINER TREE PACKING problem in this special case remains NP-complete (see Section 3.8). We also remark that the following theorem does not follow from Nash-Williams' strong orientation; the condition given below is more general than the conditions given by Nash-Williams.

Theorem 4.3.3 *Let $G = (V, E)$ be an undirected graph with terminal set $S \subseteq V(G)$. If every Steiner vertex (vertices in $V(G) - S$) 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*

$$e_{\mathcal{P}} \geq k(t-1) \tag{4.1}$$

for every partition $\mathcal{P} = (V_1, \ldots, V_t)$ of $V(G)$ such that each V_i contains a terminal vertex, where $e_{\mathcal{P}}$ denotes the number of crossing edges.

Proof. The plan is to reduce this problem to the DEGREE-SPECIFIED STEINER ROOTED ORIENTATION problem. We first argue that the condition (4.1) is necessary. For every partition $\mathcal{P} = \{V_1, \ldots, V_t\}$ such that each V_i contains a terminal vertex, let V_1 be the class that contains the root. Then, if there exists a Steiner rooted k-arc-connected orientation, each V_i for $i \neq 1$ must have at least k incoming edges. Hence there must be at least $k(t-1)$ crossing edges in $\mathcal{P}$ for this to happen. So condition (4.1) is necessary.

Now, assuming (4.1), we prove the existence of a Steiner rooted k-edge-connected orientation by choosing a suitable h, m and showing that the conditions of the min-max formula of Theorem 4.2.4 are satisfied. First we set up our problem in the language of Theorem 4.2.4. Since we are looking for a Steiner rooted k-arc-connected orientation, for each $X \subseteq S$, we set $h(X) = k$ if $r \notin X$ and $h(X) = 0$ otherwise. Then an orientation that covers the Steiner extension h^* of h is a Steiner rooted k-arc-connected orientation (see the proof of Theorem 4.1.3 in Section 4.2.1). By Lemma 4.3.2, we can set $m(v) = 1$

for each $v \in V(G) - S$.

Now let us verify that the condition of Theorem 4.2.4 is satisfied. That is, we need to show that $e_{\mathcal{Q}} \geq \sum_{i=1}^{t} h'(X_i)$ holds for every subpartition $\mathcal{Q} = \{X_1, \ldots, X_t\}$ of V. We can assume that either $X_i \subseteq S$ or $X_i \in V(G) - S$; otherwise $h'(X_i) = -\infty$ or $X_i = V(G)$ and so the inequality holds trivially. Let $\mathcal{Q} = \{X_1, \ldots, X_p, Z_1, \ldots, Z_q\}$ be a subpartition of V where $X_i \subseteq S$ for all i and $Z_j \in V(G) - S$ for all j.

Let $X_0 := S - X_1 - \cdots - X_p$. Recall that $h'(X_i) = h(X_i) + \max_{Y \subseteq V(G)-S}\{d(X_i, Y) - m(Y) + E(Y)\} = h(X_i) + \max_{Y \subseteq V(G)-S}\{d(X_i, Y) - m(Y)\}$ since there are no edges between Steiner vertices. For each X_i, let Y_i be a set which yields the maximum for $h'(X_i)$. From the definition of $h'(X_i)$, if a vertex $v \in V(G) - S$ has at least two neighbours in X_i, then it must belong to Y_i; if v has no neighbour in X_i, then it must not belong to Y_i. If v has only one neighbour in X_i, then we can assume $v \in Y_i$ or we can assume $v \notin Y_i$. Since each vertex $v \in V(G) - S$ is of degree at most 3 and satisfies $N(v) \subseteq X_0 \cup \ldots \cup X_p$, we can assume $\{X_0 \cup Y_0, X_1 \cup Y_1, \ldots, X_p \cup Y_p\}$ has the property that each v belongs to exactly one Y_i. That is, $\{Y_0, \ldots, Y_p\}$ is a partition of $V(G) - S$. So $\mathcal{P} = \{X_0 \cup Y_0, X_1 \cup Y_1, \ldots, X_p \cup Y_p\}$ is a partition of $V(G)$.

By (4.1), we have $e_{\mathcal{P}} \geq \sum_{i=1}^{p} h(X_i)$; notice that this holds even if $X_0 = \emptyset$. Let W_i be the set of vertices in Y_i which have an edge to X_0. Consider the subpartition $\mathcal{P}' = \{X_1, \ldots, X_p\}$. Now we compare $e_{\mathcal{P}'}$ to $e_{\mathcal{P}}$. Notice that the edges in $\delta(X_i, Y_i)$ for $1 \leq i \leq p$ are counted in $e_{\mathcal{P}'}$ but not in $e_{\mathcal{P}}$, and the edges in $\delta(X_0, Y_1 \cup \ldots \cup Y_p) = \cup_{i=1}^{p} \delta(X_0, W_i)$ are counted in $e_{\mathcal{P}}$ but not in $e_{\mathcal{P}'}$. Also, the edges in $\delta(Y_0, X_1 \cup \ldots \cup X_p)$ are counted in both $e_{\mathcal{P}}$ and $e_{\mathcal{P}'}$, and the edges in $\delta(Y_0, X_0)$ are not counted in both $e_{\mathcal{P}}$ and $e_{\mathcal{P}'}$. Since each vertex in W_i has exactly one edge to X_0 (otherwise it would belong to Y_0), we have $e_{\mathcal{P}'} = e_{\mathcal{P}} - d(X_0, Y_1 \cup \ldots \cup Y_p) + \sum_{i=1}^{p} d(X_i, Y_i) = e_{\mathcal{P}} + \sum_{i=1}^{p}(d(X_i, Y_i) - d(X_0, W_i)) = e_{\mathcal{P}} + \sum_{i=1}^{p}(d(X_i, Y_i) - |W_i|) \geq \sum_{i=1}^{p}(h(X_i) + d(X_i, Y_i) - |W_i|) = \sum_{i=1}^{p}(h'(X_i) + m(Y_i) - |W_i|) = \sum_{i=1}^{p}(h'(X_i) + |Y_i - W_i|)$ since $m(v) = 1$ for each $v \in V(G) - S$. Let $Z' := Z_1 \cup \ldots \cup Z_q - (Y_1 - W_1) - \ldots - (Y_p - W_p)$. So every vertex in Z' is either in Y_0 or in

W_i for some i. Therefore, by the definition of Y_0 and W_i, each vertex in Z' has at least one edge to X_0. Notice that the edges from Z' to X_0 are counted in $e_{\mathcal{Q}}$ but not in $e_{\mathcal{P}'}$. So $e_{\mathcal{Q}} \geq e_{\mathcal{P}'} + |Z'| \geq \sum_{i=1}^{p}(h'(X_i) + |Y_i - W_i|) + |Z'| \geq \sum_{i=1}^{p} h'(X_i) + |Z_1 \cup \ldots \cup Z_q| = \sum_{i=1}^{p} h'(X_i) + \sum_{j=1}^{q} m(Z_i) = \sum_{i=1}^{p} h'(X_i) + \sum_{j=1}^{q} h'(Z_i)$, as required.

Since $\mathcal{Q}$ is an arbitrary subpartition, this implies that the conditions of Theorem 4.2.4 are satisfied. Hence there exists a Steiner rooted k-arc-connected orientation (with every Steiner vertex having indegree exactly 1). This completes the proof. ∎

4.4 Steiner Rooted Orientations of Hypergraphs

This section contains the proof of our main result of this chapter (Theorem 4.1.1). The general scheme is to similar to the approach used in Section 2.2.1. We shall consider a minimal counterexample $\mathcal{H}$ of Theorem 4.4.3 with the minimum number of edges and then the minimum number of vertices. Note that Theorem 4.4.3 is a stronger version of Theorem 4.1.1 with an "extension property" introduced (Definition 4.4.1). The extension property allows us to apply a graph decomposition procedure to simplify the structures of $\mathcal{H}$ significantly (Corollary 4.4.5, Corollary 4.4.6). 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 4.4.10). To better illustrate the proof idea, we also include a proof of Theorem 4.4.3 in the special case of *rank* 3 hypergraphs (Lemma 4.4.7), where every hyperedge is of size at most 3.

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

Definition 4.4.1 (EXTENSION PROPERTY FOR STEINER ROOTED-ORIENTATIONS.) *Given $H = (V, \mathcal{E})$, $S \subseteq V(H)$ and a vertex $s \in S$, a Steiner rooted-orientation D of H* extends s *if:*

1. $d^{in}_D(s) = d_H(s)$;
2. $d^{in}_D(Y; \overline{s}) \geq \overrightarrow{d}_D(Y, s)$ *for every* $Y \subseteq V(G)$ *for which* $Y \cap S = \emptyset$;

We call s the special sink *of D.*

Condition (2) of the extension property is a technical condition that will be needed in the proof of Lemma 4.4.4. 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 4.4.2 *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 $\{\overrightarrow{P_1}, \ldots, \overrightarrow{P_k}\}$ between r and v. Now, let D' be an orientation with the same orientation as D except the orientations of all the hyperarcs in $P_1 \cup \ldots \cup P_k$ are reversed. To be more precise, let $\overrightarrow{P_i} = \{v_0, a_0, v_1, a_1, \ldots, a_{l-1}, v_l\}$ where a_i has v_i as the tail and v_{i+1} as a head, then $\overleftarrow{P_i} = \{v_l, \overleftarrow{a_{l-1}}, \ldots, \overleftarrow{a_0}, v_0\}$ where $\overleftarrow{a_i}$ has v_{i+1} as the tail and v_i as a head. For a directed path $\overrightarrow{P} = \{v_0, a_0, v_1, a_1, \ldots, a_{l-1}, v_l\}$, recall that 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 $\overrightarrow{P}$ 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 $d^{in}_{D'}(X) \geq k$ for every $X \subseteq V(H)$ which satisfies

$v \notin X$ and $X \cap S \neq \emptyset$. If $r \in X$, then $\{\overleftarrow{P_1}, \ldots, \overleftarrow{P_k}\}$ are k hyperarc-disjoint paths from v to r in D', where $\overleftarrow{P_i}$ denotes the reverse path of $\overrightarrow{P_i}$. Hence $d^{in}_{D'}(X) \geq k$ for such X. So we assume $r \notin X$. As D is a Steiner rooted k-hyperarc-connected orientation, we have $d^{in}_{D}(X) \geq k$. Recall that D and D' differ only on the orientations of the paths in $\{P_1, \ldots, P_k\}$. Notice that each path $\overrightarrow{P_i}$ has both endpoints outside of X, and thus $\overrightarrow{P_i}$ enters X the same number of times as it leaves X. Therefore, by reorienting $\overrightarrow{P_i}$ to $\overleftarrow{P_i}$ for all i, we have $d^{in}_{D'}(X) = d^{in}_{D}(X) \geq k$ 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 4.4.1. 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 4.4.1 is satisfied. For a subset $Y \subseteq V(H)$ with $Y \cap S = \emptyset$, $\overrightarrow{P_i}$ 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 $d^{in}_{D'}(Y; \overline{s}) = d^{in}_{D}(Y; \overline{s}) \geq \overrightarrow{d}_D(Y, s) = \overrightarrow{d}_{D'}(Y, s)$.
- $s \in a_1$ and $s \notin a_2$. Then $d^{in}_{D'}(Y; \overline{s}) = d^{in}_{D}(Y; \overline{s}) + 1 \geq \overrightarrow{d}_D(Y, s) + 1 = \overrightarrow{d}_{D'}(Y, s)$.
- $s \notin a_1$ and $s \in a_2$. Then $d^{in}_{D'}(Y; \overline{s}) = d^{in}_{D}(Y; \overline{s}) - 1 \geq \overrightarrow{d}_D(Y, s) - 1 = \overrightarrow{d}_{D'}(Y, s)$.
- $s \notin a_1$ and $s \notin a_2$. Then $d^{in}_{D'}(Y; \overline{s}) = d^{in}_{D}(Y; \overline{s}) \geq \overrightarrow{d}_D(Y, s) = \overrightarrow{d}_{D'}(Y, s)$.

Since we have $d^{in}_{D}(Y; \overline{s}) \geq \overrightarrow{d}_D(Y, s)$ to start with, by reorienting $\overrightarrow{P_i}$ to $\overleftarrow{P_i}$, we still have $d^{in}_{D'}(Y; \overline{s}) \geq \overrightarrow{d}_{D'}(Y, s)$. Hence the second condition in Definition 4.4.1 is also satisfied. Therefore, D' is a Steiner rooted k-hyperarc-connected orientation that extends s. This proves the lemma. ∎

We shall prove the following theorem which clearly implies Theorem 4.1.1.

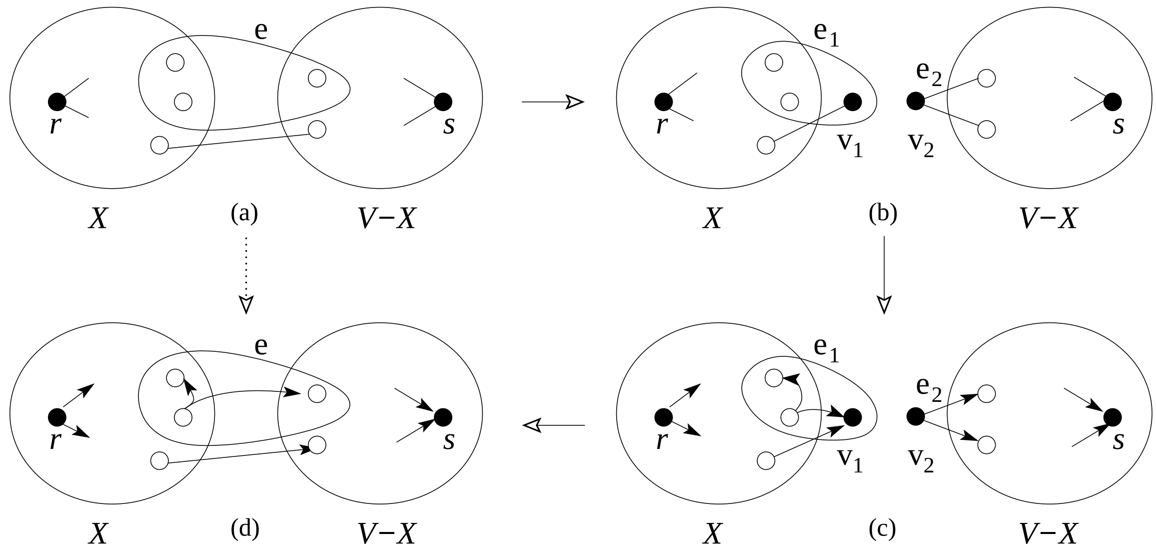

Figure 4.1: **An illustration of the proof of Lemma 4.4.4.**

Theorem 4.4.3 *Suppose $H = (V, \mathcal{E})$ is a hypergraph. If S is $2k$-hyperedge-connected in H, then there is a Steiner rooted k-hyperarc-connected orientation of H. In fact, given any vertex $s \in S$ of degree $2k$, there is a Steiner rooted k-hyperarc-connected orientation that extends s.*

Let $\mathcal{H}$ be a minimal counterexample of Theorem 4.4.3. In the following we say a set X is *tight* if $d(X) = 2k$; 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 Figure 4.1 for an illustration).

Lemma 4.4.4 *There is no nontrivial tight set in $\mathcal{H}$.*

Proof. Suppose there exists a nontrivial tight set X, i.e. $d_{\mathcal{H}}(X) = 2k$, $|X| \geq 2$ and $|V(\mathcal{H}) - X| \geq 2$. Apply the cut decomposition operation as defined in Section 3.4 on X to obtain two graphs H_1 and H_2. So, $V(H_1) = X \cup \{v_1\}$, $V(H_2) = (V(\mathcal{H}) - X) \cup \{v_2\}$ 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 precise, for a hyperedge e 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 to them as the *corresponding hyperedges* of e in H_1 and H_2 respectively. Note that $|e_1|, |e_2| \geq 2$.

Since X is non-trivial, 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 special sink of H_1 to be v_1 and the special 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 choice of $\mathcal{H}$, H_2 has a Steiner rooted k-hyperarc-connected orientation D_2 that extends s. By Lemma 4.4.2, we can assume the root of D_2 is v_2. Similarly, by the choice 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.

Now we claim that the concatenation D of the two orientations gives a Steiner rooted k-hyperarc-connected orientation of $\mathcal{G}$ that extends s. Notice for an hyperedge e in $\delta_{\mathcal{H}}(X)$, its corresponding hyperedge in H_1 is oriented with v_1 as a head (by the extension property of D_1), and its corresponding hyperedge 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. Now we show that D is a Steiner rooted k-hyperarc-connected orientation. By using Menger's theorem (Proposition 2.2.2), it suffices to show that $d^{in}_D(X) \geq k$ for every $X \subseteq V(\mathcal{H})$ for which $r \notin X$ and $X \cap S \neq \emptyset$. There are two cases to consider:

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

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

$$d^{in}_D(X) \geq d^{in}_{D_1}(X_1; \overline{v_1}) + d^{in}_{D_2}(X_2) - \overrightarrow{d}_D(X_1, X_2).$$

   Note that $\overrightarrow{d}_{D_1}(X_1, v_1) \geq \overrightarrow{d}_D(X_1, X_2)$. Therefore, by property (ii) of Definition 4.4.1, $d^{in}_{D_1}(X_1; \overline{v_1}) \geq \overrightarrow{d}_{D_1}(X_1, v_1) \geq \overrightarrow{d}_D(X_1, X_2)$. And hence $d^{in}_D(X) \geq$

$d^{in}_{D_2}(X_2) \geq k$; the second inequality is by the properties of D_2.

This implies that D is a Steiner rooted k-hyperarc-connected orientation of $\mathcal{H}$. Finally, we need to check that D extends s. The first property of Definition 4.4.1 follows from our construction. It remains to verify that property (ii) of Definition 4.4.1 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:

$$d^{in}_D(Y; \overline{s}) \geq d^{in}_{D_1}(Y_1; \overline{v_1}) + d^{in}_{D_2}(Y_2; \overline{s}) - \overrightarrow{d}_D(Y_1, Y_2; \overline{s}).$$

By property (ii) of Definition 4.4.1 applied to the extension property of H_1, we have $d^{in}_{D_1}(Y_1; \overline{v_1}) \geq \overrightarrow{d}_{D_1}(Y_1, v_1) \geq \overrightarrow{d}_D(Y_1, Y_2; \overline{s}) + \overrightarrow{d}_D(Y_1, s)$. Therefore, $d^{in}_D(Y; \overline{s}) \geq \overrightarrow{d}_D(Y_1, s)$ $+ d^{in}_{D_2}(Y_2; \overline{s})$. By property (ii) of the extension property of D_2, we have $d^{in}_{D_2}(Y_2; \overline{s}) \geq \overrightarrow{d}_{D_2}(Y_2, s)$ and hence $d^{in}_D(Y; \overline{s}) \geq \overrightarrow{d}_D(Y_1, s) + \overrightarrow{d}_{D_2}(Y_2, s) = \overrightarrow{d}_D(Y_1, s) + \overrightarrow{d}_D(Y_2, s) = \overrightarrow{d}_D(Y, s)$, as required. This verifies that D extends s, which contradicts that $\mathcal{H}$ is a counterexample. ∎

From Lemma 4.4.4, we obtain the following two important corollaries.

Corollary 4.4.5 *Each hyperedge of $\mathcal{H}$ of size at least 3 contains only terminal vertices.*

Proof. Suppose e is a hyperedge of $\mathcal{H}$ where $t \in e$ 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. Otherwise, there exists a set X which separates two terminals with $d_{\mathcal{H}}(X) = 2k$ and $d_{H'}(X) < 2k$. So $e \in \delta_{\mathcal{H}}(X)$. Without loss of generality, we assume $t \in X$. 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 4.4.4. ∎

Corollary 4.4.6 *There is no edge between two Steiner vertices in $\mathcal{H}$.*

Proof. This follows from a similar argument as in Corollary 4.4.5. 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 4.4.4. ∎

4.4.1 The Bipartite Representation of $\mathcal{H}$

Using Corollary 4.4.5 and Corollary 4.4.6, 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 4.4.5, 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; all other vertices are non-terminal vertices in B. By Corollary 4.4.5 and Corollary 4.4.6, there is no edge between two non-terminal vertices in B. Hence B is a bipartite graph. To distinguish the non-terminal vertices correspond to Steiner vertices in $\mathcal{H}$ and the non-terminal vertices correspond to hyperedges in $\mathcal{E}'$, we call the former the Steiner vertices and the latter the *hyperedge vertices*. See Figure 4.2 for an illustration.

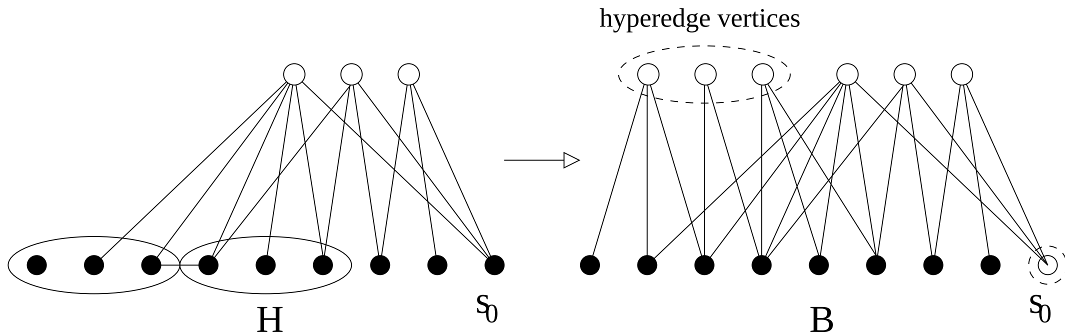

Figure 4.2: **The bipartite representation B of $\mathcal{H}$.**

4.4.2 Rank 3 Hypergraphs

To better illustrate the proof idea for the general case, we first prove the case for rank 3 hypergraphs. This motivates the proof for general hypergraphs, which is considerably more complicated.

Lemma 4.4.7 *$\mathcal{H}$ is not a rank 3 hypergraph.*

Proof. Suppose, by way of contradiction, that $\mathcal{H}$ is of rank 3. Since $\mathcal{H}$ is of rank 3, all hyperedge vertices in B are of degree at most 3. The crucial use of the rank 3 assumption is the following, which allows us to relate the hyperedge-connectivity of $\mathcal{H}$ to edge-connectivity in B.

Proposition 4.4.8 *S is $2k$-hyperedge-connected in $\mathcal{H}$ iff 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 $\mathcal{H}$. ∎

With Proposition 4.4.8, we can apply Mader's splitting-off lemma (Lemma 3.5.1) to prove the following.

Lemma 4.4.9 *Every Steiner vertex v of $\mathcal{H}$ is 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 lemma (Lemma 3.5.1) 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 4.4.6, s_1 and s_2 are terminal vertices. We add a new Steiner vertex e to $V(B)$ and replace the edge s_1s_2 by two new edges es_1 and es_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 4.4.8, 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 es_1 and es_2 as $\overrightarrow{s_1e}$ and $\overrightarrow{es_2}$ 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 $\mathcal{H}$, a contradiction. ∎

Now we are ready to finish the proof of Lemma 4.4.7. Construct $B' = B - s$, where we remove all edges in B which are incident with s. Treat all the vertices in $V(B) - S$ as Steiner vertices. We shall use Theorem 4.3.3 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'. Hence B' is k-partition-connected.

By Theorem 4.3.3 and Lemma 4.3.2, 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 $\mathcal{H}$. By this construction, property (i) of Definition 4.4.1 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 4.4.9, y can have at most one outgoing arc to s; otherwise $d_{\mathcal{H}}(\{s, y\}) < 2k$ which contradicts our assumption (recall that $d_{\mathcal{H}}(s) = 2k$ as s is the special sink). Since Y induces an independent set by Corollary 4.4.6 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 4.4.5, so we have $d_D^{in}(Y; \overline{s}) \geq \overrightarrow{d}(Y, s)$; this implies that D satisfies property (ii) of Definition 4.4.1 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 $d_D^{in}(X) \geq k$ 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 $d_D^{in}(X) \geq d_D^{in}(s) = 2k$ as s is the special 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. ∎

Notice that the above result implies a special case of Nash-Williams' strong orientation theorem: If S is $2k$-edge-connected in G, then G has a Steiner rooted k-arc-connected orientation. Our proof of this special case is different from that of Nash-Williams, as we use a different induction hypothesis. As far as the algorithmic aspects are concerned, a naive implementation of the two approaches have similar running times. This is because both methods involve reducing the original instance into instances with no tight sets (using the cut-decomposition operation), and then applying combinatorial techniques to those instances (in our case similar techniques as in Theorem 2.2.26 are used, and in Nash-Williams' case edge splitting-off techniques are used). A more efficient implementation of Nash-Williams' theorem can be found in [35].

4.4.3 General Hypergraphs

For the proof of Theorem 4.4.3 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 4.4.9. 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 4.4.7, then every Steiner vertex would end up with indegree $\lfloor d(v)/2 \rfloor$ in the resulting orientation of B. This motivates us to apply the Degree-Specified Steiner Rooted-Orientation problem by "hardwiring" $m(v) = \lfloor d(v)/2 \rfloor$ to simulate the splitting-off process. Also, we "hardwire" the indegree of the sink to be $2k$ for the extension property. (In the example of Figure 4.2, 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 4.4.3, which shows that a minimal counterexample of Theorem 4.4.3 does not exist.

Theorem 4.4.10 *Suppose that S is $2k$-hyperedge-connected in H, there is no edge between two Steiner vertices, and no hyperedge 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 4.2.4). To get an instance of that problem, we construct a bipartite graph G from the hypergraph H as in the proof of Lemma 4.4.7. Let the set of terminals be $S' := S - s_0$. All other vertices are non-terminal vertices. To distinguish the three types of non-terminal vertices (vertices in $V(H)-S$, the hyperedge vertices, and the special sink s_0), we use the term Steiner vertex exclusively for vertices in $V(H) - S$. So, the set of terminal vertices S' and s_0 is on one side of G, and the set of all Steiner vertices and hyperedge vertices is on the other side of G. Notice that edges in $\delta(s_0)$ are the only edges between non-terminal vertices.

The in-degree specification m is defined by:

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

Let $r \in S$ be the root vertex. Set $h(X) := k$ for every $X \subseteq S'$ with $r \notin X$, and $h(X) := 0$ otherwise. Recall (from the proof of Theorem 4.1.3 in Section 4.2.1) that finding a Steiner rooted k-arc-connected orientation in G which satisfies the indegree specification $m : (V(G) - S) \to Z^+$ is equivalent to finding an orientation in G that covers the Steiner extension h^* of h which satisfies the indegree specification $m : (V(G) - S) \to Z^+$. By Lemma 4.2.2, the latter problem is equivalent to finding an orientation in G that covers h' (with no indegree specification) which defined as follows.

- $h'(Z) = m(Z)$ for $Z = \{v\}, v \in V(G) - S'$.
- $h'(S_i) = k + \max_{Z \subseteq V(G)-S'} \{d(S_i, Z) - m(Z) + E(Z)\}$ if $r \notin S_i \subseteq S'$;
- $h'(S_i) = \max_{Z \subseteq V(G)-S'} \{d(S_i, Z) - m(Z) + E(Z)\}$ if $r \in S_i \subseteq S'$;
- $h'(\emptyset) = h'(V(G)) = 0$;
- $h'(X) = -\infty$ otherwise.

We remark that if we do not know whether $r \in S_i$, then we just write $h'(S_i) = h(S_i) + \max_{Z \subseteq V(G)-S'} \{d(S_i, Z) - m(Z) + E(Z)\}$.

By the theorem on degree-specified orientations (Theorem 4.2.4), G has a Steiner rooted k-arc-connected orientation with the specified in-degrees if and only if

$$e_Q \geq \sum_{i=1}^{t} h'(X_i) \tag{4.2}$$

holds for every subpartition $Q = \{X_1, \ldots, X_t\}$ of $V(G)$.

Since $h'(s_0) = m(s_0) = 2k$ and $d(s_0) = 2k$, from the definition of h' above, we can assume that s_0 does not belong to the set Z in the formula of $h'(S_i)$ for any S_i. Also, since $h'(s_0) = m(s_0) = d(s_0)$, we can assume that $\{s_0\}$ is in the subpartition of (4.2) (by putting s_0 in Q can only make (4.2) more difficult to hold).

We use the following notation. Here we only consider a set X_i if $h(X_i) \neq -\infty$; otherwise (4.2) holds trivially; thus $X_i \subseteq S'$ or $X_i \in V(G) - S'$ (note that $X_i \subseteq V(G) - S'$ can be assumed to be a singleton). If $X_i \subseteq S'$, we denote it by S_i. We assume this is the case for $X_1, \ldots, X_p$. If X_i is a hyperedge vertex, we denote it by a_i; we define $A := \cup a_i$. If X_i is a Steiner vertex, we denote it by b_i; we define $B := \cup b_i$. We denote $S^* := \cup S_i \cup \{s_0\}$. For any subset $S'' \subseteq S' \cup \{s_0\}$, we denote $\overline{S''} := S' - S''$. So, for example, $\overline{S^*} = S' - S^*$.

By (4.2), we need to prove that:

$$\begin{aligned} e_Q &= \sum_{i=1}^{p} d(S_i) + d(s_0) + \sum_{v \in A} d(v, \overline{S^*}) + \sum_{v \in B} d(v, \overline{S^*}) \\ &\geq \sum_{i=1}^{p} h'(S_i) + h'(s_0) + \sum_{v \in A} h'(v) + \sum_{v \in B} h'(v). \end{aligned}$$

Since $d(s_0) = h'(s_0)$ and $h'(A_i) = 1$ and $h'(B_i) = \lfloor d_H(B_i)/2 \rfloor$, the above is equivalent to:

$$e_Q = \sum_{i=1}^{p} d(S_i) + \sum_{v \in A} d(v, \overline{S^*}) + \sum_{v \in B} d(v, \overline{S^*}) \geq \sum_{i=1}^{p} h'(S_i) + |A| + \sum_{v \in B} \lfloor d(v)/2 \rfloor. \tag{4.3}$$

By comparing the two sides of (4.3), we can assume the following without loss of generality.

Proposition 4.4.11 *A hyperedge vertex a is in A if and only if $d(a, \overline{S^*}) = 0$. A Steiner vertex b is in B if and only if $d(b, S^*) > \lfloor d(b)/2 \rfloor$.*

Therefore, by Proposition 4.4.11, we can take

$$e_Q = \sum_{i=1}^{p} d(S_i) + \sum_{v \in B} d(v, \overline{S^*}).$$

We denote by U the set of all Steiner vertices, and denote by E the set of all hyperedge vertices. Define $Y := U \cup E$. Recall that for S_i, $h'(S_i) = h(S_i) + \max_{Z \subseteq V(G) - S'}\{d(S_i, Z) - m(Z)\}$; let $Y_i \subseteq V(G) - S'$ be a set that yields the maximum value for $h'(S_i)$. We denote the Steiner vertices which belong to Y_i by T_i, and $T := \cup T_i$. For a subset $T'' \subseteq U$, we denote $\overline{T''} := U - T''$. One should distinguish the difference between $\overline{S''}$ and $\overline{T''}$, the former is defined to be $S' - S''$ while the latter is defined to be $U - T''$. Informally, the former means the complement on the terminal vertices while the latter means the complement on the Steiner vertices. To avoid confusion, the former notation will only be used for sets which have S as a prefix, while the latter notation will only be used for sets which have T as a prefix.

From the definition of h', we can assume the following without loss of generality.

Proposition 4.4.12 *A hyperedge vertex e belongs to Y_i if and only if $d(e, S_i) > 0$. A Steiner vertex v belongs to Y_i if and only if $d(v, S_i) > \lfloor d(v)/2 \rfloor$.*

Remark: We claim that we can assume that $p \geq 1$. Suppose $p = 0$. Then $S^* = \{s_0\}$. Since S is $2k$-hyperedge-connected in H and $s_0 \in S$ is of degree exactly $2k$, there is no hyperedge vertex a with $N(a) = \{s_0\}$ and there is no Steiner vertex b with $d(b, s_0) > \lfloor d(b)/2 \rfloor$. So, by Proposition 4.4.11, $A = \emptyset$ and $B = \emptyset$, and thus (4.3) and hence (4.2) hold. So we can assume $p \geq 1$. If $p = 1$, then we claim that we can assume that $S_1 \neq S'$. Otherwise, $\overline{S^*} = \emptyset$. By Proposition 4.4.11, we have $A = E$ and $U = B$. So, by (4.2), $e_Q = d(S' + s_0, E \cup U) = d(S' + s_0, Y)$. Similarly, by Proposition 4.4.12, $Y_i = Y$. So $h'(S') = d(S', Y) - h'(Y)$. Hence, $e_Q = d(S' + s_0, Y) = d(S', Y) + d(s_0) - h'(Y) + h'(Y) = h'(S') + h'(s_0) + \sum_{v \in Y} h'(v)$, and thus (4.2) holds.

Notice that $T \subseteq B$ by Proposition 4.4.11 and Proposition 4.4.12. Since $d(S_i) =$

$d(S_i, Y_i) + d(S_i, Y - Y_i)$ and $h'(S_i) = h(S_i) + d(S_i, Y_i) - m(Y_i)$, we have

$$d(S_i) = h'(S_i) - h(S_i) + d(S_i, Y - Y_i) + m(Y_i)$$

$$\geq h'(S_i) - h(S_i) + d(S_i, Y - Y_i) + |N_E(S_i)| + \sum_{v \in T_i} \lfloor d(v)/2 \rfloor,$$

where $N_E(S_i)$ denotes the set of hyperedge vertices which are adjacent to S_i.

Hence,

$$e_Q = \sum_{i=1}^{p} d(S_i) + \sum_{v \in B} d(v, \overline{S^*})$$

$$= \sum_{i=1}^{p} \Big(h'(S_i) - h(S_i) + |N_E(S_i)| + \sum_{v \in T_i} \lfloor d(v)/2 \rfloor + d(S_i, Y - Y_i) \Big) + \sum_{v \in B} d(v, \overline{S^*}). \quad (4.4)$$

Focus on the last three terms of (4.4) which is

$$\sum_{i=1}^{p} \sum_{v \in T_i} \lfloor d(v)/2 \rfloor + \sum_{i=1}^{p} d(S_i, Y - Y_i) + \sum_{v \in B} d(v, \overline{S^*}). \quad (4.5)$$

Since all vertices in T are in B, the last term of (4.5) becomes

$$\sum_{v \in B} d(v, \overline{S^*}) = \sum_{v \in T \cap B} d(v, \overline{S^*}) + \sum_{v \in \overline{T} \cap B} d(v, \overline{S^*}) = \sum_{v \in T} d(v, \overline{S^*}) + \sum_{v \in \overline{T} \cap B} d(v, \overline{S^*}).$$

By Proposition 4.4.12, there is no hyperedge vertex in $Y - Y_i$ which is adjacent to S_i. So the middle term of (4.5) is

$$\sum_{i=1}^{p} d(S_i, Y - Y_i) = \sum_{v \in \overline{T}} d(v, S^* - s_0) + \sum_{v \in T} d(v, S^* - s_0 - S(v)),$$

where $S(v) := S_i$ for $v \in T_i$.

Adding up, (4.5) is equal to:

$$\sum_{v \in T} \lfloor d(v)/2 \rfloor + \sum_{v \in \overline{T}} d(v, S^* - s_0) + \sum_{v \in T} d(v, S^* - s_0 - S(v)) + \sum_{v \in T} d(v, \overline{S^*}) + \sum_{v \in \overline{T} \cap B} d(v, \overline{S^*})$$

$$= \sum_{v \in T} \lfloor d(v)/2 \rfloor + \sum_{v \in \overline{T}} d(v, S^* - s_0) + \sum_{v \in T} d(v, \overline{S(v)}) + \sum_{v \in \overline{T} \cap B} d(v, \overline{S^*})$$

$$= \sum_{v \in T} \lfloor d(v)/2 \rfloor + \sum_{v \in \overline{T} \cap B} d(v, S') + \sum_{v \in \overline{T} - B} d(v, S^* - s_0) + \sum_{v \in T} d(v, \overline{S(v)}). \quad (4.6)$$

Plug back (4.6) into (4.4). To prove that it is at least the right hand side of (4.3), one must show:

$$-\sum_{i=1}^{p} h(S_i) + \sum_{i=1}^{p} |N_E(S_i)| - |A|$$

$$-\sum_{v\in B}\lfloor d(v)/2\rfloor + \sum_{v\in T}\lfloor d(v)/2\rfloor + \sum_{v\in \overline{T}\cap B} d(v,S') + \sum_{v\in\overline{T}-B} d(v,S^*-s_0) + \sum_{v\in T} d(v,\overline{S(v)}) \geq 0. \tag{4.7}$$

Since $T \subseteq B$, the bottom line of (4.7) is

$$\begin{aligned}
&-\sum_{v\in\overline{T}\cap B}\lfloor d(v)/2\rfloor + \sum_{v\in\overline{T}\cap B} d(v,S') + \sum_{v\in\overline{T}-B} d(v,S^*-s_0) + \sum_{v\in T} d(v,\overline{S(v)}) \\
&= \sum_{v\in\overline{T}\cap B}\lceil d(v)/2\rceil - \sum_{v\in\overline{T}\cap B} d(v,s_0) + \sum_{v\in\overline{T}-B} d(v,S^*-s_0) + \sum_{v\in T} d(v,\overline{S(v)}) \\
&= \sum_{v\in\overline{T}\cap B}\lceil d(v)/2\rceil + \sum_{v\in\overline{T}-B} d(v,S^*-s_0) + \sum_{v\in T} d(v,\overline{S(v)}+s_0) - \sum_{v\in B} d(v,s_0) \\
&= \sum_{v\in\overline{T}\cap B}\lceil d(v)/2\rceil + \sum_{v\in\overline{T}-B} d(v,S^*) + \sum_{v\in T} d(v,\overline{S(v)}+s_0) - \sum_{v\in U} d(v,s_0),
\end{aligned}$$

since $\overline{T} - B = U - B$.

Plugging back into (4.7), we need to show:

$$\sum_{i=1}^{p} |N_E(S_i)| - |A| + \sum_{v\in\overline{T}\cap B}\lceil d(v)/2\rceil + \sum_{v\in\overline{T}-B} d(v,S^*) + \sum_{v\in T} d(v,\overline{S(v)}+s_0) - \sum_{v\in U} d(v,s_0)$$

$$\geq \sum_{i=1}^{p} h(S_i). \tag{4.8}$$

Let $\mathcal{F} = \{S_1,\ldots,S_p\}$. We say a hyperedge vertex $v \in E$ is contained in a set S_i if $N(v) \subseteq S_i$. Let $A_2 \subseteq A$ be the set of hyperedge vertices which are not contained in S_i for some i. Let $E_2 \subseteq E$ be the set of hyperedge vertices which are not contained in S_i for some i and are not contained in $\overline{S^*}$. Informally, these hyperedge vertices correspond to “crossing” hyperedges, which are used to connect different S_i.

Two times the first two terms of (4.8) is:

$$2(\sum_{i=1}^{p} |N_E(S_i)| - |A|)$$

$$= 2(\sum_{e \in E} |\{X \in \mathcal{F} : e \cap X \neq \emptyset\}| - |A|)$$

$$= 2\Big(\sum_{e \in A} |\{X \in \mathcal{F} : e \cap X \neq \emptyset\}| - |A| + \sum_{e \in E-A} |\{X \in \mathcal{F} : e \cap X \neq \emptyset\}|\Big)$$

$$= 2\sum_{e \in A_2} (\{|X \in \mathcal{F} : e \cap X \neq \emptyset\}| - 1) + 2\sum_{e \in E-A} |\{X \in \mathcal{F} : e \cap X \neq \emptyset\}|$$

$$\geq \sum_{e \in A_2} |\{X \in \mathcal{F} : e \cap X \neq \emptyset\}| + 2\sum_{e \in E-A} |\{X \in \mathcal{F} : e \cap X \neq \emptyset\}|$$

$$= \Big(\sum_{e \in A_2} |\{X \in \mathcal{F} : e \cap X \neq \emptyset\}| + \sum_{e \in E-A} |\{X \in \mathcal{F} : e \cap X \neq \emptyset\}|\Big)$$

$$+ \sum_{e \in E-A} |\{X \in \mathcal{F} : e \cap X \neq \emptyset\}|$$

$$\geq \sum_{i=1}^{p} |N_{E_2}(S_i)| + |N_{E_2}(\overline{S^*})|,$$

since every hyperedge vertex in $E - A$ has an edge to $\overline{S^*}$ by Proposition 4.4.11.

With this, *two times* the left hand side of (4.8):

$$\geq \sum_{i=1}^{p} |N_{E_2}(S_i)| + |N_{E_2}(\overline{S^*})| + \sum_{v \in \overline{T} \cap B} d(v) + 2\sum_{v \in \overline{T} - B} d(v, S^*)$$

$$+2\sum_{v \in T} d(v, \overline{S(v)} + s_0) - 2\sum_{v \in U} d(v, s_0)$$

Rewrite $\sum_{v \in \overline{T} \cap B} d(v)$ as $\sum_{v \in \overline{T} \cap B}(d(v, s_0) + d(v, S^* - s_0) + d(v, \overline{S^*}))$, rewrite 2 times $\sum_{v \in \overline{T} - B} d(v, S^*)$ as $\sum_{v \in \overline{T} - B}(d(v, S^*) + d(v, s_0) + d(v, S^* - s_0))$, and also rewrite 2 times $\sum_{v \in T} d(v, \overline{S(v)} + s_0)$ as $\sum_{v \in T}(d(v, \overline{S(v)} + s_0) + d(v, s_0) + d(v, S^* - S(v) - s_0) + d(v, \overline{S^*}))$. Regrouping the terms, we have two times the left hand side of (4.8) is

$$\geq \Big(N_{E_2}(\overline{S^*}) + \sum_{v \in \overline{T} \cap B} d(v, \overline{S^*}) + \sum_{v \in \overline{T} - B} d(v, S^*) + \sum_{v \in T} d(v, \overline{S^*})\Big)$$

$$+\Big(\sum_{i=1}^{p} |N_{E_2}(S_i)| + \sum_{v \in \overline{T} \cap B} d(v, S^* - s_0) + \sum_{v \in \overline{T} - B} d(v, S^* - s_0)$$

$$+ \sum_{v \in T} (d(v, \overline{S(v)} + s_0) + d(v, S^* - S(v) - s_0))\Big)$$

$$+\Big(\sum_{v \in \overline{T} \cap B} d(v, s_0) + \sum_{v \in \overline{T} - B} d(v, s_0) + \sum_{v \in T} d(v, s_0) - 2\sum_{v \in U} d(v, s_0)\Big) \tag{4.9}$$

The last line of (4.9) is the easiest, which is equal to $-\sum_{v \in U} d(v, s_0) \geq -d(s_0) = -2k$.

The second line of (4.9) can be rewritten as follows:

$$\sum_{i=1}^{p} \Big(|N_{E_2}(S_i)| + \sum_{v \in \overline{T}} d(v, S_i) + \sum_{v \in T_i} d(v, \overline{S_i} + s_0) + \sum_{v \in T - T_i} d(v, S_i) \Big)$$

$$= \sum_{i=1}^{p} d_H(S_i \cup \{v \in T_i\}),$$

where H is the original hypergraph. If $p \geq 2$, then clearly $\overline{S_i} \neq \emptyset$. If $p = 1$, as remarked earlier, we also have $\overline{S_i} = \overline{S^*} \neq \emptyset$. So, each $S_i \cup \{v \in T_i\}$ is an S-separating set in H. As S is $2k$-hyperedge-connected in H, the second line of (4.9) is at least $2kp$.

Suppose $\overline{S^*} = \emptyset$. Then two times the right hand side of (4.8) is equal to $2\sum_{i=1}^{p} h(S_i) = 2k(p-1)$, since one S_i must contain the root r. Clearly, the first line of (4.9) is non-negative. So, two times the left hand side of (4.8) is at least $0 + 2kp - 2k = 2k(p-1)$, and thus (4.8) holds.

Henceforth we assume $\overline{S^*} \neq \emptyset$. The first line of (4.9) is equal to $d_H(\overline{S^*} \cup \{v \in \overline{T} - B\})$, where H is the original hypergraph. Notice that $(\overline{S^*} \cup \{v \in \overline{T} - B\})$ is an S-separating set in H, since $S^* \neq \emptyset$ and $\overline{S^*} \neq \emptyset$. As S is $2k$-hyperedge-connected in H, we must have $d_H(\overline{S^*} \cup \{v \in \overline{T} - B\}) \geq 2k$, which implies the first line is at least $2k$. So, two times the left hand side of (4.8) is at least $2k + 2kp - 2k = 2kp$. Two times the right hand side of (4.8) is at most $2kp$. Hence (4.8) holds. This implies (4.3) and hence (4.2) are true, which is exactly what we want.

So, we have a k-arc-connected orientation of G with the specified indegrees, which corresponds to a k-hyperarc-connected orientation of H since each hyperedge vertex has indegree 1. It remains to check that this extends s_0. The first property of the extension property (Definition 4.4.1) follows immediately from our construction, since the indegree of s_0 is $2k$. To check the second property of the extension property, we use a similar argument as in Lemma 4.4.7. 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 4.4.5, and so do not intersect s_0. Hence, $d^{in}(Y; \overline{s}) \geq \overrightarrow{d}(Y, s)$, as required. ∎

Theorem 4.4.10 shows that a minimal counterexample does not exist, and thus finishes the proof of Theorem 4.1.1.

We remark that in the proof of Theorem 4.4.10, the indegree specifications on the Steiner vertices have two uses. The major use is to apply Theorem 4.2.4 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 4.4.1).

4.5 Element-Disjoint Steiner Rooted Orientations

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 (please refer to Section 4.1.2 for definitions). The main result of this section is a proof of Theorem 4.1.2.

In [14] the problem of PACKING ELEMENT-DISJOINT STEINER TREES is considered, where the goal is to find a maximum number of element disjoint Steiner trees of an undirected graph. Cheriyan and Salavatipour [14] show that this problem is $\Omega(\log n)$-hard to approximate, and at the same time present a randomized $O(\log n)$-approximation algorithm. In particular, they prove that if S is $O(\log n)k$-element-connected in G, then G has k element-disjoint Steiner trees [14]; this is best possible as shown by the examples in [7]. Clearly, if G has k element-disjoint Steiner trees, then G has a Steiner rooted k-element-connected orientation. Therefore, the result in [14] implies an $O(\log n)$-approximation algorithm for the ELEMENT-DISJOINT STEINER ROOTED ORIENTATION problem as well.

Here we prove a tight connectivity upper bound for the problem. By applying the min-max formula of Theorem 4.2.4, we will prove that if S is $2k$-element-connected in G, then G has a Steiner rooted k-element-connected orientation. This implies the first polynomial time constant factor approximation algorithm for the problem as well.

The proof 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 was also shown in [43, 14] but we will give a proof here for completeness. 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 constrain the indegree of each Steiner vertex to be 1, then a Steiner rooted k-edge-connected orientation is also a Steiner rooted k-element-connected orientation; this is because each Steiner vertex cannot be in two edge-disjoint paths. This turns out to give a tight connectivity bound, which again demonstrates the strength of the DEGREE-SPECIFIED STEINER ORIENTATION problem.

We remark that the property that every Steiner vertex is of indegree 1 in the orientation will be used twice - once in Lemma 4.5.2 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 [43, 14].

Lemma 4.5.1 *(See also [43, 14].) Given an undirected graph $G = (V, E)$ and a set S of terminal vertices. Suppose S is k-element-connected in G. Then we can construct in polynomial time a graph $G' = (V', E')$ 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 l-element-connected orientation in G' with every Steiner vertex of indegree 1, then there is a Steiner rooted l-element-connected orientation in G.*

Proof. Given G, if there is no edge uv between two Steiner vertices, then $G' := G$ and we have nothing to prove. In the following, we will show that we can construct G' from G by deleting and/or contracting edges between Steiner vertices. Let $G_0 := G$. By assumption, G_0 satisfies properties (1) and (3). Suppose G_t satisfies properties (1) and (3) for $t \geq 0$. If G_t also satisfies (2), then $G' := G_t$ as desired. Otherwise, we shall construct a graph G_{t+1} which still satisfies properties (1) and (3) and has fewer edges between Steiner vertices than that in G_t. Let uv be an edge in G_t between two Steiner vertices. If S is k-element-connected in $G_t - uv$, then we simply set $G_{t+1} := G_t - uv$. Clearly, G_{t+1} still satisfies properties (1) and (3) and has fewer edges between Steiner vertices than that in G_t, as required.

So suppose S is not k-element-connected in $G_t - uv$; we shall show that $G_{t+1} := G_t/\{uv\}$ would have the desired properties (recall that $G_t/\{uv\}$ means contracting the edge uv in G_t). Property (1) is trivial. It is also clear that G_{t+1} has fewer edges between Steiner vertices than that in G_t. It remains to show that S is k-element-connected in G_{t+1} (i.e. property (3) is satisfied). Since S is not k-element-connected in $G_t - uv$, by Menger's theorem, in G_t, there is a set T of k elements which contains uv and whose removal disconnects a pair of terminal vertices a, b. Suppose P_{ab} is an arbitrary set of k element-disjoint paths between a and b. Then P_{ab} must contain a path that uses the edge uv. Suppose, by way of contradiction, that S is not k-element-connected in $G_t/\{uv\}$. By Menger's theorem, in G_t, there is a set R of k elements which contains $\{u, v\}$ and whose removal disconnects a from another terminal vertex c. Since P_{ab} must contain a path that uses the edge uv and R contains $\{u, v\}$, R cannot intersect all k element disjoint paths in P_{ab} and hence R cannot disconnect a and b. So $c \neq b$. Suppose P_{ac} is an arbitrary set of k element-disjoint paths between a and c. Then P_{ac} must contain a path that uses u but not v, and a path that uses v but not u. In particular P_{ac} does not use the edge uv. Since a, b are in the same component, by the same argument, any set of k element-disjoint paths between b and c does not use the edge uv. This implies a and

b are connected in $G_t - uv$, through c, and thus yields a contradiction. Therefore, S is k-element-connected in G_{t+1}, as required.

By repeating the above procedure, we will eventually obtain a graph G_m such that it satisfies properties (1) and (3), and also has no edges between two Steiner vertices. We set $G' := G_m$, and hence (1)-(3) hold.

Finally we prove (4) by showing that if we have a Steiner rooted l-element-connected orientation of $G' = G_m$ with every Steiner vertex of indegree 1, then there is a Steiner rooted l-element-connected orientation of $G = G_0$. In the following, we say a graph G is *good* if G has a subgraph H such that H has a Steiner rooted l-element-connected orientation of G with every Steiner vertex of indegree 1. Clearly, if G is *good*, then G has a Steiner rooted l-element-connected orientation by orienting the edges without an orientation arbitrarily. By assumption, $G' = G_m$ is good. Suppose G_{t+1} is good, then we shall show that G_t is good too. Suppose we delete an edge ab between two Steiner vertices a, b in G_t to obtain G_{t+1}. In this case we do not assign an orientation to the edge ab in G_t, while all other edges in G_t has the same orientation as in G_{t+1} (including the edges without an orientation). Clearly G_t is good.

Suppose we contract an edge ab between two Steiner vertices a, b in G_t to one Steiner vertex c in G_{t+1}. By the assumption that G_{t+1} is good, G_{t+1} has a subgraph H_{t+1} for which there is a l-element-connected orientation D_{t+1} with every Steiner vertex of indegree 1. If c has no incoming arc in D_{t+1}, then c is not useful in the orientation D_{t+1}, and hence G_t is good by using the same orientation as in G_{t+1} (with all the edges between a and b unoriented). So assume that x is the only vertex adjacent to c with xc oriented as $\overrightarrow{xc}$ in D_{t+1}. If the preimage of the edge xc in G_{t+1} is the edge xa, then we orient ab as $\overrightarrow{ab}$ in G_t; if the preimage of the edge xc in G_{t+1} is xb, then we orient ab as $\overrightarrow{ba}$ in G_t. If there are multiple edges between a and b, then only one of them is assigned an orientation. All other edges in G_t have the same orientation as in G_{t+1} (including the edges without an orientation). We set H_t to be the subgraph of G_t with edges having an orientation, and D_t be the orientation of H_t. It is easy to see that if there are l element-disjoint paths between the root and a terminal vertex in D_{t+1}, then there are l element-disjoint paths between the root and a terminal vertex in D_t (since c is of indegree 1 in D_{t+1}). Furthermore, every Steiner vertex is still of indegree 1 in D_t. So, G_t is good. Repeating the same argument, $G = G_0$ is good, and we are done. ∎

The previous lemma shows that we can assume there is no edge between Steiner vertices.

Remark 1: In the following, for convenience, we will assume there is no edge between terminal vertices as well. This can be achieved by introducing a new Steiner vertex of degree 2 for each edge between two terminal vertices.

Remark 2: The following lemma is actually a special case of Theorem 4.4.10; we give a proof here because it is much simpler than the proof of Theorem 4.4.10.

Lemma 4.5.2 *Given an undirected graph $G = (V, E)$ and a set S of terminal vertices. Suppose that G is a bipartite graph with S on one side and $V(G) - S$ on the other side. If S is $2k$-element-connected in G, then G has a Steiner rooted k-element-connected orientation with every Steiner vertex of indegree 1.*

Proof. Our plan is to reduce this problem to the Degree-Specified Steiner Rooted Orientation problem. The observation is that if we set $m(v) = 1$ for all $v \in V(G) - S$, then a Steiner rooted k-arc-connected orientation is also a Steiner rooted k-element-connected orientation. So, now, we just need to prove the existence of a Steiner rooted k-arc-connected orientation with the specified indegrees. To do so, we apply Theorem 4.2.4 to show that $2k$-element-connectivity is enough to guarantee such a degree-specified orientation.

Let $\mathcal{Q} = \{X_1, \ldots, X_p, Z_1, \ldots, Z_q\}$ be a subpartition of V, where $X_i \subseteq S$ for all i and $Z_j \in V(G) - S$ for all j. We need to show that

$$e_{\mathcal{Q}} \geq \sum_{i=1}^{p} h'(X_i) + \sum_{j=1}^{q} h'(Z_j). \tag{4.10}$$

For each X_i, $h'(X_i) \leq k + \max_{Y_i \subseteq V(G)-S}\{d(X_i, Y_i) - m(Y_i) + E(Y_i)\} = k + \max_{Y_i \subseteq V(G)-S}$ $\{d(X_i, Y_i) - |Y_i|\}$ since there are no edges between Steiner vertices and $m(v) = 1$ for all $v \in V(G) - S$. In fact, we can set $Y_i = \Gamma(X_i)$ where $\Gamma(X_i)$ denotes the set of Steiner vertices with an edge to X_i, as it attains the maximum of $h'(X_i)$. Let $d(X_i)$ denote the number of edges with precisely one endpoint in X_i. Since $Y_i = \Gamma(X_i)$, we have $d(X_i) = d(X_i, Y_i)$. Recall that $h'(X_i) \leq k + d(X_i, Y_i) - |Y_i|$, so $d(X_i) \geq h'(X_i) + |Y_i| - k = h'(X_i) + |\Gamma(X_i)| - k$. Let I be the set of vertices in $Z_1 \cup \ldots \cup Z_q$ with all its neighbours in $X_1 \cup \ldots \cup X_p$, and $Z^* := Z_1 \cup \ldots \cup Z_q$. So,

$$e_{\mathcal{Q}} \geq \sum_{i=1}^{p} d(X_i) + |Z^*| - |I| \geq \sum_{i=1}^{p}(h'(X_i) + |\Gamma(X_i)| - k) + \sum_{Z_j \in Z^* - I} h'(Z_j)$$

To prove (4.10), it remains to show that

$$\sum_{i=1}^{p}(|\Gamma(X_i)| - k) \geq \sum_{Z_j \in I} h'(Z_j) = |I|. \tag{4.11}$$

For $v \in V(G) - S$, we say v is contained in X_i if $N_G(v) \subseteq X_i$. Let $\Gamma'(X_i)$ be the subset of vertices of $\Gamma(X_i)$ which are not contained in X_i, and I' be the subset of I so that $v \in I'$ is not contained in X_i for some i. A Steiner vertex which is contained in some X_i contributes one to both sides of (4.11). So to prove (4.11), it suffices to show

$$\sum_{i=1}^{p} |\Gamma'(X_i)| \geq |I'| + kp. \tag{4.12}$$

Since S is $2k$-element-connected in G, we have $|\Gamma'(X_i)| \geq 2k$. Hence $\sum_{i=1}^{p} |\Gamma'(X_i)| \geq 2kp$. If $|I'| \leq kp$, then (4.12) is satisfied. So assume $|I'| > kp$. Since each element v in I' is not contained in any X_i, so v contributes at least 2 on the left hand side. Therefore, $\sum_{i=1}^{p} |\Gamma'(X_i)| \geq 2|I'| > |I'| + kp$ and (4.12) is again satisfied. So (4.12) always holds and this completes the proof. ∎

As discussed above, the following theorem (Theorem 4.1.2) follows immediately from Lemma 4.5.1 and Lemma 4.5.2.

Theorem 4.1.2: Given an undirected graph $G = (V, E)$ and a set S of terminal vertices. If S is $2k$-element-connected in G, then G has a Steiner rooted k-element-connected orientation.

4.6 Steiner Strongly Connected Orientations

Recall that a set of vertices S is strongly k-arc-connected in a directed graph D if there are k-arc-disjoint paths from u to v for every ordered pair of vertices $u, v \in S$. In this section we give a simple proof of the following special case of Nash-Williams' strong orientation theorem: If S is $2k$-edge-connected in G, then G has a Steiner strongly k-arc-connected orientation.

As in the case of the Steiner Tree Packing problem and the Steiner Rooted-Orientation problem, we define an appropriate extension property and show that this helps to decompose the problem into simpler instances. The rest of the proof is similar to Lovász's proof of Nash-Williams' weak orientation theorem as shown in Section 2.2.4.

Definition 4.6.1 (The Extension Property for Strong Orientations) *Given $G = (V, E)$, a vertex $s \in S$, and an orientation to each edge adjacent to s, an orientation D of G* extends *s if it is consistent with the orientations of those edges adjacent to s.*

In the following we say a terminal vertex v is tight if its degree is equal to $2k$. Notice that a tight terminal vertex must have indegree k and outdegree k in a Steiner strongly k-arc-connected orientation; we say such an orientation is a *valid* orientation of $\delta_G(v)$.

Theorem 4.6.2 *Let G be an undirected graph and $S \subseteq V(G)$. If S is $2k$-edge-connected in G, then G has a Steiner strongly k-arc-connected orientation. Furthermore, given a tight terminal vertex v and a valid orientation of $\delta_G(v)$, there is a Steiner strongly k-arc-connected orientation D of G that extends v.*

Proof. Let $\mathcal{G}$ be a counterexample with the minimum number of edges. Without loss of generality, we can assume that $\mathcal{G}$ is 2-edge-connected. First we show that, by using the extension property, we can restrict our attention to graphs with simpler structures. We say a subset of vertices X is a tight set if X is a S-separating set and $d_G(X) = 2k$. A tight set is non-trivial if $|X| \geq 2$ and $|V(G) - X| \geq 2$. The extension property will be used to establish the following lemma.

Lemma 4.6.3 *$\mathcal{G}$ has no non-trivial tight set.*

Proof. Let X be a nontrivial cut of G, and v be the vertex to be extended. Apply the cut decomposition operation on X to obtain two graphs G_1 and G_2. So, $V(G_1) = X \cup \{v_1\}$, $V(G_2) = (V(\mathcal{G}) - X) \cup \{v_2\}$ and $\delta_{\mathcal{G}}(X) = \delta_{G_1}(v_1) = \delta_{G_2}(v_2)$. Since X is non-trivial, both G_1 and G_2 are smaller than G. We assume without loss of generality that $v \in G_1$. Let $S_1 := (S \cap V(G_1)) \cup \{v_1\}$ and $S_2 := (S \cap V(G_2)) \cup \{v_2\}$. Since S is $2k$-edge-connected in $\mathcal{G}$, S_1 and S_2 are $2k$-edge-connected in G_1 and G_2 respectively (by Proposition 3.4.1 (1)). By the minimality of $\mathcal{G}$, there is a Steiner strongly k-arc-connected orientation D_1 of G_1 that extends v. Notice D_1 defines a valid orientation on $\delta_{G_1}(v_1)$. Set the orientation of $\delta_{G_2}(v_2)$ to be consistent with the orientation of $\delta_{G_1}(v_1)$. So this is a valid orientation of $\delta_{G_2}(v_2)$. By the minimality of $\mathcal{G}$, there is a Steiner strongly k-arc-connected orientation D_2 of G_2 that extends v_2. By setting D to be the concatenation of D_1 and D_2, we claim that D is a Steiner strongly k-arc-connected orientation that extends v. By construction, D extends v.

We need to verify that D is a Steiner strongly k-arc-connected orientation of G. By Menger's theorem, it is equivalent to verify that $d_D^{in}(Y) \geq k$ for every S-separating set Y. Suppose $Y \subseteq V(G_1)$, then $d_D^{in}(Y) = d_{D_1}^{in}(Y) \geq k$ as required; and similarly for $Y \subseteq V(G_2)$. The interesting case is when $Y_1 := Y \cap V(G_1) \neq \emptyset$ and $Y_2 := Y \cap V(G_2) \neq \emptyset$. By the fact that Y is an S-separating set and X is an S-separating set, we can assume (renaming if necessary) that there is a terminal vertex y in Y_1 and a terminal vertex x in $V(G_2) - Y$. Since D_1 and D_2 are Steiner strongly k-arc-connected orientations, we have $d^{in}_{D_1}(Y_1) \geq k$ and $d^{in}_{D_2}(Y_2 + v_2) \geq k$. Notice that $d^{in}_{D_2}(Y_2 + v_2) = d^{in}_{D_2}(Y_2) + d^{in}_{D_2}(v_2) - \overrightarrow{d}_{D_2}(Y_2, v_2) - \overrightarrow{d}_{D_2}(v_2, Y_2)$. As v_2 is a tight vertex, $d^{in}_{D_2}(v_2) = k$. So $d^{in}_{D_2}(Y_2) - \overrightarrow{d}_{D_2}(v_2, Y_2) \geq \overrightarrow{d}_{D_2}(Y_2, v_2)$. Note that $\overrightarrow{d}_{D_2}(v_2, Y_2) \geq \overrightarrow{d}_{D}(Y_1, Y_2)$. Hence, $d^{in}_{D}(Y_2) - \overrightarrow{d}_{D}(Y_1, Y_2) \geq d^{in}_{D_2}(Y_2) - \overrightarrow{d}_{D_2}(v_2, Y_2) \geq \overrightarrow{d}_{D_2}(Y_2, v_2) \geq \overrightarrow{d}_{D_2}(Y_2, Y_1)$. Since $d^{in}_{D}(Y) = d^{in}_{D}(Y_1) + d^{in}_{D}(Y_2) - \overrightarrow{d}_{D}(Y_1, Y_2) - \overrightarrow{d}_{D}(Y_2, Y_1)$, we have $d^{in}_{D}(Y) \geq d^{in}_{D}(Y_1) \geq k$, as required. This completes the proof. ∎

By applying Mader's splitting-off lemma as in Section 3.5 (or Theorem 4.3.1, or Lemma 4.4.9), we can assume the following.

Lemma 4.6.4 *Every Steiner vertex is of degree 3 in $\mathcal{G}$ and is adjacent to exactly 3 vertices.*

The following are the consequences of Lemma 4.6.3.

Lemma 4.6.5 *There is no edge in $\mathcal{G}$ between two terminal vertices a, b with degree greater than $2k$. Also, there is no edge in $\mathcal{G}$ between a terminal vertex of degree greater than $2k$ and a Steiner vertex.*

Proof. Suppose $e = ab$ is in $\mathcal{G}$. By the minimality of $\mathcal{G}$, S is not $2k$-edge-connected in $\mathcal{G} - e$. So, by Menger's theorem, there is a S-separating set X with $e \in \delta(X)$ and $d(X) = 2k$. If a, b are terminal vertices of degree greater than $2k$, both $|X|, |V(\mathcal{G}) - X| \geq 2$, which contradicts Lemma 4.6.3. The proof of the second statement is analogous. ∎

Lemma 4.6.6 *There is no edge between two odd-degree vertices in $\mathcal{G}$. In particular, there is no edge between two Steiner vertices.*

Proof. The proof of the first statement is similar to the proof of Lemma 4.6.5, and then the second statement follows from Lemma 4.6.4. ∎

With the above lemmas, we can prove that $\mathcal{G}$ does not exist. If $|S| = 2$, then the theorem follows from Menger's theorem; we find $2k$ edge-disjoint paths and orient them consistently with the valid orientation of $\delta_{\mathcal{G}}(v)$. So assume $|S| \geq 3$. Then we claim that there must exist a terminal vertex $s \neq v$ which is of degree $2k$, where v is the vertex to be extended. Suppose, by way of contradiction, that every vertex in $S - v$ is of degree greater than $2k$. Then, by the first statement of Lemma 4.6.5, there is no edge between vertices in $S-v$. Also, by the second statement of Lemma 4.6.5, there is no edge between vertices in $S - v$ and $V(\mathcal{G}) - S$. So we must have $\delta(S - v) = \delta(v)$, but $d(S - v) > 4k$ (as $|S - v| \geq 2$) and $d(v) = 2k$ (by assumption), a contradiction. So there is a terminal vertex $s \neq v$ of degree $2k$ in $\mathcal{G}$.

Now, apply Mader's splitting-off lemma to completely split-off s; let the resulting graph be G'. By the choice of $\mathcal{G}$, there is a Steiner strongly k-arc-connected orientation D' of G'. If xy is an edge in G' resulting from splitting-off sx and sy of $\mathcal{G}$ and it is oriented as $\overrightarrow{xy}$ in D', then we orient sx and sy in $\mathcal{G}$ as $\overrightarrow{xs}$ and $\overrightarrow{sy}$ and call the resulting orientation D. We claim that D is a Steiner strongly k-arc-connected orientation of $\mathcal{G}$. By Menger's theorem, we just need to check that $d^{in}_D(X) \geq k$ for every S-separating set X.

Suppose $X \cap S = \{s\}$. We claim that $k = d^{in}_D(s) = \overrightarrow{d}(V(\mathcal{G}) - X, s) + \overrightarrow{d}(X - s, s) \leq \overrightarrow{d}(V(\mathcal{G}) - X, s) + \overrightarrow{d}(V(\mathcal{G}) - X, X - s) = d^{in}_D(X)$ holds. The equalities are easy, we will prove the inequality. By assumption, each vertex w in $X - s$ is a Steiner vertex, and thus w has degree 3 and is adjacent to 3 vertices. So each vertex w in $X - s$ has at most one outgoing arc to s. For each such w which has an outgoing arc to s, since D' is a strongly-connected orientation and $X - s$ is an independent set (by Lemma 4.6.6), w must have an incoming arc from $V(\mathcal{G}) - X$. Hence $\overrightarrow{d}(V(\mathcal{G}) - X, X - s) \geq \overrightarrow{d}(X - s, s)$ and the inequality in the above argument follows. The case when $X \cap S = S - s$ is similar; we consider $d^{out}_D(V(\mathcal{G}) - X) = d^{in}_D(X)$ and apply the previous argument on $d^{out}_D(V(\mathcal{G}) - X)$ since $(V(\mathcal{G}) - X) \cap S = \{s\}$.

Finally, suppose $X \cap (S - s) \neq \emptyset$ and $(V(\mathcal{G}) - X) \cap (S - s) \neq \emptyset$. If $s \in X$, then $d^{in}_D(X) \geq d^{in}_{D'}(X - s) \geq k$; if $s \notin X$, then $d^{in}_D(X) \geq d^{in}_{D'}(X) \geq k$. So, $d^{in}_D(X) \geq k$ as required. Therefore, D is a Steiner strongly k-arc-connected orientation of $\mathcal{G}$. This completes the proof. ∎

Finally we remark that our approach taken is similar to Mader's [71] proof of Nash-Williams' strong orientation theorem, but is considerably simpler in this special case.

4.7 Hardness Results

Nash-Williams' strong orientation theorem (Theorem 2.2.13) 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 [32]. 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 4.7.1 *The* Steiner Rooted-Orientation *problem is NP-complete.*

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 $R : V \times V \rightarrow Z^+$ 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 4.7.2 *[34] 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 $a_{u,v}$, $b_{u,v}$ 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. $R(u, v)$ edges from r to $a_{u,v}$ for every pair u, v,
2. $R(u, v)$ edges from $a_{u,v}$ to u for every pair u, v,
3. $R(u, v)$ edges from v to $b_{u,v}$ for every pair u, v,
4. for every pair of pairs (u, v) and (x, y) for which $u \neq x$ or $v \neq y$, $R(u, v)$ edges from $a_{u,v}$ to $b_{x,y}$.

Let

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

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

$$A := \{a_{u,v} : \ u, v \in V, \ u \neq v\}.$$

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

Lemma 4.7.3 *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 type 3 and 4 must be oriented towards S.

For any pair $(u, v) \in V \times V$, let us consider the set $X = V \cup (A - a_{u,v}) + b_{u,v}$. X must have in-degree at least k in D', which means, in the light of the above facts, that the edges from $a_{u,v}$ to u must be oriented towards u. Thus, all edges of type 2 are oriented towards V.

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 4. The remaining $R(u, v)$ paths necessarily start with the edges $ra_{u,v}$ and $a_{u,v}u$, 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'. ∎

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. ∎

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). In the following we prove that the corresponding in-hypergraph orientation problem is NP-complete even for $k = 1$; this contrasts with the out-hypergraph orientation problem that we studied in Theorem 4.1.1 for which the case $k = 1$ is easy.

Theorem 4.7.4 *Given a hypergraph $H = (V, E)$ with a root node $r \in V$ and a terminal set $S \subseteq V - r$, it is NP-complete to decide whether H has an in-hypergraph orientation* *such that there is a path from r to every terminal node.*

Proof. Let $H = (V, E)$ be a hypergraph. The hyperedge cover problem is to decide if V can be covered by k hyperedges. This is known to be NP-complete. We reduce this problem to our orientation problem.

We define a hypergraph $H' = (V', E')$ in the following way. Let $V' := V \cup \{v_e : e \in E\} + r$, let $S := V$ be the set of terminal vertices, and let E' consist of the following two types of hyperedges:

1. k parallel hyperedges consisting of the vertices $r, \{v_e : e \in E\}$,

2. for every $v \in V$, one hyperedge consisting of the vertices $v, \{v_e : v \in e\}$.

Let D' be an orientation of H'. There is a path from r to every terminal vertex if and only if the following hold:

- The hyperedges of type 2 are oriented towards V,

- Let T be the set of vertices of H' which are the heads of some type 1 hyperedges. Let E_T be the set of hyperedges in H which corresponds to T. Then $|E_T| \leq k$ and E_T covers V.

It follows that the hyperedge cover problem for H is solvable if and only if the orientation problem is solvable for H'. This proves the in-hypergraph Steiner rooted-orientation problem is NP-complete even for $k = 1$. ∎

For element-connectivity, we show that the STEINER ROOTED-ORIENTATION problem is NP-complete.

Theorem 4.7.5 *The* ELEMENT-DISJOINT STEINER ROOTED ORIENTATION *problem is NP-hard.*

Proof. We show that 3-SAT can be reduced to the Steiner element-connected orientation problem. Suppose that we are given an instance of 3-SAT with variables $x_1, \ldots, x_k$ and clauses $c_1, \ldots, c_l$. We construct a graph G on the following set of nodes:

- A root r,
- Two Steiner nodes v_{x_i} and $v_{\neg x_i}$ for every variable x_i,
- Two terminal nodes s_j and s'_j for every clause c_j,
- 8 Steiner nodes for every clause c_j: $a_j^0, a_j^1, a_j^2, a_j^3$, $b_j^0, b_j^1, b_j^2, b_j^3$.

Let S be the set of terminal nodes, and let $k := 4l$, where l is the number of clauses. Let the graph G consist of the following edges:

- An edge between v_{x_i} and $v_{\neg x_i}$ for every i,
- An edge between r and a_j^α for every j and every α,
- Edges from b_j^1, b_j^2 and b_j^3 to b_j^0 and to s'_j for every j,
- An edge between b_j^0 and s_j for every j,
- If x is the α-th literal in c_j, then edges from v_x to a_j^α and b_j^α $(\alpha \in \{1, 2, 3\})$, and an edge between $v_{\neg x}$ and a_j^0,
- Edges from a_j^0 to every terminal node except for s_j,
- Edges from a_j^1, a_j^2 and a_j^3 to every terminal node except for s'_j.

We shall show that the graph G has a Steiner rooted k-element-connected orientation if and only if the 3-SAT formula is satisfiable.

Let D be a Steiner rooted k-element-connected orientation of G. Since the degree of r and of every terminal node is k in G, the terminals must have in-degree k and r must have out-degree k in D. This means that we already have $k-1$ paths of length 2 from

r to each s_j (through a_i^α except a_j^0), and we have $k-3$ paths of length 2 from r to each s'_j (through a_i^α except a_j^1, a_j^2, a_j^3).

As for the remaining 3 paths from r to each s'_j, their second nodes must be a_j^1, a_j^2 and a_j^3, and their last non-terminal nodes must be b_j^1, b_j^2 and b_j^3. This means that for each literal $x \in c_j$, the edge between v_x and a_j^α must be oriented towards v_x, and the edge between v_x and b_j^α must be oriented towards b_j^α (there is no other way to complete the paths).

Let us consider the remaining one path from r to s_j. The second node of the path is a_j^0, the last non-terminal node is b_j^0, and the node before that is b_j^1, b_j^2 or b_j^3. By taking into account what we have already proved about the orientations of the edges, and the fact that all other nodes $a_{j'}^0$ $(j' \neq j)$ are used by some other path, the path can only be the following for some $\alpha \in \{1,2,3\}$:

$$\{r, a_j^0, v_{\neg x}, v_x, b_j^\alpha, b_j^0, s_j\},$$

where x is the α-th literal in c_j. It follows that our 3-SAT formula is satisfied if we set x_i to be true if the edge $(v_{x_i}, v_{\neg x_i})$ is oriented towards v_{x_i} in D, and we set x_i to be false otherwise.

Now we prove that if the 3-SAT formula can be satisfied, then there is a Steiner rooted k-element-connected orientation. As we have shown in the above paragraphs, the orientation of several edges is forced, and they give k element-disjoint paths from r to each s'_j, and $k-1$ element-disjoint paths from r to each s_j.

We orient the edge $(v_{x_i}, v_{\neg x_i})$ towards v_{x_i} if x_i is true in the valuation satisfying the formula, and orient it towards $v_{\neg x_i}$ if x_i is false. The edges of type (b_j^α, b_j^0) are oriented towards b_j^0, and the edges of type (a_j^0, v_x) are oriented towards v_x.

Suppose that x is the α-th literal in c_j, and it is true in the valuation satisfying the formula. The following path is element-disjoint from the $k-1$ paths already given from r to s_j:

$$\{r, a_j^0, v_{\neg x}, v_x, b_j^\alpha, b_j^0, s_j\}.$$

This shows that there are k element-disjoint paths from r to each terminal. This completes the proof of the theorem. ∎

One can consider minimum cost versions of the orientation problems discussed in this chapter. 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 4.7.6 *The* MINIMUM COST STEINER ROOTED ORIENTATION *problem is NP-hard to approximate within a factor of* $\Omega(\log(n))$*, even for* $k = 1$*.*

Proof. We reduce the SET COVER problem (which is NP-hard to approximate within a factor of $\Omega(\log(n))$ [20]) to the Min Cost Steiner Rooted Orientation problem, such that the number of sets in the cover corresponds to the cost of the orientation.

Given an instance of the set cover problem with ground set V and a family $\mathcal{F}$ of sets with union V, we define a graph $G' = (V', E')$, and edge costs for both orientations of each edge. Let V' consist of the following nodes:

- the nodes in V,
- a node v_Z for each $Z \in \mathcal{F}$,
- a root r.

The set of terminal nodes is $V \cup \{r\}$. The graph G' consists of two types of edges, with the following costs for their orientations:

1. An edge rv_Z for each $Z \in \mathcal{F}$. The cost is 0 if oriented towards r, and 1 if oriented towards v_Z.

2. Edges between v_Z and each node in Z, for every $Z \in \mathcal{F}$. The cost is 0 if the edge is oriented towards V, and the cost is $|V|$ if the edge is oriented towards v_Z.

Since the union of the sets in $\mathcal{F}$ is V, there is a Steiner rooted connected orientation of cost at most $|V|$: for each node $u \in V$ we select an arbitrary set $Z_u \in \mathcal{F}$ containing u; we orient the edges rv_{Z_u} towards v_{Z_u}, orient the other edges of type 1 towards r, and orient each edge of type 2 towards V. We can thus assume that in a minimum cost orientation every edge of type 2 is oriented towards V.

Such an orientation is Steiner rooted connected if and only if the family

$$\{Z \in \mathcal{F} : \text{the edge } rv_Z \text{ is oriented towards } v_Z\}$$

is a set cover. So the cost of the minimum cost orientation equals the number of sets in a minimum cover. ∎

Chapter 5

Concluding Remarks

In this thesis, we have studied two graph connectivity problems: the Steiner Tree Packing problem and the Steiner Rooted-Orientation problem. The main results are approximate min-max relations which show that the optimal values of our problems are within a constant factor to the natural connectivity upper bound. The proofs of these approximate min-max relations also give polynomial time constant factor approximation algorithms to these NP-complete problems.

For the Steiner Tree Packing problem, there are several directions for future work. The most exciting open problem is to settle Kriesell's conjecture. Although I believe that the conjecture is true, the current technique does not seem to be strong enough to prove it. Another direction is to prove an approximate min-max relation for the Steiner Network Packing problem as discussed in Section 3.11; this also seems to be very challenging. Motivated by the hypergraph orientation result (Theorem 4.1.1), it would be interesting to study the Steiner Tree Packing problem in hypergraphs. We know that it is NP-hard to approximate this problem to within a factor of $\Omega(\log n)$. On the other hand, we do not have any algorithmic result on this problem.

The main result (Theorem 4.1.1) on the Steiner Rooted-Orientation problem makes the first step to generalize Nash-Williams' strong orientation theorem to hypergraphs, but a full generalization is still far away. The main idea of Theorem 4.1.1 is to reduce the STEINER ROOTED-ORIENTATION problem in hypergraphs to a graph orientation problem concerning edge-connectivity, on which the submodular flow technique can be applied. However, this approach would not work for more general hypergraph orientation problems (e.g. obtaining a Steiner strongly k-hyperarc-connected orientation). Moreover, we do not know of any tool to deal with orientation problems beyond edge-connectivity. This is also the reason that there is very little progress in the literature on orientation problems concerning vertex-connectivity and hyperarc-connectivity. It is a common belief 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 G, is it true that G has a Steiner strongly k-element-connected orientation? I believe that settling it would be a major step towards other orientation problems concerning vertex-connectivity and hyperarc-connectivity.

Finally, the proofs of the main results in this thesis use a new technique of graph decomposition by introducing some appropriate extension properties. It would be nice to see this technique used to solve more problems.

Bibliography

[1] A. Agarwal, M. Charikar. On the Advantage of Network Coding for Improving Network Throughput. *Proceedings of the IEEE Information Theory Workshop*, 2004.

[2] R. Ahlswede, N. Cai, S.-Y.R. Li and R.W. Yeung. Network information flow. *IEEE Trans. on Information Theory*, vol. 46, pp. 1004-1016, 2000.

[3] M. Bern, P. Plassmann. The Steiner problem with edge lengths 1 and 2. *Information Processing Letters*, 32, pp.171-176, 1989.

[4] D. Bienstock, A. Bley. Capacitated network design with multicast commodities. *Proceedings in the 8th International Conference on Telecommunication Systems*, 2000.

[5] J. Bang-Jensen, G. Gutin. *Digraphs: Theory, Algorithms and Applications*. Springer-Verlag, London, 2000.

[6] J. Bang-Jensen, S. Thomassé. Decompositions and orientations of hypergraphs. *manuscript*, 2001.

[7] J. Bang-Jensen, S. Thomassé. Highly connected hypergraphs containing no two edge-disjoint spanning connected subhypergraphs. *Discrete Applied Mathematics*, 131(2), pp. 555-559, 2003.

[8] D.M. Blough, M. Leoncini, G. Resta, P. Santi. The K-neigh protocol for symmetric topology control in ad hoc networks. *Proceedings of ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, June 2003.

[9] S. Chen, O. Günlük, B. Yener. Optimal packing of group multicastings. *IEEE INFOCOMM'98*, 1998.

[10] S. Chen, O. Günlük, B. Yener. The multicast packing problem. *IEEE/ACM Transactions on Networking (TON)*, 8, pp.311-318, 2000.

[11] R. Cole, R. Hariharan. A fast algorithm for computing Steiner edge connectivity. *Proceedings of the thirty-fifth ACM symposium on Theory of computing (STOC)*, pp.167-176, 2003.

[12] C. Chekuri, B. Shepherd. Approximate integer decompositions for undirected network design problems. *Manuscript*, 2004.

[13] J. Cheriyan, M. Salavatipour. Hardness and approximation results for packing Steiner trees problems. *Proceedings of ESA, Lecture Notes of Computer Science*, Vol. 3221, pp.180-191, 2004.

[14] J. Cheriyan, M. Salavatipour. Packing element-disjoint Steiner trees. *proceedings of APPROX 2005, Lecture Notes of Computer Science*, Vol. 3624, pp.52-61, 2005.

[15] R. Diestel. *Graph Theory.*. Springer-Verlag, Heidelberg, 2000.

[16] 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. *Proc. International Workshop on Wireless Ad-hoc Networks (IWWAN) 2005*, May 2005.

[17] J. Edmonds. Minimum partition of a matroid into independent sets. *J. Res. Nat. Bur. Standards Sect.*, 869, pp.67-72, 1965.

[18] J. Edmonds, R. Giles. A min-max relation for submodular function on graphs. *Annals of Discrete Mathematics*, 1, pp. 185-204, 1977.

[19] http : //www.cs.elte.hu/egres/problems/prob_02/prob_02.html

[20] U. Feige, M. Halldorsson, G. Kortsarz, A. Srinivasan. Approximating the domatic number. *SIAM Journal on Computing*, 32, pp. 172-195, 2002.

[21] L.K. Fleischer, K. Jain, D.P. Williamson. An Iterative Rounding 2-approximation Algorithm for the Element Connectivity Problem. *Proceedings of the 42nd annual conference on the Foundation of Computer Science (FOCS)*, pp. 339-347, 2001.

[22] A. Frank. On the orientations of graphs. *J. Combin. Theory Ser. B*, 28(3), pp. 251-261, 1980.

[23] A. Frank. An algorithm for submodular functions on graphs. *Ann. Discrete Math.*, pp. 97-120, 1982.

[24] A. Frank. Finding feasible vectors of Edmonds-Giles polyhedra. *J. Combin. Theory Ser. B*, 36(3), pp. 221-239, 1984.

[25] A. Frank. Submodular flows. *Progress in Combinatorial Optimization*, pp. 147-165. Academic Press, Toronto, Ontario, 1984.

[26] A. Frank. On a theorem of Mader. *Annals of Discrete Mathematics*, 101, 1992, pp. 49-57.

[27] A. Frank. Applications of submodular functions. In *Surveys in combinatorics, 1993 (Keele)*, volume 187 of *London Math. Soc. Lecture Note Ser.*, pp. 85-136. Cambridge Univ. Press, Cambridge, 1993.

[28] A. Frank. Submodular functions in graph theory. *Discrete Mathematics, Graph Theory and Combinatorics*, 111, pp. 231-243, 1993.

[29] A. Frank. Connectivity and network flows. *Handbook of Combinatorics*, Elsevier, Amsterdam, pp. 111-177, 1995.

[30] A. Frank. Orientations of graphs and submodular flows. *Congr. Numer.*, 113, pp. 111-142, 1996. Festschrift for C. St. J. A. Nash-Williams.

[31] A. Frank. Edge-connection of graphs, digraphs, and hypergraphs. *Egerváry Research Group, Budapest*, TR-2001-11, 2001.

[32] A. Frank, T. Király. Combined connectivity augmentation and orientation problems. *Discrete Applied Mathematics*, 131, pp. 401-419, 2003.

[33] A. Frank, T. Király, Z. Király. On the orientation of graphs and hypergraphs. *Discrete Applied Mathematics*, 131, pp. 385-400, 2003.

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

[35] H. Gabow. Efficient splitting off algorithms for graphs. *Proceedings of the 26th annual ACM Symposium on Theory of Computation (STOC94)*, pp. 696-705, 1994.

[36] N. Garg, R. Khandekar, K. Kunal, V. Pandit. Bandwidth maximization in multicasting. *11th Annual European Symposium on Algorithms (ESA)*, 2003.

[37] M. Grötschel, A. Martin, R. Weismantel. Packing Steiner trees: further facets. *European J. Combin.*, 17, pp.39-52, 1996.

[38] M. Grötschel, A. Martin, R. Weismantel. Packing Steiner trees: polyhedral investigations. *Mathematical Programming*, 72, pp.101-123, 1996.

[39] M. Grötschel, A. Martin, R. Weismantel. Packing Steiner trees: a cutting plane algorithm and computational results. *Mathematical Programming*, 72, pp.125-145, 1996.

[40] M. Grötschel, A. Martin, R. Weismantel. Packing Steiner trees: separation algorithms. *SIAM J. Discrete Math.*, 9, pp.233-257, 1996.

[41] M. Grötschel, A. Martin, R. Weismantel. The Steiner tree packing problem in VLSI-design. *Mathematical Programming*, 78, pp.265-281, 1997.

[42] M. Goemans, D. Williamson. A general approximation technique for constrained forests problems. *SIAM Journal on Computing*, 24, pp.296-317, 1995.

[43] H.R. Hind, O. Oellermann. Menger-type results for three or more vertices. *Congressus Numerantium*, 113, pp.179-204, 1996.

[44] I. Holyer. The NP-completeness of edge-coloring. *SIAM Journal on Computing*, 10, pp. 718-720, 1981.

[45] M. Hosseini, N.D. Georganas. Design of a multi-sender 3D videoconferencing application over an end system multicast protocol. *Proceedings of the eleventh ACM international conference on Multimedia*, pp.480-489, 2003.

[46] K. Jain, I. Mãndoiu, V.V. Vazirani, D.P. Williamson. A Primal-dual Schema Based Approximation Algorithm for the Element Connectivity Problem. *Proceedings of the 10th annual Symposium on Discrete Algorithms (SODA)*, pp. 484-489, 1999.

[47] K. Jain. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica*, 21, pp.39-60, 2001.

[48] T. Jordán. On the existence of k edge-disjoint 2-connected spanning subgraphs. *J. Combin. Theory Ser. B*, 95(2), pp. 257-262, 2005.

[49] K. Jain, M. Mahdian, M.R. Salavatipour. Packing Steiner trees. *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pp.266-274, 2003.

[50] K. Jain, I. Mãndoiu, V.V. Vazirani, D.P. Williamson. A Primal-dual Schema Based Approximation Algorithm for the Element Connectivity Problem. *Proceedings of the 10th annual Symposium on Discrete Algorithms (SODA)*, pp. 484-489, 1999.

[51] G.W. Jeong, K. Lee, S. Park, K. Park. A branch-and-price algorithm for the Steiner tree packing problem. *Computers and Operations Research*, 29, pp.221-241, 2002.

[52] D. Karger. Minimum cuts in near-linear time. *Journal of the ACM (JACM)*, 47, pp.46-76, 2000.

[53] T. Király. *Edge-connectivity of undirected and directed hypergraphs*. Ph.D. thesis, Eötvös Loránd University, 2003.

[54] T. Király, L.C. Lau. Tight approximate min-max relations for Steiner rooted-orientations of graph and hypergraphs. *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006.

[55] Z. Király, Z. Szigeti. Simultaneous well-balanced orientations of graphs. *J. Combin. Theory Ser. B*, in press, 2006.

[56] B. Korte, H.J. Prömel, A. Steger. Steiner trees in VLSI-layout. *Paths, Flows, and VLSI-layout, Springer-Verlag*, pp.185-214, 1990.

[57] M. Kriesell. Local spanning trees in graphs and hypergraph decomposition with respect to edge-connectivity. *Technical Report 257, University of Hannover*, 1999.

[58] M. Kriesell. Edge-disjoint trees containing some given vertices in a graph. *J. Combin. Theory, Series B*, 88, pp.53-63, 2003.

[59] M. Kriesell. Disjoint Steiner trees in graphs without large bridges. *manuscript*, 2004.

[60] L.C. Lau. An approximate max-Steiner-tree-packing min-Steiner-cut theorem. *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 61-70, 2004.

[61] L.C. Lau. Packing Steiner Forests. *Proceedings of the Eleventh Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pp. 362-376, 2005.

[62] T. Leighton, S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM (JACM)*, 46, pp.787-832, 1999.

[63] Z. Li, B. Li, D. Jiang, L.C. Lau. On achieving optimal throughput with network coding. *IEEE INFOCOM'05*, 2005.

[64] Z. Li, B. Li, L.C. Lau. On Achieving optimal multicast throughput in undirected networks. *Joint Special Issue on Networking and Information Theory, IEEE Transactions on Information Theory (IT) and IEEE/ACM Transactions on Networking (TON)*, to appear in June 2006.

[65] S.Y.R. Li, R.W. Yeung, N. Cai. Linear Network Coding. *IEEE Trans. on Information Theory*, vol. 49, pp. 371-381, 2003.

[66] L. Lovász. On two minimax theorems in graph. *Journal of Combinatorial Theory (B)*, 21, 1976, pp. 96-103.

[67] L. Lovász. *Combinatorial Problems and Exercises*, North-Holland, 1993.

[68] D.S. Lun, M. Médard, R. Koetter. Network coding for efficient wireless unicast. *IEEE International Zurich Seminar on Communications*, February 2006.

[69] D.S. Lun, N. Ratnakar, M. Médard, R. Koetter, D.R. Karger, T. Ho, E. Ahmed, F. Zhao. Minimum-cost multicast over coded packet networks. *Joint Special Issue*

on Networking and Information Theory, IEEE Transactions on Information Theory (IT) and IEEE/ACM Transactions on Networking (TON), to appear in June 2006.

[70] D.S. Lun, M. Médard, R. Koetter, M. Effros. On Coding for Reliable Communication over Packet Networks. *submitted to IEEE Trans. Inform. Theory*, 2006.

[71] W. Mader. A reduction method for edge-connectivity in graphs. *Ann. Discrete Math.*, 3, pp.145-164, 1978.

[72] A. Martin, and R. Weismantel. Packing paths and Steiner trees: routing of electronic circuits. *CWI Quarterly*, 6, pp.185-204, 1993.

[73] K. Menger. Zur allgemeinen Kurventheorie. *Fund. Math.*, 10, pp.95-115, 1927.

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

[75] C.St.J.A. Nash-Williams. Edge disjoint spanning trees of finite graphs. *J. London Math. Soc.*, 36, pp.445-450, 1961.

[76] Y. Ofek, B. Yener. Reliable concurrent multicast from bursty sources. *IEEE Journal of Selected Areas in Communications*, 15, pp.434-444, 1997.

[77] H. Okamura. Paths and edge-connectivity in graphs. *Journal of Combinatorial Theory (B)*, 37, pp.151-172, 1984.

[78] L. Petingi, J. Rodriguez. Bounds on the maximum number of edge-disjoint Steiner trees of a graph. *Congressus Numerantium*, 145, pp.43-52, 2000.

[79] O. Pikhurko, *personal communication*, 2005.

[80] W.R. Pulleyblank. Two Steiner tree packing problems. *Proceedings of the twenty-seventh annual ACM Symposium on Theory of computing (STOC)*, pp.383-387, 1995.

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

[82] G. Robins, A. Zelikovsky. Improved Steiner tree approximation in graphs. *Proceedings of the tenth annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp.770-779, 2000.

[83] P. Sanders, S. Egner, L. Tolhuizen. Polynomial time algorithms for network information flow. *Proceedings of the fifteenth annual ACM symposium on Parallel algorithms and architectures*, pp.286-294, 2003.

[84] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency.* Springer-Verlag, Berlin, 2003.

[85] W.T. Tutte. On the problem of decomposing a graph into n connected factors. *J. London Math. Soc.*, 36, pp.221-230, 1961.

[86] E. Uchoa and M. Poggi de Aragão. Vertex-disjoint packing of two Steiner trees: polyhedra and branch-and-cut. *Mathematical Programming*, 90, pp.537-557, 2001.

[87] C.F. Wang, C.T. Liang, R.H. Jan. Heuristic algorithms for packing of multiple-group multicasting. *Computers and Operations Research*, 29, pp.905-924, 2002.

Index